

A Philosophers' Guide to Discounting¹

[Final draft for forthcoming OUP volume *Philosophy and Climate Change* (eds. Mark Budolfson, Tristram McPherson and David Plunkett). Feel free to cite with that information. You can contact me at mintzwoo@princeton.edu. Thanks!]

Abstract: I introduce several distinctions relevant to what is called the *discounting problem*, since the issue is how (future) costs and benefits are “discounted” to make them comparable in present terms. I defend the claim that there are good reasons to adopt Ramsey-style discounting in the context of climate change. I discuss an important distinction involved in discounting: *descriptivism* and *prescriptivism*. I argue that, even if we adopt prescriptivism, and accept that this means there is need for moral experts in parameter assignments, there is a significant issue. The type of moral expertise required for the discounting problem will not involve knowledge of moral theory—thus making moral philosophy unhelpful in terms of making particular parameter assignments, despite these being substantive moral judgments.

One of the challenges of climate change lies in the temporal distances between policies and effects. This is not only because of the lag between emissions and effects of greenhouse gases on the climate system but also because some of these gases are very long-lived and continue to have effects over long timespans. For these reasons, one prominent focus of climate ethics has been on various problems in intergenerational ethics.² My discussion will focus on the distinctions relevant to what is called the *discounting problem*, since the issue is how (future) costs and benefits are “discounted” to make them comparable in present terms.³ More precisely, the discounting problem is about which weightings to place on those costs and benefits as a function of time.⁴

An important moment in the history of discounting was the release of the so-called “Stern Review” of climate change, where Nicholas Stern (2006) argued that moral or normative justifications needed to be brought into long-term cost-benefit climate analyses. In particular,

¹ Thanks to Francis Dennig, Eike Düvel, Maddalena Ferrana, Marc Fleurbaey, Lukas Meyer, Paul Kelleher, Jeff Sebo, Karl Steininger, Frej Klem Thomsen, and to the editors of this volume for their extensive and helpful commentary. I also appreciate comments from audiences at the Mercator Research Institute on Global Commons and Climate Change, especially commentators Martin Kowarsch and Linus Mattauch, and at the University of Graz. This work was funded by the Austrian Science Fund (FWF) under research grant W 1256-G15 (Doctoral Programme Climate Change Uncertainties, Thresholds and Coping Strategies).

² Sometimes, we frame these problems in terms of models with overlapping generations, but here I am less concerned with the generational aspects and more concerned with temporal ones.

³ Helpful recent surveys of this question can be found in Davidson 2015, Greaves 2017, Heal 2007, and Kelleher 2017a.

⁴ Note that discounting in this context is a very technical economic term. It is, for instance, completely unrelated to the everyday usage of the term “discounting”, which is about *lowering prices* (such as in retail). Also, note that although the weightings are a function of time in the technical sense—they are a function of a variable t for time—the justifications for different weightings may be about things which are only correlated with time (see Parfit 1984: Appendix F).

he held that a form of utilitarian impartiality implied that the pure rate of time preference should be (near-)zero.⁵ This led to a spate of objections (Dasgupta 2008, Nordhaus 2007a, 2007b, Weitzman 2007), many of which criticized Stern's introduction of moral considerations. However, there are some important subtleties that are often mistaken about this debate. The main intention of this chapter is to introduce and explain the most important distinctions with respect to the discounting problem for a philosophical audience. A secondary intention is to defend the theory underlying this problem and to point out limitations of moral theory when addressing this problem.

In Section 1, I defend the claim that there are good reasons to adopt Ramsey-style discounting in the context of climate change. I emphasise that Ramsey-style discounting need not be intrinsically wedded to utilitarianism. In Section 2, I introduce two important distinctions involved in discounting. The first is between *pure discounting*⁶ and what I call *impure discounting*⁷ (i.e. between discounting in utility/welfare terms and discounting in monetary terms). The second is between *consumption-based* and *investment-based* discounting (i.e. depending on whether we are measuring the costs and benefits in terms of displaced social value or in terms of financial flows). In Section 3, I distinguish between *descriptivism* and *prescriptivism*. The descriptivist about discounting assigns values to the relevant parameters on the basis of observed data grounded in stated or revealed individual preferences; the prescriptivist assigns values to the parameters on the basis of factors beyond stated or revealed individual preferences, including social preferences and considerations from moral theory. In Section 4, I argue that, even if we adopt prescriptivism, and accept that this means there is need for moral experts in parameter assignments, there is a significant problem. The type of moral expertise required for the discounting problem, I argue, will not involve knowledge of moral theory—thus making moral philosophy unhelpful in terms of making particular parameter assignments, despite these being substantive moral judgments. In Section 5, I conclude with a few thoughts about which kind of moral expertise might be relevant to the discounting problem.

⁵ In particular, he said that the pure rate of time preference should be positive not because future outcomes are weighted less heavily, but for probabilistic reasons having to do with the uncertainty of exogenous human extinction. However, properly speaking, human extinction is just a sub-class of a broader class of exogenous events where the action does not generate the intended outcome, so Stern's reasoning should actually generate a higher pure rate of time preference than whatever is implied merely by the probability of human extinction (Mintz-Woo 2019).

⁶ Pure discounting is also known as "pure rate of time preference" or "utility discounting".

⁷ This is a term which I am introducing. Usually we distinguish pure discounting with discounting in consumption terms. However, since I also want to draw attention to the distinction between investment-based and consumption-based discounting, I suspect that using standard locutions like "discounting in consumption terms" and "consumption-based discounting" will be unnecessarily confusing, since these are entirely different things. Since "impure" discounting is the obvious opposite of "pure" discounting, I expect it to be an intuitive contrast. For prioritariness, there will be another level for something like "super-pure discounting" which involves the translation between utility/welfare and social value, an even purer concept.

1. Introducing and Defending the Ramsey Rule

Integrated assessment models (IAMs) perform an important role in climate economics; they allow us to do things like assess the value of particular policies by combining (very) basic climate models with various macroeconomic measurements. As mentioned before, various aspects of the discounting problem are of particular importance in climate policy evaluations as the costs and benefits of climate action are temporally dispersed. The results of IAMs are thus quite sensitive to the parameters which govern time preference.

The key to time preference on standard models is the Ramsey Rule.⁸ The Ramsey Rule involves discounting in terms of the consumption of individuals. It governs the optimal consumption path with a particular social welfare function.⁹ It can be stated in a very compact form which relates the *social discount rate* r on the left-hand side to what is sometimes called the *social rate of time preference* on the right-hand side. The social rate of time preference is the sum of the *pure rate of time preference* (the Greek letter δ , “delta”) and a term which is the product of the *growth rate of consumption* g and the *elasticity of marginal utility of consumption* (the Greek letter η , “eta”):

$$r = \delta + (\eta \times g)$$

Intuitively, discounting of consumption depends on both pure discounting, that is, factors which are not particular to consumption but time itself (or things correlated with time) (δ) and the diminishing marginal value of consumption as consumption grows ($\eta \times g$).¹⁰ It may also be helpful to think of pure discounting as discounting applied in the pure units of utility or welfare, i.e. as a utility discount rate. To give a sense of potential values, δ is usually taken to be between 0 and 0.015 (i.e. 0 to 1.5% per year); η is usually taken to be anywhere between 0.5 and 4; and g is usually taken to be between 0.01 and 0.03 (i.e. 1 to 3% per year).¹¹

⁸ Derivations can be found in Dasgupta 2008 and Goulder and Williams 2012. Note that the Ramsey Rule is a *conclusion* following from various assumptions.

⁹ Beyond the consumption path being optimal, the Ramsey Rule is optimal also in a second sense that it does not describe actual—rather, idealised—economies, due to inter alia market imperfections and distortions (Groom et al. 2005). Many of the issues, such as the divergence between descriptivism and prescriptivism, depend on the assumption we do not live in an optimal economy. Since we do not, I take these issues to be live.

¹⁰ As Kelleher (2017a) points out, it may be “highly misleading” to take $\delta > 0$ to be due simply to the fact that some cost or benefit is in the future. Beyond being justified because there is uncertainty about the future, it may also be justified by considering the demandingness of $\delta = 0$ for current people. Some philosophers take discounting which is particular to time *itself* to be the only real types of discounting. However, this is incorrect; discounting is simply differential weighting. In this context, most discounting is indexed by time, but this need not be the case. Caney (2009) and Mintz-Woo (2017), for instance, point out that discounting for uncertainty is also a form of discounting, a form which inherently has nothing to do with time.

¹¹ The growth rate is the change in consumption given as a percentage of current consumption, that is, $g(t) = \frac{\dot{c}(t)}{c(t)}$, where the dot indicates the first derivative, but if we take some long-term stable growth rate, then we can drop the argument of g (i.e. t).

The normative parameters are taken to be δ and η , with r being of—so to speak—derivative normative importance.¹² The former factor, δ , governing pure discounting, is much more intuitive and has held the bulk of philosophical attention (Broome 1994, Caney 2014, Parfit 1984, Sidgwick 1907/1981). However, given a growth rate g of more than 0.01 (or 1% growth per year), changes in η affect r more than changes in δ , since η is *multiplied* by g . So η is in general a greater driver of the social discount rate than δ . It deserves philosophical attention to a greater extent than it has received.

While uncertainty about future growth (i.e. appropriate values for g) is also of importance, I take it to be of less distinctly *ethical* importance; at least in principle, uncertainty about consumption growth can be eventually resolved given sufficient empirical (macro)economic data, which distinguishes this parameter from η and δ in social contexts, a point I will return to below.¹³

The parameter η governs utility functions given that they increase in a particular predictable manner with respect to increases in consumption.¹⁴ By “predictable”, I mean that η does not vary throughout the utility function as we change consumption levels; a 1% increase in consumption is worth the same *relative* utility gain regardless of whether you started with a lot of consumption or a little.

The main alternative in discounting to the Ramsey Rule is setting the social discount rate to equal the (highest) rate of return on investment, i , focusing on the left-hand equation:

$$i = r = \delta + (\eta \times g)$$

Sometimes, these alternatives are contrasted by calling the Ramsey approach “consumption-based discounting” as opposed to this “investment-based discounting” (or “opportunity cost discounting”).¹⁵ In an ideal economy, both of these approaches will converge because each of these three terms converge. However, in a non-ideal economy (e.g. with distortions due to incomplete markets and taxes), we can either take the left-hand equation ($i = r$) or the right-hand equation ($r = \eta \times g$). The Ramsey Rule, or consumption-based, discounting adopts the right-hand side whereas investment-based discounting depends on the left-hand side equation. I expand on this distinction in Section 3.

However, before explaining how we can assign values to the normative parameters in the Ramsey Rule in the following section, I will close this section by discussing three reasons to

¹² In other words, we can say the social discount rate r is not itself of direct ethical importance, because any value of r underdetermines what the normative parameters are.

¹³ For more on risk and uncertainty, see **Maddalena's chapter**.

¹⁴ Technically, that utility functions are written as a function of consumption in a special *isoelastic* or *constant relative risk aversion* (CRRA) form that looks like $U(c) = \frac{c^{1-\eta}}{1-\eta}$ (with $U(c) = \log c$, for $\eta = 1$), which has the pragmatic implication that derivatives of the utility function are easy to work with.

¹⁵ The Ramsey approach is called “social welfare equivalent” discounting, as opposed to “finance equivalent” discounting, by Goulder and Williams (2012).

work with the right-hand equation in the context of (climate) public policy evaluation. In other words, here are three reasons for adopting consumption-based discounting. First, it is robust in that it can be generated under different ethical approaches, such as intuitionism, not just discounted utilitarianism (Dasgupta 2005). It is also flexible; for instance, the Ramsey Rule can be “extended” or modified to include terms for certain types of risk or uncertainty (e.g. Fleurbaey et al. 2019, Groom et al. 2005) and can model situations where consumption growth is negative. Finally, after having been through several waves of theoretical work, the framework is well-understood and is required in many sources of public policy. For instance, in the United States, regulations have to be subject to cost-benefit analyses, and these require some type of discounting, at least some justifications of which are grounded in a version of the Ramsey Rule. So three reasons that we should work within this framework are that it is robust, flexible, and well-understood. I present these three reasons just in order to make the case; I lack the space here to fully address competing reasons for investment-based discounting although I do not think that similarly convincing reasons can be presented in turn.¹⁶

Let me expand on these three points in turn. The Ramsey Rule is robust in that it arises from different moral theories. In particular, there are both, in the Rawlsian parlance, teleological and intuitionist ways of generating the Ramsey Rule (Rawls 1971). For Rawls, teleological moral theories took the good to be conceptually prior to the right, with the right defined in terms of increasing the good in some manner. Intuitionists rely upon putatively self-evident or prima facie principles. As Ramsey was of the traditional British utilitarian tradition, it is unsurprising that the Ramsey Rule accords with a teleological approach. However, as Dasgupta (2005) points out, the Ramsey Rule also applies under intuitionist approaches to moral theory. As well as deriving this framework from an approach which is concerned with maximizing utility, we can also develop it from an approach which is concerned with putatively self-evident axioms used for comparing consumption streams. In particular, Koopmans (1960) showed that certain axioms imply a social welfare function that looks like Ramsey's (pure) discounted utilitarianism—and implies discounting which complies with something like a Ramsey Rule. Koopmans was interested in ranking infinite streams. The axioms used on the rankings of those streams are part of an intuitionist approach to ethics if we consider those axioms to be subject to evaluation by ethical intuitions, whereupon we work out the implications of adopting those particular axioms.¹⁷ If we interpret the streams as over consumption, then the functional form implied by these axioms has some monotonically increasing function of consumption. (If we interpret the streams as over utility, then Koopmans' axioms yield some type of prioritarianism!) It is worth noting that this differs from the project of many intuitionists, who were concerned both with self-evident or prima facie duties or

¹⁶ Interesting questions arise when moving beyond discounted utilitarianism. For instance, if you disentangle risk aversion and inequality aversion and introduce variable populations, then the results can diverge quite dramatically. I lack the space to explore these non-standard options but refer the interested reader to Fleurbaey (2014) and Fleurbaey et al. (2019) for more. Another alternative is restricting the scope of the evaluation over space, by limiting impacts nationally instead of globally. Interesting moral issues arise here as well (Mintz-Woo 2018b).

¹⁷ The scientific interest of Koopmans' results is that the axioms required to generate these social welfare functions are surprisingly weak.

rules—not axioms—and with the ways these could conflict—not the implications that could be drawn from them. However, if we take axioms to be subject to the same kind of tests of self-evidentness as the intuitionists employed, this method also generates similar functional forms to a teleological approach, making the results more robust to moral methodology than is commonly recognised. It is also appropriate for theories which include non-welfarist considerations, depending on what we take to be what Kelleher (2017b) calls the “evaluative scope”. For instance, we can adopt this framework and modify recommendations with additional side-constraints from the point of view of justice or deontology. For these reasons, Ramsey-style discounting is robust with respect to a variety of potential moral commitments.

The second reason to adopt Ramsey-type discounting is its flexibility. The basic Ramsey Rule applies under deterministic conditions and without distinctions for different individuals. These are simplifications that can reasonably be criticised as incorrect in reality. However, they are not intrinsic to the general framework and we can extend the framework to relax these conditions (obviously, with costs incurred for simplicity and computability) (Fleurbaey et al. 2019). To give a sense of potential extensions, I describe a couple here. Extensions to Ramsey discounting have been made to accommodate risk over consumption growth (e.g. Groom et al. 2005: formula (28)) as well as having differential intratemporal and intertemporal discounting (e.g. Dasgupta 2012). The former introduce a term which lowers the social discount rate r by appealing to risk over discount rates; the latter distinguishes between intragenerational discounting (which is subject to generational preference and is not ruled out) and intergenerational discounting (which is thought to be morally impermissible and set to 0). While the resultant formulations are, unsurprisingly, ungainly, they are still within the same general framework, showing how flexible it can be.

Besides extending the Ramsey Rule, we can generate very different requirements simply by varying the values of the parameters in the Ramsey Rule itself. For instance, undiscounted utilitarianism can be modelled as a special case of discounted utilitarianism simply by setting $\delta = 0$. We can also model different intertemporal approaches by considering extreme values of η : we can demand that consumption streams must only contain maximised (purely discounted) consumption ($\eta = 0$) or that they must be arranged such that we maximise minimum utility ($\eta = \infty$).¹⁸ It is also worth pointing out that the Ramsey Rule applies even in cases where expected consumption growth is negative (i.e. that there is expected consumption decline and $g < 0$). In such cases, the marginal value of current consumption is less than that of future time periods (where consumption is lower), which would be reflected in $\eta \times g < 0$, decreasing the social discount rate (r). Intuitively, if future generations are poorer, then utilitarianism tells us that we should value their consumption more highly than our own given that they will be poorer. This can lead to negative social discount rates, indicating a bias towards the future and away from the present. Because of the range of views that can be formulated in this framework, the framework's flexibility allows for a straightforward conversation while avoiding the challenges inherent in translating positions originating in different frameworks.

¹⁸ Solow (1974) was among the first to point out that Rawlsian distribution can be given as a limit case of utilitarianism where $\eta = \infty$.

The final point is that the Ramsey Rule is well-understood. This applies both to the assumptions that are required to generate it, as well as economic methods for gathering data that might be relevant to assigning values to the normative parameters within it. Unlike the other two reasons, this is meant to be a pragmatic reason; it is easier to act within a well-understood framework which also forms the basis of climate policy analysis. Because we know the implications of changing different parameters, our moral judgments can be more informed. Simple intuitions about distributions, whether over risks or equity, lack this coordinated aspect. By considering risk or equity in isolation from the intertemporal trade-offs, such connections can be obscured. For instance, we are used to thinking of distribution as zero-sum; with a particular quantity of goods, we can distribute them more or less equally (for instance). When considering goods over time, however, we have to take into account potential growth of invested capital. This is considerably less intuitive and, while not the only way, Ramsey-type discounting can help structure these judgments.

2. Why the Terms of Measurement Matter

When considering the Stern Review (2006), people often point to his moral claim: that future people (more specifically, pure benefits accruing to the future) should be counted equally to our own. This implies that *pure discounting* (δ) should be near-zero.¹⁹ The pure rate of time preference *is* a kind of discount rate; a positive value for it means that future benefits (respectively, costs) are to be counted less than the present, and even less the further in the future they are. It is a *pure* discount rate. However, it is not the only kind of discount rate.

In contrast, William Nordhaus, an influential economic modeller who rejected this introduction of moral judgments, suggested that the social discount rate (r) should reflect (risk-free, real) market interest rates. He straightforwardly adopted the investment-based discounting approach since he thought of climate policies as investments and, as such, they needed to be competitive with other available investments. Since there are often well-paying investment opportunities in growing economies, his values for the social discount rate (r) were relatively high (and certainly much higher than near-zero rates), in the range of 4% to 6%.

From this discussion, non-experts may draw the conclusion that moral judgments (of Stern's type) lead to social discount rates that are near-zero and market appeals (of Nordhaus' type) lead to rates that are relatively high. But this is a conflation, since the rates that Stern was discussing using these moral judgments are not the same *kind* of rates that Nordhaus is considering; Stern was talking directly about δ whereas Nordhaus about r . Although δ and r are both discount rates, the former applies to (pure) costs and benefits (i.e. as measured in utility/welfare) whereas the latter applies to consumption equivalents. For this reason, the relatively high values of the social discount rate that Nordhaus advocates are not directly comparable to the near-zero values of the pure rate of time preference that Stern advocates.

¹⁹ The value actually suggested was $\delta = 0.001$, for an assumed 0.1% annual risk of extinction. For our purposes, we can consider this zero or near-zero.

To make comparable claims, as Broome (2008) points out, Stern advocated $r = 1.4$, for $\delta = 0.001$, $\eta = 1$ and $g = 0.013$.²⁰ This is, of course, lower than Nordhaus' r values, but Stern's near-zero δ value is only a component of Stern's r , as is clear from the Ramsey Rule. The difference in application is what is meant to be suggested by the term “pure” as opposed to what I am calling “impure”. Unfortunately, however, the economic terminology is rather opaque, so what the “purity” in “pure rate of time preference” refers to may not be obvious. Furthermore, this opacity is compounded by the fact that the term “social discount rate” is not helpful as both rates may be social. The names are more reflective of historical contingencies than transparent labelling.

Here is another way of making this point: if we believe that the marginal value of consumption gains decline the more initial consumption one has (that consumption generates “decreasing marginal utility”) and that there will be consumption growth, then one should be in favour of a positive (impure) social discount rate *even if one thinks there should be no pure discounting*. Why is this? Formally, for $r = \delta + \eta \times g$, if $\eta > 0$ and $g > 0$, then $r > 0$, even given $\delta = 0$. Intuitively, if we consider how to distribute extra consumption between (by assumption) on average richer future people and poorer present people, we should prefer distributing it to the present. Due to the relative poverty of the present, the marginal utility generated by additional consumption would be greater for present people than future people. In other words, this argument shows we should discount consumption, or have positive impure discount rates even if we do not accept any pure discounting at all.²¹

So what is the important difference? Moral philosophers will be more familiar and used to working in what I'm calling pure units, such as utility/welfare, whereas economists tend to work with impure units. There are important reasons for these differences; pure units are of immediate moral importance, whereas impure units are more readily observable and therefore more amenable to empirical treatment. Furthermore, moral philosophers usually assume that we can talk about pure units in cardinal terms; that is, we can make well-defined comparisons between ratios of interpersonal utility levels. This is a significant assumption; if we do not make it, then it makes more sense to talk in impure units, since these are observable and comparable. This is one of the reasons that modern economics tends towards discussion only in terms of increasing returns in financial terms. Since the advent of positive economics, economists have tended to stay away from such assumptions (Dennig 2018). The relevance for this difference in the context of climate economics is that philosophers focus on pure discounting (which naturally apply to pure units) while neglecting the importance of impure rates which apply to, for instance, consumption discounting. Much of this focus by economists is due to the fact that we cannot empirically determine the cardinality of pure units (only their individual ordinality). In contrast, philosophers are not required to restrict themselves to the empirically verifiable, so are free to talk in pure units. This practice is justifiable in moral philosophy, but it is important to remember that it requires a contentious assumption.

²⁰ To contrast, Nordhaus (1994) takes $\delta = 0.03$ and $\eta = 1$ and Cline (1992) takes $\delta = 0$ and $\eta = 1.5$.

²¹ Recall that the opposite result, negative discounting (weighting future consumption *more heavily* than present consumption), occurs if we expect consumption levels to decline, i.e. if $g < 0$, given $\eta > 0$, $g \times \eta < 0$.

However, keeping the philosophical discussion attuned only to pure units has an important cost. As discussed in the previous section, δ is only one component in r , the social discount rate. Moreover, the pure rate of time preference is a relatively *small* component, since $g > 0.01$ implies that changes in the elasticity of marginal utility of consumption, η , are more important. In other words, r is driven more by changes in η than in changes to δ . For these reasons, it is important to pay attention to the units that a given discount rate is being applied to, since the meaning changes dramatically and pure discount rates are (less important) drivers of impure social discount rates. One might worry that η is merely an empirical unit that does not lend itself to theory but I argue below both that there are reasons to take it to be moral and that there are interesting things that can be said about η if one takes it to be moral.

The next point to make is about the role of “consumption”. Since there are a variety of goods that one can consume in order to satisfy one's preferences, a complete model would incorporate heterogeneity in goods. However, for tractability in modelling contexts, we often assume that any goods have equivalents in terms of a single homogenous good called “consumption” (conventionally represented by the variable c). There are important difficulties here, since climate change may differentially affect various categories of goods—in particular, natural goods might not grow as quickly under climate change as artificial goods, so we might need to take the relative prices of natural goods into account (Stern 2008). This debate is sometimes conducted in terms of “strong” versus “weak” sustainability, depending on whether one thinks that there are some goods that cannot be (or should not be) converted into consumption equivalents, or whether one thinks that such goods are worth large amounts of consumption equivalents. A more philosophical objection is that goods might not be tradable against each other in terms of a single homogenous good, since they might have values which are agent-relative or intrinsic. So, for instance, a given object's particular history might make it matter in agent-relative terms (e.g. this cutlery was passed down from my parents) or it could be that some goods are valuable independent of anyone's preferences (e.g. an ecosystem's dynamics being maintained might matter even if no one were aware of them). One version of such claims is that one cannot replace the good without losing that historical or personal connection, which is true, but it does not follow that there are no replacements which would be of equivalent value to the person. I would argue that one can develop connections or history with replacement goods as well, although of course this may cost time and involve regret, things that should be factored in. These are important and complicated discussions, but I wish to set them aside here, since they are not directly linked to the discounting problem. Goulder and Williams (2012: fn 9), for instance point out that, in theory, introduction of different goods which are imperfect substitutes can be incorporated into what is gained through mitigation using a single discount rate. Intuitively, we can get equivalent results by assuming that the (relative) value of mitigation benefits rise over time, showing that the problem of multiple goods can be isolated from the discounting problem.

In the previous section, in the context of public policy evaluation, I advocated consumption-based discounting. This is because the purpose of public decision-making is to increase social *welfare* and not to maximise social value *in monetary* terms. However, this can also be

bolstered by the descriptive claim that, since national income primarily comes from tax revenue, and that revenue primarily displaces consumption, consumption-based discounting is appropriate (Arrow 1999b). However, if one wishes to adopt a(n impure) social discount rate which is a weighted average of consumption-based and investment-based discounting, this discussion can be framed as discussing the consumption-based component of the weighted average. After all, the investment-based discounting component is of considerably less philosophical interest and should be approximated by market rates of return in well-functioning markets.²² Since the social welfare function is an obvious candidate for moral discussion, it is worth considering how moral considerations can be adopted in consumption-based discounting. In the next section, I argue that there are surprising difficulties in doing so.

3. Descriptivism and Prescriptivism in Discounting Methodology

An important distinction between families of discounting methodologies is between *descriptive* and *prescriptive* methodologies (or *descriptivism* and *prescriptivism*, sometimes also referred to as “positive” and “normative” economics, respectively). In the context of discounting and climate economics, this distinction comes from Arrow (1996), following Manne (1995). Descriptivism has been the dominant historical consensus in economics from the late 1970s (Dennig 2018). However, the theoretical differences potentially encompass welfare economics whenever there is the possibility of utility being construed in either normative or positive fashions. One of the important claims in this section is that both are motivated by some form of impartiality; descriptivism about impartiality between revealed (or stated) preferences of different individuals and prescriptivism about impartiality between temporally dispersed pure costs and benefits. Note that, in this context, the term “preference” is very broad, and need not be confined to what philosophers sometimes call *mere* preferences; in other words, preferences may be informed by moral values, be other-regarding, and/or be highly consequential.

The first thing to recognise is that the distinction between descriptivism and prescriptivism is kept even if we accept the distinction between investment- and consumption-based discounting (although the conjunction of prescriptivism and investment-based discounting would be hard to defend). As I have argued for consumption-based discounting here, it is worth considering how the methodologies would differ between descriptivism and prescriptivism.²³ The claim I argue for is that, no matter which way we go, there will be an important role for moral experts.

²² I say “approximated” because the market rates of discount need not converge to observed returns on investment.

²³ Kelleher (2017b) points out that we can distinguish between descriptive approaches which are concerned with ranking potential outcomes in terms of a social welfare function (which is what I have been calling consumption-based discounting as opposed to investment-based) with those that are simply concerned with searching for potential Pareto improvements over the status quo. While this is in fact a distinction between outputs, I do not see the need here to be so fine-grained about distinguishing methodologies.

If considering prescriptivism in consumption-based discounting, the relevant questions are to what extent society *should* discount future consumption. Stern (2006) and Cline (1992) take moral considerations to imply that $\delta \approx 0$, based on temporally impartial principles possibly stemming from classical utilitarian assumptions. Insofar as these are moral considerations, these judgments clearly require some type of (moral) expertise.

If considering descriptivism in consumption-based discounting, the relevant questions are to what extent society *does* value greater over lesser consumption. Sometimes this is phrased as “inequality aversion”, but more properly we should think of it just as the generated utility. If the actual utility generated by those with higher income from a given marginal increase in consumption is lower than those with a lower income—i.e. if consumption has decreasing marginal value, then we do want to allocate it in such a manner that it does generate more utility.

With the advent of positive economics in the twentieth century, it became clear that there was no empirically objective manner with which to evaluate such interpersonal comparisons of utility (cf. Dennig 2018 for a discussion of this historical debate). In the context of standard IAMs, the problem is even more acute since any decreasing marginal value is captured by a single (constant) parameter, η .²⁴ So there is a very blunt parameter which governs this phenomenon which in practice could be very complicated.

However, there are very good pragmatic reasons to take this parameter to be constant: keeping it constant makes the calculations tractable and making it a function would generate too many degrees of freedom for the modeller without clear benefits in accuracy. For instance, if the function was non-differentiable everywhere, that would make analyses far more complicated. Also, it would require determination of which kind of function to adopt (e.g. square root, cube root) as well as a variety of data points to fit curves to.

In this constant, but tractable, form, η is not directly connected to observation. It is highly unlikely that society has preferences over consumption that can be modelled by a constant elasticity of marginal utility of consumption. On the one hand, we can take the modelling as tractable *approximations* of individual preferences. On the other hand, by adopting these simplified forms, we divorce parameter assignments to η from straightforward observed preferences. In other words, in theory the parameter assignments could be directly responsive to empirical results but in practice they are approximations—justifiable because computationally tractable, but approximations, nonetheless. So my claim is that, even under descriptivism, we require (expert!) judgment in order to amalgamate and approximate the empirical information. This means that descriptivist motivations also justify, albeit in a more roundabout way, a demand for expertise. Thus, the first parallel between these methodological families is both require expertise.

²⁴ By constant, I mean that it does not vary depending on the income level; it is not expressed as a function of consumption, which it could theoretically be.

There is another interesting parallel between the families. If we are being charitable, both sides are appealing to impartiality, but for the descriptivist it is impartiality over (existing) people's preferences and for the prescriptivist it is impartiality over (pure) costs and benefits.²⁵ Since, in fact, people's preferences are not temporally impartial, we have a conflict between the two. Many people's preferences exhibit pure discounting (and often even *hyperbolic discounting*, which means significant discounting near the present and diminishing or even no discounting in the future (Frederick et al. 2002)), which can be exemplified by *temporal myopia* in the individual case. In other words, many people prefer that pleasant experiences occur closer to the present, even when this comes at a cost.

To illustrate, a famous example of temporal myopia is the marshmallow experiment for children (Mischel et al. 1972). In this experiment, preschool children (aged three to five) were asked to delay a less preferred reward (such as eating a pretzel) in order to get a more desired reward (such as eating a marshmallow). The experiment was about what strategies children used to delay gratification, such as distracting themselves with the toys which were provided to some of them. The difficulty in delaying gratification, even when doing so yields a preferred outcome, is a well-known way in which temporal myopia can be instantiated.

We now turn to explaining the appeal of descriptivism.²⁶ Why should we think that existing people's preferences are normatively important? The standard economic response is that policy choices are about getting people *what they want*, and there are at least three reasons for trying to do so (I do not endorse all of these points).²⁷ First, because introducing evaluation of people's extant preferences is thought of as just applying one's own preferences and there is no general way to show that one's own preferences are superior. Second, because consumer preferences are subject to *consumer sovereignty*, and should be accorded priority due to democratic concerns. So, for instance, we have Lerner (1972: 258) writing that “This view I find very close to the idea of democracy or freedom—the idea of normally letting each member of society decide what is good for himself, rather than have someone else play a paternal role”. Third, because, unlike some alternatives, it is empirically applicable. A rough caricature of the position is: preferences are measurable, and the only preferences we have access to are the preferences of extant people. So, for instance, philosophers prefer to appeal to reasons or arguments, but these may not be the right kind of tool for such quantitative questions, or, to the extent they are, they may just feed into the preferences of those who endorse those reasons or arguments.

These claims lead to the discussions in the final sections. While I think the first two reasons can be answered, my own view is that this third point is important and not sufficiently recognised. I think that a practical methodology has to recognise that there are only extant preferences available to consider. However, in the context of social policy, I deny that this

²⁵ I think in the most charitable interpretation of the descriptivist position, it is impartiality over preferences at all times (and places) but since we cannot observe non-current preferences, we make do with what is observable.

²⁶ I thank Tristram McPherson for pushing for this explanation.

²⁷ Dennig (2018: 44) notes that the standard economic view has “widespread agreement in the profession”.

implies that all preferences should count equally. My claim is that some social preferences can justifiably be weighted more highly than others. So although we have only preferences to contend with in setting these parameters, some of them may be better-informed in a morally relevant manner.

4. The Role of Moral Experts in Parameter Assignments

My contention in this section is that, if one is being prescriptivist about discounting, a position which has been appealing to philosophers (e.g. Caney 2008, Moellendorf 2014), there are unrecognised challenges in appealing to moral theory in this context.²⁸ I think there is a deep reason for it not being moral philosophers who should decide the values of the normative parameters η and δ . By and large, moral philosophers do not have the training to recognise the implications of changing these parameter values. While they have deep familiarity with the moral theories that make these parameter values important in the evaluation of policy, they are not trained in how to assign values to these parameters or which data in the real world might suggest certain values (or ranges of values). More contentiously, I am going to argue that moral theory does not help make up for this lack. To be precise, I am conceiving of moral theory in a narrow manner, where it is the construction of valid and invalid arguments based on moral principles. While this is far from all that moral philosophers do, it is a large and important part.²⁹ Furthermore, it is obviously the case that some moral philosophers *are* well-informed about the welfare assumptions in modern economics and, indeed, about the relevant empirical work. But, insofar as there is a trade-off between knowledge of moral theory and the relevant domains, my view will be that the latter dominates in assigning values. In my view, applied ethics does require combining moral theory with empirical information.

Note that these claims may not be true *in general* with respect to applied ethics. In most contexts, unlike the discounting problem, one is not working within a systematic framework but trying to balance intuitions about cases with principles. However, the point is that this will be challenging in this context where cost-benefit analysis already has significant presumptions about the normatively relevant factors—namely, costs and benefits. Furthermore, there is a dearth of principles that might point to specific values for the normative parameters. This is especially clear in the case of η . Utilitarianism, for instance, tells us to maximize utility but provides no guidance on how society gains utility with changes in consumption. But, again, the empirical data tell us that different approaches to measuring those changes provides divergent estimates for this parameter. So it is not a straightforwardly empirical question that can be simply considered the role of economists or other social scientists to solve.³⁰

²⁸ Prescriptivism is not only of appeal to philosophers; for instance, Arrow (1999a) explicitly adopts prescriptivism.

²⁹ In Mintz-Woo forthcoming, I argue that it is the *key* role of moral philosophers; the evaluation of arguments is plausibly the role of the public instead.

³⁰ I will suggest that the preferences of relevant domain experts should have their preferences privileged but not because it is a merely empirical question; rather that it is a *moral* question which is difficult to answer using moral argumentation.

Prescriptivism is the view that the assignment of parameter values should involve moral considerations. Introducing moral considerations into discounting is the role of (some type of) moral expert. It would be reasonable to assume that these will be moral philosophers, but I will argue that this is not the case in this context. Here, I provide the argument in brief. Let us define moral philosophers as those who are knowledgeable about, and typically trained in the use of, moral theories.³¹ In other words, they know the implications of various moral principles or judgments and which moral theories cohere with different principles or judgments. So my contention is that prescriptivism requires a moral expert. However, I will argue that the first approximation—that the relevant moral experts will be moral philosophers—is problematic in this context, since the parameter value assignments can both have infinite ranges and take what I call “non-critical values”. While I endorse this strong claim, it is worth noting that it is not necessarily needed here. I could point to the contingent point that philosophical training is unlikely to provide the mathematical acuity or the familiarity with the relevant empirical facts which are clearly relevant to having informed views on appropriate values for these normative parameters. While I think the conceptual point is correct, the contingent point is enough to suggest the limitations of moral philosophers playing the role of moral experts.

This is not to gainsay the role that moral philosophers can and should play in addressing intergenerational issues. Although I deny that philosophical argumentation can be adduced directly for particular parameter assignments, we can consider distinct roles for philosophers. A first is a standard, but more modest, role of evaluating or considering scientific methodologies and justifications. In other words, this role is not in discussing any particular parameter assignments, but instead discussing potential methodologies for generating parameter assignments. I think this is a helpful role for philosophers to play (indeed, I view the contributions of this chapter to be of this type). A second role concerns substantive moral *presuppositions*. So, for instance, the kinds of social welfare functions at issue seem to presuppose claims including, but not limited to, aggregationism, welfarism, and consequentialism. Moral theory and argumentation have an important role to play in determining whether to adopt these controversial moral claims.

However, even considering this more limited second role, I think appearances could be misleading. The first point to consider is that it is possible that these controversial moral claims are *not* presupposed. As Kelleher (2017b) points out, we can imagine that the discounting problem has limited scope. For instance, the resultant judgments about which policies to pursue might only be value-concerning, and not be all-things-considered judgments. There might be, for instance, countervailing justice-related or deontological concerns. So, regardless of the theory adopted, insofar as we are concerned with avoiding negative consequences—a consideration that most theories have even if they take value to be one of multiple aspects of

³¹ In the limit case, those theories may not be systematic; for my purposes, a non-systematic theory like moral particularism counts as a theory—just one with no theoretical implications for particular judgments. In other words, it generates no principled arguments which can be used to assess domains which satisfy the two criteria mentioned below.

morality—the problem of discounting is relevant for the axiological questions which could arise.³²

The second, and more important, point is that we can ask these questions internal to the moral theory while bracketing the external question of whether the theory is true. Considering these points conditional on the truth of the appropriate value claims, we see that the moral theory still only supports particular parameter assignments in a very few critical cases.

My remaining worry is that, since moral philosophy is conducive to arguing for critical parameter values, one will take those to be the *correct* values. The worry is a version of the law of the instrument, usually credited to Abraham Kaplan. He wrote that “It comes as no particular surprise to discover that [one] formulates problems in a way which requires for their solution just those techniques in which [one] is especially skilled” (Kaplan 1964: 28).³³ Philosophical argumentation is too coarse-grained to assign parameter values outside of these critical values, since there are not philosophical arguments for $\delta = 1$ or $\eta = 2$. The worry is that this could lead to endorsement of critical values *simply because they are amenable to philosophical argumentation*.³⁴

So, to be clear, while prescriptivism requires appeal to moral theory and argumentation, it is not so clear that moral philosophy can actually fulfil this role outside of a few critical values. My claim is not that these critical values are necessarily false, but that focus upon them could lead us to neglect the rest of the possibility space and that simply because these critical values are subject to philosophical argumentation, they might appear more plausible to philosophers.

An objection to this position might be that it is *in light of* the possibility of applying moral argumentation to critical values that we know they are the correct ones. So, for instance, many philosophers think that pure discounting is impermissible and that $\delta = 0$ is required by morality,³⁵ either by consequentialist impartiality (Broom 1994, Gardiner 2011, Jamieson 2014, Parfit 1984) or by broadly deontological concerns (Caney 2008, 2009, Davidson 2006).³⁶ However, this is sometimes accompanied by the claim that η should be revealed from consumer behaviour—or, more often, ignoring the role of η altogether, despite its relative importance. I will argue that this is more difficult to defend than is usually recognized.

³² There are of course moral theories in which consequences have no weight whatsoever. I lack space here to discuss theories of this type.

³³ The first recorded use of the law was a spoken report from a medical conference, where Kaplan said, “Give a boy a hammer and everything he meets has to be pounded” (Horowitz 1962: 637).

³⁴ I may be granting too much to the zero pure discounters; Heath (2017: 453) suggests that most philosophers don't argue for zero pure discounting but instead act as if “the commitment to a zero rate of pure time preference serv[es] as some sort of framework principle”.

³⁵ Sometimes, such as with Parfit (1984) or Kelleher (2012), there is imprecision about whether philosophers argue that $\delta = 0$ or that $r = 0$, but I think it more charitable to read them all as arguing for the former (and see Heath 2017 for similar concerns).

³⁶ However, rejecting probability discounting, as Caney (2009) does, becomes problematic because one's theory risks becoming too coarse-grained to address small changes in risk (Mintz-Woo 2017).

Notice that this is the position taken by both Stern (2006) and Cline (1992), who argue that $\delta \approx 0$ is implied by utilitarianism but that η should be revealed through observed choices.³⁷ As Dasgupta (2008: 159) argues:

Cline and Stern would appear not to have bothered at all about consistency. They chose η on the basis of estimates obtained from consumer behaviour, but ignore consumer behaviour entirely when it came to the choice of δ and sought the advice of moral philosophers instead.

When arguing for a consistent methodology, we can view this as a reductio.³⁸ Suppose that these moral philosophers should determine the value of the normative parameter δ . Then, to maintain methodological consistency, they should determine the value of the normative parameter η . It is implausible that moral philosophers should be the source for parameter assignments for η . So, assuming we want to treat different normative parameters consistently, this could be seen as a reductio on moral philosophers determining the value of δ .³⁹

There is a second reason to take η and δ to be both treated as normative parameters (or neither). If we adopt the Ramsey Rule and fix the social discount rate r (for instance, by appeal to observed rates of return) and the growth rate of consumption g (for instance, keeping this low to keep intergenerational consumption similar), then η and δ end up being ethical *duals* of each other (Dasgupta 2005). In other words, the Ramsey Rule is observed using infinitely many pairs of η and δ , and by reducing one and increasing the other, we could respect the rule. So, in practice, we have two moral judgments we can adjust as pairs of parameter values. Since we can manipulate one normative parameter by changing the value of the other, this places them in some sense morally on a par.

The upshot is that these parameter value assignments are normative, broadly speaking, but, I claim, not the especial realm of moral theory or moral philosophers. This point has gone unrecognised by many working in this area.

³⁷ The motivation is that η is dependent on how actual people smooth consumption, revealing how their utility is affected by temporally unequal consumption. However, in the context of social policies, the parameter η doesn't just govern intrapersonal utility, but interpersonal social utility. Since we cannot directly measure interpersonal utility, η is not simply a descriptive parameter in the intergenerational context.

³⁸ Note that we are assuming that the normative parameters, η and δ to be subject to prescriptivism and moral considerations. My point is that doing so in a parallel manner for both is, at the very least, challenging.

³⁹ Frej Klem Thomsen suggests that we can run this as a ponens rather than a tollens: if one thinks non-moral philosophers should have no role in determining the appropriate value of δ then, by parity, they should have no role in determining values for η . I am sympathetic insofar as these parameters should be taken as ethical duals and, while I think the tollens is more plausible, am happy to grant that the ponens would be more plausible for some.

5. Conclusion

While I have argued that consumption-based discounting is appropriate for social projects, and that prescriptivism about the discounting problem is appropriate, the major claim of this chapter has been negative: moral philosophy is of limited use in the context of quantitative issues such as these. However, this can also be seen as an interesting theoretical claim: practical ethical judgments may require normative considerations that are not a part of traditional philosophical training. Although I do think that moral experts are needed to inform prescriptivism about the discounting problem, I have left it open which moral experts they are—i.e. if they are not moral philosophers. In other work, I discuss the role that those not especially trained in moral or ethical issues can bring to moral questions.⁴⁰ In particular, I would suggest that those who know about the theory and practice of the implications of changing the moral parameter assignment values are better placed to be experts in the context of the discounting problem. While those who have not worked directly with the parameter values will have little intuitive feel for differences between parameter value assignments, those who do will. My view is that means that their intuitions should be weighted more heavily, perhaps even by taking them to be the moral experts in this context.

6. Works Cited

- Arrow, Kenneth. 1996. Intertemporal equity, discounting, and economic efficiency. In J. Bruce, H. Lee and E. Haites (eds.), *Climate Change 1995—Economic and Social Dimensions of Climate Change*. Cambridge: Cambridge University Press.
- Arrow, Kenneth. 1999a. Discounting, Morality, and Gaming. In P. R. Portney, J. P. Weyant (eds.), *Discounting and Intergenerational Equity* (13–22). Washington: Resources for the Future.
- Arrow, Kenneth. 1999b. Inter-generational Equity and the Rate of Discount in Long-Term Social Investment. In M. R. Sertel (ed.), *Contemporary Economic Issues, Volume 4* (chapter 5, 69–88). Houndmills: Macmillan Press.
- Broome, John. 1994. Discounting the Future. *Philosophy and Public Affairs* 23(2):128–156. doi:10.1111/j.1088-4963.1994.tb00008.x
- Caney, Simon. 2008. Human rights, climate change, and discounting. *Environmental Politics* 17(4):536–555. doi:10.1080/09644010802193401
- Caney, Simon. 2009. Climate change and the future: Discounting for time, wealth, and risk. *Journal of Social Philosophy* 40(2):163–186.
- Caney, Simon. 2014. Climate change, intergenerational equity and the social discount rate. *Politics, Philosophy & Economics* 13(4):320–342. doi:10.1177/1470594X14542566
- Cline, William R. 1992. *The Economics of Global Warming*. Washington, DC: Institute for International Economics.
- Dasgupta, Partha. 2005. Three Conceptions of Intergenerational Justice. In H. Lillehammer and D. H. Mellor (eds.) *Ramsey's Legacy*. Oxford: Clarendon Press, pp. 149–169.

⁴⁰ This thesis work is summarized in Mintz-Woo 2018a.

- Dasgupta, Partha. 2008. Discounting Climate Change. *Journal of Risk and Uncertainty* 37:141–169. doi: 10.1007/s11166-008-9049-6
- Davidson, Marc D. 2006. A Social Discount Rate for Climate Damage to Future Generations Based on Regulatory Law. *Climatic Change* 76(1):55–72. doi: 10.1007/s10584-005-9018-x
- Davidson, Marc D. 2015. Climate change and the ethics of discounting. *Wiley Interdisciplinary Reviews: Climate Change* 6(4):401–412. doi:10.1002/wcc.347
- Dennig, Francis. 2018. Climate change and the re-evaluation of cost-benefit analysis. *Climatic Change* 151(1):43–54.
- Fleurbaey, Marc et al. 2019. The Social Cost of Carbon: Valuing Inequality, Risk, and Population for Climate Policy. *The Monist* 102(1):84–109. Available at: <https://academic.oup.com/monist/article/102/1/84/5255707>
- Fleurbaey, Marc and Zuber, Stéphane. 2015. Discounting, Beyond Utilitarianism. *Economics* 9(12).
- Frederick, Shane, Loewenstein, George and O'Donoghue, Ted. 2002. Time Discounting and Time Preference: A Critical Review. *Journal of Economic Literature* 40(2):351–401.
- Goulder, Lawrence H., and Williams, Roberton C. 2012. The Choice of Discount Rate for Climate Change Policy Evaluation. *Climate Change Economics* 03(04):1250024
- Greaves, Hilary. 2017. Discounting for Public Policy: A Survey. *Economics and Philosophy* 33(3):391–439.
- Groom, Ben, Hepburn, Cameron, Koundouri, Phoebe, and David Pearce. 2005. Declining discount rates: The long and the short of it. *Environmental and Resource Economics*, 32(4):445–493.
- Heath, Joseph. 2017. Climate ethics: Justifying a positive social time preference. *Journal of Moral Philosophy*, 14(4):435–462.
- Heal, Geoffrey M. 2007. Discounting: A review of the basic economics. *The University of Chicago Law Review* 74(1):59–77.
- Horowitz, Milton J. 1962. Trends in Education. *Journal of Medical Education* 37(6):634–637.
- Jamieson, Dale. 2014. *Reason in a Dark Time: Why the Struggle to Stop Climate Change Failed—and What it Means for Our Future*. Oxford: Oxford University Press.
- Kaplan, Abraham. 1964. *The Conduct of Inquiry*. San Francisco: Chandler.
- Kelleher, J. Paul. 2012. Energy Policy and the Social Discount Rate. *Ethics, Policy & Environment*, 15(1):45–50. Available at: <http://www.tandfonline.com/doi/full/10.1080/21550085.2012.672684>
- Kelleher, J. Paul. 2017a. Descriptive Versus Prescriptive Discounting in Climate Change Policy Analysis. *Georgetown Journal of Law & Public Policy*, 15:957–977.
- Kelleher, J. Paul. 2017b. Pure Time Preference in Intertemporal Welfare Economics. *Economics and Philosophy*, 33(3):441–473. Available at: <https://doi.org/10.1017/S0266267117000074>
- Koopmans, Tjalling C. 1960. Stationary Ordinal Utility and Impatience. *Econometrica: Journal of The Econometric Society*, 28(2):287–309.
- Lerner, Abba P. 1972. The Economics and Politics of Consumer Sovereignty. *The American Economic Review* 62(1/2):258–266. Available at: <http://www.jstor.org/stable/1821551>.
- Manne, Alan S. 1995. The Rate of Time Preference. *Energy Policy* 23(4/5):391–394.

- Mischel, Wischel, Ebbesen, Ebbe B. & Zeiss, Antonette Raskoff. 1972. Cognitive and attentional mechanisms in delay of gratification. *Journal of Personality and Social Psychology* 21(2):204–218.
- Mintz-Woo, Kian. 2017. A new defence of probability discounting. In A. Walsh, S. Hormio and D. Purves (eds.), *The Ethical Underpinnings of Climate Economics*. Oxford: Routledge. doi: 10.4324/9781315649153
- Mintz-Woo, Kian. 2018a. Moral Uncertainty Over Policy Evaluation. *Erasmus Journal for Philosophy and Economics* 11(2):291–294. Available at: <http://www.ejpe.org/journal/article/view/351>
- Mintz-Woo, Kian. 2018b. Two Moral Arguments for a Global Social Cost of Carbon. *Ethics, Policy & Environment*, 21(1):60–63. doi:10.1080/21550085.2018.1448038 Available at: <https://www.tandfonline.com/doi/full/10.1080/21550085.2018.1448038>
- Mintz-Woo, Kian. 2019. Principled Utility Discounting Under Risk. *Moral Philosophy and Politics* 6(1):89–112. doi: 10.1515/mopp-2018-0060 Available at: <http://www.degruyter.com/view/j/mopp.2019.6.issue-1/mopp-2018-0060/mopp-2018-0060.xml>
- Mintz-Woo, Kian. Forthcoming. Public Values in the Right Context. *Australasian Philosophical Review*.
- Moellendorf, Darrel. 2014. *The Moral Challenge of Dangerous Climate Change*. Cambridge: Cambridge University Press.
- Nordhaus, William D. 1994. *Managing the Global Commons: The Economics of Climate Change*. Cambridge: MIT Press.
- Nordhaus, William D. 2007a. A Review of the “Stern Review on the Economics of Climate Change”. *Journal of Economic Literature* 45(3):686–702.
- Nordhaus, William D. 2007b. Critical Assumptions in the Stern Review on Climate Change. *Science* 317(5835):201–202.
- Parfit, Derek. 1984. *Reasons and Persons*. Oxford: Oxford University Press.
- Rawls, John. 1971. *A Theory of Justice*. Cambridge: Belknap Press.
- Sidgwick, Henry. 1907/1981. *Methods of Ethics (7th Edition)*. Indiana: Hackett Publishing.
- Solow, Robert M. 1974. Intergenerational equity and exhaustible resources. *The Review of Economics and Statistics* 41:29–45.
- Stern, Nicholas. 2006. *The Economics of Climate Change: The Stern Review*. Cambridge: Cambridge University Press.
- Sterner, Thomas and Persson, U. Martin. 2008. An Even Sterner Review: Introducing relative prices into the discounting debate. *Review of Environmental Economics and Policy* 2(1):61–76.
- Weitzman, Martin Lawrence. 2007. A Review of the “Stern Review on the Economics of Climate Change”. *Journal of Economic Literature* 45(3):703–724. doi: 10.1257/jel.45.3.703