

Temporal Discounting and Climate Change

Prepared for inclusion in Routledge Companion to Philosophy of Time

J. Paul Kelleher

April 18, 2024

Abstract

Temporal discounting is a technical operation in climate change economics. When discount rates are positive, economic evaluation treats future benefits as less important than equivalent present benefits. This chapter explains and critically evaluates four different reasons economists have given for tying discount rates to the interest rates we observe in real-world markets. I suggest that while philosophers have correctly criticized three of these reasons, their criticisms of the fourth miss the mark. This is because philosophers have not taken heed of the distinct analytical framework in which the fourth reason arises.

A shopkeeper may discount a clothing item, and a hiker may discount the chance of a storm. In the context of climate change economics, discounting is an operation that makes essential reference to time. Like the shopkeeper and the hiker, when an economist discounts, she is typically reducing something in a specific sort of way, for a specific reason. Often, economists' reasons relate to the interest rates that prevail in real-world financial markets. Typically, moral philosophers condemn economic frameworks that tie discount rates to observed interest rates.

This chapter explains how discounting figures in climate economics and critically evaluates four different reasons economists give for calibrating discount rates by reference to observed interest rates.

1 Temporal Discounting in a Ranking Framework

One analytical framework used in climate change economics is centered around the idea of a normative ranking of policy outcomes. Commonly, the outcomes are infinite *consumption paths* of the following form:

$$(c_t, N_t, c_{t+1}, N_{t+1}, \dots)$$

where c_t denotes per capita consumption at time t and N_t denotes the size of the global population at t . “Consumption” refers to the enjoyment of goods and services. (I will use the term “goods” to refer to both goods and services.) Economic models commonly simplify the representation of consumption at a time by expressing it as a single dollar figure. To achieve this simplification, one identifies a pre-specified reference bundle of goods whose total value (at some set of reference prices) is \$1. One then assumes that any additional dollar of consumption will be allocated across these goods in the same proportion as the first dollar, so that the amount of each good signified by \$2 of consumption is double the amount signified by \$1.

The ranking framework's focus on *per capita* consumption is sometimes justified by the (obviously false) simplifying assumption that the world will never exhibit *intra*temporal inequalities in consumption. I shall adopt that assumption here.

The goal of the framework is to take all technologically feasible consumption paths and place them into a ranked ordering. This is done by running each path through a *social welfare function* (SWF), which is a function that assigns a numerical score to each path. These scores are then used to rank paths from best to worst.

The most common SWF in climate economics is the *discounted utilitarian* SWF:

$$W = \sum_{t=0}^{\infty} N_t \cdot w(c_t) \cdot \frac{1}{(1 + \delta)^t} \quad (1)$$

Here w is a *well-being* function that converts a given level of per capita consumption into a corresponding level of well-being. The well-being function is commonly assumed to take the following *isoelastic* form:

$$w = \frac{c_t^{1-\eta}}{1 - \eta} \quad (2)$$

The parameter η determines the degree to which the extra well-being that flows from an extra dollar of consumption declines as an individual's consumption level increases. When $\eta = 0$, an extra dollar yields the same well-being increase regardless of how much a person is already consuming; as η gets larger and larger, the increase in well-being from any additional dollar becomes smaller and smaller.

A SWF defined by equations (1) and (2) assigns a ranking score to a consumption path by taking the per capita level of consumption at each time and finding the level of well-being afforded by the time's per capita consumption level. It then multiplies each time's per capita well-being level by the size of the population at that time. Before adding up all these time-specific well-being aggregates, the SWF weights each by its respective *pure time discount factor*; this is the expression $\frac{1}{(1+\delta)^t}$ in (1). A pure time discount factor is built around the pure time discount *rate*, δ . The value one chooses for δ determines a specific version of what I shall call the "well-being hurdle rate rule":

Begin with an arbitrary baseline consumption path and use (2) to convert it into a well-being path of the form $(w_t, N_t, w_{t+1}, N_{t+1}, \dots)$. Now imagine a policy that reduces period t well-being by one unit and increases period $t + 1$ well-being by z units. For this policy to yield a new consumption path that the SWF ranks at least as highly as the baseline path, z must equal at least $1 + \delta$. That is, the well-being-based return on the intertemporal investment must be at least δ percent. Similarly, for the policy to put the world onto a *higher*-ranked consumption path, the well-being-based return on investment must be greater than—must *hurdle over*— δ percent.

When an SWF is outfitted with a pure time discount factor that features a positive δ , the SWF in effect holds policies that benefit future generations to a higher standard than policies that benefit present generations. If one's SWF has the form of (1), then the SWF will be indifferent between

the status quo and a policy that imposes a cost of one well-being unit on a present person in order to give an extra unit of well-being a different present person; but it will *not* be indifferent between the status quo and a policy that is identical to the first except for the fact that the one-unit well-being benefit accrues to a *future* person. In this way a positive pure time discount rate discriminates against future well-being; and the higher the value of δ , the stronger the discrimination.

The term “social discount rate” does not usually refer to the pure time discount rate. It more commonly refers to what is called the *consumption discount rate*. To explain this second concept, select an arbitrary baseline consumption path. Now suppose we are considering a policy that would reduce aggregate *consumption* by \$1 today while increasing consumption by some amount next year. The consumption discount rate is the rate of return on this consumption investment that would leave the world on a consumption path that is level with the baseline path in the SWF’s ranking. The consumption discount rate can therefore be used to formulate a consumption hurdle rate rule analogous to the well-being-based version stated above.

Discounted utilitarian SWFs of (1)’s form entail a specific formula for computing the consumption discount rate for each time along a given consumption path. This is sometimes called the Ramsey formula (after philosopher and economist Frank Ramsey):

$$\rho_t = \delta + \eta \cdot g_t \quad (3)$$

where δ is the SWF’s pure time discount rate, η is the parameter from the well-being function (2), and g_t is the rate of growth of per capita consumption from t to $t + 1$ along the chosen consumption path. The consumption discount rate, ρ_t , is also a hurdle rate—the rate of return that a consumption investment at t must yield at $t + 1$ for the investment to move the world from status quo consumption path to one that is ranked higher by the SWF. The Ramsey formula shows how consumption discount rates reflect (1) the degree to which future well-being is treated as worth less than present well-being by the SWF (which is determined by δ), (2) the degree to which average consumption is rising over time (reflected in g_t), and (3) the degree to which the marginal well-being of consumption declines with rising consumption (which is determined by η).

Consumption discount rates can be used to construct consumption discount factors, which are analogous to pure time discount factors. The product of a chain of consumption discount factors, one for each year, can be used to weight a consumption change in a future year, in order to determine the change in consumption today to which it is equivalent—equivalent, that is, in terms of its impact on the the score assigned by the SWF. The process is this. Begin with an SWF and an arbitrary baseline consumption path. Then use (3) to construct a time path of consumption discount rates along that path. Next, imagine a policy that would increase consumption by Δ_τ in a future time period τ , and which would cost q dollars in the present period, t . The policy will move the world onto a higher-ranked consumption path just in case the following is true:

$$\Delta_\tau \cdot \left(\frac{1}{(1 + \rho_t)} \cdot \frac{1}{(1 + \rho_{t+1})} \cdot \dots \cdot \frac{1}{(1 + \rho_{\tau-1})} \right) > q \quad (4)$$

In words, (4) says that the policy moves the world onto a higher ranked consumption path just in case the *present value* of the future consumption benefit Δ_τ is greater than the policy’s present

cost, q . To compute the present value of the future consumption benefit, one weights that benefit by the product of the appropriate chain of consumption discount factors.

I have called (3) the Ramsey formula. There is another equation named after Ramsey that is called the Ramsey rule:

$$r_t = \delta + \eta \cdot g_t \quad (5)$$

The Ramsey rule holds along the top-ranked consumption path. Along that path, the economy's rate of return on an investment made in period t , r_t , will equal the consumption discount rate at t . The reasoning is simple: as a hurdle rate, the consumption discount rate gives the rate of return an investment must beat to move the world to a higher-ranked path. When the rate of return equals the consumption discount rate, no more (or less) investment will move the world onto a higher-ranked path. So if the economy is on the highest-ranked path, the economy's time path of consumption discount rates will be identical to its time path of consumption rates of interest.

Some economists working within this ranking framework believe that the SWF should be chosen to yield a philosophically defensible ranking over consumption paths. Examples include Stern (2007) and Dasgupta (2008). On this approach, δ and η are calibrated to reflect whatever a priori ethical considerations bear on a normative ranking of consumption paths. This is also the stance taken by virtually all moral philosophers writing on climate economics (Broome (2012) is a leading example).

Other economists believe that the SWF's ranking should reflect the values of the population whose consumption is at issue. For example, Anthoff et al. (2009, 2) look to "the behaviours of democratically elected governments to infer distributions of these critical factors of Ramsey discounting that are actually used in practice." Their method uses the Ramsey rule, combined with "data on nominal per capita consumption growth rates, inflation rates, and nominal interest rates", to infer values for δ and η from observed behavior (15). That is, they assume that governments have made investments right up until the point at which prevailing interest rates equal their nationally determined discounted utilitarian consumption discount rates. Anthoff et al. then purport to use empirical observations about r_t and g_t in equation (5) to make inferences regarding the values of δ and η .

At best, the method of Anthoff et al. can produce rankings of consumption paths that reflect nations' actual value judgments. But this method is compatible with the alternative philosophical method that aims to construct an ethical ranking of paths that *should* guide nations' policies. In this alternative, today's observed interest will have no conceptual connection to the proper values for δ and η . Many economists, however, hold that a priori ethical inquiry is irrelevant to the proper calibration of those two parameters. Their arguments for this view are usually breezy or unstated altogether; sometimes (e.g. Weitzman (2007, 712)) they make reference to what "the discipline of economics" cares about, as if that alone could render irrelevant what other disciplines study or care about. Even if the revealed intertemporal priorities of actual societies are relevant to public policy, this does not make philosophical ethics irrelevant.

2 Temporal Discounting in a General Equilibrium Framework

In contrast with the social welfare function framework's concern to rank consumption paths, general equilibrium theory is concerned (among other things) to predict outcomes in market economies and to explain how certain putatively desirable outcomes can be brought about through market mechanisms. As explained in the previous section, a SWF embodies a set of normative commitments concerning both the desirability of consumption and its temporal location. Meanwhile, the sole normative standard in general equilibrium theory is *Pareto efficiency*: outcome x is Pareto efficient if and only if there is no feasible alternative outcome y such that at least one person prefers y to x and nobody disprefers y to x .

The notion of Pareto efficiency makes essential reference to people's preferences. In general equilibrium economics, it is common to assume that individuals' preferences can be represented by a numerical function. Such functions are called *utility functions*. A person's preferences over outcomes are represented by a utility function u when, for any two outcomes x and y , $u(x) \geq u(y)$ if and only if x is at least as high as y in the person's preference ranking. If each person in an economy has preferences that can be represented by a utility function, then Pareto efficiency can be defined in terms of these: outcome x is Pareto efficient if and only if there is no alternative outcome y such that $u(y) > u(x)$ for at least one person and $u(y) \geq u(x)$ for every person.

Utility functions can differ radically from the well-being functions used in the SWF framework. First, there might be more to well-being than preference satisfaction. Second, the SWF framework's well-being functions must be interpersonally comparable across individuals in a way that enables benefits to some to be weighed against costs to others. No such requirement applies to the utility functions used in general equilibrium economics: there is no need to be able to say that one person's preferences are more or less satisfied than another's, or that a policy improves one person's preference satisfaction by more than it reduces another's.

In a textbook economy in which there is a competitive market for every desired good and in which there are no pollution-like external effects, no informational asymmetries, and no costs to bargaining, the end result of free exchange between individuals will be Pareto efficient: it will not be possible to better satisfy some person's preferences without imposing a change that is dispreferred by someone else (Varian 1992, 323ff.).

This result does not hold, however, when the economy exhibits pollution externalities. Externalities undermine Pareto efficiency because of missing markets. Take the example of greenhouse gases: there are no markets in which the agents who have rights over atmospheric quality can sell those rights to others who are willing to pay more for them than the rightholder is willing to accept to relinquish the right. Suppose, for instance, that presently living individuals have a right to emit however much pollution they like. Suppose they emit right up until the private benefit of engaging in a bit more of the emitting activity equals the private cost of doing so. Then they may well emit so much that when future people arrive on the scene they would prefer to have inherited a better atmosphere at the cost of being materially poorer. If there had been a way for those future people to signal this preference to present emitters, and if there had been a way to transfer future material goods to present people as payment, present people might have been willing to emit less in exchange for greater material consumption. But since no such market exists, no bargaining takes place. The result will be Pareto *inefficient* if there is indeed a

way to bring about the envisaged intertemporal transfer. I shall assume that there is a way, and that it involves present people taking out a loan to fund more material consumption in exchange for emitting less pollution. The loan would then be carried into the future, with interest, until future people are legally required to pay it off.

A theorem in general equilibrium theory states that the Pareto efficient result of bargaining over pollution (if it were possible) would satisfy the following condition: the cost to consume a “dirty” good would equal the good’s market cost plus the good’s “external cost” (Sandmo 1975). The notion of external cost at work here is easiest to grasp if one assumes that atmospheric rights have been granted to future generations and that we somehow know future people’s preferences. Then we could work out how much control over atmospheric quality future people would be willing to sell to present people, who would like to engage in dirty consumption. Once that amount of control had been sold to present people (by putting the cost into an interest-bearing investment that future people will eventually cash in), the allocation of atmospheric rights will be such that the cost of engaging in a little more dirty consumption today will equal the sum of the market cost of dirty consumption (e.g. the market cost of the car and the gas that together emit greenhouse gases) *plus* the amount that the car buyer would have to pay to future people to acquire a slightly expanded right to pollute. In general equilibrium theory, this second cost is the “external cost” of dirty consumption.

The same result holds when it is present people who are endowed with rights to control the atmosphere. In that case, a faithful trustee of future people would be willing to execute a transfer from future people to present people to induce the latter to scale back their dirty consumption. The inducing will continue until what future people would be willing to pay for a little less pollution is equal to what present people would accept to forego the consumption that generates that extra pollution. At that point, any present person who chooses to engage in dirty consumption pays two costs to do so: the market cost (e.g. the cost of the car and the gas), and the *opportunity cost* of *foregoing* the payment that future people offer to her to refrain from emitting.

Thus when an allocation of goods and rights is Pareto efficient, the cost of dirty consumption will be the sum of the market cost of the dirty good plus its external cost. Because of missing markets, real-world intertemporal allocations are rarely Pareto efficient. But further results in general equilibrium theory entail that even when it is infeasible to create a market for an externality, any given Pareto efficient allocation can be brought about by (1) taxing dirty consumption at an amount equal to the allocation’s external costs of dirty consumption, and (2) rearranging intergenerational wealth with tax-financed transfers of income (Feldman and Serrano 2006, chs. 5, 7).

Typically, when economists use general equilibrium models to estimate the pollution taxes that must be imposed on dirty consumption to bring about a given Pareto efficient allocation, they assume that it is present people who have the right to pollute. This implies that the relevant external costs are future people’s willingness to pay for pollution-reduction. In order to impose these cost on present emitters in the form of a tax, the willingnesses to pay of all people who are harmed by the emissions must be aggregated together into a single amount.

This aggregation must be done carefully, because from the standpoint of a present emitter, the willingness to pay of a person who will live in one-hundred years is not equivalent to the

willingness to pay of a different person who will live in two-hundred years. That is because present people are not indifferent between a dollar today and a dollar in a year's time. Instead, they are indifferent between a dollar today and some larger amount of money next year. After all, a present person can always trade a dollar today for a larger number of dollars next year by investing today's dollar in an interest-bearing savings account. That she does not invest more than she does indicates that, at the margin, she is indifferent between the dollar she can consume today and the larger number of dollars she could consume next year. Her rate of indifference is just the interest rate.

You might already see that the way to aggregate different future people's willingnesses to pay for pollution-reduction—the way to compute today's external cost of dirty consumption—is to discount future willingnesses to pay using interest rates as discount rates, and then to add up the discounted values. Given a present person's preference for present over future dollars, the way for her to compare what she would accept as a bribe not to pollute with future people's time-specific willingnesses to pay is to discount the latter using the intervening rates of interest.

Importantly, the external cost that is relevant to determining the proper tax on dirty consumption is the external cost *at the intertemporal Pareto efficient allocation at issue*. Likewise, the interest rates to be used in aggregating future people's willingnesses to pay are the interest rates that would prevail along that same intertemporally Pareto efficient economic trajectory. Different intertemporally Pareto efficient trajectories will exhibit different interest rate schedules, and very likely none of these will be identical to the schedule of interest rates that prevails at an allocation (such as the *actual* allocation) that is *inefficient* because of the presence of an unpriced externality.

Nevertheless, some climate economists claim that observed interest rates are relevant to the general equilibrium framework. The reasons why are complex, but I can provide a high-level overview.

When certain assumptions are made about the agents in a market economy, as well as the structure of these agents' intertemporal preferences, a theorem in general equilibrium theory entails that every Pareto efficient allocation maximizes a function of the form:

$$W = \sum_{t=0}^{\infty} N_t \cdot u(c_t) \cdot \frac{1}{(1 + \theta)^t} \quad (6)$$

with the utility function u taking the following form:

$$u = \frac{c_t^{1-\sigma}}{1-\sigma} \quad (7)$$

I do not have space here to explain the full derivation; see Kelleher (forthcoming, §3.13). But the central assumption is fairly simple: the economy's consumers are all assumed to have a utility function in which the value of σ expresses the degree to which consumers care less and less about having more of a consumption good, the more of that good they are already consuming; and the consumers (which are typically assumed to live for just one period) also care about what happens to their descendants, with the degree of this care declining exponentially at the rate θ . Thus they care about their children less than they care about themselves, and they care about

their grandchildren less than they do about their children, and so on. Given these and a few other assumptions, all Pareto efficient allocations in the intertemporal economy will maximize a function of (6)'s form. Different versions of (6), characterized by different σ and θ values, will pick out different Pareto efficient allocations.

Now, even though functions of (6)'s form *look* like the social welfare functions from section 1, their role in the general equilibrium framework is radically different. Instead of representing a normative ranking of intertemporal consumption paths, different versions of (6) are merely mathematically useful ways of referring to different intertemporally Pareto efficient allocations (Kelleher forthcoming, ch. 3). Yet since these referring equations are mathematically isomorphic with section 1's social welfare functions, certain formal results concerning the latter also hold for the former. In particular, we know that the Ramsey rule (equation (5)) holds at any allocation that maximizes a function having (6)'s form. This means that the annual interest rates along any Pareto efficient consumption path will equal the righthand side of the Ramsey rule, with the parameters changed from η and δ to σ and θ . Still, these interest rates differ from those we observe in the Pareto inefficient real world. So how might the latter be relevant?

Some climate economists claim they are relevant because (the economists say) inefficient allocations can also be represented as maximizing solutions to equations of (6)'s form. The work of William Nordhaus is one important example. Nordhaus is concerned to identify the Pareto efficient consumption path the world economy would achieve if the present generation had the right to pollute as much as it wanted and if (contrary to fact) there were a well-functioning intertemporal market for pollution abatement (Nordhaus 1994, 9). He then assumes that the function that this Pareto efficient allocation maximizes is identical to the function that would be maximized by the real-world inefficient allocation, if this latter maximization exercise were subject to the constraint that no mitigation occur. Nordhaus imposes this constraint to mimic, in his model, the fact that intertemporal markets for pollution abatement do not exist in the real world.

If these efficient and inefficient allocations both maximize the same version of equation (6), then the Ramsey rule would also hold at the real-world inefficient allocation. In that case, the lefthand side of (5) could be set equal to last year's observed interest rate, and the g_t value on the righthand side could be set equal to the observed rate of growth in per capita consumption from last year to this year. This would allow one to identify pairs of parameters (σ, θ) that solve the Ramsey rule's equation. Then, if one could defend the choice of one of the two unknown parameters, one could then infer the value of the other. This is Nordhaus's calibration method (Nordhaus 1994, 10–12; 2008, 50).

With a fully fleshed out version of (6) in hand, Nordhaus then maximizes it without the artificial constraint on the level of mitigation, which (he says) yields the Pareto efficient consumption path he is after. From here, he estimates the external costs of emissions along the efficient trajectory, and aggregates these together by discounting them at the efficient path's interest rates. He presents the result as the pollution tax that must be imposed on today's emissions if the economy is to be guided from its current inefficient path to Nordhaus's targeted Pareto efficient path. Importantly, Nordhaus neither reports nor recommends the transfers of income that would be needed to guide the economy toward his target Pareto efficient allocation. This is why he portrays the present generation as net losers from efficient climate policy, rather

than the winners they would be if policy mimicked the outcome of well-functioning intergenerational markets for abatement (Nordhaus 2008, 179f.).

Moral philosophers have spent decades now criticizing Nordhaus's approach to discounting, but they have always assumed that Nordhaus works within the social welfare function framework of section 1. If he instead works within a general equilibrium framework, philosophers' criticisms have been misplaced. A more appropriate line of inquiry would have concerned Nordhaus's initial assumption that public policy should target the Pareto efficient allocation that would have resulted, given the prevailing intergenerational distribution of wealth, had there been a well-functioning intertemporal market for pollution-abatement. Nordhaus never really defends this assumption, and several philosophical theories of intergenerational ethics conflict with it. Nor does Nordhaus defend his view that climate policy should focus only on the *emissions reductions* that correspond to that Pareto efficient allocation, rather than the intergenerational transfers of wealth that would be needed, together with pollution taxes, to steer the economy toward that allocation.

3 Two Other Frameworks

Two other common analytical frameworks employ observed interest rates as discount rates. The first is the *opportunity cost* framework. Imagine an environmental policy that would generate a future consumption benefit in period τ of Δ_τ dollars and which costs q dollars in the present period 0. Proponents of the opportunity cost approach recommend that Δ_τ be discounted by a time path of discount factors that use observed interest rates as discount rates. They claim that if the resulting present value of Δ_τ is lower than q , an even larger consumption benefit could be given to people in period τ if q were invested in the money market instead. Alternatively, a consumption benefit of the exact size that is offered by the environmental policy could be given to people in period τ by investing some amount less than q . This shows (proponents argue) that by discounting a policy's prospective benefits using real-world interest rates, one can sometimes identify conventional investments that would leave everyone better off than they would be with the environmental policy (Weisbach and Sunstein 2008).

But there is one big problem: observed interest rates are *private* rates of return on conventional investments. Yet investments that have large private rates of return for investors can also have significant *external* effects on third-parties. The opportunity cost framework ignores these entirely, and therefore cannot reliably indicate when a financial investment really would improve everyone while worsening no one (Broome 1992, 90–91).

Conventional cost-benefit analysis (CBA) offers a final framework in which observed interest rates figure as discount rates. CBA considers a policy a good thing if, after it is implemented, its winners could fully compensate its losers and still be better off than without the policy. Sometimes this test is articulated in terms of *potential Pareto improvements*. A change is Pareto-improving if at least one person prefers it to the status quo and no one disprefers it to the status quo. A *potential* Pareto improvement is a change that *would be* Pareto-improving if the winners in fact fully compensate any losers. CBA does not, however, say that a change must be actually Pareto-improving to be a good thing.

To test for potential Pareto improvements, economists model the outcome of a policy change, and then determine how each person's financial situation would have to change—after the fact—to leave them as well off as they were before the change. In the case of winners, these amounts are their willingnesses to pay for the policy; in the case of losers, these amounts are their willingnesses to accept compensation. Economists then discount all of these willingnesses to pay/accept at observed or projected interest rates, and work out whether the aggregate present value of what winners are willing to pay is more than the aggregate present value of what losers are willing to accept. If it is, the economists claim that the policy constitutes a potential Pareto-improvement.

This use of prevailing market interest rates to convert willingnesses to pay/accept into present values might seem familiar from the discussion of the general equilibrium framework in section 2. In that framework, a divergence between total discounted willingnesses to pay and total discounted willingness to accept indicates that an outcome is not Pareto efficient, which entails that *some* Pareto-improvement is possible. But this is different from saying that whenever the willingness to pay for a policy is greater than its willingness to accept, then *that policy itself* makes a Pareto improvement possible. And in fact, it has been proven that policies can pass a conventional cost-benefit test even when they do not make Pareto improvements possible (Boadway 1974). For this reason alone CBA lacks a cogent foundation. But even if economists' test were a reliable indicator of potential Pareto improvements, no moral philosopher (that I know of) accepts the framework's potential-Pareto-improvement conception of what makes a policy change a good thing. If a rich person is benefited at the expense of a poor person, that outcome is not redeemed by the mere fact the poor person *could* have been fully compensated. This is another reason to question the CBA framework and its rationale for discounting future benefits at real-world interest rates.

*

I have sought to explain and critically evaluate four reasons climate economists give for using observed interest rates to calibrate discount rates. I have suggested that while philosophers correctly criticize three of these, their analysis of the fourth is wanting. This is because philosophers do not realize this fourth reason arises within the context of general equilibrium theory, and not the SWF-based ranking framework.

Further Reading

Portney, P.R., and Weyant, J.P., Eds. (1999) *Discounting and intergenerational equity*. Resources for the Future, New York, NY. (A classic collection of papers on discounting by leading environmental economists.)

Broome, J. (1992) *Counting the costs of global warming*. White Horse Press. (A classic book by a philosopher that addresses many dimensions of discounting.)

Nesje, F., Drupp, M.A., Freeman, M.C., and Groom, B. (2023) Philosophers and economists agree on climate policy paths but for different reasons. *Nature Climate Change*, 1–8. (A paper reporting survey results on economists' and philosophers' beliefs about the Ramsey discounting parameters.)

Related Topics

Time Biases — Meghan Sullivan
Time and Well-being — Eden Lin

References

- Anthoff, D., Tol, R.S., and Yohe, G. (2009) Discounting for climate change. *Economics: The Open-Access, Open-Assessment E-Journal*, Vol. 3.
- Boadway, R.W. (1974) The welfare foundations of cost-benefit analysis. *The Economic Journal*, 84, 926–939.
- Broome, J. (1992) *Counting the costs of global warming*. White Horse Press.
- (2012) *Climate Matters: Ethics in a Warming World*. W. W. Norton & Company.
- Dasgupta, P. (2008) Discounting climate change. *Journal of Risk and Uncertainty*, 37, 141–169.
- Feldman, A.M., and Serrano, R. (2006) *Welfare Economics and Social Choice Theory*, 401 p. Springer Science & Business Media.
- Kelleher, J.P. (forthcoming) *The social cost of carbon: ethics and the limits of climate change economics*. Oxford University Press.
- Nordhaus, W.D. (1994) *Managing the global commons: The economics of climate change*. MIT Press, Cambridge, MA.
- (2008) *A Question of Balance*. Yale University Press.
- Sandmo, A. (1975) Optimal taxation in the presence of externalities. *The Swedish Journal of Economics*, 86–98.
- Stern, N. (2007) *The Economics of Climate Change: The Stern Review*. Cambridge University Press, Cambridge, UK.
- Varian, H.R. (1992) *Microeconomic Analysis*, 506 p. Norton.
- Weisbach, D., and Sunstein, C.R. (2008) Climate change and discounting the future: a guide for the perplexed. *Yale Law & Policy Review*, 27, 433–457.
- Weitzman, M.L. (2007) A review of the Stern Review on the economics of climate change. *Journal of Economic Literature*, 45, 703–724.

WORD COUNT (excluding Title, Name, Date, Abstract, and Related Topics): 5,480