Applying Asset Pricing Theory to Calibrate the Price of Climate Risk

Kent D. Daniel*, Robert B. Litterman†, and Gernot Wagner‡

25 October 2016

Abstract

Pricing greenhouse gas emissions is a risk management problem. It involves making trade-offs between consumption today and unknown and potentially catastrophic damages in the (distant) future. The optimal carbon price is based on society's willingness to substitute consumption across time and across uncertain states of nature. A large body of work in macroeconomics and finance has attempted to infer societal preferences using the observed behavior of asset prices, and has concluded that the standard preference specifications are inconsistent with observed asset valuations. This literature has developed a richer set of preferences that are more consistent with asset price behavior.

In this paper, we explore the implications of these richer preference specifications for the Social Cost of Carbon (SCC), the expected discounted damage of each marginal ton of carbon emissions at an optimal emissions reductions pathway. We develop a simple discrete-time model in which the representative agent has an Epstein-Zin preference specification, and in which uncertainty about the effect of carbon emissions on global temperature and on eventual damages is gradually resolved over time. In our model the SCC is equal to the value of the carbon emissions price at any given point in time that maximizes the utility of the representative agent at that time. We embed a number of features including tail risk, the potential for technological change, and backstop technologies. When coupled with the potential for low-probability, high-impact outcomes, our calibration allows us to decompose the SCC into the expected damages and the risk-premium. In contrast to most modeled carbon price paths, our calibration suggests a high SCC today that is expected to decline over time. It also points to the importance of backstop technologies and, in contrast to standard specifications, to potentially very large deadweight costs of delay. We find, for example, that with damage distributions calibrated to an SCC of $40, a value associated with only a small risk premium, the deadweight loss in utility associated with delaying the implementation of optimal pricing by 15 years is equivalent to a 6% loss of consumption.

Keywords: Asset pricing, discounting, climate change, global warming; Epstein-Zin utility.

JEL code: D81, G11, Q54.

---

* Columbia Business School, New York, NY, and National Bureau of Economic Research, Cambridge, MA (kd2371@columbia.edu)

† Kepos Capital, New York, NY, USA (blitterman@gmail.com)

‡ Harvard John A. Paulson School of Engineering and Applied Sciences, Cambridge, MA, USA (gwagner@fas.harvard.edu).

1. Introduction

Evidence continues to mount that the atmosphere's capacity to absorb greenhouse gases (GHGs) is limited, and that additional GHG emissions lead to global warming and societal damages. The relationship between these damages and GHG emissions is uncertain. Following Coase (1960), because property rights to the atmosphere are poorly defined, self-interested individuals have too little incentive to curtail emissions of GHGs. However, optimal usage of the atmosphere's capacity to absorb GHGs can be obtained when individuals are charged the full social cost of each ton of carbon dioxide (CO_2) they emit into the atmosphere, or conversely the benefits that accrue to society with the reduction of CO_2 emissions by one ton. That cost of putting an additional ton of CO_2 into the atmosphere at any given time t, assuming an optimal emissions reductions pathway throughout, is commonly known as the Social Cost of Carbon (SCC).[1] This paper addresses the determination of the SCC.

The modern approach to asset pricing recognizes that the SCC is determined by appropriate discounting of the marginal benefits of reducing emissions by one ton at all future times and across all states of nature (Duffie, 2010; Hansen and Richard, 1987). In practice this can be done by discounting those future benefits not by a discount factor which is invariant across states of nature, but rather by a stochastic discount factor which is appropriate to each possible outcome.

Until recently, the climate-economic literature has largely ignored the pricing of the risk in the payoffs resulting from the mitigation of climate emissions, and where it has done so it has used a constant relative risk aversion (CRRA)/constant-elasticity of substitution (CES)/power utility preference specification inconsistent with the evidence from financial economics.[2] We show here that the risk premium embedded in the SCC is likely to be large and an important component when calculated with a preference specification consistent with the historical asset returns.

The valuation assigned to different traded assets suggests that society is willing to pay only a small premium to substitute consumption across time, but a large premium to substitute across different states of nature. For example, between 1871 and 2012, a

---

[1] The assumption of an optimal emissions reductions pathway beginning at time 0 is an important, albeit often implicit, assumption in defining the SCC_t path across time. Note, however, that the U.S. SCC explicitly assumes no such path. Instead, it focuses on pricing the marginal ton of emissions given the current trajectory (U.S. Government Interagency Working Group on Social Cost of Carbon, 2015).

[2] Notable recent exceptions include Ackerman, Stanton, and Bueno (2013) and especially work by Christian Traeger and also Derek Lemoine; e.g. Crost and Traeger (2014), Jensen and Traeger (2014), Lemoine and Traeger (2014), Lemoine and Traeger (2016), Lemoine (2015) and especially also Traeger (2015). Others have considered climate risk in a different context; e.g. Nordhaus and Boyer (2000) incorporating risk based on an expert survey (Nordhaus, 1994). See Kopp *et al.* (2016) for a discussion of the latter. Nordhaus (2014, 2013) presents the latest results, based on the well-known DICE model going back to Nordhaus (Nordhaus, 1992, 1991). See e.g. Convery and Wagner (2015), Pindyck (2013), Stern (2013), and Wagner and Weitzman (2015) for extensive critiques of prior treatments of risk in SCC calculations. Golosov *et al.* (2014), among others, pursues another extension of standard climate-economy models, employing a dynamic stochastic general-equilibrium (DSGE) model, while still relying on a CRRA utility function.

portfolio of U.S. bonds earned an average annual real return of 1.6 percent, and a diversified portfolio in U.S. stocks earned an average annual real return of 6.4 percent. Society presumably discounts equity payoffs at a far higher discount rate because equities earn large returns in good economic times (when marginal utility is low) but often perform poorly precisely when economic growth is low (and marginal utility is high).

Conversely, society is willing to pay handsomely for the right pattern of cash flows across states: these numbers imply that a portfolio which was short US equities, providing insurance against bad economic outcomes, earned an annual return of -4.8%/year over this long period. Society is clearly willing to pay a large premium for insurance against risky outcomes, and discounts payoffs that are "risky" at a high rate.[3]

The *high* historical equity premium, combined with the low historical volatility of consumption growth, suggests that society is unwilling to substitute consumption across states of nature at some future point in time. In contrast, the *low* risk-free rate in combination with the high average consumption growth rate over the past 150 years suggests that agents are far more willing to substitute consumption across time. These two empirical regularities are not consistent with a CRRA model of preferences.

We approach climate change as a standard asset pricing problem. Carbon in the atmosphere is an 'asset'—albeit one with negative payoffs—and ought to be treated as such. Our model uses a state-contingent discount rate, calibrated to the returns over time of financial assets. In contrast to the standard CRRA utility function used in most climate studies, we use here a utility function proposed by Epstein and Zin and used throughout the asset pricing literature (Epstein and Zin, 1991, 1989). It has CRRA utility as a special case and also allows for differences between the intertemporal marginal rate of substitution (IMRS)/intertemporal elasticity of substitution (IES) and risk aversion, which allows us to calibrate to standard financial returns, in particular the equity risk premium and risk-free interest rates.

An important property of non-time separable models such as the Epstein-Zin model is that agents' utility depends on not just levels of consumption in each state, but also on the way in which uncertainty is resolved over time. To capture the resolution of uncertainty over time, we employ a discrete time binomial tree model like that employed in many pricing financial-economic modeling applications (see Cox, Ross and Rubinstein (1979) for an early example), building on a related approach in Summers and Zeckhauser (2008). Different states in the tree represent different degrees of fragility of the environment which, when combined with the level of greenhouse gases in the atmosphere, imply different damages, different consequences for the utility of the representative agent. Information about the state is revealed over time, and in each period the agent chooses a level of emissions mitigation that maximizes his expected discounted utility, based on the information available at that time. The optimal level of emissions mitigation is obtained when the reduction in utility from additional

[3] Based on data collected by Shiller (2000) and since continuously updated: http://www.econ.yale.edu/~shiller/data.htm

expenditure on mitigation at each point in time is exactly offset by the probability-weighted increase in utility from reductions in damages in future states.

Figure 1 shows the profound implications. We choose damage parameters such that under a standard constant relative risk-aversion utility function with the IMRS calibrated to 0.9 in order to reflect a 2.74% yield on a zero-coupon bond that matures at the end of the last period, the optimal carbon price today is around $38 per ton.[4]

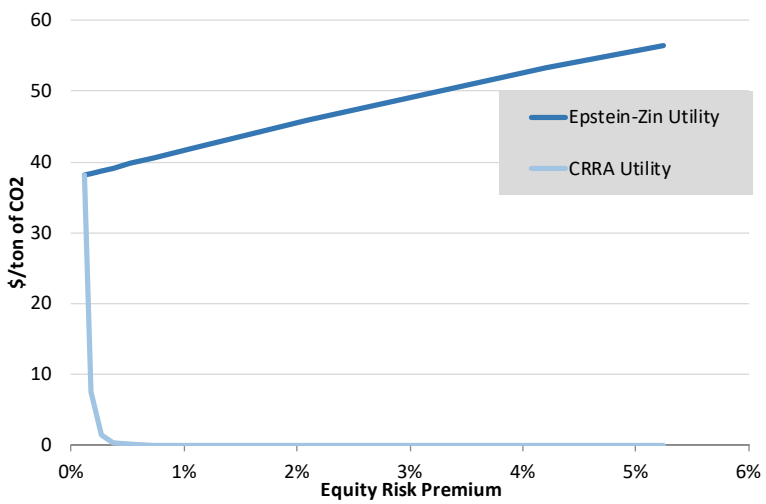

Figure 1—Using Epstein-Zin utility functions results in increasing carbon prices with increasing risk aversion translated into the implied equity risk premium using Weil (1989)'s conversion, while holding the implied market interest rate stable at 2.74%

The problem with the CRRA specification used here is that this specification embeds the assumption that agents' willingness to substitute consumption across states of nature is the same as their willingness substitute consumption over time. Thus, an increase in the coefficient of risk aversion (or stated differently, a decreased elasticity of substitution across states), is necessarily linked to a decreased IMRS. Given the fact that consumption grows at a rate of about 2%/year, an unwillingness to substitute across time leads to a (counterfactually) high risk-free discount rate. Since consumption damages occur far into the future, a CRRA utility function with a high level of risk-aversion (and a reasonable rate of time preference) necessarily implies a high discount rate for these damages, and a low SCC.

In contrast, Epstein-Zin utility allows for separation of the coefficient of risk-aversion and the IMRS, consistent with equity-premium/risk-free rate puzzle. With an Epstein-Zin specification, holding the IMRS fixed at 0.9 and increasing the degree of risk aversion, the SCC increases, while the real interest rate remains at around 2.74%/year.[5]

[4] This is close to the official SCC of $40 for a ton of CO_2 released in 2015, in 2015 US$ (U.S. Government Interagency Working Group on Social Cost of Carbon, 2015).

[5] The exact interest rate in our Epstein-Zin calibration is almost independent of the risk aversion coefficient. It is 2.742%/year with a coefficient of relative risk aversion of 1.11 (=1/0.9) (equivalent to

As the level of risk aversion is raised from a very low level to a level consistent with the historically observed equity-risk premium, the optimal carbon price increases from $38 to over $55 per ton.

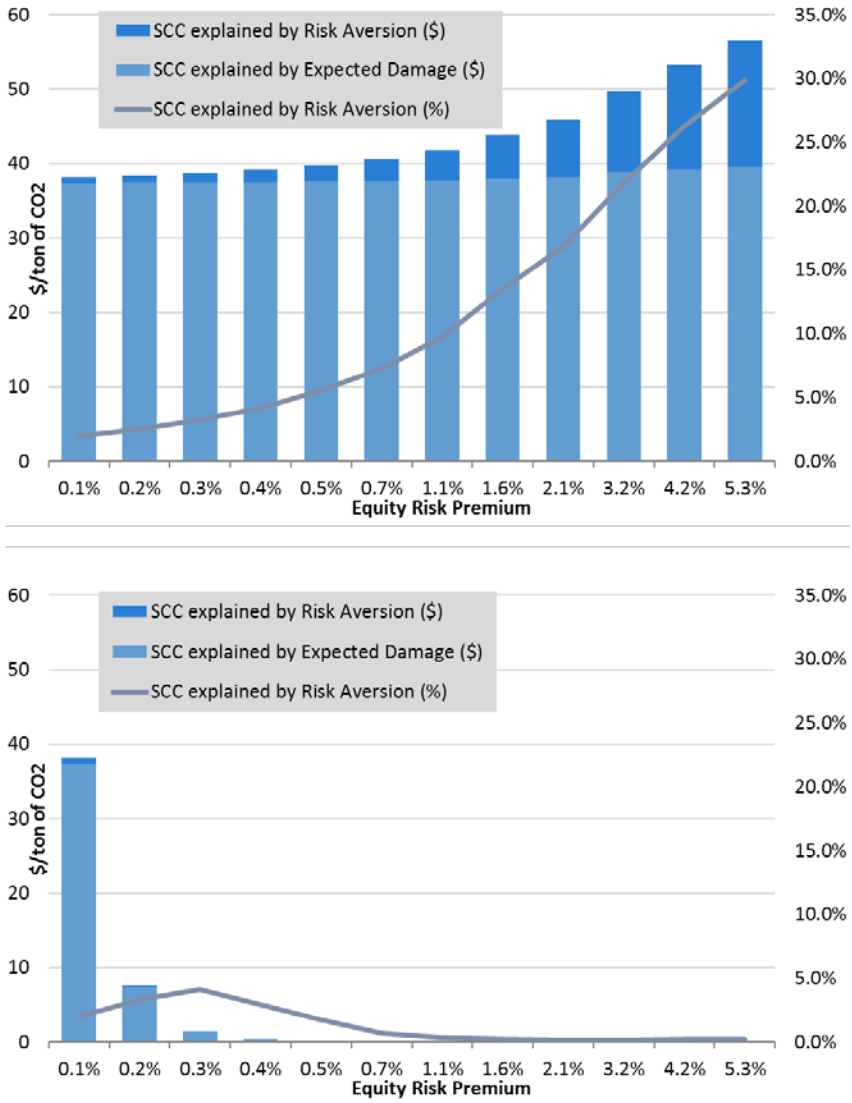

Figure 2—Decomposition of the Social Cost of Carbon (SCC) into risk aversion and expected damage portions with Epstein-Zin (top panel) and constant relative risk aversion (CRRA, bottom panel) preferences

---

CRRA), and reaches 2.744%/year with a coefficient of relative risk aversion of 50. Bansal and Yaron (2004) are able to match the equity premium with a far lower coefficient of risk-aversion, owing to the presence of shocks to the long-term growth rate of consumption in their model, which are correlated with equity returns. Similarly, in the model we present here, a link between higher climate fragility and lower consumption growth rates would lead to a higher SCC with a lower coefficient of risk-aversion. In general, climate damages hitting growth rates rather than levels of GDP can have a significant effect on the SCC (Convery and Wagner, 2015; Heal and Park, 2016; Wagner and Weitzman, 2015).

We further decompose the optimal SCC into a risk aversion and an expected damages component (Figure 2).[6] This decomposition similarly varies widely with the assumed equity risk premium—and it crucially depends on the distinction between using Epstein-Zin versus CRRA preferences (top and bottom panel, respectively).

2. The model

We now proceed to solve for the optimal carbon price as a function of time and of information about the earth's fragility. We begin with a 'business as usual' (BAU) scenario as our baseline, assuming constant consumption growth and GHG emissions that grow over time without mitigation. Mitigating emissions is costly. Hence, assuming no government action to price carbon, atomistic agents do zero mitigation. However, as GHGs build up in the atmosphere, temperatures rise. As a result, a fraction of the baseline consumption is lost to damages. The damages as a function of mitigation are not known ex-ante, but are rather a function of the earth's fragility at that time, θ_t. Each period of the model, agents learn more about the level of fragility, but they only know the actual fragility in the final two periods of the model.

These assumptions simplify reality in two important ways: As the only unknown in our model is the earth's fragility, θ_t, we do not allow for interactions of shocks to fragility with those to other state variables (e.g., productivity). The second simplification is the assumption of full knowledge of θ in period $T-1$ (in 2300 in our base case). Important aspects of climate science are deeply and persistently uncertain, and science may not learn the true θ_t at a time scale relevant to policy (Wagner and Zeckhauser, 2016; Zeckhauser, 2006). We attempt to solve the second problem by delaying the complete resolution of uncertainty until 2300.

To determine the optimal level of mitigation, we next ask what level of mitigation a single, optimizing, representative agent would choose at each point in the tree (i.e., at each time and state). Note that, in a representative agent framework, there are no externalities; the agent internalizes any damage done to the atmosphere by failing to mitigate. Thus, solving for the level of mitigation the optimizing representative agent selects gives us the socially optimal level of mitigation at each point in time and for each level of fragility.

Once we know the optimal level of mitigation by time and state, we move back to the atomistic agent setting by calculating the SCC for that time and state. To do this, we determine the price on carbon that has to be imposed in the atomistic agent setting that will implement the optimal level of mitigation (as determined in the representative agent setting). Intuitively, the optimal price is the marginal benefit to society of reducing carbon emissions, where the marginal benefit is evaluated at the social optimum. We determine this optimal price by inverting the function that gives the level of mitigation as a function of the price of carbon.

[6] Section 2.4 provides details of the risk decomposition presented in Figure 2.

The setting of our model is a discrete time, endowment economy with a single representative agent. In each period $t\epsilon\{0,1,2,\dots,T\}$, the agent is endowed with a certain amount of the consumption good, $\bar{c}_t$. However, the agent is not able to consume the full endowed consumption for two reasons: climate change and climate policy. First, in periods $t\epsilon\{1,2,\dots,T\}$, some of the endowed consumption may be lost due to climate change damages. Second, in periods $t\epsilon\{0,1,2,\dots,T-1\}$, the agent may elect to spend some of the endowed consumption to reduce his impact on the climate. The resulting consumption c_t, after damages and mitigation costs are taken into account, is given by:

(1) $$c_0 = \bar{c}_0 \cdot \big(1 - \kappa_0(x_0)\big)$$

(2) $$c_t = \bar{c}_t \cdot \big(1 - D_t(X_t, \theta_t) - \kappa_t(x_t)\big), for\ t\epsilon\{1,2,\dots,T-1\}$$

(3) $$c_T = \bar{c}_T \cdot (1 - D_T(X_T, \theta_T))$$

In equations (2) and (3), the climate damage function $D_t(X_t, \theta_t)$ captures the fraction of endowed consumption that is lost due to damages from climate change. If $D_t(X_t, \theta_t) = 0$, the agent would receive the full consumption endowment. However, damages from climate change can push D_t above zero. D_t depends on two variables: X_t, which we define as the cumulative GHG mitigation up to time t, and θ_t, a parameter that characterizes the uncertain relation between the level of GHGs in the atmosphere and consumption damages. θ_t evolves stochastically as described in section 2.3.

Cumulative mitigation X_t, in turn, depends on the level of mitigation in each period from 0 to t, which is given by:

(4) $$\mathrm{X_t} = \frac{\sum_{s=0}^{t} g_s \cdot x_s}{\sum_{s=0}^{t} g_s},$$

where g_s is the flow of GHG emissions into the atmosphere in period s, for each period up to t, absent any mitigation.[7] The level of mitigation at any time s is given by x_s, where $x_s = 0$ denotes no climate action at time s, and $x_s = 1$ denotes full mitigation, or equivalently that there zero net flow of new GHG emissions into the atmosphere in period s. One might imagine that mitigation x_s should be restricted to be below 1. However, in our baseline analysis we allow for the use of a backstop technology, a technology for pulling CO_2 directly out of the atmosphere, typically called carbon dioxide removal (CDR) or, confusingly, direct carbon removal (DCR). In our baseline simulation the backstop technology is employed in large scale, which results in mitigation above 100% in those future states in which the climate turns out to be fragile (see the discussion in Sections 2.2.1 and 3.1).

Mitigation reduces the stock of GHGs in the atmosphere and leads to lower climate damages, and hence to higher future consumption. However, mitigating GHG emissions is costly. Mitigating a fraction x_t of emissions costs a fraction $\kappa_t(x_t)$ of the endowed consumption. We describe the details of the cost function, and our calibration, in Section 2.2.

[7] This means that the cumulative GHG emissions that must be absorbed into the atmosphere or oceans is G_t*(1-X_t), where $G_t = \sum_{s=0}^{t} g_s$ denotes the cumulative emissions under the BAU scenario.

In the framework we propose, the representative agent's optimization problem involves trading off the (known) costs of climate mitigation against the uncertain future benefits associated with mitigation. The agent does this by solving the dynamic optimization problem to determine the optional path of mitigation, $x_t^*(\theta_t)$, so as to maximize lifetime utility at each time and for each state of nature.

To make the solution tractable, we model the resolution of uncertainty about climate damage with a binomial tree, discussed in detail in Section 3.1 (see also Figure 10). Our baseline analysis uses a 7-period tree, beginning in 2015. An initial mitigation decision is made in 2015, and subsequent mitigation decisions are made after information is revealed about climate fragility and the resulting damages in years 2030, 2060, 2100, 2200, and 2300. The final period, in which consumption simply grows at a constant rate, begins in 2400 and lasts forever. At each node of the tree, more information about the consumption damage function is revealed (as reflected in the parameter θ_t), but uncertainty is not fully resolved until the beginning of the next-to-last period in 2300. The agent's utility in each state is calculated based on interpolated consumption flows at five-year sub-periods, as discussed in Section 3.1. We solve for mitigation levels over time that maximize expected utility, looking forward, at the start of each period (except the final period), and in each fragility state. The resulting SCC in each period and state is the carbon price that implements this level of mitigation.

In the next section, we describe the agent's preferences, and provide some motivation for the preferences specification we employ. In Sections 2.2 and 2.3, we lay out the cost and damage functions and describe their calibration. Section 3 presents the results of a set of simulations designed to illustrate the effects that various parameters have on optimal climate policies. Section 4 concludes.

2.1 Preferences

As noted earlier, the CRRA specification typically used in climate studies embeds the assumption that agents' willingness to substitute consumption across states of nature is the same as their willingness to substitute consumption over time. However, this is inconsistent with the observed low risk-free rate and high equity premium (Mehra and Prescott, 1985; Weil, 1989). To resolve this puzzle, financial economists have begun to employ the preference specification suggested by Epstein and Zin that allows for different rates of substitution across time and states.[8] This is the specification we use here.

[8] See Bansal and Yaron (2004) and Hansen, Heaton, and Li (2008) for more detailed discussions. Bansal and Ochoa (2009) and Bansal and Ochoa (2011) use this preference specification in combination with a framework in which temperature shocks affect future consumption growth. Ackerman, Stanton, and Bueno (2013) and Crost and Traeger (2014) use this utility function in DICE.

In an Epstein-Zin utility framework, the agent maximizes at each time t:

(5) $$U_t = \left[(1-\beta)c_t^{\rho} + \beta\left[\mu_t\left(\widetilde{U}_{t+1}\right)\right]^{\rho}\right]^{1/\rho},$$

where $\mu_t\left(\widetilde{U}_{t+1}\right)$ is the certainty-equivalent of future lifetime utility, based on the agent's information at time t, and is given by:

(6) $$\mu_t\left(\widetilde{U}_{t+1}\right) = \left(E_t[U_{t+1}^{\alpha}]\right)^{1/\alpha}.$$

In this specification, $(1-\beta)/\beta$ is the pure rate of time preference. The parameter ρ measures the agent's willingness to substitute consumption across time. The higher is ρ, the more willing the agent is to substitute consumption across time. The elasticity of intertemporal substitution is given by $\sigma = 1/(1-\rho)$. Finally, α captures the agent's willingness to substitute consumption across (uncertain) future consumption streams. The higher is α, the more willing the agent is to substitute consumption across states of nature at a given point in time. The coefficient of relative risk aversion at a given point in time is $\gamma = (1-\alpha)$. This added flexibility allows for calibration across states of nature and time.

Note that with $\rho = \alpha$, equations (5) and (6) are equivalent to the standard CRRA utility specification. Plugging (6) into (5) for our model generates:

(7) $$U_0 = \left[(1-\beta)c_0^{\rho} + \beta\left(E_0\left[\widetilde{U}_1^{\alpha}\right]\right)^{\rho/\alpha}\right]^{1/\rho}$$

(8) $$U_t = \left[(1-\beta)c_t^{\rho} + \beta\left(E_t[\mathrm{U}_{t+1}^{\alpha}]\right)^{\rho/\alpha}\right]^{1/\rho}, \text{ for } t\epsilon\{1,2,\dots,T-1\}.$$

with c_0 and c_t, respectively, given by equations (1) and (2).

In the final period, which in our base case is the period starting in 2400, the agent receives the utility from all consumption from time T forward. Given our assumption that all uncertainty has already been resolved at this point, consumption grows at a constant rate g from T through infinity (i.e., $c_t = c_T(1+g)^{t-T}$ for $t \geq T$), and produces a utility to the agent of:

(9) $$U_T = \left[\frac{1-\beta}{1-\beta(1+g)^{\rho}}\right]^{1/\rho} c_T,$$

and with c_T given by equation (3)

2.2 Mitigation Cost Function Specification and Calibration

In this section we discuss the specification and the calibration of the mitigation cost function.

To calibrate the model, we need to find a relationship between τ, g, and x (where τ is the tax rate per ton of emissions, g is the resulting flow of emissions in gigatonnes of CO_2-equivalent emissions per year, Gt CO_2e, and x is the fraction of emissions reduced). To do so, we follow Pindyck (2012), which calibrates gamma distributions for temperature levels given greenhouse gas concentrations, and for economic damages given temperature levels.

McKinsey (2009) constructs a marginal abatement cost curve for GHGs that allows us to deduce τ, g, and x for the year 2030. We take McKinsey's estimates but assume no mitigation ($x(\tau) = 0$) at $\tau = 0$; i.e. no net-negative or zero-cost mitigation. Table 1 shows the resulting calibration.[9]

Table 1—Marginal abatement cost curve for 2030, from McKinsey (2009)

GHG taxation rate τ	GHG emissions flow $g(\tau)$	Fractional GHG reduction $x(\tau)$
€0/ton	70 Gt CO_2e/year	0
€60/ton	32 Gt CO_2e/year	0.543
€100/ton	23 Gt CO_2e/year	0.671

Fitting McKinsey's point estimates (in $US) from Table 1 to a power function for $x(\tau)$ yields:

(10) $$x(\tau) = 0.0923 \cdot \tau^{0.414}.$$

The corresponding inverse function, solving for the appropriate tax rate to achieve x is:

(11) $$\tau(x) = 314.32 \cdot x^{2.413}.$$

We are interested in $\kappa(\tau)$, the cost to the society when a GHG tax rate of τ is imposed. We can calculate $\kappa(\tau)$ using the envelope theorem. Intuitively, GHG emissions are an input to the production process that generates consumption goods. At any tax rate τ, assuming agents choose the level of GHG emissions $g(\tau)$ so as to maximize consumption given τ, then the marginal cost of increasing the tax rate must be the quantity of emissions at that tax rate, that is:

(12) $$\frac{dc(\tau)}{d\tau} = -g(\tau),$$

Thus, to calculate the consumption associated with a GHG tax rate of τ we integrate this expression, giving:

(13) $$c(\tau) = \bar{c} - \int_0^{\tau} g(s)\, ds,$$

[9] We have emissions stabilize at 57% above current levels. In our unmitigated baseline scenario, GHG concentrations reach approximately 1,000 ppm by 2200.

where $\bar{c}$ is the endowed level of consumption (assuming zero damages). However, this equation is correct only if the GHG tax is purely dissipative—that is, if the government were to collect the tax and then waste 100% of the proceeds. In our analysis, we instead assume that the tax is non-dissipative, meaning that the proceeds of the tax ($g(\tau) \cdot \tau$) would be refunded lump-sum, making the decrease in consumption just equal to the distortionary effect of the tax (in dollars) which is:[10]

(14) $$K(\tau) = \int_0^\tau g(s)\, ds - g(\tau) \cdot \tau.$$

Writing $g(\tau) = g_0\big(1 - x(\tau)\big)$, where g_0 is the baseline level of GHG emissions, we can rewrite $K(\tau)$ as:

$$\begin{aligned} K(\tau) &= g_0 \int_0^\tau \big(1 - x(s)\big) ds - \tau g_0 \big(1 - x(\tau)\big) \\ &= g_0 \left[\tau - \int_0^\tau x(s) ds\right] - \tau g_0 + \tau g_0 x(\tau) \\ &= g_0 \big[\tau x(\tau) - \textstyle\int_0^\tau x(s) ds\big] \end{aligned} \tag{15}$$

Substituting (10) into (15) and simplifying gives the total cost K as a function of the tax rate τ:

(16) $$K(\tau) = g_0[0.09230 \cdot \tau^{1.414} - 0.06526 \cdot \tau^{1.414}]$$
(17) $$= g_0 \cdot 0.02704 \cdot \tau^{1.414},$$

Substituting (11) into (17) gives K as a function of fractional-mitigation x:

(18) $$K(x) = g_0 92.08 \cdot x^{3.413},$$

where total cost $K(x)$ is expressed in dollars. Finally, we divide by current (2015) aggregate consumption to determine the cost as a fraction of baseline consumption:

(19) $$\kappa(x) = \left(\frac{g_0 \cdot 92.08}{c_0}\right) \cdot x^{3.413},$$

where $g_0 = 52$ Gt CO_2e/year represents the current level of global emissions, and $c_0 =$ \$31 trillion/year is current global consumption. This function expresses the total cost of a given level of mitigation as a percentage of consumption, and we hold that fixed in all periods except for the impact of technological change. We further assume that, absent technological change, the function $\kappa(x)$ is time invariant.

[10] Note that were the proceeds from the (Pigouvian) GHG tax used to reduce other distortionary taxes, the effective cost of the carbon tax would be still lower than what we calculate here, and thus would justify a higher optimal τ. For a summary of this "double-dividend" argument, see Goulder (1995).

2.2.1 Backstop Technology Specification

The McKinsey estimates on which our cost function $\kappa(x)$ are based reflect the cost of traditional mitigation only. However, in addition to standard mitigation, technologies are available for pulling CO_2 directly out of the atmosphere, such as carbon dioxide removal (CDR) or direct carbon removal (DCR) (National Research Council, 2015). We label these as *backstop technologies.*

We assume our backstop technology is available at a marginal cost of τ^*, for the first ton of carbon that is removed from the atmosphere. The marginal cost increases as extraction increases. We assume that unlimited amounts of CO_2 can be removed as the marginal cost approaches $\tilde{\tau} \geq \tau^*$. Under the most aggressive backstop scenario presented in the results section, we assume a price of \$350 per ton today for τ^* and a price of \$400 per ton for $\tilde{\tau}$. Given our underlying cost curve for emissions mitigation, these values imply that the backstop technology kicks in at mitigation levels above 104%.

In fitting the marginal cost curve to these lower and upper bounds for the backstop technology we build a marginal cost function for the backstop technology of the form:

(20) $$B(x) = \tilde{\tau} - \left({}^{k}/_{x}\right)^{1/b}.$$

The upper bound of the cost function is, thus, $\tilde{\tau}$. Moreover, we calibrate (18) for the backstop technology to be used once the mitigation level, x_0, is such that:

(21) $$B(x_0) = \tilde{\tau} - \left({}^{k}/_{x_0}\right)^{1/b} = \tau^*,$$

which allows us to express:

(22) $$k = x_0(\tilde{\tau} - \tau^*)^b.$$

Second, we impose a smooth-pasting condition; i.e. the derivative of the marginal cost curve is continuous at x_0. This allows us to solve for parameter b:

(23) $$b = \frac{\tilde{\tau} - \tau^*}{(\alpha - 1)\tau^*}.$$

Figure 3 and Figure 4 show, respectively, the total cost of mitigation as a fraction of consumption, $\kappa(x)$, and the marginal cost, $\tau(x)$, assuming a backstop technology at \$400 per ton.

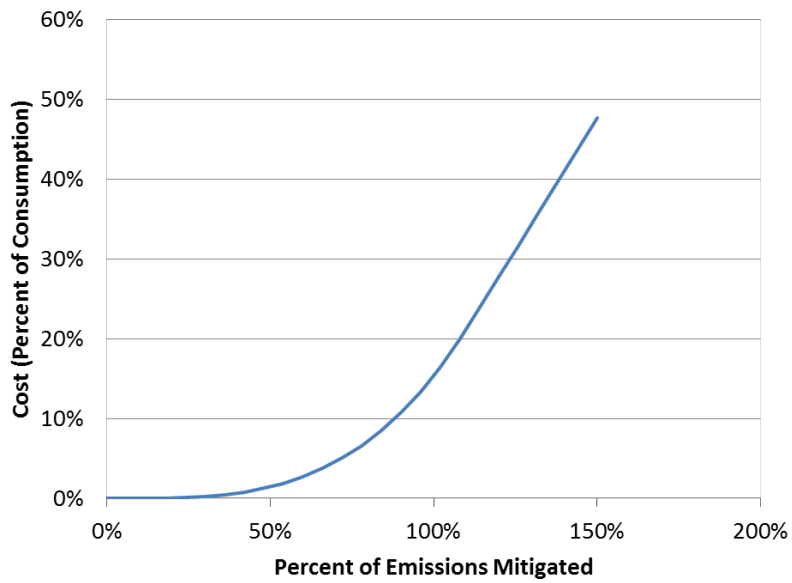

Figure 3—Total cost of abatement as a fraction of consumption, $\kappa(x)$

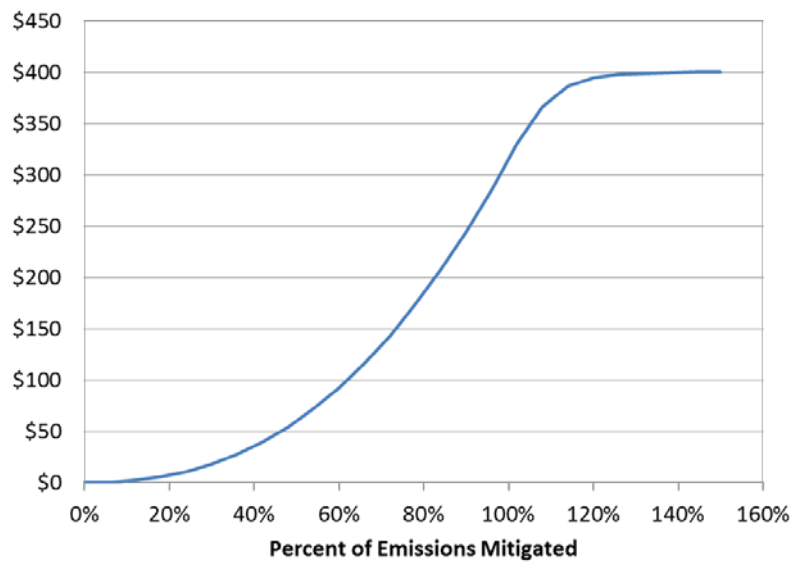

Figure 4—Marginal cost of abatement (in \$/ton), $\tau(x)$

2.2.2 Technological Change Specification

These cost curves are calibrated to $t = 0$. In subsequent periods, we allow the marginal cost curve to decrease at a rate determined by a set of technological change parameters: a constant component, φ_0, and a component linked to mitigation efforts to date, $\varphi_1 X_t$, where X_t is the average mitigation up to time t (equation (4)). Thus, at time t, the total cost curve is given by:

(24) $$\kappa(x,t) = \kappa(x)[1 - \varphi_0 - \varphi_1 X_t]^t.$$

This functional form allows for easy calibration. For example, if $\varphi_0 = 0.005$ and $\varphi_1 = 0.01$, then with average mitigation of 50%, marginal costs decrease as a percentage of consumption at a rate of 1% per year.

2.3 Damage Function Specification

We next specify the climate damage function $D_t(X_t, \theta_t)$. Damages are a function of temperature change, which is in turn a function of greenhouse gas concentrations which, in our setting, is defined by average level of mitigation up to that point in time, X_t. The only way to affect the level of damages, then, is to change mitigation X_t, including both decreased carbon emissions and direct carbon dioxide removal from the air.

The specification of damages has two components: a non-catastrophic component and an additional catastrophic component triggered by crossing a particular threshold. The hazard rate associated with hitting that threshold increases with temperature. If the threshold is crossed at any time, additional damages decrease consumption in all future periods.

The overall damage function $D_t(X_t, \theta_t)$ is calculated via Monte-Carlo simulation. As we describe in detail below, we run a set of simulations for each of three mitigation levels X_t. In each run of the simulation, we draw a set of random variables: (1) the temperature change; (2) the parameter characterizing damages as a function of temperature, and (3) for each period on each path an indicator variable which determines whether or not the atmosphere hits a tipping point at that time, and (4) the tipping point damage parameter. The state variable θ_t indexes the distribution resulting from these sets of simulations, and interpolation across these three mitigation levels gives us a continuous function of X_t.

2.3.1 The Specification of Temperature as a Function of GHG Levels

The distribution of temperature outcomes as a function of mitigation strategies is calibrated to three carbon scenarios, indexed by a maximum level of CO_2 in the atmosphere. For the original calibration, we follow Weitzman (2009) and Wagner and Weitzman (2015) in calibrating a log-normal distribution for equilibrium climate sensitivity—the eventual temperature rise as atmospheric concentrations of CO_2 double. The calibration uses a conservative interpretation of the IPCC's "likely" range, as well as statements around extreme outcomes.

Specifically, Wagner and Weitzman (2015) calibrate a log-normal function assuming a 78% probability of climate sensitivity being in the 1.5-4.5°C range. (The IPCC says that range is "likely," which it defines as having at least a 66% probability. The IPCC's "very likely" designation implies at least a 90% probability. Wagner and Weitzman (2015) split the difference to arrive at 78%.) Moreover, the Intergovernmental Panel on Climate Change (IPCC)'s Fifth Assessment Report (IPCC, 2013) judges climate sensitivity above

6°C to be "very unlikely," giving it a 0-10% probability. Wagner and Weitzman's (2015) calibration assigns it a roughly 5% chance.

Wagner and Weitzman (2015) then use this calibration to translate the International Energy Agency's projections for concentrations of CO_2-equivalent tons into final temperature outcomes. Under the assumptions of their "new policies scenario," International Energy Agency (2013) projects that atmospheric concentrations will reach 700 ppm CO_2e by 2100. That concentration would result in a projected, eventual median temperature increase of 3.6°C. Wagner and Weitzman (2015) present eventual median temperature outcomes for concentrations of between 400 and 800 ppm. We take their calibration and extrapolate to 1000 ppm, which we assume to be the zero-mitigation scenario, marking an upper bound of sorts. We similarly assume that 100% mitigation over time leads to a maximum GHG level of 400 ppm. Other levels of average mitigation are assumed to lead to damages associated with GHG levels linearly interpolated between those levels. Thus, an average mitigation of 50% through any point in time leads to the interpolated damages associated with a maximum GHG level of 700 ppm at that time. We then use assumptions akin to Pindyck (2012) to fit a displaced gamma distribution around final GHG concentrations, while setting levels of GHG 100 years in the future equal to equilibrium levels.

Table 2 gives the probability of different levels of ΔT_{100} – the temperature change over the next 100 years – for given maximum levels of GHGs in atmosphere. The 450 ppm, 650 ppm, and 1000 ppm maximum levels of CO_2 equivalents in the atmosphere reflect, respectively, a strict, a modest, and an ineffective mitigation scenario.

Table 2—Probability of $\Delta T_{100} > T$

	Maximum GHG Level (ppm of CO_2)		
T	450	650	1000
2°C	**0.400**	**0.850**	**0.990**
3°C	**0.125**	**0.540**	**0.860**
4°C	**0.040**	**0.300**	**0.655**
5°C	**0.015**	**0.145**	**0.455**
6°C	**0.002**	**0.072**	**0.303**

We then fit a displaced gamma distribution to each of these sets of probabilities. Table 3 gives the parameters for these distributions, and the probabilities from the fitted displaced gamma distributions, which line up well with the numbers in Table 2.

Table 3—Fitted values of $Prob(\Delta T_{100} > T)$ for three specified gamma distributions

	Maximum GHG Level (ppm of CO2)		
T	450	650	1000
2°C	**0.396**	**0.870**	**0.994**
3°C	**0.139**	**0.566**	**0.910**
4°C	**0.042**	**0.289**	**0.696**
5°C	**0.011**	**0.124**	**0.443**
6°C	**0.003**	**0.047**	**0.242**
Gamma distribution parameters			
Alpha	2.810	4.630	6.100
Beta	0.600	0.630	0.670
Displace	-0.25	-0.5	-0.9

To obtain the temperature distribution at other times, we follow Pindyck (2012), and specify that the time path for the temperature change at time t (in years) is given by:

$$\Delta T(t) = 2\,\Delta T_{100}\left[1 - 0.5^{\frac{t}{100}}\right]. \tag{25}$$

The temperature paths are plotted for different levels of ΔT_{100}. As time increases, the temperature change asymptotes to double the value of ΔT_{100}. We would like to emphasize that both the distribution of ΔT_{100} and the functional form for the path in equation (25) merit further scientific scrutiny.

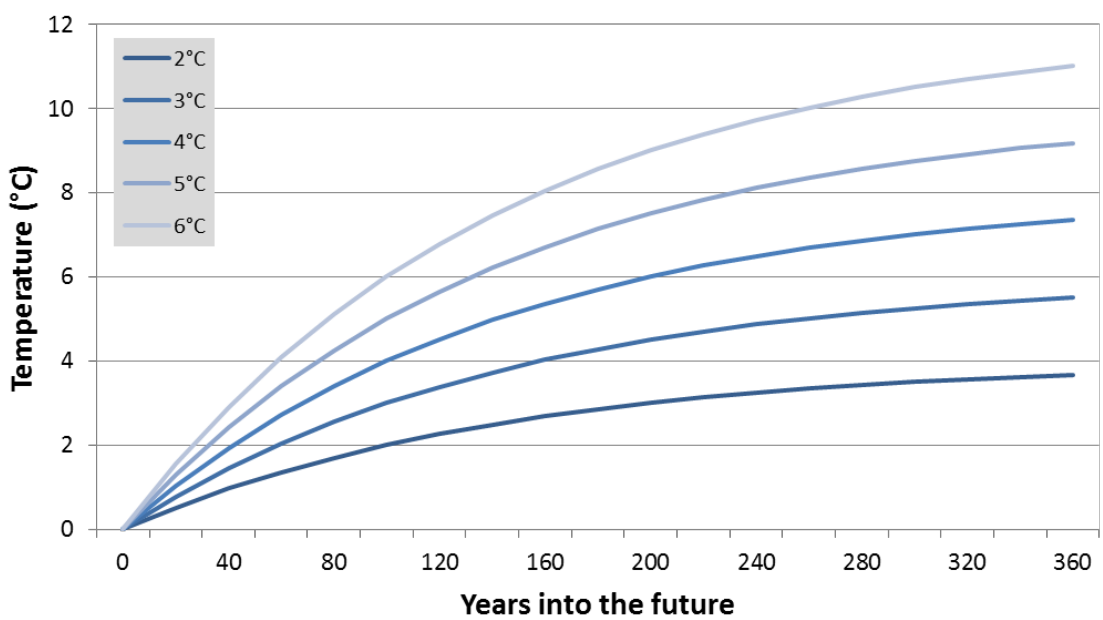

Figure 5—Calibrated time path for temperature increases given assumed temperature increases by 2100

2.3.2 The Specification of Damages as a Function of Temperature

Our next step is to translate average global surface warming into global mean economic losses via our damage function. There are two components to our damage function: a *non-catastrophic* and *catastrophic* component. The functional form of each component is known to the agent. However, as with the GHG-ΔT_{100} relationship discussed in the previous section, the functional form for each damage function component contains a parameter that characterizes the high uncertainty in our present understanding of this relationship. In our model, the agent knows the form of the distribution of this parameter at the initial date, and in each period learns more about the distribution of the parameter. However, the final realization of the parameter is not known until the next-to-last period.

The non-catastrophic component of our damages is based on Pindyck (2012), who fits a functional form to data from the IPCC's Fourth Assessment Report (IPCC, 2007), and obtains a loss function of the form:

(26) $$\mathrm{L}(\Delta T(t)) = e^{-13.97 \cdot \gamma \cdot \Delta T(t)^2},$$

where γ is drawn from a displaced gamma distribution with parameters r = 4.5, λ= 21341, and θ = −0.0000746.

Based on non-catastrophic damages, consumption in any time t is reduced as follows:

(27) $$\mathrm{CD}_t = \bar{c}_t \cdot \mathrm{L}(\Delta T(t)).$$

A major concern with the damage function above is that it effectively rules out catastrophic risks, even at high temperature changes. Take an 8°C temperature change, well outside the range typically assumed to be 'safe'. If per capita consumption is assumed to grow in real terms by 2% annually, then such damage applied to consumption 50 years hence would reduce the average consumption from 2.7 times today's value to 2.2 times, a significant reduction, but hardly a catastrophe of significant concern today. Even the 1% point in the outcome distribution conditional on an 8°C average temperature change is assumed here to be a reduction in consumption of only 32% which implies people are still 1.8 times wealthier than today. We hence augment Pindyck's (2012) damage function with the possibility of catastrophic events after reaching a particular temperature threshold, which itself creates the potential for a much larger impact on consumption.

While the possibility of climate tipping elements is receiving considerable attention in the scientific community, there is no single right specification (Kopp et al., 2016). We employ an *ad-hoc* specification as part of our broader calibration effort to approximate the US government estimate of $40 for the SCC (U.S. Government Interagency Working Group on Social Cost of Carbon, 2015). Our results, hence, ought to be considered what they are: sensitivity analyses probing which factors contribute the most to the structure of the SCC. In our specification, Prob(TP) denotes the probability of hitting a 'Tipping

Point' over a given interval of length "*period*" as a function of the global temperature change as of that time ($\Delta T(t)$), and of a parameter, $peakT$:

$$\text{(28)} \qquad \text{Prob(TP)} = \left\{1 - \left(1 - \left[\frac{\Delta T(t)}{\max[\Delta T(t), peakT]}\right]^2\right)^{(period/30)}\right\}$$

Figure 6 plots Prob(TP) as a function of $\Delta T(t)$ for a 30-year *period* and a set of values of $peakT$. As $peakT$ increases, the probability of reaching a climatic tipping point decreases for a given $\Delta T(t)$. Our subsequent base case specification, calibrated to an SCC of around \$40 in 2015 employs $peakT = 11$.

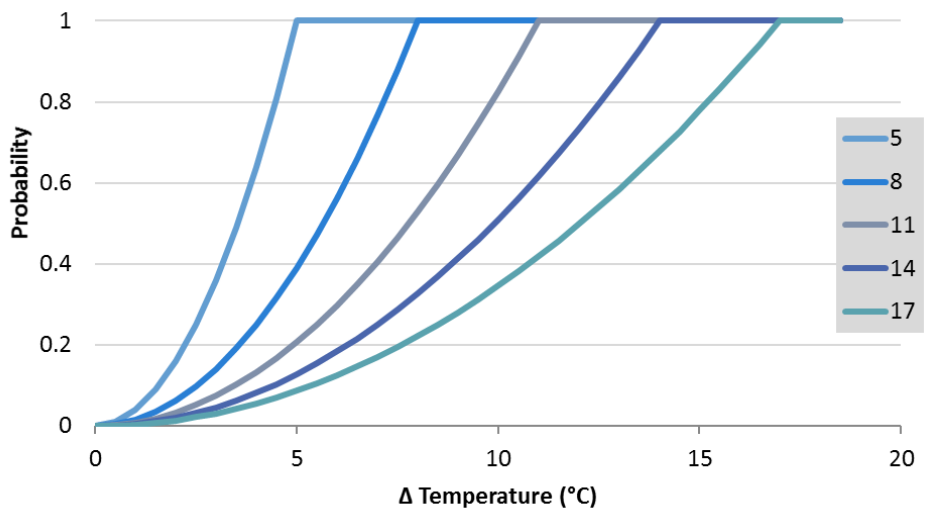

Figure 6—Probability of reaching a climatic tipping point as a function of *peakT*

In our simulations, in each period p and for each state, there is a probability Prob(TP) that a tipping point will be hit (given $\Delta T(t)$ and $peakT$). Conditional on hitting a tipping point at time t^* in a given run of the simulation, the level of consumption for each period $t \geq t^*$ is then at a level of:

$$\text{(29)} \qquad \text{CDTP}_t = \text{CD}_t \cdot \text{e}^{-\text{TP\_damage}} = \bar{c}_t \cdot \text{L}(\Delta T(t)) \cdot \text{e}^{-\text{TP\_damage}} \text{ for } t \geq t^*,$$

where TP_damage is a random variable drawn from a gamma distribution with parameters $\alpha = 1$ and $\beta = disaster\_tail$. The cumulative distribution for tipping point damage (i.e., $\left(1 - \text{e}^{-\text{TP\_damage}}\right)$) for values of $disaster\_tail$ ranging from 2 to 10 is plotted below:

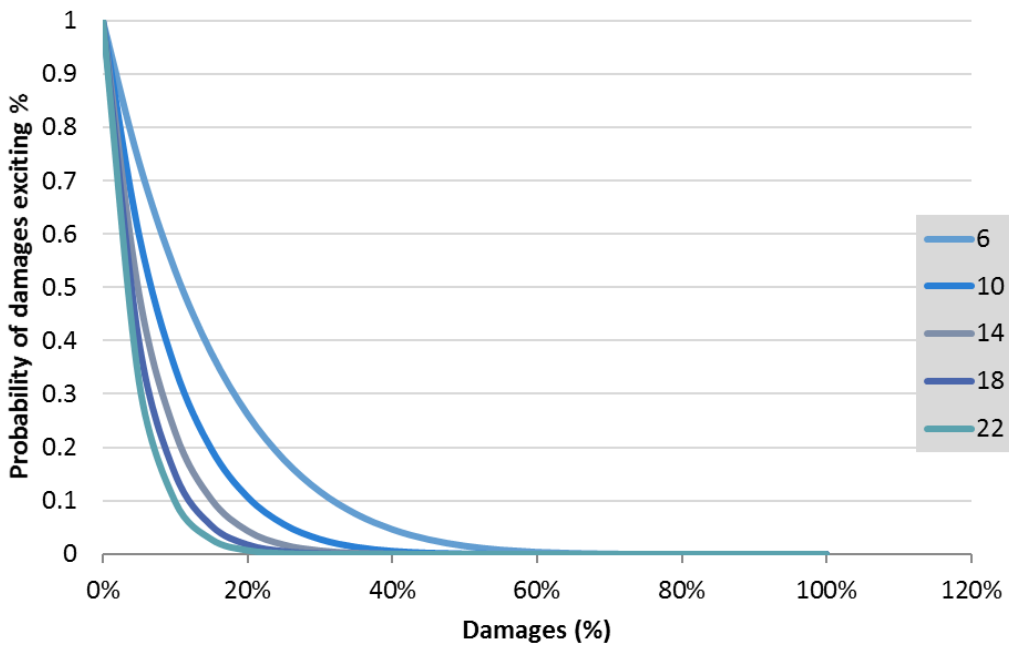

Figure 7—Probability of damage greater than a particular percentage of output, given different *disaster_tail* assumptions

2.3.3 Interpolation and Incorporation of Uncertainty in the Damage Function

Comparing equation (29) with equation (2) shows that the damage function for a given level of mitigation and in a given state of nature is:

$$\mathrm{D}_t = (1 - \mathrm{L}(\Delta T(t)) \cdot \left(1 - \mathrm{I}_{TP}\left[\, 1 - \mathrm{e}^{-\mathrm{TP\_damage}}\right]\right), \tag{30}$$

where I_{TP} is an indicator variable which is equal to one if a tipping point has been hit, and zero otherwise. However, recall that $\mathrm{L}(\Delta T(t))$, I_{TP}, and $\mathrm{e}^{-\mathrm{TP\_damage}}$ are each dependent on the specific realization of the draws of random numbers in our simulations.

Therefore, for each of three values for the maximum GHG—450, 650, and 1000 ppm—we run a set of simulations to generate a distribution of D_t for each period. We order the simulations based on D_T, the damage to consumption in the final period. We choose states of nature with specified probabilities to represent different percentiles of this distribution. For example, if the first state of nature is the worst 1% of outcomes, then we assume the damage coefficient at time *t* for the given level of mitigation is the average damage at time *t* for the worst 1% of values for D_t.

More generally, if the kth state of nature represents the simulation outcomes in the range [prob(k-1) , prob(k)], then the damage coefficient for the kth state of nature is the

average damage in that range of simulations in which the distribution for D_t lies within those percentiles.

The next step is to calculate damages in any particular period for any particular state of nature and any chosen mitigation action. We do this by interpolating smoothly with respect to the average percentage mitigated up to each point in time. Zero mitigation corresponds to the 1000 ppm maximum GHG scenario, whereas 100% average mitigation is assumed to correspond to a 400 ppm maximum GHG scenario. Since there is a potential total of 600 ppm additional GHG in the atmosphere to be mitigated, the 450 ppm maximum GHG scenario corresponds to a 91.7% mitigation (=550/600) and the 650 ppm maximum GHG scenario corresponds to a 58.3% mitigation (=350/600).

Our task is to calculate an interpolated damage function between the three scenarios where we have damage coefficients (for a given state and period) to find a smooth function that gives damages for any particular average mitigation percentage up to each point in time.

We first calculate a quadratic section of the damage function which starts (for a given state and period) at the level of damages in the 1000 ppm maximum GHG scenario and is assumed to have a zero derivative at that point. The curvature as a function of mitigation is calculated such that the damage function matches the damage coefficient at the 650 ppm maximum GHG scenario. For emissions mitigation percentages less than 58.3% we use this quadratic curve to interpolate damages.

We next calculate a quadratic section of the damage function which starts at the level of damages in the 650 ppm maximum GHG scenario and is assumed to have a derivative equal to that of the first quadratic where they meet at the 58.3% emissions mitigation point. The curvature of the second quadratic is then calculated such that the damage function matches the damage coefficient at the 450 ppm maximum GHG scenario. We use this quadratic curve to interpolate damages when emissions mitigation is greater than 58.3% and less than 100%.

We allow for the possibility of net GHG removal from the atmosphere, in which case emissions mitigation can become greater that 100%. In that case we extend the second quadratic interpolation but decay it toward zero by dividing by $2^{10(\%mitigation-1)}$. Thus at 110% mitigation we divide by 2; at 120% mitigation we divide by 4; etc. The purpose of this decay is to cause the quadratic curve to smoothly decay toward zero damages.

As an example, consider the 10% worst case in period 5, which is calculated here for our base case, using a *peakT* of 11 and *disaster_tail* of 18 to result in the following interpolated damage functions shown in Figure 8:

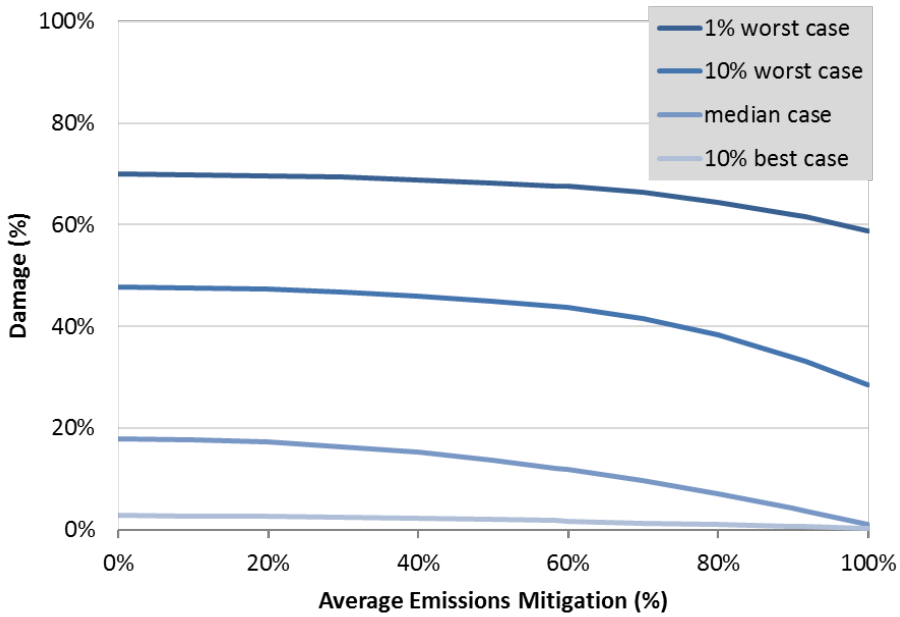

Figure 8—Interpolated final period damage functions.

The climate sensitivity—summarized by state of nature θ_T—is not known prior to the final period ($t = T$). Rather, what the agent knows is the distribution of possible final states. We specify that the damage in period t, given average mitigation of X_t up to time t, is the probability weighted average of the interpolated damage function over all final states of nature reachable from that node. Specifically, the damage function at time t, for the node indexed by θ_t is assumed to be:

(31) $$\mathrm{D}_t\left(X_{t,}\theta_t\right) = \sum_{\theta_T} \Pr(\theta_T|\theta_t) \cdot \mathrm{D}_t\left(X_{t,}\theta_T\right),$$

where the sum is taken over all states that are possible from the node indexed by θ_t (i.e., for which $\Pr(\theta_T|\theta_t) > 0$).

2.4 Risk decomposition

Figure 2 above decomposes the SCC into a risk aversion and an expected damages component. We present here the mathematical derivation of these results: Let $D_{s,t}$ denote the marginal damage, that is the loss of consumption in state s in future period t that results from putting one more ton of carbon into the atmosphere today (at time 0). The SCC is then:

(32) $$SCC = \sum_{t=1}^{T} \sum_{s=1}^{S(t)} \pi_{s,t} m_{s,t} D_{s,t} \left(= \sum_{t=1}^{T} E_0\left[\widetilde{m}_t \widetilde{D}_t\right]\right).$$

where $m_{s,t}$ is the pricing kernel in state s at time t, which is the marginal value today of one additional unit of consumption in state s at time t,[11] $\pi_{s,t}$ denotes the probability of state s at time t, and $S(t)$ denotes the number of states at time t. That is, to calculate the cost to the representative agent of an additional ton of carbon emissions, we sum over all the consumption damages that result from this, in every state of nature at every future time, multiplied by the value of an additional unit of consumption in that state at that time.

Equation (32) can be decomposed as:

(33) $$SCC = \sum_{t=1}^{T} E_0[\widetilde{m}_t] \cdot E_0\left[\widetilde{D}_t\right] + \sum_{t=1}^{T} cov_0\left(\widetilde{m}_t, \widetilde{D}_t\right).$$

Note that:

(34) $$E_o[\widetilde{m}_t] = \frac{1}{R^f(0,t)}$$

where $R^f(0,t)$ is the time t payoff to an investment at time 0 of \$1 in a risk-free bond that matures at time t.[12] This implies we can rewrite the first component of (33) as the sum of the marginal damages, discounted back to the present at the risk-free rate. The second component is the premium over the expected damages that society is willing to pay because of the 'risk' of the damages, defined as the covariance of the marginal damages with marginal utility.

We can, thus, label the first component of (33) as discounted expected damages $ED = \sum_{t=1}^{T} \frac{E_0[\widetilde{D}_t]}{R^f(0,t)}$ and the second component as the risk premium $RP = \sum_{t=1}^{T} cov_0\left(\widetilde{m}_t, \widetilde{D}_t\right)$. Rewriting (33) then gives the risk premium as the difference between the social cost of carbon and the expected-damages, both of which are easily calculated in our model:

(35) $$RP = SCC - ED.$$

3. Results

The main model output is the price of carbon, both today, and at the beginning of each of the next five periods. These are the times in the model when mitigation decisions are made. Figure 9 shows the results for the CRRA model run (top left) and two additions: a risk aversion coefficient of 7 (lower panel), calibrated to observed financial asset prices; and an inclusion of extreme events, what we call the 'disaster' scenario (right panel). It also shows the implications of using three different climate sensitivity distributions: one

[11] Equivalently, $m_{s,t}$ is defined as the ratio of marginal utility with respect to current consumption in that state to the marginal utility today, that is $m_{s,t} = \left(\frac{\partial U}{\partial c_{s,t}}\right) / \left(\frac{\partial U}{\partial c_0}\right)$, where $c_{s,t}$ denotes the agent's consumption in state s at time t. We present the pricing kernel for each time and state, for our base case, in Figure 14.
[12] Alternatively, $E_0[\widetilde{m}_t]$ is the risk-free discount factor between today and time t.

following Pindyck (2012), the second Roe and Baker (2007), and the third the log-normal calibration employed by Wagner and Weitzman (2015).

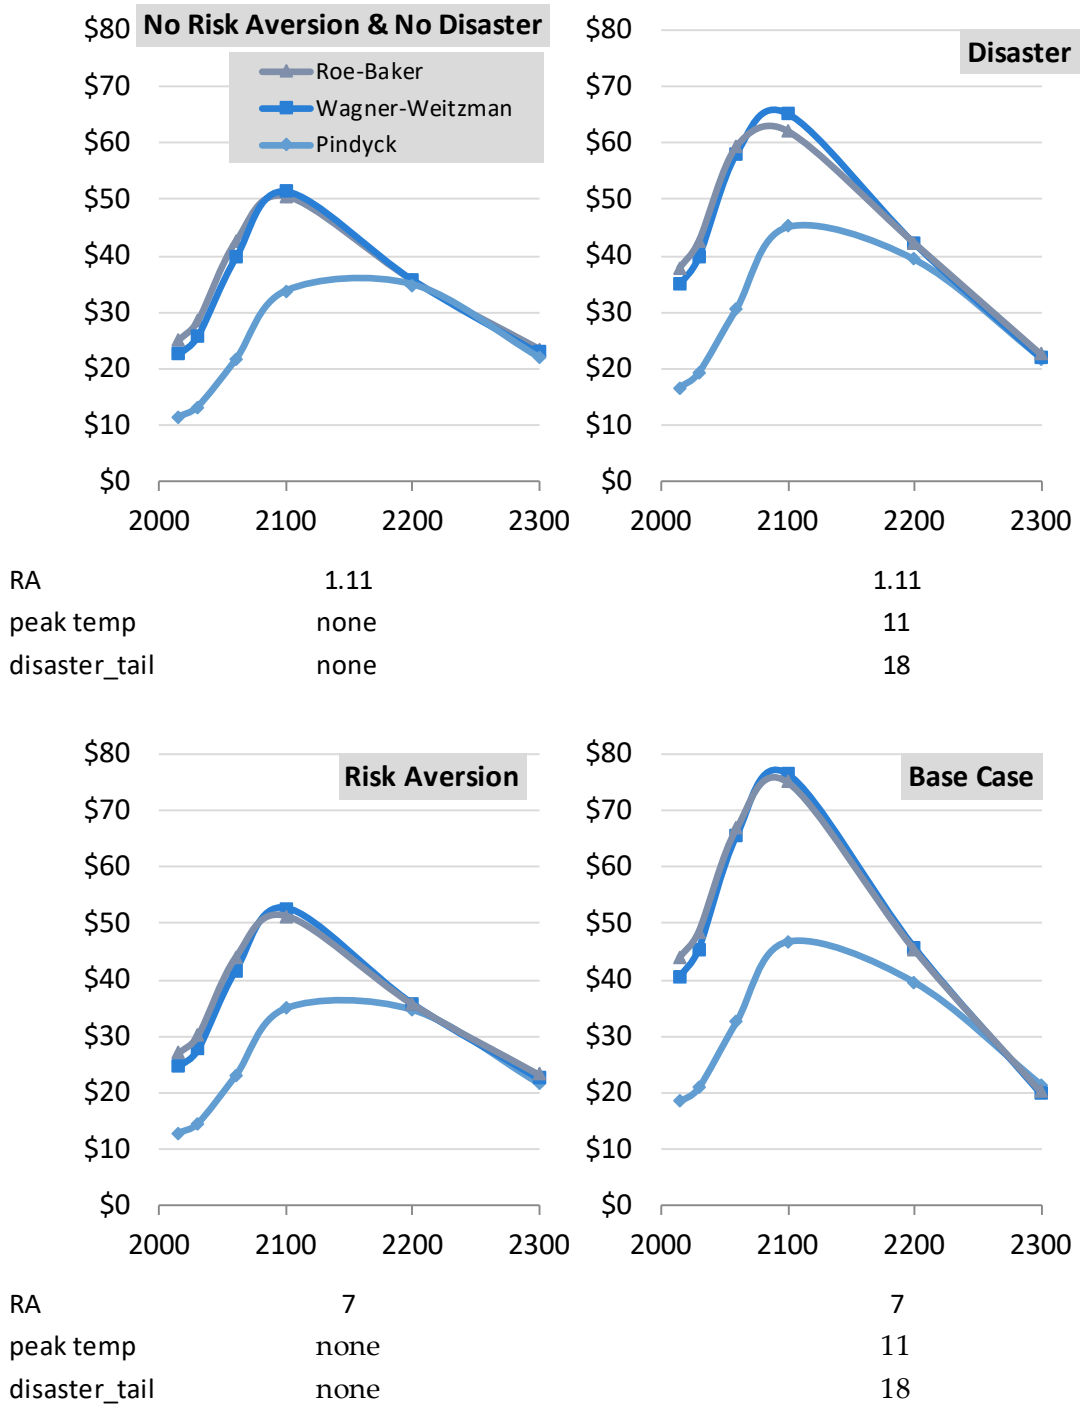

Figure 9—Expected price per ton of carbon under four different scenarios and three different assumed climate sensitivity distributions.

Risk aversion alone increases prices slightly in the early periods, though barely noticeably. Disasters alone increase prices dramatically, an effect that is further

magnified by the inclusion of risk aversion.[13] Using a Pindyck (2012) climate sensitivity calibration likely leads to conservative estimates. By instead using a heavy-tailed probability distribution function such as Roe-Baker or Wagner-Weitzman, implied prices increase dramatically. We proceed to use the Wagner-Weitzman (2015) climate sensitivity distribution, combined with Pindyck's (2012) loss function, equation (26), as our base case for the remainder of our runs.

3.1 Tree structure

Figure 10 illustrates the tree structure employed in our baseline analysis. At the start of the model (i.e., in 2015), the agent is assumed to know the information filtration (i.e., the structure of the tree) meaning he knows the state probabilities and the damage function in each future state of the world. Period zero runs from 2015 until 2030. In 2030, the agent learns whether the world is in state 'u' or state 'd'. There is a 50% probability of each of the two states. Similarly, at the end of period two (in year 2060) the agent learns whether the world is in state 'uu', 'ud', or 'dd', etc. Notice that at the end of period four, all uncertainty is resolved, in that the agent will learn which of the six final states the world is in and what the true damage function is. Following this point, in period five, the agent has one final period in which he can do mitigation. However, in period six, which in our base case runs from 2400 on to infinity, the agent can no longer mitigate. Consumption continues to grow deterministically from this point forward at a rate g, meaning that consumption after $T = 2400$ is given by $c_t = c_T(1+g)^{t-T}$, and the period six utility is given by equation (9).

In the baseline model, where a move up or down in each period is equally likely, the probabilities of the final states are given by a binomial distribution.

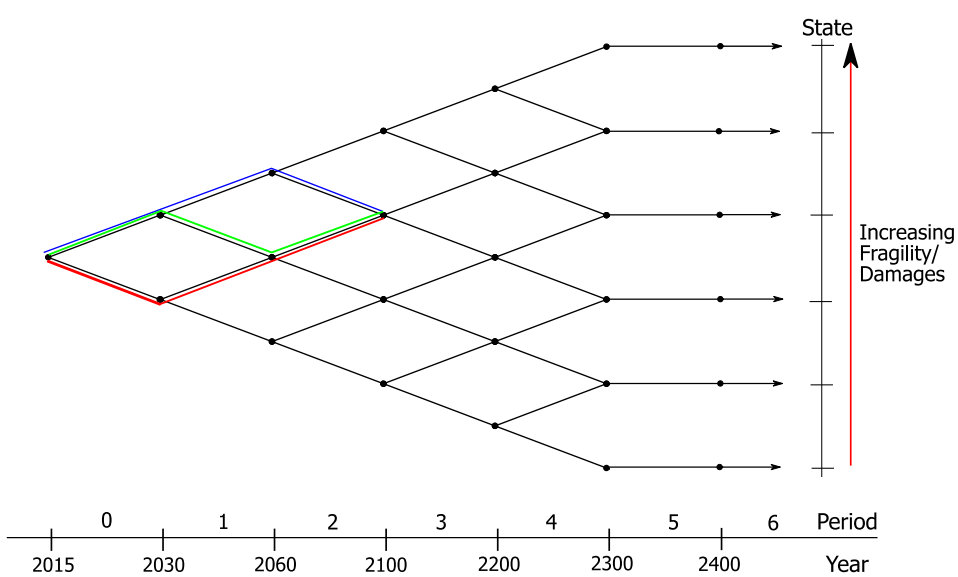

Figure 10—Diagram of tree structure used in solving the model for each state of nature across time.

[13] Note that the 2015 price comes from a single node in the tree. In each subsequent year, that price is set in expectation over all possible states of nature in that given year.

Another feature of the tree is particularly important given our use of the Epstein-Zin preference specification. We employ a recombining tree structure, meaning that the damage function in state '$uuudd$', for example, is independent of the way in which information was revealed at the end of each period. However, the agent's utility is path-dependent, as the history of mitigation depends on the process by which the agent learns the state. Thus, consumption, and mitigation, will depend upon the path. Consequently, in solving for the agent's utility along each of these paths, we need to keep track of the path by which the agent learned about the damage function.

While the agent learns more about the nature of the damage function in discrete "chunks," we calculate the agent's consumption over five-year sub-periods using an interpolation method. Specifically, as noted earlier, in our baseline model endowed consumption (before climate damages and before mitigation costs) is assumed to grow at a rate of 2% per year. Following equation (1), the consumption flow is c_0, the endowed consumption, less the cost of mitigation. Also, from equation (2), the consumption flow at the start of period 1 is given by $c_1 = \bar{c}_0 \cdot e^{0.02\times 15}\big(1 - D_1(X_1, \theta_1) - \kappa_1(x_1)\big)$. That is, the consumption at the start of period one (in 2030), c_1, is equal to endowed consumption (endowed 2015 consumption plus 2% growth for 15 years), minus the fractional cost of damages and of mitigation chosen at the beginning of period one. Mitigation is optimally chosen by the agent, and is therefore a function of the state—mitigation will be lower if the agent learns that the world is in state 'd' rather than state 'u'.

Thus, this analysis gives us the consumption flow in 2015 and in 2030, for the two states. What we do to calculate the agent's consumption between 2015 and 2030 is to fit an exponential growth function to the consumption at these two points of time. Note that this is equivalent to assuming that immediately after choosing the period zero mitigation, the agent's consumption starts to reflect climate damages from the first revealed state ('u' or 'd'). However, the agent is not allowed to change the period zero mitigation to reflect this knowledge.

The purpose of introducing this interpolation scheme is to ensure that the agents consumption path is relatively smooth. This will clearly lead to approximation errors in our solution. Adding more periods (as opposed to sub-periods) at which the agent can choose a new level of mitigation would result in a smaller approximation error, but would result in far higher computational costs: with T periods, we have a "$2^{T+1} - 1$"-dimensional optimization problem, in that we must choose this number of optimal mitigation levels (as a function of the path of states) to calculate the SCC.

In our base case each state has equal probability. The current optimal price of emissions is around \$40/ton for the representative agent. At the start of period 1, in 2030, the agent learns that he has moved either up or down in the tree. The fragility of the environment is a function of the number of up moves as the agent traverses through the tree structure. Thus, moving up in the tree leads to states which on average have greater fragility and worse outcomes. It makes sense for the agent to spend more on mitigation as he moves further up in the tree structure. The roughly \$45 expected price for

emissions in period 1 which we show in the lower right quadrant of Figure 9 is the average of the price in the up state, $58, and the price in the down state, $33.

Figure 11 shows probabilities and emissions prices in each node, through the start of period five. The lines connecting the boxes indicate the paths that information about the earth's fragility has taken. All grouped nodes at a given time have the same degree of fragility and thus the same damage for a given amount of greenhouse gas in the atmosphere. Consider period 2, starting in 2060, which has four equally probable nodes. The fragility in the up-down node is the same as in the down-up node, but the prices are slightly different: $60.47 for up-down, and 60.68 for down-up.

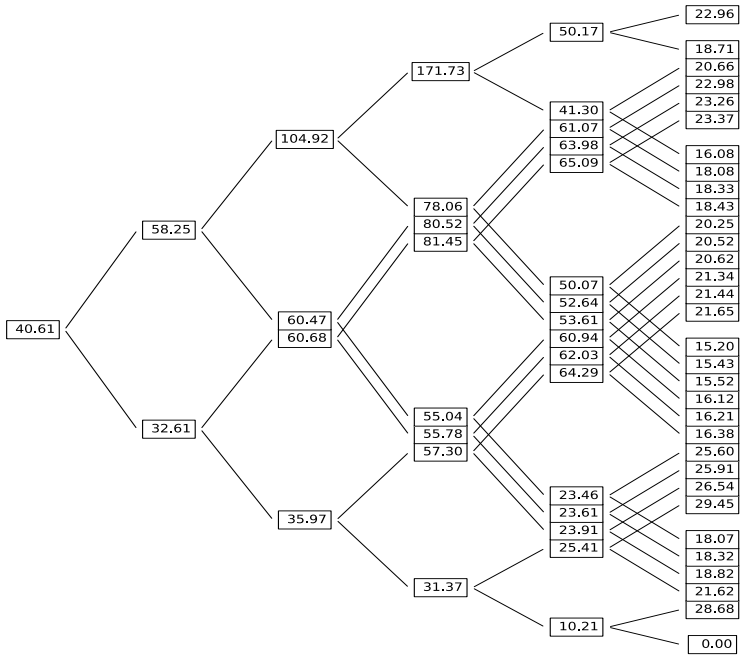

Figure 11—Emissions prices (in $ per ton of CO_2) across time and states

Figure 12 shows the fractional mitigation for each state, and reveals the reason for the different prices. At the start of period one, when there is bad news about fragility, the representative agent chooses to mitigate more (0.5463) than when he receives good news (0.4296). To induce (atomistic) agents to mitigate more, the optimal emissions price in the period 1 up state is $58.25/ton, relative to only $32.61/ton in the down state. Then, at the start of period 2 "up-down" state the representative agent mitigates slightly less (0.6695) than in the "down-up" state (0.6705), not because the damage function is different, but rather because, on the "up-down" path, the agent does so much more mitigation in period 1. Thus, there are path dependencies in the tree and we keep track of each path-dependent node separately.

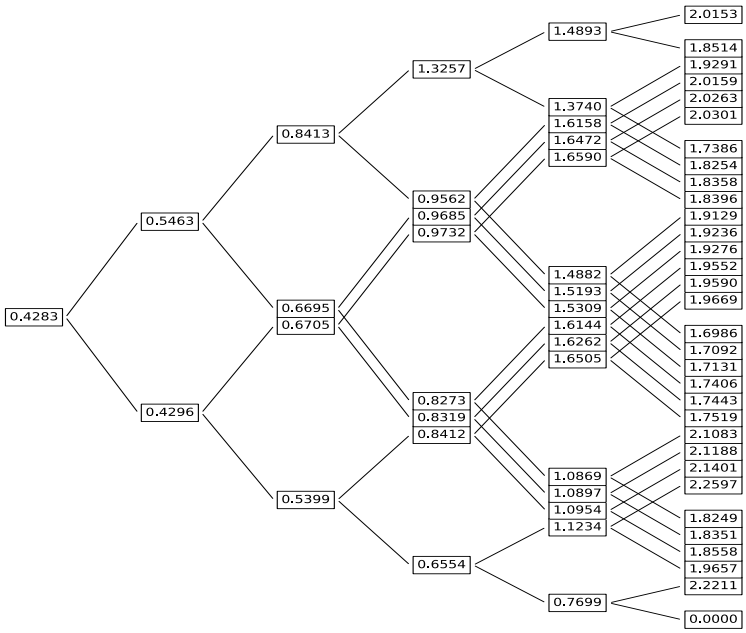

Figure 12—Fractional mitigation across time and states

In our baseline setting, mitigation becomes cheaper over time as technological change occurs, and as a result the optimal level of mitigation increases over time. For example, 42.83% of emissions are mitigated in period 0. If the agent gets two good draws (ie., two down moves in the tree) then in period 2, the period from 2060 to 2100, the agent mitigates a larger fraction, 53.99% of emissions. On the other hand, if the agent gets two bad draws (up moves) then in period 2 he mitigates 84.13% of emissions. In the base case scenario, the cost of the backstop technology has become so low by the fourth period that it is employed in the highest fragility state in the fourth period. This is reflected in the fractional mitigations which are greater than 1.0. By the fifth period, the cost is sufficiently low that it is optimal to extract carbon from the atmosphere in all but the lowest fragility state.

Figure 13 shows the consumption (as a multiple of consumption today) and the level of greenhouse gases in ppm in the atmosphere at each point in the tree. As would be expected, moving up in the tree (to higher levels of fragility) leads to decreased consumption as a result of both higher damages and higher mitigation. Note also that, if the earth turns out to be more robust than anticipates (i.e., at the bottom right of the tree) then the optimal policy leads to a higher overall level of GHGs in the atmosphere.

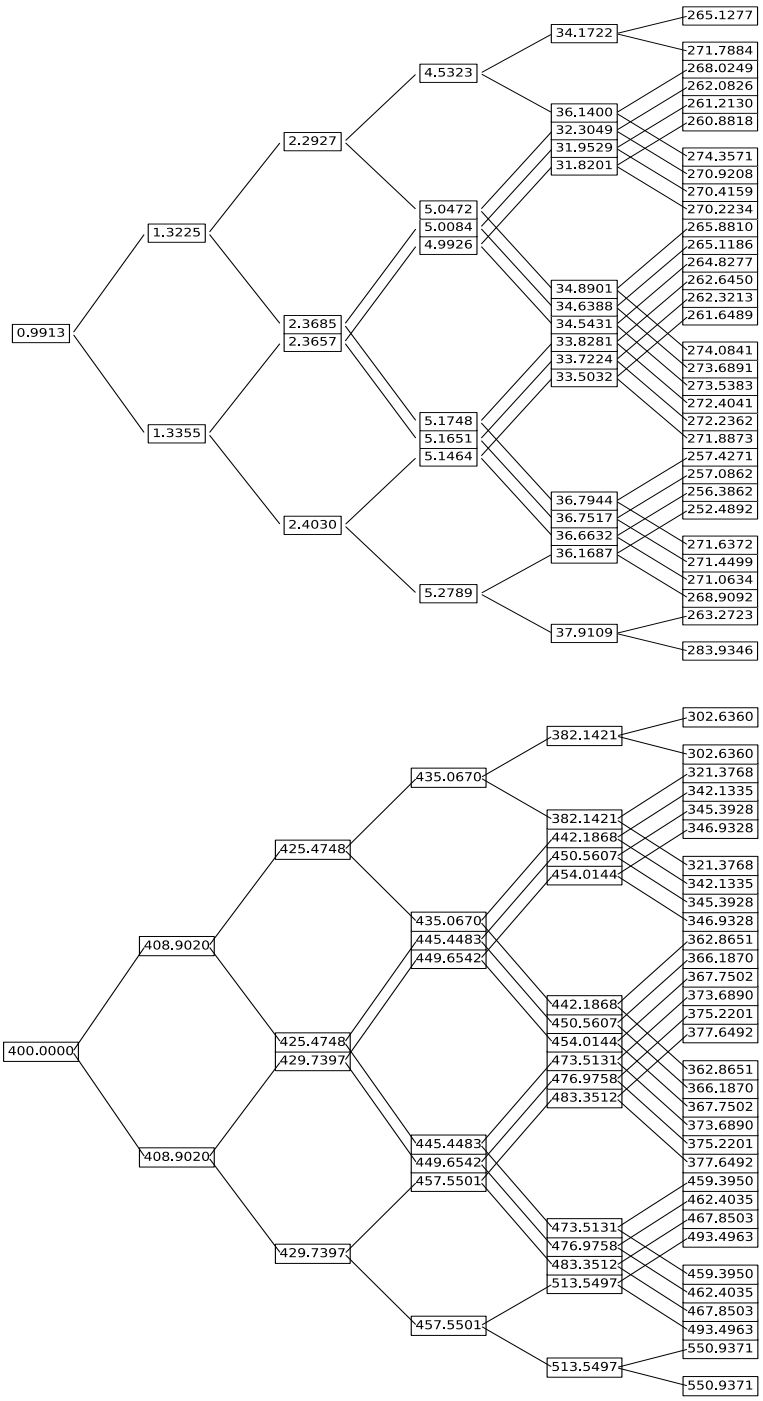

Figure 13—Consumption (upper panel, as multiple of today's consumption) and greenhouse gases (lower panel, in ppm) across time and states

Figure 14 shows how the pricing kernel (Hansen and Richard, 1987) evolves over time and across fragility states. The pricing kernel in a given state is the ratio of marginal utility with respect to current consumption in that state to the marginal utility today, that is $m_{s,t} = \left(\frac{\partial U}{\partial c_{s,t}}\right) / \left(\frac{\partial U}{\partial c_0}\right)$. So, for example, the pricing kernel in the bottom right state

in Figure 14 is $9.87 \cdot 10^{-6}$. Intuitively this means that, to make the representative agent in at least as well off in that state, \$1 (real) in decreased consumption today, as a result of mitigation, has to result in *at least* an increase in consumption in the lower right state of about \$100,000 $[= 1/(9.87 \cdot 10^{-6})]$. This large number reflects the fact that the agent's consumption, between 2015 and 2300, increases so dramatically (c.f. Figure 13).

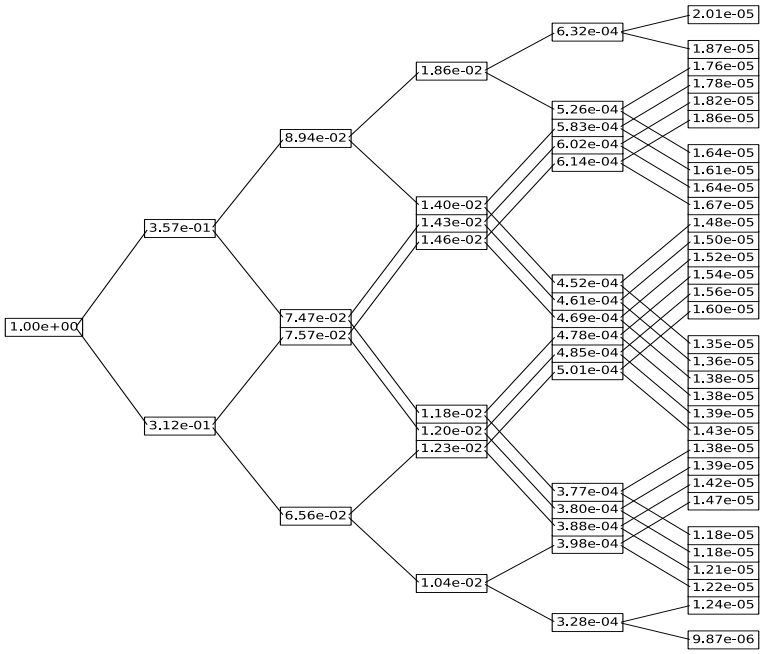

Figure 14. Pricing Kernel as a function of time and state

Another item of interest in Figure 14 is that the pricing kernel does not change much across states. Going from the least to the most fragile state in the final period increases the pricing kernel by about a factor of 2. This is again consistent with the consumption across states shown in Figure 13. The climate damages drive down consumption considerably in the most-fragile state (see, e.g., the upper right corner of Figure 13) but consumption is still very high relative to today's level.

Finally, Figure 15 presents the cost to society of mitigation, expressed as a fraction of total consumption. Perhaps the most striking set of numbers are the costs at $t = 3$: In the most fragile state, 11.4% of consumption is "spent" on mitigation. In contrast, in the least fragile state, only about 1% of consumption needs to be spent on mitigation efforts.

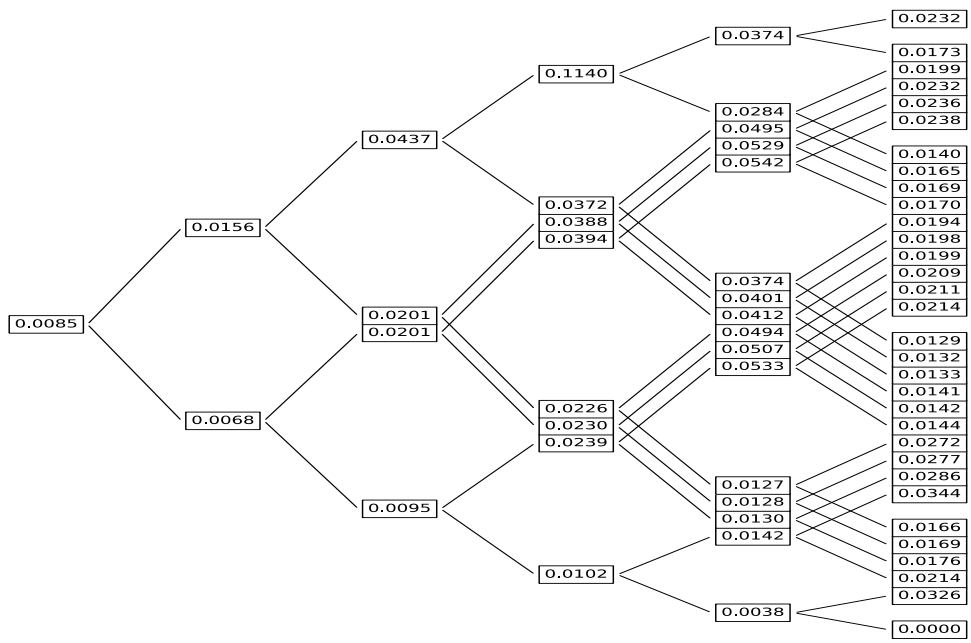

Figure 15. Mitigation cost as a fraction of consumption

3.2 Sensitivity analyses

Note that in our model the SCC at any t reflects the expected discounted damage of the marginal ton of emissions at time t, conditional on an optimal policy beginning now, at $t = 0$; i.e. that emissions are priced immediately at the appropriate level. Since those incentives are not currently in place, and might not be put in place soon, one could also be interested in the marginal damage created by emissions today conditional on a delay in the implementation of a pricing policy.

Define the constrained SCC as the expected discounted damage of the marginal ton of emissions today, conditional on a specified delay in the implementation of an emissions price. For example, in our model we can calculate the cost of a constraint of not pricing emissions during the first period, which is 15 years. It turns out that this constraint almost triples the cost of emissions, it raises the expected discounted damage of the marginal ton of emissions today from $40 to $112. In the context of our model, we can make this calculation by first finding the optimal emissions plan, subject to the constraint of zero emissions mitigation in the first period, and then finding the increase in consumption required to equal the change in utility from a marginal increase in emissions.

We can also calculate the deadweight loss created by such a constraint. We find the increase in consumption today required to bring the utility of an agent inhabiting a constrained world up to the utility of an unconstrained world. The answer is that an increase in consumption of 6% throughout the first period is required to compensate for the deadweight loss created by not pricing emissions during the first period.

The mitigation choices that optimize the utility of the representative agent in the tree structure and determine the SCC depend on a host of factors. We test them in turn. For one, two key assumptions concern the rate of technological change and the potential for a backstop technology. We analyze both exogenous and endogenous technical change. The former has some surprising results, as the SCC first increases before decreasing again (Figure 16). As technical improvements change from 0% to 1%, the SCC rises, only to fall again with annual technical improvements at 1.5% or higher.

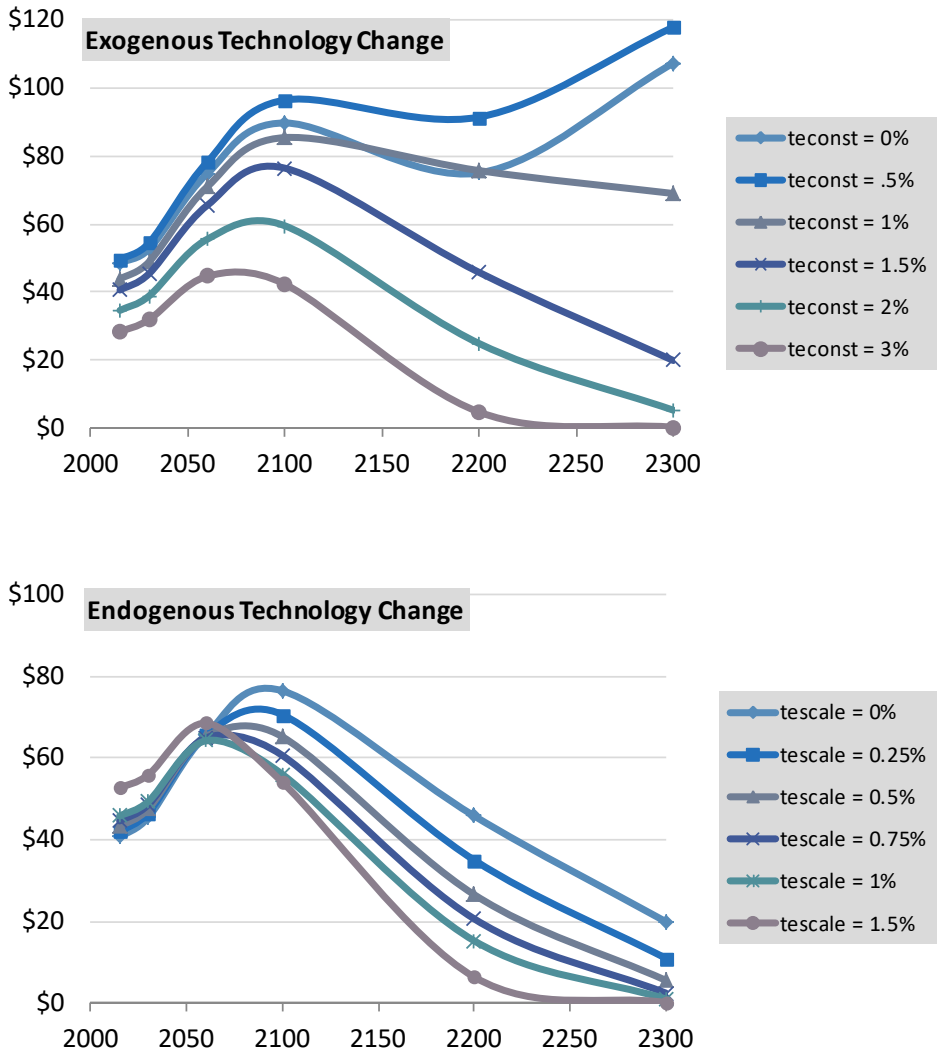

Figure 16—SCC decreases with increased exogenous or endogenous technological change

Endogenous technological change also interacts with a backstop technology, even one as high as, for example, $350 per ton of CO_2 (Figure 17).[14] A backstop plus induced,

[14] See Figure 24 on page 37 for an exploration of different backstop technology thresholds, under a high-risk scenario.

endogenous technological change at first leads to a slightly higher SCC. Over time, however, endogenous technological change helps decrease the SCC as well.

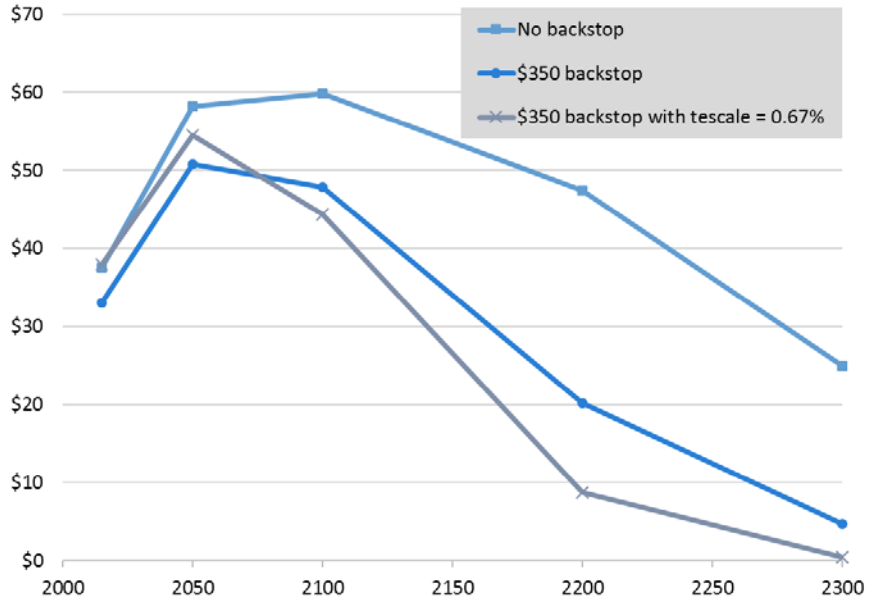

Figure 17—SCC decreases with backstop, with our without endogenous technological change

Next we investigate the effect of growth and interest rates, with corresponding changes in the IMRS. Figure 18 shows SCC values with increasing growth rates around the central case of 2.0%. In this set of results we adjust the IMRS as we change the assumed growth rate in order to hold constant the real interest rate at 2.74%.

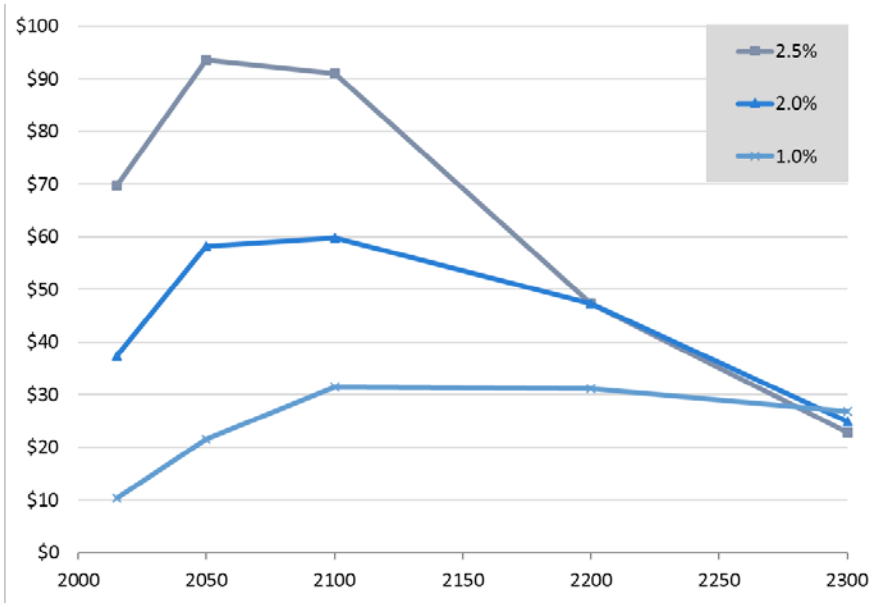

Figure 18—SCC increases with higher assumed economic growth rates

Figure 19 shows results for different rates of interest around the central 2.74% case, generated by choosing different values for the IMRS. High growth rates (holding the IMRS constant) lead to higher SCC values, whereas higher interest rates (holding the economic growth rate constant) have the opposite effect.

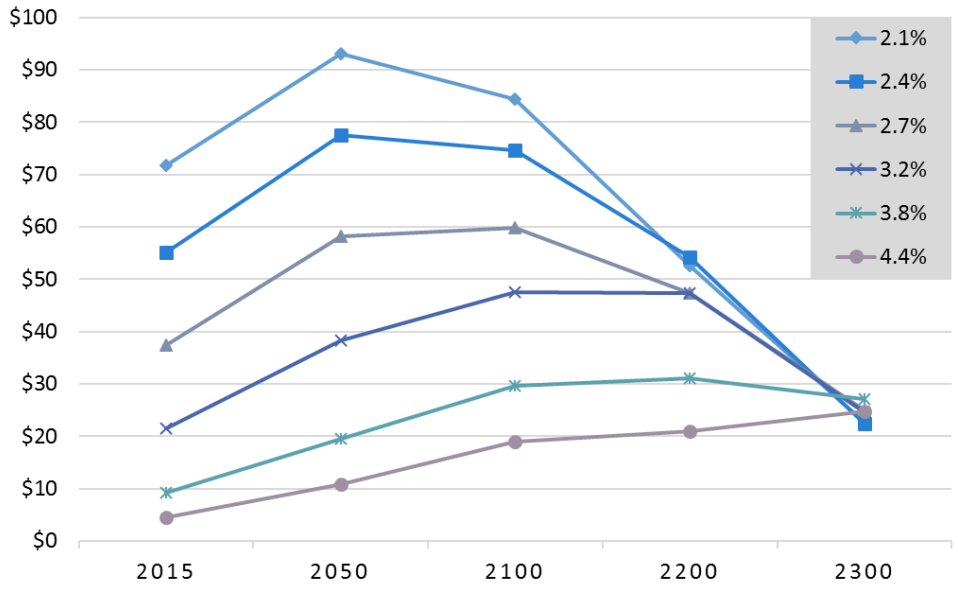

Figure 19—SCC decreases with higher real interest rates

Figure 20 shows the implications of changing the pure rate of time preference. The upper graph, testing implications of changing the pure rate of time preference, holds fixed the implied real yield on a zero-coupon bond that matures at the end of the last period. In early years, a lower discount associated with pure time preference increases the SCC, while the reverse is true in later years. This perhaps counter-intuitive result occurs because we hold the interest rate constant which implies a lower elasticity of substitution and a greater desire to smooth consumption. When the climate model is calibrated to observed interest rates the pure rate of time preference is not a particularly important determinant of the emissions price.

The lower graph of Figure 20 changes the rate of time preference and holds the IMRS fixed: in this case the higher the rate of time preference, the higher is the interest rate and the lower is the SCC.

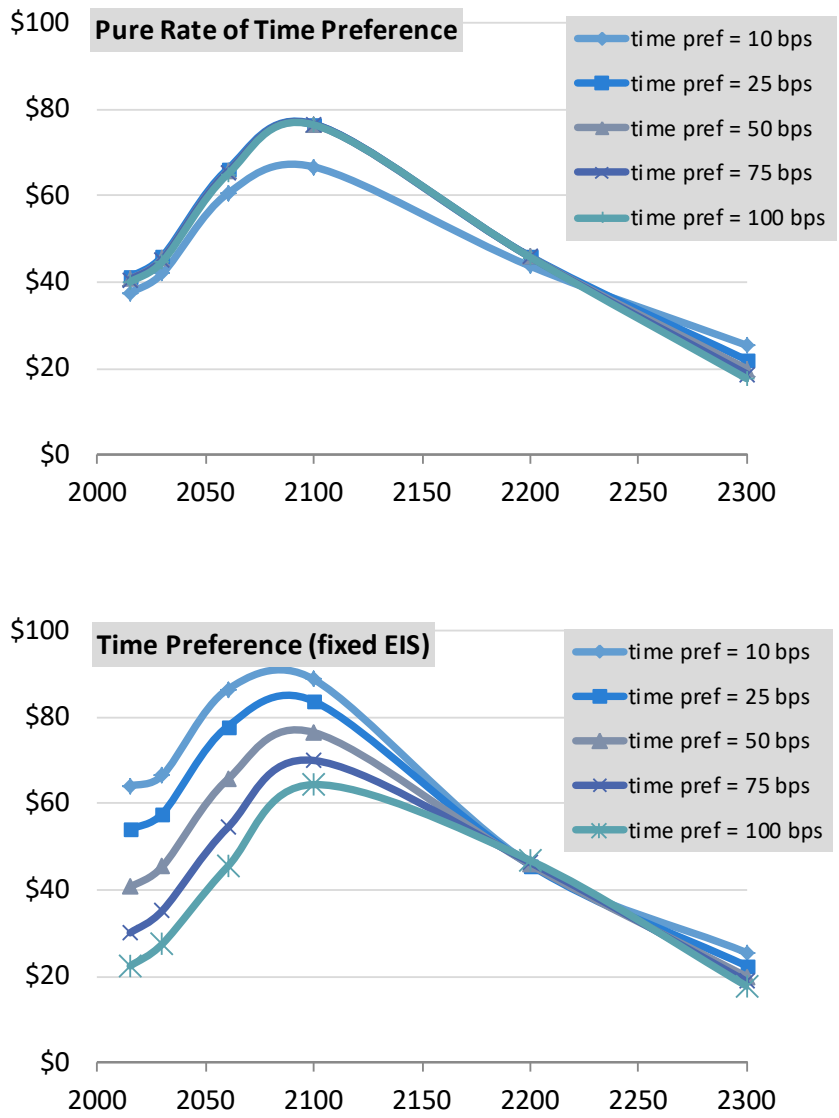

Figure 20—As time preference increases and the IMRS is held constant, the SCC decreases

We test for the importance of the peak temperature and disaster tail parameters in Figure 21 and Figure 22, respectively. As *peakT* decreases—and, thus, as the probability of reaching a climatic tipping point at any particular temperature increases—the optimal carbon cost generally increases (except for very low levels of *peakT*, when the increased risk causes higher mitigation early on, thus resulting in lower expected rates of mitigation and cost in the distant future).

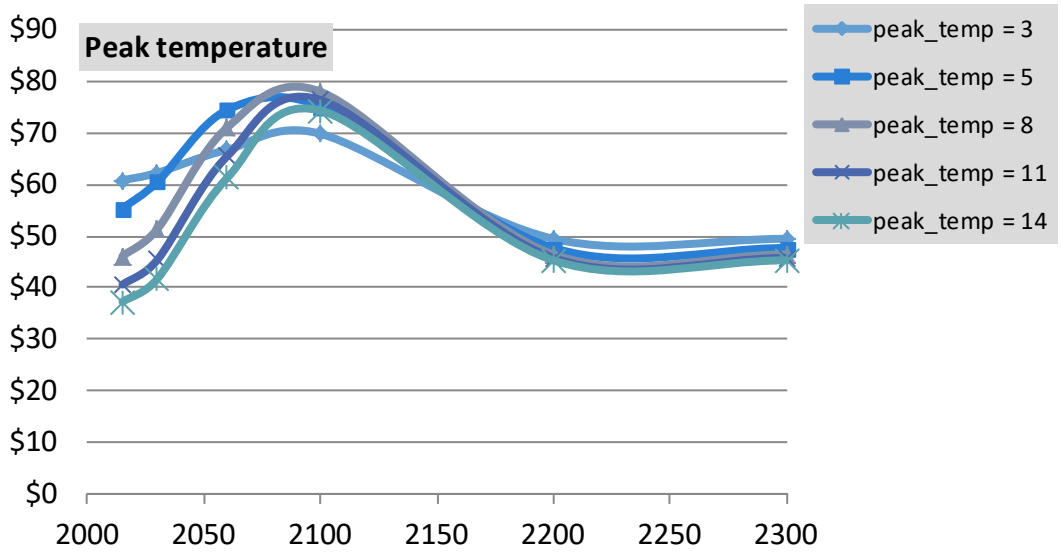

Figure 21—As the peak temperature parameter increases, and the likelihood of reaching a tipping point at lower temperatures decreases, the SCC decreases

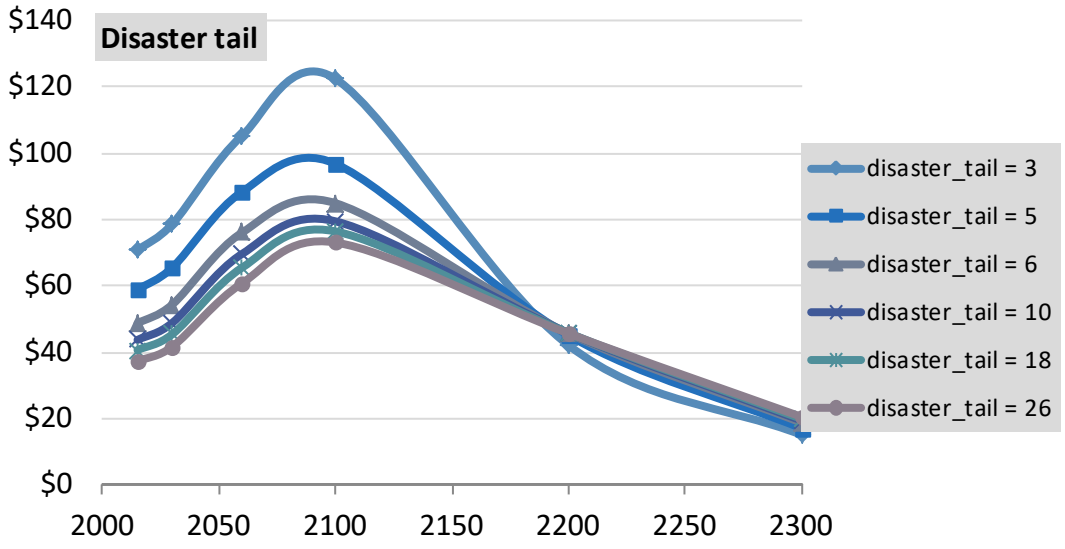

Figure 22—As the disaster tail parameter increases, thinning the tail of the damage distribution, the SCC decreases

Our standard case includes periods of increasing length, but in Figure 23 we display the effect of varying the duration of fixed length periods: as period length increases, the expected price path pivots downward and optimal carbon price today increases.

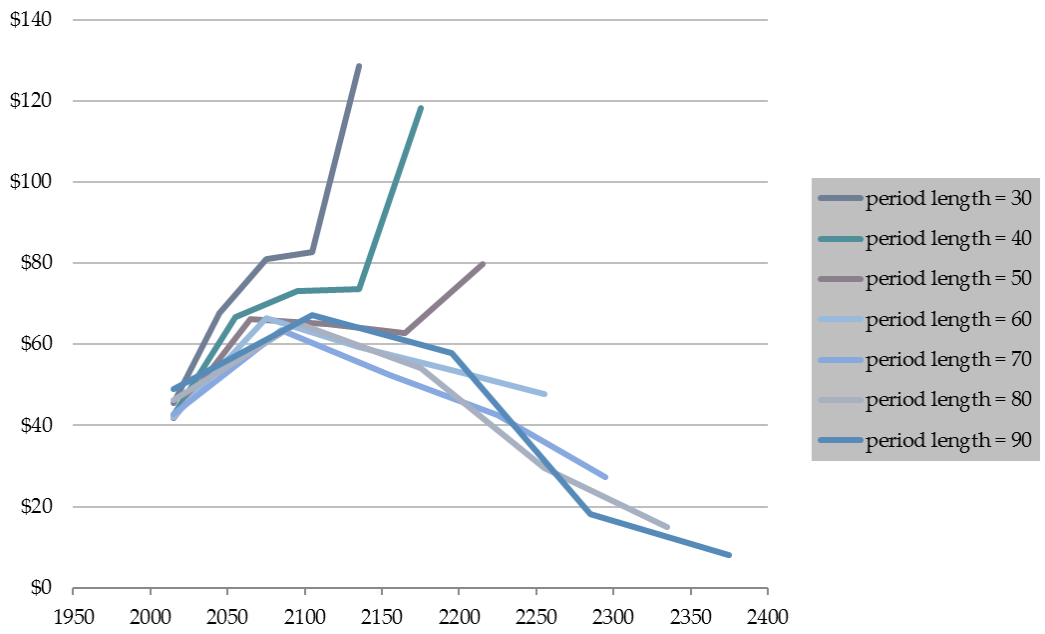

Figure 23—As the period length increases, the SCC tends to decrease

3.3 'High risk' decomposition

Our base case in this paper has been a model that has economic parameters roughly calibrated to match historical real interest rates, at 2.74%, a low equity risk premium of less than 1%, and with a damage distribution calibrated to match a value of the SCC of $40. As seen in Figure 2, with this calibration the risk premium component of the SCC is only 7% of the total. We have not attempted to calibrate the appropriate damage distribution as a function of emissions.[15] We have instead chosen parameters that lead to a value of the SCC that approximately matches the current US government estimate (U.S. Government Interagency Working Group on Social Cost of Carbon, 2015). Above we report on the impact of changing these parameters one at a time.

We also report here on what we call a "high risk" case, and show the impact of choosing a different set of parameters designed to increase the risk component of the SCC in our model. In the high risk case these parameters are $peakT = 4$ and $disaster\_tail = 4.5$. compared with $peakT = 11$ and $disaster\_tail = 18$ for the base case. The implications in terms of increased probabilities for tipping points and increased damage conditional on a tipping point can be seen in Figure 6 and Figure 7 above.

To further increase the risk component of the SCC decomposition, we increased the degree of risk aversion from 7 to 9, reduced the IMRS from 0.9 to 0.6, lowered the exogenous potential growth rate of consumption from 2% annually to 1%, and lowered the assumed exogenous rate of technological progress with respect to emissions

[15] One notable efforts in this area comes courtesy of the Risky Business Project and its work on the American Climate Prospectus: http://climateprospectus.org/ (Houser et al., 2015). Kopp et al. (2016) attempts to lay out a path to integrating possible climate-economic tipping points in such an analysis.

reduction from 1.5% per year to 0.8% per year. With these changes the real interest rate drops from 2.74% to 2.27%, and the SCC rises from around $40 to $156. The risk decomposition changes from 93% expected damage and 7% risk premium in the base case to 82% expected damage and 18% risk premium in the high risk case.

The 'high risk' parameterization has a particularly dramatic impact on the cost of delay. In the base case, the constraint of waiting 15 years to price emissions causes the SCC to almost triple from $40 to $112. With the 'high risk' parameterization, the same constraint causes the SCC to jump from $156 to $451 per ton. And finally, rather than the 6% of consumption deadweight cost of delay, in the high risk parameterization the deadweight cost of a 15-year delay in pricing emissions is a loss of utility equivalent to a 36% reduction in consumption during that period.

As expected, the high-risk specification also highlights, once again, the importance of the backstop technology. The availability of a binding backstop technology decreases the SCC in all time periods, providing a firm upper bound (Figure 24).

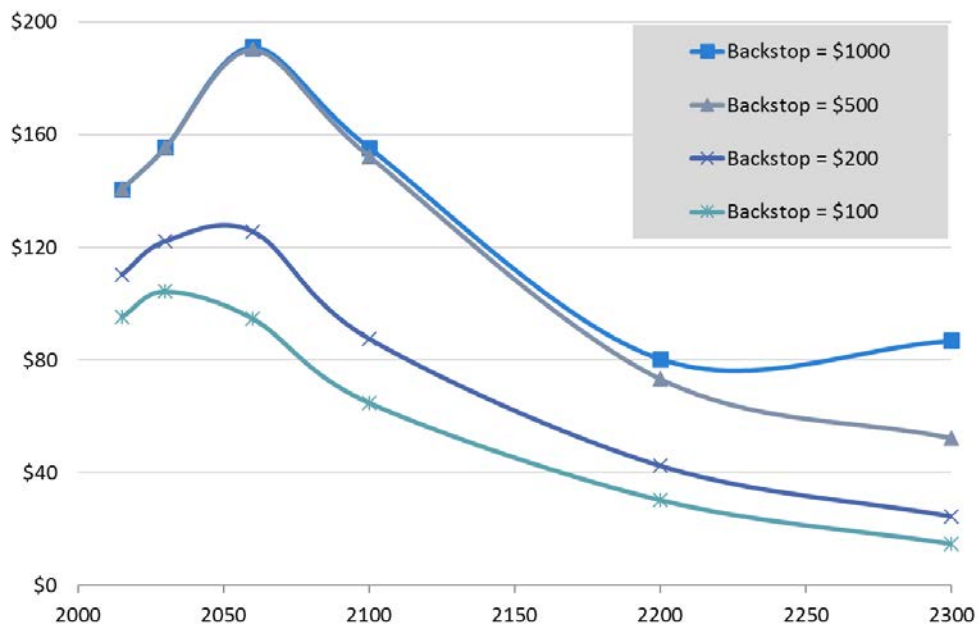

Figure 24—A binding air capture backstop technology lowers the SCC in the high-risk case

3.4 A 'Low growth' scenario

Another key parameter in our analysis is the rate of economic growth of 2% per year (ignoring damages and mitigation costs). Our assumption is consistent with empirically observed growth rates over the last century. However, a number of scholars now argue that we should anticipate far lower growth going forward than we have observed in the last century (e.g., Gordon, 2016, 2012).

This growth rate is a particularly interesting parameter to explore in the context of a carbon pricing analysis. As noted earlier, a large fraction of the damages that from climate change are likely to occur far into the future. With the assumption of a 2%/year consumption growth rate, the level of per capita consumption in 2300, even the worst state of nature, is more than 250 times higher than what it is today (see Figure 13). As a result of this large disparity in current and future consumption levels, future damages are discounted at a very high rate.[16] However, if future real consumption growth is 1%/year instead of 2%, the resulting level of consumption in 2300 will be about 95% lower. This implies that the any damages will be far more important to the current generation factor (ie., the discount factor applied to these damages will be far closer to unity).

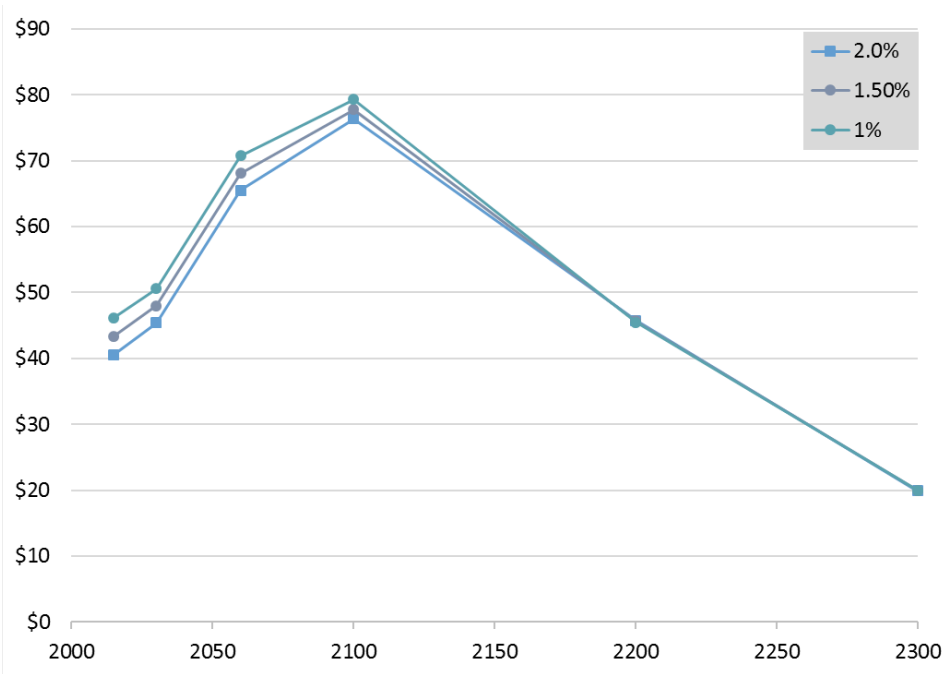

Figure 25—SCC, by year, as a function of the assumed base-case consumption growth rate.

Figure 25 plots the SCCs, as a function of time, that result from running our simulation with different economic growth rates, and holding preferences constant.[17] Given our specification, changing the assumed growth rate from 2%/year to 1%/year results in only a small increase in the 2015 SCC—from $41 to $46/ton. Given the considerably smaller discount rate, why is it that the SCC does not increase more in response to this change? The answer lies in our damage specification in equation(2), where damages, for a given level of fragility and cumulative mitigation, are proportional to the baseline level of consumption $\bar{c_t} = \bar{c_0} \cdot (1+g)^t$. Thus, in this alternative specification with a 1% annual growth rate, and consumption in year 2300 that is lower by a factor of 20, we are also implicitly specifying that damages are lower by a factor of 20. Whether it is reasonable

[16] In the base case, with a growth rate $g = 2\%/year$, the riskfree rate is 2.74%/year. With a lower $g = 1\%/year$, the risk-free rate falls to 1.62%/year.

[17] For each of the simulations illustrated in Figure 25, the base case preference specification is employed, meaning that the elasticity of intertemporal substitution is 0.90.

to assume that, if people are poorer in the future, they will be hurt proportionally less by climate change is worthy of future research.

3.5 Future research and extensions

We have emphasized throughout this paper that the model and model parameterization employed here, both in the base case and in the 'high risk' case, are *ad hoc* and merit considerable future scrutiny by the scientific community. What we believe is unique and innovative is the methodology and solution approach outlined here. The framework better takes account of climatic risks, without putting undue computational burden on the analyst.

There are many possible avenues for future work and research. First and foremost is a better calibration of some of the key parameters, perhaps led by the probability and impact of tail risks. Our base case assumes a *peakT* of 11 and *disaster_tail* of 18. While we explore various specifications, including the 'high risk' case, none is empirically grounded—i.e. linked to actual probabilities and impacts of low-probability, high-consequence climatic events. A starting point might be some of the guidance provided by Kopp *et al.* (2016): start with candidate 'climatic tipping elements' and take the growing body of empirical, econometric analyses to estimate their effects. Second, employ experiments with empirical, process-based impact models to assess the relative importance of different tipping elements. Third, look at social tipping elements such as civil conflict and their associated costs. Projects like Risky Business are beginning to quantify the impacts of some of these climatic and social tipping elements (Houser et al., 2015). Finally, where data are scarce, conduct structured expert elicitations to generate probability functions around tipping elements and their possible economic impacts (e.g., Cooke, 2013). These calibrations are much needed to derive defensible estimates of the SCC that stand independent of current U.S. SCC numbers.

Even now, though, the model lends itself to many extensions. For example, if climate damages hit growth rates as opposed to levels of economic output, one could similarly assume the opposite, likely raising the appropriate SCC today (Bansal and Ochoa, 2011; Burke et al., 2015; Dell et al., 2012; Heal and Park, 2013; Moore and Diaz, 2015). None of those exploring this important question have done so within the context of an Epstein-Zin utility framework.

Another possible extension is around climate impacts on poor versus rich. Our representative agent framework models the average, globally representative consumer. One could easily imagine different climate based on wealth and income. While the rich may have more capital at risk, it would likely be the poor who will be hurt proportionally more, as they are less able to adapt (e.g., Wagner and Weitzman, 2015).

4. Conclusion

An oft-told analogy in climate economics represents the climate system as a hard-to-navigate ocean liner. This is used to argue for early action through a slow and gradually increasing carbon price. Too strong a policy early on would be overly costly; a small course correction now will save us from hitting the far-off proverbial iceberg. There are clearly costs of action, but as our simulations show, once we do include a proper accounting of risk aversion and extreme events, this standard logic gets turned on its head: The optimal carbon price may, in fact, be high today, declining over time.[18]

That decline reflects the rate at which information is revealed going forward, the degree of risk aversion, and the potential for technological progress and backstop technologies. Either way, the 'ocean liner' logic doesn't hold. Or perhaps it gets completed: for turning a large ship long down the line takes bearing off decisively and early, especially in a world of uncertain obstacles. The less certain we are about the risks facing us in future states of the world, the higher the optimal price on carbon today.

Acknowledgments

For helpful comments and discussions, and without any implications, we thank Jeffrey Bohn, V.V. Chari, Christian Gollier, Cameron Hepburn, Christian Traeger, Martin Weitzman, Matthew Zaragoza-Watkins, Richard Zeckhauser, Stanley Zin, and seminar participants at: the 2015 American Economic Association meetings; the Environmental Defense Fund; the Journal of Investment Management (JOIM) conference on Long-run Risks, Returns and ESG Investing; and the University of Minnesota. For excellent research assistance, we thank Katherine Rittenhouse.

Bibliography

Ackerman, F., Stanton, E.A., Bueno, R., 2013. Epstein–Zin Utility in DICE: Is Risk Aversion Irrelevant to Climate Policy? Environ. Resour. Econ. 56, 73–84. doi:10.1007/s10640-013-9645-z

Bansal, R., Ochoa, M., 2011. Temperature, Aggregate Risk, and Expected Returns (Working Paper No. 17575). National Bureau of Economic Research.

Bansal, R., Ochoa, M., 2009. Temperature, Growth, and Asset Prices.

Bansal, R., Yaron, A., 2004. Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles. J. Finance 59, 1481–1509. doi:10.1111/j.1540-6261.2004.00670.x

Burke, M., Hsiang, S.M., Miguel, E., 2015. Global non-linear effect of temperature on economic production. Nature 527, 235–239. doi:10.1038/nature15725

Coase, R.H., 1960. The Problem of Social Cost, in: Gopalakrishnan, C. (Ed.), Classic Papers in Natural Resource Economics. Palgrave Macmillan UK, pp. 87–137.

[18] Ulph and Ulph (1994) similarly derive conditions for a declining carbon price, though driven by optimal resource extraction in the context of climate policy.

Convery, F.J., Wagner, G., 2015. Reflections–Managing Uncertain Climates: Some Guidance for Policy Makers and Researchers. Rev. Environ. Econ. Policy 9, 304–320. doi:10.1093/reep/rev003
Cooke, R.M., 2013. Expert judgement assessment: Quantifying uncertainty on thin ice. Nat. Clim. Change 3, 311–312. doi:10.1038/nclimate1860
Cox, J.C., Ross, S.A., Rubinstein, M., 1979. Option pricing: A simplified approach. J. Financ. Econ. 7, 229–263. doi:10.1016/0304-405X(79)90015-1
Crost, B., Traeger, C.P., 2014. Optimal CO2 mitigation under damage risk valuation. Nat. Clim. Change 4, 631–636. doi:10.1038/nclimate2249
Dell, M., Jones, B.F., Olken, B.A., 2012. Temperature Shocks and Economic Growth: Evidence from the Last Half Century. Am. Econ. J. Macroecon. 4, 66–95. doi:10.1257/mac.4.3.66
Duffie, D., 2010. Dynamic Asset Pricing Theory. Princeton University Press.
Epstein, L.G., Zin, S.E., 1991. Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: An Empirical Analysis. J. Polit. Econ. 99, 263–286.
Epstein, L.G., Zin, S.E., 1989. Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework. Econometrica 57, 937–969. doi:10.2307/1913778
Golosov, M., Hassler, J., Krusell, P., Tsyvinski, A., 2014. Optimal Taxes on Fossil Fuel in General Equilibrium. Econometrica 82, 41–88. doi:10.3982/ECTA10217
Gordon, R.J., 2016. The Rise and Fall of American Growth: The U.S. Standard of Living since the Civil War. Princeton University Press.
Gordon, R.J., 2012. Is U.S. Economic Growth Over? Faltering Innovation Confronts the Six Headwinds (Working Paper No. 18315). National Bureau of Economic Research.
Goulder, L.H., 1995. Environmental taxation and the double dividend: A reader's guide. Int. Tax Public Finance 2, 157–183. doi:10.1007/BF00877495
Hansen, L.P., Heaton, J.C., Li, N., 2008. Consumption Strikes Back? Measuring Long-Run Risk. J. Polit. Econ. 116, 260–302. doi:10.1086/588200
Hansen, L.P., Richard, S.F., 1987. The Role of Conditioning Information in Deducing Testable Restrictions Implied by Dynamic Asset Pricing Models. Econometrica 55, 587–613. doi:10.2307/1913601
Heal, G., Park, J., 2016. Temperature Stress and the Direct Impact of Climate Change: A Review of an Emerging Literature. Rev. Environ. Econ. Policy rew007. doi:10.1093/reep/rew007
Heal, G., Park, J., 2013. Feeling the heat: Temperature, physiology & the wealth of nations (Working Paper No. 19725). National Bureau of Economic Research.
Houser, T., Hsiang, S., Kopp, R., Larsen, K., 2015. Economic Risks of Climate Change: An American Prospectus. Columbia University Press.
International Energy Agency, 2013. World Energy Outlook.
IPCC, 2013. Fifth Assessment Report: Climate Change.
IPCC, 2007. Fourth Assessment Report.
Jensen, S., Traeger, C.P., 2014. Optimal climate change mitigation under long-term growth uncertainty: Stochastic integrated assessment and analytic findings. Eur. Econ. Rev., Sustainability and Climate Change: From Theory to Pragmatic Policy 69, 104–125. doi:10.1016/j.euroecorev.2014.01.008

Kopp, R.E., Shwom, R., Wagner, G., Yuan, J., 2016. Tipping elements and climate-economic shocks: Pathways toward integrated assessment. Earths Future 2016EF000362. doi:10.1002/2016EF000362
Lemoine, D., 2015. The Climate Risk Premium.
Lemoine, D., Traeger, C., 2014. Watch Your Step: Optimal Policy in a Tipping Climate. Am. Econ. J. Econ. Policy 6, 137–166. doi:10.1257/pol.6.1.137
Lemoine, D., Traeger, C.P., 2016. Economics of tipping the climate dominoes. Nat. Clim. Change 6, 514–519. doi:10.1038/nclimate2902
McKinsey, G.I., 2009. Pathways to a Low-Carbon Economy Version 2 of the Global Greenhouse Gas Abatement Cost Curve. McKinsey & Company, Stockholm.
Mehra, R., Prescott, E.C., 1985. The equity premium: A puzzle. J. Monet. Econ. 15, 145–161. doi:10.1016/0304-3932(85)90061-3
Moore, F.C., Diaz, D.B., 2015. Temperature impacts on economic growth warrant stringent mitigation policy. Nat. Clim. Change 5, 127–131. doi:10.1038/nclimate2481
National Research Council, 2015. Climate Intervention: Carbon Dioxide Removal and Reliable Sequestration. National Academies Press, Washington, D.C.
Nordhaus, W.D., 2014. Estimates of the Social Cost of Carbon: Concepts and Results from the DICE-2013R Model and Alternative Approaches. J. Assoc. Environ. Resour. Econ. 1, 273–312. doi:10.1086/676035
Nordhaus, W.D., 2013. The Climate Casino: Risk, Uncertainty, and Economics for a Warming World. Yale University Press.
Nordhaus, W.D., 1994. Expert Opinion on Climatic Change. Am. Sci. 82, 45–51.
Nordhaus, W.D., 1992. An Optimal Transition Path for Controlling Greenhouse Gases. Science 258, 1315–1319. doi:10.1126/science.258.5086.1315
Nordhaus, W.D., 1991. A Sketch of the Economics of the Greenhouse Effect. Am. Econ. Rev. 81, 146–50.
Nordhaus, W.D., Boyer, J., 2000. Warming the World: Economic Models of Global Warming. MIT Press.
Pindyck, R.S., 2013. The Climate Policy Dilemma. Rev. Environ. Econ. Policy 7, 219–237.
Pindyck, R.S., 2012. Uncertain outcomes and climate change policy. J. Environ. Econ. Manag. 63, 289–303. doi:10.1016/j.jeem.2011.12.001
Robert Shiller, 2000. Irrational Exuberance [WWW Document]. URL http://press.princeton.edu/titles/10421.html (accessed 8.17.16).
Roe, G.H., Baker, M.B., 2007. Why Is Climate Sensitivity So Unpredictable? Science 318, 629–632. doi:10.1126/science.1144735
Stern, N., 2013. The Structure of Economic Modeling of the Potential Impacts of Climate Change: Grafting Gross Underestimation of Risk onto Already Narrow Science Models. J. Econ. Lit. 51, 838–859. doi:10.1257/jel.51.3.838
Summers, L., Zeckhauser, R., 2008. Policymaking for posterity. J. Risk Uncertain. 37, 115–140. doi:10.1007/s11166-008-9052-y
Traeger, C.P., 2015. Analytic Integrated Assessment and Uncertainty.
Ulph, A., Ulph, D., 1994. The Optimal Time Path of a Carbon Tax. Oxf. Econ. Pap. 46, 857–868.

U.S. Government Interagency Working Group on Social Cost of Carbon, 2015. Technical Update of the Social Cost of Carbon for Regulatory Impact Analysis Under Executive Order 12866.

Wagner, G., Weitzman, M.L., 2015. Climate Shock: The Economic Consequences of a Hotter Planet. Princeton University Press, Princeton, NJ.

Wagner, G., Zeckhauser, R.J., 2016. Confronting Deep and Persistent Climate Uncertainty. Harvard Kennedy School, Faculty Research Working Paper Series, No. 16-025.

Weil, P., 1989. The equity premium puzzle and the risk-free rate puzzle. J. Monet. Econ. 24, 401–421. doi:10.1016/0304-3932(89)90028-7

Weitzman, M.L., 2009. On Modeling and Interpreting the Economics of Catastrophic Climate Change. Rev. Econ. Stat. 91, 1–19. doi:10.1162/rest.91.1.1

Zeckhauser, R.J., 2006. Investing in the Unknown and Unknowable. Capital. Soc. 1/2.