

THE CHOICE OF DISCOUNT RATE FOR CLIMATE CHANGE POLICY EVALUATION

LAWRENCE H. GOULDER^{*,†,‡,¶} and ROBERTON C. WILLIAMS III^{†,‡,§,||}

^{*}*Department of Economics
Stanford University*

Landau Economics Building Stanford, CA 94305

[†]*Resources for the Future
1616 P St NW
Washington, DC 20036*

[‡]*National Bureau of Economic Research
1050 Massachusetts Ave
Cambridge MA 02138*

[§]*Department of Agricultural and Resource Economics
University of Maryland, Symons Hall
College Park MD 20742*
[¶]*lgoulder@stanford.edu*
^{||}*rwilliams@arec.umd.edu*

Published 31 December 2012

Nearly all discussions about the appropriate consumption discount rate for climate-change policy evaluation assume that a single discount rate concept applies. We argue that two distinct concepts and associated rates apply. We distinguish a *social-welfare-equivalent* discount rate (r_{SW}) appropriate for determining whether a given policy would augment social welfare (according to a postulated social welfare function) and a *finance-equivalent* discount rate (r_F) suitable for determining whether the policy would offer a potential Pareto improvement.

Distinguishing the two rates helps resolve arguments as to whether the choice of discount rate should be based on ethical considerations or empirical information (such as market interest rates), and about whether the discount rate should serve a prescriptive or descriptive role. Separating out the two rates also helps clarify disputes about the appropriate stringency of climate change policy.

We find that the structure of leading numerical optimization models used for climate policy analysis may have helped contribute to the blurring of the differences between r_{SW} and r_F . In addition, we indicate that uncertainty about underlying ethical parameters or market conditions implies that both r_{SW} and r_F should decline as the time-horizon increases.

Keywords: Discounting; discount rate; climate change.

1. Introduction

The choice of discount rate is critical to climate change policy assessments. Most of the climate-related benefits from current policy efforts would take the form of avoided

damages many years from now, while many of the costs would be borne in the nearer term. A high consumption discount rate¹ thus tends to shrink the present value of benefits relative to the present value of costs and weakens the case for aggressive current action. Relatively small differences in the choice of this rate can make a very large difference in the policy assessment.

The discount rate issue has become a source of significant disagreement. The *Stern Review on the Economics of Climate Change* (2007) gained considerable attention in supporting a policy of reducing greenhouse gas emissions by about 3% per year (relative to business as usual), starting more or less immediately.² To support this conclusion, the *Review* employed a consumption discount rate of 1.4%.³ But several analysts, including Nordhaus (2007) and Mendelsohn (2008), have argued that the *Review's* rate was inappropriately low and that its conclusions consequently are not well founded. Nordhaus, in particular, has argued that a considerably higher consumption discount rate has greater justification, and that once this higher rate is employed, one can no longer justify climate action nearly as aggressive as that endorsed by the *Review*. His preferred model simulations employ a discount rate of about 4.3%.³ Using his own Dynamic Integrated Climate-Economy (DICE) model, Nordhaus indicated that the differences between the Stern-endorsed and Nordhaus-supported discount rate accounted for all of the difference between the more aggressive climate policy endorsed by Stern and the considerably more modest effort supported by Nordhaus.

The disagreements about the discount rate are not merely arguments about empirical matters; there are major debates about conceptual issues as well. For example, Stern (2008) and Sterner and Persson (2008) argue that the choice of consumption discount rate should be based almost entirely on ethical considerations: there is no need, for example, to ground the consumption discount rate in observed or expected interest rates or in estimates of the opportunity cost of capital. In contrast, Nordhaus maintains that it is critical to base the choice of discount rate in observed behavior — behavior that is reflected in market interest rates. Similarly, some analysts argue that the choice of discount rate is a purely prescriptive issue, while others claim it should be a descriptive question (i.e., empirically based).

There remains relatively little agreement as to what might constitute a reasonable value for the consumption discount rate. This can leave policy analysts and decision makers confused about what conclusions can legitimately be drawn.

This paper closely examines the discount rate issue. We aim to unravel and clarify the sources of differences of viewpoint. As a result of our efforts to sort out the disagreements, we have arrived at important distinctions that we feel resolve apparent contradictions across the viewpoints. Specifically, we find that nearly all of the discussion implicitly assumes that the term “consumption discount rate” refers to a single

¹We offer definitions for this rate in Section 3.

²This policy would aim to stabilize greenhouse gas atmospheric concentrations at approximately 550 parts per million in CO₂ equivalents.

³This number derives from a particular definition of and formula for “consumption discount rate,” as discussed below.

concept, when in fact two very different concepts are involved. The same label has been used for two very different notions: what we will call the *social-welfare-equivalent* consumption discount rate (r_{SW}) and the *finance-equivalent* consumption discount rate (r_F). Distinguishing the two concepts can resolve a good part of the controversy over “the” discount rate.⁴ Distinguishing the concepts also can substantially clarify the ways in which ethical and empirical considerations become relevant to the choice of discount rate. As we discuss, r_{SW} is less directly linked with actual market behavior than is r_F . In that sense, r_F has a clearer empirical basis, though (as we discuss below) certain empirical considerations are relevant to the choice of r_{SW} as well. In addition, both rates are useful for policy evaluation, and thus both have a prescriptive role. Depending on the objective of the policy analysts involved, one or the other rate will be appropriate.

We find that the particular structure of leading numerical optimization models may have helped contribute to the blurring of the differences between the two rates. In most of the leading optimization models used for climate change policy analysis, the actual and the desirable are not clearly separated. As we discuss below, the same function serves both as a *behavioral function* (to indicate how individuals actually would behave under various conditions) and as a *social welfare function* (to indicate how individuals or societies should behave).⁵ In these models, parameters are selected to generate plausible behavioral responses. But the absence of a distinction between a behavioral function and a social welfare function means that the same parameters employed to generate a plausible behavioral function perform must be parameters of the social welfare function — since only one function is involved. This eliminates the possibility of distinguishing r_{SW} from r_F .

Separating the two consumption discount rates (and, in numerical models, the two types of objective functions) resolves disagreements about the extent to which “the” discount rate should be grounded in actual behavior. Nordhaus’s insistence that applied models reflect actual (as opposed to ideal) behavior, and Dietz and Stern’s argument that climate policy evaluation must embrace ethics (the desirable) as well as economics (narrowly defined, to refer to the actual) no longer are inconsistent.⁶ In addition, consistently accounting for these distinctions can considerably narrow the differences in viewpoint as to how ambitious current policy should be.

The rest of the paper is organized as follows. Section 2 offers some basic definitions. Section 3 defines r_{SW} and r_F , and justifies the distinction between the two. Section 4 indicates why r_{SW} and r_F cannot be distinguished in existing numerical optimization models for climate change policy, and indicates that this derives from the

⁴Our paper’s distinction between the two discount rates parallels the distinction offered in [Kaplow *et al.* \(2010\)](#) between “evaluative” and “predictive” discounting. Our paper also parallels [Kaplow *et al.*](#) in bringing out how a social-welfare function offers a basis for decision making distinct from the basis offered by market conditions. We became aware of the [Kaplow *et al.*](#) paper after having drafted the present paper and were struck by the overlap of ideas.

⁵This is a necessary consequence of the assumption, employed in many models, that the behavior of an economy can be expressed in terms of an infinitely lived representative agent. From this assumption it follows directly that whatever maximizes this agent’s utility function also maximizes social welfare.

⁶These arguments are in [Nordhaus \(2007\)](#) and [Dietz and Stern \(2008\)](#). The views are sharply contrasted in [Dasgupta \(2008\)](#).

equating of behavioral and social welfare functions in these models. This section also suggests how current practice biases the results of numerical models, and how a cleaner separation of the actual and desirable (along with a consistent use of r_{SW} and r_F) would help resolve apparent differences in policy conclusions. We show that the climate policy that maximizes a plausible social welfare function is likely to be more aggressive than one that maximizes net benefits based on r_F .

Throughout Sec. 4, our analysis ignores the issue of uncertainty. The phenomenon of uncertainty raises a large number of hugely complex issues. As discussed in Sec. 5, attending to uncertainty can influence one's choice of values for both r_{SW} and r_F . Under either approach, it calls for using lower rates when discounting over longer time horizons, which tends to imply a more aggressive climate policy. But the underlying distinctions between r_{SW} and r_F remain, as do the distinctions between behavioral and social welfare functions. In the presence of uncertainty, it remains the case that the social-welfare-maximizing climate policy can be significantly more aggressive than the one that maximizes net benefits based on r_F . Section 6 provides the conclusion.

2. Preliminaries

Before introducing a distinction between the types of consumption discount rates, it is worth noting the difference between a discount rate on *utility* and one on *consumption*. Ethicists often argue that future utility should not be discounted — that the well-being of future generations should count as much as that of the current generation in a social welfare function.⁷ This suggests a value of zero for the social rate of time preference⁸ — or perhaps a very low value to reflect the possibility that, because of a future exogenous calamity (e.g., an asteroid's hitting the earth) some future generations might not ever arrive. Using this logic, the *Stern Review* employs a value of 0.001 for the social rate of time preference, which we designate by ρ .

Consumption discount rates, in contrast, translate values of future consumption into equivalent values of current consumption. There is no necessary contradiction between employing a (positive) discount rate to future consumption and maintaining the view that future utilities should not be discounted.

3. Consumption Discount Rates

Suppose the social welfare function W is of the intertemporally additive form:

$$W_0 = \sum_{t=0}^{\infty} \left(\frac{1}{1 + \rho} \right)^t U_t(C_t) \quad (1)$$

⁷See Broome (2008) for example.

⁸The social rate of time preference is sometimes referred to as a utility discount rate. As emphasized by Kaplow *et al.* (2010), the term is potentially confusing since it does not make clear whether it represents the intertemporal trade-off of utilities in a social welfare function or the trade-off within an individual's intertemporal welfare function. The social rate of time preference refers to the former.

where W_0 is social welfare evaluated at time 0, ρ is the social rate of time preference, C_t is consumption in year t and U_t is utility at time t . Suppose that the utility function does not change through time ($U_t(C_t) \equiv U(C_t), \forall t$) and that the utility function has the constant elasticity form:

$$U(C_t) = C_t^{1-\eta}/(1-\eta) \tag{2}$$

The parameter η is the (constant) elasticity of marginal utility of consumption.⁹

3.1. The social-welfare-equivalent discount rate

Define the *social-welfare-equivalent discount rate* r_{SW} as that rate which translates a marginal change in consumption at some future date t into the social-welfare-equivalent marginal change in consumption at time 0. Thus r_{SW} must satisfy:

$$\frac{\partial W_0}{\partial C_0} = (1 + r_{SW})^t \frac{\partial W_0}{\partial C_t} \tag{3}$$

Substituting in the derivatives of Eq. (1) with respect to C_0 and C_t and rearranging yield

$$(1 + r_{SW})^t = (1 + \rho)^t \frac{\partial U_0/\partial C_0}{\partial U_t/\partial C_t} \tag{4}$$

This equation illustrates that there are two distinct reasons for discounting: the rate of time preference (represented by the first term on the right-hand-side) and the difference in the marginal utility of consumption between the two time periods (the second term). Because the marginal utility of consumption declines as consumption rises, that second term will depend on how fast consumption is rising or falling over time. Let g represent the growth rate of C , such that $C_t = (1 + g)^t C_0$. Substituting that, along with the derivative of the utility function (2) with respect to C_t , into (4), and simplifying give

$$(1 + r_{SW})^t = (1 + \rho)^t (1 + g)^{\eta t} \tag{5}$$

Assuming “small” values for ρ , η and g , taking logarithms of both sides of (5), and simplifying,¹⁰ yield¹¹:

$$r_{SW} \approx \rho + \eta g \tag{6}$$

⁹Two features of this formulation of utility and social welfare may be noted. First, it considers only a single consumption good. Recent papers (e.g., [Sterner and Persson, 2008](#); [Traeger, 2011](#)) have pointed out that with multiple consumption goods that are not perfect substitutes, relative prices are likely to change over time, with environmental goods rising in value relative to produced goods. These papers argue that this implies a different consumption discount rate for each good, with a lower discount rate for environmental than for produced goods. This is a very important point, but we view it as largely separate from discounting. The phenomenon can be incorporated into the calculation of the benefits from climate change policy (implying that value of the benefits from climate change mitigation will be rising over time), which then permits the use of a single discount rate, rather than a separate discount rate for every good. Both approaches yield the same answer, but the latter approach will generally be simpler.

Second, a more general formulation would define W_0 as $\sum_{t=0}^{\infty} (\frac{1}{1+\rho})^t V_t(U_t(C_t))$, where $V_t = (U_t(C_t))^{1-\theta}/(1-\theta)$ and $U_t = C_t^{1-\eta}/(1-\eta)$. This would allow a separate specification for the curvature of the social welfare function and that of the utility function, as determined by θ and η , respectively. [Kaplow et al. \(2010\)](#) point out where failing to split out the two elements has led to misinterpretations. [Dasgupta \(2008\)](#) discusses additional possible formulations for U .

¹⁰A key step in the simplification uses the approximation that if x is small, $\ln(1 + x) \approx x$.

¹¹[Dasgupta \(2008\)](#) offers a different but complementary derivation.

This expression again reveals that social-welfare-equivalent discounting depends on two main elements: the social rate of time preference (ρ) and the change in the marginal utility of consumption over time. The latter in turn is the product of how fast consumption is growing or falling over time (g) and how sensitive the marginal utility of income is to changes in consumption (η).¹²

It is important to recognize that r_{SW} serves to convert future consumption into a level of current consumption that is equivalent *in terms of social welfare*. r_{SW} will generally be positive to reflect both the (minimal) discounting of future utility implied by ρ , and the fact that the marginal contribution to social welfare of increased consumption is declining, as implied by η .¹³ It is also important to note that the formula relies on no information about the structure of the economy, although the expected performance of the economy is certainly implied by the assumed value of g . The critical determinants of r_{SW} are assumed parameters of the social welfare function (namely ρ) and of the utility function (namely η).

Clearly ethical considerations are relevant to the choices of values for both ρ and η . In choosing a value for ρ , one is taking an ethical stand: namely, indicating the extent to which future well-being should count relative to current well-being in the social welfare function. Likewise, in choosing a value for η , one is also making an ethical claim by stipulating how much an increase in consumption should yield in increase in the social welfare function.¹⁴

However, empirical considerations apply as well. An empirical issue relevant to the choice of η is the extent to which increments to consumption in fact lead to higher individual (as opposed to social) well-being, which can be observed via individuals' elasticity of intertemporal substitution. In addition, the choices of both ρ and η can be guided by the implications of various combinations of ρ and η for savings behavior or other economic outcomes, assuming that individuals were required to make choices consistent with maximizing the social welfare function. Along these lines, Dasgupta (2008) has shown that the optimal (according to the social welfare function) savings rate is negatively related to the value of η . Based on this relationship, he argues that certain combinations of ρ and η are difficult to justify because they would call for savings rates so high as to be ethically indefensible.¹⁵

¹²The formula is a bit more complicated if the utility function U does not exhibit constant elasticity of marginal utility of consumption or if the social welfare function is not additively separable. But the essential factors remain the same: how much the utility of a future person counts relative to a current person, and the marginal utility of changes in income.

¹³However, Dasgupta *et al.* (1999) and Dasgupta (2008) point out that for some countries, it cannot be guaranteed that growth will be positive over the relevant time-interval, particularly if there is severe climate change. This raises the possibility of a negative value for r_{SW} .

¹⁴As discussed in footnote 9, η can also incorporate society's aversion to inequality. Ethical considerations influence the choice of the value of this aversion parameter. Empirical considerations may apply as well in that the choice of η might be based on an empirical assessment of the aversion to inequality expressed by individuals.

¹⁵For example, he argues that the combination of $\rho = 0.001$ and $\eta = 1$ is indefensible because it implies an optimal savings-output ratio of approximately 97%.

Table 1. Implicit values of r_{SW} in leading climate policy evaluations.

	ρ	η	g	r_{SW}
Stern (2007)	0.1%	1.0	1.3%	1.40%
Cline (1992)	0.0%	1.5	1.3%	2.05%
Nordhaus (2007)	3.0%	1.0	1.3%	4.30%

The choices of ρ and η also have very important implications for the stringency of climate change policy. Using Eq. (6), we can calculate the implied values of r_{SW} in several leading studies. Dasgupta (2008) uses reported values for ρ and η , along with an assumed value of 1.3 for g ¹⁶ to arrive at the Table 1.

The differences in r_{SW} account for much of the difference in conclusions about the appropriate level of aggressiveness in climate change policy. For example, when the DICE model applies Nordhaus's preferred consumption discount rate of 4.3%, it yields an optimal abatement path involving CO₂ emissions reductions of 14% in 2015, 25% in 2050 and 43% in 2100 (Nordhaus, 2007). When the same model employs Stern's preferred rate of 1.4%, it yields emissions reductions of 53% in 2015. The implied difference in optimal carbon prices is very large as well: \$35/ton in 2015 for a discount rate of 4.3%, versus \$360/ton for a discount rate of 1.4%.¹⁷ These differences reflect the fact that relatively small differences in the consumption discount rate imply large differences in the discounted values attached to events in the distant future. For example, a given loss of consumption 100 years from now is 17 times smaller using a discount rate of 4.3% as compared with the result under a discount rate of 1.4%.

3.2. The finance-equivalent discount rate

The finance-equivalent discount rate r_F is different from r_{SW} . We define r_F as that rate which equates future and current consumption *in financial terms*. Put differently, this is the marginal product (or marginal opportunity cost) of capital. The rate r_F indicates how consumption levels are connected across time: if society forgoes one unit of consumption in any given period in order to increase the capital stock, this will increase the amount available for consumption in the next period by $1 + r_F$.

The basis for r_F is empirical. If capital is paid its marginal product (that is, if the capital market is undistorted), then the market rate of interest will equal r_F . Similarly, if the individual savings/consumption decision is undistorted and individuals are not liquidity constrained, then individuals will discount consumption at the rate of interest.

¹⁶The value for g was based on Dasgupta's estimate of the growth rate of consumption under business as usual in Stern (2007).

¹⁷More precisely, these two simulations incorporate the values for ρ , η and g in the rows marked "Nordhaus (2007)" and "Stern (2007)," respectively, in Table 1.

If any of these markets are distorted, however, then individuals may discount consumption at a rate that differs from r_F .¹⁸

Especially important is the fact that, irrespective of whether distortions are present or absent, r_F generally is not equal to r_{SW} . The former reflects equivalences in terms of social welfare (as defined by a given social welfare function). It is not generally equal to the rate at which an individual discounts his or her own future consumption.

As mentioned in Sec. 1, one reason for the apparent impasse in discussions about discounting in climate policy is that it is often assumed that analysts should settle on one rate. As we discuss below, this assumption lacks justification. The appropriate rate — r_{SW} or r_F — depends on the question at hand.

3.3. When should each rate be used?

Suppose a given climate change policy is being considered, and that models suggest that this policy, by preventing some climate change, would produce a benefit of ΔC_t — that is, it avoids a loss of ΔC_t in future consumption. The discount rate r_{SW} could then be used to show how much current consumption could be sacrificed without lowering social welfare. So long as the sacrifice of current consumption is less than $\Delta C_t / (1 + r_{SW})^t$ (assuming no other impacts), the policy raises social welfare.

The same policy could be evaluated using r_F — but for a different type of evaluation. r_F is appropriate if one wishes to determine whether the policy would yield a potential Pareto improvement: that is, whether the winners from the policy could in theory compensate all the losers and still be better off. In other words, it is appropriate for evaluating whether a policy would satisfy the Kaldor–Hicks criterion.¹⁹

The potential for a Pareto improvement can be illustrated (with r_F) as follows. Consider a combination of the climate change policy and a transfer from the future beneficiaries of the policy to the present generation such that the combination leaves utility unchanged in all future periods. Suppose that this transfer is accomplished by diverting resources from current investment to current consumption, and that (other things equal) it allows current-period consumption to increase by $\Delta C_t / (1 + r_F)^t$. This transfer would hold constant consumption in all intervening periods between the current period and time t , and would reduce consumption by ΔC_t in period t (which exactly offsets the gain of ΔC_t in period t from the climate policy).²⁰ Thus, the combination of climate policy and transfer leaves utility unchanged in all future periods. Now if the sacrifice of current consumption required by the climate policy

¹⁸For example, capital taxes create a substantial wedge between r_F (which, as we have defined it here, is equal to the return on capital before taxes) and the rate at which individuals discount consumption, which will equal the after-tax return on capital (under the assumption that the capital tax is the only distortion in the capital market).

¹⁹Similarly, Kaplow *et al.* (2010) discuss how the return on capital is relevant to evaluating whether a given climate-change policy represents a better use of resources than alternative investments.

²⁰This could be achieved via government borrowing to finance consumption in the current period with the borrowed funds paid back in period t . (This assumes that private saving does not change in response; if it does change, then the necessary change in the path of government borrowing would be more complicated, but could still accomplish the transfer as long as full Ricardian equivalence does not hold.)

(before the transfer) is less than $\Delta C_t/(1 + r_F)$, then (assuming no other impacts) this combination will raise utility in the current period, thus yielding a Pareto improvement.

Thus, both r_{SW} and r_F can be used for assessing climate policy, but they match up to two different criteria: social welfare in the first case, a potential Pareto improvement in the second. The objectives of increasing social welfare and of offering a potential Pareto improvement are both important — hence both r_{SW} and r_F serve important purposes. If r_{SW} and r_F differ, then these two measures can yield different results: that is, a policy can offer a potential Pareto improvement without increasing social welfare, or *vice versa*.

It is sometimes argued that the choice between using a “market interest rate” and using a “social discount rate” in climate-change policy analysis is the choice between a “descriptive” and a “prescriptive” approach to policy analysis. While it is the case that r_F has a more immediate connection with actual behavior (and in that sense has a more “descriptive” foundation), both r_F and r_{SW} can be used to evaluate or prescribe policies. The social welfare function underlying the use of r_{SW} offers a normative basis for policy analysis, and thus r_{SW} offers a basis for recommending some policy options and rejecting others. But the Kaldor–Hicks criterion that underlies the use of r_F also has a normative basis: most people would agree that satisfying the Kaldor–Hicks criterion gives a policy option greater appeal, other things equal. Hence r_F , as well as r_{SW} , can be used prescriptively.

However, it should be recognized that, by definition, the use of r_{SW} offers a broader assessment of the social-welfare implications of a policy option. When a policy’s net benefits are assessed using r_F , its potential benefits are measured according a criterion (the Kaldor–Hicks criterion) that takes no account of the distribution of impacts and other potential determinants of social welfare.

3.4. Implications of discrepancies between r_{SW} and r_F

A discrepancy between r_{SW} and r_F implies that resources (and associated levels of consumption) are not allocated across different time periods in a way that maximizes social welfare. This issue of consumption allocation can be understood without specific reference to climate policy. If, for example, $r_{SW} < r_F$, then transferring resources to future time periods by consuming less now and increasing the capital stock (e.g., by reducing the government budget deficit) will increase social welfare. Reducing consumption by one unit today and increasing the capital stock will increase the amount available for consumption next period by $1 + r_F$. If $r_{SW} < r_F$, that increase is worth more to social welfare than one unit of consumption today. If $r_{SW} > r_F$, then the opposite intertemporal reallocation of consumption — consuming more now and decreasing the capital stock — will increase social welfare. In either case, the existing policy with respect to the capital stock is not optimal because a change could increase social welfare.

A discrepancy between r_{SW} and r_F also has implications for the kind of policy that can most efficiently increase social welfare. Specifically, it can imply that climate policy will not be the most efficient approach. As we will discuss in a moment, this is not an argument against climate policy — such policy can still be welfare-improving — but it indicates that the government may have even more efficient opportunities to raise welfare.

Specifically, consider the case in which $r_{SW} < r_F$. In this circumstance, suppose that a particular climate change policy (with costs ΔC_0 now and benefit ΔC_t at time t) has positive discounted net benefits when using r_{SW} as a discount rate, but negative discounted net benefits when using r_F : that is, $\Delta C_0(1 + r_{SW})^t < \Delta C_t < \Delta C_0(1 + r_F)^t$. In this case, the policy increases social welfare. But it also in effect transfers resources from the present to the future. The government could achieve a larger increase in social welfare by making a similar transfer of resources to the future via the capital stock. Or put differently, the climate policy is a less efficient way to transfer resources to the future than a policy that increases the capital stock.

In this case, what action should the government take on the climate policy? The answer depends on what the options are. If the choice is simply whether to enact the climate policy or not, without any other policy changes, then the policy is worth enacting: it increases social welfare. If the choice is between the climate policy and a similarly costly transfer to future generations via the capital stock — one cannot do both — then using the capital stock would be better.

When the government has the opportunity to optimize the capital stock as well as consider a given proposed climate policy, the implemented changes in the capital stock will influence the values of r_{SW} and r_F . Assuming that initially r_F exceeds r_{SW} , the optimization involves increasing the capital stock, which means consuming less and saving more now. This implies both a higher r_{SW} (since consuming less and saving more now will increase the rate of growth of consumption over time, thus increasing r_{SW}) and a lower r_F (because increasing the capital stock will lower the marginal rate of return to capital). At the social welfare optimum, the two rates will be the same, at values somewhere above the original value of r_{SW} but below the original value of r_F .²¹ At the optimum, both rates give the same answer about whether the candidate climate policy is worthwhile — that is, the proposed policy will either pass both the social-welfare-improvement and the Kaldor–Hicks tests, or fail both.²²

²¹Note that this does not just mean that the two rates will be equal in the social-welfare–optimal steady state. The two rates will also be equal at all points along the social-welfare–optimizing growth path.

²²Dasgupta (2008) maintains that “only in a fully optimizing economy ... is it appropriate to discount future consumption costs and benefits at the rate that reflects the direct opportunity cost of capital.” Translated into the framework of our paper, this statement means that it is only legitimate to use r_F to measure changes in a social welfare function when all policies are optimal — in which case r_F will equal r_{SW} .

4. The Melding of the Two Rates in Numerical Optimization Models

The distinction between r_{SW} and r_F implies a need for a refinement in the structure of many optimization models used for climate change policy analysis. Fundamental assumptions in these models make it impossible to clearly separate the two rates. This has helped perpetuate confusion and unnecessary disagreement.²³

To see this, it helps to start by recognizing that optimization models have two important tasks. In these models, both tasks are served by a single function — the objective function. As indicated here, the two tasks require two functions, not one.

One key task is to indicate how the economy might *actually perform* under business-as-usual and under various alternative policies. To be taken seriously, the model needs to generate plausible behavior under business as usual, as well as realistic behavioral responses to various policy changes. In these models, the objective function drives behavior. Thus the parameters of the objective function must be chosen so as to imply plausible behavior. In Nordhaus's DICE model,²⁴ for example, the objective function is an intertemporal utility function of a representative agent. Parameters of this function are chosen so that the consumption and saving decisions of this agent seem plausible, given initial conditions and the conditions specified by policy.

A second task of numerical optimization models is (as the name implies) to reveal what policy would maximize social welfare. By definition, social welfare is maximized when the objective function is maximized. But this objective function is also the behavioral function. Thus, in the DICE model, for example, the same intertemporal utility function serves both to specify how people actually behave and is used as the metric of social welfare.²⁵ The fact that the objective function serves two roles is critical: it restricts the *social welfare function* to be the same as the *behavioral function*. Whatever parameters are chosen to make behavior realistic must also serve as parameters of a social welfare function.

Ethicists will have difficulty with this, for two reasons. For one, they might wish to employ a social welfare function that has a different functional form from that of the behavioral function in the numerical model. Second, even if they accepted the functional form, in general they would not wish to employ in the social welfare function the

²³Kaplow *et al.* (2010) arrive at similar conclusions.

²⁴The first comprehensive description and application of the model is in Nordhaus (1994). More recent applications and model refinements appear in several articles, including Nordhaus and Boyer (2000), Nordhaus (2007) and Nordhaus (2010).

²⁵In many optimization models, the identification of the social welfare function and the behavioral function stems from the assumption that the performance of the economy can be modeled by the behavior of an infinitely lived representative agent. Once this assumption is made, it is natural to assume that whatever maximizes the representative agent's welfare also maximizes social welfare (the only alternative would be to have a social welfare function that overrides the agent's preferences). It should be noted that the representative agent assumption is a strong one, and that the use of this assumption (combined with the need to parameterize the function to yield realistic behavior) severely restricts what can be considered as a social welfare function. A model with multiple agents with finite lives (such as the standard overlapping-generations model) provides far more flexibility in the choice of a social welfare function.

same parameters as the (restricted) parameters used in the behavioral function — the parameters that were restricted by the need to generate realistic time-paths.

The blurring of the distinction between the social welfare function and the behavioral function helps sustain the misimpression that there is but one consumption rate of discount. To generate realistic behavior, the optimization models must generate plausible values for, among other things, the market interest rate or opportunity cost of capital. This opportunity cost depends directly on the choice of the utility discount rate (ρ) and on the choice of the intertemporal elasticity of substitution in utility ($1/\eta$). Only certain combinations of the utility discount rate and η are consistent with plausible savings behavior. One is constrained in the choice of social welfare function.

This may seem like a fundamental dilemma. On one hand, it seems vitally important that optimization models (and other models, for that matter) generate realistic behavioral responses. On the other, it seems important to allow for a social welfare function to incorporate ethical considerations that might imply different parameters or functional forms from those that meet the demands of calibration.

The dilemma is not real. It is possible to maintain a behavioral function that leads to realistic behavioral responses *and* to have a wide choice of social welfare functions. Likewise, it is possible to disentangle — even in numerical models — the two discount rates r_{SW} and r_F . The solution is to include distinct behavioral and a social welfare functions in numerical optimization models. As is current practice, the behavioral function can be parameterized so as to generate plausible behavioral responses and plausible values for the opportunity cost of capital. Nordhaus, for example, might wish to stick with the parameters that he arrived at for the behavioral function in DICE. At the same time, a social welfare function can be superimposed on the model, to evaluate the outcomes that the behavioral function and other aspects of the model generate. The social welfare function would not alter the behavior of the model; it would only evaluate the outcomes.

This approach would yield both r_{SW} and r_F as distinct rates. The former would derive directly from the social welfare function. For example, if the social welfare function is of the form in Eq. (1), it would be calculated from Eq. (6). The choice of η and ρ would not be constrained by the interest rate or opportunity cost of capital that emerges from the model (though, as discussed above, some empirical considerations could influence these choices). The latter would derive directly from the opportunity cost of capital that emerges from the model.

Separating out the two discount rates in these models can help resolve disputes about the appropriate stringency of climate change policy. Using Eq. (6) often yields a value for r_{SW} that is lower than most estimates of r_F (the opportunity cost of capital). Thus, analysts may well support a relatively aggressive approach to emissions abatement to the extent that they are evaluating the policy in social welfare terms. At the same time, analysts can point to the higher value of r_F as an indicator that the Kaldor–Hicks criterion cannot support policies as aggressive as the social welfare criterion can.

5. Uncertainty and Discount Rates

All of the discussion thus far has assumed that all of the relevant parameters of the problem are known with certainty. This is useful for exposition but obviously highly unrealistic — particularly in the context of climate change policy. Hence it is important to consider how uncertainty affects discounting. This issue applies for both r_{SW} and r_F and uncertainty usually has qualitatively the same influence on each of the two rates. For convenience, through most of this section the discussion simply refers to “the discount rate.”

One aspect of the climate change problem that is highly uncertain is the potential benefit of any given climate policy. What effect this has depends on how the uncertainty in benefits is correlated with uncertainty about future levels of consumption. For example, a policy with benefits that are negatively correlated with consumption (i.e., one that has relatively large benefits when consumption is low and relatively low benefits when consumption is high) effectively provides some insurance, and thus is more attractive than would be suggested if one focused only on the expected benefit.²⁶

One way to handle this type of uncertainty is to use the expected benefit in the analysis, and to address risk through adjustments in the discount rate. For a policy with benefits negatively correlated with consumption, this approach would use a discount rate below the risk-free rate.²⁷ This approach can sometimes be convenient but it conflates the issues of risk and discounting. Uncertainty about the benefits of policy has exactly the same effect in a case where the benefits occur immediately as it does in a case where the benefits occur far in the future. Thus, it is generally better to take an approach that separates the two issues by using a risk-free discount rate and incorporating risk into the analysis by using certainty-equivalent benefits rather than expected benefits. For a policy with benefits negatively correlated with consumption, the certainty-equivalent benefits are higher than expected benefits (because they reflect both the expected benefit and the insurance value of the benefits). Thus, both approaches recognize that the negative correlation makes the policy more attractive, but the latter approach separates the issue of uncertainty from the issue of discounting — it avoids lumping those two issues together by adjusting a single discount rate.

The one situation where discounting and uncertainty cannot be separated is when there is uncertainty about the elements that define the discount rate. In the case of r_{SW} , there is uncertainty about the growth rate g ; for example, in the case of r_F , there is uncertainty about future opportunity costs of capital. Suppose that there is a range of possible states of nature, denoted by j , where r_j is the discount rate in state j and p_j is the probability of that state.²⁸ One way to handle this situation is to explicitly include

²⁶More generally, what matters is the correlation with the marginal utility of consumption. Uncertainty about benefits makes a policy more attractive if the policy provides higher benefits when the marginal utility of consumption is high.

²⁷In many other contexts (such as business investments), the more common case is to have benefits that are positively correlated with consumption. In such a case, this approach would use a rate above the risk-free rate.

²⁸For simplicity, assume that in any given state, the discount rate is constant over time.

this uncertainty when calculating the discounted cost and benefit of a policy. Under this approach, the expectation of the discounted value of a certain benefit of ΔC_t at time t would equal $\sum_j [p_j(1 + r_j)^{-t} \Delta C_t]$. This approach has the advantage of transparency, but may be impractical because it would require every cost–benefit analysis to consider the entire distribution of possible discount rates.

An alternative approach would be to collapse the range of possible values for the discount rate into a “certainty-equivalent rate” that would yield the same result.²⁹ This implies that

$$(1 + r^*)^{-t} = \sum_j [p_j(1 + r_j)^{-t}] \quad (7)$$

where r^* denotes the certainty-equivalent rate. Note, however, that the r^* that satisfies this equation will vary depending on t . For $t = 1$, the r^* that satisfies Eq. (6) will simply equal the expected value of r_j , and this will be approximately true for any relatively small value of t (i.e., for discounting over a relatively short time horizon). But the larger the value of t , the lower will be the value of r^* that satisfies Eq. (6). In the limit as t goes to infinity, r^* will converge to $\min_j r_j$. In other words, when discounting over very long time horizons, this certainty-equivalent discount rate equals the lowest possible discount rate.³⁰

The key to this result is that decisions depend on the expected value over discount factors [the right-hand-side of Eq. (6)], not discount rates. And as t increases, the discount factor associated with a high discount rate will decline much more rapidly than the discount factor associated with a lower discount rate. Thus, as t increases, the higher the discount rate is in a particular state, the less important that state becomes in determining the certainty-equivalent rate. So for very large values of t , the certainty-equivalent rate is determined almost entirely by the lowest of the possible discount rates. Similarly, if one explicitly uses the full distribution of possible discount rates (rather than a certainty-equivalent rate), the states with the lowest discount rates will become increasingly important as t increases.

This is a powerful result. It implies that when discounting the distant future, one should use a lower rate — potentially much lower — than the rate one would use for relatively short time horizons. And since the longer the time horizon is, the more important the discount rate becomes, this result can have dramatic consequences.³¹

²⁹This alternative approach has the virtue of simplicity, and is particularly useful for illustrating the effect of uncertainty about discount rates. But that simplicity relies on the highly unrealistic assumption that any uncertainty about benefits is uncorrelated with uncertainty about the discount rate. (One can still employ this approach when this assumption does not hold, but in such a case the certainty-equivalent rate will depend on the joint distribution of benefits and discount rates.)

³⁰This point was first made by Weitzman (1998, 2001).

³¹This time-varying discount rate can lead to cases in which the optimal decision between two policies depends on when that decision is made. These choice reversals resemble the preference reversals that occur under hyperbolic discounting, but occur for a very different reason: the choice reversals occur because of new information that is revealed over time, not because of any preference reversal or other similar irrationality.

This issue applies to both r_{SW} and r_F , because both are uncertain. As discussed earlier in the paper, r_{SW} depends on ρ , η and g , each of which are uncertain: there is a great deal of disagreement in the literature about the appropriate values to use for ρ and η , and any forecast of future economic growth rates is highly uncertain. The value of r_F is also uncertain because the future marginal product of capital is difficult to predict. Because the sources of uncertainty differ across the two rates, the magnitude of the uncertainty will probably also differ across the two rates, and thus the quantitative importance of this issue will likely differ. But the direction of the effect will be the same: whichever rate is appropriate for a given analysis, the corresponding certainty-equivalent rate will be lower the longer the time horizon.

Newell and Pizer (2003) provide an empirical analysis of this effect, using two centuries of data on U.S. interest rates.³² It shows that under a random-walk model of the interest rate uncertainty, the certainty-equivalent rate falls gradually from 4% for very short time horizons to 2% after 100 years, 1% after 200 years and 0.5% after 300 years. The analysis suggests that taking this interest-rate uncertainty in account roughly doubles the expected discounted benefits from climate mitigation policy.

6. Conclusions

There has been much debate about the appropriate consumption discount rate for use in climate change policy analysis. Nearly all discussions implicitly assume that a single discount rate concept applies. We argue that in fact two distinct rates apply, and that one or the other rate will be appropriate depending on the evaluation criterion. If the objective is to assess whether a given policy would augment social welfare (according to the postulated social welfare function), the social-welfare-equivalent discount rate r_{SW} is appropriate. If the objective is to determine whether the policy would yield a potential Pareto improvement (i.e., pass a Kaldor–Hicks test), the finance-equivalent discount rate r_F should be used.

Distinguishing the two rates resolves major disagreements about “the” consumption discount rate. For one, it resolves the debate about the extent to which the rate should be grounded in actual saving–investment behavior and the associated opportunity cost of capital, as opposed to less empirically based ethical considerations. When the evaluation criterion is the Kaldor–Hicks condition, then r_F is the appropriate rate and a focus on market conditions (the opportunity cost of capital) is justified. When the evaluation metric is the value of a social welfare function, then r_{SW} is the appropriate rate — a rate based on parameters that derive from ethical considerations with a less direct empirical basis (although empirical considerations might influence one’s views as to the appropriate values for the “ethical” parameters in the formula for r_{SW}).

³²Thus, their quantitative results are primarily relevant for r_F , though the general pattern of the results also provides some guidance for how analyses using r_{SW} might be affected.

Is one objective a better basis for decision-making than the other? By definition, the social welfare function offers the most complete measure of impacts of policy on social welfare. In principle, a social welfare function will embrace all relevant normative dimensions (including both efficiency and distributional considerations); this all-encompassing quality gives it great appeal. At the same time, the appropriate blending or weighting of these various dimensions is subjective and leads to disagreements as to the appropriate form and parameters of the social welfare function. For this reason, some analysts prefer to focus on the narrower Kaldor–Hicks criterion (which engages r_F). Although this criterion focuses only on one normative dimension (namely, the potential for a Pareto improvement), it can be more tractable.

All of this is to suggest that neither objective dominates the other: the choice between them is between an approach that is more comprehensive and one that might involve less subjectivity. Note that the choice between the social-welfare-function-based approach and the Kaldor–Hicks approach is relevant not only to climate policy evaluation but to policy assessments in many other policy contexts, particularly when there are serious distributional as well as efficiency consequences.

A closely related issue is whether the choice of discount rate should be based on “descriptive” or “prescriptive” considerations. We find that both rates are used for evaluating policy. In this sense both have a prescriptive function. At the same time, r_F is more directly tied to actual behavior, and in this respect it has a stronger descriptive element. Nevertheless, descriptive considerations also influence the choice of the parameters that determine r_{SW} . For example, the η parameter that enters the typical formula for r_{SW} can be based on considerations of individuals’ actual or revealed aversion to inequality.

We find that the particular structure of leading numerical optimization models may have helped contribute to the blurring of the differences between the two rates. In all of the optimization models we have encountered, the same function — the objective function — serves both as a behavioral function (to indicate how individuals actually would behave under various conditions) and as a social welfare function (to indicate how individuals or societies should behave). This means that the same parameters calibrated or statistically estimated to generate a plausible behavioral function must be parameters of the social welfare function — since only one function is involved. This forces the social welfare function to be directly based on actual behavior and the opportunity cost of capital. This prevents r_{SW} from being separated from r_F , and eliminates the possibility of any separation of the desirable from the actual. Fortunately, this difficulty can be overcome by introducing separate behavioral and social welfare functions in optimization models.

Uncertainty over the appropriate value for that rate implies that the rate used should decline as the time horizon increases. This applies to both r_{SW} and r_F . It has substantial practical importance for climate change policy, because of the long time horizons involved.

Separating out the two discount rates can help clarify disputes about the appropriate stringency of climate change policy. Analysts who implicitly concentrate on r_{SW} , focusing on ethical considerations, tend to call for a relatively low discount rate. This leads them to argue for more aggressive abatement efforts. Analysts who implicitly focus on r_F , drawing attention to the (relatively high) opportunity cost of capital, tend to call for a higher rate.³³ This leads them to support less aggressive action. The two views are not incompatible, r_{SW} may well be lower than r_F . In this case, a relatively aggressive climate policy might well pass the social-welfare-enhancement test yet fail the Kaldor–Hicks (potential Pareto improvement) test. Whether a given level of policy stringency is justified will depend on which of the two important evaluation criteria is being employed.

This analysis implies that practitioners need to make clear the evaluation metric in assessing climate policy. Some analysts might be more comfortable utilizing r_F since it has a closer connection to observed behavior. It may be viewed as less subjective than r_{SW} , which puts more emphasis on ethical considerations. There is no inherent problem in using r_F , but if r_F is employed it is important to interpret the results appropriately. In particular, a policy's failing the benefit–cost test with r_F only shows that the policy lacks the potential to generate a Pareto improvement: it does not rule out the possibility that the policy is social-welfare-improving. Our recommendation would be for policy analysts to make very clear the distinction between the two rates and the two associated evaluation criteria. This will lead to more informative evaluations and help avoid unnecessary disagreements.

Acknowledgments

We gratefully acknowledge helpful comments from Kenneth Arrow, Elisabeth Moyer, Nicholas Stern, David Weisbach and participants in the RFF-Stanford workshop, “The Next Round of Climate Economics & Policy Research,” Washington, DC, October 27–28, 2011.

References

- Broome, J (2008). The ethics of climate change. *Scientific American*, June 2008, pp. 97–102.
- Cline, W (1992). *The Economics of Global Warming*, Washington, DC: Institute for International Economics.
- Dasgupta, P (2008). Discounting climate change. *Journal of Risk and Uncertainty*, 37, 141–169.
- Dasgupta, P, KG Maler and S Barrett (1999). Intergenerational equity, social discount rates, and global warming. In *Discounting and Intergenerational Equity*, PR Portney and JP Weyant (eds.), pp. 51–78. Washington, DC: Resources for the Future.

³³In principle, ethical considerations could support a value for r_{SW} that exceeds r_F . In practice, however, ethical arguments for discount rates higher than r_F are rare.

- Dietz, S and N Stern (2008). Why economic analysis supports strong action on climate change: A response to the Stern Review's critics. *Review of Environmental Economics and Policy*, 2(1), 94–113.
- Kaplow, L, E Moyer and DA Weisbach (2010). The social evaluation of intergenerational policies and its application to integrated assessment models of climate change. *B.E. Journal of Economic Analysis and Policy*, 10(2).
- Mendelsohn, RO (2008). Is the Stern Review an economic analysis? *Review of Environmental Economics and Policy*, 2(1), 45–60.
- Newell, RG and WA Pizer (2003). Discounting the distant future: How much do uncertain rates increase valuations? *Journal of Environmental Economics and Management*, 46, 52–71.
- Nordhaus, W (1994). *Managing the Global Commons: The Economics of Climate Change*. Cambridge, MA and London: MIT Press.
- Nordhaus, W (2007). A review of the Stern Review on the economics of climate change. *Journal of Economic Literature*, 45, 686–702.
- Nordhaus, W (2010). Economic aspects of global warming in a post-Copenhagen environment. *Proceedings of the National Academy of Sciences*.
- Nordhaus, W and J Boyer (2000). *Warming the World: Economic Models of Global Warming*, Cambridge, MA and London: MIT Press.
- Stern, N (2007). *The Economics of Climate Change: The Stern Review*. Cambridge, UK, and New York: Cambridge University Press.
- Stern, N (2008). The economics of climate change. *American Economic Review*, 98(2), 1–37.
- Stern, T and UM Persson (2008). An even Sterner Review: Introducing relative prices into the discounting debate. *Review of Economics and Policy*, 2(1), 61–76.
- Traeger, C (2011). Sustainability, limited substitutability, and non-constant social discount rates. *Journal of Environmental Economics and Management*, 62, 215–228.
- Weitzman, ML (1998). Why the far-distant future should be discounted at its lowest possible rate. *Journal of Environmental Economics and Management*, 36(3), 201–208.
- Weitzman, ML (2001). Gamma discounting. *American Economic Review*, 91(1), 260–271.
- Weitzman, ML (2007). A review of the Stern Review on the economics of climate change. *Journal of Economic Literature*, 45, 703–734.