The Social Cost of Carbon from Theory to Trump

J. Paul Kelleher

To appear in Ravi Kanbur and Henry Shue (eds.) *Climate Justice: Integrating Economics and Philosophy* (Oxford University Press, under contract)

Abstract

The social cost of carbon (SCC) is a central concept in climate change economics. This chapter explains the SCC and investigates it philosophically. As is widely acknowledged, any SCC calculation requires the analyst to make choices about the infamous topic of discount rates. But to understand the nature and role of discounting, one must understand how that concept—and indeed the SCC concept itself—is yoked to the concept of a value function, whose job is to take ways the world could be across indefinite timespans and to rank them from better to worse. A great deal, therefore, turns on the details of the value function and on just what is meant by "better" and "worse." This chapter seeks to explicate these and related issues, and then to situate them within the evolving landscape of federal climate policy in the United States.

Keywords

climate change economics, climate change ethics, climate change policy, social discount rate, time preference, future generations, intergenerational equity

12.1. Introduction

The *social cost of carbon* (SCC) is a central concept in climate change economics. My aim in this chapter is to explain the concept and then to look under its philosophical hood, so to speak. It is widely acknowledged that calculating the SCC involves making choices about the infamous topic of discount rates. But to understand the nature and role of discounting, it is crucial to understand how that economic concept—and indeed the SCC concept itself—is yoked to the concept of a *value function*, whose job is to take ways the world could be across indefinite timespans and to rank them from better to worse. A great deal, therefore, turns on the details of the value function and on just what is meant by "better" and "worse." I seek to explicate these and related issues, and then to situate them within the evolving landscape of federal climate policy in the United States.

12.2. Value Functions, Discount Rates, and the Road from Value to Policy Choice

Climate change economics employs two foundational concepts: the *feasible* and the *valuable* (Brennan, 2007). Feasibility concerns how things could be given technological constraints. For example, it focuses on what forms of conventional and environmental investment are possible, and how a given type and degree of investment would change the current trajectory of economically important phenomena.

The most salient phenomenon for mainstream climate change policy analysis is the consumption of commodities whose contribution to individual well-being is usually measured by individuals' willingness to pay for them.[1] Economists try hard to understand in detail how climate change will impact the consumption of various commodities, although the economic models I'll focus on are concerned with averages and aggregates. For example, the most simple models treat the issue of

[1] This is the standard *economic* approach to measurement. Philosophers routinely criticize this way of doing things. See, e.g., Hausman (2012).

feasibility as that of identifying the range of *time paths of per capita global consumption* that are possible given various levels of financial and environmental investment. That is, if we assume that time, t, is measured in discrete periods ($t = 0, 1, 2, \ldots$), and that global per capita consumption at t is an aggregate index C_t, then climate economics is concerned with which time paths ($C_0, C_1, C_2, \ldots$) are feasible, given current and future technologies.[2]

The second fundamental concept for climate economics, *value*, is typically operationalized in the form of a *value function* (this is sometimes referred to as a *social welfare function*). The job of a value function is to place feasible time paths of consumption into a ranked ordering.[3] In this way, value functions are tools for ranking the different ways the world could be.

Without yet getting into the details of any particular value function, suppose we have one that we like—call it V—which takes each feasible consumption path and assigns to it a real number. (The entire path is the

---

[2] In this simple framework, environmental quality is relevant only insofar as it affects aggregate consumption. It is, however, consistent with the economic foundations of climate policy analysis to treat environmental quality as having fundamental importance. For more on this, see Sterner and Persson (2008) and Heal (2009). It is also possible to switch from a focus on global per capita consumption to the per capita consumption in each of several geographic regions. See, e.g., Nordhaus and Yang (1996).

[3] Strictly speaking, a value function *represents* an ordering, which suggests that the ordering is somehow prior to the function. But I will speak as if the function *determines* the ordering, since a great deal of climate economics concerns first choosing the form and content of the value function and only then running models to discover which ordering over feasible streams results.

V-function's "argument," we say.) For us to "like" a value function is for us to think it does a good job of using such number assignments to place the paths into a betterness ordering, with the best sequence being assigned a number higher than any other. Setting aside some key technical details concerning measurement,[4] we can use the concept of a value function to define the SCC as follows.

First, determine which feasible consumption path is the business-as-usual (BAU) path, i.e. the path that would prevail in the absence of additional climate policy. Now find the number, *V*(BAU), that *V* assigns to this BAU path. Next, determine which consumption path would result if the BAU path were changed in only the following respect: there is one extra ton of carbon dioxide (CO_2) emitted in the first period. Call this the *emission-perturbed path*, or PER_E. Now find the number, $V(PER_E)$, that *V* assigns to PER_E. Next, subtract $V(PER_E)$ from *V*(BAU) to find the *V*-based measure of the difference made by one extra ton of CO_2 emissions (where the BAU consumption path is the comparison baseline). Now use a parallel calculation to find the difference made by adjusting the BAU path in only the following respect: reduce the first time period's per capita consumption by one dollar. Call the result the *consumption-perturbed path*, or PER_C, and calculate the difference it makes (again using BAU as

[4] See Broome (1991, pp. 70–5).

the baseline) by finding $V(\text{BAU}) - V(PER_C)$. The formula for calculating the SCC is then:

$$SCC = \frac{V_{BAU} - V_{PER_E}}{V_{BAU} - V_{PER_C}} \quad (1)$$

In words, the SCC is the ratio of the *V*-based difference made by a marginal unit of present emissions to the difference made by a marginal unit of present consumption. This ratio expresses the marginal impact on *V* that present CO_2 emissions make in terms of the marginal impact that present consumption makes. It is a way of making these impacts comparable, so that (for example) the benefits of greenhouse gas abatement can be compared with the monetary costs associated with abatement.

I have focused here on the SCC for the first time period, which expresses the difference made (across time) by marginal emissions in the first period in terms of the difference made by marginal consumption in that period. But the SCC can in principle be calculated for any designated time period. We will soon see why the SCC in all time periods is of special importance to climate economics.

When it comes to the SCC, clearly a great deal turns on the form of the value function, *V*, and on its parameter values. As regards form, the standard framework in climate economics stems from an approach set out by Frank Ramsey in 1928 (Ramsey, 1928). Ramsey's approach was utilitarian in two senses. First, Ramsey's value function had the following utilitarian form:

$$V_{Ramsey} = \sum_{t=0}^{MAX} N_t \cdot u(c_t) \tag{2}$$

where $u(c_t)$ expresses the utility of per capita consumption level c at t and N_t is the number of people alive at t. Ramsey's approach was therefore to rank consumption paths in terms of the sum total of utility they contain.[5] The second sense in which Ramsey's framework was utilitarian is that he held that the *betterness ordering* given by a utilitarian *V* is *ipso facto* a *rightness ordering*, where a rightness ordering ranks feasible paths in terms of *everything* that matters to policy choice. Let us refer to such a tight relationship between value (as defined by a given *V*-function) and choice using the following phrase:

$$Value \rightarrow Choice \tag{3}$$

Many philosophers will be comfortable with (and all philosophers will be familiar with) the idea—which is just the rejection of (3)—that value considerations cannot *by themselves* fully determine proper choice from

---

[5] This not fully accurate. For technical reasons stemming from problems with the idea of an infinite sum, Ramsey sought to minimize the difference between actual realized utility and a maximum possible amount he defined in terms of a concept he called *bliss*.

among feasible policy options. Climate economists, in my experience, find this a harder idea to make sense of.[6] I'll return to it shortly.

As already defined, the SCC is a *marginal* concept: it is the ratio of the difference made by a *very small* increase in CO_2 emissions to the difference made by a *very small* decrease in first-period consumption. Because of these marginality assumptions, equation (1) is not the only formula one can use to calculate the SCC, and Ramsey once again helps us see why.

Ramsey in effect showed that when the value function takes the following *generalized utilitarian* form,

$$V_{GenUtil} = \sum_{t=0}^{MAX} N_t \cdot u(c_t) \cdot \frac{1}{(1+\delta)^t} \tag{4}$$

where δ represents the *utility discount rate*, one can derive from *V* the following formula:

$$\rho_t = \delta + \eta \cdot g_t \tag{5}$$

Here the utility discount rate, δ, is the rate at which a given increment of utility declines in value (as measured by its impact on *V*) as its enjoyment is delayed by one period. (Ramsey believed this parameter should be set to zero to yield equation (2), which is a special case of equation (4).) η signifies

[6] This separates climate economists from health economists, since the latter are familiar with the idea that economists' focus—e.g. maximizing quality-adjusted life years (QALYs)—is not the only consideration relevant to health policy.

the *elasticity of the marginal utility of consumption*, which represents the "curvature" of the utility function $u(\cdot)$, or the degree to which enjoying more and more consumption yields less and less utility. Finally, g_t is the rate of growth in per capita consumption between two periods, t and $t+1$, along a given time path of consumption. What Ramsey showed is that for any given time period along any given path of consumption, one can use equation (5) to derive that period's *consumption discount rate* (ρ_t). A consumption discount rate is a *rate of indifference*: for example, period 0's consumption discount rate it is the rate of return required in period 1 in order for *V* to be indifferent between investing a marginal unit of consumption in period 0 and someone's consuming it in period 0. Thus if period 0's consumption discount rate is 0.05, then a 5 percent return is required in period 1 in order for *V* to be indifferent between consuming a marginal unit in period 0 and investing it instead.

Consumption discount rates offer a sort of shortcut around using equation (1) to calculate SCC values. So long as the incremental emissions associated with the SCC concept create incremental consumption changes in future time periods that are genuinely marginal[7] (when compared with the BAU consumption path), we can use equation (5) to derive each period t's consumption discount rate. We then define

[7] See Dietz and Hepburn (2013) for the requirements of genuine marginality.

period $t + 1$'s *consumption discount factor* as $\frac{1}{(1+\rho_t)}$. By multiplying each period's incremental change in consumption by its consumption discount factor, we can express period $t + 1$'s consumption in terms of its period t equivalent. Then, once we have a full time path of consumption discount factors, we are in a position to translate incremental consumption in any period into its equivalent in another period. This enables us to use a time path of consumption discount factors to express a time stream of incremental consumption changes as a "present value," that is, as the change in period 0 consumption that is equivalent to it. In this way, we can avoid having to use equation (1) to calculate an SCC value: so long as the relevant consumption changes are truly marginal, we can use the consumption discount factor method instead.

Equation (5) is sometimes called the *Ramsey rule*,[8] but that name is also frequently given to a conceptually quite distinct proposition. To see this, note that equation (5) can be derived entirely from equation (4). That is, (5) gives us the marginal rate of indifference *with respect to the value considerations embodied in V*. Nothing in the manner of claim (3)—i.e. the claim that *V* reflects all considerations relevant to policy choice—is needed to derive (5). However, if one embraces *both* (3) *and* (4), then

[8] See for example Johansson and Kriström (2015, p. 67), Adler and Treich (2015, p. 18), and Kolstad et al. (2014, p. 229).

one can derive an ostensibly similar but entirely new proposition that is also frequently dubbed the Ramsey rule:[9]

$$r_t = \delta + \eta \cdot g \tag{6}$$

Here r_t represents the social rate of return on investment, or the *productivity of capital*. If one subscribes to (3) and (4), then (6) follows for reasons that, once again, Ramsey was the first to articulate. The idea is intuitive. If at a certain time, and along a certain BAU consumption path, one's consumption discount rate (as given by (5) and the right-hand side of (6)) is below the productivity of capital, then one can increase value (as measured by *V*) by investing a bit more today in order to earn a rate of return that exceeds one's *V*-based rate of indifference. Since we are supposing for the sake of argument that one subscribes to claim (3), one will apply Ramsey's logic right up to the point where one's consumption rate of indifference is exactly equal to the productivity of capital. The same logic will apply, *mutatis mutandis*, if the productivity of capital is currently lower than the consumption rate of indifference. In that case, one maximizes *V* by increasing consumption today and investing less, right up until the point at which the world conforms to the equality expressed in (6).

As noted, (5) and (6) are each frequently called "the" Ramsey Rule. It would be useful to have different names for them, and if it were up to

---

[9] See for example Dasgupta and Heal (1979, pp. 296–7), Stern (2014, p. 456), and Arrow et al. (1996, p. 134).

me I'd call (5) the Ramsey *formula* and (6) the Ramsey *rule*. (Even then, I wouldn't want readers to assume that the Ramsey rule is a morally or even economically *sound* rule—only that it is a rule that many economists, including Ramsey, endorse.) In any case, the important point for now is that these express distinct propositions and that one will embrace the rule embodied in (6) only if one also embraces the tight connection between value and choice expressed in (3).[10]

To illustrate these conceptual points, consider the view taken by environmental economists David Pierce, Edward Barbier, and Anil Markandya (PBM).[11] PBM claim that the consumption discount rate given by (5) should reflect "current generational orientated considerations."[12] By this they mean, for example, that the utility discount rate should reflect the rate that individuals today reveal in their behavior when they trade-off utility within their own lives. Since people seem to show an intra-personal bias in favor of earlier utility over later utility, PBM's approach to (5) has the implication that climate policy analysis should discount the interests of future generations relative to those of the current generation. That is actually the standard view in climate economics, and I'll discuss it further later in the chapter. Yet to this standard view PBM add an unconventional

[10] Dasgupta 1982 clearly distinguishes these propositions, referring only to (6) as the Ramsey *Rule*. See in particular his equations (10) and (18).
[11] Pearce et al., 1990.
[12] Pearce et al., 1990, p. 46.

proviso: they claim that one should first "define the rights of future generations"—decide what levels of consumption they are morally *entitled* to, say—and only then maximize a "current generation oriented" version of *V subject to the constraint that future generations' rights are fully respected*.[13] PBM are, therefore, likely to reject (3), since they do not believe that *V* captures all of the normative considerations relevant to policy choice. They would therefore also reject (6). But they needn't—and in fact shouldn't—reject (5), since (5) is entailed by any *V*-function having the generalized utilitarian form (4), which PBM's current generation oriented *V*-function does. Equation (5) merely gives the marginal rate of indifference between consumption in different time periods *as judged by V*. This rate of indifference is *relevant* to choice just in case the ranking given by *V* is relevant to choice. But that ranking can be relevant to choice without its being the *sole* relevant consideration. In contrast, the Ramsey rule expressed by equation (6) entails that the ranking given by *V* is the sole policy-relevant consideration.

As I've said, PBM are in the minority in rejecting an airtight relationship between value (as expressed by *V*) and policy choice. Most economists working in climate economics think that it is the job of economists to find a *V*-function that can sustain a recommendation of the form, "Identify the feasible consumption path that maximizes *V* and then

[13] Pearce et al., 1990, p. 46.

implement the portfolio of policies required to put the world on that path."[14] But it is absolutely essential to see that this default position is entirely discretionary: absolutely nothing *in the economics* dictates that economists must embrace this particular relationship between proper policy choice and the conception of value embodied in the *V*-function. *V-functions merely place consumption paths into an ordering*. What that ordering signifies morally, and what public policy is to do with that ordering, are distinct and further questions.[15]

I hope it is now clear how the concept of the social cost of carbon is yoked to the idea of a value function, and how the value function might or might not be related to the practical task of policy choice. If there is indeed a gap between the policies that *V* ranks highest and the policies that we should pursue all-things-considered, then the SCC will not be an ironclad guide to policy choice. But if there is no gap, then the SCC will be a very useful policy instrument. For whenever the SCC exceeds the cost we would have to bear today to abate a ton of CO_2, we know that we would increase *V* by abating an additional ton. We also know (and this is a

[14] Or, if they insist on restricting themselves to *climate* policy only, their recommendation will take the form, "Identify the feasible consumption path that maximizes *V*, then identify the associated time-path of greenhouse gas emissions (or emissions reductions), and then implement the portfolio of policies required to put the world on that path of emissions (reductions)." This is what von Below, Denning, and Jaakkola call the "naïve implementation" of the Ramsey Rule. See von Below et al. (n.d., pp. 26–7).

[15] See (Kelleher 2017a) for an extended discussion on this point.

standard result in climate economics) that along the path that *V* ranks highest—what those who embrace (3) would call the "optimal" path—the SCC for each period will be equal to the cost of abating one more ton of CO_2 in that period. (This last "optimality condition" is used by economists to help work out what the optimal path looks like.) Finally, we also know that the SCC along the highest-ranked path gives us the "optimal tax," that is, the tax that would bring us on to the highest-ranked path if we applied it to each and every ton of CO_2 that is emitted around the world.

12.3. Investigating Value Functions Further

Not all climate economists embrace a *V*-function falling into the generalized utilitarian family captured by (4). For example, in a recent book Humberto Llavador, John Roemer, and Joaquim Silvestre (LRS) explore and defend just the second half of the two-pronged view championed by Pearce, Barbier, and Markandya. In LRS's view, climate economists should reject all *V*-functions that resemble (4) and from which the Ramsey formula (5) can be derived.[16] They argue instead that the entire framework should be built around a *V*-function that ranks consumption paths with respect to how the worst-off generation fares.[17] On such a view, larger benefits for better-off generations cannot justify smaller losses for worse-off generations, *unless* the trade-off

[16] Llavador et al., 2015, p. 205.

[17] Actually, LRS's view is at once more complicated and more lenient than this, but I will not get into the details of their view here.

is part of a policy that somehow improves the situation of the worst-off generation. But that means LRS-style *V*-rankings are not interested in adding time-stamped gains and losses and weighing them against one another, and this in turn means that they have little use for the SCC concept as I have defined it. For it is only in the context of what economists call *additive V*-functions that the SCC becomes a relevant concept both from the standpoint of economics and—if policy cares about the *V*-function—from the standpoint of public policy.

V-functions (2) and (4) are not the only ones that display additivity, however. A still different one is the *prioritarian V*-function:

$$V_{prior} = \sum_{t=0}^{MAX} g(u(c_t)) \tag{7}$$

Like (2) and (4), a Ramsey formula specifying the consumption rate of indifference can be derived from (7), though it looks somewhat different from (5) since (7) sums not discounted utilities, but utilities that have been "transformed" by a function *g*(.) that gives less and less weight to extra utility the more beneficiaries already have of it.[18] But like (2) and (4)—and quite unlike LRS's "maximin" *V*-function—the additive prioritarian *V*-function and its associated Ramsey formula can be used to calculate a (prioritarian) SCC using the formula given by (1).[19]

---

[18] Adler and Treich, 2015, p. 296.
[19] Adler et al., 2017.

These three families of value functions—generalized utilitarian, prioritarian, and maximin—share the feature that they address a heterogeneous intra- and inter-temporal demography head-on, and calibrate their parameters with an eye toward yielding betterness orderings that acknowledge trade-offs between distinct people (or at least different groups of people—e.g. generations). LRS call approaches with this feature *ethical observer* approaches, while others call them *social planner* approaches. They stand in contrast to what are called *representative agent* models, which aim to make consumption path-ranking exercises tractable by giving the societal V-function the same form as an anonymous currently living individual's V-function. Economists like this representative agent approach because there are observational data from individual behavior they can use to calibrate the value parameters, and because they are already trained to think systematically about the structure and properties of individual preference.

One significant problem for representative agent models is that they cannot adequately reflect prioritarian considerations. Prioritarian value functions like (7) rank consumption paths by applying a "concave transform," $g(.)$, to each period-specific utility and then summing the results of these operations across the relevant time horizon. The transform is used to give more weight to utility improvements that accrue to those who start with less of it. The resulting conception of betterness is therefore conceptually distinct from a conception on which betterness is represented by the sum of utilities. The problem for representative agent models—models that treat societal betterness rankings as if they were the ranking of an individual—is that it is definitional of the idea of (von Neumann–Morgenstern) utility that betterness for an individual is represented by summing utilities, and not by summing concave transformations of utilities.[20] Many representative agent proponents try to evade this objection by allowing prioritarian considerations to influence η, the value parameter representing the elasticity of the marginal utility of consumption. As noted earlier, η reflects the curvature of the utility function and thus the degree to which adding more and more consumption brings less and less utility. By increasing η, representative agent proponents hope to express the *combined* curvature of an explicit utility function and an implicit prioritarian transform.[21] Yet not only is this solution conceptually quite imperfect, but it also ignores two different

[20] For more on this point, see Greaves (2015). Traeger (2010) derives a representative agent framework that obeys the von Neumann–Morgenstern axioms and that also permits concave transformations of utility. But these transformations are an expression of risk aversion and are relevant only in contexts of uncertainty; they drop out when certain consumption paths are being evaluated. By contrast, prioritarianism assigns values to certain consumption paths by summing concave transformations of period-specific utilities. So Traeger's framework still cannot accommodate prioritarian considerations.

[21] See Stern (1977, pp. 241–2), Dasgupta and Heal (1979, pp. 279–80), and Kaplow and Weisbach (2011) for rare exceptions in the literature that make this prioritarian transform explicit.

technical issues—explained in note 22—that absolutely require the explicit decomposition the representative agent proponent was hoping to avoid.[22]

Whether they adopt a representative agent or a social planner approach to *V*-functions, most economists working in climate change economics rely heavily on *revealed preference* to calibrate their *V*-function's value parameters. Some go so far as to say that if parameters are *not* set in this "descriptivist" way—for instance, if they are instead set using "prescriptivist" *a priori* moral reasoning—then the whole path-ranking exercise would no longer fall within the discipline of economics.[23] Here, for example, is Martin Weitzman:

> [C]line and Stern are soulmates in their *cri de coeur* justifying $\delta \approx 0$ by relying mostly on a priori philosopher-king ethical judgements about the immorality of treating future generations differently from the current generation—instead of trying to back

[22] First, "implicit prioritarian" approaches work with a utility function that is unique up to positive affine transformations, whereas no prioritarian value function is invariant to positive affine transformations of the utility function. What a prioritarian value function requires is a utility function that is invariant to *ratio* transformations. See Adler and Treich (2015, §3.2). Second, integrating prioritarianism with decision theory requires an explicit choice between what Adler and Treich call *ex ante* and *ex post* versions of prioritarianism under uncertainty. Although each of these has serious theoretical costs—*ex ante* prioritarianism violates the "sure thing principle," *ex post* prioritarianism violates the *ex ante* Pareto relationship between individual betterness rankings and social betterness rankings—any bona fide prioritation must choose between them. (See Adler, 2012, ch. 7.) But implicit prioritarianism standardly leads economists who embrace it to evade the issue altogether.

[23] See Kelleher (2017b) and Kelleher and Wagner (n.d.) for more on the so-called descriptivist/prescriptivist divide in climate economics.

> out what possibly more representative members of society than either Cline or Stern might be revealing from their behavior is their implicit rate of pure time preference. An enormously important part of the "discipline" of economics is supposed to be that economists understand the difference between their own personal preferences for apples over oranges and the preferences of others for apples over oranges. Inferring society's revealed preference value of δ is not an easy task in any event (here for purposes of long-term discounting, no less), but at least a good-faith effort at such an inference might have gone some way towards convincing the public that the economists doing the studies are not drawing conclusions primarily from imposing their own value judgements on the rest of the world. (Weitzman, 2007, p. 712)

Pace Weitzman, revealed preference approaches face several serious objections from the standpoint of policy analysis. First, they violate the Humean dictum that one cannot infer an "ought" from an "is." The data that revealed preference advocates wish to use to calibrate the value parameters reveal the trade-offs that real-life agents *do in fact* make in their daily lives. But, arguably, the *V*-function required for social policy analysis should reflect societal judgments about how intergenerational trade-offs *should* or *ought* to be made. It would be one thing if revealed preference advocates warned, as

PBM did earlier, that such *V*-functions capture just *one among many other* considerations relevant to policy choice. But a great many economists who prioritize revealed preference take the results of their models to indicate what *proper* policy *should* look like. Few say what they would need to say to properly heed Hume's dictum: "The ranking of consumption paths I offer is more the product of quantitative sociology or anthropology than it is of political theory or applied ethics. So please do not assume that the ranking I end up with is necessarily indicative of *morally justifiable* social choice." Instead, revealed preference proponents describe the results of their analyses as determining "optimal policy," and they criticize policies that would deviate from the "optimum" so-defined.[24]

Even if one grants that "current generation oriented" observations are *morally relevant* to policy analysis, a second problem with the revealed preference approach is that it is arguably predicated upon the wrong observational data. For the value parameters in revealed preference models are derived from the value parameters in individuals' own, largely self-interested consumer behavior. But the path-ranking exercise relevant to climate policy analysis concerns ranking paths (and thus policies) that will impact billions of people for centuries to come. Referring to the types of consumer behavior that revealed preference advocates wish to use to set the value parameters, Nicholas Stern writes:

[24] See, for example, Nordhaus (1994, pp. 79ff.).

"as this borrowing and lending take place through private decisions . . . this does not necessarily answer the relevant question . . . namely, how do we, acting together, evaluate our responsibilities to future generations."[25] This is a strong argument against policy analysis that relies heavily upon the data standardly used to calibrate value parameters in revealed preference models. Different sorts of data could be collected, for example in the context of focus groups, but to my knowledge this work has not been done on a scale sufficient for application in policy analysis.[26]

The reliance (or not) on revealed preference will inevitably influence how a given value function trades off well-being at one time with well-being at another. In addition, note that whether one views a given *V*-function's trade-offs as defensible will also have much to do with one's prior stance on claim (3). That is, if one begins with the view that a value function's ranking of feasible consumption paths should be an *all-things-considered* rightness ranking, then one will seek a value function that incorporates all normative considerations one deems relevant to policy choice. However, if one instead uses a value function to rank paths in terms of *some but not all* normatively relevant considerations, then the trade-offs embodied in the value function need not coincide perfectly with the full set of trade-offs one ultimately wants embodied in policy. Consider

[25] Stern, 2010, p. 51.
[26] For a proof-of-concept example of what this could look like, see Frederick (2003).

the following distinction between moral principles described by John Broome:

> Very roughly, our moral duties can be divided into duties of beneficence—promoting good—and duties of justice . . . The need for this division can be illustrated by an example that comes from Judith Jarvis Thomson. A surgeon has in her hospital five patients, each in need of a different organ for transplant. One needs a new heart, one needs a new liver, and so on. Each will die unless she gets a new organ. The surgeon kills an innocent visitor to the hospital and distributes her organs to her patients. Thereby the surgeon saves five lives at the expense of one. She successfully promotes good. Yet she clearly acts wrongly . . . It must be that there is some principle of morality that can conflict with the duty to promote good, and is sometimes important enough to override it. In this particular case, it is evidently a duty not to inflict harm on someone, even for the sake of greater general goodness. I take this to be a duty of justice.[27]

Elsewhere Broome claims that the value function relevant to the economics of climate change is insensitive to justice-based considerations and instead

[27] Broome, 2016, p. 921. The Fifth Assessment Report of the Intergovernmental Panel on Climate Change (IPCC) was the first IPCC report to reference and incorporate this distinction between goodness and justice. Broome was a lead author of the relevant section. See Kolstad et al. (2014, p. 215).

aims to rank consumption paths solely in terms of their capacity to promote general goodness.[28] But if one did not understand that this is Broome's focus, one might mistakenly criticize his claims about the value parameters; after all, the parameters relevant to a path-ranking exercise concerned solely with what Broome calls "goodness" may well diverge from the parameters relevant to an exercise concerned with both goodness and justice. If one did not realize that Broome is focused on goodness only, and if one instead assumed he was focused on *all* relevant normative considerations, then one might be inclined to pick a misguided theoretical fight with him about his preferred value parameters. This in fact is my diagnosis of a great number of putative disagreements between philosophers and climate economists.[29]

The key takeaway for our purposes is that since the SCC depends so heavily on the *V*-function, it could be a giant mistake—both analytically and dialectically—to speak of "the" SCC. For if it makes sense in different contexts to be interested in different kinds of path-ranking exercise, and if each of these employs an additive *V*-function, then there will be a distinct SCC concept tied to each sort of exercise. This is especially important in the context of policy, since only an SCC tied to a *V*-function capturing all

[28] Broome, 2012, ch. 6.

[29] I develop this analysis of prominent debates in the literature in Kelleher (2017a) and Kelleher (2017b).

relevant normative considerations should be used to identify policies that are "optimal" in the sense relevant to final social choice.

The observations made in this section can be used to analyze the approach taken by the controversial *Stern Review* of the economics of climate change.[30] Stern adopted a value function in line with (2), with a utility discount rate of essentially zero.[31] Stern's value function was not of the representative agent variety because he conceived of the utility discount rate as an explicitly moral parameter serving to weight utility gains going to distinct generations. Yet in line with more mainstream analyses—including just about all representative agent analyses—Stern chose to calibrate the elasticity of the marginal utility of consumption on the basis of revealed preference, by drawing on studies of revealed consumer behavior.[32] At the same time, Stern notes that his *V*-function "has no room . . . for ethical dimensions concerning the processes by which outcomes are reached," and that it thereby ignores "different notions of ethics, including those based on concepts of rights, justice and freedoms."[33] Stern's *V*-function was therefore in line with the goodness-

---

[30] Stern, 2007.

[31] Stern actually adopted a value function with a very low positive utility discount rate, but that rate was used to reflect an exogenous risk of extinction. His approach was consistent with a utility discount rate of zero, so long as extinction risks can be captured in some other way (as indeed they can be).

[32] See Stern (2007, p. 52, n. 10). Elsewhere in the *Review* Stern suggests that this is actually a flawed approach to the elasticity of the marginal utility of consumption. See, for example, pp. 33–4, box 2.1, and p. 34, n. 6.

[33] Stern, 2007, p. 32.

focused *V*-function described and employed by Broome, and indeed Broome was acknowledged in the *Stern Review* for being an influential advisor.[34] And yet, despite working with a *V*-function that self-consciously ignores at least one set of morally crucial considerations—considerations that Stern himself seems to think are relevant to policy—Stern's headline conclusion was that "prompt and strong action is clearly warranted."[35] This is what we would expect from someone who subscribes to (3).

One can therefore find in the *Stern Review* traces of many of the approaches to the *V*-function canvassed in this section, but without a methodological discussion of whether these elements are compatible with one another, or why definitive policy conclusions can be derived from *V*-functions (and thus SCCs) that reflect only a subset of policy-relevant considerations. In flagging this I do not mean to reject the *Stern Review*'s ultimate policy conclusions. I mean only to suggest that future economic analyses of climate change will be all the stronger if they attend to the issues, distinctions, and tensions raised in this section.

12.4. The Social Cost of Carbon: From Theory to Trump

With the express aim of informing climate policy in the United Kingdom, the *Stern Review* was commissioned by the Chancellor of the Exchequer in 2005, and was released in 2006. By contrast, it was not until 2007 that the US

---

[34] Stern, 2007, p. 32.
[35] Stern, 2007, p. xv.

Supreme Court compelled the Environmental Protection Agency (EPA) to regulate greenhouse gas emissions as part of its responsibilities under the Clean Air Act.[36] One year later, the Ninth Circuit Court of Appeals declared that federal agencies could no longer fail to put a price on carbon emissions, holding that while there may be a range of reasonable numbers for the SCC, "the value of carbon emissions reductions is certainly not zero."[37] Partly as a result of these rulings, the Obama administration in 2009 formed the Interagency Working Group on the Social Cost of Carbon (IWG), which issued its first set of SCC figures in 2010.[38] These figures were later updated in 2013 and 2015, and were used by the Obama administration's EPA to underwrite its Clean Power Plan.[39] However, on March 28, 2017, President Donald Trump issued an executive order that disbanded the IWG, directed federal agencies not to use IWG's SCC results, and instead directed agencies to rely on 2003 guidance (OMB Circular A-4) issued by George W. Bush's Office of Management and Budget.[40] Most recently, an expert panel of the National Academies of Sciences, Engineering, and Medicine (NAS), responding to a previous request for guidance by the IWG, issued a major

[36] 127 S Ct 1438 (April 2, 2007).
[37] *Center for Biological Diversity vs. NHTSA*, 538F.3d 1172, 1200; 9th Cir., 2008.
[38] Interagency Working Group on Social Cost of Carbon, 2010.
[39] US Environmental Protection Agency, 2015.
[40] White House, 2017.

report on how the US's SCC values could be updated and improved in light of developments in the peer-reviewed literature.[41]

The IWG approach improved upon the *status quo ante* in 2010 by projecting SCC values for the year 2015 between $4.70 and $64.90, with a "central value" of $23.80 (in 2007 dollars). In its 2015 update, the corresponding range became $11 to $105, with a central value of $36. Much of the IWG's methodology was in line with what I have described in this chapter. Employing three of the leading "integrated assessment models" (IAMs), the analysis began with descriptions of the "business-as-usual" (BAU) consumption path and then "pulsed" that path with one ton of CO_2. Next, the models estimated the temperature increases consequent on the pulse, and then the annual decreases in consumption (i.e. "climate damages") consequent on the increase in temperature. The resulting stream of incremental damages was then discounted back to the year in which the pulse occurred in order to arrive at that year's SCC value.

In the IWG's analysis, each IAM began with five possible BAU paths. Each path was then pulsed with an extra ton of CO_2 10,000 times, with each of these 10,000 "model runs" selecting at random from a probability density function expressing climate science's best understanding of the complex relationship between greenhouse gas concentration and temperature increase. This exercise yielded 150,000 streams of climate

[41] National Academies of Sciences, Engineering, and Medicine, 2017.

damages (3 IAMs × 5 BAU paths × 10,000 runs), each of which was then discounted back to present values using the three constant discount rates of 2.5 percent, 3 percent, and 5 percent. The 150,000 SCC values for each discount rate were then averaged to yield the IWG's SCC value for that discount rate, and a fourth SCC value (the upper end of the ranges noted earlier) was produced by reporting the SCC value at the 95th percentile using the central 3 percent discount rate.[42]

The follow-up NAS report rightly criticized the IWG for using constant discount rates instead of constructing discount rates using the Ramsey methodology discussed in Section 12.2.[43] Because it is derived from the conceptually fundamental *V*-function, the Ramsey approach is theoretically appropriate and allows for the possibility (unlikely in the short term, but possible in the very long term) that consumption discount rates become negative due to the severe damages imposed by unmitigated climate change. (Consumption discount rates will be negative if g in equation (5) is negative and if the absolute value of the product of g and η is greater than the utility discount rate.) Negative discount rates have the effect of placing *more* weight on future outcomes (relative to equivalent outcomes in the present).

---

[42] Interagency Working Group on Social Cost of Carbon, 2010.

[43] National Academies of Sciences, Engineering, and Medicine, 2017, pp. 18, 169. See Kelleher and Wagner (forthcoming) for a further important correction to the NAS report related to the Ramsey discounting methodology.

Both the IWG and NAS adopt the revealed preference approach to value functions discussed in Section 12.3. In defending its highest constant consumption discount rate (5 percent), the IWG argues that "many individuals smooth consumption by borrowing with credit cards that have relatively high rates . . . [T]he high interest rates that credit-constrained individuals accept suggest that some account should be given to the discount rates revealed by their behavior."[44] And although the NAS rejects the IWG's use of constant discount rates, the NAS agrees that (certainty-equivalent) consumption discount rates should be "consistent, over the next several decades, with consumption rates of interest"—that is, they should be consistent with the consumption rates of indifference that individuals reveal in their everyday consumer behavior.[45] The IWG and NAS therefore both face the objections to revealed preference approaches set out in Section 12.3.

Nevertheless, the IWG and NAS approaches might very well be superior to the approach we are likely to see from the Trump administration. While OMB Circular A-4 was not written with climate economics in mind, its guidance suggests that streams of incremental climate damages should be discounted at two different rates: 3 percent and 7 percent. The 3 percent value, like the IWG's central value of 3

---

[44] Interagency Working Group on Social Cost of Carbon, 2010, p. 19.
[45] National Academies of Sciences, Engineering, and Medicine, 2017, p. 180.

percent, results from a revealed preference approach that calibrates society's consumption discount rate by looking to the rate at which actual individuals are willing to trade present consumption for future consumption (what the NAS calls the "consumption rate of interest"). As noted, the Ramsey-inspired approach is theoretically preferable, but at least in its focus on a 3 percent discount rate OMB Circular A-4 hews closely to the Obama administration's IWG's central value. Things change, however, with Circular A-4's introduction of a 7 percent discount rate.

The 7 percent rate is meant to represent the *opportunity cost of capital*, or the average rate of return earned by private investments. Let us set aside the question of whether the correct rate for this is 7 percent and ask: why should one use the opportunity cost of capital as a discount rate? The answer will be familiar from private investment decisions. If one is choosing between investments A and B, and investment A yields a 6 percent return and B yields 7 percent, then one should choose B—other things being equal and assuming the upfront costs of each investment are the same. A different but equivalent way to arrive at this conclusion is to discount A's future payout at the rate of return offered by the best alternative investment (in this case B), and then to subtract from this discounted value the initial cost of the investment. If the resulting value is negative, that signifies that the best alternative investment (B) is better than the investment under investigation (A). In recommending 7 percent as an acceptable discount rate, OMB Circular A-4 is essentially directing federal agencies to keep an eye on the opportunity costs—i.e. the best alternative investments—of federal projects.

Even if 7 percent is a reasonable estimate of the long-term rate of return on private investment (and there are good reasons to doubt that it is),[46] it is widely agreed in the theoretical literature that this is the wrong way to take opportunity costs into account, especially for projects having long-term effects. The superior way, as indeed Circular A-4 itself notes, is to view the benefits of foregone projects as a *cost*, rather than as a *discount rate*.[47] To the extent that investment in project A displaces project B, project B's annual payouts should be deducted from the annual payouts of A. To work out B's payouts, one will of course have to take into account the rate of return one could have received if one had invested in B instead: first one determines the incremental improvements in consumption made possible by B, and then one discounts those back to a single present value using consumption discount rates, which, as we've seen, Circular A-4 sets at 3 percent. So instead of featuring in the analysis as a discount rate, B's 7 percent rate of return should be used to calculate the consumption gains one foregoes by investing in A instead. As Dasgupta, Sen, and Marglin put it in 1972, the productivity of capital is

[46] Council of Economic Advisors, 2017.
[47] Office of Management and Budget, 2003, p. 33.

relevant because "the present value of the consumption stream forgone when such investment is displaced . . . is relevant, and this present value is relevant as a cost, not as a discount rate."[48]

The method advocated by Dasgupta, Sen, and Marglin carries even more weight in contexts where a project's impacts will be felt for many years to come. That is because OMB's 7 percent method asks us to consider only the upfront costs of making the capital investment. It does not consider the downstream costs associated with whatever productive activity makes the investment's payout possible in the first place. To see this, consider the example that compared A and B in the previous paragraph. Suppose again that B pays investors an annual 7 percent return and that when we discount A's future benefits at 7 percent and then subtract the upfront cost of A's investment, we get a negative number. According to Circular A-4's logic, that suggests B is the better investment. But now suppose that B is an investment in a paper company, and that while the paper company pays out stock dividends at 7 percent per year, the company's pollution creates consumption damages for third parties for decades to come. Surely a full accounting of A's and B's relative merits should take this into account, but Circular A-4's 7 percent method does not do this. Dasgupta, Sen, and Marglin's method does. For the latter directs us to compare each project by working out the *net* incremental

[48] Dasgupta et al., 1972, p. 171.

consumption impact of each, and discounting each net impact stream at the consumption discount rate. This method, which is broadly in line with that recommended by the IWG and NAS, has the virtue of not being unduly swayed by ostensibly high rates of return for investors that involve significant costs down the line for uninvolved parties.

In the end the Trump administration leaned on OMB Circular A-4's 7 percent method to help arrive at quite low SCC values. Almost certainly fearing a judicial ruling requiring it to retain at least *some* SCC value, the Trump EPA combined the 7 percent discount rate with the view that SCC values for US public policy should take into account only *domestic* climate damages. This resulted in a "central" SCC value of $1 (it had been $51 in the last Obama-era analysis).[49] The move from a global to a domestic SCC obviously raises its own economic ethical issues, which I lack the space to discuss here.[50] But all on its own, the use of a 7 percent discount rate was a significant mistake from the standpoint of both economics and ethics.

12.5. Conclusion

Shortly after President Trump's executive order, and in response to the recent NAS report, the venerable nonprofit organization Resources for the Future announced that it will lead "a multi-year, multidisciplinary research

---

[49] See Armstrong (2017). These figures are denominated in 2017 dollars, whereas the final Obama-era analysis used 2007 dollars.

[50] I know of no extensive philosophical discussion of this move. For the most comprehensive *economic* analysis I know of, see Ceronsky et al. (2014, pp. 5–12).

initiative that will advance the NAS recommendations and lead to a comprehensive update of the social cost of carbon estimates, as well as enhance the capabilities of decisionmakers and analysts worldwide who use the social cost of carbon to measure the benefits of emissions reductions."[51] This is heartening news at a time when federal climate policy analysis (and policy) has been anything but heartening. My aim in this chapter has been to assist this and related efforts by illuminating some of the theoretical and philosophical dimensions that require further investigation and discussion among and between economists and philosophers, especially (but not only) in the age of Trump.

References

Adler, M. (2012). *Well-Being and Fair Distribution*. Oxford, UK: Oxford University Press.

Adler, M., Anthoff, D., Bosetti, V., Garner, G., Keller, K., & Treich, N. (2017). Priority for the Worse-off and the Social Cost of Carbon. *Nature Climate change*, *7*(6), 443-449. doi:10.1038/nclimate3298

Adler, M. D., & Treich, N. (2015). Prioritarianism and Climate Change. *Environmental & Resource Econimics*, *62*(2), 279-308. doi:10.1007/s10640-015-9960-7

[51] <http://www.rff.org/research/collection/updating-and-improving-social-cost-carbon>.

Armstrong, A. (2017). Trump Administration Drops Social Cost of Carbon from $51 to $1. *S&P Global Marketplace Intelligence*. Available online at: < https://www.snl.com/InteractiveX/articleabstract.aspx?id=42286458&KPLT=8>

Arrow, K. J., Cline, W. R., Maler, K. G., Munasinghe, M., Squitieri, R., & Stiglitz, J. E. (1996). Intertemporal Equity, Discounting, and Economic Efficiency. In J. P. Bruce, H. Lee, & 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* (pp. 127-144). Cambridge University Press.

Brennan, G. (2007). Discounting the Future, Yet Again. *Politics, Philosophy & Economics*, *6*(3), 259-284.

Broome, J. (1991). *Weighing Goods: Equality, Uncertainty and Time*. Oxford, UK: Basil Blackwell.

Broome, J. (2012). *Climate Matters: Ethics in a Warming World* (1St Edition ed.). W. W. Norton & Company.

Broome, J. (2016). The Well-being of Future Generations. In M. D. Adler & M. Fleurbaey (Eds.), *The Oxford Handbook of Well-being and Public Policy*. Oxford, UK: Oxford University Press.

Ceronsky, M. et al. (2014) Comments on the U.S. Social Cost of Carbon. Available online at: < https://academiccommons.columbia.edu/catalog/ac:182531>

Council of Economic Advisors. (2017). Discounting for public policy: theory and recent evidence on the merits of updating the discount rate. Available online at: < https://obamawhitehouse.archives.gov/sites/default/files/page/files/201701_cea_discounting_issue_brief.pdf>

Dasgupta, P., Sen, A., & Marglin, S. (1972). *Guidelines for Project Evaluation*. United Nations.

Dasgupta, P. S., & Heal, G. M. (1979). *Economic Theory and Exhaustible Resources*. Cambridge University Press.

Dasgupta, P. 1982. “Resource Depletion, Research and Development, and the Social Rate of Discount.” In Discounting for Time and Risk in Energy Policy edited by R. C. Lind, 273-305. Washington, DC: Resources for the Future.

Dietz, S. and Hepburn, C. (2013). Benefit–cost analysis of non-marginal climate and energy projects. *Energy Economics* 40: 61-71.

Frederick, S. (2003). Measuring Intergenerational Time Preference: Are Future Lives Valued Less. *Journal of Risk and Uncertainty*, *26*(1), 39-53.

Greaves, H. (2015). Antiprioritarianism. *Utilitas*, *27*(01), 1-42. doi:10.1017/S0953820814000193

Hausman, D. (2012). *Preference, Value, Choice, and Welfare*. Cambridge University Press.

Heal, G. (2009). The Economics Of Climate Change: A Post-Stern Perspective. *Climatic Change*, *96*(3), 275-297. doi:10.1007/s10584-009-9641-z

Interagency Working Group on Social Cost of Carbon (2010). Social Cost of Carbon for Regulatory Impact Analysis under Executive Order 12866. Retrieved from https://obamawhitehouse.archives.gov/sites/default/files/omb/inforeg/for-agencies/Social-Cost-of-Carbon-for-RIA.pdf

Johansson, P.-O., & Kriström, B. (2015). *Cost-benefit Analysis for Project Appraisal*. Cambridge University Press.

Kaplow, L., & Weisbach, D. (2011). Discount Rates, Social Judgments, Individuals' Risk Preferences, and Uncertainty. *Journal of Risk and Uncertainty*, *42*(2), 125-143.

Kelleher, J. P. (2017a). Pure Time Preference In Intertemporal Welfare Economics. *Economics and Philosophy* 33: 441-473.

Kelleher, J. P. (2017b). Descriptive versus Prescriptive Discounting in Climate Change Policy Analysis. *Georgetown Journal of Law & Public Policy* 15: 957-977.

Kelleher, J. P., & Wagner, G. (Forthcoming). Ramsey Discounting Calls for Subtracting Climate Damages from Economic Growth Rates. *Applied Economics Letters*. Available online at: < https://gwagner.com/ramsey-discounting-calls-for-subtracting-climate-damages-from-economic-growth-rates/>

Kelleher, J.P., & Wagner, G (n.d.). Prescriptivism, Risk Aversion, and Intertemporal Substitution in Climate Economics. Available online at: < https://gwagner.com/kelleher-wagner-prescriptivism-risk-aversion-and-intertemporal-substitution-in-climate-economics/>.

Kolstad, C., Urama, K., Broome, J., Bruvoll, A., Cariño-Olvera, M., Fullerton, D., . . . Jotzo, F. (2014). Social, economic and ethical concepts and methods. *Climate Change 2014: Mitigation of Climate Change*, 173-248.

Llavador, H., Roemer, J. E., & Silvestre, J. (2015). *Sustainability for a Warming Planet*. Cambridge, MA: Harvard University Press.

National Academies of Sciences, E., and Medicine. (2017). Valuing Climate Damages: Updating Estimation of the Social Cost of Carbon Dioxide. Retrieved from https://doi.org/10.17226/24651

Nordhaus, W. D. (1994). *Managing The Global Commons: The Economics Of Climate Change* (31). Cambridge, MA: MIT Press.

Nordhaus, W. D., & Yang, Z. (1996). A Regional Dynamic General-Equilibrium Model Of Alternative Climate-Change Strategies. *The American Economic Review*, *86*(4), 741-765.

Office of Management and Budget (2003). Circular A-4. Retrieved from https://www.transportation.gov/sites/dot.gov/files/docs/OMB%20Circular%20No.%20A-4.pdf

Pearce, D., Barbier, E., & Markandya, A. (1990). *Sustainable Development: Economics And Environment in the Third World*. Hants, UK: Edward Elgar.

Ramsey, F. P. (1928). A Mathematical Theory of Saving. *The economic journal*, *38*(152), 543-559.

Stern, N. (1977). The Marginal Valuation of Income. In M. J. Artis & A. R. Nobay (Eds.), *Proceedings of the Association of University Teachers of Economics, Edinburgh Meeting of April 1976* (pp. 209-258). Oxford, UK: Basil Blackwell.

Stern, N. (2007). *The Economics of Climate Change: The Stern Review*. Cambridge, UK: Cambridge University Press.

Stern, N. (2014). Ethics, Equity And The Economics Of Climate Change Paper 2: Economics And Politics. *Economics and Philosophy*, *30*(03), 445-501. doi:10.1017/S0266267114000303

Nicholas, S. (2010). The Economics of Climate Change. In S. Gardiner, Caney, S., Jamieson, D., Shue, H., & Pachauri, R. K. (Ed.), *Climate Ethics: Essential Readings* (pp. 39-76). Oxford, UK: Oxford University Press.

Sterner, T., & Persson, U. M. (2008). An Even Sterner Review: Introducing Relative Prices into the Discounting Debate. *Review of Environmental Economics and Policy*, *2*(1), 61-76.

White House (2017). Presidential Executive Order on Promoting Energy Independence and Economic Growth. Retrieved from https://www.whitehouse.gov/the-press-office/2017/03/28/presidential-executive-order-promoting-energy-independence-and-economi-1

Traeger, C. (2010). Intertemporal Risk Aversion: or Wouldn't it be Nice to Tell Whether Robinson Crusoe is Risk Averse? *Unpublished ms.*

U.S. Environmental Protection Agency (2015). Regulatory Impact Analysis for the Clean Power Plan Final Rule. Retrieved from https://www3.epa.gov/ttnecas1/docs/ria/utilities_ria_final-clean-power-plan-existing-units_2015-08.pdf von Below, David, Francis Denning, and Nikko Jaakkola. (n.d.). The Climate Pension Deal: An Intergenerational Bargain. Unpublished manuscript. URL = < https://www.nottingham.ac.uk/climateethicseconomics/documents/papers-workshop-2/jaakkola-et-al.pdf>

Weitzman, M. L. (2007). The Stern Review of the Economics of Climate Change. *Journal of Economic Literature*, *45*(3), 703-724.