Applying Asset Pricing Theory to Calibrate the Price of Climate Risk

Kent D. Daniel, Robert B. Litterman, and Gernot Wagner[1]
December 29, 2014

Working Draft – Please do not cite without permission
Comments Welcome

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 price is necessarily 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. The climate-economy literature by and large has not adopted this richer set of preferences.

In this paper, we explore the implications of these richer preference specifications for the optimal pricing of carbon emissions. We develop a simple discrete-time model with Epstein-Zin utility in which uncertainty about the effect of carbon emissions on global temperature and on eventual damages is gradually resolved over 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 to historical real interest rates and the equity risk premium suggests a high price for carbon emissions today which is then expected to decline over time. This is in contrast to most modeled carbon price paths, which tend to start low and rise steadily over time.

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

JEL code: D81, G11, Q54.

---

1. Introduction

Following Coase (1960), external damages from greenhouse gas emissions into the atmosphere can be accounted for by assigning full property rights to the atmosphere and allowing trading of those rights. While property rights to the atmosphere cannot reasonably be assigned in this way and enormous transactions costs prevent any true Coasian solution, we can nonetheless achieve the optimal level of emissions by setting a carbon price equal to its full social cost. The modern approach to asset pricing recognizes that the value of the full benefits achieved by reducing emissions is generally determined by appropriate discounting of those benefits at future times and states of nature (Hansen and Richard, 1987; Duffie, 2001). In practice this can be done by weighing the future benefits not only by a pure rate of time discount, but also by the probability and the marginal utility of consumption appropriate to each state.

Meanwhile, climate modelers and environmental economists more broadly have long focused on ethical justifications for different values for the pure rate of time preference to value the costs of unmitigated climate change and the benefits of investments leading to lower greenhouse gas (GHG) pollution.[2] Another more recent focus has been on the implication of a declining term-structure of discount rates associated with uncertainty of future interest rates.[3] While relatively new in the climate economic literature, the notion that higher uncertainty about future interest rates causes the term structure of interest rates to decline is well-understood in the asset pricing literature (e.g. Litterman, Scheinkman, and Weiss, 1991). However, this climate-economic literature for the most part all but ignores the pricing of the risk in the payoffs arising from the mitigation of climate emissions.[4] As we show here the risk premium can easily become the dominant influence on the appropriate value for the current price of emissions because climate science cannot rule out low-probability, high-impact outcomes.

Improper accounting for low-probability, high-impact events in standard economic approaches to valuing returns on equities is one explanation for large equity premium puzzles (Weitzman, 2007).[5] The valuation assigned to different traded assets suggests that society requires a far higher rate of return on investments in risky assets, and is willing to tolerate far lower returns on assets that provide insurance against unfavorable outcomes. For example, between 1871 and 2012, an investment in a portfolio in U.S. bonds would have earned an average annual real return of 1.6 percent. In contrast, over the same period an investment in a diversified portfolio in U.S. stocks would have earned an average annual real return of 6.4 percent [6]

Society requires a higher return for equities because equities have large payoffs in good economic times but often perform poorly precisely when the economy is doing poorly.

---

[2] For some of the most prominent examples, see Stern (2007) and also Nordhaus (2013, 2014).
[3] E.g., Arrow et al. (2013, 2014). See Gollier (2012) for a comprehensive survey of discounting under uncertainty. Weitzman (2001) first formalized the idea of declining long-term discount rates in the context of climate policy. Gollier and Weitzman (2010) similarly argue for a declining rate.
[4] Two notable exceptions are Ackerman, Stanton, and Bueno (2012) and Crost and Traeger (2014).
[5] See our discussion in section 2.1 and references therein.
[6] Based on data collected by Shiller (2000) and since continuously updated:
http://www.econ.yale.edu/~shiller/data.htm

Conversely, society is willing to pay handsomely for the right pattern of cash flows across states. These numbers also reveal that the pattern across states of the payoffs associated with climate change mitigation will have a dramatic effect on the appropriate discount rate. Because the potential benefits from mitigation accrue many centuries into the future, even small changes in the discount rate can have a large effect on the appropriate emissions price.

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 constant-elasticity of substitution (CES) utility function used in most climate studies, we use here a utility function proposed by Epstein and Zin (1989, 1991) and used throughout the asset pricing literature. It has CES utility as a special case and also allows for differences between the intertemporal marginal rate of substitution (IMRS) and risk aversion, which allows us to calibrate to standard financial returns, in particular the equity risk premium and risk-free interest rates.

To examine the implications we use a simple model that showcases the tradeoff between (known) climate action and (uncertain) climate damages. We employ a discrete time binomial tree model, building on one introduced by 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, based on the information available at that time, that maximizes his expected discounted utility. The optimal level of emissions mitigation is obtained when the reduction in utility from additional 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. The use of Epstein-Zin utility allows us to explore the impact on the optimal emissions price of the higher elasticity of marginal utility to uncertain outcomes associated with a risk-averse society.

Figure 1 shows the profound implications. Under a standard constant relative risk-aversion utility function with the IMRS calibrated to 0.643 in order to reflect a 2.5% yield on a zero-coupon bond that matures at the end of the last period, the optimal carbon price today is $36.28 per ton. With this standard utility specification the optimal carbon price today declines rapidly to zero with increasing risk aversion as the implied interest rates increase and less weight is given to future outcomes. In contrast, under Epstein-Zin utility, holding the IMRS fixed at 0.643 and increasing the degree of risk aversion, the carbon price increases, while the real interest rate remains at around 2.5%/year.[7] 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 $36.28 to around $90 per ton.

---

[7] The exact interest rate in our Epstein-Zin calibration with a risk aversion of 1/.643 starts at 2.506%/year (equivalent to CES), reaching 2.541%/year with risk aversion = 50.
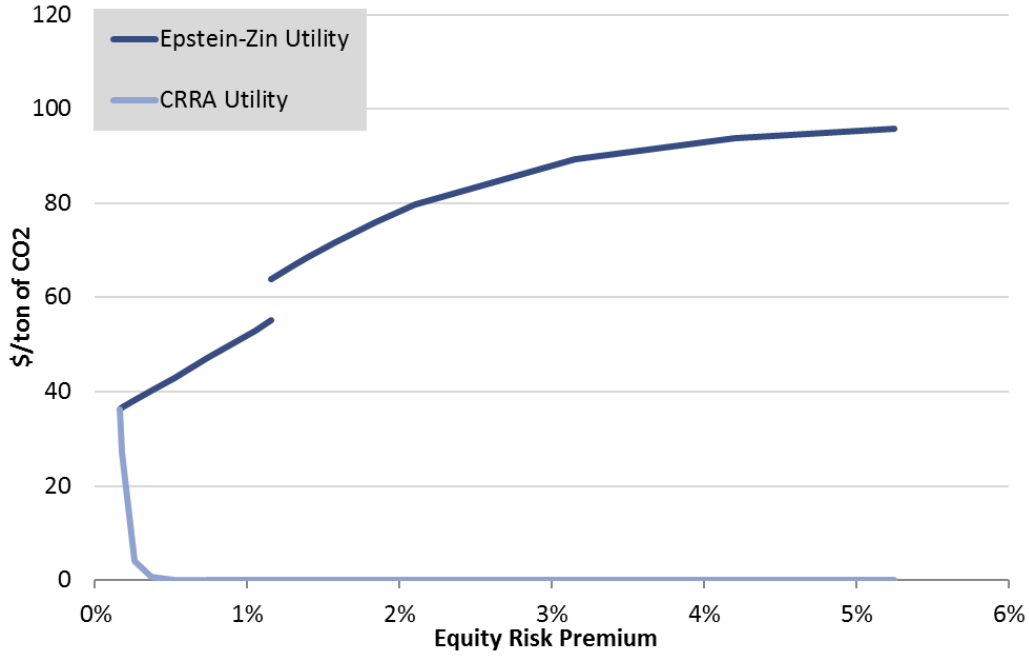
Figure 1 — **Using Epstein-Zin utility functions results in increasing carbon prices with increasing risk aversion translated into the implied equity risk premium using Weil (1989)'s conversion, while holding the implied market interest rate stable at our target rate of 2.5%.[8]**

2. The model set-up

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 is given by:

(1) $\qquad C_0 = \bar{C}_0 \cdot \left(1 - \kappa_0(x_0)\right)$
(2) $\qquad C_t = \bar{C}_t \cdot \left(1 - D_t(X_t, \theta_t) - \kappa_t(x_t)\right)$, for $t\epsilon\{1,2,\dots,T-1\}$
(3) $\qquad C_T = \bar{C}_T \cdot (1 - D_T(X_T, \theta_T))$

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

---

[8] The discontinuity at an equity risk premium of 1.16 reflects a different local maximum becoming the global maximum in our optimization exercise. See Section 4 for a fuller explanation.

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 at any time from 0 to t, which is given by:

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

where g_s is the flow of GHG emissions into the atmosphere in period s, for each period up to t, absent any mitigation. 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 zero net flow of new GHG emissions into the atmosphere at time s. In our base specification, $x_s \epsilon [0,1]$. In Section 2.2.1, we extend this to include a *backstop technology,* which allows for CO_2 capture, and which can result in mitigation above 100%.

Mitigation reduces the stock of GHGs in the atmosphere and leads to lower climate damages and hence higher future consumption. However, mitigating GHG emissions is costly. To mitigate 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.

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, $x_t^*(\theta_t)$, to maximize lifetime utility at each time and for each state of nature. In our baseline analysis, we start with a 5 period tree with 30-year intervals, beginning in 2015. The subsequent years for calibration are: 2045, 2075, 2105, and 2135, followed by $t = T$ at 2200. At each node of the tree, more information about the consumption damage function is revealed (and reflected in the parameter θ_t), but uncertainty is not fully resolved until the final period.

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 discusses the discontinuity in Figure 1, and Section 5 concludes.

2.1 Preferences

The *equity premium puzzle* refers to the empirical regularity that the high average return and low volatility of U.S. equity return is not matched by a corresponding high covariance with consumption growth (Hansen and Singleton, 1982; Mehra and Prescott,

1985; Weil, 1989). Put simply, investors act more risk averse than standard economic theory prescribes they should. These observations are particularly problematic for Constant Relative Risk Aversion (CRRA) utility functions (see, e.g., Hansen and Jagannathan, 1991; Hansen, Heaton, Lee and Roussanov 2007).

To resolve this puzzle, financial economists have long employed an alternative and more flexible preference specification. Epstein and Zin (1989, 1991) attempt to escape the hard link between risk aversion and the intertemporal discount rate by introducing a utility function that allows for different rates of substitution across time and states.[9] This is the specification we use here.

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

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

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:

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

Note that with $\rho = \alpha$, equations (5) and (6) are equivalent to the standard CRRA/power/isoelastic utility specification. Setting $\rho = 0$ and $\alpha = 0$ independently generates standard log-utilities across time and across states of nature, respectively.

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 (as yet unknown) states. $\gamma = (1-\alpha)$ is the coefficient of relative risk aversion at a given point in time. This added flexibility allows for calibration across states of nature and time.

Plugging (6) into (5) for our tree-structure model generates:

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

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

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

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

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 mimic 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, the middle of period 1. 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.[10]

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:

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

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

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

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:

---

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

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

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:[11]

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

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

$$K(\tau) = g_0 \int_0^\tau \left(1 - x(s)\right) ds - \tau g_0 \left(1 - x(\tau)\right)$$
$$= g_0 \left[\tau - \int_0^\tau x(s) ds\right] - \tau g_0 + \tau g_0 x(\tau)$$
(15) $$= g_0 \left[\tau x(\tau) - \int_0^\tau x(s) ds\right]$$

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:

---

[11] 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 (1993). Separately, note that in this analysis we are necessarily not in a representative agent framework—in that we are assuming that individuals choice of the optimal level of GHG emissions takes into account the marginal effect of the tax τ they pay on GHG emissions, but not the marginal effect of the refund of the tax proceeds.

(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.

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). 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 2 and Figure 3 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.
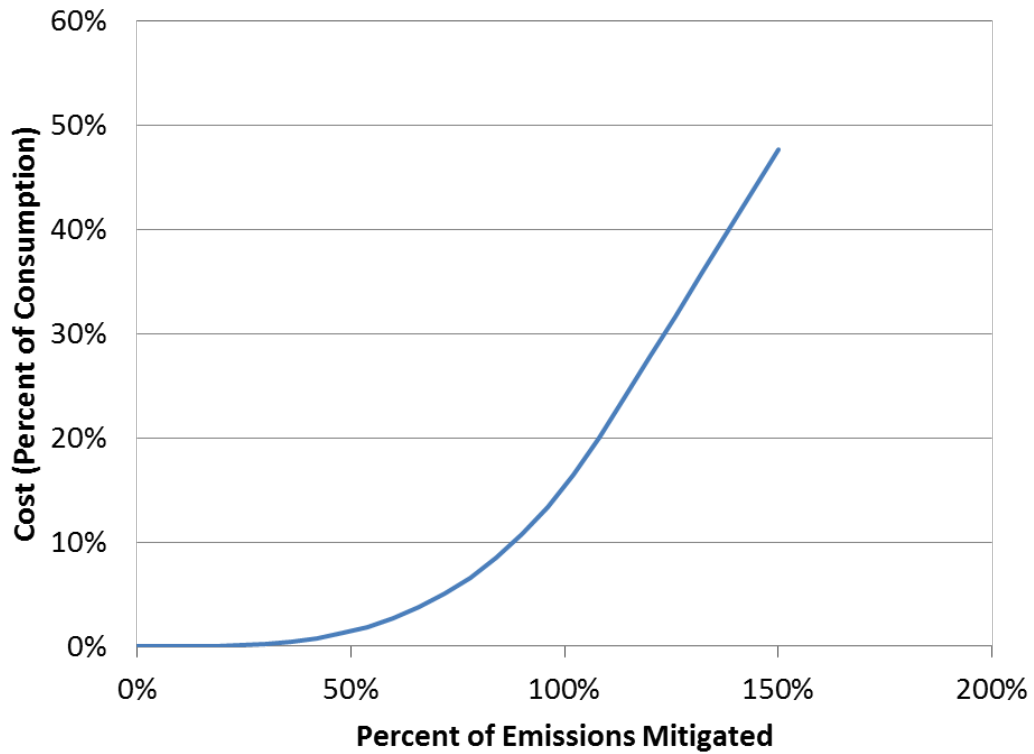
Figure 2—Total cost of abatement as a fraction of consumption, $\kappa(x)$.
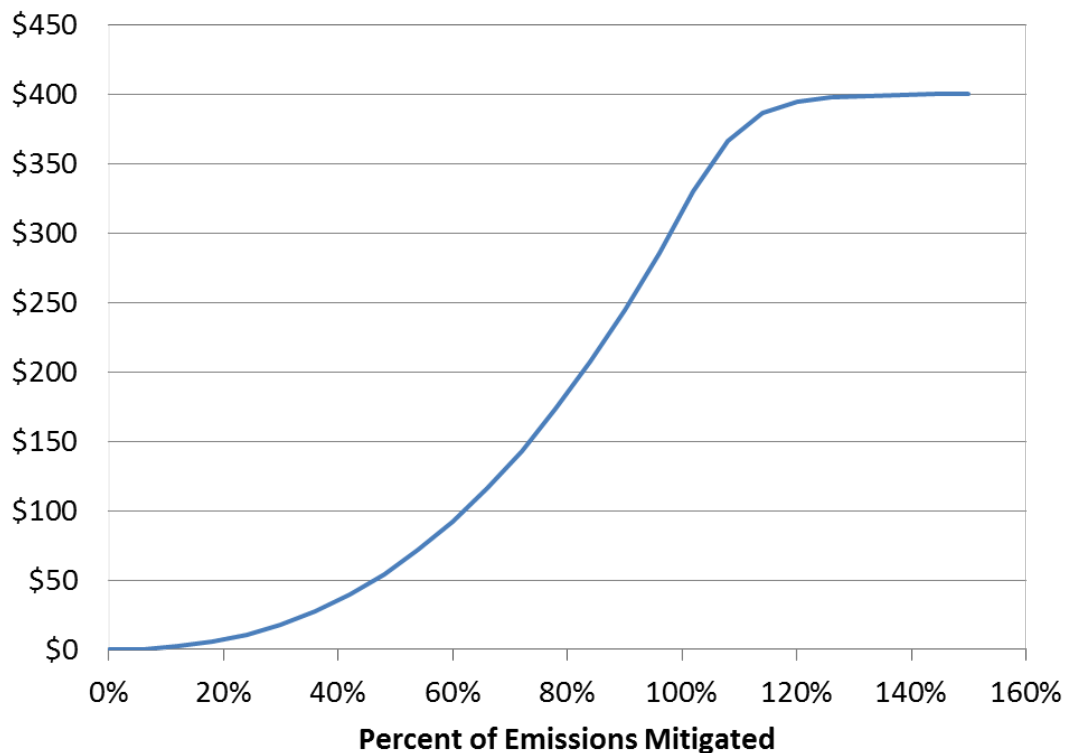
Figure 3—Marginal cost of abatement (in $/ton), $\tau(x)$.

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:

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

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 climate the damage function $D_t(X_t, \theta_t)$. Damages are defined by temperature change, which is defined by greenhouse gas concentrations, which in turn is defined by levels of mitigation. The only way to affect the level of damages, then, is to change mitigation.

We first specify the distribution of temperature change as a function of mitigation. We then add a function that provides a mapping from temperature change to economic damages. The specification of damages has two components: a non-catastrophic component and an additional *tipping-point* element. The hazard rate associated with hitting a tipping point increases with temperature. If a tipping point is hit at any time, additional damages force down 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 lean on Wagner and Weitzman (2015), who calibrate a log-normal distribution for equilibrium climate sensitivity—the eventual temperature rise as atmospheric concentrations of CO_2 double—using 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 IPCC Assessment Report Five (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," IEA World Energy Outlook (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.

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 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

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:

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

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 note that both the distribution of ΔT_{100} and the functional form for the path in equation (25) merit further scrutiny and a more careful calibration to equilibrium temperature outcomes.
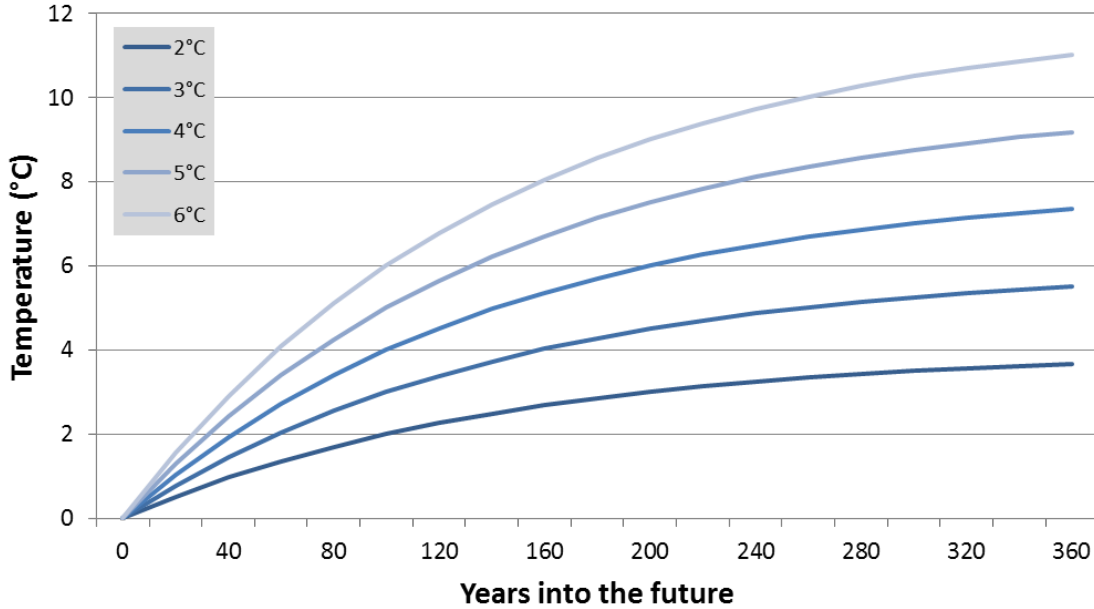
Figure 4—Calibrated time path for temperature increases given assumed temperature increases by 2100.

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 final period.

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

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

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) $\qquad CD_t = \bar{C}_t \cdot L(\Delta T(t)).$

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 damage function with the possibility of *tipping points*, which themselves create the potential for a much larger impact on consumption.

While the possibility of climate tipping points is receiving considerable attention in the scientific community, to our knowledge there is no accepted specification for tipping points. Hence, we employ an *ad-hoc* specification, which can be modified once more progress is made on this issue. In our specification, Prob(TP) denotes the probability of hitting a tipping point over a 30 year period as a function of the global temperature change as of that time ($\Delta T(t)$), and of a parameter, *peakT:*

(28) $\qquad \text{Prob(TP)} = \dfrac{\Delta T(t)}{\max[\Delta T(t), peakT]^2}$

Figure 5 plots Prob(TP) as a function of $\Delta T(t)$ for a set of values of *peakT*. As *peakT* increases, the probability of reaching a climatic tipping point increases as well.

Figure 5—Probability of reaching a climatic tipping point as a function of _peakT_

In our simulations, in each 30-year 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 in a period centered 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:

(29) $\qquad \text{CDTP}_t = \text{CD}_t \cdot e^{-\text{TP\_damage}} = \bar{C}_t \cdot L(\Delta T(t)) \cdot e^{-\text{TP\_damage}}$ 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., $(1 - e^{-\text{TP\_damage}})$) for values of _disaster_tail_ ranging from 2 to 10 is plotted below:
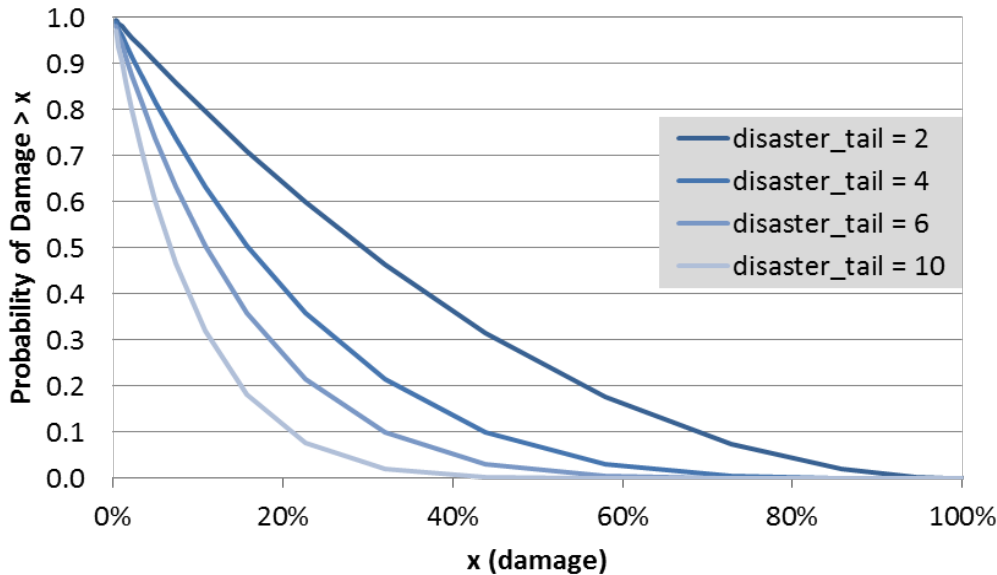
Figure 6—Damage conditional on crossing a tipping point probability of damage > x

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:

$$(30) \qquad D_t = (1 - L(\Delta T(t)) \cdot \left(1 - I_{TP}\left[1 - e^{-TP\_damage}\right]\right),$$

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 $L(\Delta T(t))$, I_{TP}, and e^{-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, 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 last 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 *k*-th 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.

Below we show graphically the average damage functions $D_t(X_t)$ for the 450 and 1000 ppm maximum GHG scenarios, with parameters *peakT* = 8 and disaster_tail = 6.

Given the damage coefficients for these three scenarios, 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.
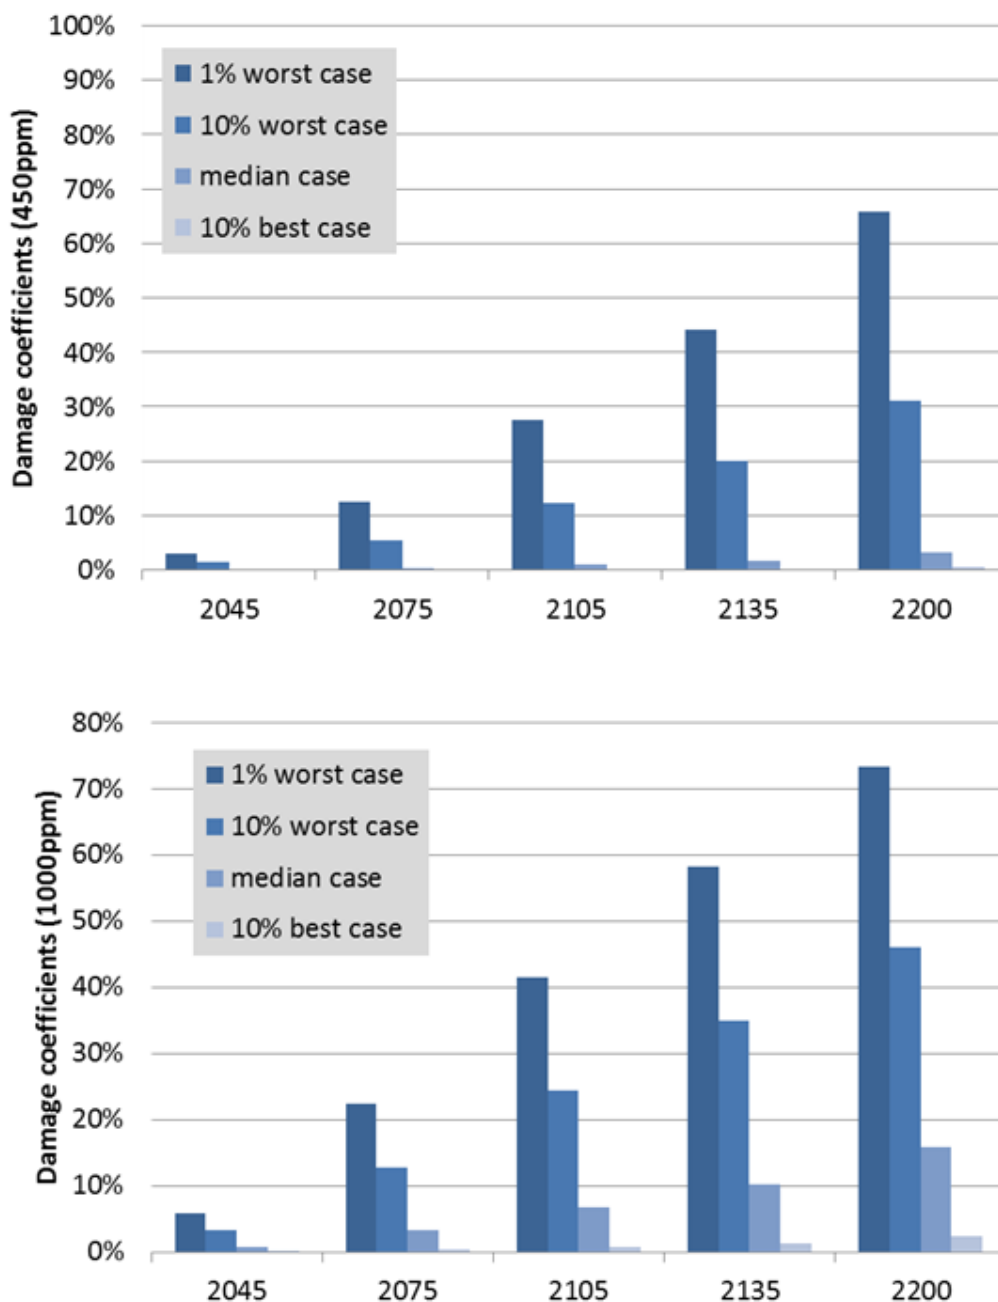
Figure 7—Damage coefficients over time in the 450ppm and 1000ppm scenarios.

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 for a *peakT* of 6 and disaster_tail of 4 to have the following damage coefficient values (Table 4).

Table 4—Damage coefficients associated with *peakT* of 6 and disaster_tail of 4.

GHG-level	Mitigation	1%	10%	50%	90%
450	.9167	.6590	.3098	.0326	.0050
650	.5833	.7132	.4211	.1022	.0150
1000	.0000	.7329	.4598	.1580	.0240

These points determine the interpolated damage functions shown in Figure 8.
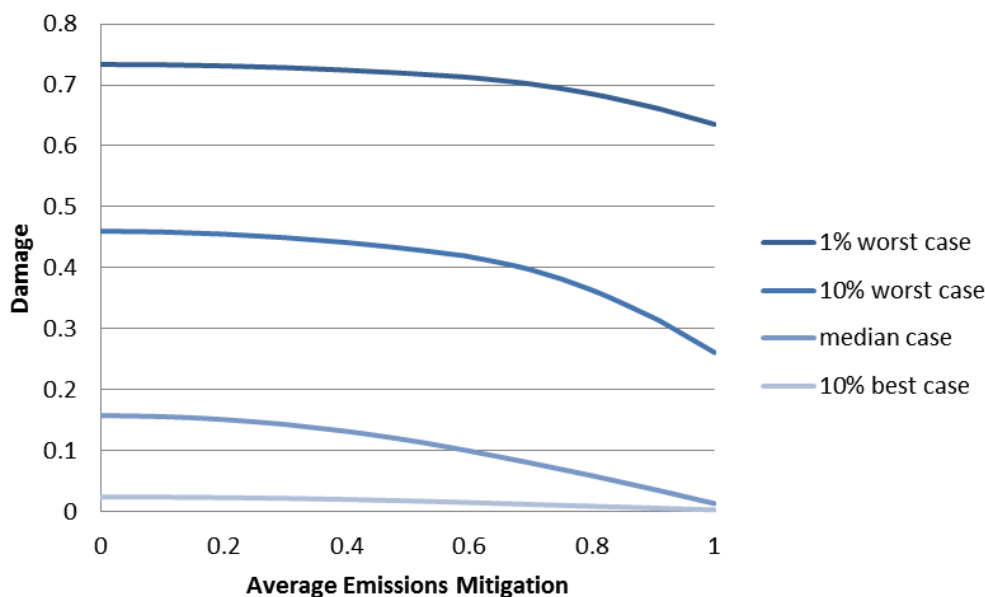
Figure 8—Interpolated final period damage functions.

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) $\qquad D_t(X_{t,}\theta_t) = \sum_{\theta_T} \Pr(\theta_T|\theta_t) \cdot D_t(X_{t,}\theta_T),$

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

3. **Results**

The main model output is the price of carbon in each of five periods. 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).[12] It also shows the implications of using three different climate sensitivity distributions: one mimicking Pindyck (2012), the Roe and Baker (2007) distribution, and the log-normal calibration employed by Wagner and Weitzman (2015).

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. Using a Pindyck (2012) 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 distribution as our base case for the remainder of our runs.

---

[12] 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.

No Risk Aversion & No Disaster

	Roe-Baker	Wagner-Weitzman	Pindyck

6.3   14.0   20.0   25.2   27.2
1.2   2.7   4.0   5.2   6.1

Disaster

86.1
80.7
68.7
76.0
36.3
46.8   43.8
30.1   41.0
13.6
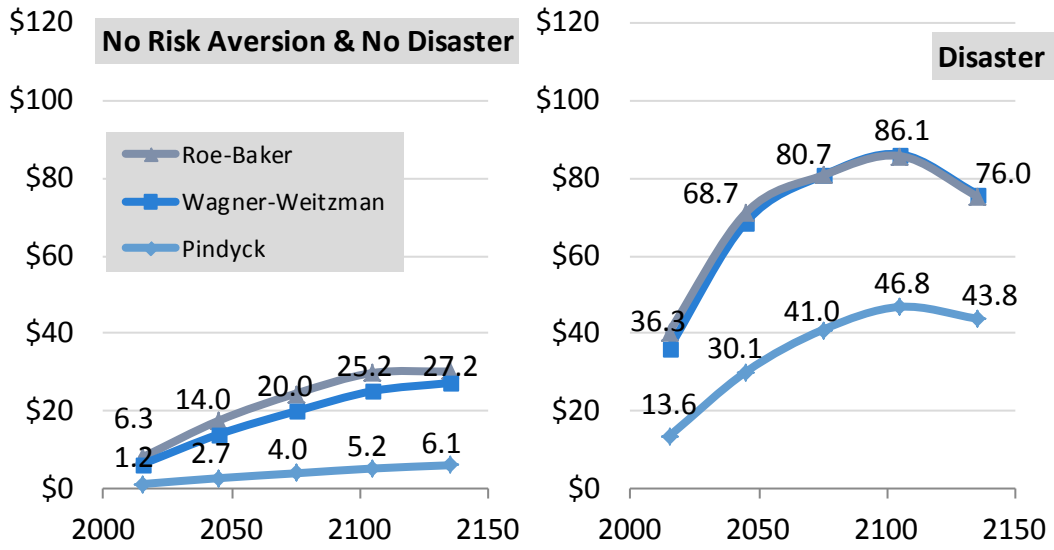
RA	1.56	1.56
peak temp	none	8
disaster_tail	none	6

Risk Aversion

7.9   17.0   23.2   27.6   27.6
1.3   3.0   4.2   5.3   5.5

Base Case

94.2   94.1
84.6   77.5
47.2
55.9   55.4
48.1   45.4
24.3
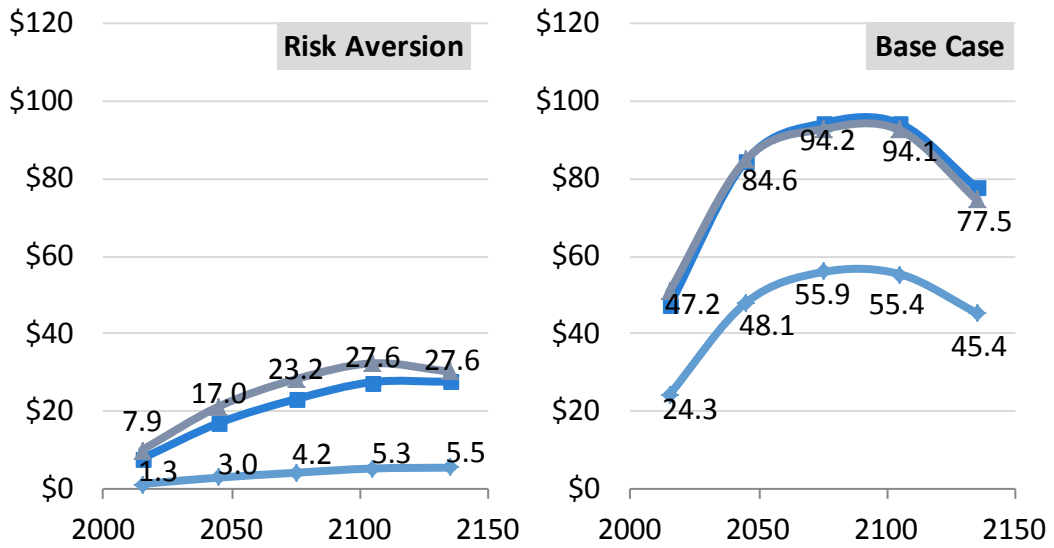
RA	7.00	7.00
peak temp	none	8
disaster_tail	none	6

Figure 9 – Expected price per ton of carbon under four different scenarios and three different assumed climate sensitivity distributions. (Data labels are for Wagner-Weitzman and Pindyck distributions.)

We illustrate the tree structure in the diagram below. In our base case each state has equal probability. The current optimal price of emissions is $47.18/ton for the representative agent. In period 1, in 2045, 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 the agent moves further up in the tree structure. The $84.62 expected price for emissions in period 1 which we show in the lower quadrant of Figure 9 is the average of the price in the up state, $117.91, and the price in the down state, $51.32. Below we show probabilities and emissions prices in each node, and color-code the nodes by the number of up moves. All nodes at a given time with the same number of up moves have the same degree of fragility and thus the same damage for a given amount of greenhouse gas in the atmosphere. Consider period 2, in 2075, which has four equally probable nodes. The fragility in the up-down node is the same as in the down-up node, but because the agent mitigated more emissions in period 1 in the up state the damage is lower and he mitigates slightly less in period 2 (leading to a slightly lower price) when he takes the up-down path than if he takes the down-up path. Thus, there are path dependencies in the tree and we keep track of each path-dependent node separately.

Probabilities

2015	2045	2075	2105	2135
				0.0625
			0.125	0.0625
		0.25		0.0625
				0.0625
	0.50		0.125	0.0625
			0.125	
			0.125	0.0625
				0.0625
1.00		0.25		0.0625
		0.25		0.0625
				0.0625
			0.125	0.0625
	0.50		0.125	
			0.125	0.0625
		0.25		0.0625
				0.0625
			0.125	0.0625
				0.0625

Emissions Price

2015	2045	2075	2105	2135
				$215.73
			$223.91	$169.66
		$174.15		$128.03
				$119.91
	$117.91		$114.51	$113.46
			$112.66	
			$110.55	$ 54.56
				$ 57.31
$ 47.18		$ 80.22		$ 59.78
		$ 83.97		$ 59.25
				$ 61.33
			$ 51.25	$ 62.85
	$ 51.32		$ 56.23	
			$ 60.32	$ 29.17
		$ 38.61		$ 31.80
				$ 34.38
			$ 23.40	$ 36.19
				$ 7.24
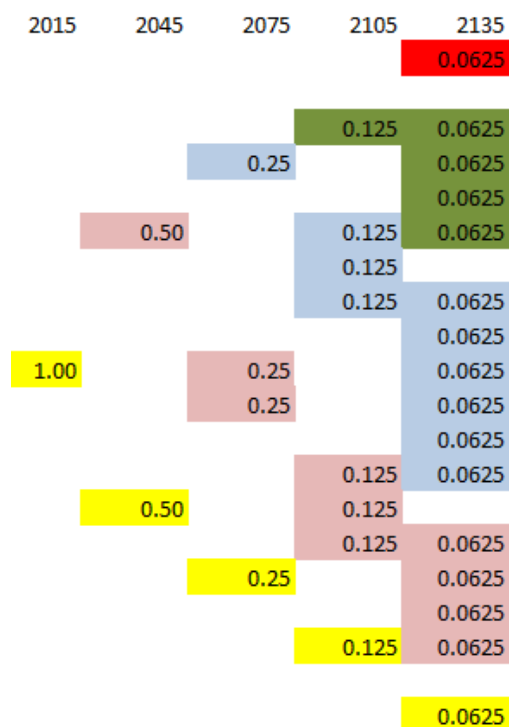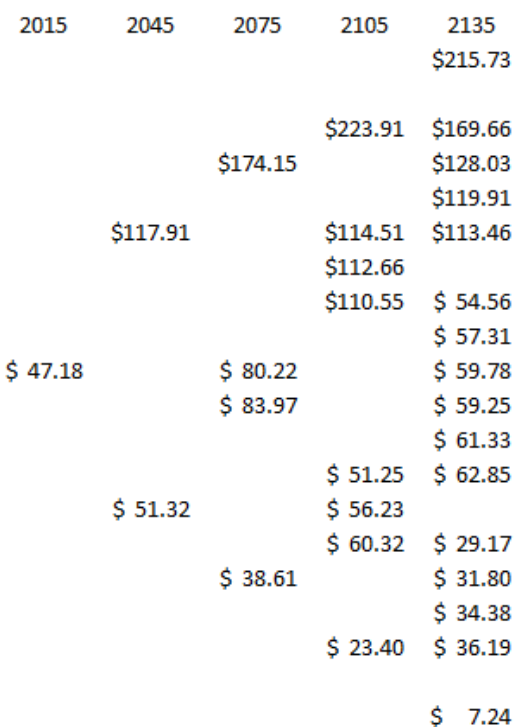
Below we show the amount of mitigation and the damage in each node of the tree. Note that mitigation becomes cheaper over time as technological change occurs. Thus, 46% of emissions are mitigated in the 30-year period from 2015 to 2045 with an emissions price of $47.20. If the agent gets two good draws (down moves in the tree) then in

period 2, the 30-year period from 2075 to 2105, the agent mitigates a larger fraction, 54% of emissions, at a lower price of \$38.61. On the other hand, if the agent gets two bad draws (up moves) then in period 2 he mitigates 101% of emissions at a price of \$174.15. The last mitigation decision is made at the beginning of the final period in 2135, with final damages occurring at the end of the period in 2200.

Mitigation

2015	2045	2075	2105	2135
				1.41
			1.26	1.28
		1.01		1.14
				1.11
	0.75		0.96	1.08
			0.95	
			0.94	0.80
				0.81
0.46		0.73		0.83
		0.74		0.83
				0.84
			0.69	0.85
	0.53		0.71	
			0.73	0.62
		0.54		0.64
				0.66
		0.50		0.67
				0.35

Damages

2015	2045	2075	2105	2135	2200
				23.9%	15.2%
			11.8%	8.3%	5.2%
		5.3%		10.9%	13.0%
				13.1%	15.9%
	1.4%		4.3%	14.7%	18.0%
			6.0%		
			7.1%	2.9%	6.2%
				3.8%	7.1%
0.0%		2.1%		4.8%	7.9%
		3.0%		4.6%	7.8%
				5.7%	8.5%
			1.7%	6.6%	9.2%
	0.6%		2.3%		
			3.1%	1.2%	2.8%
		1.0%		1.5%	3.0%
				1.9%	3.3%
			0.8%	2.4%	3.8%
				0.3%	0.6%

Next we show the consumption (as a multiple of consumption today) and the level of greenhouse gases in the atmosphere at each point in the tree.

Consumption

| 2015	2045	2075	2105	2135	2200
				6.5	33.1
		4.4		8.7	37.0
	2.8			8.8	33.9
				8.6	32.8
	1.7		5.4	8.5	32.0
			5.3		
			5.2	10.2	36.6
				10.1	36.2
1.0		3.1		10.0	35.9
		3.1		10.0	36.0
				9.9	35.7
			5.7	9.8	35.4
1.8			5.7		
			5.6	10.5	37.9
		3.2		10.5	37.8
				10.4	37.7
			5.9	10.4	37.5
				10.7	38.7

Greenhouse gases (in ppm)

| 2015	2045	2075	2105	2135	2200
				433	342
			461	433	371
	461			465	435
				494	471
		436	461	516	498
			489		
			511	465	510
				494	536
400		461		516	555
		484		522	561
				540	577
			489	559	594
	436		511		
			532	522	608
		484		540	621
				559	636
			532	584	657
				584	730

Finally, we show the cost of emissions mitigation as a percentage of consumption and the stochastic discount factor (SDF), that is the ratio of the marginal utilities in the node to the marginal utility in the previous period times the rate of pure time preference.

Mitigation cost

2015	2045	2075	2105	2135
				15.2%
		14.2%		10.8%
	8.8%			7.3%
				6.6%
	4.5%		5.5%	6.1%
			5.4%	
			5.2%	2.2%
				2.3%
1.1%		2.9%		2.5%
		3.1%		2.4%
				2.6%
			1.8%	2.7%
1.4%			2.0%	
			2.2%	0.9%
		1.0%		1.0%
				1.1%
			0.6%	1.2%
				0.1%

Stochastic Discount Factors

| 2015	2045	2075	2105	2135	2200
				27.37%	0.74%
			8.67%	6.13%	1.53%
		5.32%		22.65%	2.26%
				10.02%	3.03%
	3.74%		2.88%	23.63%	2.40%
			6.82%		
			3.71%	9.38%	3.00%
				18.62%	2.96%
		2.58%		12.69%	2.97%
		4.25%		24.35%	2.51%
				9.03%	2.97%
			6.93%	18.79%	2.92%
	2.47%		3.59%		
			5.64%	12.38%	2.95%
		2.87%		18.89%	2.87%
				12.01%	2.92%
			4.17%	16.93%	2.90%
				13.49%	2.91%

The mitigation choices that optimize the utility of the representative agent in the tree structure and determine the social cost of carbon depend on a host of factors. We test them in turn. For one, our analysis so far puts equal weight on each state of the world. If instead we zero in on the tail and explore tail events more carefully, the optimal price goes up (Figure 10).

The higher is $prob\_scale$, the greater is the focus on higher damage tail events in the sense that states are no longer equally spaced. The probability of each state i is given by:

(32) $\qquad Prob(i) = Prob\,(i-1)prob\_scale^{\frac{1}{i}}.$

With $prob\_scale = 1$, each of 16 states has equal chance of 6.25% of occurring. With $prob\_scale = 3$, the first state (with the highest damage) includes outcomes with the worst 0.33% probability, the second state includes the next worst outcomes with a 0.99% probability, etc., increasing the focus on tails.

The focus on states also allows us to explore the value of early information. The top panel of Figure 10 shows results when information is revealed earlier, while the bottom panel shows results after re-ordering states deliberately to model later revelation of information. In this order, for example, the worst case is in the top half of the tree structure, whereas the next worst case is included in the bottom half of the tree. Later revelation of information leads to lower prices and less mitigation early, with higher expected mitigation in the middle-term and lower in the very long term.
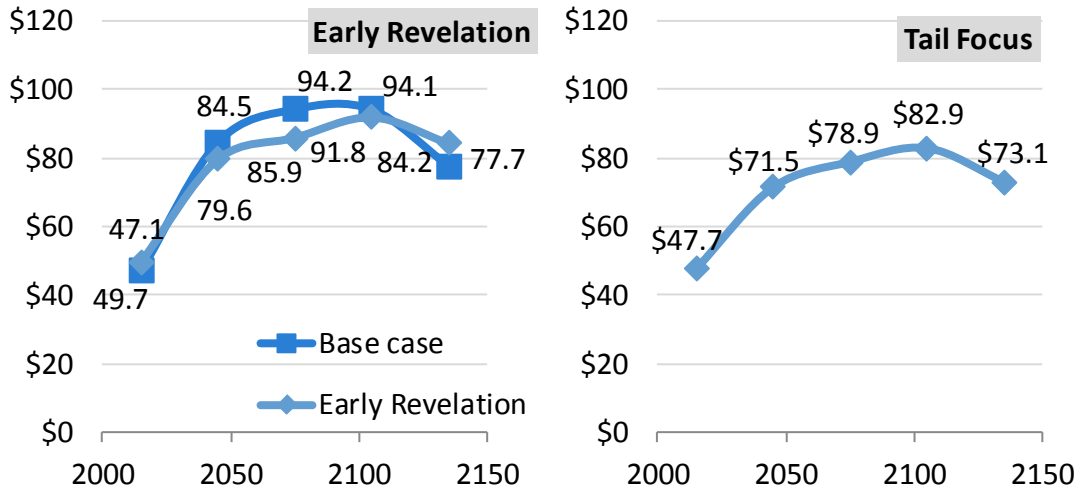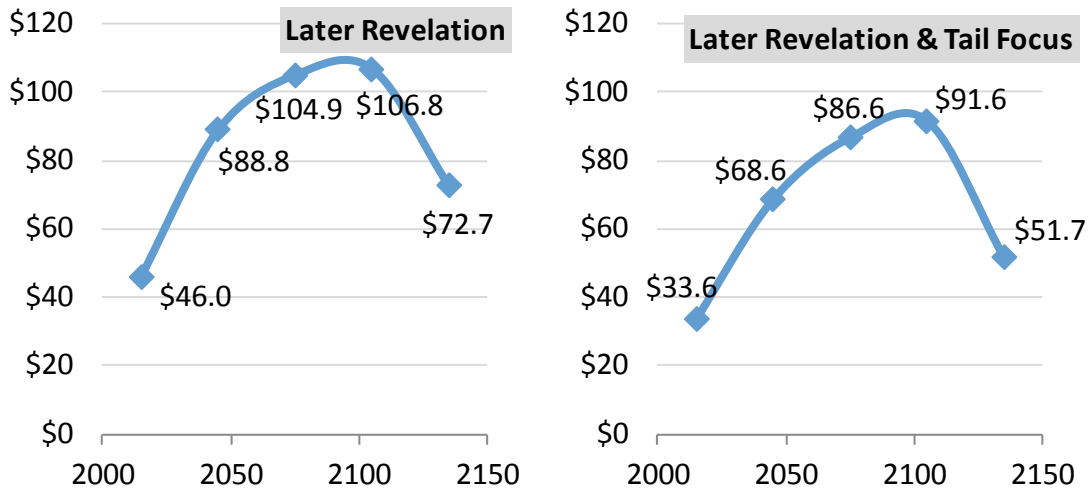
Figure 10—Social cost declines with increased tail focus. Early social cost declines with late revelation of information, but increases in the longer run.[13]

Two key assumptions concern the rate of technical change and the potential for a backstop technology. We analyze both exogenous and endogenous technical change. The former has some surprising results, as the social cost first increases before decreasing

---

[13] Note that in the base case, recombining trees are used in place of early or late revelation. The slight difference between the base case in Figure 10 (top left graph) and Figure 9 (lower right graph) is due to sampling differences from the Monte Carlo simulation.

again (Figure 11). As technical improvements change from 0% to 1, the social cost rises, only to fall again with annual technical improvements at 1.5% or higher.
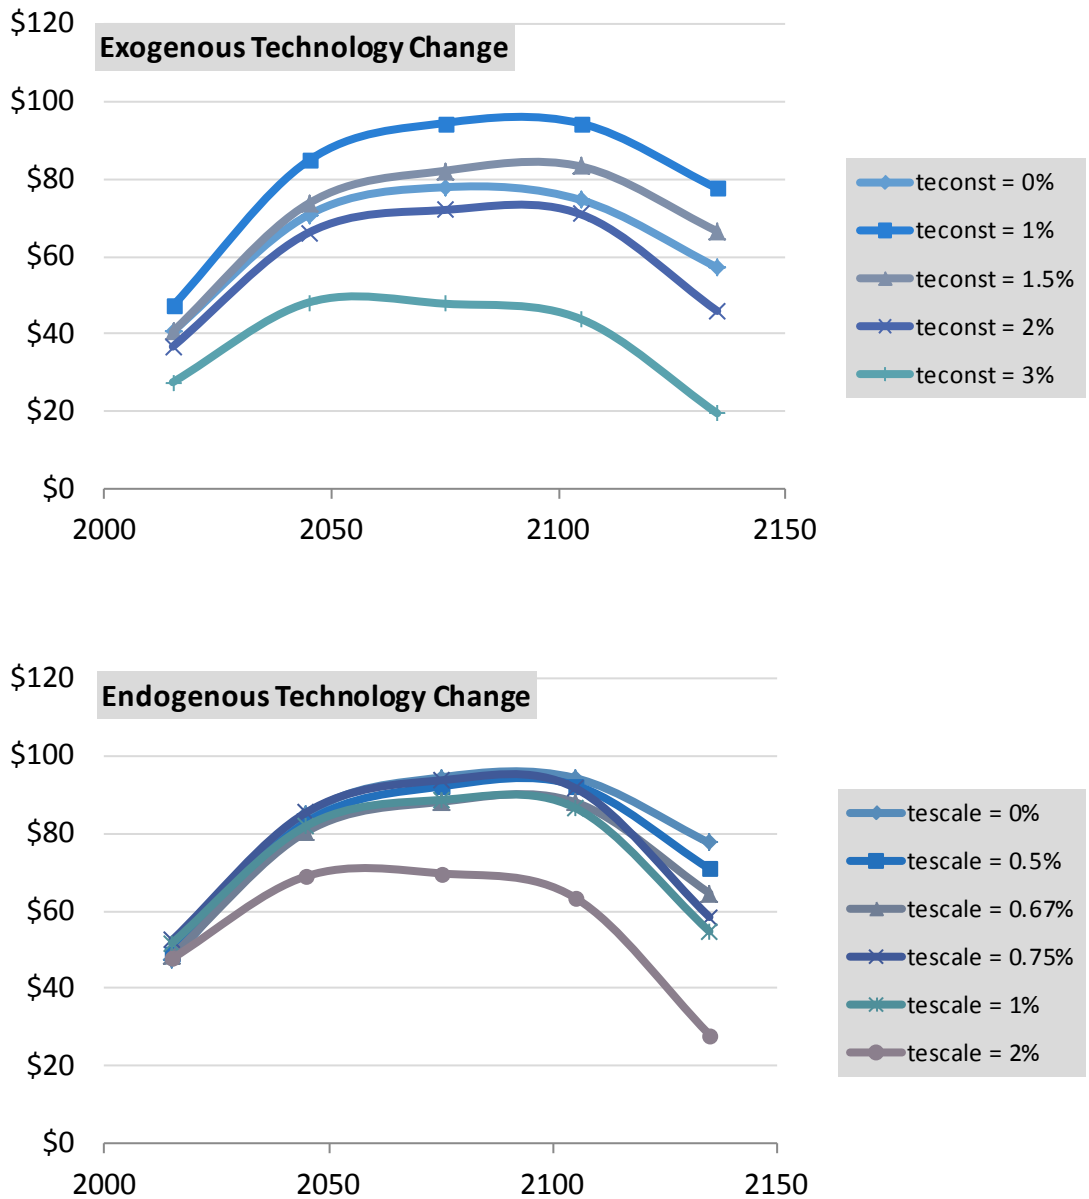
Figure 11—Social cost rises and then falls with increased exogenous technological change. Increased endogenous technical change decreases the social cost.

We now explore the impact of the backstop technology. Including a backstop technology at a *maxPrice* of 350 decreases the social cost of carbon in both the exogenous and endogenous technology change scenarios (Figure 12).

Base Case

$94.4 $94.3

$84.8

$77.5

$47.2

tescale	0
teconst	1
join	2000
maxPrice	2500

Exogenous vs endogenous tech change

Endogenous

Exogenous

|		
tescale	0	0.75
teconst	1.5	1
join	2000	2000
maxPrice	2500	2500
	(exo)	(endo)

Exogonenous change with and without backstop technology

No backstop

Backstop

tescale	0	0
teconst	1.5	1.5
join	2000	300
maxPrice	2500	350
	(no b.s.)	(b.s.)

Endogenous change with and without backstop technology

No backstop

Backstop

|		
tescale	0.75	0.75
teconst	1	1
join	2000	300.0
maxPrice	2500	350.0
	(no b.s.)	(b.s.)

Figure 12—Air capture backstop technology lowers the social cost.

Next we investigate the effect of growth and interest rates, with corresponding changes in the elasticity of intertemporal substitution. Figure 13 shows social cost values with increasing growth rates around the central case of 2.0%.
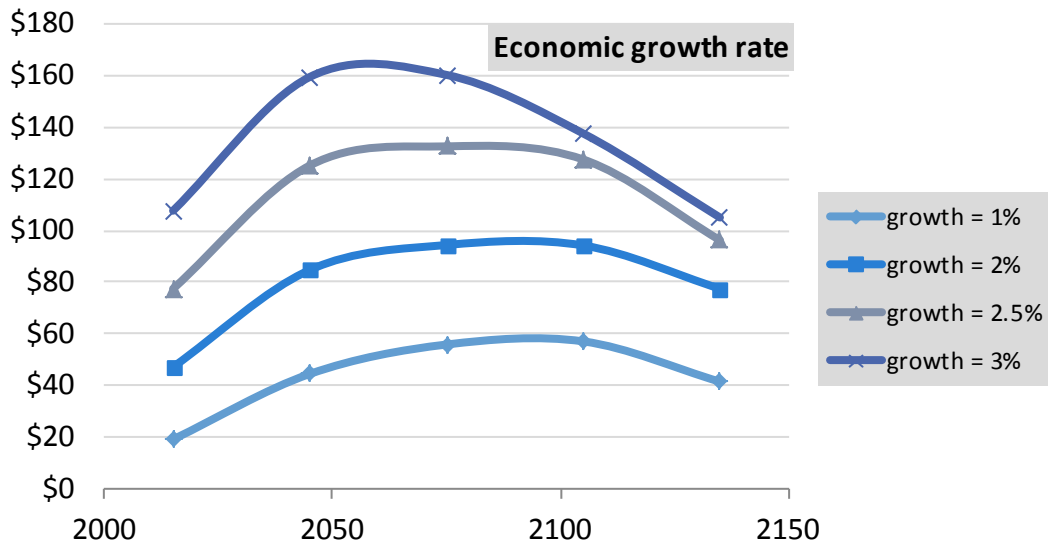
growth		0.01	0.02	0.025	0.03
eis		0.304	0.643	0.8106	0.976
interest rate		2.5	2.5	2.5	2.5
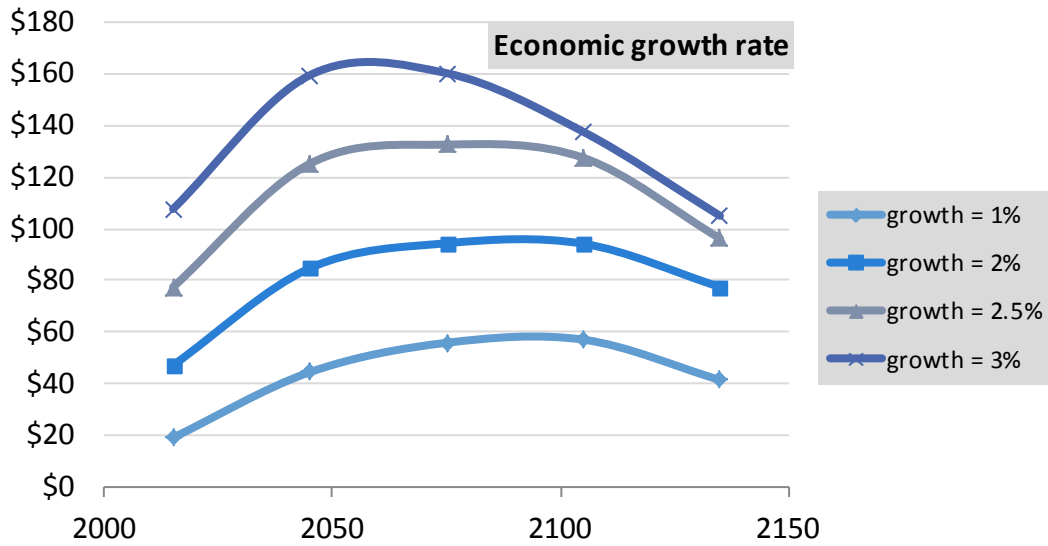
growth		0.01	0.02	0.025	0.03
eis		0.304	0.643	0.8106	0.976
interest rate		2.5	2.5	2.5	2.5

Figure 13—Higher growth rates lead to higher social cost.

Figure 14 shows results for different rates of interest around the central 2.5% case. High growth rates lead to higher social cost values, whereas higher interest rates have the opposite effect.

| eis | | 1.28 | 0.773 | 0.643 | 0.552 | 0.42 |
| interest rate | | 1.0 | 2.0 | 2.5 | 3.0 | 4.0 |

Figure 14—High interest rates lead to lower social cost.

One caveat in all simulations is their inherent probabilistic outcome. All results shown so far are based on Monte Carlo simulations with 4 million draws. To test the accuracy of the model, we ran the 'Risk Aversion & Disaster' case with 4 million draws twenty times. The $47.18 number presented in Figure 10, which comes from a particular run, happens to equal the result for prices in 2015 averaged over all twenty model runs, but there is a standard deviation of $.04 for this value across the different runs. The average and standard deviations for all periods are shown in Table 5.

Table 5—Means and standard deviations for twenty runs of the 'Risk Aversion & Disaster' base case with 4,000,000 each.

4,000,000 draws		
	Mean	*Std. Dev.*
2015	47.18	0.04
2045	84.66	0.06
2075	94.29	0.08
2105	94.16	0.10
2135	77.58	0.07

Figure 15 shows the implications of changing the pure rate of time preference. The upper graph 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 social cost of carbon, 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 changes the rate of time preference and holds the EIS fixed: in this case the higher the rate of time preference, the higher the interest rate and the lower the social cost.
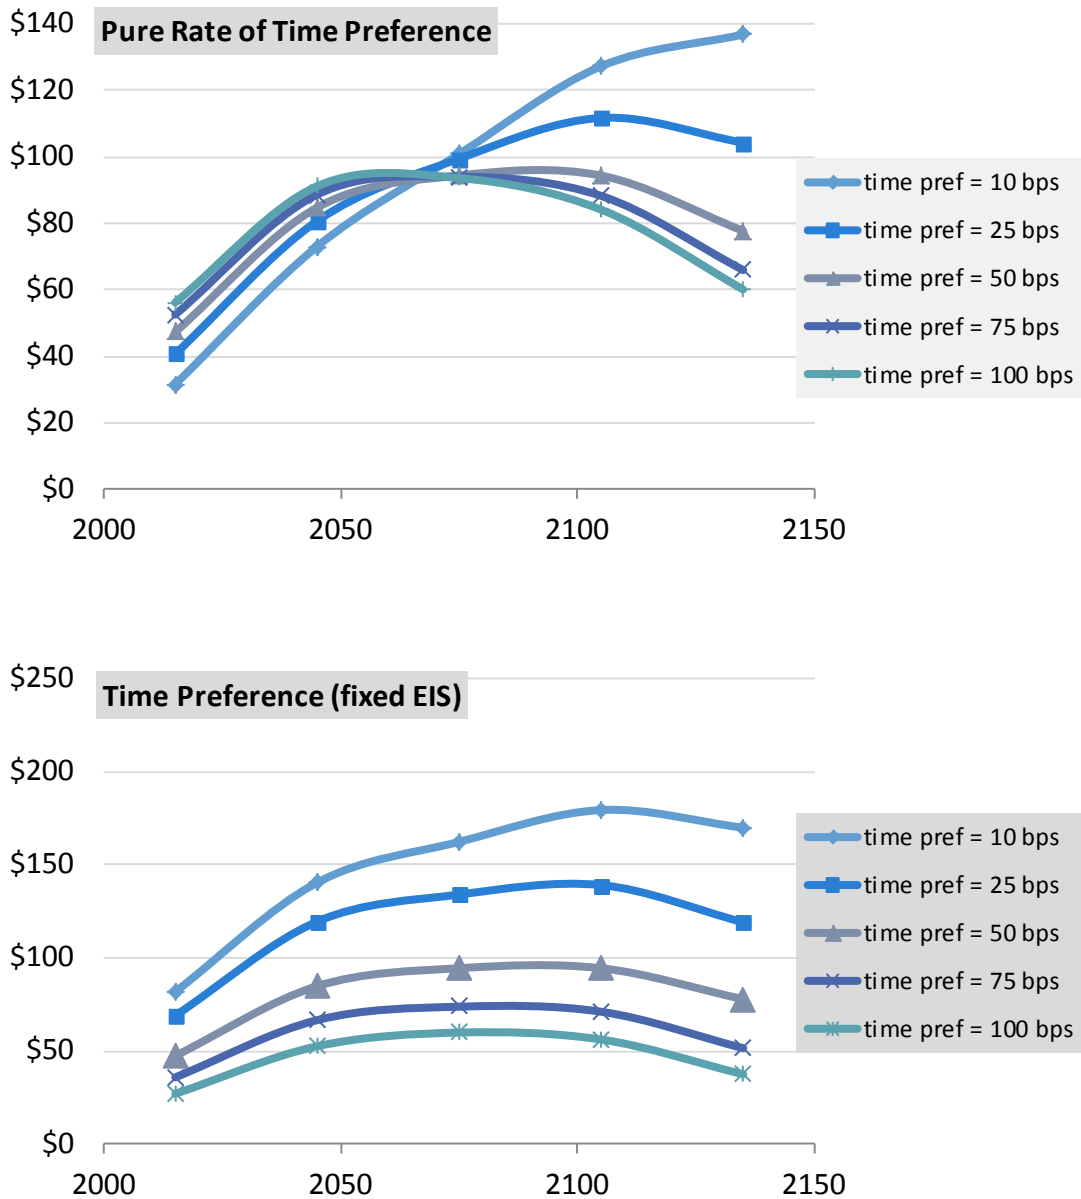
Figure 15 – As time preference increases and the EIS is held constant, the social cost decreases.

4. Optimization

We optimize the utility function of the representative agent over choices of emissions mitigation in each of the 31 nodes in the tree. We use the optimization function, "fmin_l_bfgs_b" in the python library scipy, a quasi-Newton algorithm which makes use of analytic derivatives. In Figure 1 there is a discontinuity in the price of carbon as risk aversion is increased. The discontinuity occurs because as risk aversion is increased the strategy that maximizes the utility of the representative agent at lower risk aversion levels reaches a point where it equals the utility of a different strategy, one that requires a higher degree of emissions mitigation today. The latter strategy has higher mitigation and lower consumption in early periods, but the higher emissions reductions allow for higher consumption in high damage scenarios at later dates.

Below we show the consumption in each node for the two strategies which have equal utility at the point where the price jumps up. The right hand side represents the lower bound of a strategy which is globally optimal for higher rates of risk aversion. The left hand side is the upper bound of a strategy with lower initial emissions mitigation which is globally optimal for lower rates of risk aversion.

2015	2045	2075	2105	2135	2200	2015	2045	2075	2105	2135	2200
				6.72	33.46					7.76	34.90
				8.86	37.15					9.55	37.80
			4.42	8.88	34.39				4.08	8.76	37.03
		2.80		8.70	33.22			2.69		8.75	33.53
			5.35	8.58	32.37				5.34	8.59	32.44
	1.69		5.26				1.67		5.28		
			5.20	10.25	36.69				5.21	10.32	36.93
				10.15	36.37					10.18	36.46
		3.11		10.02	36.01			3.12		10.05	36.09
0.99	3.07			10.05	36.10	0.98	3.08			10.06	36.13
				9.92	35.77					9.93	35.79
			5.75	9.81	35.52				5.76	9.82	35.55
			5.70						5.71		
	1.77		5.63	10.56	37.95		1.77		5.64	10.57	37.98
				10.52	37.85					10.52	37.86
		3.21		10.46	37.74			3.21		10.47	37.75
			5.86	10.39	37.60				5.86	10.40	37.61
				10.72	38.76					10.72	38.76

5. 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.[14]

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 greater is the need for climate action today.

Bibliography

Ackerman, Frank, Elizabeth A. Stanton, and Ramón Bueno. 2012. "Epstein-Zin utility in DICE: Is risk aversion irrelevant to climate policy?" E3 working paper.

Arrow, K., M. Cropper, C. Gollier, B. Groom, G. Heal, R. Newell, W. Nordhaus, Robert S. Pindyck, William A. Pizer, Paul R. Portney, Thomas Sterner, Richard S. J. Tol, and Martin L. Weitzman, 2013."Determining Benefits and Costs for Future Generations." *Science* 341, no. 6144: 349-350.

Arrow, Kenneth J., Maureen L. Cropper, Christian Gollier, Ben Groom, Geoffrey M. Heal, Richard G. Newell, William D. Nordhaus, Robert S. Pindyck, William A. Pizer, Paul R. Portney, Thomas Sterner, Richard S. J. Tol, and Martin L. Weitzman. 2014. "Should Governments Use a Declining Discount Rate in Project Analysis?" *Rev Environ Econ Policy* 8: 145-163.

Bansal, Ravi, and Marcelo Ochoa. 2009. "Temperature, growth and asset prices." *Duke University Working Paper.*

Bansal, Ravi, and Marcelo Ochoa. 2010. "Temperature, aggregate risk, and expected returns." *Duke University Working Paper.*

Bansal, Ravi, and Amir Yaron. 2004. "Risks for the long run: A potential resolution of asset pricing puzzles." *The Journal of Finance,* 59, 1481–1509.

Coase, R. H. 1960. "The problem of social cost." *Journal of Law & Economics*, *3*, 1-69.

Crost, Benjamin and Christian P. Traeger. 2014. "Optimal CO2 mitigation under damage risk valuation." *Nature Climate Change,* 4, 631–636.

Duffie, Darrel. 2001. *Dynamic Asset Pricing Theory* (3rd ed.). Princeton University Press.

Epstein, Larry, and Stanley Zin. 1989. "Substitution, risk aversion, and the temporal behavior of consumption growth and asset returns I: A theoretical framework." *Econometrica,* 57, 937–969.

Epstein, Larry, and Stanley Zin. 1991. "Substitution, risk aversion, and the temporal behavior of consumption growth and asset returns II: An empirical analysis." *Journal of Political Economy,* 99, 263–286.

Gollier, Christian. 2012. *Pricing the Planet's Future: The Economics of Discounting in an Uncertain World.* Princeton University Press.

Gollier, Christian, and Martin L. Weitzman. 2010. "How Should the Distant Future Be Discounted When Discount Rates Are Uncertain?" *Economics Letters* 107.3: 350–53.

Goulder, Lawrence H. 1995. "Environmental taxation and the double dividend: a reader's guide." *International Tax and Public Finance* 2 157–183.

---

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

Hansen, L. P., John C. Heaton, and Nan Li. 2008. "Consumption strikes back? Measuring long-run risk." *Journal of Political Economy,* 116, 260–302.

Hansen, L. P., J. Heaton, J. Lee, and N. Roussanov. 2007. Intertemporal substitution and risk aversion. *Handbook of econometrics*, *6*, 3967-4056.

Hansen, L. P., and K. J. Singleton. 1982. "Generalized instrumental variables estimation of nonlinear rational expectations models." *Econometrica*, 50, 1269-1286.

Hansen, L. P., and Ravi Jagannathan. 1991. "Implications of Security Market Data for Models of Dynamic Economies." The Journal of Political Economy, 99(2), 225-262.

Hansen, L. P., and Scott F. Richard. 1987. "The role of conditioning information in deducing testable restrictions implied by dynamic asset pricing models." *Econometrica,* 55, 587–613.

"World Energy Outlook 2013." *International Energy Agency* (2013). http://www.worldenergyoutlook.org/publications/weo-2013/.Kreps, David M., and Evan L. Porteus. 1978. "Temporal resolution of uncertainty and dynamic choice theory." *Econometrica,* 46, 185-200.

Litterman, Robert B., José Scheinkman, and Laurence Weiss. 1991. "Volatility and the yield curve." *Journal of Fixed Income,* 1(1), 49-53.

McKinsey. 2009. "Pathways to a low-carbon economy: Version 2 of the global greenhouse gas abatement cost curve." *Discussion paper McKinsey and Company*. Available at http://www.mckinsey.com/clientservice/ccsi/pathways low carbon economy.asp.

Mehra, Rajnish, and Edward C. Prescott. (1985). "The equity premium: A puzzle." *Journal of Monetary Economics 15*(2), 145-161.

Nordhaus, William D. 2013. *The Climate Casino: Risk, Uncertainty, and Economics for a Warming World*, Yale University Press.

Nordhaus, William. 2014. "Estimates of the social cost of carbon: concepts and results from the DICE-2013R model and alternative approaches." *Journal of the Association of Environmental and Resource Economists* 1, no. 1.

Pindyck, Robert S. 2012. "Uncertain outcomes and climate change policy." *Journal of Environmental Economics and Management,* 63(3), 289-303.

Stern, Nicholas. 2007. Stern Review Report on the Economics of Climate Change. Cambridge and New York: Cambridge University Press.

Summers, Lawrence, and Richard Zeckhauser. 2008. "Policymaking for posterity." *Journal of Risk and Uncertainty* 37, 115–140.

Shiller, Robert. 2000. *Irrational Exuberance*. Princeton University Press.

Ulph, A., & Ulph, D. 1994. The optimal time path of a carbon tax. *Oxford Economic Papers*, 857-868.

Wagner, Gernot and Martin L. Weitzman, 2015. *Climate Shock: the Economics Consequences of a Hotter Planet*. Princeton University Press.

Weil, Philippe. 1989. "The equity premium puzzle and the risk-free rate puzzle." *Journal of Monetary Economics*, 24(3), 401–421.

Weitzman, Martin L. 2001. "Gamma discounting." *American Economic Review*: 260-271.

Weitzman Martin L. 2007. "Subjective Expectations and Asset-Return Puzzles." *American Economic Review*, 97(4), 1102-30.