CPB Netherlands Bureau for Economic Policy Analysis

CPB Discussion Paper | 296

# The social discount rate under a stochastic A2 scenario

Rob Aalbers
Marjon Ruijter
Kees Oosterlee

# The social discount rate under a stochastic A2 scenario

Rob Aalbers* Marjon Ruijter†
Kees Oosterlee‡
CPB Netherlands Bureau for Economic Policy Analysis

December 3, 2014

**Abstract**

Using a general equilibrium model in which both capital productivity and temperature are uncertain, we show that the social discount rate (SDR) will decline from 1% in 2010 to 0.6% in 2300 under the conventional, quadratic specification of the damage function, and to −2.0% under the reactive specification of the damage function. Moreover, interaction between economic and climate risks further lowers this estimate of the SDR by 0.9%. Surprisingly, the decline of the SDR never starts before 2100. We attribute this to the slow response of the earth's climate to increases in radiative forcing, thus highlighting the critical importance of properly taking into account the long-term dynamics of the climate system for the SDR. Interestingly, a substantial part of the decrease in the SDR under the reactive specification can be attributed to the presence of a term premium in long-run bonds.

*Keywords*: Discounting; Risk-free rate; Social discount rate; IPCC Scenario; Distant future; Climate change.

*JEL Codes*: C61; G12; H43; Q51; Q54.

*Corresponding author: CPB Netherlands Bureau for Economic Policy Analysis, P.O. Box 80510, 2508 GM The Hague, The Netherlands. E-mail: r.f.t.aalbers@cpb.nl. The responsibility for the contents of this CPB Discussion Paper remains with the authors. The authors gratefully acknowledge the very valuable comments and suggestions by Ben Groom, which have substantially improved this paper. We would also like to thank Peter Broer, Casper van Ewijk and numerous seminar participants at the 2013 AERE Conference in Banff and the 2013 EAERE Conference in Toulouse.

†CWI - Centrum Wiskunde & Informatica, Amsterdam, The Netherlands. E-mail: m.j.ruijter@cwi.nl

‡CWI - Centrum Wiskunde & Informatica, Amsterdam, The Netherlands. E-mail: c.w.oosterlee@cwi.nl, and Delft University of Technology, Delft, The Netherlands.

# 1 Introduction

Recent work shows that the social discount rate may escape the logic of exponential discounting, when shocks on capital productivity are permanent instead of transitory. Here, the social discount rate is defined as the certainty-equivalent rate of return on a take-it-or-leave-it marginal investment at time 0 with a certain payoff at time $t$. For example, using a simple Ramsey optimal growth model with an immediate and once-and-forever shock to capital productivity, Weitzman (2010) finds that the social discount rate for three centuries hence equals 0.6% even though the expected rate of return in his model equals 6% per year. Intuitively, permanent shocks imply that the riskiness of consumption is exponentially increasing over time, giving even moderately risk-averse consumers a strong incentive to increase their savings. Although the message from Weitzman's paper is crystal clear - in cost-benefit analysis, the far-distant future should be discounted far less heavily than is done by standard exponential discounting at a constant rate - any practical implementation within climate-change CBA's hinges crucially on the availability of social discount rates that incorporate actual climate risks.

To that end, this paper employs the stochastic A2 scenario of Roe and Bauman (2012) to study the impact of climate risk on the social discount rate between 2010 and 2300. This stochastic A2 scenario emulates the climate forcing of the IPCC A2 scenario in the 21st century, after which a decline in forcing is assumed at a rate which approximately stabilizes temperature for the median trajectory. It not only recognizes that the thermal inertia of the ocean will prevent too abrupt warming of the earth's climate, but also that the earth's climate sensitivity and its response time are positively related: whereas it will take 50 years to reach two-thirds of the equilibrium impact when climate sensitivity equals 1.5 °C, it will take 5000 years when climate sensitivity equals 15 °C. Taken together, these considerations imply that the transient response of the earth's climate system to increases in greenhouse gases is limited (Roe and Bauman, 2012). Nonetheless, it is conceivable, although unlikely, that we might see double-digit increases in temperature three centuries hence. For example, the 2-sigma temperature range for the stochastic A2 scenario employed in this paper equals $[3.2, 6.2]$ °C in 2100, but $[3.3, 10.8]$ °C in 2300.

To access the impact of the 'gradual' dynamics of the earth's temperature on the social discount rate, we employ a stochastic version of a Ramsey optimal growth model, which was originally developed by Cox, Ingersoll and Ross (1985*a*). Our paper thus fits into a small number of papers that have used this CIR-model as a motivation to establish a schedule of declining discount rates by using historical data (Newell and Pizer, 2003; Groom et al., 2007; Masoliver, Montero and Perelló, 2013). In contrast, our model is forward looking and not only incorporates the above-mentioned climate-change-related stylized facts, but also uses the idea of rare macro-economic disasters (Rietz, 1988; Barro, 2009) to adequately capture key features of financial markets data. Thus, in a world unaffected by climate change, the expected rate of return on capital and the social discount rate are constant over time and equal to 6% and 1% respectively, implying an equity premium of 5%.

Our main results are fourfold. First, climate-change related risks will not have a sizeable effect on the social discount rate in this century. Intuitively, the earth's inertia prevents 'catastrophic' states of nature to be reached with sufficient probability before 2100. Second, the social discount rate may see a rapid decline after the turn of this century, where the extent of the decline primarily depends on the size of the damage in 'catastrophic' states of nature. For example, under the reactive specification of the damage function proposed by Weitzman (2012), the social discount rate reaches sub-zero levels as of 2180, after which it declines to -2.0% in 2300. The decrease of the social discount rate is more moderate under the more conventional quadratic specification of the damage function used by Nordhaus (2008), reaching 0.7% in 2200 and 0.6% in 2300. All in all, our simulations reveal that relatively moderate deviations from the quadratic damage function are sufficient for a sub-zero social discount rate by 2300. Third, the addition of interaction between economic and climate risks significantly lowers our estimate of the social discount rate by about 0.9% by 2300. Fourth, the social discount rate under the stochastic scenario deviates significantly from its deterministic counterpart, signifying that the proper treatment of far-distant uncertainty is critically important. To illustrate, by 2300 the social discount rate under the stochastic scenario is 0.1 percentage point (13%) lower for the quadratic damage function, but more than 2.3 percentage points (767%) lower for the reactive damage function.

Surprisingly, our results show that Weitzman's stylized example with an immediate and once-and-forever shock to capital productivity is actually at the upper end of our estimates of the social discount rate. Under the stochastic A2 scenario, any deviation from the frequently used quadratic damage function will result in lower, and in many cases negative, estimates of the social discount rate by 2300. Notice that negative real discount rates are relatively common and may occur over extended periods of time (Reinhart and Sbrancia, 2011; Masoliver, Montero and Perelló, 2013). For example, in times of a crisis, investors may have such a desire to avoid risk that they are willing to accept a negative yield on riskless investments.[1] Our results do not necessitate unrealistically low assumptions on either the pure rate of time preference or the coefficient of relative risk aversion, as, in our base case, these have values of 0.015 and 4, respectively. Instead, it is the *combination* of the reactive damage function and the risk of high temperatures under the A2 scenario that is crucial for obtaining negative social discount rates three centuries hence.

Interestingly, a substantial part of the difference between our results and Weitzman's stylized example can be attributed to the fact that the expected rate of return on the take-it-or-leave-it investment with a certain payoff at time-$t$ contains a term premium. This term premium measures an investments' potential to hedge against shifts in temperature and can be traced back to the desire of non-log-utility consumers to hedge against changes in temperature (cf. Cox, Ingersoll and Ross (1981)). As a result, the discount rate to be used at each instant for this take-it-or-leave-it-investment will in general not be equal to the prevailing risk-free rate. In our results, this desire to hedge results in a decrease of the

[1] In the aftermath of the financial crisis of 2007, both 2-year German government bonds and 1-month US Treasury bills have been sold at a negative real rates (U.S. Department of the Treasury, 2014).

SDR of 0.04 percentage points for the quadratic and 0.9 percentage points for the reactive damage function.

To date, there is, to the best of our knowledge, no study tying the stochastic development of temperatures under a realistic emissions scenario, such as the IPCC A2 scenario, to the level of the social discount rate. Using the parable of a once-and-forever shock to capital productivity, Weitzman (2010) obtains an analytical solution for the social discount rate in the context of a simple Ramsey optimal growth model. In a standard Lucas tree economy, Gollier (2014) shows that efficient discount rates are decreasing with maturity, i.e., contain a term premium, when growth rates are serially correlated, but he does not explicitly model temperature. Nordhaus (2008) and Stern (2007) use assumptions on the pure rate of time preference and the coefficient of relative risk aversion (which they refer to as the elasticity of marginal utility) to obtain a social discount rate, but they do not establish a formal link between uncertainty and the social discount rate. In a study on the optimal level of environmental investment within a setting of economic and environmental disasters, Barro (2013) stresses the importance of gauging (fat-tailed) uncertainty based on empirical evidence. Using a database of 185 economic rare disasters for 40 countries over periods going back as far as 1870, he is able to explain the observed equity premium on financial markets without invoking unrealistically low assumptions on either the pure rate of time preference and the coefficient of relative risk aversion. He extends the model to include environmental rare disasters, but the present-day probability of an environmental disaster under the baseline calibration does not concord with the fact that the thermal inertia of the ocean will prevent too abrupt warming of the earth's climate. Cai, Judd and Lontzek (2013) extend the DICE model of Nordhaus (2008) to a stochastic setting, but they do not include fat-tailed uncertainty based on empirical evidence on rare macro-economic disasters.

The outline of this article is as follows. In Section 2, we present our model of choice and derive the corresponding social discount rate. Section 3 presents and discusses the calibration of the model's economic and physical parameters. The results are presented in Section 4. Section 5 discusses some policy implications and concludes.

# 2 The optimal portfolio model and the social discount rate

In this section, we explain our stochastic general equilibrium model, which is based on the general equilibrium model of asset pricing by Cox, Ingersoll and Ross (1985*a*). Consider an economy in which a single production sector produces the consumption good $C(t)$. The time-$t$ value of an initial investment in this production sector is denoted by $K(t)$. Under continuous reinvestment of the output, this value evolves according to the following stochastic differential equation

$$dK(t) = (\alpha(T(t)) - \kappa\lambda)K(t)dt + K(t)G(t)d\omega(t) + K(t)\kappa dq(t), \tag{1}$$

$$G(t)d\omega(t) = \begin{bmatrix} G_1(t) & G_2(t) \end{bmatrix} \begin{bmatrix} d\omega_1(t) \\ d\omega_2(t) \end{bmatrix}. \tag{2}$$

Here, $\omega(t)$ denotes a two-dimensional Wiener process. The increments $d\omega_i$ are independent normally distributed with mean zero and standard deviation $\sqrt{dt}$. This Wiener process evolves continuously over time: whereas random process $\omega_1(t)$ represents 'normal' economic fluctuations, random process $\omega_2(t)$ represents fluctuations in the earth's climate. $G(t)$ denotes the return volatility, which may be time-dependent. We extend the model of Cox, Ingersoll and Ross (1985*a*) with jumps, as in Ahn and Thompson (1988). Let $q(t)$ be a Poisson process with intensity rate $\lambda$, while $\kappa < 0$ is its fixed jump size. This Poisson process is discontinuous and reflects that from time to time macroeconomic disasters, such as wars, depressions and revolutions, may hit the economy (cf. Barro (2009)). In our model, these investments have stochastic constant returns to scale with expected value $\alpha(T(t))$. Note that we explicitly allow the expected rate of return $\alpha$ to depend on the temperature $T(t)$, which follows the following stochastic differential equation

$$dT(t) = \mu(t, T(t))dt + s_2(t, T(t))d\omega_2(t), \quad T(t_0) = T_0. \tag{3}$$

Here, $\mu$ is the drift of temperature and $s_2$ is the temperature volatility. Note that unexpected changes in the economy do not affect temperature as we set $s_1 = 0$. Thus, equations (1) to (3) describe an economy that is confronted with exogenous temperature risk: unexpected changes in temperature and unexpected changes in the rate of return on capital are correlated with covariance $\text{Cov}(dK(t), dT(t))/dt = K(t)G(t)s(t, T(t))'$, with $s = [0,\ s_2]$.

The consumer allocates a share $a$ of his wealth $W(t)$ to the physical production opportunity (with expected return $\alpha(T)$) and the remaining part $1 - a$ to the riskless opportunity for borrowing and lending (with a risk-free return $r$).[2] Thus, $aW(t)$ denotes the amount of time-t wealth invested in the physical production sector. In addition, he chooses his consumption flow $C(t)$ as to maximize the expected discounted utility over the interval $[t, t^*]$, where $t^*$ denotes the terminal time. We use the univariate CRRA utility function

$$U(C) = C^{1-\eta}/(1-\eta), \quad \eta > 0,\ \eta \neq 1, \tag{4}$$

where $\eta$ is the constant coefficient of relative risk aversion. The value function is

$$J(t, W, T) = \max_{\{a, C\}} \mathbb{E}^{t,W,T} \left[ \int_t^{t^*} e^{-\delta(\tau - t)} U(C(\tau))d\tau + e^{-\delta(t^* - t)} U(W(t^*)) \right], \tag{5}$$

[2] In general, the risk-free rate will depend on temperature $T(t)$, wealth $W(t)$ and time $t$.

with $\delta$, the pure rate of time preference, capturing the consumer's impatience. The wealth dynamics evolve according to

$$\begin{aligned} dW(t) &= aW(t)\frac{dK(t)}{K(t)} + (1-a)rW(t)dt - C(t)dt \\ &= \left(a(\alpha(T(t)) - \kappa\lambda - r) + r - p(t)\right) W(t)dt + aW(t)G(t)d\omega(t) + aW(t)\kappa dq(t), \end{aligned} \quad (6)$$

where $p := C/W$ denotes the proportion of wealth consumed. The corresponding Hamilton-Jacobi-Bellman equation reads (Cox, Ingersoll and Ross, 1985*a*; Ahn and Thompson, 1988))

$$\begin{aligned} 0 = \max_{C\geq 0, a>0} \Big[ & U(C) + J_W\left(a(\alpha - \kappa\lambda - r)W + rW - C\right) + \tfrac{1}{2}J_{WW}a^2W^2GG' \\ & + J_T\mu + \tfrac{1}{2}J_{TT}ss' + J_{WT}aWGs' + J_t - \delta J + \left(J(t, W + a\kappa W, T) - J\right)\lambda \Big] \\ := \max_{C\geq 0, a>0} & \psi(C, a|t, W, T), \end{aligned} \quad (7)$$

where subscripts of the value function denote partial derivatives. The optimal values of the control variables are denoted by $\hat{C}$ and $\hat{a}$. An economic equilibrium is defined as a set of stochastic processes $(r, \hat{a}, \hat{C})$ satisfying the first order conditions $\psi_C = 0$ and $\psi_a = 0$, and the market clearing condition $a = 1$. We get

$$0 = U_C - J_W, \quad (8)$$

$$0 = J_W(\alpha - \kappa\lambda - r)W + J_{WW}aW^2GG' + J_{WT}WGs' + J_W(t, W + a\kappa W, T)W\kappa\lambda. \quad (9)$$

The first order condition for consumption gives $\hat{C} = J_W^{-1/\eta}$. With separation of variables, it follows from Cox, Ingersoll and Ross (1985*b*) that the value function takes the form

$$J(t, W, T) = f(t, T)U(W) \quad (10)$$

and

$$-\frac{WJ_{WW}}{J_W} = \eta, \quad -\frac{J_{WT}}{J_W} = -\frac{f_T}{f} = \eta\frac{C_T}{C}. \quad (11)$$

This implies that the optimal consumption policy $\hat{C}(t, W, T)/W = f(t, T)^{-1/\eta} = \hat{p}(t, T)$ does not depend on the wealth level. The choice of the utility function implies that the equilibrium risk-free rate $r(t, T)$ (hereafter simply denoted by the risk-free rate) is independent of the wealth level too. It can be derived as in Cox, Ingersoll and Ross (1985*a*) and Ahn and Thompson (1988), and equals

$$r(t, T) = \alpha - \eta GG' - \eta\frac{\hat{C}_T}{\hat{C}}Gs' + [(1+\kappa)^{-\eta} - 1]\kappa\lambda, \quad (12)$$

where $\hat{C}_T$ represents the derivative of optimal consumption with respect to temperature. Inspection of Eq. (12) reveals that the risk free rate is equal to the expected rate of

return $\alpha$ minus the risk premium on aggregate wealth $\eta GG' + \eta \frac{\hat{C}_T}{\hat{C}} Gs' - [(1+\kappa)^{-\eta} - 1]\kappa\lambda$. The risk premium on aggregate wealth can be associated with the presence of 'normal' economic fluctuations $\eta GG'$, temperature risk $\eta \frac{\hat{C}_T}{\hat{C}} Gs'$, and the presence of macroeconomic disasters $-[(1+\kappa)^{-\eta} - 1]\kappa\lambda$. When wealth and temperature are uncorrelated, that is $Gs' = 0$, the second term in the risk premium will cancel out. The risk premium associated with macroeconomic disasters is proportional to the disaster probability $\lambda$, but depends nonlinearly on the disaster size $\kappa$ and the coefficient of relative risk aversion $\eta$.

**The social discount rate** The prime interest of this paper lies with the social discount rate, which is defined as the certainty-equivalent rate of return on a take-it-or-leave-it investment at time $t_0$ with a certain payoff at time $t$ (cf. Weitzman (2010)). Notice that this social discount rate is *not* equal to the risk-free rate $r(t, T(t))$ in Equation (12), which is the instantaneous rate of return on an riskless investment in the interval $(t, t + dt)$. Instead, the social discount rate $R(t)$ is given by

$$R(t) = -\frac{1}{t - t_0} \ln(\Phi(t)). \tag{13}$$

Here, the time-$t$ expected discount factor $\Phi(t)$ denotes the current (time $t_0$) price of a pure discount bond promising to pay one euro at time $t$. According to Cox, Ingersoll and Ross (1985*a*, Lemma 3), the value of this bond is given by

$$\Phi(t) = \mathbb{E}\left[\exp\left(-\int_{t_0}^{t} \beta(\tau, T(\tau))d\tau\right)\right], \tag{14}$$

where the expectation is taken over the actual temperature process (3) and $\beta(\tau, T(\tau))$ denotes the expected rate of return on our pure discount bond. Equation (14) states that the present value of this bond equals the expected discounted value of the promised payment of one euro at time $t$, where the relevant discount rate is given by its rate of return $\beta(\tau, T(\tau))$. Unfortunately, it does not provide a constructive way of finding $\Phi(t)$, since the randomly varying rate of return $\beta(\tau, T(\tau))$ is in general not known in advance (cf. Cox, Ingersoll and Ross (1985*a*)). In particular, notice that the rate of return $\beta(\tau, T(\tau))$ will in general not be equal to the risk-free rate $r(\tau, T(\tau))$, since the rate of return on long-term pure discount bonds will in general contain a term premium. This term premium can either be positive of negative and measures the extra compensation for the bond's potential to hedge against shifts in the expected rate of return $\alpha(T(t))$ resulting from shifts in temperature $T(t)$ (cf. Cox, Ingersoll and Ross (1981)). Only in the special case where we have equality of local expected rates of return on all bonds, we have $\beta(\tau, T(\tau)) = r(\tau, T(\tau))$.[3] In fact, equation (14) provides a generalization of the social discount concept

[3]In the asset pricing literature, this condition is known as the Local Expectations Hypothesis, see Cox, Ingersoll and Ross (1981). They show that a primary requirement for the Local Expectations Hypothesis to hold is that utility is logarithmic, i.e., $U(C) = \ln(C)$. In that case, the value function is separable in wealth and temperature, i.e., $J_{WT} = 0$ (see Cox, Ingersoll and Ross (1985*b*)), and consumers have no desire to hedge against changes in temperature.

employed by Weitzman (1998). To see this, consider his case where the risk-free rate $r_j(\tau)$ is drawn from a known distribution at time $t_0 = 0$ for $\tau \geq t_0$ and all uncertainty is resolved immediately after the draw. From (14), we have that $\Phi(t) = \mathbb{E}[e^{-\int_{t_0}^{t} r_j(\tau)d\tau}]$, which is equal to Weitzman's formula (5) on page 203.

Fortunately, Lemma 4 of Cox, Ingersoll and Ross (1985*a*) does provide us with a practical way of determining the price of a pure discount bond promising to pay one euro at time $t$. It reads

$$\Phi(t) = \mathbb{E}\left[\exp\left(-\int_{t_0}^{t} r(\tau, T(\tau))d\tau\right)\right], \tag{15}$$

where the expectation in (15) is taken with respect to

$$dT(t) = [\mu(t, T(t)) - \phi_T]dt + s_2(t, T(t))d\omega_2(t), \quad T(t_0) = T_0, \tag{16}$$

and where $\phi_T = \eta G s' + \eta \hat{C}_T / \hat{C} s s'$ is the temperature risk premium.[4] Intuitively, equation (15) states that the price of a pure discount bond promising to pay one euro at time $t$ is equal to its expected discounted value, where the discount rate is equal to the risk-free rate and the expectation is taken with respect to the risk-adjusted temperature process in (16). The risk adjustment is done by subtracting the temperature risk premium $\phi_T$ from the drift of the actual temperature process (3). Since the expressions in (14) and (15) measure the same value, the special case of $\phi_T = 0$ implies that the expected rate of return on our pure discount bond must be equal to the risk-free rate in that case, i.e., we have $\beta(\tau, T(\tau)) = r(\tau, T(\tau))$. This confirms our claim that equation (14) provides indeed a generalization of the social discount rate concept of Weitzman (1998).

# 3 Calibration of the Model

The basic idea underlying the calibration of our model is that it should be able to explain a number of important stylized facts regarding asset returns or, equivalently, discount rates. More specifically, we require that - given reasonable values for the pure rate of time preference and the coefficient of relative risk aversion - the model is able to get into the right ballpark regarding the historically observed values for both the risk-free rate and the expected rate of return on assets. In other words, we impose the not unreasonable assumption that at the initial time $t_0 = 2010$, climate change has not yet affected the rate of return on assets. Thus, the expected rate of return $\alpha(T_0)$ is set equal to 6% per year. This is somewhat below the expected rate of return in Barro (2006, 2009), but concords with the value chosen by Weitzman (2010). In addition, the risk-free rate $r(t_0, T_0)$ is set equal to 1% per year. The risk premium is then equal to 5% per year.

[4] From Cox, Ingersoll and Ross (1985*a*, equation (20)), we have $\phi_T = -\frac{J_{WW}}{J_W} W G s' - \frac{J_{WT}}{J_W} s s'$. Using (11), gives $\phi_T = \eta G s' + \eta \hat{C}_T / \hat{C} s s'$.

| Parameter | | Value |
|---|---|---|
| $\eta$ | : coefficient of relative risk aversion | 4 |
| $G$ | : return volatility vector | $[0.02, 0]$ |
| $s$ | : temperature volatility vector | $[0, s_2]$ |
| $\lambda$ | : macroeconomic disaster probability | 0.017 |
| $\kappa$ | : effective macroeconomic disaster size | -0.406 |
| $\alpha(T_0)$ | : expected rate of return absent climate change | 0.06 |
| $\delta$ | : pure rate of time preference | 0.015 |
| $\delta_K$ | : rate of depreciation | 0.05 |

Table 1: Parameters in the baseline calibration.

Next, we calibrate the parameters in Equation (12) for the case of $Gs' = 0$. As noticed by Barro (2006), the usual view in the finance literature is that $\eta$ is in the range of something like 2 to 5. Weitzman (2010) uses a value of 3, while Nordhaus (2008) uses a value of 2. However, Barro (2009) finds that values lower than 3 do not accord with observed equity premia and risk-free rates. In our baseline calibration, we use $\eta$ equal to 4. This and subsequent calibration parameters are collected in Table 1. We take the return volatility vector $G$ equal to [0.02, 0]. The value of $G_1$ accords with Barro (2006, 2009), but notice that its value is quantitatively unimportant in the calibrations, as $\eta G_1 G_1 = 0.16\%$.[5] This is the famous risk-premium puzzle. However, the disaster probability $\lambda$ and the contraction proportion $\kappa$ are nontrivial in the calibrations: $-[(1+\kappa)^{-\eta} - 1]\kappa\lambda$ equals 4.85%, when $\lambda = 0.017$ per year and $\kappa = -0.406$. Here, $\lambda = 0.017$ corresponds to the observed frequency at which macroeconomic disasters, such as wars, depressions, and financial crises, have occurred during the last century (Barro, 2009). Instead of using Barro's (2009) empirical distribution of $\kappa$ with impacts ranging from 15% to 64% and a mean of 29%, we use the effective average value of a contraction, which is equal to 40.6%. As noted by Barro (2009), this effective average value generates about the same equity premium and welfare effects as the empirical observed frequency distribution. Figure 1 displays the iso-risk-free-rate lines for $r(t_0, T_0)$ and shows which choices of $\kappa$ and $\eta$ keep the level of the risk-free rate constant (given the calibrated values of the other parameters $G_1 = 0.02, \alpha = 0.06$ and $\lambda = 0.017$). In the figure, our baseline calibration is indicated by a star. In the sensitivity analysis, we will vary $\kappa$ and $\eta$ such that the risk-free rate $r(t_0, T_0)$ remains unchanged. Finally, we set the pure rate of time preference equal to 0.015, but notice that, in the baseline calibration, its choice is immaterial for the level of the risk-free rate, because $Gs' = 0$.[6]

To visualize the role of the return volatility $G_1$, the disaster probability $\lambda$ and the contrac-

[5] We have not been able to find any empirical information on the value of $G_2$. Hence, we set this parameter to zero in the baseline calibration.

[6] Notice that this is an artifact of our specific modeling assumptions: with constant stochastic returns to scale and power utility, changes in the pure rate of time preference have no effect on the expected rate of return, the risk-free rate, and the equity premium (cf. Barro (2009)).

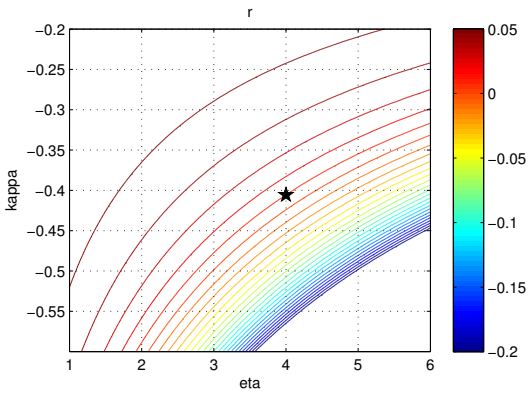


Figure 1: Risk-free rate $r(t_0, T_0)$ as function of $\kappa$ and $\eta$.

tion proportion $\kappa$ in the level of the risk premium, Figure 2 displays twenty-five randomly drawn wealth paths for the period 2010 to 2200 satisfying Eq. 6 and the parameter values of Table 1.[7] The Figure clearly displays the contrast between the almost negligible influence of the return volatility $G_1$, which is associated with the 'wiggling' behavior of the wealth paths, and the huge impact of contractions, which is associated with the large discontinuous jumps in wealth.

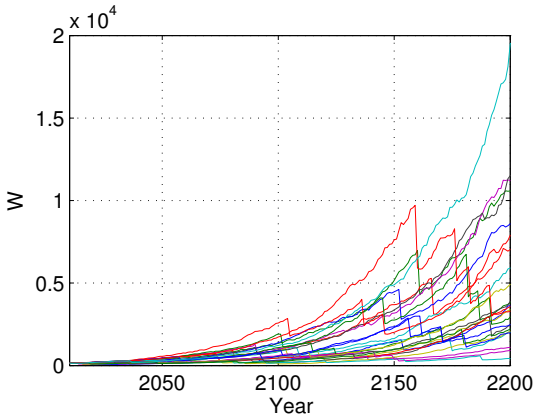


Figure 2: Simulated wealth paths.

Next, we specify the dependence of the expected rate of return $\alpha(T)$ on temperature $T$. Tol (2009) provides a comprehensive list of the welfare effects of climate change. None of the thirteen studies in his review has looked at the welfare effects of temperature increases beyond 3 °C. Moreover, these studies show that considerable uncertainty exists regarding the welfare impact of climate change, even for 'moderate' temperature increases. For

[7]Notice that in these simulations (only), we assumed that climate change has no effect on the rate of return $\alpha(T)$, i.e., $\forall T : \alpha(T) = \alpha(T_0)$. Consequently, optimal consumption is a fixed percentage of wealth.

example, the estimates for a 2.5 °C warming range from -1.9 percent to +0.9 percent of GDP. The absence of more-than-three-degrees studies implies that any estimate of the welfare impacts of climate change beyond 3 °C is, by necessity, based on extrapolation. And although we may be fairly confident about a damage extrapolation for a warming of 4 °C, that confidence will be greatly reduced for damage extrapolations to 10 °C or even 15 °C. Essentially, we have no objective way to determine the magnitude of high-temperature damages (Weitzman, 2012). Therefore, in this study, we use a 'range' of damage functions, encompassing both the low- and high-end of the available appraisals. Specifically, the damage function employed by Nordhaus (2008) serves as our low-end damage function, whereas the damage function employed by Weitzman (2012) serves as our high-end damage function, i.e.,

$$\Omega^N(T) = 1/(1 + (T/20.46)^2), \tag{17}$$

$$\Omega^W(T) = 1/(1 + (T/20.46)^2 + (T/6.081)^{6.754}). \tag{18}$$

Above, $\Omega^N(T)$ and $\Omega^W(T)$ denote the 'Nordhaus'- and 'Weitzman'-damage function respectively. Notice that these damage functions are not readily comparable, as Nordhaus defines his damage function $\Omega^N(T)$ over gross output, whereas Weitzman defines his damage function $\Omega^W(T)$ over net output. For practical reasons that will become clear below, we choose to retain these differences. Hence, the expected rates of return for our low- and high-end cases are given by

$$\alpha^N(T) = -\delta_K + (\alpha(T_0) + \delta_K)\Omega^N(T), \tag{19}$$

$$\alpha^W(T) = \alpha(T_0)\Omega^W(T), \tag{20}$$

where $\delta_K$ denotes the rate of depreciation in the economy and is taken to be equal to 5%. Figure 3 displays the range of our assumptions on the expected rate of return. Observe that the expected rates of return $\alpha^N$ and $\alpha^W$ are almost indistinguishable for temperature increase less than 3 °C, but differ markedly for larger increases in temperature. For example, an increase of 6 °C results in an expected rate of return of 5.1% (Nordhaus) and 3.0% (Weitzman), whereas an increase of 10 °C results in an expected rate of return of 3.9% (Nordhaus) and 0.2% (Weitzman). Finally, notice that these differences would have been much larger, had we chosen to define the Weitzman damage function $\Omega^W(T)$ over gross, instead of net, output.

We calibrate our temperature process on the A2 scenario of Roe and Bauman (2012), who employ a simple climate model to describe the interactions between the atmosphere, the surface and deep ocean layer, thereby capturing three material characteristics of global warming. First, the thermal inertia represented by the deep ocean slows down the response of surface temperatures on global warming. Second, climate sensitivity is assumed to be normally distributed, which results in a skewed density for the temperature response. Third, the response time, i.e., the time needed to reach the equilibrium temperature given a certain level of forcing, is positively related to climate sensitivity: whereas it will take

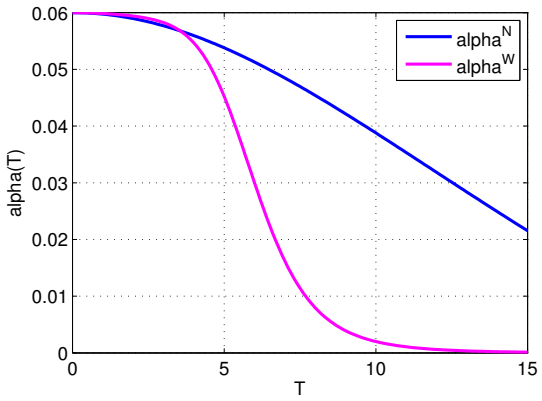


Figure 3: The low- and high-end expected rates of return $\alpha^i(T)$.

approximately one hundred years to achieve equilibrium, when climate sensitivity is low, it will take thousands of years to achieve equilibrium when climate sensitivity is high. Figure 4 presents the simulation results of Roe and Bauman (2012). In their scenario, the climate forcing is chosen to emulate the IPCC A2 scenario (Nakićenović et al., 2000), where after a decline in forcing is assumed at a rate that approximately stabilizes the temperature for the median trajectory. In Figure 4, the shadings represent the 1-, 2-, and 3-sigma ranges for climate sensitivity.

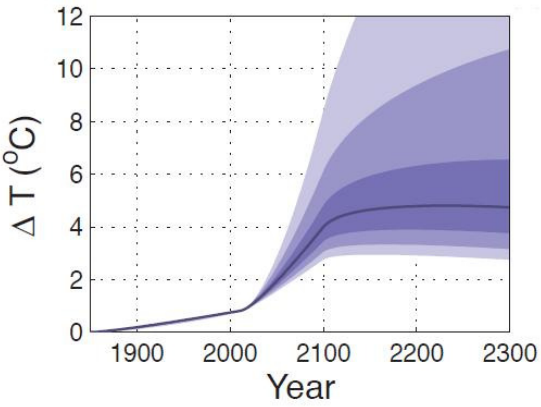


Figure 4: Quantile functions of Roe and Bauman (2012), their Fig. 4b.

We take initial time $t_0 = 2010$ and terminal time $t^* = 2300$ and calibrate functions $\mu(t, T)$ and $s_2(t, T)$ in Eq. (3) such that the quantiles of our simulations resemble the results of Roe and Bauman (2012).[8] Figure 5a shows a number of simulated temperature paths together with our simulated ranges for climate sensitivity, which resemble the results of Roe and Bauman (2012) rather well. The shape of the temperature distribution at particular times

[8]Details on the calibration procedure are provided in Appendix A.

is shown in Figure 5b, which clearly illustrates the time evolution of uncertainty under the stochastic A2 scenario: even though temperature increases of more than 6 °C are very unlikely to occur before 2100, such increases are much more likely to occur after 2100.

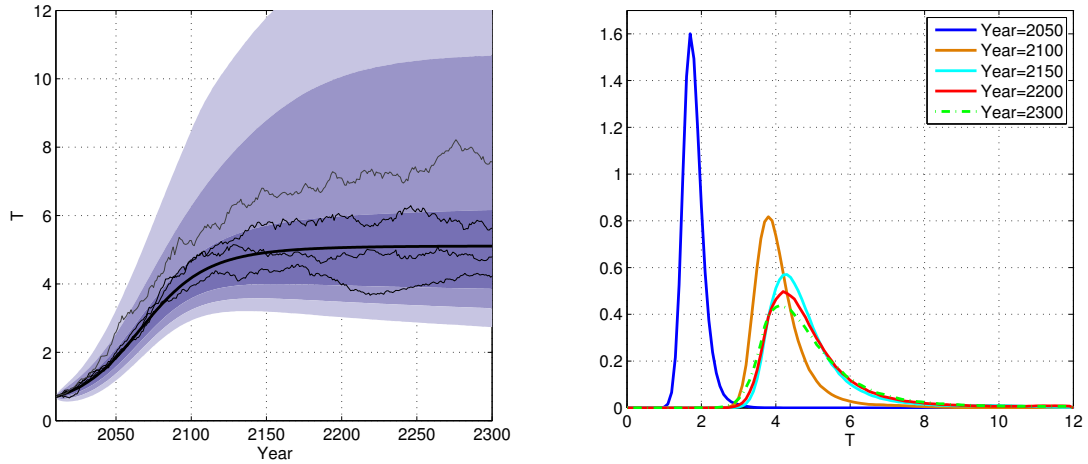


(a) Temperature paths $T(t)$ and sigma ranges. (b) Density function $T(t)$ for different times $t$.

Figure 5: Calibrated temperature process.

# 4 Results

Using the parameters in the baseline calibration displayed in Table 1, we determine the social discount rate for both our low-end and high-end damage functions.[9] For sake of reference, the dashed blue lines in Figure 6 show the development of the social discount rate at the median value of temperature, i.e., the temperature on the (deterministic) A2 path of the IPCC, under both the low-end (left figure) and high-end damage function (right figure). These figures reveal immediately that the social discount rate on the deterministic A2 path does not differ markedly between these damage functions. Intuitively, this follows from the observation that on the IPCC A2 path the temperature increase is limited to 5 °C. At these temperature increases, the low- and high-end damage functions are still rather similar (see Figure 3). Finally, notice that the social discount rate on the IPCC A2 path decreases only mildly from 1% in 2010 (our starting year) to respectively 0.7% ('low-end' damage function) and 0.3% ('high-end' damage function) in 2300.

[9]Our procedure to obtain the social discount rate is as follows. First, we solve the optimal consumption problem of Eq. (5) by using the 2D-COS method, which is described in detail in Appendix B. Substitution of the optimal consumption policy in Eq. (12) gives the risk-free rate $r(t, T)$. Finally, we approximate the expected value in equation (15) by means of a Monte Carlo simulation of the temperature and corresponding risk-free rate paths.

Next, the red lines in Figure 6 show the development of the social discount rate under our stochastic A2 scenario for both the low-end and high-end damage function. A number of striking results appear from this Figure. First of all, comparison of Figure 6a and 6b shows that in both cases the social discount rate is almost stable over the course of this century. The intuition is that the earth's inertia prevents strong increases in temperature - and thereby high damages - before the end of this century. From this, however, it does not follow that climate change damages are immaterial for the social discount rate. To the contrary, Figure 6a shows that under the low-end damage function, the social discount rate slowly declines from 0.9% in 2100 to 0.6% in 2300, which is just under the SDR on the (deterministic) IPCC A2 path of 0.7%. However, under the high-end damage function the decline is much more pronounced: the social discount rate declines from 0.9% in 2100 to -2.0%(!) in 2300 (see Figure 6b). Thus, under our stochastic A2 scenario and the high-end damage function, the net present value of 1 euro in 2300 would be equal to 350 euros today. Comparison of the social discount rates in Figure 6a and 6b reveals that this 'discounting reversal' must be attributed to the choice for the high-end damage function, as the social discount rate is still positive under the low-end damage function. Notice, however, that even under the low-end damage function, the logic of exponential discounting is strongly diminished: 1 euro in 2300 would be equal to 0.18 euro today, instead of 0.13 (using the social discount rate that corresponds to the median value of temperature as the relevant discount rate) or zero (using the expected rate of return of 6% as the relevant discount rate).[10]

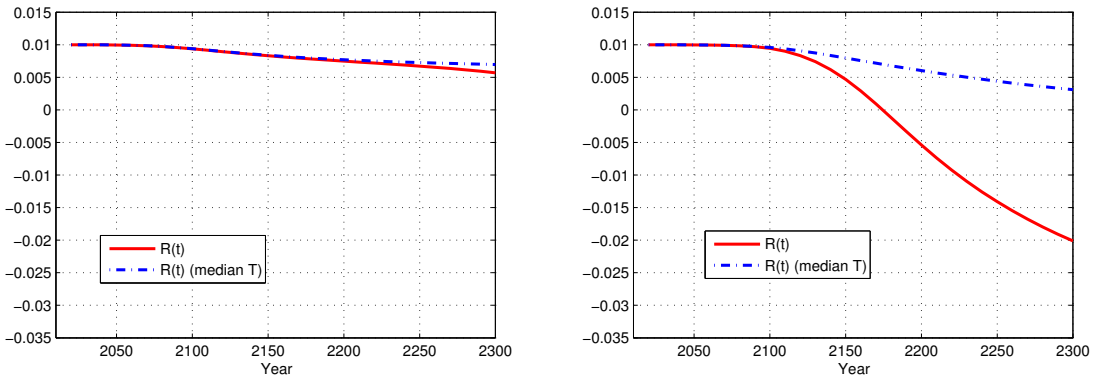


(a) low-end case, $\alpha^N$

(b) high-end case, $\alpha^W$.

Figure 6: The social discount rate $R(t)$ in the baseline calibration.

To further apprehend our results, Figure 7a presents 25 randomly drawn temperature paths from Equation (3). On these paths, the global mean temperature in 2300 varies between 3 and 8.5 °C. The corresponding risk-free rates are shown in Figures 7b and 7c for our low-end and high-end damage function, respectively. From these figures, it is immediately

[10] All these present value calculations follow by taking the relevant discount rate $R(2300)$, and calculating $e^{-290R(2300)}$. We have $e^{-290*-0.02} = 350$, $e^{-290*0.006} = 0.18$, $e^{-290*0.007} = 0.13$, and $e^{-290*0.06} \approx 0$.

apparent that the negative values of the risk-free rate on high-temperature paths are driving our results. As in Weitzman (2010), low-probability states of the world play an important role in long-term discounting, if they are associated with a sufficiently high and persistent impact. Under our stochastic A2 scenario, the persistency of the impacts follows from the inertia of the climate system and the shape of the damage function: when today's temperature and damages are high, it is very likely that tomorrow's temperature and damages will be high as well. Surprisingly, our estimates of the social discount rate turn out to be much lower than in Weitzman's stylized example, even though in our case - and in contrast to the immediate and once-and-forever shock considered by Weitzman (2010) - high-impact states of the world do not occur before the end of this century. Intuitively, this can be explained by the observation that under our stochastic A2 scenario low-probability, high impact states of the world are actually much more likely than in Weitzman's stylized example. For example, the probability that the expected rate of return is no larger than 1%, equals 0.49% in the case considered by Weitzman (2010),[11] but nearly 10% under the stochastic A2 scenario and the high-end damage function.

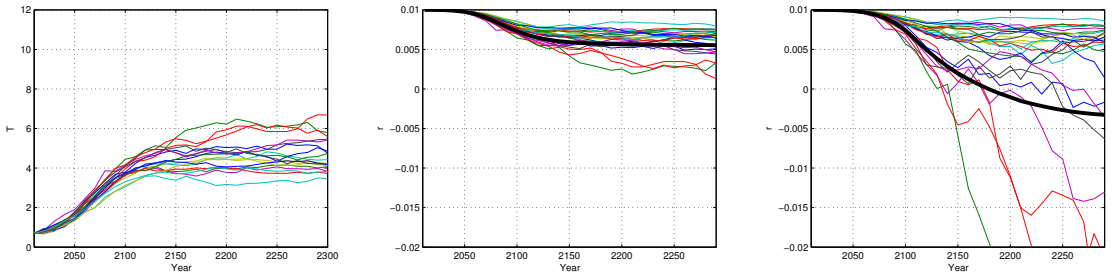


(a) $T(t)$ paths. (b) low-end damage function (c) high-end damage function

Figure 7: Simulated temperature paths (left) and the corresponding risk-free rates (middle and right).

To access the sensitivity of our results with respect to the damage function, Figure 8 shows the social discount rate for the intermediate expected rate of return $\alpha^I = 0.5(\alpha^N + \alpha^W)$. This figure reveals that - even for this intermediate case - the social discount rate becomes negative just after 2250. This can be explained by the fact that the probability that the expected rate of return is no larger than 1% is still 1.9% under the stochastic A2 scenario and the intermediate case, which remains much higher than Weitzman's 0.49%. Hence, the prime factor driving our results is not the damage function per se (although damage must surely be sizeable), but the relatively large probabilities of extreme climate change under the stochastic A2 scenario.

Table 2 shows that the term premium in the pure discount bond underlying the SDR explains a substantial part of the SDR's decreasing term structure, at least for the inter-

[11] See his Table 2 on page 9.

| Damage function | 2010 | 2050 | 2100 | 2150 | 2200 | 2250 | 2300 |
|---|---|---|---|---|---|---|---|
| Low-end | 0.0 | 0.0 | 0.0 | 0.0 | 0.0 | 0.0 | 0.0 |
| Intermediate | 0.0 | 0.0 | 0.0 | 0.0 | −0.1 | −0.2 | −0.4 |
| High-end | 0.0 | 0.0 | 0.0 | −0.3 | −0.9 | −1.0 | −0.9 |

Table 2: The decrease in the SDR resulting from the term premium in the pure discount bond.

mediate and the high-end damage functions.[12] Whereas, even in 2300, the contribution of the term premium to the SDR is almost negligible for the low-end damage function, it is sizeable for both the intermediate and high-end damage functions. For example, in 2300, the impact of the term premium in the pure discount bond on the SDR is 0.9 percentage points for the high-end and 0.4 percentage points for the intermediate damage function. Thus the term premium accounts for almost 29% of the total decrease in the SDR for the high-end and 25% for the intermediate damage function.[13] Intuitively, the term premium will only effect the SDR when temperature has a sizeable impact on aggregate risk, since, under those circumstances, a bond's potential to hedge against shifts in temperature has value. Obviously, this is the case under both the intermediate and the high-end damage function, but not so under the low-end damage function. This can also be confirmed directly from Figure 6a, where the narrow margin between the SDR on the deterministic and stochastic paths is a clear signal that aggregate risk is only changing slowly in that case.

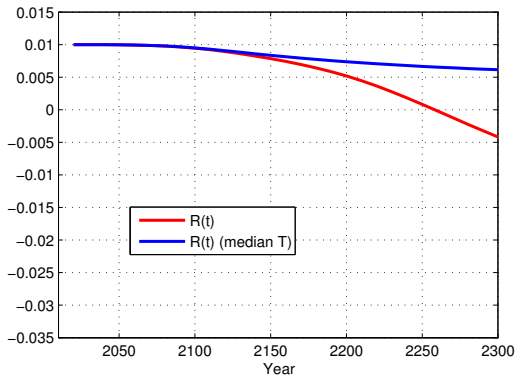


Figure 8: The social discount rate $R(t)$ under the intermediate case, $\alpha^I$.

Under our baseline assumption that unexpected changes in temperature are uncorrelated with unexpected changes in wealth, changes in the coefficient of relative risk aversion $\eta$ *do not* affect the social discount rate. To see why, use $Gs' = 0$ and rewrite (12) as

$$r(t,T) = r(t_0,T_0) + \alpha(T) - \alpha(T_0), \tag{21}$$

[12] In the Table, a '-0.0' denotes a negative number which is rounded of to one decimal.
[13] As measured from the no-climate-change value of the SDR of 1.0%.

which states that the risk-free rate at time $t$ equals the risk-free rate at time $t_0$, $r(t_0, T_0)$, plus the expected change in the rate of return on assets between $t_0$ and $t$, $\alpha(T) - \alpha(T_0)$. Subsequently, substitution of (21) into (15) and (13) gives

$$R(t) = r(t_0, T_0) - \alpha(T_0) - \frac{1}{t - t_0} \ln \mathbb{E}\left[\exp\left(\int_{t_0}^{t} -\alpha(T(\tau))d\tau\right)\right], \tag{22}$$

from which it immediately follows that the social discount rate $R(t)$ is indeed independent of the level of risk aversion $\eta$. Eq. (22) states that the time-$t$ social discount rate equals the time-zero risk-free rate $r(t_0, T_0)$ *minus* the time-0 expected rate of return on assets $\alpha(T_0)$ *plus* the time-$t$ certainty-equivalent rate of return on assets. Intuitively, the surprising result that the social discount rate is independent from the coefficient of relative risk aversion can be understood by realizing first of all that the risk-free rate at $t_0$ is not affected by changes in the coefficient of relative risk aversion $\eta$: the requirement that $r(t_0, T_0) = 0.01$ implies that we are moving along the iso-risk-free-rate line indicated by a star in Figure 1. Hence, any change in the coefficient of relative risk aversion must be fully offset by an appropriate change in the disaster size $\kappa$, leaving the risk-free rate at $t_0$ unchanged.[14] In addition, changes in risk aversion do not affect the expected rate of return on assets $\alpha(T)$. Hence, the social discount rate is unaffected by changes in risk aversion, whenever unexpected changes in temperature are uncorrelated with unexpected changes in wealth.

Inspection of Eq. (22) shows that high values of the expected rate of return will discount themselves out of existence, which implies that the certainty-equivalent rate of return on assets tends to the lowest possible value when $t$ goes to infinity (cf. Weitzman (1998)). In terms of our stochastic A2 scenario, this means that the expected rates of return belonging to low temperature paths become less important over time, whereas the expected rates of return belonging to high temperature paths become more important over time. Eventually, the expected rates of return on the highest temperature path will come to dominate the social discount rate $R(t)$.

Finally, we explore the sensitivity of our results with respect to $G_2(t)$, which captures the impact of unexpected changes in temperature on the rate of return on assets. Given the lack of empirical evidence on this parameter, we let $G_2(t)$ decline linearly from zero in 2010 to -0.025 in 2050, after which its value remains constant. This choice implies that, as of 2050, an unexpected increase of temperature by 0.5 °C decreases the rate of return by 12.5%. So, in case the expected rate of return would have been 0.06, an unexpected temperature increase of 0.5 °C would have resulted in a $0.06 * 0.125 = 0.0075$ percentage point decrease of the rate of return on assets. Figure 9 shows that - under the high-end damage function and compared to our baseline calibration - such a return volatility decreases the social discount rate in 2300 from -2.0% to -2.9%.[15] Moreover, the impact of

[14] Notice that the derivative of $r(t, T)$ with respect to $\eta$ equals $-GG' - \ln(1 + \kappa)\kappa\lambda(1 + \kappa)^{-\eta}$. Hence, in general, higher risk aversion will lower the risk-free rate.

[15] Under the low-end damage function, the correlation between unexpected changes in temperature and unexpected changes in wealth has a negligible impact on the social discount rate. Results are available from the authors upon request.

the return volatility $G_2(t)$ on the social discount rate is increasing over time until 2200, after which it roughly stabilizes. Since the return volatility itself is constant from 2050 onwards, this implies that the additional uncertainty associated with this return volatility matters more when wealth is low (and climate damages are high). Given that the risk-free rate $r(t,T)$ now depends on the level of risk aversion $\eta$, as can be seen from Eq. (12), it is to be expected that the social discount rate will depend on the level of risk aversion as well. Indeed, Figure 10 confirms that higher risk aversion is associated with a (somewhat) lower social discount rate.[16]

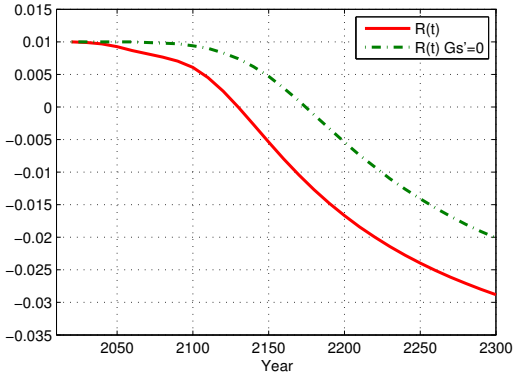


Figure 9: The social discount rate $R(t)$ under the high-end damage function for the case of time-dependent (red continuous line) and non-time-dependent (dashed green line) return volatility $G_2(t)$.

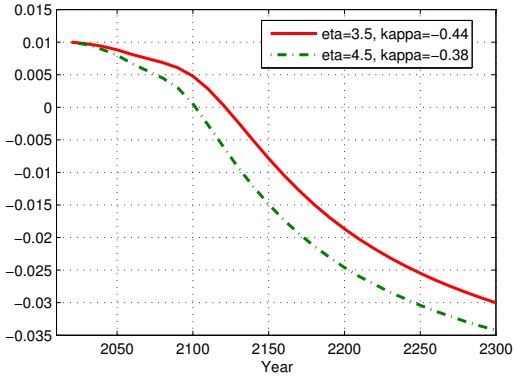


Figure 10: The social discount rate $R(t)$ under the high-end damage function and time-dependent return volatility $G_2(t)$ for the case of low risk aversion $\eta = 3.5; \kappa - 0.44$ (continuous red line) and high risk aversion $\eta = 4.5; \kappa - 0.38$ (dashed green line).

[16] In order to preserve the notion that the model should be in accordance with historically observed levels of the risk-free rate, we simultaneously vary the level of risk aversion $\eta$ and the effective disaster size $\kappa$, keeping the risk-free rate at time $t_0$ at 1%, i.e., we are moving along the iso-risk-free-rate line indicated by a star in Figure 1.

# 5 Discussion and Conclusion

This paper has shown that the far-distant future may be crucially important for long-term discounting. Even though the inertia of the earth's climate prevents a too rapid increase of surface temperature, our results suggest that a small, but non-negligible probability of an extreme temperature response in the far-distant future may drive the social discount rate to sub-zero levels. As a result, the present value of one euro three centuries hence may be everywhere between 0.18 and 350 euros.[17] Notwithstanding the extremely wide margin in this present value, what stands out is that it is much, much larger than the present value under standard exponential discounting at a constant rate. For example, taking the discount rate three centuries hence based on our low-end (Nordhaus) damage function results in a present value that is almost ten times larger than the present value that is obtained by using the much-criticized, and supposedly too low, discount rate of 1.4% of the Stern Review.[18]

Importantly, our results suggest that the key issue determining the level of the social discount rate is not how large damages are or might be (although they surely must be substantial to invoke a sizeable response in the social discount rate), but how large the probability of extreme climate change is. Under our stochastic A2 scenario, this probability is sufficiently large to make the logic of exponential discounting in the far distant future (almost) completely irrelevant. Importantly, our limited state of knowledge regarding the high-temperature damages does not interfere with this conclusion, as the logic of exponential discounting is severely diminished under both the quadratic and reactive specification of the damage function. The shape of the damage function is, however, critically important for the occurrence of negative discount rates. Whereas the social discount rate will be strictly positive under our low-end, quadratic damage function, it may reach sub-zero levels under more extreme specifications of the damage function, such as the reactive damage function proposed by Weitzman (2012). In that respect, it must be noticed that in the 'short' run, i.e., this century, the impact of the damage function on the social discount rate seems to be severely constrained by the inertia of the earth's climate system. In particular, our results suggest that the probability of a high-enough temperature response under our stochastic A2 scenario is insufficient to drive this century's social discount rate away from its conventional, no climate change, level.

Thus, our results seem to imply that, by and large, exponential discounting will still produce consistent valuations in *this* century. This surprising result is, however, counterbalanced by the observation that the SDR in the next century will be hugely affected by climate change unless we are more or less certain that damages from climate change will be relatively low. In many instances, it might even become negative. If it does, the social cost of carbon will explode to infinity, thus providing a clear signal that society should reduce emissions whatever the cost. Of course, the price signal of a discount rate on the IPCC A2

[17] From $R(2300) = 0.006$ and $-0.02$, we have $e^{-290*0.006} = 0.18$ and $e^{-290*-0.02} = 350$, respectively.
[18] $e^{-290*0.0061} = 0.17$ is about ten times larger than $e^{-290*0.014} = 0.017$.

path is valid only for marginal mitigation projects, i.e., for investment projects that do not affect aggregate welfare, aggregate emissions and global temperature. When considering non-marginal investment projects, the probability of extreme climate change will decline, and it is likely that the SDR will increase from the levels reported in this paper. As a result, the social cost of carbon will decline, thereby reducing the value of the marginal mitigation project on the new emissions path. The extent to which our results carry over to more moderate emission scenarios, such as the IPCC B2 scenario, is therefore crucial for cost-benefit analysis of climate change and remains a interesting topic for further research.

# Appendix

## A Construction temperature process

In this section we explain our procedure to calibrate the temperature dynamics,

$$dT(t) = \mu(t, T(t))dt + s_2(t, T(t))d\omega_2(t). \tag{23}$$

We consider a time horizon $[t_0, t^*] = [2010, 2300]$ years and set $T(t_0) = 0.7$ °C. For the construction of functions $\mu(t, T)$ and $s_2(t, T)$ we use so-called 'reference paths'. The reference paths are represented by the black lines in Figure 11. The slopes of the reference paths are known and for each year and temperature the value of function $\mu(t, T)$ is interpolated using the three reference paths and their slopes.

During the first years, the volatility $s_2(t_0, T)$ equals 0.03. For temperature values close to the lowest reference path, volatility will stay close to this value. For temperatures along the mid reference path, volatility will increase moderately and linearly in time, up to $s_2(2100, T) = 0.045$ after 90 years and will stay constant thereafter. Along the highest reference path, volatility will increase linearly in time up to the high value $s_2(2100, T) = 0.2$. After 90 years the volatility is assumed to decreases to the value $s_2(t^*, T) = 0.08$ at the terminal time $t^*$.

The lowest and mid reference paths tend to an equilibrium temperature, while equilibrium is not reached yet on the highest reference path. Again interpolation with three reference paths is used to find the value $s_2(t, T)$. With this, we capture the feature of high uncertainty on high temperature paths (Roe, 2009).

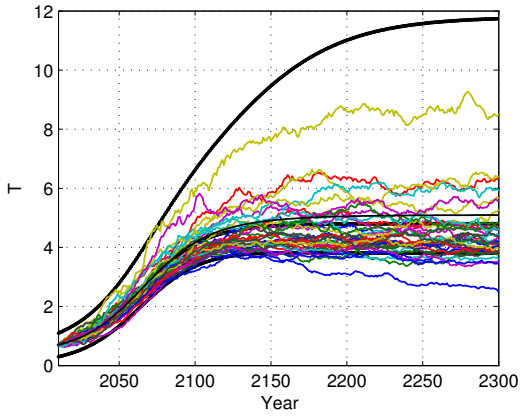


Figure 11: The three reference paths (black lines) used for construction of the temperature model together with twenty five simulated temperature paths.

# B Numerical method

Our numerical method to solve the optimal portfolio problem is based on the dynamic programming principle and Fourier cosine series expansions, called the COS method (Fang and Oosterlee, 2008, 2009; Ruijter and Oosterlee, 2012). In this section we give some details of this method.

We use utility function $U(C) = C^{1-\eta}/(1-\eta)$, so $p = p(t,T)$ (see Section 2). The value function under economic equilibrium is given by

$$J(t,W,T) = \max_{\{C\geq 0\}} \mathbb{E}^{t,W,T}\left[\int_t^{t^*} e^{-\delta(\tau-t)}U(C(\tau))d\tau + e^{-\delta(t^*-t)}U(W(t^*))\right]. \tag{24}$$

$W$ and $T$ represent the current wealth and temperature level, respectively, and the wealth dynamics evolve according to

$$dW(t) = \left(\alpha(T(t)) - \kappa\lambda - p(t,T(t))\right)W(t)dt + W(t)Gd\omega(t) + W(t)\kappa dq(t). \tag{25}$$

We switch to the log-domain, $X(t) := \ln W(t)$, with

$$dX(t) = \left(\alpha(T(t)) - \kappa\lambda - p(t,T(t)) - \tfrac{1}{2}GG'\right)dt + Gd\omega(t) + \ln(1+\kappa)dq(t). \tag{26}$$

The corresponding value function is

$$I(t,x,T) := \max_{\{p\geq 0\}} \mathbb{E}^{t,x,T}\left[\int_t^{t^*} e^{-\delta(\tau-t)}U(e^{X(\tau)}p(\tau))d\tau + e^{-\delta(t^*-t)}U(e^{X(t^*)})\right]. \tag{27}$$

$x$ and $T$ represent the current log-wealth and temperature level, respectively. We take an equidistant grid of control times $t_0 < t_1 < \ldots t_m < \ldots < t_{\mathcal{M}} = t^*$, with time-step $\Delta t := t_{m+1} - t_m$ and approximate processes by their discrete variants. The log-wealth and temperature processes are discretized by an Euler scheme (see, for example, (Platen, 2010; Korn, Korn and Kroisandt, 2010)).

We approximate the value function by (see (Kushner and Dupuis, 2000))

$$I(t_m, x, T) = \max_{\{p_i\}} \mathbb{E}^{t_m, x, T} \left[ \sum_{i=m}^{\mathcal{M}-1} e^{-\delta(t_i - t_m)} U(e^{X_{t_i}} p_i) \Delta t + e^{-\delta(t_{\mathcal{M}} - t_m)} U(e^{X_{t_{\mathcal{M}}}}) \right]. \tag{28}$$

The discrete-time value function converges in $\mathcal{M}$ to the continuous-time variant. The dynamic programming principle gives us

$$\begin{aligned} I(t_m, x, T) &= \max_{p_m} \mathbb{E}^{t_m, x, T} \left[ U(e^x p_m) \Delta t + e^{-\delta \Delta t} I(t_{m+1}, X(t_{m+1}), T(t_{m+1})) \right] \\ &:= \max_{p_m} \Big[ F(x, p_m) + c(t_m, x, T, p_m) \Big]. \end{aligned} \tag{29}$$

The first part in the maximization operator is called the running utility function, the second part is the continuation value. The continuation value can be approximated by using the 2D-COS formula (Ruijter and Oosterlee, 2012) and the problem is solved backwards in time (Bellman's principle of optimality). The method parameters are

$$[a_1, b_1],\ [a_2, b_2] \quad \text{computational domain log-wealth } (1^{th} \text{ dimension}) \text{ and temperature } (2^{nd} \text{ dimension}), \text{ respectively,} \tag{30}$$

$$N_1,\ N_2 \quad \text{number of Fourier cosine coefficients in } 1^{th} \text{ and } 2^{nd} \text{ dimension, respectively,} \tag{31}$$

$$J_1, J_2 \quad \text{number of control regions in } 1^{th} \text{ and } 2^{nd} \text{ dimension, respectively.} \tag{32}$$

Function $c(t_m, x, T, p_m)$ is called the continuation value and is approximated by a two-dimensional COS formula

$$\begin{aligned} \hat{c}(t_m, x, T, p_m) = e^{-\delta \Delta t} \sum_{k_1=0}^{N_1-1}{}' \sum_{k_2=0}^{N_2-1}{}' \frac{1}{2} \Bigg[ & \operatorname{Re}\left( \varphi_{levy}\left( \tfrac{k_1 \pi}{b_1 - a_1}, + \tfrac{k_2 \pi}{b_2 - a_2} \middle| x, T, p_m \right) e^{i k_1 \pi \frac{x - a_1}{b_1 - a_1}} e^{i k_2 \pi \frac{T - a_2}{b_2 - a_2}} \right) \\ & + \operatorname{Re}\left( \varphi_{levy}\left( \tfrac{k_1 \pi}{b_1 - a_1}, - \tfrac{k_2 \pi}{b_2 - a_2} \middle| x, T, p_m \right) e^{i k_1 \pi \frac{x - a_1}{b_1 - a_1}} e^{-i k_2 \pi \frac{T - a_2}{b_2 - a_2}} \right) \Bigg] V_{k_1, k_2}(t_{m+1}). \end{aligned} \tag{33}$$

$\operatorname{Re}(.)$ denotes taking the real part of the input argument. $\varphi(., .|x, T, p_m)$ is the *bivariate conditional characteristic function* of $(X(t_{m+1}), T(t_{m+1}))$, given $(X(t_m), T(t_m)) = (x, T)$. The Fourier cosine coefficients of the value function are given by

$$V_{k_1, k_2}(t_m) := \tfrac{2}{b_1 - a_1} \tfrac{2}{b_2 - a_2} \int_{a_2}^{b_2} \int_{a_1}^{b_1} I(t_m, y, \varsigma) \cos\left( k_1 \pi \tfrac{y - a_1}{b_1 - a_1} \right) \cos\left( k_2 \pi \tfrac{\varsigma - a_2}{b_2 - a_2} \right) dy d\varsigma. \tag{34}$$

We divide the domain $[a_1, b_1] \times [a_2, b_2]$ into rectangular subdomains $\mathcal{D}^{q_1,q_2} = [z_{q_1}, z_{q_1+1}] \times [w_{q_2}, w_{q_2+1}]$. For each subdomain we determine the optimal control value $p^{q_1,q_2}$, then

$$\begin{aligned} V_{k_1,k_2}(t_m) &= \tfrac{2}{b_1-a_1}\tfrac{2}{b_2-a_2} \sum_{q_1,q_2} \iint_{\mathcal{D}^{q_1,q_2}} F(y, p^{q_1,q_2}) \cos\left(k_1\pi \tfrac{y-a_1}{b_1-a_1}\right) \cos\left(k_2\pi \tfrac{\varsigma-a_2}{b_2-a_2}\right) dyd\varsigma \\ &+ \tfrac{2}{b_1-a_1}\tfrac{2}{b_2-a_2} \sum_{q_1,q_2} \iint_{\mathcal{D}^{q_1,q_2}} c(t_m, y, \varsigma, p^{q_1,q_2}) \cos\left(k_1\pi \tfrac{y-a_1}{b_1-a_1}\right) \cos\left(k_2\pi \tfrac{\varsigma-a_2}{b_2-a_2}\right) dyd\varsigma \\ &:= \sum_{q_1,q_2} U_{k_1,k_2}(\mathcal{D}^{q_1,q_2}, p^{q_1,q_2}) + \sum_{q_1,q_2} C_{k_1,k_2}(t_m, \mathcal{D}^{q_1,q_2}, p^{q_1,q_2}). \end{aligned} \tag{35}$$

The terms $U_{k_1,k_2}$ are time-independent and known analytically or approximated. The term $C_{k_1,k_2}$ are approximated by using an FFT algorithm. With this we can recover the coefficients $V_{k_1,k_2}(t_m)$ backwards in time and solve the optimal portfolio problem. All parameters are chosen such that the numerical approximations give rise to converged solutions, values and functions of interest.

# References


**Ahn, C.M., and H.E. Thompson.** 1988. "Jump-Diffusion Processes and the Term Structure of Interest Rates." *The Journal of Finance*, 43(1): 155–174.

**Barro, R.J.** 2006. "Rare Disasters and Asset Markets in the Twentieth Century." *The Quarter Journal of Economics*, 121(3): 823–866.

**Barro, R.J.** 2009. "Rare Disasters, Asset Prices, and Welfare Costs." *American Economic Review*, 99(1): 243–264.

**Barro, R.J.** 2013. "Environmental Protection, Rare Disasters, and Discount Rates." *NBER Working Paper 19258.*

**Cai, Y., K.L. Judd, and T.S. Lontzek.** 2013. "The Social Cost of Stochastic and Irrreversible Climate Change." *NBER Working Paper 18704.*

**Cox, J. C., J.E. Ingersoll, and S.A. Ross.** 1981. "A Re-Examination of Traditional Hypotheses about the Term Structure of Interest Rates." *The Jou*, 36(4): 769–799.

**Cox, J. C., J.E. Ingersoll, and S.A. Ross.** 1985*a*. "An Intertemporal General Equilibrium Model of Asset Prices." *Econometrica*, 53(2): 363–384.

**Cox, J. C., J.E. Ingersoll, and S.A. Ross.** 1985*b*. "A Theory of the Term Structure of Interest Rates." *Econometrica*, 53(2): 385–407.

**Fang, F., and C.W. Oosterlee.** 2008. "A Novel Pricing Method For European Options Based On Fourier-Cosine Series Expansions." *SIAM Journal on Scientific Computing*, 31(2): 826–848.

**Fang, F., and C.W. Oosterlee.** 2009. "Pricing Early-Exercise and Discrete Barrier Options by Fourier-Cosine Series Expansions." *Numerische Mathematik*, 114(1): 27–62.

**Gollier.** 2014. "Gamma Discounters are Short-Termist." *TSE Working Paper Series 14-499*. mimeo.

**Groom, B., P. Koundori, K. Panipoulou, and T. Pantelides.** 2007. "Declining Discount Rates: How Much Does Model Selection Affect the Certainty Equivalent Discount Rate?" *Journal of Applied Econometrics*, 22(3).

**Korn, R., E. Korn, and G. Kroisandt.** 2010. *Monte Carlo Methods and Models in Finance and Insurance. Financial Mathematics*, CRC Press/Taylor & Francis.

**Kushner, H.J., and P.G. Dupuis.** 2000. *Numerical Methods for Stochastic Control Problems in Continuous Time. Applications of Mathematics*, Springer.

**Masoliver, J., M. Montero, and J. Perelló.** 2013. "Uncertain Growth and the Value of the Future." *Cowles Foundation Discussion Papers 1930, Yale University.*

**Nakićenović, N., J. Alcamo, G. Davis, B. De Vries, J. Fenhann, S. Gaffin, K. Gregory, A. Grübler, T. Jung, T. Kram, et al.** 2000. "IPCC special report on emissions scenarios (SRES)."

**Newell, R., and W. Pizer.** 2003. "Discounting the Benefits of Climate Change Mitigation: How Much do Uncertain Rates Increase Valuations?" *Journal of Environmental Economics and Management*, 46(1): 52–71.

**Nordhaus, W.D.** 2008. *A Question of Balance: Weighing the Options on Global Warming Policies.* Yale University Press, New Haven & London.

**Platen, E.** 2010. *Numerical Solution of Stochastic Differential Equations with Jumps In Finance. Stochastic Modelling and Applied Probability*, Springer-Verlag.

**Reinhart, C. M., and M. B. Sbrancia.** 2011. "The Liquidation of Government Debt." *NBER Working Paper 16893.*

**Rietz, T.A.** 1988. "The Equity Risk Premium: a Solution." *Journal of Monetary Economics*, 22(1): 117–131.

**Roe, G.H.** 2009. "Feedbacks, Timescales, and Seeing Red." *Annual Review of Earth and Planetary Sciences*, 37: 93–115.

**Roe, G.H., and Y. Bauman.** 2012. "Climate Sensitivity: Should the Climate Tail wag the Policy Dog?" *Climatic Change*, published online 26 September.

**Ruijter, M.J., and C.W. Oosterlee.** 2012. "Two-Dimensional Fourier Cosine Series Expansion Method for Pricing Financial Options." *SIAM Journal on Scientific Computing*, 34(5): B642–B671.

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

**Tol, R.S.J.** 2009. "The Economic Effects of Climate Change." *Journal of Economic Perspectives*, 23(2): 29–51.

**U.S. Department of the Treasury.** 2014. *http://www.treasury.gov/resource-center/data-chart-center/interest-rates/Pages/TextView.aspx?data=billRatesYear&year=2011*.

**Weitzman, M.** 1998. "Why the Far-Distant Future Should Be Discounted at Its Lowest Possible Rate." *Journal of Environmental Economics and Mangament*, 36(3): 201–208.

**Weitzman, M.L.** 2010. "Risk-adjusted Gamma Discounting." *Journal of Environmental Economics and Management*, 60(1): 1–13.

**Weitzman, M.L.** 2012. "GHG Targets as Insurance Against Catastrophic States Climate Damages." *Journal of Public Economic Theory*, 14(2): 221–244.

Publisher:

CPB Netherlands Bureau for Economic Policy Analysis
P.O. Box 80510 | 2508 GM The Hague
T (070) 3383 380

December 2014 | ISBN 978-90-5833-669-9