Learning and Climate Feedbacks: Optimal Climate Insurance and Fat Tails*

David L. Kelly
Department of Economics
University of Miami
dkelly@miami.edu

Zhuo Tan
Department of Economics
University of Miami
z.tan@umiami.edu

April 1, 2013

Abstract

We study the effect of potentially severe climate change on optimal climate change policy, accounting for learning and uncertainty in the climate system. In particular, we test how fat upper tailed uncertainty over the temperature change from a doubling of greenhouse gases (the climate sensitivity), affects economic growth and emissions policy. In addition, we examine whether and how fast uncertainties could be diminished through Bayesian learning. Our results indicate that while overall learning is slow, the mass of the fat tail diminishes quickly, since observations near the mean provide evidence against fat tails. We denote as "partial learning" the case where the planner rejects high values of the climate sensitivity with high confidence, even though significant uncertainty remains. Fat tailed uncertainty without learning reduces current emissions by 38% relative to certainty, indicating significant climate insurance, or paying to limit emissions today to reduce the risk of very high temperature changes, is optimal. However, learning reduces climate insurance by about 50%. The optimal abatement policy is strongly influenced by the current state of knowledge, even though greenhouse gas (GHG) emissions are difficult to reverse. Non-fat tailed uncertainty is largely irrelevant for optimal emissions policy.

JEL classification: Q54; Q58; D83; D81; H43

*We would like to thank Derek Lemoine, Alex Marten, and seminar participants at the Association for Environmental and Resource Economists Second Annual Summer Conference, the 13th Occasional Workshop on Environmental and Resource Economics, University of California at Santa Barbara, 2011 Camp Resources XVIII, the University of Miami, Notre Dame University, the University of Stavanger, and Washington University in St. Louis for helpful comments and suggestions. An earlier version of this paper circulated with the title "Learning, Growth and Climate Feedbacks: Does the Fat Tail Wag the Dog?"

1 Introduction

Uncertainty is a dominant feature of climate change. Recent research highlights a particular aspect of climate change uncertainty: there exists a relatively small chance of severe climate change. In particular, a doubling of greenhouse gases (GHGs) above preindustrial levels may cause a very large steady state increase in temperature.[1] Uncertainty creates an insurance motive for reducing emissions, in that paying to limit GHG emissions today prevents GHG concentrations from rising, which reduces the probability of very high temperature changes. Further, the prior distribution of the temperature change caused by a doubling of GHGs is known to have a fat upper tail, meaning the upper tail of the distribution of temperature changes declines at a rate slower than exponential. The existence of a fat tail significantly increases the insurance value of current GHG abatement, since households are willing to pay more up front abatement costs to eliminate fat tailed risk of severe climate change.

However, climate change uncertainty differs from a standard insurance problem in that learning reduces climate change uncertainty over time. If learning resolves climate uncertainty relatively quickly, then the initial insurance premium might be small, as the planner would still have time to increase abatement if learning quickly indicated the climate sensitivity was large. However, if learning is slow to resolve climate uncertainty, then the optimal policy calls for aggressive initial abatement for insurance purposes. The central question for climate policy is then: how fast will learning resolve uncertainty about potential large steady state temperature changes caused by GHG emissions, and what is the optimal climate policy with fat tailed climate uncertainty and learning?

The prior literature finds that learning is a slow process. Kelly and Kolstad (1999b) consider uncertainty regarding the heat capacity of the ocean. In their integrated assessment model, stochastic weather shocks obscure the climate change signal in the temperature data, which slows Bayesian learning. This result has since been confirmed in models with other types of climate uncertainty and different distributional assumptions. In particular, Leach (2007) considers uncertainty over the climate sensitivity, which is the steady state temperature change per unit of radiative forcing,[2] a measure of the elasticity of the climate

[1]The Intergovernmental Panel on Climate Change (2007) reviews many studies and finds values higher than 4.5°C cannot be ruled out, although the best estimate is closer to 3°C. Weitzman (2009) averages 22 studies and finds a 5% chance that a doubling of GHGs will cause temperatures to rise more than 7°C. Other papers which estimate the current scientific uncertainty regarding the climate sensitivity include: Lemoine (2010), Newbold and Daigneault (2009), Roe and Baker (2007), Schwartz (2007), and Baker and Roe (2009).

[2]The more familiar $\Delta T_{2\times}$, or the steady state temperature change from a sustained doubling of GHG

to changes in GHG concentrations, and finds Bayesian learning about the climate sensitivity is extremely slow.[3] Roe and Baker (2007) argue that resolving uncertainty regarding the climate sensitivity is difficult, because small uncertainties in climate feedbacks magnify the uncertainty about the climate sensitivity.[4] Keller, Bolker, and Bradford (2004) show that slow learning about the climate sensitivity combined with an uncertain climate threshold, implies significant near term abatement is optimal, to avoid accidently exceeding the threshold. Lemoine and Traeger (2013) study alternative uncertain thresholds with learning and find somewhat smaller effects on near term abatement.

However, it is possible that the planner learns enough to reject severe climate change with a high degree of confidence quickly, even though the climate sensitivity is difficult to pin down precisely. We define this case as "partial learning." To investigate partial learning, we develop a quantitative integrated assessment model of the world economy.[5] In the model, the planner faces stochastic weather shocks and uncertainty over the first order autoregressive coefficient in the equation governing evolution of temperature, which is the climate feedback parameter.

Because the climate feedback parameter is uncertain, the climate sensitivity is also uncertain. Further, if the climate feedback parameter is close to one, then GHG "shocks" to temperature are long lived, and therefore an increase in GHG emissions causes very high steady state temperature changes. Hence, although uncertainty in the feedback parameter is normally distributed, uncertainty in the climate sensitivity has a fat upper tail (see for example, Roe and Baker 2007). The social planner learns the feedback parameter, and therefore the climate sensitivity, by updating prior beliefs given stochastic temperature data using Bayes rule.

We define the lower bound of the fat tail as when a doubling of GHGs implies steady state temperatures increase by 1.5°C more than the mean of the current prior distribution. For example, we calibrate the mean of the current prior equal to 2.76°C, so the tail of the distribution equals values greater than or equal to 4.26°C.[6] When the planner can reject the

concentrations, is proportional to the climate sensitivity.

[3]Both of these papers and our paper consider observational learning in the sense that the planner learns from the data on temperature and GHG concentrations. An alternative is to allow learning where the planner pays for R&D. Nonetheless, to fully resolve uncertainty, all R&D must eventually be confirmed in the data.

[4]Climate feedbacks are changes in the climate system brought on by higher temperatures which amplify or diminish the relationship between GHGs and temperature (climate forcing). For example, higher temperatures melt ice, which in turn implies less heat is radiated back into space, which amplifies climate forcing. The magnitude of many climate feedbacks are uncertain (Forest, Stone, and Sokolov 2005).

[5]An integrated assessment model is broadly defined as a model which combines scientific and socioeconomic aspects of climate change to assess policy options for climate control (Kelly and Kolstad 1999a).

[6]No generally agreed upon value for what constitutes the tail of the prior distribution exists. Nonetheless,

hypothesis that the climate sensitivity implies a steady state temperature increase greater than or equal to the critical temperature at the 1% or 0.1% level, we say that partial learning is complete.[7] Such learning is partial in the sense that significant uncertainty typically remains even after a high climate sensitivity is rejected.

Our results show that the social planner rejects that the climate sensitivity is in the upper tail of the prior distribution very quickly. That is, although we confirm results in the previous literature that learning the actual true value *precisely* is a relatively slow process, the planner is able to reject values of the climate sensitivity in the upper tail of the prior distribution quickly. In fact, the planner can reject very high values of the climate sensitivity (e.g. 1.5°C or more above the mean estimate) at 99% confidence interval in less than a decade, if the true climate sensitivity is moderate. First, observations near the moderate true value provide evidence against the tails of the distribution. In addition, the density of even a fat tail is not large, so Bayes rule requires relatively few observations to reduce the mass of the fat tail below the critical confidence level. This result is surprising given the common intuition in the literature that reducing uncertainty in the tail of the climate sensitivity prior distribution must be a slow process since climate disasters are rare (see for example, Weitzman 2009, page 12).

If the true climate sensitivity turns out to be relatively high, learning progresses more slowly. First, Bayes rule requires more observations to move the mean estimate from the prior of a moderate climate sensitivity to the true high value. Second, Bayes rule requires more observations to resolve the difference between a climate sensitivity that is relatively high and a climate sensitivity which is very high. Nonetheless, because a high climate sensitivity is relatively unlikely according to the prior, the possibility that learning is slower due to a high climate sensitivity receives relatively little weight when computing the expected time until partial learning is complete.

For example, we find that the expected learning time when the lower bound of the fat tail is 4.26°C is only about 8 years at the 0.1% level. However, if the lower bound is 10°C, partial learning is not complete for almost 50 years. Integrating over the prior distribution, we find that the expected time, conditional on prior information, until partial learning is complete is only about 17 years at the 0.1% level.

Like Weitzman (2009), our model considers a high climate sensitivity as a possible scenario with high damages from climate change. Other potentially high damage scenarios exist,

much of the literature uses higher values (e.g. Weitzman, 2009 discusses values above 7°C). A larger lower bound would only strengthen our results.

[7] See Kelly and Kolstad (1999b) for a justification for using hypothesis tests to measure learning.

including high sea level rise (Nicholls, Tol, and Vafeidis 2008), thermohaline circulation collapse (Keller, Bolker, and Bradford 2004), and a reduction in the decay rate of carbon into the ocean (Lemoine and Traeger 2013).[8] Learning in each of these contexts may differ from our results. For example, if variation in the size of ice sheets unrelated to temperature was significant, learning about sea level rise would proceed more slowly. Nonetheless, fat tailed uncertainty in the climate sensitivity is most commonly analyzed high damage scenario in the literature, presumably because the fat tail is clearly evident in the scientific priors.

In terms of optimal policy, we quantify the effect of uncertainty on near term emissions and abatement policy. With uncertainty but without learning, in the initial period emissions are about 38% lower, and the carbon tax is $22.94 higher, than under certainty. The planner insures by reducing emissions, paying for more abatement to reduce the probability of high damages that occur if the climate sensitivity is high. However, in the current period, emissions with uncertainty and learning are only about 19% lower than under certainty. The optimal carbon tax with uncertainty and learning is only $8.84 per ton higher than under certainty. Therefore, learning reduces emissions abatement for insurance purposes by about 50%. Further, optimal emissions with uncertainty and learning converge quickly to emissions given perfect information, typically in about 16 years. Uncertainties remain after 16 years, but the remaining uncertainty is not relevant for the optimal emissions policy. The fat tail drives policy, and learning shrinks the mass of the fat tail quickly.[9]

Fat tailed uncertainty arises naturally when multiple uncertainties exist,[10] or, in the approach taken here, when an uncertain parameter has a multiplicative effect through feedbacks in the climate system. Climate sensitivity uncertainty is then fat tailed, even if the prior distribution for the uncertain parameter is normal (Roe and Baker 2007). In turn, Weitzman (2009) shows that, given a coefficient of risk aversion greater than or equal to one, the risk premium required to accept an uncertain climate sensitivity with fat tails is infinite (the dismal theorem).[11]

However, Costello, Neubert, Polasky, and Solow (2010) shows that the dismal theorem is

[8]Tol (2009) reviews the findings of the literature on various damages caused by climate change.

[9]Optimal policy under learning converges quickly to the perfect information case even if the true climate sensitivity is very high. This is because with a higher mean estimate, deviations from certainty are driven by the mass of the tail of the new distribution with a higher mean, and the mass of the tail of the new distribution still shrinks quickly.

[10]For example, if the prior distribution for the climate sensitivity conditional on the variance of the weather shocks is the thin tailed normal distribution and the variance of the weather shocks is also unknown, with a gamma prior, then the unconditional prior distribution for the climate sensitivity is the fat tailed t distribution.

[11]See also Geweke (2001).

an asymptotic result. They find that truncating the distribution of uncertainty at a high level invalidates the dismal theorem. They argue that temperature changes should be truncated, since infinite temperatures are not physically possible, which we adopt here.[12]

Our results do not contradict the dismal theorem. Fat tails do in fact remain in our climate sensitivity distribution for any finite number of observations, and without truncation, the risk premium is infinite.

Our results therefore show the importance of results by Costello, Neubert, Polasky, and Solow (2010) for policy. We show that the fat tail is important for near term policy, even with truncation. Nonetheless, we show that with a small number of observations, learning reduces the upper tail of the distribution to close to an exponential. Therefore, although learning has no effect on the infinite risk premium without truncation, with truncation learning does significantly reduce the risk premium.

Put differently, our results highlight the difference between *variance* and *fat tails*. It is possible, through learning, to reduce the variance so that the mass in the upper tail is arbitrarily small. For example, we show that in many cases the planner can reduce the mass in the tail to 1% or 0.1% in just 10-16 years. The planner then easily rejects the hypothesis of a high climate sensitivity. However, tails are still fat in that the rate of decline in the upper tail is eventually slower than exponential. Therefore, with truncation, abatement policy is sensitive to the mass of the fat tail, which learning reduces by reducing the variance.

Considerable debate exists in the literature on the importance of fat tailed uncertainty for near term climate policy when learning is possible. For example, Weitzman (2011) argues that strong inertia in the climate should mean learning is less relevant for near term policy. Given that the stock of GHG emissions is difficult to reduce quickly, it will be difficult to reduce GHG concentrations if we learn climate change is more severe than expected. Conversely, Nordhaus (2011) argues that severe climate change should be evident in the data within the next 50 years, and so time exists to reduce GHG concentrations. Pindyck (2011) points out that the question is inherently quantitative and depends on the cost of insurance, the probability of severe climate change, etc.

This paper provides a quantitative answer. It is indeed important to reduce emissions initially due to the fat tailed uncertainty and the difficulty in reversing GHG stocks. Nonetheless, consistent with Nordhaus' idea, learning is indeed fast enough so that mid course corrections are possible. For example, we show that if the true climate sensitivity implies a steady state temperature change of 5°C for a doubling of GHG concentrations, then the

[12] Millner (2011) discusses other key assumptions of the dismal theorem.

planner reduces emissions to within 1% of the certainty level in only 14 years. Emissions under learning remain slightly below (0.6-0.7%) the certainty level for 53 years as the planner corrects for initially over-emitting, to bring GHG concentrations to the perfect information trajectory.

Note that in all of above results, convergence to the perfect information level of emissions requires only partial learning. Indeed we show that learning is not complete (in the sense that the planner can reject values plus or minus 5% away from the true value with 95% confidence, see for example, Leach 2007) for 50-100 years depending on the true value of the climate sensitivity.[13] Learning is slow, but the remaining uncertainty after the fat tail is statistically rejected is not important for policy. Only the fat tail provides an insurance motivation for near term abatement.

The structure of this paper is as follows. Section 2 presents our integrated assessment model with learning. The model is solved and calibrated in Section 3, and Sections 4 and 5 presents simulation results. Section 6 concludes.

2 Model

The model is similar in spirit to a simplified DICE model (Nordhaus 2007). Economic growth generates GHG emissions, which in turn cause temperatures to rise, reducing productivity. However, we use the abatement cost function from Bartz and Kelly (2008), and use different assumptions about future improvements in the emissions intensity of output. The learning model and stochastic temperature change follows Kelly and Kolstad (1999b) and Leach (2007) except that we look at learning the fat tailed climate sensitivity rather than learning the primitives of the temperature model. The damage function is from Weitzman (2009).

2.1 Economic system

The population of L_t identical households have preferences over consumption C_t/L_t given by the period constant relative risk aversion utility function:

$$U\left(\frac{C_t}{L_t}\right) = \frac{\left(\frac{C_t}{L_t}\right)^{1-\sigma} - 1}{1-\sigma}. \tag{1}$$

[13]In this sense, our results are consistent with Kelly and Kolstad (1999b), who estimate learning about the heat capacity was complete after about 90 years.

A constant returns to scale technology exists that produces output Q_t, from capital K_t, and productivity-augmented labor A_tL_t. Here A_t is labor productivity, which grows exogenously at rate ϕ. Population grows at exogenous rate η. The production technology is such that:

$$Q_t = F\left(K_t, A_tL_t\right) = K_t^{\psi}\left(A_tL_t\right)^{1-\psi}. \tag{2}$$

Unabated GHG emissions are an exogenous proportion $1/B_t$ of output. Let u_t denote the fraction of emissions abated, then $(1-u_t)/B_t$ is the emissions intensity of output, and emissions, E_t, are:

$$E_t = \left(1-u_t\right)\frac{Q_t}{B_t}. \tag{3}$$

The cost of abatement is $\Lambda\left(u_t\right)Q_t$. Hence, output net of abatement costs, Y_t, is

$$Y_t = \left(1-\Lambda\left(u_t\right)\right)Q_t. \tag{4}$$

We assume a convex cost function:

$$\Lambda\left(u_t\right) = 1-\left(1-u_t\right)^{\epsilon}. \tag{5}$$

The abatement cost function (5), differs from standard cost functions in the literature. For example, Nordhaus (2008) uses a two parameter function: $\Lambda\left(u_t\right) = \epsilon_1 u_t^{\epsilon_2}$. The abatement cost function (5) has a particular advantage in that it is consistent with a balanced growth path (Bartz and Kelly 2008), which simplifies the computations considerably. Further, without a balanced growth path, either emissions goes to zero or infinity, or the growth rate in labor productivity must be forced to go to zero.

Using equations (2)-(5) to substitute out for Q_t and u_t implies output net of abatement costs is Cobb-Douglas:

$$Y_t = F\left(K_t, B_tE_t, L_tA_t\right) = K_t^{\theta}\left(B_tE_t\right)^{\epsilon}\left(A_tL_t\right)^{1-\theta-\epsilon}. \tag{6}$$

Here $\theta = \psi(1-\epsilon)$ is the capital share and ϵ can thus be interpreted as the emissions share.[14]

A balanced growth path is a steady state where aggregate capital, output, and consumption all grow at the same constant rate $\left(1+\eta\right)\left(1+\phi\right)-1$. A balanced growth path exists with constant emissions if the exogenous growth rate of B_t equals the growth rate of

[14] Bartz and Kelly (2008) calibrate the emissions share for four air pollutants.

output:[15]

$$B_{t+1} = (1+\eta)(1+\phi) B_t. \tag{7}$$

Exogenous growth of B_t captures both technological change in abatement and compositional changes in output.

Let capital depreciate at rate δ_k. The resource constraint then sets consumption plus net investment equal to production net of abatement costs after damages, $D(T)$, due to climate change:

$$C_t = (1 - D(T_t)) Y_t + (1-\delta_k) K_t - K_{t+1}, \tag{8}$$

where the damage function is:

$$D(T_t) = 1 - e^{-b_1 T_t^{b_2}}. \tag{9}$$

Here b_1 and b_2 are damage parameters.

2.2 Climate system

Let M_t represent the current accumulation of carbon-equivalent GHG in the atmosphere, and MB is the preindustrial stock. We assume the ocean and biosphere absorb atmospheric carbon at a constant rate δ_m. Let $\gamma = A_0 L_0 E_0 / Q_0$ be the initial emission intensity coefficient. The stock of pollution accumulates according to:

$$M_{t+1} - MB = (1-\delta_m)(M_t - MB) + \gamma E_t. \tag{10}$$

We use a one equation physical model for temperature:

$$\hat{T}_t = \hat{T}_{t-1} + \frac{1}{\alpha}\left(F_t - \frac{\hat{T}_{t-1} - \Gamma}{\lambda}\right) + \nu_t. \tag{11}$$

Here $\hat{T}_t$ is the annual global temperature (difference in °C between year t and the 1961-1990 average temperature); Γ is preindustrial temperature difference from the 1961-1990 average temperature; α is the heat capacity of the upper ocean; λ is the climate sensitivity;

[15] Note that if B_t grew slower than the rate of output then the returns to emissions savings innovation would approach infinity, while the returns to labor productivity would go to zero, and the reverse if B_t grew faster than the rate of output.

$\nu \sim N(0, \sigma_\nu^2)$ is the stochastic weather shock; and F_t is radiative forcing of GHGs:

$$F_t = \Omega \log_2 \left(\frac{M_t}{MB} \right). \tag{12}$$

Here Ω is the radiative forcing parameter. Equations (11) and (12) vastly simplify large physical models of climate, known as general circulation models (GCMs). Nonetheless, models similar to (11) and (12) are frequently estimated and used for policy analysis.[16]

The uncertain parameter is the climate sensitivity λ, which measures how responsive temperature is to GHG concentrations. The climate sensitivity amplifies the effect of radiative forcing on temperature. To see this, rewrite equations (11) and (12) as:

$$T_t = \beta_1 T_{t-1} + \beta_2 \log_2 \left(\frac{M_t}{MB} \right) + \nu_t. \tag{13}$$

Here $T_t = \hat{T}_t - \Gamma$ is the annual global temperature deviation from preindustrial level, $\beta_1 = (1 - \frac{1}{\lambda\alpha})$ is the climate feedback parameter, and $\beta_2 = \frac{\Omega}{\alpha}$. The climate feedback parameter is positively related to the climate sensitivity. Since the climate sensitivity is uncertain, the climate feedback parameter β_1 is also uncertain.

Let $\Delta T_{2\times}$ be the steady state temperature change that results from a steady state doubling of the GHG concentrations, then from equations (11) and (13):

$$\Delta T_{2\times} = \Omega\lambda = \frac{\beta_2}{1 - \beta_1}. \tag{14}$$

Since λ is uncertain $\Delta T_{2\times}$ is also uncertain. The steady state temperature change from a doubling of CO_2 is most straightforward to understand, so we will report most results in terms of the equivalent $\Delta T_{2\times}$.

Emissions, E_t, temperature, T_t, and GHG concentrations, M_t, are constant along the balanced growth path.

2.3 Learning

Each period the social planner observes new statistical records of the climate system and updates beliefs on the uncertain feedback parameter. Bayesian learning characterizes this

[16] See for example, Andronova and Schlesinger (2001) and Schwartz (2007). Kelly, Kolstad, Schlesinger, and Andronova (1998) discuss some weaknesses of one equation physical models. Traeger (2012) calibrates a one equation model to match near term temperature changes predicted by the Nordhaus DICE model.

process.

Assume the planner has prior beliefs that the true β_1 is drawn from a normal distribution, $N(\mu_0, S_0)$, where S_0 is the variance of the prior distribution. Let:

$$H_t = T_t - \beta_2 \log_2 \left(\frac{M_t}{MB} \right) = \beta_1 T_{t-1} + \nu_t. \tag{15}$$

Then $H_{t+1} \sim N\left(\mu_{H,t}, \sigma^2_{H,t}\right)$ combines the stochastic weather shocks and feedback uncertainty into a single random variable, where $\mu_{H,t} = \mu_t T_t$ and $\sigma^2_{H,t} = T_t^2 S_t + \sigma^2_\nu$.

The $t+1$ weather shock occurs at the start of the period, before decisions are made. The social planner thus observes H_{t+1}, and T_t and updates the prior on β_1. Let $p_\nu = 1/\sigma^2_\nu$ be the precision of ν. Then Bayes rule implies that the posterior distribution of β_1 is also normally distributed with:

$$\mu_{t+1} = \frac{\mu_t + S_t p_\nu T_t H_{t+1}}{1 + S_t p_\nu T_t^2}, \tag{16}$$

$$S_{t+1} = \frac{S_t}{1 + S_t p_\nu T_t^2}. \tag{17}$$

Note that from equation (17) that the variance estimate on β_1 is monotonically non-increasing with time. We use variance instead of the usual precision since the variance is bounded above by the prior, while the precision is unbounded above.[17] Perfect information implies $S = 0$ and $\mu = \beta_1$.

Roe and Baker (2007) compute the probability density function (PDF) for the climate sensitivity from a Jacobian transformation. Let, $\beta_1 = 1 - \frac{1}{\alpha\lambda} \sim N(\mu, S)$, then the density for the climate sensitivity is:

$$h_\lambda(\lambda) = h_{\beta_1}(\beta_1(\lambda)) \left(\frac{\partial \beta_1}{\partial \lambda} \right) = \frac{1}{\sqrt{2\pi S}} \frac{\frac{1}{\alpha}}{\lambda^2} \exp\left[-\frac{1}{2S} \left(1 - \frac{1}{\alpha\lambda} - \mu \right)^2 \right]. \tag{18}$$

The prior distribution of the climate sensitivity has a fat upper tail if and only if the tail probability declines to zero at a rate slower than exponential:

$$\lim_{\lambda \to \infty} \frac{h_\lambda(\lambda)}{\exp(-a\lambda)} > 0, \quad a > 0. \tag{19}$$

It is straightforward to verify that $h_\lambda(\lambda)$ satisfies condition (19). The fat upper tail is clearly visible in Figure 1, which plots $h_\lambda(\lambda)$.

[17]The computational solution method requires a bounded state space.

2.4 The Recursive Problem

The social planner maximizes the expected present discounted value of the stream of period utilities weighted by the population size:

$$W = \max \sum_{t=0}^{\infty} \beta^t L_t U\left(\frac{C_t}{L_t}\right). \tag{20}$$

We normalize the problem so that variables are in per labor productivity unit terms. Let $f(K,E) = F(K,E,1)$, $k_t = K_t/(A_t L_t)$ and similarly for c_t and y_t, $m_t = \frac{M_t}{MB}$ and assume $\hat{\beta} \equiv \beta(1+\eta)(1+\phi)^{1-\sigma} < 1$. Then we can write the recursive version of the social planning problem as:

$$\begin{aligned} v(k,m,T,\mu,S) \;\; = \;\; & \max_{k',E} \left\{ \frac{c^{1-\sigma}}{1-\sigma} + \right. \\ & \left. \hat{\beta} \int_{H'} v(k',m',T',\mu',S') \,\mathrm{N}\left(H', \mu T, T^2 S + \sigma_\nu^2\right) dH' \right\}, \end{aligned} \tag{21}$$

subject to:

$$c = (1 - D(T)) f(k,E) + (1-\delta_k) k - (1+\eta)(1+\phi) k', \tag{22}$$

$$m' = 1 + (1-\delta_m)(m-1) + \left(\frac{\gamma}{MB}\right) E, \tag{23}$$

$$T' = H' + \beta_2 \log(m'), \tag{24}$$

$$\mu' = \frac{\mu + S p_\nu T H'}{1 + S p_v T^2}, \tag{25}$$

$$S' = \frac{S}{1 + S p_v T^2}. \tag{26}$$

Equation (21) condenses the double expectation over unknown variables β_1 and ν' into an expectation over a single unknown variable H'. Table 1 gives the variable definitions for the above problem.

3 Calibration and Solution Method

Table 2 gives the parameter values. For the economic parameters, we chose a risk aversion coefficient of $\sigma = 1.5$. The discount factor is consistent with a 5% rate of time discount, which implies the model economy matches the US capital to output ratio. The depreciation rate of capital is 6%, which implies the model economy matches the US investment to output ratio. The population growth rate is 1% and the growth rate of per capita GDP is 1.8%, which match US data. These values are broadly consistent with the business cycle literature (Bartz and Kelly 2008, Cooley and Prescott 1995).

The emissions share parameter is also the parameter of the abatement cost function. In general, the abatement cost function has only one parameter and therefore cannot match both the low cost of abatement when u_t is low, and the convexity of the Nordhaus (2008) cost function. We chose the calibrated value to match well for low values of u_t, which results in an abatement cost function which is more convex. Therefore, initial control rates, emissions, and carbon taxes will tend to be similar to Nordhaus, but in later years, when larger control rates are optimal, our results will generally show lower abatement costs and carbon taxes.

The damage parameters are taken from Weitzman (2009). The damage parameters and discount factor are set relatively conservatively. A discount factor closer to one or a high damage scenario (for example, Weitzman also considers $b_2 = 4$) would increase both the value of abatement as climate insurance, and the value of learning. We consider the high discount factor case in Section 6.

One issue with the calibration is that the model sets the growth rate in the inverse of the emissions intensity of output (B_t) equal to the growth rate of output. This is done so that the balanced growth path exists with constant emissions. In the long run, the growth rate of B_t should equal the growth rate of output otherwise the returns to labor productivity innovations will either be infinite or zero. That is, any other value is inconsistent with a more general model in which R&D expenditures flow towards the sector (emissions or labor) with the highest marginal returns to innovation. Nonetheless, in the short run the rate of growth in B_t can differ from that of output (for example, if random innovations in one sector are more successful than the other). In fact, the growth rate of B_t has exceeded the growth rate of output in recent years. Allowing for different short run growth rates (as is done in Nordhaus, 2008, for example), would considerably complicate the analysis by adding additional state variables, however.

The variance of the weather shock is taken from Leach (2007). The remaining climate and forcing parameters are set consistent with Nordhaus (2008).

We use a multidimensional spline approximation of the value function and value function iteration to solve the problem. After we solve the model, we approximate the policy function using another spline approximation and do the simulations. Our numerical integration truncates the temperature distribution at 26.8°C, for computational reasons.[18]

[18]See Costello, Neubert, Polasky, and Solow (2010) for a justification. Further, the numerical integration

4 Results: Learning

4.1 Partial Learning

The solution of the dynamic program is a value function $v(k, m, T, \mu, S)$ and the corresponding policy functions giving optimal investment and emissions $k'^{*}(k, m, T, \mu, S)$ and $E^{*}(k, m, T, \mu, S)$. The evolution of the state and control variables over time follows:

$$\left[\begin{array}{cccccc} k' & m' & T' & \mu' & S' & E \end{array}\right] \equiv G\left(\left[k, m, T, \mu, S\right], \lambda, \nu'\right). \tag{27}$$

Here the function G is comprised of the policy functions and the laws of motion for the state variables, equations (23)-(26). A simulated time path thus requires a set of random draws for ν over time and a true value of λ. Each experiment consists of 150 simulations, over 6,500 years, for a given true value of λ. We report simulations for various hypothetical true values for λ, and then take expectations over the results using the prior distribution to get expected results given the current state of information.

Table 3 gives the initial conditions, which are set to the year 2005. The initial conditions for the learning parameters are from Roe and Baker (2007), who survey various global circulation models (GCMs). Initial conditions for capital, temperature, and GHG concentrations are from Nordhaus (2008).

We say partial learning is complete if the planner can reject values of the climate sensitivity which are in the tail of the distribution at a given confidence level. That is, partial learning is complete if the planner rejects the hypothesis $\lambda < T^{L}/\Omega$ or equivalently $\Delta T_{2\times} < T^{L}$, where T^{L} is the lower border of the tail of the distribution (°C).

The lower border of the climate sensitivity distribution which constitutes the "tail" has no precisely agreed upon definition. Further, from equation (21) the optimal decision takes into account the entire distribution in a continuous way. Our hypothesis is that the tail of the distribution drives current optimal abatement policy. If the mass of the tail shrinks quickly, the remaining uncertainty is irrelevant and the optimal policy converges to the certainty case. The first step is therefore to show that learning about the tail of the distribution differs from learning about the mean, which requires a definition of what constitutes the tail of the distribution.

routine also truncates so that values of the climate feedback parameter outside the unit interval have zero probability. In this sense, a prior distribution bounded by zero/one, such as the beta distribution may be more appropriate. We use the normal distribution since any computational solution method requires a conjugate family of distributions, but also to follow the literature (see for example, Roe and Baker 2007).

Therefore, we set the lower border of the tail of the distribution conservatively as $T^L = \Delta T_{2\times,t} + 1.5$°C, where $T_{2\times,t}$ is the current mean estimate of the climate sensitivity.[19] The initial prior mean equals 2.76°C, so the initial lower border of the tail is $T^L = 4.26$°C, which constitutes the upper 16.7% of the mass of the distribution. If the true value is large, say 5°C, then learning will move the mean of the prior higher, and the uncertainty can eventually be partitioned into uncertainty about the exact value of the climate sensitivity (near 5) and uncertainty about the probability of still higher values (even more disastrous values greater than 6.5 are possible).[20]

The planner's optimization problem (21) does not specify such a test as part of the optimal policy. Instead, we are developing a hypothesis that the effect of uncertainty on optimal policy can be approximately partitioned into two parts: the effect of uncertainty over the exact value of the climate sensitivity, and the effect of uncertainty over the mass of the tail of the distribution.

We consider two confidence levels, 99% and 99.9%. Given the high damages associated with severe climate change, we assume the planner requires a relatively high level of confidence that the climate sensitivity is not large before rejecting the hypothesis that the true climate sensitivity is in the tail of the distribution.

We consider 50 possible true values, $\Delta T^{i*}_{2\times}$, indexed by i. For each $\Delta T^{i*}_{2\times}$, for each simulation j, and associated random vector of weather shocks ν_j, we record the first period, $n^*\left(\Delta T^{i*}_{2\times}, \nu_j\right)$, for which the hypothesis that $\Delta T_{2\times,jt} > T^L$ is rejected and not subsequently not-rejected. We then say that the planner achieves partial learning in period n^*, for true value of $\Delta T^{i*}_{2\times}$ and simulation j.

We then average over all simulations, and weight all n^* using the prior distribution. Mathematically:

$$\mathrm{E}\left[n^* \middle| \mu_0, S_0\right] = \int_{1/\alpha}^{\infty} \int_{\nu} n^*\left(\lambda, \nu\right) N\left(0, \sigma_\nu^2\right) h_\lambda\left(\lambda\right) d\nu d\lambda. \tag{28}$$

Table 4 shows the results. The expected time to complete partial learning varies from 8.99-14.58 periods, depending on the required confidence level. Partial learning is relatively

[19]Values above 1.5° will speed up the learning, values less than 1.5° mean the tail is relatively likely to occur, which is inconsistent with the idea that the true value is in the tail with relatively low probability.

[20]Alternatives are less attractive. If the lower border of the tail is fixed for all t (say $T^L = 5$°C), then the planner must essentially learn the exact value of the climate sensitivity for true values near the lower border, since in that case the tail of the distribution converges to the upper half of the distribution. As $\Delta T_{2\times,t} \to T^L$, this definition of partial learning requires an arbitrary large number of periods to complete, which is not consistent with the optimal policy in Section 5.

quick for three reasons. First, by definition the initial mass of the upper tail is not large. Therefore it takes relatively few observations away from the tail to reduce the mass to 1% or 0.1%. Second, although the global temperature is subject to random weather shocks which makes the exact value of the climate sensitivity difficult to pin down, the calibrated standard deviation of the shocks is only 0.11°C, which makes it relatively easy to reject high values of the climate sensitivity if the true value is not too large. Third, most of the mass of the prior distribution is relatively close to the mean. Therefore, true values with smaller learning times receive more weight when calculating the expected learning time conditional on the prior distribution. As expected, partial learning at the 99.9% confidence level requires more periods than at the 99% confidence level.

Figure 2 shows the learning time as a function of $\Delta T^*_{2\times}$ for both confidence levels. For true values near the mean of the prior distribution (2.76°C), learning takes less than 10 periods, but rises to 90 years or more as $\Delta T^*_{2\times}$ increases. Partial learning becomes increasingly difficult as $\Delta T^*_{2\times}$ increases. Since $\lambda = 1/(\alpha(1-\beta_1))$, for values of β_1 near one, small differences in β_1 generate large differences in λ. Therefore, the planner must learn β_1 with increased precision to reject values of λ in the tail of the distribution when β_1 is near one. For example, with β_1 equal to the prior, the lower bound of the tail ($T^L = 4.26$) corresponds to $\beta_1^L = 0.77$, which is rejected in reasonable time given a true value of $\beta_1^* = 0.65$ and a calibrated standard deviation of the weather shocks equal to 0.11. Conversely, if $\Delta T^*_{2\times} = 5$ and $T^L = 6.5$ then $\beta_1^* = 0.81$ and $\beta_1^L = 0.85$, which is much harder to reject.

Put differently, for β_1 large, precise learning is much more important, because small differences in β_1 imply large differences in the decay rate of GHG "shocks" to temperature. Small differences in β_1 therefore imply large differences in steady state temperature for a given concentration of GHGs, making a precise estimate of β_1 more important.

Figures 1 and 2 together show that, for most of the prior distribution, partial learning requires less than 20 periods to complete. The expected number of periods until partial learning is complete conditional on prior information is relatively small. Longer learning times are possible but unlikely.[21]

[21]Note that the priors are based on physical, rather than statistical models of temperature change. Therefore, the results may be interpreted as saying that to confirm or refute physical models which predict a relatively high climate sensitivity will require relatively little additional data.

4.2 Alternative Assumptions

We next report learning times given alternative assumptions. First, we simulate the model with an ocean layer.[22] Let $\hat{O}_t$ be the ocean temperature in deviations from the recent average and ζ and τ be heat transfer coefficients, then the climate system (11) now becomes:

$$\hat{T}_t = \hat{T}_{t-1} + \frac{1}{\alpha}\left(F_t - \frac{\hat{T}_{t-1} - \Gamma}{\lambda}\right) - \zeta\left(\hat{T}_{t-1} - \hat{O}_{t-1}\right) + \nu_t, \tag{29}$$

$$\hat{O}_t = \hat{O}_{t-1} + \tau\left(\hat{T}_{t-1} - \hat{O}_{t-1}\right). \tag{30}$$

Because heat transfer to the ocean is a slow process, the model now takes much longer to reach the steady state. However, Table 4 shows that partial learning is complete in 15.84-24.39 periods, which is only slightly longer than with the one equation model. To learn the true climate sensitivity, the planner must statistically isolate the upward trend in the climate from the stochastic weather shocks. The ability of learning to do so is only marginally affected by the ocean and other climate processes with long lags, since the ocean temperature is essentially constant on a year to year basis. Thus our results are robust to the addition of a more complicated climate model.

A very large weather shock can cause a hypothesis which was rejected in period n to be no longer rejected in $n+1$. We also considered an alternative criterion in which n^* is the first period for which the hypothesis is first rejected, even if the hypothesis is no longer rejected in a subsequent period. Table 4 gives the results, which are only slightly smaller than the base case given in Table 4. As expected, learning is faster in this case, but still of a similar magnitude.

4.3 Other Criteria

The above results constitute partial learning in that the only goal is to reject values of the climate sensitivity in the tail of the prior distribution, not to fully learn the true value. Similar to previous work (Leach 2007), we define two hypothesis tests, $\Delta T_{2\times} \leq 0.95 \cdot \Delta T^*_{2\times}$ and $\Delta T_{2\times} \geq 1.05 \cdot \Delta T^*_{2\times}$, where $\Delta T^*_{2\times}$ is the true value, and a desired confidence level of 95%. If the planner rejects both hypothesis, we say that the planner has fully learned the uncertain climate sensitivity, or that full learning is complete. Figure 3 plots the mean number of

[22]For computational reasons, we use the emissions and investment policy functions given in equation (27), which assumes the ocean layer is constant. Since the ocean layer changes slowly over time, this assumption is reasonable for the short term.

periods required to fully learn the climate sensitivity. The learning time is increasing in the true value. First, as the true value gets farther from the prior, Bayes rule must completely reject the prior information in favor of the new observations. Second, as noted above, $\Delta T_{2\times}$ is a nonlinear function of β_1. When $\Delta T_{2\times}$ is large, the range of values of β_1 that constitute the plus or minus 5% of the true value becomes narrower, which increases the expected learning time. Integrating over the prior distribution, we find an average learning time equal to 79.67 years, This is roughly consistent with the previous literature (Kelly and Kolstad 1999b, Leach 2007).

Figure 4 plots the learning dynamics for the mean of the feedback parameter μ for three simulations, when the true value is 0.65 ($\Delta T_{2\times} = 2.76$). In all simulations, learning converges rather slowly to the true value (about 55-150 years). However, even for simulation 5, where the first few weather shocks are positive, the mean estimate of $\Delta T_{2\times}$ is well below the border of the tail, $T^L = 4.26$. The hypothesis $T^L \geq 4.26$ is first rejected at the 99.9% level in periods 7, 10, and 5, respectively. Figure 5 plots the density for simulation 5 after 10 periods. Considerable uncertainty remains. Values of $\Delta T_{2\times}$ between 2.5 and 4 are plausible. However, values above 4 are very unlikely. Partial learning is complete, although full learning is not complete.

The last 3 rows of Table 5 presents the same information for the mean of 150 simulations. The initial mass of the tail is 16.5%, and the 95% confidence interval admits a wide range of possible values for the climate sensitivity, from the benign 0.97 to the disastrous 7.09°C temperature change for a doubling of GHGs. By 2015, however, partial learning is complete at the 99.9% confidence level, whereas full learning is not complete: the 95% confidence interval is still fairly wide at almost 1°C. By 2050, the 95% confidence is fairly narrow, although full learning is still not quite complete.

Figure 6 reports the learning dynamics for the standard deviation of the prior distribution of the feedback parameter, $\sqrt{S}$, for simulation 5 when the true climate sensitivity equals the prior. The standard deviation is monotonically decreasing as expected, and converges to zero.

Learning the climate sensitivity precisely is a slow process. According to the prior literature with thin tails, slow learning indicates the optimal policy under learning is unlikely to be much different than the optimal policy without learning. Since learning is slow, the planner acts using current information. However, the above analysis shows that with learning the planner rejects extreme values relatively quickly, unless the true value is large. With fat tails, the extreme values drive current policy. Therefore, the learning is potentially much more policy relevant with fat tails. The next section makes these ideas precise.

5 Results: Optimal Policy

5.1 Optimal Insurance

In this section, we examine how learning and fat tailed uncertainty affect the optimal emissions policy. Figure 7 plots optimal emissions for the case where the true climate sensitivity is equal to the initial prior value. The circle line corresponds to perfect information, where the initial variance of the prior is set to zero. In this case, optimal emissions increases for a short period of time and is then decreasing. The initial world capital stock is only 68% of its steady state level. Therefore, the planner postpones most emissions control until the capital stock has converged and more resources are available. Both damages and costs are a fraction of world GDP, so an increase in GDP affects both damages and costs equally. However, more wealth implies more consumption, which decreases the marginal utility of consumption. Therefore, emissions control becomes more attractive. In addition, each year's emissions have only a small effect on the GHG concentration, so the planner does not incur much additional damage by waiting.

The line with squares shows the optimal policy under uncertainty with learning. As the planner learns the true value, the emissions under learning and uncertainty converges to emissions given perfect information. Notice that emissions under learning are initially below the perfect information case. The planner insures, emitting less than under perfect information just in case climate change turns out to be severe. Emissions are initially 19.3% lower under learning than under perfect information (Table 5), but are only 1.1% lower by 2021. Uncertainty matters for a relatively short period of time. In Figure 7, the true value is the initial guess. Therefore, the planner quickly rejects values of the climate sensitivity in the fat tail, and the policy approximately converges to the perfect information case by 2025.[23]

The plus line corresponds to the no learning case. In this case, the learning parameters μ and S are not updated, despite the new information. In addition, the model is solved so that the planner knows μ and S will not be updated. Therefore, the no learning case differs from the learning case, even in the initial period when the state vector is identical

[23]Note that emissions are slightly above the perfect information case for a short period of time. This is because the planner has under-emitted relative to perfect information during the learning process. The planner can therefore over-emit after rejecting the fat tail to reach the same steady state stock of GHGs.

for both policies. Emissions are lowest under no learning. The planner must insure more by reducing emissions without learning, because the planner knows that she cannot adjust later as more information is revealed. Therefore, learning reduces the need for climate insurance. Emissions for the no learning case are initially 38% below emissions given perfect information, whereas emissions under learning are 19.3% below perfect information. Therefore, learning reduces climate insurance by about 50% (Table 5).

Emissions under no learning are below the true optimal emissions for the entire time path. The planner under no learning must continue to insure, whereas under learning climate insurance is required for only a short time. Emissions for the no learning case are 48% below perfect information in 2020, whereas emissions under learning are only 1.42% below certainty at this point.

Given a true value equal to the prior ($\Delta T_{2\times} = 2.76$), the planner rejects the tail at the 99.9% confidence level after 7.2 periods (see Figure 2 and Table 5). Emissions under learning is only 4.64% lower than emissions under certainty after 8 periods, and is only 1.07% lower after 16 periods (Figure 7 and Table 5). In contrast, full learning is not complete for 63.8 years if $\Delta T_{2\times} = 2.76$ (see Figure 3 and Table 5). Therefore, optimal emissions policy is more sensitive to partial learning. Although uncertainty is present and the climate sensitivity is difficult to pin down precisely, the planner rejects values in the tail quickly and thereafter proceeds as if the planner was certain that $\Delta T_{2\times}$ equals the mean of the current prior.

Figure 8 gives the emissions control rate, which is increasing over time under certainty. This is similar to the "ramp up" strategies found for example by Nordhaus (2008). The control rate under learning is initially more stringent, but converges to the certainty case at approximately the same rate as emissions. Without learning, the control rate remains elevated as the planner continues to insure. The initial control rate under learning is 19.3%, which is similar to values found in the literature. In contrast, without learning the initial control rate is 38.3%, approximately twice as high (Table 5). Figure 9 shows the carbon tax. The initial carbon tax is \$46.1 per ton, also within the range of typical estimates.

Figure 10 gives the path of temperature increases. The policies with learning and perfect information differ very little in terms of the GHG stock, since the planner adjusts emissions to keep the economy on the same GHG stock trajectory after learning takes place. Emissions and the GHG stock are lower under no learning due to the insurance, resulting in a smaller temperature increase.

Figures 11-13 repeat the experiment with the true $\Delta T_{2\times} = 2$ rather than 2.76. In this case, the planner learns over time that the climate has only small feedback effects, and therefore that GHG concentrations cause smaller steady state temperature increases. The planner then emits considerably more under perfect information than when $\Delta T_{2\times} = 2.76$, 3.55 gigatons more in the steady state. Figure 12 indicates the planner under perfect information reduces emissions by less than one percent initially, which increases to a maximum reduction of only 1.5%. Under learning, the planner begins with the same information set as when the true $\Delta T_{2\times} = 2.76$, and therefore chooses the same initial policy as in Figure 7. As new information arrives which decreases the prior, the planner under learning begins to increase emissions.

The planner rejects the fat tail sooner here, since a relatively small λ implies β_1^L is not close to β_1^*. Emissions under learning converge to within 0.4% of emissions under perfect information after 4 periods. Similarly, partial learning is complete after 3.65-5.07 periods at the 99% and 99.9% confidence levels, respectively (Figure 2). In contrast, full learning is not complete until after 57.0 periods (Figure 3). After ruling out the fat tail, the planner proceeds along a path very close to perfect information, even though uncertainty remains.

Without learning, the initial conditions are identical as when the true value is 2.76. Therefore, the initial policies without learning are identical in Figures 7 and 11. Since the planner does not update the prior, emissions are below perfect information indefinitely. In fact, emissions without learning are quite similar in Figures 7 and 11. The insurance motivation is the main determinant of emissions policy without learning, and differences in emissions policy caused by the different temperature trajectories are minor.

Finally, Figures 14-16 repeat the experiment with the true $\Delta T_{2\times} = 5$. Given that GHG concentrations are projected to more than double,[24] this represents a high damage case. Under certainty, the planner responds to the high $\Delta T_{2\times}$ by severely limiting emissions. Figure 15 indicates the planner reduces emissions initial by 52.4% and by 72.2% in the steady state under certainty.

Initial optimal emissions policy with learning is unchanged from the previous cases, since the initial beliefs are unchanged. Emissions fall over time as the planner increases the mean belief of the climate sensitivity over time in response to higher than expected temperatures. Emissions policy under learning converges to within 1.1% of emissions policy under certainty in 17 periods. From Figure 2, partial learning is complete after 16.09 periods at the 99.9% confidence level. In contrast, learning is not complete for about 105.6 years (see Figure 3) for a true value of $\Delta T_{2\times} = 5$. The planner learns that the climate sensitivity is much higher

[24]However, the optimal emissions path under certainty limits GHG concentrations to only 55% above preindustrial levels when $\Delta T_{2\times} = 5$.

than the prior, but also learns that still higher values are unlikely, even if the exact value is difficult to pin down. Optimal emissions converges to the certainty case, and the remaining uncertainty is not relevant for policy.

Emissions under learning is about 0.6%-0.7% less than emissions under certainty from 2022 to 2270. Under learning the planner has over-emitted relative to perfect information in the first 16 years, and must therefore under emit relative to certainty to bring GHG concentrations to the same steady state trajectory. The planner smoothes out the error correcting over a considerable period because emissions control costs are convex and because of consumption smoothing.

For all three true values of $\Delta T_{2\times}$, the upper tail of the distribution is initially policy relevant. Emissions under learning are significantly below emissions if the prior is known with certainty (Figure 11). However, learning quickly reduces the mass of the upper tail of the distribution, and emissions quickly converge to the certainty case.

5.2 Fat Tails and Variance

Figures 1, 17, and 18 shows the evolution of the posterior PDF for $\Delta T_{2\times}$ when the true value is 2.76 (equal to the current prior), after 0, 10, and 50 periods, respectively. We also plot the normal distribution with identical parameters on each graph,[25] to emphasize the fat tail of the distribution of the climate sensitivity. Contrasting Figures 1 and 17, we see that after only 10 periods the mass of the fat tail shrinks considerably. Values of $\Delta T_{2\times}$ above 3.5 are very unlikely.

However, as Weitzman (2009) points out for the case of the t distribution, the tail remains fatter than normal for any finite number of observations. Here our distribution is not t, but the same result applies. Condition (19) holds for any $S > 0$, so the tails remain fat for any finite number of observations. After 50 years, the mass of the tail barely visible on Figure 18, but nonetheless, the tail is still fatter than normal.

Section 5.1 shows fat tails are relevant initially for optimal emissions policy. Emissions with learning are 19.3% below the certainty case. Further, when a climate sensitivity above 4.26 is rejected at the 99.9% level, emissions are within 1% of certainty, even though fat tails remain for all finite observations. Therefore, what is important for emissions policy is not the existence of fat tails, but the mass of the fat tail, which is a function of the variance. Learning is relevant since learning affects the variance.

[25] Note that the variance of λ is infinite. Figures 1, 17, and 18 plot $h(\lambda; \mu, S)$ along with a normal distribution with mean μ and variance S.

6 Sensitivity Analysis

The climate sensitivity is considered a major source of uncertainty in integrated assessment models (Kelly and Kolstad 1999b). Nonetheless, other sources of uncertainty also exist, especially the level and convexity of damages (Weitzman 2009) and the discount factor (Nordhaus 2008). Higher or more convex damages or a discount factor closer to one all increase the benefits of abatement. We therefore consider a representative sensitivity analysis, decreasing the pure rate of time preference from 0.05 to 0.03, or alternatively increasing β to 0.97.

Figures 19-21 graph the mean optimal emissions for the high discount factor case. Contrasting Figures 7 and 19, we see that optimal emissions are lower with the higher discount factor, as expected. Optimal emissions under certainty in the initial period fall by 24.33% when the discount factor increases from 0.95 to 0.97. Emissions under learning and no learning fall by 39.63% and 35.37%, respectively, when the discount factor increases.

Since damages are potentially greater, optimal insurance also increases. Table 6 indicates current emissions under no learning are 47.31% lower than emissions under certainty for the high discount factor case, versus 38.31% for the low discount factor case. Similarly, emissions under learning are 35.62% lower under learning than emissions under certainty for the high β case, versus 19.3% for the low β case. Although damages are potentially greater, learning becomes somewhat less important, since the high discount factor means large emissions reductions are optimal for a wide range of outcomes of the learning process. Learning reduces the insurance premium by only 24.71% in the high β case, versus about 50% in the low β case.

Both partial and full learning are slightly slower for the high β case. Kelly and Kolstad (1999b) prove that more restrictive climate policies slow learning because the climate change signal is less pronounced amidst the noisy weather. Table 6 shows that this effect is small, however. The mass of the fat tail is 1.4% in 2010 when $\beta = 0.97$, versus 1.3% when $\beta = 0.95$. The 95% confidence interval is $[2.37, 3.31]$ in 2015 when $\beta = 0.97$ versus $[2.39, 3.30]$ when $\beta = 0.95$. Although learning is slightly slower, emissions under learning converges faster to perfect information when $\beta = 0.97$. Table 6 and Figure 19 indicate emissions under learning converges to within 1% of emissions under certainty in 7 years, versus about 17 period when $\beta = 0.95$.

Overall, sensitivity analysis indicates that if the discount factor is closer to one or damages are greater or more convex, the main results remain. Fat tails are initially policy relevant, with a large insurance premium, learning significantly reduces the insurance premium, and learning is fast in that the fat tail is rejected quickly, although significant uncertainty remains for decades, that uncertainty is not relevant for optimal policy.

7 Concluding Remarks

In this paper, we study the effect of a possible high climate sensitivity on near term optimal climate change policy, accounting for learning and uncertainty in the climate system. We find three major results. First, fat tails are initially policy relevant in that near term GHG emissions policy is much more restrictive when the planner accounts for fat tailed uncertainty in the climate sensitivity (a 38% reduction in emissions). Second, when the planner accounts for learning, the near term emissions reduction falls by half to only 19.3%. Third, although full learning is slow, learning quickly reduces the mass of the fat tail. Optimal emissions policy is much more sensitive to the mass of the fat tail than the uncertainty in the prior around the mean. Therefore, optimal emissions policy converges quickly to perfect information, even though some uncertainty remains for decades.

The planner knows values of the climate sensitivity in the tail of the prior distribution will be rejected quickly at a high level of confidence if the true climate sensitivity is moderate. If the climate sensitivity is high the planner can quickly reject still higher values, and quickly adjusts emissions to get back on the optimal temperature trajectory. The planner has an option to essentially purchase climate insurance: by paying to limit GHG emissions today, the planner prevents GHG concentrations from rising, which in turn prevents the possibility of very high temperature changes. Without learning, the planner takes out a significant amount of insurance. However, with learning the planner insures about 50% less. First, learning quickly rejects values of the climate sensitivity in the fat tailed part of the prior distribution, if the true climate change is moderate. Second, the planner has time to adjust emissions to keep the economy on the same GHG stock trajectory. Climate insurance under learning in most cases falls to less than 1% after about 17 years as the planner reduces the mass of the tail end of the distribution and the remaining uncertainty is not important for emissions policy.

Several caveats are in order. First, for computational reasons, our model of the climate system is highly simplified. For example, we do not include a separate ocean layer. Nonetheless, we computed an optimal policy assuming the ocean temperature is constant, but simulated the model and learning with an ocean layer. The results are not much different since learning here is about isolating the upward trend in the atmospheric temperature from the stochastic weather shocks. The ability of learning to do so is only marginally affected by the ocean and other climate processes with long lags, since the ocean temperature is essentially constant on a year to year basis. Thus our results are robust to the addition of longer lags to the climate model.

Second, we consider only a single uncertainty, the climate sensitivity. Climate change has many uncertainties, including the parameters of the damage function, the heat capacity of the ocean, etc. In general, multiple simultaneous uncertainties slows learning. It is unclear, however, how much the partial learning we consider here would slow.

Third, our model has no irreversibilities, tipping points, etc. The existence of irreversibilities makes the planner much more cautious, which increases insurance with or without learning. Learning would certainly still reduce climate insurance in this case, but by less as the planner may not be able to correct a mistake of initially over-emitting. We leave this interesting extension to future research.

Fourth, in our model the planner estimates climate feedbacks using current data. If the climate is subject to regime shifts which occur in the far future, then it might be difficult to learn about the existence of regime shifts today. However, if the process which causes the regime shift is observable today, then our model still applies. Suppose for example, the climate sensitivity is different if the polar ice caps melt as sunlight no longer reflects back into space as efficiently (the albedo effect). If one can estimate the albedo effect by observing changes in ice cover and changes in temperature, then the planner can learn the albedo effect before a regime shift to a world without polar ice caps occurs.

Regardless, our main results are likely robust to any of these extensions. Fat tails matter for climate policy, even if the distribution has a truncation point. Nonetheless, we show that learning significantly reduces the influence of fat tails, especially over the near term. Given these results it is important for policy makers to maintain policy flexibility, and to stand ready to quickly adjust the emissions policy as new information arrives.

Finally, our results have interesting potential implications for recent research which finds fat tailed uncertainty in other contexts (equity market returns, banking crises, etc.). Fat tails in other contexts is typically modeled as an exogenous property of the return distribution rather than an endogenous implication of parameter uncertainty. Our results show that if the exogenously imposed fat tails are in fact the result of parameter uncertainty, then learning has the potential to reduce fat tailed uncertainty over time, which limits the risk premium of fat tailed uncertainty. We leave this interesting possibility to future research.

8 Figures and Tables

	State Variables
k	Productivity adjusted capital stock per capital
m	GHG concentration over preindustrial level
T	Atmospheric temperature
μ	Mean estimate of feedback parameter
S	Variance of the prior distribution of feedback parameter
	Control Variables
E	GHG emissions
k'	Gross investment
	Random Variables
ν'	Weather shock
β_1	Feedback parameter

Table 1: Variable definitions.

Parameter	Description	Calibrated Value
β	Discount factor	0.952
σ	Coefficient of risk aversion	1.5
ψ	Capital share	0.402
ϵ	Emissions share	0.0057
δ_k	Capital depreciation rate	0.046
η	Population growth rate	0.011
ϕ	Productivity growth rate	0.018
δ_m	GHG stock decay	0.0083
b_1	Damage function parameter	0.003
b_2	Damage function parameter	2
γ	Initial emissions intensity (GtC/trillion 2005 \$)	4.66
MB	Preindustrial GHG concentrations (GtC)	596.4
α	Heat capacity of the ocean	0.22^{-1}
Γ	Preindustrial temperature	-0.4607
Ω	Forcing parameter (W/M^2)	4.39
σ_ν	Standard deviation of weather shocks	0.11
$\max \Delta T_{2\times}$	Maximum temperature change	26.83
ζ	Heat transfer from atmosphere to ocean	0.05
τ	Heat transfer from ocean to atmosphere	0.30

Table 2: Calibrated parameter values.

State Variable	Initial Value	Units
K	137	Trillions of 2005 dollars
m	1.3563	Fraction relative to preindustrial
T	0.731	°C above preindustrial
μ	0.65	Watts per square meter (W/M^2)
$\sqrt{S}$	0.13	Watts per meter squared (W/M^2)

Table 3: Initial conditions corresponding to 2005. Dollar units are adjusted for purchasing power parity. Initial values for K, m, and T are from Nordhaus (2008). Initial values for μ and S are from Roe and Baker (2007).

Test	99%	99.9%	Full
First reject, not subsequently not rejected	8.99	14.58	85.35
First reject, not subsequently not rejected, with ocean	15.84	24.39	110.52
First reject	6.86	10.43	71.06
First Reject, with ocean	11.18	17.43	94.94

Table 4: Expected learning time in years conditional on current information. Expected number of years until the hypothesis $\Delta T_{2\times} \geq \Delta T^*_{2\times} + 1.5$ is rejected at the given confidence level, where $\Delta T^*_{2\times}$ is the true value. Column 4 is the expected number of years until full learning is complete.

Year	2005	2010	2015	2020	2050
Emissions: Learning (% vs certainty)	19.30	7.68	3.69	1.42	-0.49
Emissions: No learning (% vs certainty)	38.31	46.3	48.17	48.02	45.72
Control rate: Learning (% vs no learning)	49.55	81.36	81.36	75.54	54.76
Carbon tax: Learning (% vs no learning)	23.44	40.71	45.50	46.72	45.96
Probability $T \geq 4.26$	0.165	0.013	0.0002	$< 10^{-6}$	$< 10^{-6}$
95% confidence interval lower bound, $\Delta T_{2\times}$	0.97	2.07	2.39	2.51	2.71
95% confidence interval upper bound, $\Delta T_{2\times}$	7.09	3.57	3.30	3.19	3.00

Table 5: Difference in optimal emissions policy, learning, no learning, and perfect information. The true value is the prior. The first two rows give the percent difference between emissions under certainty and emissions under learning and no learning. That is the first cell of the table indicates emissions under learning are 19.3% lower than emissions under certainty. Rows 3-4 give the percent difference in policies under no learning and the policies under learning. The last two rows give the progress of partial learning versus full learning. All results are the mean of 150 simulations.

Year	2005	2010	2015	2020	2050
Emissions: Learning (% vs certainty)	35.62	1.54	0.26	-0.11	-0.24
Emissions: No learning (% vs certainty)	47.31	46.28	45.31	44.63	42.96
Control rate: Learning (% vs no learning)	11.99	38.08	34.22	31.42	25.07
Carbon tax: Learning (% vs no learning)	18.07	42.91	44.11	44.01	42.95
Probability $T \geq 4.26$	0.165	0.014	0.0003	$< 10^{-5}$	$< 10^{-6}$
95% confidence interval lower bound, $\Delta T_{2\times}$	0.97	2.05	2.37	2.5	2.69
95% confidence interval upper bound, $\Delta T_{2\times}$	7.09	3.58	3.31	3.21	3.02

Table 6: Difference in optimal emissions policy, learning, no learning, and perfect information, when $\beta = 0.97$. The true value is the prior. The first two rows give the percent difference between emissions under certainty and emissions under learning and no learning. Rows 3-4 give the percent difference in policies under no learning and the policies under learning. The last two rows give the progress of partial learning versus full learning. All results are the mean of 150 simulations.

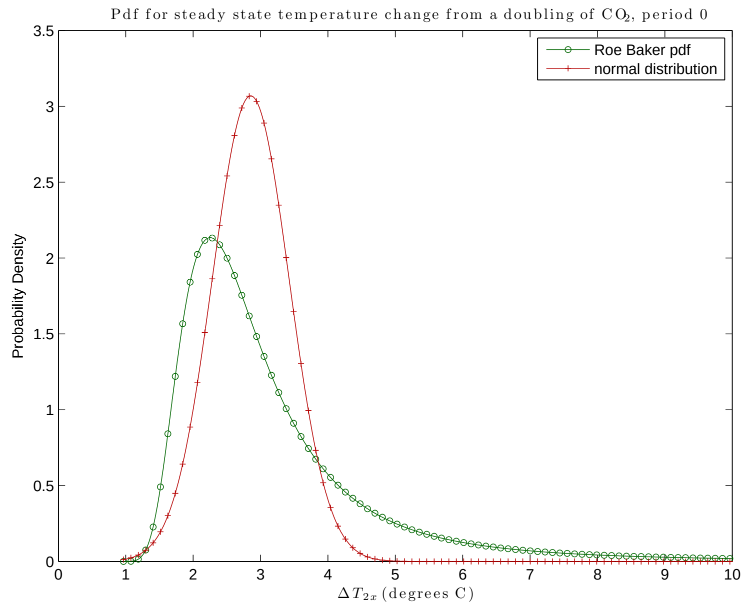

Figure 1: Prior PDF for the climate sensitivity. $\beta_1 \sim N\,(0.65, 0.013)$.

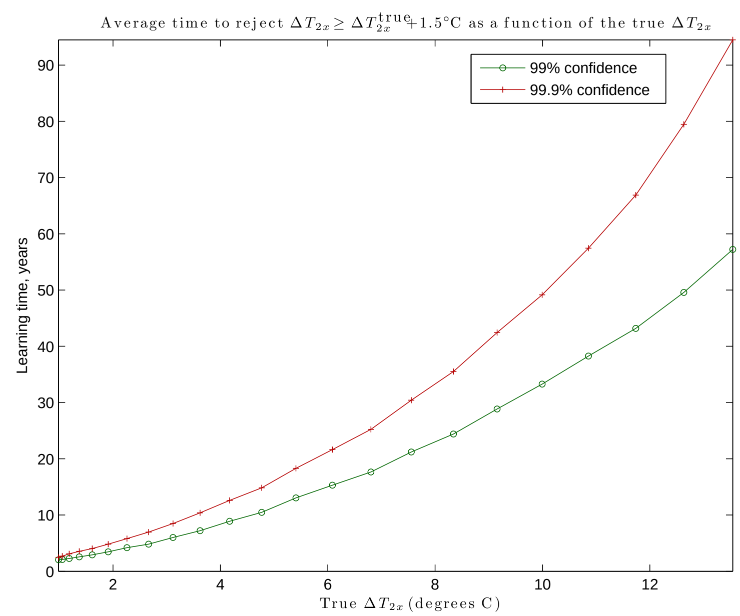

Figure 2: Mean learning time to reject $\Delta T_{2\times} \geq \Delta T^{*}_{2\times} + 1.5$°C at 1% and 0.1% critical values as a function of the true $\Delta T_{2\times}$. Average of 150 simulations for each value of $\Delta T^{*}_{2\times}$.

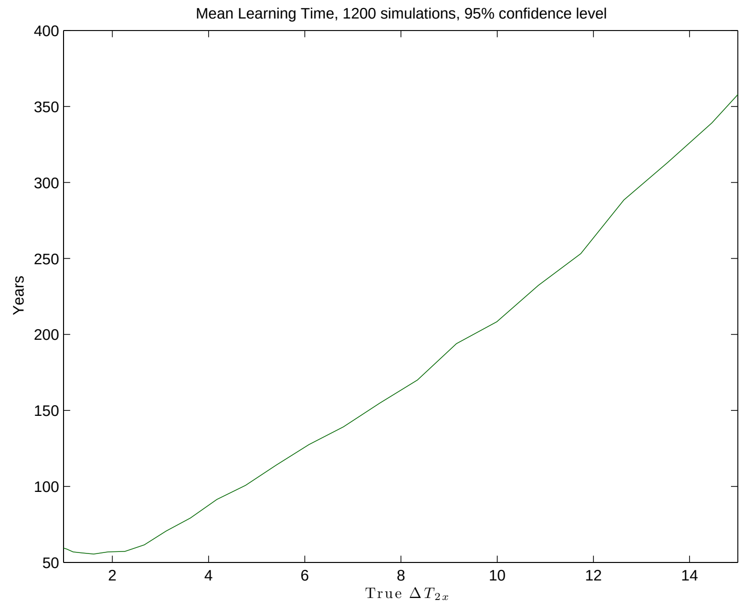

Figure 3: Mean learning time required to reject the hypothesis that $\Delta T_{2\times} \leq 0.95 \cdot \Delta T^*_{2\times}$ and $\Delta T_{2\times} \geq \Delta 1.05 \cdot T^*_{2\times}$ with 95% confidence. Mean of 1200 simulations.

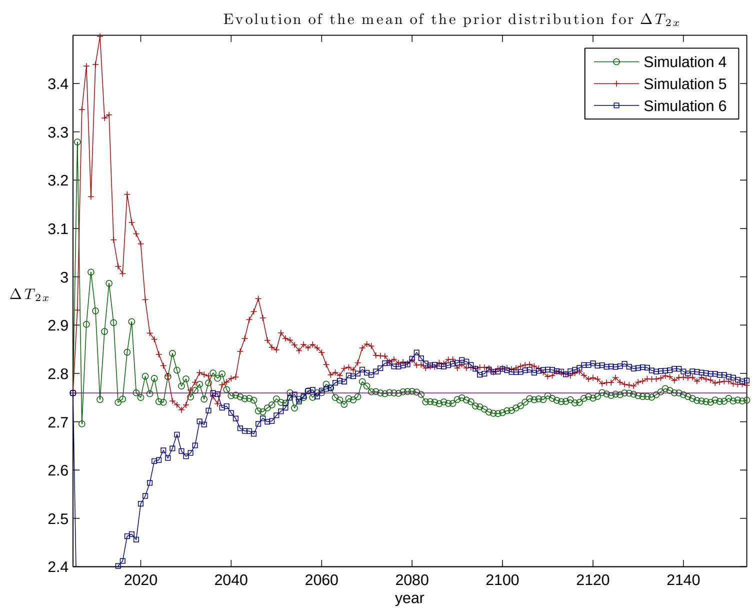

Figure 4: Learning dynamics for the mean of the feedback parameter. The solid line is the true value and is also the ideal case of learning when all weather shocks are zero.

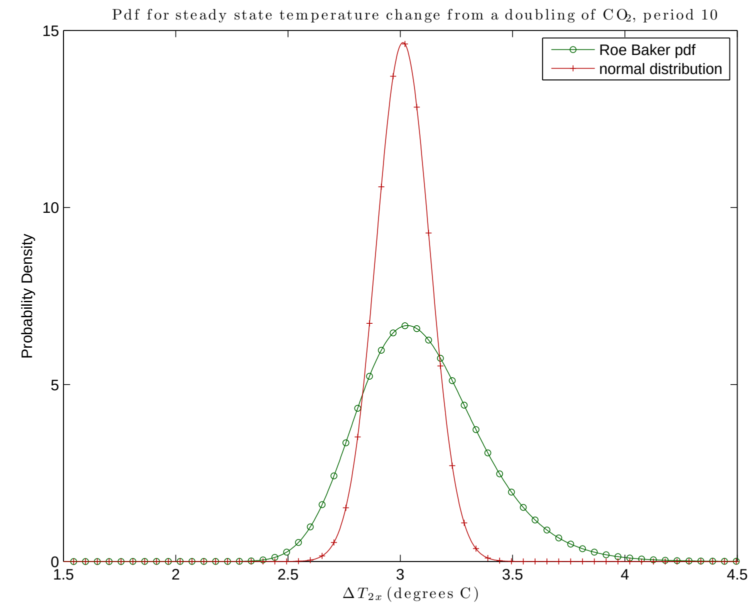

Figure 5: Posterior PDF for the climate sensitivity after 10 observations. The true value equals the prior, simulation 5.

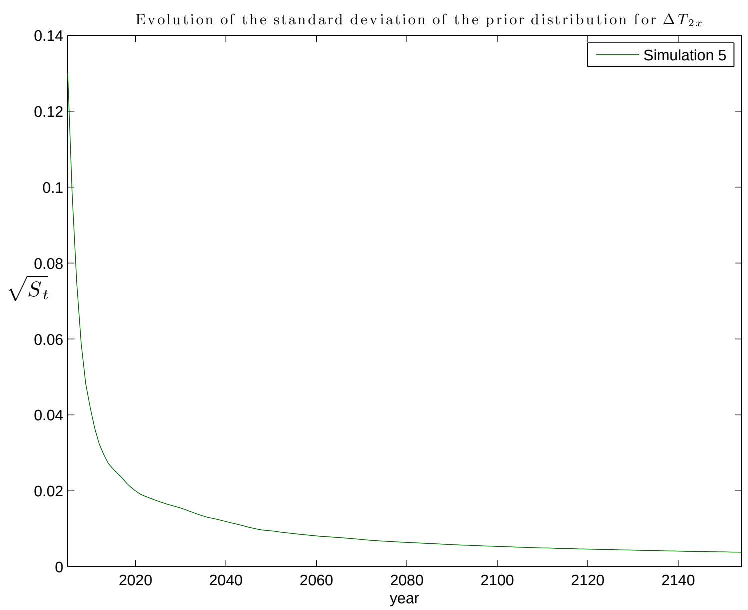

Figure 6: Learning dynamics for the variance of the feedback parameter. The true value equals the prior.

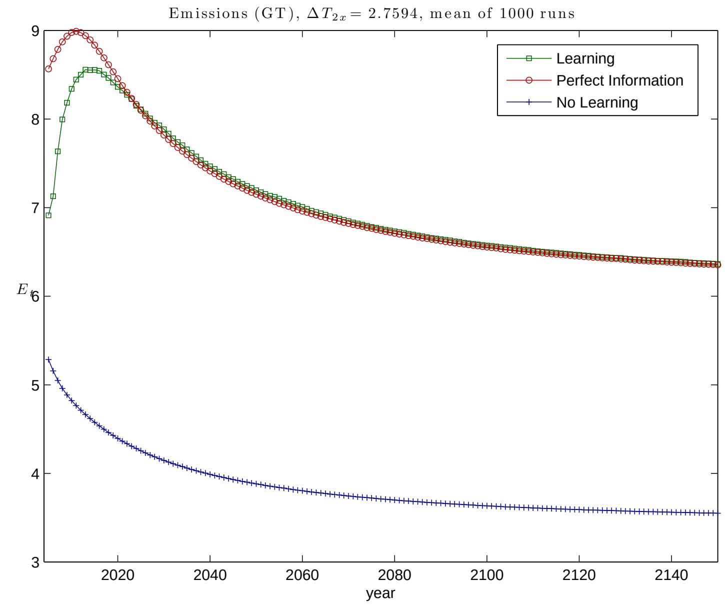

Figure 7: Optimal emissions policy in gigatons (GT), true $\Delta T_{2\times} = 2.76$. Mean of 1000 runs.

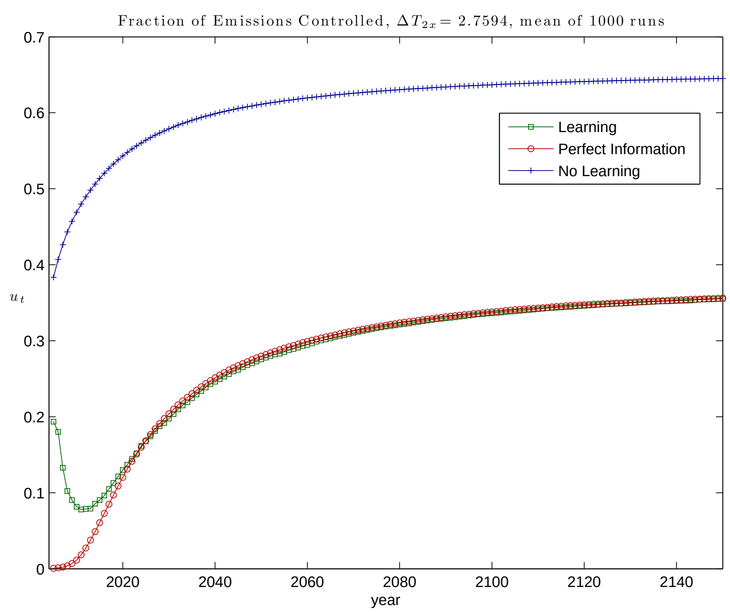

Figure 8: Optimal emissions control rate, true $\Delta T_{2\times} = 2.76$. Mean of 1000 runs.

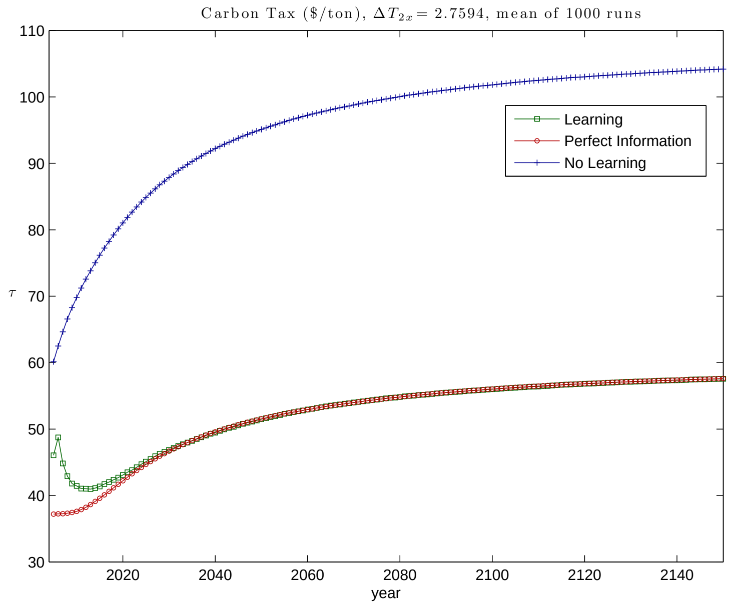

Figure 9: Optimal carbon tax ($/ton), true $\Delta T_{2\times} = 2.76$. Mean of 1000 runs.

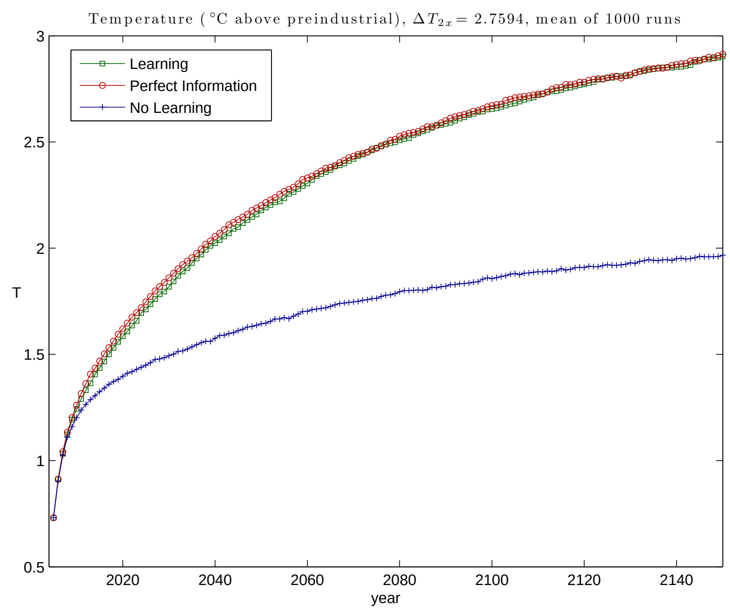

Figure 10: Temperature, degrees C above preindustrial, true $\Delta T_{2\times} = 2.76$. Mean of 1000 runs.

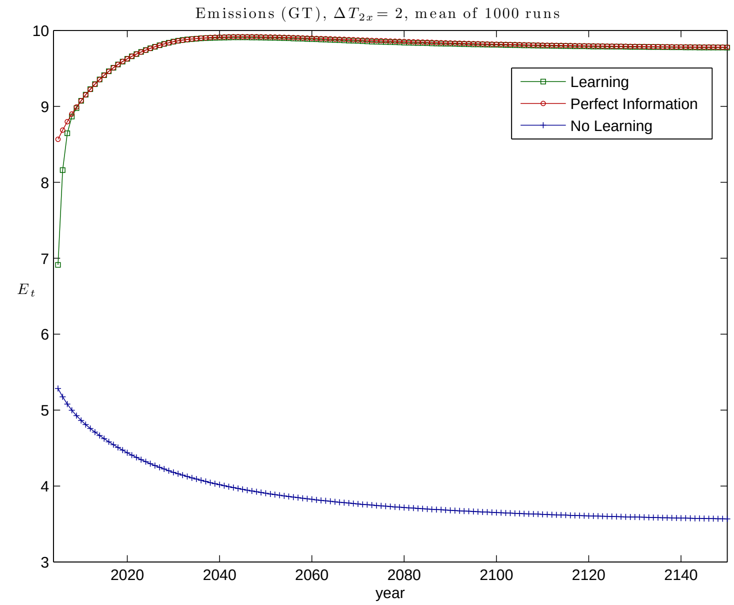

Figure 11: Optimal emissions policy (GT), true $\Delta T_{2\times} = 2$. Mean of 1000 runs.

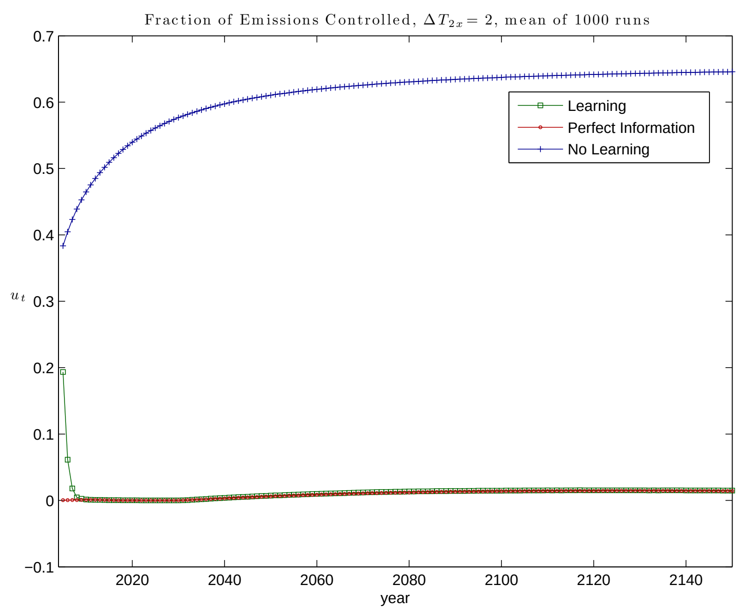

Figure 12: Optimal emissions control rate, true $\Delta T_{2\times} = 2$. Mean of 1000 runs.

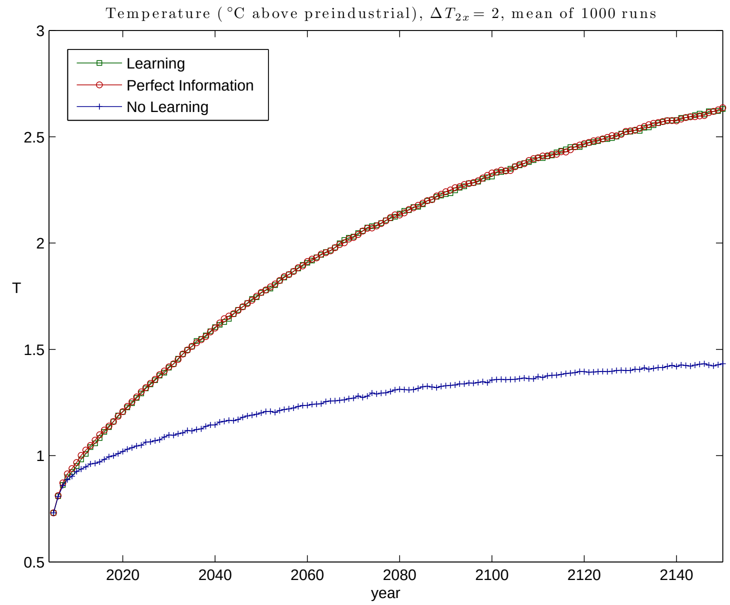

Figure 13: Temperature, degrees C above preindustrial, true $\Delta T_{2\times} = 2$. Mean of 1000 runs.

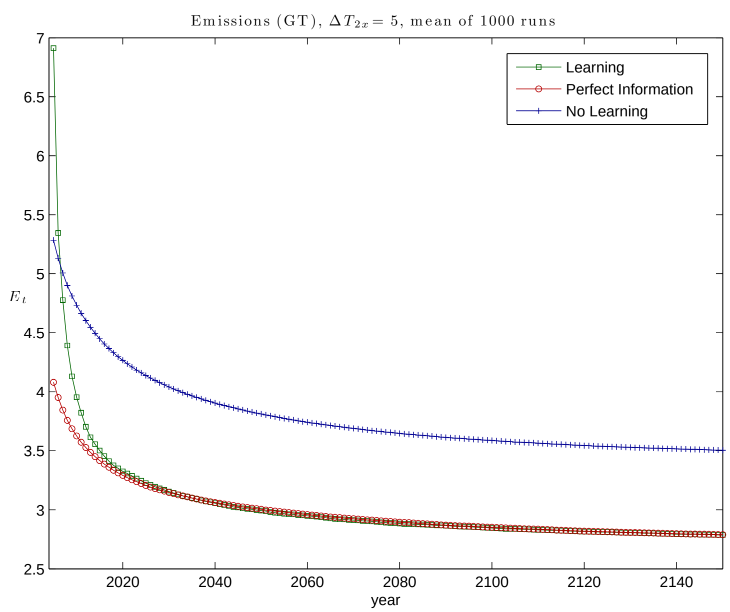

Figure 14: Optimal emissions policy (GT), true $\Delta T_{2\times} = 5$. Mean of 1000 runs.

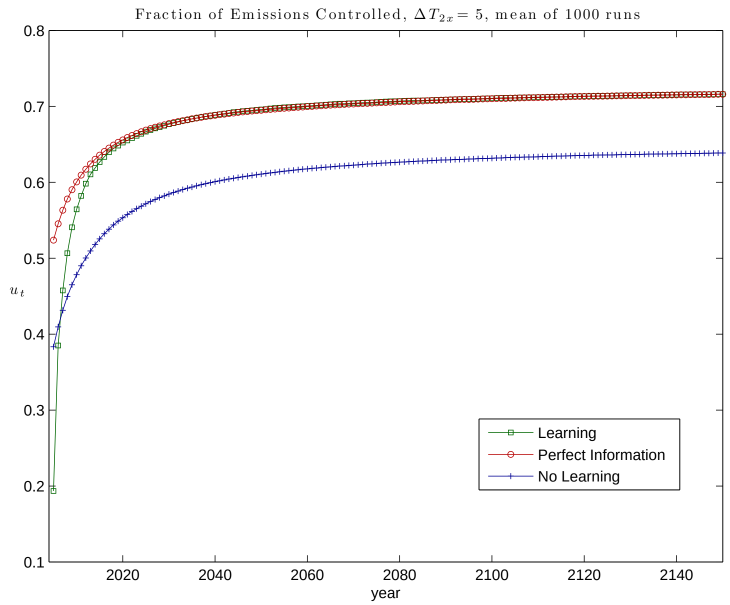

Figure 15: Optimal emissions control rate, true $\Delta T_{2\times} = 5$. Mean of 1000 runs.

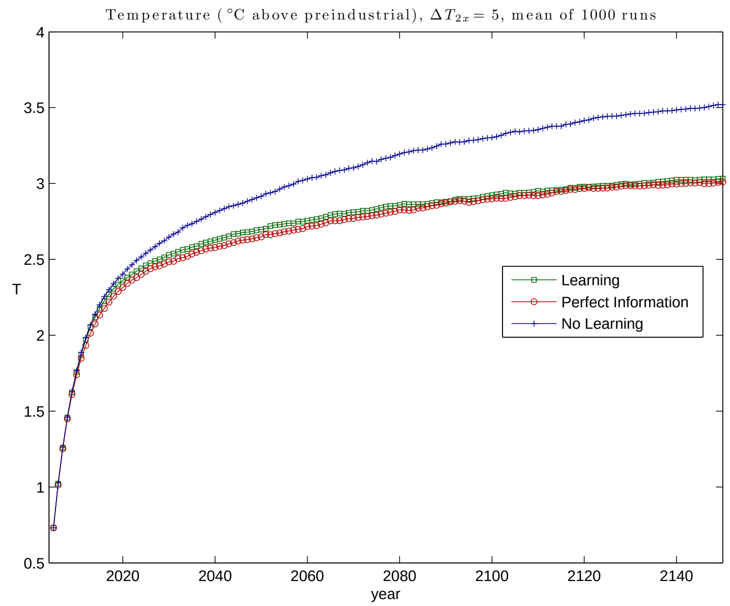

Figure 16: Temperature, degrees C above preindustrial, true $\Delta T_{2\times} = 5$. Mean of 1000 runs.

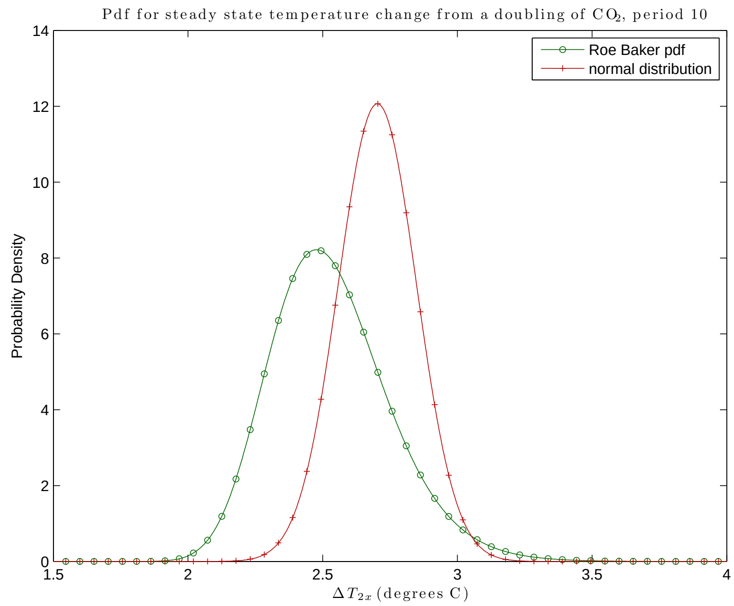

Figure 17: Posterior PDF of $\Delta T_{2\times}$, true value is 2.76 after 10 periods.

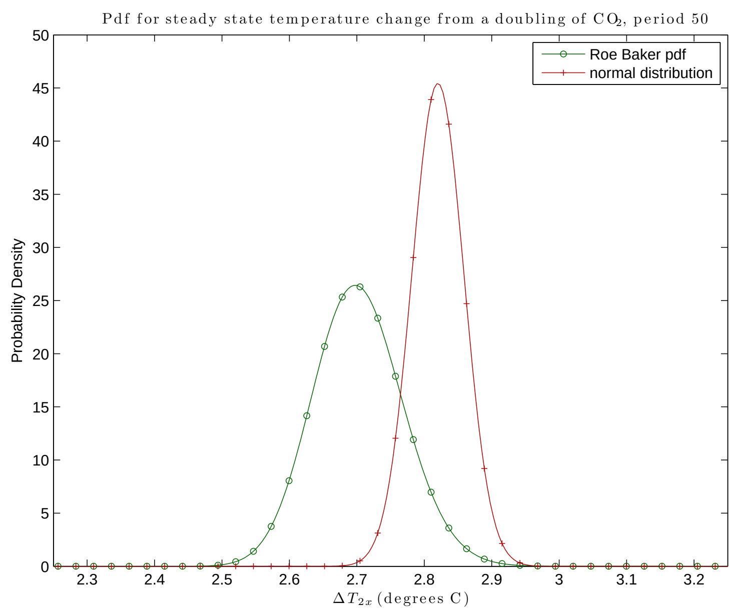

Figure 18: Posterior PDF of $\Delta T_{2\times}$, true value is 2.76 after 50 periods.

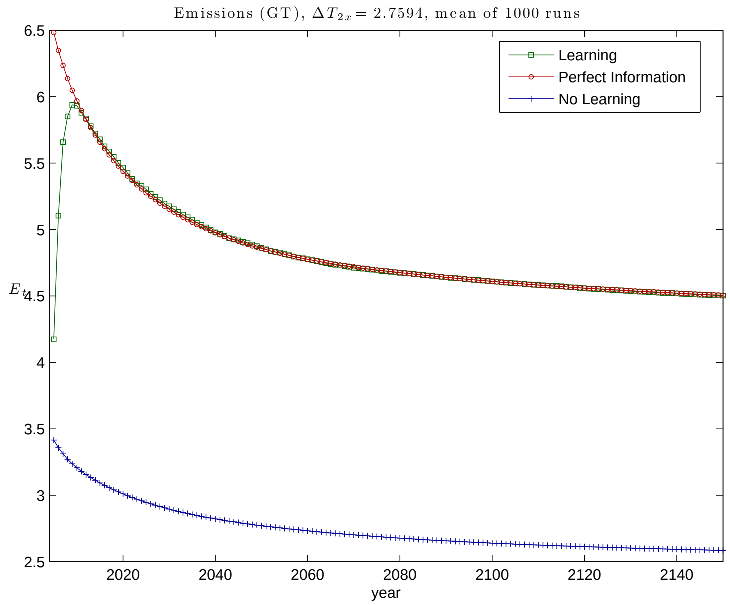

Figure 19: Optimal emissions policy in gigatons (GT), true $\Delta T_{2\times} = 2.76$, $\beta = 0.971$. Mean of 1000 runs.

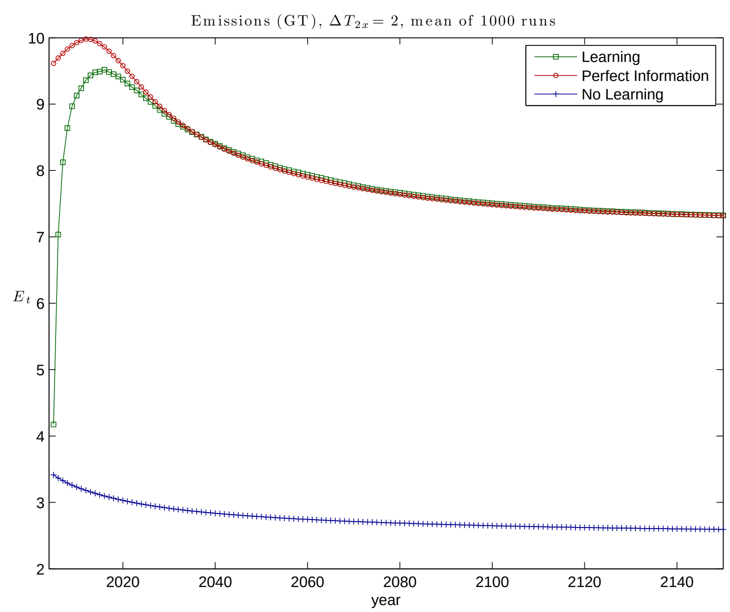

Figure 20: Optimal emissions policy in gigatons (GT), true $\Delta T_{2\times} = 2$, $\beta = 0.971$. Mean of 1000 runs.

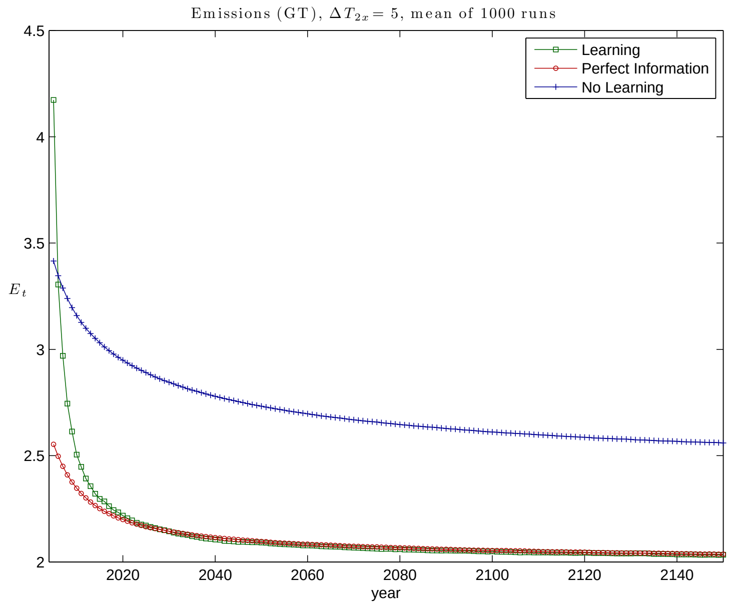

Figure 21: Optimal emissions policy in gigatons (GT), true $\Delta T_{2\times} = 5$, $\beta = 0.971$. Mean of 1000 runs.

References

Andronova, N., and M. Schlesinger, 2001, "Objective Estimation of the Probability Density Function for Climate Sensitivity," *Journal of Geophysical Research*, 106, 605–22.

Baker, M., and G. Roe, 2009, "The Shape of Things to Come: Why Is Climate Change So Predictable?," *Journal of Climate*, 22, 4574–89.

Bartz, S., and D. Kelly, 2008, "Economic Growth and the Environment: Theory and Facts," *Resource and Energy Economics*, 30, 115–49.

Cooley, T., and E. Prescott, 1995, "Economic Growth and Business Cycles," in Thomas Cooley (ed.), *Frontiers of Business Cycle Research* . chap. 1, pp. 1–38, Princeton University Press, Princeton, NJ.

Costello, C., M. Neubert, S. Polasky, and A. Solow, 2010, "Bounded Uncertainty and Climate Change Economics," *Proceedings of the National Academy of Sciences*, 107, 8108–10.

Forest, C., P. Stone, and A. Sokolov, 2005, "Estimated PDFs of Climate System Properties Including Natural and Anthropogenic Forcings," Discussion Paper Report 125, MIT Joint Program on the Science and Policy of Global Change.

Geweke, J., 2001, "A Note on Some Limitations of CRRA Utility," *Economic Letters*, 71, 341–5.

Intergovernmental Panel on Climate Change, 2007, "Climate Change 2007: The Physical Science Basis: Summary for Policy Makers, Contribution of Working Group I to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change," Discussion paper, IPCC Secretariat, World Meteorological Organization, Geneva, Switzerland, `http://www.ipcc.ch/ipccreports/ar4-wg1.htm`.

Keller, K., B. Bolker, and D. Bradford, 2004, "Uncertain Climate Thresholds and Optimal Economic Growth," *Journal of Environmental Economics and Management*, 48, 723–41.

Kelly, D. L., and C. Kolstad, 1999a, "Integrated Assessment Models for Climate Change Control," in T. Tietenberg, and H. Folmer (ed.), *International Yearbook of Environmental and Resource Economics 1999/2000: A Survey of Current Issues* . chap. 5, pp. 171–97, Edward Elgar, Cheltenham, UK.

Kelly, D. L., and C. D. Kolstad, 1999b, "Bayesian Learning, Pollution, and Growth," *Journal of Economic Dynamics and Control*, 23, 491–518.

Kelly, D. L., C. D. Kolstad, M. Schlesinger, and N. G. Andronova, 1998, "Learning About Climate Sensitivity From the Instrumental Temperature Record," University of Mimai.

Leach, A. J., 2007, "The Climate Change Learning Curve," *Journal of Economic Dynamics and Control*, 31, 1728–52.

Lemoine, D. M., 2010, "Climate Sensitivity Distributions Depend on the Possibility that Models Share Biases," *Journal of Climate*, 23, 4395–4415.

Lemoine, D. M., and C. P. Traeger, 2013, "Watch Your Step: Optimal Policy in a Tipping Climate," *American Economic Journal: Economic Policy*, forthcoming.

Millner, A., 2011, "On Welfare Frameworks and Catastrophic Climate Risks," University of California at Berkeley Working Paper.

Newbold, S., and A. Daigneault, 2009, "Climate Response Uncertainty and the Benefits of Greenhouse Gas Emissions Reductions," *Environmental and Resource Economics*, 44, 351–77.

Nicholls, R., R. Tol, and A. Vafeidis, 2008, "Global Estimates of the Impact of a Collapse of the West Antarctic Ice Sheet: an Application of FUND," *Climatic Change*, 91, 171–91.

Nordhaus, W., 2007, "A Review of the Stern Review on the Economics of Climate Change," *Journal of Economic Literature*, 45, 686–702.

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

Nordhaus, W., 2011, "The Economics of Tail Events with an Application to Climate Change," *Review of Environmental Economics and Policy*, 5, 240–57.

Pindyck, R., 2011, "Fat Tails, Thin Tails, and Climate Change Policy," *Review of Environmental Economics and Policy*, 5, 258–74.

Roe, G., and M. Baker, 2007, "Why is Climate Sensitivity so Unpredictable?," *Science*, 318, 629–32.

Schwartz, S., 2007, "Heat Capacity, Time Constraint, and Sensitivity of the Earth's Climate System," *Journal of Geophysical Research*, 112, D24S05.

Tol, R., 2009, "The Economic Effects of Climate Change," *Journal of Economic Perspectives*, 23, 29–51.

Traeger, C. P., 2012, "A 4-Stated DICE: Quantitatively Addressing Uncertainty Effects in Climate Change," Discussion Paper 1130, CUDARE Working Paper.

Weitzman, M., 2009, "On Modeling and Interpreting the Economics of Catastrophic Climate Change," *Review of Economics and Statistics*, 91, 1–19.

Weitzman, M., 2011, "Fat Tailed Uncertainty and the Economics of Catastrophic Climate Change," *Review of Environmental Economics and Policy*, 5, 275–92.