

The Social Cost of Carbon: Valuing Inequality, Risk, and Population for Climate Policy

Marc Fleurbaey, Maddalena Ferranna, Mark Budolfson, Francis Dennig, Kian Mintz-Woo, Robert Socolow, Dean Spears, and Stéphane Zuber

ABSTRACT

We analyze the role of ethical values in the determination of the social cost of carbon, arguing that the familiar debate about discounting is too narrow. Other ethical issues are equally important to computing the social cost of carbon, and we highlight inequality, risk, and population ethics. Although the usual approach, in the economics of cost-benefit analysis for climate policy, is confined to a utilitarian axiology, the methodology of the social cost of carbon is rather flexible and can be expanded to a broader set of social-welfare approaches.

1. INTRODUCTION

The Social Cost of Carbon (SCC) is likely the most important economic concept in the design and implementation of climate change mitigation policies. The SCC provides a monetized value of the present and future damages caused by the emission of a ton of CO₂. The SCC depends on ethical premises because it depends on how climate change impacts are valued and aggregated.

Estimates of the SCC are a crucial component in the cost-benefit analysis of climate policies. Policy makers can use the SCC in either of two ways: to set the optimal carbon tax, or to fix the emission reduction target in a cap-and-trade system. The higher the value of the SCC, the larger the optimal mitigation effort. More broadly, the SCC provides regulators with guidance for the introduction of climate-affecting policies, such as energy efficiency standards in buildings and transportation, incentives for the installation of low-carbon technologies, and measures for the protection of forests.¹

Efforts to compute the SCC have given rise to a lively discussion in the literature and policy arena. The Stern-Nordhaus debate is probably the most notable example. Nordhaus (2008) was in favor of a very low SCC around \$8 per ton of CO₂,² but Stern (2006) proposed a much larger value of \$85 per ton of CO₂. Because the two analyses share comparable assumptions about the evolution of the climate and the

MFl: Princeton University, MFe: Harvard University, MB: University of Vermont, FD: Yale-NUS, KMW: Graz University, RS: Princeton University, DS: University of Austin, SZ: Paris School of Economics.

size of its impacts, the divergence in results is mainly attributable to different ethical values, in particular concerning discounting, i.e., the reduced weight assigned to future monetary values (Nordhaus 2007).³ The range of cost estimates stemming from different ethical assumptions is also reflected in the official U.S. SCC estimates by the Environmental Protection Agency (EPA), which up to recently⁴ provided a range of values depending on the discount rate and the size of impacts (ranging from \$11 to \$105 for the year 2015, with a central value of \$36).

This paper analyzes the role of ethical values in the determination of the SCC. In particular, we go beyond the familiar debate about discounting. Other ethical issues are equally important to computing the SCC: namely, inequality, risk, and population ethics. An important question we examine in the process is whether the SCC methodology is tied to an excessively narrow (e.g., utilitarian, economic) ethical approach, or whether it can accommodate a variety of relevant principles and values. The thesis defended here is that while usual practice is indeed unduly narrow, the methodology itself is rather flexible and can be expanded.

The SCC belongs to a cost-benefit approach to climate policy which, in general, relies on a consequentialist and even more specifically a welfarist axiology. There are other approaches involving notions of harm, rights, or duties which may reach different conclusions about the appropriate mitigation efforts by individuals, organizations, and states. Broome (2012) argues that states can focus on the pursuit of goodness (understood, in his view, in broadly welfarist terms) while individuals can focus on avoiding actions that harm other people (in this case, future people), including by compensating potential harms through carbon offsetting. This moral and practical division of labor can be debated (individuals may also contribute to the good; groups and states may also be liable for harmful actions). However, it is widely accepted that coordinating the pursuit of the common good in the presence of diffuse externalities is difficult without the aid of public agencies suitably interconnected across the affected jurisdictions—in the case of climate, the whole world. The SCC is an important tool for such coordination and clearly pertains to this axiological public morality.

The paper is structured as follows. Section 2 describes how to compute the SCC and why it is important for policy-making. Section 3 introduces the most widely-studied ethical aspects of the SCC and focuses in particular on inter- and intragenerational inequalities. The section is divided into two parts: the first focuses on the discounting problem, the second on the evaluation of nonmonetary damages. Section 4 discusses the additional considerations that should be taken into account when risk and uncertainty are introduced. Section 5 explains why ethical questions about population size matter for the SCC. Section 6 concludes.

2. WHY AND HOW TO COMPUTE THE SCC

2.1. Cost-benefit analysis and externality pricing

The “social cost” of carbon is defined in opposition to the “private cost,” which is borne directly by the economic decision makers who choose how much fossil fuel to use. The private cost reflects the scarcity of fossil fuels and the economic resources needed to produce the final carbon-emitting product that is ultimately consumed. In

the absence of an externality, the private cost would embody the full extent of the cost to others of one individual consuming a unit of carbon-emitting goods. But due to the greenhouse effect, the consumption of that unit imposes a further cost on others: the damages due to climate change. This is the social cost.

Against these costs, there are the benefits of activities that emit carbon, such as heating or traveling. Rational individuals who maximize a private objective will consume up to the point where their marginal private benefit is equal to the marginal private cost they bear. If there is a social cost they are not internalizing, individuals will overconsume relative to what would maximize overall welfare. Overall welfare maximization requires equalizing the marginal benefit with the total (i.e., private *plus* social) marginal cost. The socially optimal allocation can be achieved by setting up a mechanism, tax or cap-and-trade, by which the social cost is internalized by private agents.

In their classical approach to the social cost, economists focus on overall welfare. One might insist on including concerns of *justice* in climate affairs—beyond overall welfare—and worry that the social-cost approach appears unable to do that. The fact that the SCC is aimed at equalizing marginal benefits and costs sheds light on this issue. Equalizing marginal benefits and costs makes sense if morality is defined in terms of maximizing a certain magnitude—a “value” function—but makes less sense if morality involves side constraints, such as respecting rights or avoiding harm. Side constraints (such as rights) can be approached as a limit case of the social cost approach: if the social cost were set to be larger than any other consideration in the value function, then internalizing the social cost would act as a side-constraint (Jackson and Smith 2006). However, the social-cost approach is not well suited to agent-specific considerations which are common in justice approaches. For instance, consider a justice approach according to which no individual should emit more than a threshold: injustice occurs as soon as one individual trespasses, even if global emissions are low. This would require an individualized SCC that would depend on who emits the marginal ton of CO₂.

Another concern with the standard approach is that it suggests that a unique tax, corresponding to the magnitude of the SCC, should be imposed on all private agents in order to maximize overall welfare. But this recommendation is not always correct. A unique tax equal to the SCC is optimal if the distribution of resources among individuals is socially optimal. If inequalities are excessive and will not be corrected by inequality-reducing transfers implemented in addition to climate policy, then a single carbon tax is, actually, socially worse than differentiated taxes that favor disadvantaged populations (Chichilniski and Heal 1994).

It is often objected that differentiated taxes would generate an inefficient distribution of abatement, because the marginal cost of abatement would be greater where the tax is greater, thereby failing to minimize the total abatement cost.⁵ But this objection assumes that minimizing the total abatement cost is a good objective, which is not the case if the cost is distributed among individuals with unequal social priority (because some are poorer).⁶

To be sure, direct transfers from the rich to the poor would be a more effective redistributive policy than a set of different carbon taxes for the rich and the poor. But as long as the politics of taxes and transfers remain complex and constraining,

differentiated taxes may be the best available policy. In practice, given the reality and the political salience of the development needs of poor countries, advocacy for a globally uniform carbon price generates predictable and understandable political opposition. Pottier (2016) argues that by insisting on a uniform carbon price, economists have been an obstacle to the progress of climate negotiations.⁷

The fact that climate mitigation policy must consider existing inequalities is an instance of a general rule in public economics: the principle of “second-best” policy. It states that it is impossible to separate equity and efficiency considerations, and impossible to treat any market imperfection without taking account of how the other imperfections are being treated. It is tempting to treat climate externalities exclusively as an inefficiency problem due to private agents failing to internalize the costs they impose on others. But even if climate policy is not meant to redistribute resources, it will not maximally contribute to overall welfare if it does not take account of the prevailing distributional imperfections. It can be significantly welfare superior to tax carbon emissions differently in rich versus poor regions, in the realistic case that global economic inequalities will persist (Anthoff 2009; Budolfson and Dennig forthcoming).

This principle of second-best policy is vexing because it enormously complicates policy analysis. It is seldom respected by analysts, and even this paper will only partly apply it, because it will focus on climate policy without considering other imperfections than a suboptimal distribution of resources.

2.2 Integrated assessment and discounted utilitarianism

The SCC is normally computed with an Integrated Assessment Model (IAM),⁸ which is a model combining descriptions of the economy and of the climate system. Climate-economy interactions go both ways. Production and consumption processes generate carbon emissions, which perturb the climate system and induce an increase in temperature. The perturbed climate generates damages which can range from agricultural and productivity losses to health and biodiversity impacts. These damages affect the economy side of the model, by shrinking production and consumption possibilities, thus reducing society’s well-being.

IAMs keep track of the evolution of the economy and the climate over time and across individuals. A social welfare function (SWF) is the part of the model that aggregates the well-being of different generations and individuals. The SWF takes the role of society’s value function: it is the objective to be maximized by the choice of policy.

Standard IAMs adopt a “discounted utilitarian” SWF, which gives an exponentially decreasing weight to individuals living in the future. Individuals derive utility only from a consumption variable that includes the value of nonmonetary climate impacts, measured in monetary equivalent. In mathematical terms, this SWF reads as follows:

$$\sum_t \frac{1}{(1 + \rho)^t} \left(\sum_{i \in N_t} u(c_i) \right)$$

where t indexes time periods (one period per generation—in applied models a period is often a decade), ρ is the “time discount rate” which makes the weight of the

future decrease exponentially (the rate ranges from 0.1% to 1.5% per year in the literature), N_t is the population of generation t , c_i is the consumption of individual i , and u is the utility function.

The mathematical expression above can represent more ethical approaches than only utilitarianism. The individual utility function can be interpreted either as a measure of well-being, or as the social valuation of individual consumption. Its diminishing marginal value, depending on the interpretation, can reflect either the classical “diminishing marginal utility” pattern, or the ethical principle that the poor have greater priority. In this paper we use the word “utility” to refer to the value of the utility function in the SWF, independently of whether it is interpreted as measuring well-being or as another ethical valuation.⁹

In IAMs, inequalities are often treated in questionable ways. A frequent simplifying assumption is that there is no inequality in consumption across individuals belonging to the same generation; this assumption simply ignores the fact that such inequality exists. In the models with heterogeneous regions of the world, the intragenerational distributive problem is sometimes assumed away by introducing so-called “Negishi weights” (e.g., Nordhaus 2008), which weight the utility in different regions. These weights are set proportionally to the inverse of marginal utility, in order to make the distribution across regions of a given total of consumption optimal (no matter how unequal).¹⁰ This practice has been criticized in policy circles as antagonistic to the idea of equitable burden-sharing (Ackerman et al. 2011). For their supporters, Negishi weights are meant to insulate the analysis of climate policy from distributive problems, which are supposed to be a separate issue, belonging to a different branch of government or civil society. Yet, as explained in 2.1, in a second-best setting—which realistically describes our unequal world—such insulation is not possible in general and may be counterproductive.

The issue of intergenerational inequality is similarly dealt with by assuming optimal capital accumulation. Like Negishi weights, the discount parameter ρ in the SWF is adjusted so that the relative value of present and future consumption for the social objective reproduces observed market rates of return on investment. This goes some way toward making the observed and extrapolated trends appear optimal according to the SWF.¹¹

Even though discounted utilitarianism is the most common welfare criterion, other aggregation methods have been proposed, and sometimes used, in the climate change literature. For example, Adler et al. (2016) drop the time-discounting term in the SWF and introduce an additional value function applied to individual utility and embodying a special priority for the worse-off in addition to the already present diminishing marginal utility of consumption. The rationale is to better represent the prioritarian approach in climate policy analysis. However, as already noted in n. 9, the formula for the utilitarian SWF introduced above already accommodates this approach. Formally, its “utility” function can be interpreted as the combination of two functions, a well-being function depending on consumption, and a value function applied to well-being with a diminishing marginal value. Dietz and Asheim (2012) have also studied optimal climate policy (and indirectly the SCC) for a sustainable discounted utilitarian SWF that discounts future utility only if future generations are

better-off, introducing a form of prioritarianism through the discount rate (see also Zuber and Asheim [2012] for a similar idea).

2.3 The social cost of carbon

The value of the SCC today represents the present value, as assessed by the SWF, of the costs borne by present and future generations due to the emission of an additional ton of CO₂ today. If the future impacts of a ton of CO₂ emitted now reduce the SWF by a magnitude M and a reduction of present consumption by \$1 reduces the SWF by a magnitude m , then the SCC is M/m , i.e., the number of times a reduction of present consumption by \$1 needs to be done to have the same social welfare impact as emitting one ton of CO₂.

It is helpful to decompose this as a four steps procedure:

1. projecting the impact of an extra ton of CO₂ on the climate at all future dates by using the climate module of the IAM;
2. identifying the physical consequences of those climate alterations on the future environment, on production and consumption possibilities, and on nonmarket sources of value;
3. assessing those physical impacts in monetary terms;
4. computing the present value of these future monetary damages.

The first two steps are empirical, as they depend on our knowledge about the evolution of the climate and the consequences of letting it change as average temperature rises. The last two steps, instead, involve ethical considerations, since they depend on: i) how we value nonmarket goods relative to commodities; and ii) how we value future dollars in terms of present dollars.

The fourth step crucially relies on the SWF as the metric making future impacts on consumption and utility commensurate to costs on present consumption and utility: if \$1 taken from generation t decreases the SWF by M_t , then the present value of this dollar is equal to M_t/m , where m is again the decrease in SWF induced by taking \$1 from the present generation. The conversion factor M_t/m directly reflects the “social discount rate” which is discussed later in 3.1.

Two points are worth noting. First of all, IAMs generally simplify the description of the population by proceeding as if every decade contained a different generation. The damages occurring in a particular decade are thus counted as affecting the generation of that decade. A more realistic approach would have overlapping generations, each individual living for a few decades. Individual utility would then be *lifetime utility*, and damages occurring in a decade would modify the utility of all individuals living in that decade who incur some impact. In particular, a thorough modeling of overlapping generations would make it possible to depict inequalities in longevity, not just consumption inequalities.¹²

Second, the conversion of everything into monetary values is to some extent a matter of convention, since in theory a more general form of SWF could allow for individual utility functions depending on all multidimensional aspects of people’s lives.

Modifying the SWF in this way would increase the dimensionality of the computation, however, so in practice this option is set aside for convenience.

The SCC methodology is as general as the SWF on which it relies. Different assumptions about the ethical parameters of the SWF affect the level of both the social discount rate and monetary damages borne by generation t . To this we turn in the next section.

3. ETHICAL VALUES AND INEQUALITIES

This section highlights the importance of different ethical assumptions for the level of the SCC. Two main ethical variables determine the SCC: the social discount rate and the monetary value of climate damages.

3.1 The social discount rate (SDR)

The value of future monetized impacts in terms of current dollars is computed by applying the social discount rate (SDR) to future dollars. The social discount rate is expressed as an annual rate. If the rate is, for instance, 1% per year for the evaluation of dollars accruing in 100 years, then the weight applied to impacts in 100 years is $1/(1.01)^{100}$, amounting to 37¢ today per dollar in 100 years.

The SDR need not be independent of the date at which the future dollars are evaluated, so that the sequence of discount factors over the years need not take a simple exponential form such as $1/(1.01)^t$, where t denotes the year. As explained below, the SDR depends in part on the average growth rate between now and the future period for which it is computed; this average growth rate varies over time when the annual growth rate itself varies.

The exact value of the SDR is crucial for the evaluation of values far in the future, as in the typical century-like horizon of climate policy. For example, changing the SDR from 3% to 4% reduces the present value of a future dollar by 10% in 10 years and by 50% in 70 years. Such seemingly small modifications can dramatically alter the SCC.

The SDR can also, potentially, be negative, meaning that future dollars then have *greater* value than present dollars. One reason this can happen, as explained later in this section, is when the prediction of catastrophes makes the future appear likely to be poorer than the present generation.

There is a large literature analyzing the determinants of the SDR (see, e.g., Portney and Weyant [1999]; Gollier [2013]; Fleurbaey and Zuber [2015a]) as well as its ethical underpinnings (Fleurbaey and Zuber 2013; Greaves 2017). The following list summarizes the main reasons why future dollars should (or should not) be discounted:

A. Time discounting. Recall that the discounted utilitarian SWF has a discount term which bears on future utilities. This is a traditional assumption in economics, introduced mainly for practical reasons in models of optimal growth with an infinite horizon. An infinite horizon creates a summation problem for utilitarianism because the total population over time is then infinite. Pure time discounting is a way to make the sum of utilities finite.

Ramsey (1928) was adamantly opposed to introducing a pure present bias in the utilitarian sum and proposed an alternative solution.¹³ Nevertheless, time discounting ultimately became popular after Koopmans's (1960) influential axiomatization of discounted utilitarianism. The reason why Koopmans's solution won over Ramsey's may have to do with its mathematical simplicity. But, in a way that foreshadowed debates over climate mitigation policy, Koopmans's success may also have reflected the fact that the optimal growth paths obtained in numerical solutions required less staggering investment efforts from present generations under the time discounting approach than under an objective function with no discount.

In another vein, Arrow (1999) provided an important but controversial argument that a present bias is not only computationally convenient but is also ethically justified, on the grounds that it reflects a permissible agent-relative preference for ourselves and our own projects (Arrow cited Scheffler [1982] as an exemplar of the idea in the philosophical literature).

As noted in Stern (2006), there is an ethically uncontroversial reason to adopt a form of Koopmans's exponential discounting when there is an exogenous risk of catastrophic extinction. Such a risk of exogenous extinction does exist, because there are serious threats coming from meteors and from infectious diseases, which could exterminate a species like ours in a few generations. Stern assumes that the hazard rate is 1/1,000 per year, which may seem pessimistic (it implies a probability of extinction in the coming 100 years close to 10%), but is sufficiently low that, for SCC computations that are typically limited to a few centuries, it is not very different from a true zero.

However, most IAMs adopt a stronger time discount than does the Stern Review. The motivation, as argued for instance by Nordhaus (2007), comes from the issue of accumulation already alluded to in section 2.2. If one wants to obtain a SDR close to investors' private time preferences as revealed by financial markets, one may find it expedient to incorporate a strong present bias in the discounted utilitarian SWF. This choice is questionable for two reasons. First, practically, it is not obvious which market interest rate is relevant for the IAM computation. The stock market rate of 5–6% incorporates a risk premium that should not be confused with a present bias in investors' time preferences. The rate on sure assets such as Treasury bonds is much lower, and would therefore justify a SDR that would be closer to Stern's than to Nordhaus's (Gollier 2013).

The second reason why this is questionable is ethical. The market rate on any asset reflects the preferences of investors whose calculations may include some altruism for future generations. Yet, such private preferences have no reason in general to be up to the ethical standards of a social welfare evaluation. Investors are expected to be selfish in their decisions, thinking more of their own future and their children than worrying about their great-great-grandchildren. It is very dubious that social policy should embrace the same ethical and temporal myopia when assessing trade-offs across distant generations. One can accept that social discounting should indeed reflect the population's time preferences in the case of short-term investments in which the payers and beneficiaries are the same people, so public policy is merely acting on behalf of private individuals who cannot obtain these investment services (say, in infrastructure) on private markets. But one can simultaneously argue that for

trade-offs across generations, impartiality across individuals becomes imperative—and excludes a pure time discount other than the level derived from considering the risk of extinction.¹⁴

Finally, one interesting complication is that if one accepts extinction risk as the justification of a small present bias in SWF, then one needs to face the question of how to evaluate risks in a more comprehensive way. Extinction risk is only one of many relevant risks. It is awkward to make a computation in a deterministic (non-risky) framework while justifying certain parameters by reference to risk.

B. Growth effect. The growth effect refers to the fact that the SWF may assign less priority to future dollars than present dollars, hence a greater SDR, if future generations are richer than the present one. Indeed, along a growth path, a mitigation policy which transfers resources from the present (mitigation costs) to the future (benefits of a preserved climate) looks regressive. The ethical feature of the SWF which governs this effect is its degree of “inequality aversion,” a technical term referring to how socially undesirable it is to transfer money from a poor to a rich population.¹⁵ For the discounted utilitarian SWF, the degree of inequality aversion depends on how concave the utility function u is, i.e., how quickly marginal utility diminishes when consumption increases.

The growth and the inequality aversion considerations reinforce one another, because when inequality aversion is greater, the impact of a given increase in the growth rate of consumption on the SDR is magnified. Relying on the discounted utilitarian SWF, economists have derived a simple formula—known as the Ramsey formula—approximating the value of the SDR as follows:

$$\text{SDR} = \text{time discount rate} + \text{growth rate} \times \text{coefficient of inequality aversion}$$

In this formula, the growth effect discussed here is captured by the product of the last two terms. The time discount rate is the ρ parameter of the discounted utilitarian SWF. The growth rate is the average growth rate of consumption between the present period and the contemplated future period.

The Ramsey formula is widely used. It offers a convenient way to understand the basic components of the relative social value of dollars between different time periods. It also helps in grasping how different authors reason. For instance, what has been called the ethical, or “prescriptive” approach employed by the Stern Review takes the pure time discount rate to be determined by the hazard rate of extinction, and considers the coefficient of inequality aversion to be primarily an ethical parameter reflecting how one should trade off consumption between rich and poor.¹⁶ The “descriptive” approach preferred by Nordhaus takes the SDR from the market interest rate; this approach adopts a coefficient of inequality aversion reflecting an empirical estimate of the curvature of utility functions in the population, either from risk aversion or from aversion to time fluctuations of consumption, and deduces the value of the discount term that closes the Ramsey equation.

The “prescriptive-descriptive” terminology, due to Arrow et al. (1996), is commonly used but misleading because even the descriptive approach involves the

ethical choice to adopt the discounted utilitarian SWF and also the ethical choice to adopt for policy-making purposes the time preferences of the investors as observed in their market decisions—a choice that is questionable, as explained earlier, for trade-offs between distant generations (see Kelleher [2017] for further discussion of the prescriptive vs. descriptive debate over discounting).

C. Adjusting intertemporal inequality for intratemporal inequality. The Ramsey formula is usually applied in frameworks in which inequalities within generations are ignored and the analysis is done as if there were perfect equality in every generation.¹⁷ However, not everyone in the future is expected to be equally rich, and thus richer than the current population. For example, while an individual living in a developing country in 50 years is likely to be richer than his compatriots living today, it is questionable to assume, for example, that all future South African people will all be richer than current US people. Knowing that not all future individuals will be wealthier than the present ones might warrant a greater effort in their favor.

This “inequality effect” depends on an empirical prediction (the degree of inequality in the future world) and an ethical value (the degree of inequality aversion, already introduced in the previous paragraph).¹⁸ The more unequal the future economy is expected to be, the larger the motive to provide resources (or a good climate) to future generations.

To obtain a quantitative estimate of this inequality effect, one can perform the following simple thought experiment with the help of the Ramsey formula. The Ramsey formula applies to a framework in which every generation is perfectly equal in the distribution of consumption. So, let us ask what level of consumption, if equally given to everyone in the present generation, would produce the same social welfare as the current, unequal distribution. This hypothetical level is commonly called the “equally-distributed equivalent” (EDE). It is useful to know that most commonly used inequality indices can be read as providing the EDE in the following way: if the index is 35%, this means that the EDE is 100 minus the index, i.e., 65%, of the income per capita. One can do the same for any future generation and compute its EDE level of consumption. In the Ramsey formula, one can simply replace the average growth rate of consumption per capita, between the two periods, by the average growth rate of their EDE (between the two periods). If the inequality index goes up by 5 percentage points in the future, for instance, this reduces the future EDE by 5% of per capita income, and decreases the SDR by this amount times the coefficient of inequality aversion and divided by, roughly, the time distance between the two periods.

So far, we have assumed that there was only one coefficient of inequality aversion, which embodied an attitude (which could be utilitarian, prioritarian, or egalitarian in origin) to inequality in consumption across and within generations. But Schelling (1995) famously suggested that distance in time and distance in space might justify different degrees of receding priority. Do people worry more about a poor child alive today in a remote country than about their grand-grand-grand-child, or the opposite? From a philosophical standpoint, it seems hard to make a difference between individuals on the basis of where and when they live, but the debate about global justice

and cosmopolitanism has been the occasion to hear voices justifying some bias against aliens in defining social priorities (Rawls 1999; Nagel 2005). Arguably, however, although the existence of certain local and national institutions may justify different types and levels of action due to the unequal effectiveness of action in different institutional settings, and while relations between individuals may also have normative implications, the pure ethical evaluation of individuals' situations should be impervious to their mere location and time, just as ethical evaluation disregards race, gender, sexual orientation, and social background. The existence of local and national bonds and institutions is itself endogenous to human decisions and cannot be taken as a basis for ethical assessment in evaluations where it is itself an object of assessment and possible reform (Beitz 1999).¹⁹

That being said, there is no formal difficulty in transforming the SWF to incorporate a different degree of inequality aversion within and across generations. This requires abandoning the simple additive form of utilitarianism (discounted or not), computing a separate value of social welfare for each generation, and then computing intergenerational social welfare as an aggregate of these generational social welfare values. The generational social-welfare function would have its inequality aversion (applied to inequalities within each generation), and the intergenerational aggregator would have a different inequality aversion (Dennig 2009; Anthoff and Emmerling 2016). Such a transformation would modify either the intergenerational component of the SDR, or the intragenerational inequality component, depending on which coefficient of inequality aversion is greater.

D. Intratemporal inequality in the costs and benefits of mitigation policy. So far we have been working under the assumption that adding a ton of CO₂ today will produce climate damages but will have no large impacts on the underlying economy (in particular, not on its distribution of income). However, climate damages may affect not only the *level* of consumption but also the *distribution*.

The empirical relationship between damages from climate change and income is hard to predict. In particular, would climate damages increase or decrease with wealth? The SDR depends on the correlation between wealth and damages. Because rich countries and individuals have more assets at risk, we might expect wealthy individuals (and individuals living in rich countries) to be more exposed to climate damages simply because they have more to lose. On the other hand, rich individuals can afford investments in adaptation and in risk-sharing that will reduce the size of the damages they bear. In contrast, because poor individuals have lower resources, their ability to adapt is reduced, which would lead us to expect a negative relationship between damages and income. Recent work seems to support the view that climate impacts will disproportionately hurt the future poor, even if there is currently no good estimate of the disproportion (World Bank 2016; Hsiang et al. 2017). If the poor will bear the brunt of climate damages, then climate change is expected to increase the degree of inequality in the future population. As a consequence, introducing a mitigation policy would not only prevent economic and noneconomic losses, but it would also help to reduce the level of inequality. Once again, the magnitude of this effect depends on the degree of intragenerational inequality aversion in the SWF.

Similar considerations apply also to the distribution of costs of climate policy among members of the present generation, because this affects inequality in living standards now. A carbon tax or the elimination of fuel subsidies (which are common in poor countries as part of social policy) may be regressive if the mitigation policy raises energy costs and makes energy less accessible to the present poor. Redistributing the proceeds of the tax to the population, in cash or in public goods and services, may, in principle, counter this problem and render the policy progressive within the present generation.

Therefore, quite generally, a distribution-sensitive computation of the SDR depends on the distribution of the cost among members of the present generation and the distribution of the benefits among members of the future generations. A policy that makes the rich today pay for the benefit of the future poor may appear to be a progressive transfer if the future poor remain below the level of the current rich. In this case, such a policy would be evaluated with a negative SDR, because the future dollar for the poor has greater social value than the present dollar for the rich. In contrast, a policy that would cost the present poor for the benefit of the future rich would have a large SDR because it would amount to a starkly regressive transfer. Applying this methodology, [Dennig et al. \(2015\)](#) and [Budolfson et al. \(2017\)](#) show that, depending on the distribution of present costs and future benefits, the optimal carbon tax can go from very low to very high levels, dwarfing the debate between Stern and Nordhaus on discounting, in which this intragenerational inequity issue was ignored altogether.

3.2 Evaluation of damages

In order to compute the SCC, we need to be able to monetize the physical impacts of a changing climate. Clearly, some damages can be easily turned into monetary terms. For example, losses in the agricultural sector can be valued by using the market price of agricultural products. Some categories of damages, instead, are more difficult to evaluate. Think, for instance, of impacts on human health and mortality risk, or the extinction of species: what should be the price of preserving a human or non-human life or ecosystem?

In IAMs, nonmarketable goods are monetized and incorporated into the consumption figures. For instance, the value-of-statistical-life is used to value improvements in health and housing prices are relied upon to determine the value of environmental services. However, these measures have two important drawbacks: first of all, they are based on individuals' willingness to pay for improvements in nonmarketable goods. In other words, an anthropocentric perspective is taken, where goods with no market value matter only if they matter to individuals.²⁰ Second, an "average" willingness to pay is usually considered, which ignores the problem of how to measure and compare individual well-being when preferences are heterogeneous and income is unequally distributed.

Let us illustrate the latter issue with the example of the value of a life. Suppose that one has estimates of every individual's willingness-to-pay (WTP) to extend life by one year. This WTP depends on the individuals' wealth as well as on their preferences over longevity. It may make sense to give less importance to the life of

individuals who genuinely care less about living long, but it seems repugnant to discount the lives of the poor; yet, unfortunately, both determinants are entangled into the WTP. An SWF, fortunately, offers an elegant solution. It may give a greater priority to the plight of the disadvantaged population through its inequality aversion feature. Therefore, when evaluating the protection of a poor person's longevity, the SWF will weight the WTP that provides the dollar equivalent of longevity for that person with the "marginal social value of money" for that person, which will be greater the worse-off the person is.

Usual methods, however, do not involve marginal-social-value-of-money weighting, and therefore fail to appreciate the importance of giving special attention to the worse-off. But relying on the average WTP rather than the actual distribution of WTP goes some way toward alleviating the repugnance of an unweighted use of WTP. The average WTP being greater than a poor person's WTP, valuing a poor person's life at this greater value is equivalent to applying a special weight to this person's actual WTP. But this is obviously less than rigorous.

One further difficulty for the SCC is to determine the scope of damages induced by a ton of CO₂ emitted. The standard approach, on which we have focused here, is restricted to the *climate change* effects of the emission. This is sometimes referred to as the social cost of "carbon dioxide." Other effects of the emission include *health* effects due to air pollution by particulates. One could also consider visual effects on blurred landscape views, for example. This broader range of impacts would then be named the social cost of "carbon" proper. In principle, one should take account of all impacts but this is a daunting task, though health cobenefits of mitigation have attracted a lot of interest in policy debates (IPCC 2014, chapter 6) and are quantified in Scovronick et al. (2017b).

Section 3 has shown that the usual approach to the SCC does not deal with ethical issues in a satisfactory way when time discounting is adopted without an ethically justified reason and when inequalities are neglected, both in the determination of the social discount rate and in the evaluation of damages. However, the SCC methodology would be compatible with limiting time discounting to a reasonable range determined by the risk of extinction, and with giving due attention to inequalities both between and within generations.

4. RISK

So far, we have considered deterministic consequences of policies, without any risk. However, climate change is an inherently risky phenomenon, as the size and distribution of its impacts are highly uncertain. Moreover, the long-term future is deeply uncertain, because various threats and opportunities may unfold in many ways. Beyond the most immediate future, the uncertainty is enormous, especially as regards technological innovation and the climate feedbacks in response to increasingly large departures from previously known states of the Earth.

This section's question, then, is how the introduction of risk will affect the SCC. We will first retain the utilitarian framework adopted in the bulk of this paper, then we will discuss alternative ethical approaches.

4.1 Risk in the standard utilitarian SWF

The social discount rate is the main channel through which risk will have an impact on the SCC. What is convenient is that, formally, risk and inequalities play out in a similar way in the analysis. While inequality introduces heterogeneity along the dimension of population, risk introduces heterogeneity along the dimension of “states of the world.” The concepts and methods built to deal with one can be used to tackle the other.

More specifically, in parallel with the presence of inequality, risk has two main effects:

A. **Precautionary effect.** This effect is the counterpart of the inequality effect of the previous section. If future economic conditions are uncertain, one should favor a transfer from the present to the future so as to protect the living conditions of the future individuals who potentially end up with a bad outcome. Just as the inequality effect could be quantified through the replacement of average per capita consumption by the EDE within a time period, the precautionary effect suggests replacing the expected value of consumption (per capita or EDE) by its certainty equivalent (CE), which is defined as the riskless level of consumption that provides the same social welfare as the risky prospect under consideration.

The effect is similar to the inequality effect that was previously defined, but, from an ethical point of view, it depends on risk attitudes, rather than inequality aversion. In economists’ standard approach to incorporating risk into utilitarianism—a method historically influenced by [Harsanyi \(1955\)](#)—the social objective is the expected value of the sum of utilities, or equivalently the sum of individual expected utilities. As a result, the individual utilities then embody risk aversion, implying that risk aversion and inequality aversion coincide, because the coefficient of inequality aversion and the coefficient of risk aversion both depend on the curvature of the same utility function. Therefore, aversion to heterogeneity in the distribution of income across generations (or more generally, people) is valued the same as heterogeneity in the distribution of income across different states of the world. This is a serious restriction that motivates departures from utilitarianism discussed at the end of this section.

B. **Insurance effect.** The insurance effect is the effect on the SCC of the correlation between income and climate damages, across possible states of the world. In particular, if damages are expected to be larger in states of the world characterized by low income, then one has more reason to invest for the future, and the SCC will increase. Like the effect related to the intratemporal distribution of costs and benefits from mitigation policy (subsection 3.1d), what matters here is the relationship between income and damages, now across states of the world.

Indeed, if damages and income are negatively related, climate mitigation has an insurance effect, i.e., introducing a climate policy will reduce the risk that society is facing. As with inequality in climate impacts, the sign of the relationship between damages and aggregate income risk remains an open question. For example, [Nordhaus \(2017\)](#) thinks that the correlation is positive, because a richer economy will emit more and have more infrastructures exposed to extreme weather, which will

tend to raise the SDR and, consequently, reduce the SCC. Others, instead, favor a negative relationship and a large SCC, arguing that economic and environmental systems are jointly determined, so that larger climate impacts arise in situations where economic growth is low (Sandmark and Vennemo, 2007).²¹

Weitzman (2009) observes that if there is a small probability of a catastrophic form of climate change that impoverishes humanity and puts it on the brink of starvation, and if the marginal social value of consumption in such a bad state is extremely high, then the possibility of catastrophic risk may swamp the whole cost-benefit analysis. If so, the optimal policy for the present generation would depend substantially on the exact value of consumption in the disaster scenario. Weitzman's analysis is combined with a statistical analysis of fat tails, but the key point is really this simple observation about the overwhelming weight of extreme values, even if they have small probabilities (Millner 2013). It is worth emphasizing that with sufficient inequality aversion and sufficient inequalities within generations, this point can be extended so that if there is a small probability that even a small group of the population will suffer terribly from climate impacts, this may loom large in the cost-benefit computation (see Fleurbaey and Zuber, 2015a).²²

4.2 Risk in nonutilitarian SWFs

So far, we have considered the consequences of risk for a utilitarian SWF, by which we mean Harsanyi's sum of expected utilities. One could argue that in the absence of risk, assuming that all relevant dimensions of well-being are suitably encapsulated in a consumption variable and that one can choose the utility function freely, an additive SWF can accommodate many different ethical views, including prioritarianism and egalitarianism. Some philosophers reject the idea that an additively separable function can represent an egalitarian approach, which they perceive as intrinsically nonseparable, but this is debated.²³

Once risk is introduced, Harsanyi's utilitarianism limits the analysis because, as already noted, it conflates risk aversion and inequality aversion. This is a problem because many philosophers and economists argue that risk aversion in the SWF should primarily reflect the population's risk attitudes. Some prioritarians (e.g., Rabinowicz 2002), however, do consider that an extra moral risk aversion is advisable, even though such a view imposes a paternalistic precautionary attitude on the population. However risk aversion is parameterized, it is hard to see why the ethical choice of *inequality* aversion should coincide with a population's *risk* attitude, if only because risk attitudes are heterogeneous in the population whereas a single coefficient of inequality aversion is needed in the SWF. The challenge, then, is to combine a sound evaluation of inequalities in the SWF with a respectful incorporation of individual risk attitudes.

The literature has explored two escapes from utilitarianism. By Harsanyi's (1955) theorem, the utilitarian SWF is the only one that respects individual risk attitudes when both the individual and the social situations are assessed in terms of expected utility (expected social welfare, more accurately, for the latter). Trying to escape or generalize utilitarianism thus requires either respecting risk attitudes in a weaker fashion, or abandoning the expected utility approach (at the social level at least).

The former option is the so-called ex post approach, which, in a prominent example, computes the EDE of consumption in every state of the world, and then computes the (possibly weighted) sum of the individuals' expected utilities applied not to their personal consumptions but to the risky EDE. To fix ideas, the mathematical form of this SWF is as follows:

$$\sum_i Eu_i(\text{EDE}),$$

where Eu_i denotes the expected utility of individual i . The ex post label comes from the fact that this approach is sensitive to ex post inequalities, and proceeds by first assessing the social situation in every final consequence before aggregating over states of the world. It does not respect individual risk attitudes in general, but does so for allocations which are egalitarian in every state of the world—because in this case, the EDE is the allocation itself. Fleurbaey and Zuber (2017a) provide an axiomatic justification for this approach, and propose an argument for an extreme variant of this criterion, which equates the EDE with the lowest consumption in the distribution (as with the maximin criterion) and then applies the most risk-averse of the individuals' utility functions to it.

Note that, if one disregards time as a variable justifying a differential treatment of individuals, the EDE should be computed not separately for every generation, but for the whole population of all generations under consideration. This extensive computation generates a strong nonseparability across generations. In particular, the consumption level of all our ancestors then affects the level of the EDE, and therefore influences the assessment of future prospects even if the fate of our ancestors cannot be affected by our actions. This nonseparability raises practical questions about the scope of the criterion. Should one include all past generations, whose mere existence may have an impact on the computation? What about future generations beyond the typical horizon of IAMs? We conjecture that the nonseparability is not so burdening, and simply boils down to adding specific weights which modify the probabilities of states of the world to take into account whether these states are overall good or bad for the whole population under consideration.²⁴

The other option is the ex ante approach, which consists in first computing the expected utility for each individual, and then incorporating the distribution of these expected utilities into an SWF. One particularly attractive variant, in the current context, takes the certainty-equivalent (CE) for every individual and incorporates the distribution of individual CEs into the SWF. This approach is popular in economics (see, e.g., Epstein and Segal [1992]), perhaps because it satisfies the Pareto principle and thereby fully respects individual risk attitudes, but it does not take the form of an expected value of social welfare. Inequality aversion over expected utilities or over individual CEs implies that this approach will often condone risk-taking activities that will produce final consequences that inequality aversion of the SWF dislikes once social welfare is recomputed later in time, after risk is resolved. This violates a key rationality condition named dominance, which stipulates that if one prefers an

option in all states of the world, one should still prefer it when the state of the world is unknown.

One complication in the application of these approaches to climate policy is that in the far future, different states of the world are likely to be inhabited by different people, raising a particular variant of the nonidentity problem. These individuals' risk attitudes may not be relevant to assessing the risk to society, because they do not face the whole risk. In the extreme, imagine that every state of the world has totally different populations. In that case, individuals in the future do not actually face any personal risk. Their individual risk attitudes, then, would not constrain social evaluation in this case, and one can appeal to prudential or moral considerations to choose a social attitude to risk.

The study of the implications of these *ex post* and *ex ante* criteria for climate policy is still largely an open field for research. Ferranna (2015) examines the computation of the SDR under such criteria, and shows that the SDR can be modified in any direction depending on the nature of the risks and the relative values of the coefficients of risk aversion and inequality aversion.²⁵

Finally, it should be mentioned that uncertainty surrounding climate change is sufficiently deep that criteria relying on probabilistic calculus require forming probabilistic beliefs out of very sparse information about the likelihood of various phenomena. Heal and Millner (2014) have argued that one should refrain from doing so, on the ground that uncertainty about probabilities should be acknowledged in a more thorough way in the cost-benefit analysis. They recommend adopting “ambiguity-averse” approaches, which focus on the less favorable probabilities, among those which are considered plausible, when evaluating policy options. To illustrate the notion of ambiguity aversion, suppose that a mitigation policy A may be successful with probability ranging from 0.4 to 0.9, whereas an adaptation policy B may be successful with probability ranging from 0.5 to 0.8. Then one should weight the successful outcome of policy A with probability close to 0.4, and retain a value close to 0.5 for policy B. Ambiguity aversion recommends doing this pessimistic assessment even if there is little plausibility in the idea that probabilities could be so low simultaneously for both policies. This “inconsistent” pessimism induced by ambiguity aversion has generated debates about whether ambiguity aversion was irrational, unlike expected utility. Indeed, several authors²⁶ have shown that ambiguity aversion can induce violations of the dominance principle.

In conclusion, debates continue about how to deal with the risk associated with climate change and climate policy. We have argued that Harsanyi's utilitarianism is not flexible enough to separate considerations of risk aversion and inequality aversion in the evaluation, and that, among the more flexible approaches, *ex post* probabilistic criteria which avoid the irrational decision patterns of *ex ante*, recursive, and ambiguity-averse approaches appear more attractive, although their nonseparability raises some concerns.

5. POPULATION

Estimating a social cost of carbon, we have seen, requires a method of valuing social outcomes in which benefits and harms are distributed to different groups of people.

The possible groups of people affected by climate change will live at different times, experience different health, and have different levels of wealth. These groups could also be of different sizes. Further, climate change and climate policy could importantly influence the group of people who will live and its size, by influencing patterns of fertility and mortality. Therefore, estimating a social cost of carbon requires a method to incorporate population sizes into social valuations (Broome 2012). Population ethics is the search for such a method.

5.1 Endogenous population size and the SCC

Population ethics is known for its difficult puzzles: no single approach has emerged that is consistent with all attractive intuitions (Parfit 1984; Blackorby et al. 2005; Broome 2004; Arrhenius forthcoming). A broad division among approaches to population ethics separates those that regard a larger population living good lives as a social improvement, all else equal, from those that do not value such an increase in population size. The classic example of the former is Total Utilitarianism (TU), which seeks to maximize the sum of utility over the whole population. TU would socially prefer a sufficiently large increase in the size of the population, even at a cost to average well-being. The latter approach is exemplified by Average Utilitarianism (AU), which disregards population size and seeks to maximize average utility within the population. AU would reject an opportunity to substantially increase the size of a very well-off population, if it came at the cost of even a tiny decrease in average well-being. In empirical surveys of social choice, both views receive support (Spears 2017). Beyond these, a wide range of population-sensitive social objectives have been introduced and characterized in the literature (Ng 1989; Blackorby et al. 2005; Asheim and Zuber 2014; Fleurbaey and Zuber 2015b).

Climate change may influence the size of the future human population. For example, climate change may cause a large disaster that radically reduces the size of the future population by changing the habitability of the planet for humans (Sherwood and Huber 2010). Alternatively, climate change may have more moderate effects on mortality and fertility rates that change the number of people born. If a larger global temperature change is likely to cause a larger reduction in the size of the future population, then that reduction would increase the SCC under a TU-type social objective, but not under an AU-type social objective if there were no further consequences, such as for average well-being.

If the risk of extinction or of a population-size disaster is large enough, then how to value the endogenous effect of climate change on population size could be a first-order issue for the SCC. In the case of an extinction risk, increasing carbon emissions does not only change future consumption, but it may also change the probability of a future extinction. Weitzman (2009) discussed how this could be analyzed in terms of the “value of a statistical civilization,” which is like the value of a statistical life (VSL). Bommier et al. (2012) and Méjean et al. (2017) show how the SCC methodology can be adapted to this issue. In particular, Méjean et al. (2017) show that the social valuation of population size is a critical determinant of optimal policy. Even with an inequality-averse SWF, one may prefer to increase intergenerational inequality by

reducing the consumption of near-term generations, if this mitigation effort reduces the probability of future extinction.

5.2 Exogenous population size and the SCC

Different theoretical approaches to population ethics offer different methods of aggregating impacts across groups of people of different size, living at different times (Dasgupta 2001). But the simplest forms of discounted or nondiscounted AU and TU do not differ in how they estimate the SCC, and therefore in how the SCC reacts to an *exogenous* change in the projected size of the future population.²⁷ This is because the AU function is equal to the TU function divided by a fixed term (the total population). Scovronick et al. (2017a) show that in Nordhaus's DICE model, if the size of the population is exogenously predicted to be greater, then the SCC rises by so much that the optimal carbon tax policy induces lower emissions and a lower temperature profile over the future.²⁸

However, other variants of utilitarianism may behave differently. For instance, Scovronick et al. (2017a) examine a discounted form of within-period AU that seeks to maximize the time-discounted sum, not weighted by population, of average utility within each time period.²⁹ Introducing this in the DICE model, they show that a larger projected future population increases the SCC by more under the usual discounted version of TU than under their version of AU, because these social objectives differ in whether or not future harms are weighted by population size. The consequences for the SCC are quantitatively large: they are of comparable size and importance to the time discount term.

5.3 Nonhuman species

Another set of issues related to population ethics is raised by biodiversity and the interests, welfare, and rights of nonhuman species.³⁰ Climate change is an important stressor for biodiversity due to the rapid possible pace of change. Change may overwhelm species that are slow to move or adapt.

Existing estimates of the SCC poorly incorporate the impact of emissions (and temperature increase) on nonhuman species. Some IAMs include estimates of the ecosystem services or amenities produced by nonhuman species, but they focus on services and amenities with direct economic value. As highlighted by Tol (2009), these estimates are very rough and incomplete, and probably underestimate the true costs. More fundamentally, existing work only focuses on the instrumental value for humans of nonhuman species (and more generally the environment) and completely ignores any possible intrinsic value (or even related indirect values such as cultural values that may grant a specific status to some species). In that sense, existing estimates can be viewed as conservative estimates that may understate the real gain of emission reductions.

The literature on animal ethics and climate change is developing quickly (e.g., Hsiung and Sunstein 2007; McShane forthcoming; Sebo forthcoming; Budolfson and Spears 2018). One of the key difficulties that singles out this subfield from the human-centered issues concerns cross-species comparisons. There seems to be no

way to assess the trade-offs between human and nonhuman interests without weighing the relative degree of priority of the different populations. This is known to be difficult for human beings with different plans in life, values, and preferences, but it is much harder for different species with different needs and different capabilities.

The utilitarian approach may seem to offer a practical recipe by focusing on well-being, particularly for a hedonist conception of well-being, because it may seem that sensations and emotions can be gauged independently of the type of being who experiences them. But hedonism is hardly a consensual approach, as it appears to take a very narrow perspective of a life's achievements, and psychologists emphasize that human emotions are hardly separable from cognitive attitudes (Kahneman et al. 1999). Even if one accepted hedonism, though, it is not at all obvious how to compare emotions across differently structured brains, or neural systems more generally.

6. CONCLUSION

This paper has introduced the reader to the logic of estimation of the SCC in the most common approaches. We have tried to highlight some of the key ethical assumptions of these approaches, both in terms of how future costs and benefits are being compared to current costs and benefits (discounting) and in terms of what is included in future impacts. The Ramsey formula, in particular, provides very useful guidance to the key components of the discounting of future dollars.

But this paper has also made a call for going beyond the most traditional approach, which focuses on the intergenerational trade-off and generally neglects three sets of considerations that deserve much greater attention than they currently receive: inequality, risk, and population size. Ethical issues and choices about each of these sets of considerations may ultimately have larger quantitative consequences for climate policy than the typical debates over the discount rate. In a nutshell, climate policy, and the determination of the SCC, is not merely about investing resources now for the benefit of future richer generations. It is also about altering present and future inequalities, protecting future populations against uncertain disasters jeopardizing their livelihood, and considering the consequences of population growth.

It has been argued here that the framework offered by the SCC methodology and the related social-welfare functions used in integrated assessment models is flexible enough to accommodate many ethical options which are foreclosed in the dominant discounted utilitarian approach in climate economics. It is possible to take account of inequalities between and within generations, to discount the future only insofar as its existence is uncertain, to disentangle risk aversion and inequality aversion, and to give population size a value ranging from the total to the average approach, with intermediate options in between.

However, while the question of inequalities is perhaps rather well-trodden in theory (in particular given the fact that interpersonal comparisons are not complex in the stylized settings of climate-economy models), the assessment of risks at the social level, and the evaluation of population effects, including over nonhuman species, are still full of open questions and disagreements among specialists. Devoting more philosophical attention to these ethical conundrums receives a renewed urgency from the applications waiting in climate policy analysis.³¹

NOTES

1. See, e.g., [Greenstone et al. \(2013\)](#) for the use of the SCC in U.S. federal regulations, and [Pearce \(2003\)](#) for the U.K. case.
2. All the SCC values in this paper refer to tons (i.e., 1,000 kg) of CO₂. [Nordhaus \(2017\)](#) revised his estimate upwards to \$31 per ton of CO₂ by including a more refined description of the economy.
3. The range of SCC values reported in the literature is quite dispersed, depending both on ethical assumptions and the description of the economy and the climate impacts. See for example [Van den Bergh and Botzen \(2014\)](#) for a meta-analysis of the available estimates.
4. One can now find the old EPA webpages on the SCC at https://19january2017snapshot.epa.gov/climatechange/social-cost-carbon_.html.
5. See, e.g., [Bowen \(2011\)](#). Differences in taxation may also produce “leakages” if agents who are taxed more heavily transfer their polluting activities to jurisdictions with lower taxation.
6. It is true that, starting from a situation with unequal marginal costs of abatement, introducing a market for emission permits would enable the agents to make beneficial transactions, potentially improving the situation for everyone (though this is not guaranteed when other markets are affected, altering commodity prices and wages), while the marginal abatement costs would be equalized to the equilibrium common price of permits. This does not contradict the previous point—what happens is that the market for permits implements, via the payments from buyers to sellers, transfers of money that improve the distribution.
7. This is closely connected to the large debate about the allocation of emission rights and the desirability of a global market for permits ([Pottier et al. 2017](#)).
8. Prominent IAMs for the computation of the SCC include RICE (Nordhaus 2010) and FUND (Anthoff and Tol 2013). This paper often refers to results obtained with a variant of RICE called NICE ([Budolfson et al. 2017](#)), which incorporates inequalities within regions of the world.
9. This means in particular that prioritarian approaches which can be formulated with this same mathematical formula are also covered in our discussion, at least when there is no risk (more will be said about this in the section about risk). Notice that the framework is also in principle flexible as for how individual “well-being” is defined and could embody various approaches, including objective good and capability approaches, different preference satisfaction approaches (including equivalent income) or truly welfarist approaches.
10. If marginal utilities are weighted by their inverse at a certain allocation, the weighted marginal utilities are automatically equal in the allocation that gives the inverses their value, implying that this distribution is optimal.
11. However, this does not justify the extra endogenous adjustment of optimal savings to climate impacts and policy that is also made in some models ([Rezai, Foley, and Taylor 2012](#)).
12. Mortality impacts that shorten people’s lives are currently accounted for as losses of life years valued as a certain proportion of consumption in the period in which these life years are lost.
13. His approach avoids time bias by computing a sum of the utility gaps to a bliss point. If the utility of future individuals converges to this bliss point sufficiently quickly, then the utility gaps converge to zero quickly enough that their sum can be finite. This elegant solution, however, depends on many assumptions, including a meaningful bliss point and quick convergence toward it.
14. See [Fleurbaey and Zuber \(2015a\)](#) for a computation of an SDR which combines respect for individual intertemporal preferences and impartiality between individuals.
15. It is important to note that the expression “inequality aversion” does not imply that the philosophy of the analysis is egalitarian, as opposed to utilitarian or prioritarian. Instead, inequality aversion quantifies the attitude of the SWF toward concrete trade-offs between individual consumptions at different levels. As already explained in n. 9, the “discounted utilitarian” SWF that pervades the literature can also serve as a formula for prioritarianism.
16. However, the view adopted in the Stern Review is that this coefficient should reflect diminishing marginal utility due to individuals’ risk aversion rather than pure distributional considerations.
17. A part of the literature on the valuation of climate impacts has taken intragenerational inequalities into account, introducing distributional weights that give more priority to the worse-off (in particular [Fankhauser et al. \[1997\]](#); [Tol \[2001\]](#); [Hope \[2008\]](#); [Anthoff et al. \[2009\]](#)).
18. [Caney \(2009, 172\)](#) makes a similar point.

19. In a way that is compatible with our earlier discussion of present bias, Heath (2017) notes that one pertinent difference between spatially and temporally distant individuals is that spatially distant individuals exist whereas future individuals are only possible.
20. A related issue is to which extent individual preferences can be used to make normative evaluations since individuals may be subject to cognitive errors or egoistic reasoning. See Hammitt (2013) for a discussion of positive and normative conceptions of cost-benefit analysis.
21. See also Dietz et al. (2015) for a careful analysis of the channels through which positive or negative correlation may arise, as well as a quantitative analysis using the DICE model.
22. Catastrophes might also be triggered by mitigation actions, e.g., by failed geoengineering.
23. See, e.g., Temkin (1993), Broome (2015), Fleurbaey (2015).
24. See Fleurbaey (forthcoming) and Ferranna (2018) for an elaboration of this point.
25. Another part of the literature (see, e.g., Warren [2014]; Traeger [2015]; Cai [2015]; Lontzek et al. [2015]; Daniel et al. [2016]; and Lemoine and Traeger [2016]) has been attracted by the Epstein-Zin (1989) criterion, which belongs neither to the ex post nor to the ex ante approach, and in which expected social welfare is computed recursively, as follows. Social welfare at a given date is the expected value of a function which aggregates the utility of the generation living in that period and the CE of the expected value of the same function computed at the next period. This recursive form makes it also possible to separate inequality aversion between a particular date and all following dates (as captured by the aggregation function), from risk aversion used to compute the CE. Like the ex ante approach, it can generate violations of dominance because it does not compute expected social welfare on final consequences but on intermediate evaluations.
26. See in particular Seidenfeld (1988); Al-Najjar and Weinstein (2009); Steele (2010).
27. When a time discount term appears in the SWF, it seems that AU should similarly discount population numbers at the denominator. Otherwise, social welfare could be increased simply by making people live at earlier dates, along a path of constant individual utility. This SWF is introduced by Dasgupta (2001).
28. Different population projections can also influence the SCC if the composition of the future population changes, such that more people are projected to live in climate-vulnerable regions. Budolfson et al. (2018) consider such a case of an actual revision of population projections that increased the number of people expected to live in sub-Saharan Africa.
29. For a discussion of this SWF and its sources, see Dasgupta (2001).
30. Population ethics bears on existential issues for individual animals, not only species. There are welfare issues for animals which are not existential, and for convenience this brief section covers them as well.
31. This paper has benefited from very helpful comments by J. Broome and two referees. Kian Mintz-Woo wishes to acknowledge funding from the Austrian Science Fund (FWF) under research grant W 1256-G15 (Doctoral Programme Climate Change—Uncertainties, Thresholds and Coping Strategies).

REFERENCES

- Ackerman, F., R. Bueno, S. Kartha and E. Kemp-Benedict 2011. "Comparing Climate Strategies: Economic Optimization Versus Equitable Burden-Sharing," Stockholm Environment Institute Working Paper WP-US-1104.
- Adler, M., D. Anthoff, V. Bosetti, G. Garner, K. Keller and N. Treich 2017. "Priority for the Worse Off and the Social Cost of Carbon," *Nature Climate Change* 7: 443–49.
- Al-Najjar, N.I. and J. Weinstein 2009. "The Ambiguity Aversion Literature: A Critical Assessment," *Economics and Philosophy* 25: 249–84.
- Anthoff, D. 2009. "Optimal Global Dynamic Carbon Taxation," WP278, ESRI.
- Anthoff, D. and J. Emmerling 2016. "Inequality and the Social Cost of Carbon," FEEM working paper.
- Anthoff, D., C. Hepburn and R.S.J. Tol 2009. "Equity Weighting and the Marginal Damage Costs of Climate Change," *Ecological Economics* 68: 836–49.
- Arrhenius, G. forthcoming. *Population Ethics*, Oxford: Oxford University Press.
- Arrow, K.J. 1999. "Discounting, Morality and Gaming," in P.R. Portney and J.P. Weyant, eds., *Discounting and Intergenerational Equity*, New York: Resources for the Future.

- Arrow, K.J., W.R. Cline, K.-G. Mäler, M. Munasinghe, R. Squitieri and J.E. Stiglitz 1996. "Intertemporal Equity, Discounting, and Economic Efficiency," in J.P. Bruce, H. Lee and E.F. Haites, eds., *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: Cambridge University Press.
- Asheim, G.B. and S. Zuber 2014. "Escaping the Repugnant Conclusion: Rank-Discounted Utilitarianism with Variable Population," *Theoretical Economics* 9: 629–50.
- Beitz, C. 1999. "International Liberalism and Distributive Justice: A Survey of Recent Thought," *World Politics* 51: 269–96.
- Blackorby, C., W. Bossert, and D. Donaldson 2005. *Population Issues in Social Choice Theory, Welfare Economics, and Ethics*, Cambridge: Cambridge University Press.
- Bommier, A., B. Lanz, and S. Zuber 2015. "Models-As-Usual for Unusual Risks? On the Value of Catastrophic Climate Change," *Journal of Environmental Economics and Management* 74: 1–22.
- Bowen, A. 2011. "The Case for Carbon Pricing," *Policy Brief*, Grantham Research Institute on Climate Change and the Environment.
- Broome, J. 2004. *Weighing Lives*, Oxford: Oxford University Press.
- . 2012. *Climate Matters: Ethics in a Warming World*, New York: Norton.
- . 2015. "Equality versus Priority: A Useful Distinction," *Economics and Philosophy* 31: 219–28.
- Budolfson, M. and F. Dennig forthcoming. "Optimal Climate Policy and Differential Regional Abatement," in G. Chichilnisky and A. Rezaei, eds., *Handbook on the Economics of Climate Change*, Northampton, MA: Edward Elgar.
- Budolfson, M., F. Dennig, M. Fleurbaey, A. Siebert and R. Socolow 2017. "The Comparative Importance for Optimal Climate Policy of Discounting, Inequalities, and Catastrophes," *Climate Change* 145: 481–94.
- Budolfson, M., F. Dennig, M. Fleurbaey, N. Scovronick, A. Siebert, D. Spears and F. Wagner 2018. "Optimal Climate Policy and the Future of World Economic Development," *World Bank Economic Review* doi: 10.1093/wber/lhx016.
- Budolfson, M. and D. Spears 2018. "Incorporating Non-Anthropocentric Valuation of Ecosystems and Animal Welfare into Environmental Decision-Making Models," working paper.
- Cai, Y., K.L. Judd, and T. S. Lontzek 2013. "The Social Cost of Stochastic and Irreversible Climate Change," Working Paper 18704, National Bureau of Economic Research.
- Caney, S. 2009. "Climate Change and the Future: Discounting for Time, Wealth, and Risk," *Journal of Social Philosophy* 40(2): 163–86.
- Chichilnisky, G. and G. Heal 1994. "Who Should Abate Carbon Emissions?: An International Viewpoint," *Economics Letters* 44(4): 443–49.
- Daniel, K.D., R.B. Litterman, and G. Wagner 2016. "Applying Asset Pricing Theory to Calibrate the Price of Climate Risk," *NBER Working Paper* n. 22795.
- Dasgupta P. 2001. *Human Well-Being and the Natural Environment*, Oxford University Press.
- Dennig F. 2009. *On the Welfare Economics of Climate Change*, D.Phil. University of Oxford.
- Dennig F., M.B. Rudolfson, M. Fleurbaey, A. Siebert, and R.H. Socolow 2015. "Inequality, Climate Impacts on the Future Poor, and Carbon Prices," *Proceedings of the National Academy of Sciences* 112(52): 15827–32.
- Dietz, S. and G.B. Asheim 2012. "Climate Policy under Sustainable Discounted Utilitarianism," *Journal of Environmental Economics and Management* 63(3): 321–35.
- Dietz, S., C. Gollier, and L. Kessler 2015. "The Climate Beta," Grantham Research Institute on Climate Change and the Environment, working paper #190.
- Dietz, S. and N. Stern 2015. "Endogenous Growth, Convexity of Damage and Climate Risk: How Nordhaus' Framework Supports Deep Cuts in Carbon Emissions," *The Economic Journal* 125(583): 574–620.
- Epstein, L. and U. Segal 1992. "Quadratic Social Welfare Function," *Journal of Political Economy* 100: 691–712.

- Epstein, L. and S. Zin 1989. "Substitution, Risk Aversion and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework," *Econometrica* 57(4): 937–69.
- Fankhauser, S., R.S.J. Tol, and D.W. Pearce 1997. "The Aggregation of Climate Change Damages: A Welfare Theoretic Approach," *Environmental and Resource Economics* 10(3): 249–66.
- Ferranna, M. 2015. Three Essays on Decision Making under Risk and Equity Concerns, PhD Thesis, TSE.
- . 2018. "Intergenerational Equity, Risk Aversion and the Social Cost of Carbon," working paper.
- Fleurbaey, M. 2015. "Equality versus Priority? How Relevant Is the Distinction," *Economics and Philosophy* 31: 203–18.
- . forthcoming. "Welfare Economics, Risk and Uncertainty," *Canadian Journal of Economics*.
- Fleurbaey, M. and S. Zuber 2013. "Climate Policies Deserve a Negative Discount Rate," *Chicago Journal of International Law* 13(2): 565–95.
- . 2015a. "Discounting, Risk and Inequality: A General Approach," *Journal of Public Economics* 128: 34–49.
- . 2015b. "Discounting Beyond Utilitarianism," *Economics: The Open-Access, Open-Assessment E-Journal* 9 (2015-12): 1–52.
- . 2017. "Fair Management of Social Risk," *Journal of Economic Theory* 169: 666–706.
- Gollier, C. 2013. *Pricing the Planet's Future*, Princeton: Princeton University Press.
- Greaves, H. 2017. "Discounting for Public Policy: A Survey," *Economics and Philosophy* 33(3): 391–39.
- Greenstone, M., E. Kopits, and A. Wolverton 2013. "Developing a Social Cost of Carbon for US Regulatory Analysis: A Methodology and Interpretation," *Review of Environmental Economics and Policy* 7(1): 23–46.
- Hammit, J. 2013. "Positive versus Normative Justifications for Benefit-Cost Analysis: Implications for Interpretation and Policy," *Review of Environmental Economics and Policy* 7(2): 199–218.
- Harsanyi, J. 1955. "Cardinal Welfare, Individualistic Ethics and Interpersonal Comparison of Utilities," *Journal of Political Economy* 63: 309–21.
- Heath, J. 2017. "Climate Ethics: Justifying a Positive Social Time Preference," *Journal of Moral Philosophy* 14(4): 435–62.
- Hope, C. 2008. "Discount Rates, Equity Weights and the Social Cost of Carbon," *Energy Economics* 30: 1011–19.
- Hsiang S., R. Kopp, A. Jina, J. Rising, M. Delgado, S. Mohan, D.J. Rasmussen, R. Muir-Wood, P. Wilson, M. Oppenheimer, K. Larsen, and T. Houser 2017. "Estimating Economic Damage from Climate Change in the United States," *Science* 356(6345): 1362–69.
- Hsiung, W. and C. Sunstein 2007. "Climate Change and Animals," *University of Pennsylvania Law Review* 155: 1695–740.
- IPCC 2014. *5th Assessment Report, Working Group III*, Cambridge: Cambridge University Press.
- Jackson, F. and M. Smith 2006. "Absolutist Moral Theories and Uncertainty," *Journal of Philosophy* 103(6): 267–83.
- Kahneman, D., E. Diener, and N. Schwarz, eds. 1999. *Well-Being: The Foundations of Hedonic Psychology*, New York: Russell Sage Foundation.
- Kelleher, J.P. 2017. "Descriptive Versus Prescriptive Discounting in Climate Change Policy Analysis," *Georgetown Journal of Law & Public Policy* 15: 957–77.
- Koopmans T.J. 1960. "Stationary Ordinal Utility and Impatience," *Econometrica* 28: 287–309.
- Lemoine, D. and C.P. Traeger 2016. "Economics of Tipping the Climate Dominoes," *Nature Climate Change* 6: 514–20.
- Lontzek, T.S., Y. Cai, K.L. Judd, and T.M. Lenton 2015. "Stochastic Integrated Assessment of Climate Tipping Points Indicates the Need for Strict Climate Policy," *Nature Climate Change* 5: 441–44.

- McShane, K. forthcoming. "Why Animal Welfare Is Not Biodiversity, Ecosystem Services, or Human Welfare: Toward a More Complete Assessment of Climate Impacts," *Les Ateliers de l'Éthique/The Ethics Forum*.
- Méjean, A., A. Pottier, S. Zuber, and M. Fleurbaey 2017. "Intergenerational Equity under Catastrophic Climate Change," CES working paper #2017.40.
- Millner, A. 2013. "On Welfare Frameworks and Catastrophic Climate Risks," *Journal of Environmental Economics and Management* 65(2): 310–25.
- Nagel, T. 2005. "The Problem of Global Justice," *Philosophy and Public Affairs* 33(2): 113–47.
- Ng, Y-K. 1989. "What Should We Do About Future Generations? Impossibility of Parfit's Theory X," *Economics and Philosophy* 5: 235–53.
- Nordhaus, W.D. 2008. *A Question of Balance: Weighing the Options on Global Warming Policies*, New Haven: Yale University Press.
- . 2007. "A Review of the Stern Review on the Economics of Climate Change," *Journal of Economic Literature* 45: 686–702.
- . 2017. "Revisiting the Social Cost of Carbon," *PNAS* 114(7): 1518–23.
- Parfit, D. 1984. *Reasons and Persons*, Oxford: Clarendon Press.
- Pearce, D. 2003. "The Social Cost of Carbon and its Policy Implications," *Oxford Review of Economic Policy* 19(3): 362–84.
- Portney, P.R. and J.P. Weyant 1999. *Discounting and Intergenerational Equity*, New York: Resources for the Future.
- Pottier, A. 2016. *Comment les Economistes Réchauffent la Planète*, Paris: Seuil.
- Pottier, A., A. Méjean, O. Godart and J.C. Hourcade 2017. "A Survey of Global Climate Justice: From Negotiation Stances to Moral Stakes and Back," *International Review of Environmental and Resource Economics* 11: 1–53.
- Rabinowicz, W. 2002. "Prioritarianism for Prospects," *Utilitas* 14(1): 2–21.
- Ramsey, F.P. 1928. "A Mathematical Theory of Savings," *Economic Journal* 38: 543–59.
- Rawls, J. 1999. *The Law of Peoples*, Cambridge, MA: Harvard University Press.
- Rezai, A., D. Foley, and L. Taylor 2012. "Global Warming and Economic Externalities," *Economic Theory* 49(2): 329–51.
- Sandsmark, M. and H. Vennemo 2007. "A Portfolio Approach to Climate Investments: CAPM and Endogenous Risk," *Environmental and Resource Economics* 37(4): 681–95.
- Scheffler, S. 1982. *The Rejection of Consequentialism*, Oxford and New York: Clarendon Press.
- Schelling, T. 1995. "Intergenerational Discounting," *Energy Policy* 23: 395–401.
- Scovronick, N., M.B. Budolfson, F. Dennig, M. Fleurbaey, A. Siebert, R.H. Socolow, D. Spears and F. Wagner 2017a. "Impact of Population Growth and Population Ethics on Climate Change Mitigation Policy," *Proceedings of the National Academy of Sciences* 114(46): 12338–43.
- Scovronick, N., M.B. Budolfson, F. Dennig, M. Fleurbaey, W. Peng, R.H. Socolow, D. Spears, and F. Wagner 2017b. "The Impact of Health Co-Benefits on Optimal Global CO2 Mitigation Policy," working paper.
- Sebo, J. forthcoming. "Animals and Climate Change," in M. Budolfson, T. McPherson and D. Plunkett, eds. *Philosophy and Climate Change*, Oxford University Press.
- Seidenfeld, T. 1988. "Decision Theory Without 'Independence' or Without 'Ordering'," *Economics and Philosophy* 88: 267–90.
- Sherwood, S. C. and M. Huber 2010. "An Adaptability Limit to Climate Change due to Heat Stress," *PNAS* 107(21): 9552–55.
- Spears, D. 2017. "Making People Happy or Making Happy People? Questionnaire-Experimental Studies of Population Ethics and Policy," *Social Choice and Welfare* 49(1): 145–69.
- Steele, K.S. 2010. "What Are the Minimal Requirements of Rational Choice? Arguments from the Sequential-Decision Setting," *Theory and Decision* 68: 463–87.
- Stern, N. 2006. *Stern Review: The Economics of Climate Change*, volume 30. London: Her Majesty's Treasury.

- Temkin, L. 1993. *Inequality*, Oxford: Oxford University Press.
- Tol, R.S.J. 2001. "Equitable Cost-Benefit Analysis of Climate Change Policies," *Ecological Economics* 36(1): 71–85.
- . 2009. "The Economic Effects of Climate Change," *Journal of Economic Perspectives* 23(2): 29–51.
- Traeger, C. 2015. "Analytic Integrated Assessment and Uncertainty," Department of Agricultural and Resource Economics, University of California, Berkeley.
- Vallentyne, P. 2005. "Of Mice and Men: Equality and Animals," *The Journal of Ethics* 9(3/4): 403–33.
- Van den Bergh J.C.J.M. and W.J.W. Botzen 2014. "A Lower Bound to the Social Cost of CO₂ Emissions," *Nature Climate Change* 4: 253–58.
- Warren, R. 2014. "Environmental Economics: Optimal Carbon Tax Doubled," *Nature Climate Change* 4: 534–35.
- Weitzman, M. 2009. "On Modeling and Interpreting the Economics of Catastrophic Climate Change," *The Review of Economics and Statistics* 91(1): 1–19.
- World Bank 2016. (S. Hallegatte, M. Bangalore, L. Bonzanigo, M. Fay, T. Kane, U. Narloch, J. Rozenberg, D. Treguer, and A. Vogt-Schilb) *Shock Waves: Managing the Impacts of Climate Change on Poverty*, Climate Change and Development Series, Washington, DC: World Bank.
- Zuber, S. and G.B. Asheim 2012. "Justifying Social Discounting: The Rank-Discounted Utilitarian Approach," *Journal of Economic Theory* 147: 1572–601.