Probabilistic Stabilization Targets*

Luke G. Fitzpatrick
Department of Economics
University of Miami
luke.g.fitzpatrick@gmail.com

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

February 6, 2015

Abstract

We study stabilization targets: common environmental policy recommendations that specify a maximum probability of an environmental variable exceeding a fixed target (e.g. limit climate change to at most 2°C above preindustrial). Previous work generally considers stabilization targets under certainty equivalence. Using an integrated assessment model with uncertainty about the sensitivity of the temperature to greenhouse gas (GHG) concentrations (the climate sensitivity), learning, and random weather shocks, we calculate the optimal GHG emissions policy with and without stabilization targets. We characterize the range of feasible targets and show that in general, climate change has too much uncertainty and inertia to be controlled with the precision implied by stabilization targets.

We find that uncertainty exacerbates the welfare cost of stabilization targets. First, the targets are inflexible and do not adjust to new information about the climate system. Second, the target forces the emissions policy to overreact to transient shocks. These effects are present only in a model with uncertainty. Total welfare costs in the baseline model are 4.7%, which is 66% higher than the welfare cost under certainty.

*We would like to thank Larry Karp, Derek Lemoine, Christian Traeger and seminar participants at the University of California at Berkeley, the University of California at Davis, Fordham University, University of Miami, the Southern Economics Association Meetings, and the 14th Occasional Workshop on Environmental and Resource Economics, University of California at Santa Barbara for useful comments and suggestions.

1 Introduction

Stabilization targets are common environmental policy recommendations that specify the maximum allowable change of some environmental variable tied to pollution emissions. The most common application is in climate change policy. For example, many atmospheric scientists and policy makers recommend limiting greenhouse gas (GHG) concentrations so that the global mean temperature rises by at most 2°C above its preindustrial level. Stabilization targets call for limiting temperature changes to reduce the likelihood of irreversible and catastrophic climate change, or to prevent the temperature from changing more than an amount deemed harmful to society.[1] Once the target is specified, economists calculate the least-cost emissions path for which the temperature does not exceed the target under certainty.

However, random weather shocks and uncertain parameters, such as the sensitivity of the climate to GHG concentrations (the climate sensitivity), cause the target to be exceeded with positive probability. If the climate sensitivity is unexpectedly high, or following a series of unexpectedly high weather shocks, the temperature target may be exceeded regardless of emissions policy. Further, if the temperature increases still further above the target, the temperature target becomes infeasible. Learning reduces uncertainties over time, which allows some fine tuning of the emissions policy as the variable approaches the target. Nonetheless, slow learning and strong inertia in the climate makes fine tuning difficult. We show that even if GHG emissions are immediately and permanently reduced to zero, enough inertia exists in the climate that the global mean temperature will exceed 2°C with 15% probability.

An alternative is probabilistic stabilization targets (see for example, Held, Kriegler, Lessmann, and Edenhofer 2009), which instead require that the environmental variable be beneath the target with a given probability.[2] Probabilistic stabilization targets are feasible if the allowable probability of exceeding the target is sufficiently high. In general, however, we show that climate change has far too much uncertainty and inertia to be controlled to the degree assumed by probabilistic stabilization targets: uncertain parameters cause the temperature to drift to levels for which the target is either non-binding, or infeasible.

Using an integrated assessment model of the climate and economy with Bayesian learning and random weather shocks,[3] we calculate the welfare cost of stabilization targets. We

[1]Stabilization targets are also relatively straightforward to communicate to the public.

[2]One can think of a standard stabilization target as a special case of a probabilistic target where the probability of exceeding the target is set to zero. Therefore, we will henceforth without loss of generality consider only probabilistic targets.

[3]See Kelly and Kolstad (1999a) for a survey of integrated assessment models.

find that uncertainty exacerbates the welfare cost of stabilization targets. First, as new information arrives, the optimal temperature changes. For example, if the temperature turns out to be very sensitive to GHG concentrations, then achieving a given temperature requires more abatement expenditures. Since the cost of meeting a given temperature rises, and the benefit is unchanged, the optimal temperature rises. By definition, however, the target temperature does not change with a stabilization target. Therefore, a welfare cost ensues because the stabilization target is *inflexible*.

Second, we show that stabilization targets force overly stringent policy responses to transient shocks. For example, when a random weather fluctuation causes the environmental variable to exceed the target, stabilization expenditures must rise in response. However, the environmental variable naturally reverts back to the target as the shock dies out. Therefore, the stabilization target forces a costly policy response that provides only a relatively mild benefit of causing the environmental variable to return to the target more quickly than would occur naturally. Therefore, a second welfare cost occurs because stabilization targets cause an *overreaction* to transient shocks.

Both welfare losses are new to this paper and are present only when the temperature is a function of uncertain and/or random variables such as the climate sensitivity and weather shocks, and the model does not assume certainty equivalence. In models which assume certainty, a welfare cost only occurs if the target is set below the temperature resulting from the optimal emissions policy without the target. For example, Nordhaus (2007) calculates the welfare loss of restricting the temperature to 2°C under certainty, relative to an optimal temperature change of 3.5°. This result is sensitive to assumptions about the discount rate, the lack of threshold effects, and other assumptions of the Nordhaus DICE model. Indeed, we show that when uncertainty about the climate sensitivity is added to the DICE model, the optimal temperature without a target falls to 3.26°C,[4] which reduces the welfare loss associated with the 2°C target. In contrast, welfare losses associated with inflexibility and overreaction depend only on the interaction of uncertainty and a rigid target and are therefore robust to alternative modeling assumptions.

We show that stabilization targets cause a welfare loss of 4.7% given the baseline stabilization policy, which is 66% higher than the welfare cost under certainty. We find that most of the increase in welfare loss comes from the inflexibility of the target, but overreaction accounts for 8.2% of the increase in welfare loss. Further, the total welfare loss is relatively insensitive to the policy parameter which specifies the maximum probability of exceeding

[4]With uncertainty, the risk averse planner becomes more conservative, reducing emissions which are more damaging if the climate sensitivity turns out to be higher than expected.

the target.

There is a long history in the environmental economics literature of evaluating the welfare costs of sub-optimal policies. Most of the literature either compares market with inefficient, non-market based regulation (see for example, Stavins 1993), or evaluates whether or not a given sub-optimal policy improves welfare relative to a baseline without the policy (see for example, Portney, Parray, Gruenspecht, and Harrington 2003). In contrast, the welfare losses we focus on arise solely due to uncertainty. Indeed, under certainty one could always choose a temperature target high enough so that no welfare loss ensues. In our model with uncertainty, inflexibility and overreaction cause a welfare loss even if the target equals the average optimal temperature.

Stabilization targets are ubiquitous in climate change policy. Policy makers recommending a 2°C stabilization target include the European Commission (2007), the Copenhagen Accord (2009), and the German Advisory Council on Global Change (Schubert et al. 2006). Many atmospheric scientists (Hansen 2005, O'Neill and Oppenheimer 2002) also advocate for the 2°C limit; Hansen in particular is a vocal advocate. Other climate-related stabilization targets are also common. For example, the German Advisory council recommends limiting sea level rise to at most 1 meter and ocean acidification to at most 0.2 units of pH below its preindustrial level.[5]

Economists (Nordhaus 2007, Richels, Manne, and Wigley 2004, Lemoine and Rudik 2014) then compute the least cost GHG emissions path which stabilizes the climate at 2°C under certainty. However, it is well known that parameters of the climate system are uncertain. For example, the climate sensitivity, which measures the elasticity of the global mean temperature with respect to GHG concentrations, is notoriously uncertain (Intergovernmental Panel on Climate Change 2007, Kelly and Kolstad 1999b). Therefore, following the least cost pathway calculated under certainty can, under uncertainty actually result in warming that exceeds the target by a considerable margin.[6] Indeed, a branch of the literature focuses on the likelihood of meeting current targets for various emissions scenarios proposed by policy makers, or what emissions paths satisfy the target for various values of the climate sensitivity.

[5] Some authors refer to a policy which stabilizes GHG concentrations at a particular level as a stabilization target (see for example den Elzen, Meinshausen, and van Vuuren 2007, Lemoine and Rudik 2014). Policy groups such as 350.org also favor a target GHG concentration. Since carbon cycle uncertainties are smaller than uncertainty in climate models, our results are less applicable to GHG concentration targets. Nonetheless, our results make clear that a fixed GHG target is not optimal because optimal GHG concentrations change with the resolution of uncertainty.

[6] Paradoxically, stabilization targets evolved as method of dealing with uncertainty in integrated assessment models. The idea was to propose limits on temperature changes which, if exceeded, would cause damages high enough so that uncertainties in the cost of abatement and other parameters are less relevant.

For example, Hare and Meinshausen (2006) and Keppo, O'Neill, and Riahi (2007) compute temperature changes for various emissions scenarios; Harvey (2007) proposes allowable CO_2 emissions paths for different ranges of the climate sensitivity. One robust result from the aforementioned studies is that emission paths following the upper boundary of emissions scenarios are less likely to meet targets than those which follow the lower boundary. The above research provides an important first step in estimating the range of feasible probabilities of exceeding the target, given current information. Here we take the next step and consider stabilization targets when learning allows the emissions policy to adjust as new information arrives.

Implementing a stabilization target with uncertainty is non-trivial. For any given emissions policy time path, a possibility exists such that the climate sensitivity is high enough so that the temperature eventually exceeds the target with arbitrarily high probability. Therefore, to ensure feasibility, the optimization problem must pre-specify an emissions policy when the target cannot be met with the required probability. In turn, the stringency of that pre-specified emissions policy affects near term optimal emissions policy: the more stringent and costly is the pre-pecified policy, the more cautious the planner is when setting near term optimal emissions. A stringent pre-specified emissions policy is most consistent with the idea of stabilization targets. If the target is exceeded, then presumably the policy should be to set a very low or zero emissions level until the target is feasible. To avoid the high cost of zero emissions, the planner sets optimal near term emissions conservatively to reduce the chance of exceeding the target. Indeed, we show that controlling temperature to the degree implied by a stabilization target often implies a very high near term emissions abatement rate. For example, we find that if the climate sensitivity is such that a doubling of GHGs causes a 3.9°C temperature change, rather than the prior estimate of 3.08°C, the planner raises the abatement rate from 35% to 71% as early as 2015.[7]

Uncertainty makes controlling the temperature more difficult. Learning, by reducing uncertainty, allows the planner to more easily meet a target by quickly adjusting emissions if new information indicates the climate sensitivity is different than previously thought. Therefore, learning allows the planner to move closer to the target, and still remain below the target with the same probability. However, Kelly and Kolstad (1999b), Leach (2007), and Roe and Baker (2007) show that learning about the climate sensitivity is a slow process, because noisy weather fluctuations obscure the climate change signal. Further, small uncertainties about the climate sensitivity ultimately have large effects on the temperature

[7]The initial year is set to 2005 for consistency with the literature.

through feedback effects. Therefore, controlling the temperature requires a precise estimate of the climate sensitivity, which takes time. Therefore, we find the optimal near term policy is similar to the case without learning. Our model has Bayesian learning about the climate sensitivity, and random weather fluctuations. We find that learning moves the optimal target under uncertainty closer to the target under certainty, but the effect is marginal since learning is slow.

Our model considers targets when the climate sensitivity is unknown. Rudik (2014) considers targets with damage uncertainty and learning. He shows that if the learning model is misspecified, targets become attractive since they prevent the planner from increasing the temperature when believing incorrectly that damages are not very convex. Here we show stabilization targets more problematic when the climate sensitivity is uncertain. The relationship between the policy variable, emissions, and the target, temperature, is unknown and nonlinear, which makes the target difficult to achieve.

Lorenz, Schmidt, Kriegler, and Held (2012), Schmidt, Lorenz, Held, and Kriegler (2011), and Neubersch, Held, and Otto (2014) debate the magnitude of the value of information which allows more precise control of the climate with probabilistic stabilization targets. These studies use frameworks where learning occurs once and emissions paths are adjusted once and focus on the value of information. Our focus is instead on the welfare costs of stabilization targets. For this purpose, we use Bayesian learning, which is endogenous and incremental: beliefs about unknown parameters and decisions based on beliefs adjust each period as new information arrives.[8] Bayesian learning allows decisions to adjust to new information each period, which allows some smoothing of the cost of changes in abatement expenditures over time. However, learning is often considerably slower than assumed by studies with exogenous learning.

Our integrated assessment model computes both the cost of emissions abatement and the damages from higher temperatures. Therefore, we can calculate the welfare cost of a stabilization target. The mean welfare cost ranges from 2.85%, when the maximum probability of exceeding the target is 0.975, to 5.2%, when the maximum probability of exceeding the target is 0.2. Neubersch, Held, and Otto (2014) calibrate a maximum probability of 0.33, based on an interpretation of IPCC statements calling for policies for which achieving the 2° target is likely. Using this value as a baseline policy recommendation, the welfare loss is 4.7%. The welfare loss is not very sensitive to the choice of the maximum probability of exceeding the target. The probability of exceeding the target in any period depends mostly

[8]Other examples of Bayesian learning in climate models include Kelly and Kolstad (1999b), Leach (2007), and Lemoine and Traeger (2014).

on factors such as the resolution of the uncertainty and the climate inertia, and is affected very little by the current emissions. Therefore, the probability of exceeding the target tends to be either zero or one and the maximum allowable probability has little effect.

The welfare loss increases with the unknown true value of the climate sensitivity. If the climate sensitivity turns out to be higher than expected, the abatement cost of keeping the temperature at the target increases, which increases the optimal temperature. We show that the welfare loss can increase to 14% or more depending on true value of the climate sensitivity.

Other authors compute optimal emissions paths under certainty which keep temperatures below a threshold, beyond which specific irreversible and disastrous consequences occur. Keller et al. (2005) propose emissions paths which prevent coral bleaching or the disintegration of ice sheets. Kvale et al. (2012) propose emission paths which limit ocean level rise and acidification. Additionally, Bruckner and Zickfield (2009) compute emission paths that reduce the likelihood of a collapse of the Atlantic thermohaline circulation. The aforementioned studies employ the tolerable windows approach, an inverse modeling method that asks: in order to limit GHG concentrations or warming below a threshold at all future dates, how should emissions be controlled in every period moving forward? These studies have some similarities to stabilization targets with certainty in that the optimal policy stabilizes the climate below the threshold. Lemoine and Traeger (2014) compute the optimal emissions path in a model with various climate tipping points in an environment with uncertainty and learning, but do not study stabilization targets. While we do not specifically model irreversibilities, our model combines convex damages with uncertainty. Therefore, the planner insures against very high damages by pursuing a conservative emissions policy. Thus, the literature with irreversibilities and thresholds consider either certainty, or do not consider stabilization targets. We show that the stabilization targets are fundamentally different with uncertainty: control of the climate becomes difficult and additional welfare losses occur.

2 Model

We consider an infinite horizon version of the Nordhaus DICE model (Nordhaus 2007). In the DICE model, economic production causes GHG emissions, which raise the global mean temperature. Higher temperatures reduce total factor productivity (TFP). The social planner chooses capital investment and an emissions abatement rate to maximize welfare. Our model has four differences from the DICE model. First, we use an annual time step rather than the 10 year step in DICE. Second, we use the simplified model of the atmosphere/climate due to Traeger (2014), in which the ocean temperature changes exogenously and GHG concentrations immediately mix evenly in the atmosphere. Third, the model is stochastic, with an uncertain climate sensitivity and random weather shocks that obscure the effect of GHGs on temperature. The planner learns about the uncertain climate sensitivity over time by observing temperature changes. Fourth, we impose stabilization targets to ascertain their effects on welfare, temperature, and economic growth. Sections 2.1-2.2 describe the economic and climate models briefly (refer to Traeger 2014, for a detailed discussion).

2.1 Economic system

The global economy produces gross output, Q, from capital K and labor L according to:

$$Q_t = A(t) K_t^{\gamma} L(t)^{1-\gamma}. \tag{1}$$

Here variables denoted as a function of t, such as $L(t)$ and TFP, $A(t)$, grow exogenously. Appendix A.1 gives the growth rates for all variables which change exogenously over time. Variables with a t subscript are endogenous.

An emissions abatement technology exists which can reduce emissions by a fraction x_t at a cost of $\Lambda(x_t) = \Psi(t) x_t^{a_2}$ fraction of gross output. Here Λ is the cost function and $\Psi(t)$ is the exogenously declining cost of a backstop technology which reduces emissions to zero. Further, increases in global mean temperatures above preindustrial levels, T_t, reduce TFP by a factor $1/(1 + D(T_t))$, where $D(T_t) = b_1 T^{b_2}$ is the damage function. Therefore, output net of abatement spending and climate damages, Y_t, is:

$$Y_t = \frac{1 - \Psi(t) x_t^{a_2}}{1 + b_1 T_t^{b_2}} A(t) K_t^{\gamma} L(t)^{1-\gamma}. \tag{2}$$

Let C_t be consumption and let capital depreciate at rate δ_k. Then the resource constraint is:

$$Y_t = C_t + K_{t+1} - (1 - \delta_k) K_t. \tag{3}$$

Period utility is constant relative risk aversion:

$$u = \frac{(C_t/L_t)^{1-\phi} - 1}{1 - \phi}. \tag{4}$$

We follow Costello, Neubert, Polasky, and Solow (2010) and assume $T_t \leq T^{max} << \infty$ to prevent expected utility from being unbounded from below.

The discount factor for future utility is $\exp(-\delta_u)$, where δ_u is the pure rate of time preference.

2.2 Climate System

Current period GHG emissions, E_t, from production depend on the planner's choice of emissions abatement rate x_t, the emissions intensity of output $\sigma(t)$, exogenous emissions from land use changes, $B(t)$, and gross global output:

$$E_t = (1 - x_t)\,\sigma(t)\,Q_t + B(t)\,. \tag{5}$$

The stock of GHG equivalents, M_t, depends on current period emissions and the natural decay rate of GHGs into the biosphere and ocean. Let $\delta_m(t)$ denote the decay rate (which changes exogenously) and M_B denote the stock of GHGs during pre-industrial times. Then M_t accumulates according to:

$$M_{t+1} - M_B = (1 - \delta_m(t))\,(M_t - M_B) + E_t. \tag{6}$$

We normalize GHG stocks relative to pre-industrial. Let $m_t \equiv \frac{M_t}{M_B}$, then:

$$m_{t+1} - 1 = (1 - \delta_m(t))\,(m_t - 1) + \frac{E_t}{M_B}. \tag{7}$$

Radiative forcing of GHGs, F_t, increases the temperature:

$$F_{t+1} = \Omega \log_2(m_{t+1}) + EF(t)\,. \tag{8}$$

Here $EF(t)$ is forcing from other sources, which grows exogenously.

The global mean temperature evolves according to:

$$\hat{T}_{t+1} = \hat{T}_t + \frac{1}{\alpha}\left(F_{t+1} - \frac{\hat{T}_t - \Gamma}{\tilde{\lambda}} + \xi\left(\hat{O} - \hat{T}\right)(t) - \Gamma\right) + \tilde{\nu}_{t+1} \tag{9}$$

Here $\hat{T}$ and $\hat{O}$ denote the absolute global atmospheric and oceanic temperatures in °C, respectively; α is the thermal capacity of the upper oceans; Γ is the pre-industrial atmospheric temperature; ξ is a coefficient of heat transfer from the upper oceans to the atmosphere; $\tilde{\nu}_t \sim N(0, 1/\rho)$ is the random weather shock; $\tilde{\lambda}$ is the uncertain climate sensitivity. We assume the ocean-atmosphere temperature differential changes exogenously. The climate sensitivity $\tilde{\lambda}$ describes how sensitive the atmospheric temperature is to GHG concentrations, and is the subject of great uncertainty.

Let $\Delta T_{2\times}$ be the steady state atmospheric temperature deviation from pre-industrial time resulting from a doubling of GHG concentrations, also relative to pre-industrial levels. Then:

$$\Delta T_{2\times} = \Omega\tilde{\lambda}. \tag{10}$$

Since the climate sensitivity is uncertain, $\Delta T_{2\times}$ is also uncertain. Stocker, Dahe, and Plattner (2013) estimate that $\Delta T_{2\times}$ is most likely to lie somewhere between 1.5°C and 4.5°C. The initial mean of the prior distribution is 3.08, taken from the mean of estimates in the atmospheric science literature (Roe and Baker 2007).[9]

Let $T_t = \hat{T}_t - \Gamma$ and $O_t = \hat{O}_t - \Gamma$ be the current deviations from pre-industrial temperatures, $\tilde{\beta}_1 = 1 - 1/\tilde{\lambda}\alpha$ denote the climate feedback parameter, $\beta_2 = \frac{1}{\alpha}$, and $\beta_3 = \xi/\alpha$. The climate system simplifies to:

$$T_{t+1} = \tilde{\beta}_1 T_t + \beta_2 F_{t+1} + \beta_3 (O - T)(t) + \tilde{\nu}_{t+1}. \tag{11}$$

Since $\tilde{\lambda}$ is uncertain, the climate feedback parameter is also uncertain. The climate feedback parameter is increasing in $\tilde{\lambda}$. For example, if GHG induced warming reduces ice cover, which reduces the amount of sunlight reflected back into space (the albedo effect), causing still higher temperatures, we have a positive feedback (higher $\tilde{\beta}_1$ and $\tilde{\lambda}$).[10] Uncertain climate feedbacks cause uncertainty in the climate sensitivity (Stocker, Dahe, and Plattner 2013).

2.3 Learning

Assume the planner has prior beliefs that the climate feedback parameter is drawn from a normal distribution with mean μ_t and precision η_t. The weather shock ν occurs at the beginning of each period, before the abatement rate is chosen. We combine the two uncertain terms in equation (11) and denote the sum $\tilde{H}$:

$$\tilde{H}_{t+1} = \tilde{\beta}_1 T_t + \tilde{\nu}_{t+1}. \tag{12}$$

Since $\tilde{H}_{t+1}$ is the sum of two normally distributed random variables, it is also normally distributed with mean $\mu_t T_t$ and variance $\sigma^2_{\tilde{H}} = T_t^2/\eta_t + 1/\rho$. The planner observes $H_{t+1} = T_{t+1} - \beta_2 F_{t+1} - \beta_3 (O - T)(t)$ at the beginning of $t + 1$ and updates beliefs of $\tilde{\beta}_1$. Bayes' Rule implies that the posterior distribution of $\tilde{\beta}_1$ is also normal, with mean and precision:

$$\mu_{t+1} = \frac{\eta_t \mu_t + \rho H_{t+1} T_t}{\eta_t + \rho T_t^2}, \tag{13}$$

[9]More precisely, the initial mean of the climate feedback parameter, defined below, is identical to Roe and Baker (2007). Our implied $\Delta T_{2\times}$ is slightly lower because of small differences in the other parameters of the model.

[10]Other feedbacks include changes in cloud cover and water vapor.

$$\eta_{t+1} = \eta_t + \rho T_t^2. \tag{14}$$

Perfect information implies that $\mu = \tilde{\beta}_1$ and $\eta = \infty$. The information set used by the planner to select x_t includes μ_t.

2.4 Recursive Problem

The planner chooses emission abatement each period to maximize social welfare.

$$W = \max_{k_{t+1}, x_t} \mathrm{E}\left[\sum_{t=0}^{\infty} \exp\left(-\delta_u t\right) L_t u\left(\frac{C_t}{L_t}\right)\right]. \tag{15}$$

Let $k \equiv K/\left(LA^{\frac{1}{1-\gamma}}\right)$ denote normalized capital per productivity adjusted person (Kelly and Kolstad 2001, Traeger 2014) and the same for y and c, and $s = [k, T, m, t, \mu, \eta]$. The recursive version of the social planning problem is:

$$V(s) = \max_{k', x \in [0,1]} \left\{ u(c) + \beta(t) \int_{-\infty}^{\infty} V[s'] \, N\left(\mu T, \frac{T^2}{\eta_t} + \frac{1}{\rho}\right) d\tilde{H} \right\}, \tag{16}$$

subject to:

$$C = y - \exp\left(-g_A(t) - g_L(t)\right) k' + \left(1 - \delta_k\right) k. \tag{17}$$

$$T' = \beta_2 F' + \beta_3 \left(O - T\right)(t) + \tilde{H}', \tag{18}$$

$$F' = \Omega \log_2\left(m'\right) + EF(t), \tag{19}$$

$$m' = 1 + \left(1 - \delta_m(t)\right)(m - 1) + \frac{E}{M_B}, \tag{20}$$

$$E = (1 - x)\,\sigma(t)\, A(t)^{\frac{1}{1-\gamma}} L(t)\, k^{\gamma} + B(t). \tag{21}$$

$$\mu' = \frac{\eta\mu + \rho \tilde{H}' T}{\eta + \rho T^2}, \tag{22}$$

$$\eta' = \eta + \rho T^2. \tag{23}$$

$$t' = t + 1. \tag{24}$$

Equation (16) condenses the double expectation over $\tilde{\beta}_1$ and $\tilde{\nu}_{t+1}$ into one expectation over the random variable H_{t+1}. Appendix A.1 gives the equations which govern the evolution of variables that change exogenously over time, including the growth rates of productivity and population, g_A and g_L. Therefore time, t, is a state variable. The discount factor accounts for growth in population and productivity. Because the growth rates change over time, the normalized discount factor $\beta(t)$ is not constant, but is exogenous.

In the model, two state variables, t and η, are non-stationary. Therefore, the computational solution replaces the precision η with the variance $1/\eta$, and replaces time with a bounded, monotonic increasing function.[11] Table 1 gives parameter values and definitions for the above problem.

2.5 Stabilization Targets

The probabilistic stabilization target is a constraint on temperature which satisfies:

$$Pr\left(T_{t+1} \geq T^*\right) \leq \omega, \ \forall t. \tag{25}$$

A pure stabilization target is a special case of equation (25) with $\omega = 0$, if $\omega = 1$ the constraint is always satisfied. A constraint exists for each period, $t = 0, 1, \ldots$, which restricts the probability that $T_{t+1} \geq T^*$.

The stabilization target is effectively a constraint on the abatement rate x, since x affects T_{t+1}. Rewriting the left hand side of (25) gives:

$$Pr\left(T_{t+1} \geq T^*\right) = Pr\left(\beta_2 F_{t+1} + \beta_3 \left(O - T\right)(t) + \tilde{H}_{t+1} \geq T^*\right), \tag{26}$$

$$= Pr\left(\tilde{H}_{t+1} \geq T^* - \beta_2 F_{t+1} - \beta_3 \left(O - T\right)(t)\right), \tag{27}$$

$$= 1 - \text{NCDF}\left[\frac{\left(T^* - \beta_2 F_{t+1} - \beta_3 \left(O - T\right)(t) - \mu_t T_t\right)}{\sqrt{\frac{T_t^2}{\eta_t} + \frac{1}{\rho}}}\right]. \tag{28}$$

Here NCDF is the cumulative distribution function of the standard normal distribution. Let:

$$P_{t+1} \equiv \beta_2 F_{t+1} + \beta_3 \left(O - T\right)(t) + \mu_t T_t, \ \text{and} \tag{29}$$

$$\sigma_{H,t} \equiv \sqrt{\frac{T_t^2}{\eta_t} + \frac{1}{\rho}} \tag{30}$$

be the mean and standard deviation of T_{t+1}, respectively, then:

$$Pr\left(T_{t+1} \geq T^*\right) = 1 - \text{NCDF}\left[\frac{T^* - P_{t+1}}{\sigma_{H,t}}\right]. \tag{31}$$

[11] See for example, Kelly and Kolstad (1999b) or Traeger (2014).

Constraint (25) is therefore equivalent to:

$$\mathrm{NCDF}\left[\frac{T^{*}-P_{t+1}}{\sigma_{H,t}}\right]\geq 1-w, \tag{32}$$

$$\frac{T^{*}-P_{t+1}}{\sigma_{H,t}}\geq \mathrm{NCDF}^{-1}\left[1-w\right], \tag{33}$$

$$F_{t+1}\leq \frac{1}{\beta_2}\left(T^{*}-\beta_3\left(O-T\right)(t)-\mu_t T_t-\sigma_{H,t}\cdot \mathrm{NCDF}^{-1}\left[1-w\right]\right), \tag{34}$$

$$x_t\geq 1+\frac{M_B}{\sigma(t)\,L(t)\,A(t)^{\frac{1}{1-\gamma_t}}k_t^{\gamma}}\Bigg(1+\left(1-\delta_m(t)\right)\left(m_t-1\right)- \exp\left\{\frac{\log(2)}{\Omega\beta_2}\left[T^{*}-\beta_3\left(O-T\right)(t)-\mu_t T_t-\sigma_{H,t}\cdot \mathrm{NCDF}^{-1}\left[1-w\right]\right]\right\}\Bigg). \tag{35}$$

$$x_t\geq PC\left(s_t,T^{*},\omega\right). \tag{36}$$

Here PC is the right hand side of (35). A stabilization target is therefore equivalent to a minimum abatement rate.

Let θ denote the Lagrange multiplier on the probabilistic constraint. The recursive version of the problem, which includes the probabilistic constraint, is then:

$$V(s)=\max_{k',x\in[0,1]}\left\{u(c)+\theta\left[x-PC\left(s,T^{*},\omega\right)\right]+\beta(t)\int_{-\infty}^{\infty}V\left[s'\right]N\left(\mu T,\sigma_H\right)d\tilde{H}\right\}, \tag{37}$$

subject to equations (17)-(24). In period t, the planner anticipates facing constraints in periods $t+i$, which restrict the probability that $T_{t+i+1}\geq T^{*}$ for all $i=1,2,\ldots$. Therefore, constraints in period t on the probability that $T_{t+i+1}\geq T^{*}$ are implicit in V, since the planner anticipates in period t facing a probabilistic constraint $t+i$.

3 Feasibility

3.1 A Feasible Constraint

Assuming emissions cannot be negative, constraint (36) may not be feasible; the set of $x\in[0,1]$ which also satisfy (36) may be empty, violating a necessary condition for the existence of a maximum. The problem is not feasible when $PC\left(s,T^{*},\omega\right)>1$, since the maximum abatement rate is one. This occurs when the temperature rises close to or above T^*. In this case, given the inertia of the climate, even an abatement rate of 1 cannot reduce the temperature enough to satisfy the constraint. Feasibility is also affected by ω: as $\omega \to 1$, the planner is allowed to exceed the target with high probability, and so the model has a feasible solution even if the temperature is relatively high.

To solve the infeasibility problem while keeping with the spirit of a stabilization target, we assume that $x = 1$ is always feasible.[12] From today's perspective, we pre-specify a policy of $x = 1$ whenever the target is infeasible. Let θ_t denote the Lagrange multiplier for the probabilistic constraint in t, then the constrained problem is:

$$V(s) = \max_{k', x \in [0,1]} \left\{ u(c) + \theta \left[x - \min \left\{ PC\left(s, T^*, \omega\right), 1 \right\} \right] + \beta(t) \int_{-\infty}^{\infty} V\left[s'\right] N\left(\mu T, \sigma_H\right) d\tilde{H} \right\}, \tag{38}$$

subject to equations (17)-(24).

Problem (38) always has a feasible solution. Further, when the probabilistic constraint is exceeded due to a large random weather shock or an unexpectedly high realization of β_1, the planner must return as quickly as possible (by setting $x = 1$) to the range where $PC < 1$. This is in keeping with the idea that exceeding the target causes damages and should be avoided if possible. The planner anticipates future constraints in (38), through the future value $V[s']$. For example, a future value s' that has a relatively high T' has a lower value in the constrained problem, because the planner anticipates having to set $x = 1$ in the future to satisfy the constraint. The planner faces a tradeoff: choose a relatively low emissions abatement rate and take the risk that the temperature exceeds the target and be forced to set $x = 1$, or raise the current abatement rate as a precaution. Since abatement costs are convex, the planner prefers to smooth abatement costs over time and therefore raises current abatement somewhat as a precaution. However, given the normal distribution, the planner cannot eliminate all risk (even with zero emissions), and so the computational question is

[12] Other options are possible, but less attractive for welfare analysis. Ignoring infeasible constraints is not attractive because the planner would have an incentive to push the climate change as close as possible to T^* in the hope that the climate will go over the target so that the constraint can be ignored. The other option would be to include a penalty function for going over T^* (Neubersch, Held, and Otto 2014). However, a penalty for high temperatures already exists (the damage function), so it is unclear from a welfare perspective what the penalty represents. In contrast, here exceeding the constraint causes two penalties: first damages increase, and second the planner must incur the cost of returning as quickly as possible back to the target. Both penalties are therefore implicitly consistent with the welfare costs and damages in the model.

how much extra risk is optimal.

3.2 Tightness of the Constraint

Here we calculate probabilities in period $t+i$, $\omega_{t+i}^{\min}$ and $\omega_{t+i}^{\max}$, which are the probabilities of exceeding the target when the abatement rate is set to 1 and 0, respectively, for all i, conditional on current information. Values of ω between $\omega_{t+i}^{\min}$ and $\omega_{t+i}^{\max}$ are feasible policies from periods t to $t+i$. For $\omega > \omega_{t+i}^{\max}$, the constraint is non-binding at period $t+i$, since the planner can set $x_{t+i} = 0$ for all periods up to $t+i$ and still expect to satisfy the constraint. Conversely, any value of $\omega < \omega_{t+i}^{\min}$ implies $PC\left(s_{t+i}\right) > 1$, even if the abatement rate is set to one immediately for all periods up to $t+i$. By doing this exercise we glean intuition as to how tight the constraint is as a function of ω, which helps to explain the results in the next section. The tightness of the constraint may equivalently be controlled by altering T^*, but we assume here T^* is a given policy.

The highest possible probability for which $T_{t+1} \geq T^*$ occurs with a zero abatement rate. In this case, we have from (20):

$$E_t^{\max} = \sigma\left(t\right) A\left(t\right)^{\frac{1}{1-\gamma}} L\left(t\right) k_t^{\gamma} + B\left(t\right), \tag{39}$$

$$m_{t+1}^{\max} - 1 = \left(1 - \delta_m\left(t\right)\right)\left(m_t - 1\right) + \frac{E_t^{\max}}{M_B}, \tag{40}$$

$$F_{t+1}^{\max} = \Omega \log_2\left(m_{t+1}^{\max}\right) + EF\left(t\right). \tag{41}$$

Next, from (18):

$$T_{t+1}^{\max} = \tilde{H}_{t+1} + \beta_2 F_{t+1}^{\max} + \beta_3\left(O - T\right)\left(t\right), \tag{42}$$

We can then calculate $\omega^{\max}$ as:

$$\omega_{t+1}^{\max} = Pr\left(T_{t+1}^{\max} \geq T^*\right), \tag{43}$$

$$\omega_{t+1}^{\max} = Pr\left\{\tilde{H}_{t+1} \geq T^* - \beta_2 F_{t+1}^{\max} - \beta_3\left(O - T\right)\left(t\right)\right\}, \tag{44}$$

$$\omega_{t+1}^{\max} = 1 - \text{NCDF}\left(\frac{T^* - P_{t+1}^{\max}}{\sigma_{H,t}}\right), \text{ where} \tag{45}$$

$$P_{t+1}^{\max} \equiv \mu_t T_t + \beta_2 F_{t+1}^{\max} + \beta_3\left(O - T\right)\left(t\right). \tag{46}$$

Since the difference between the target and the expected temperature in period $t+1$ under maximum emissions is not infinite, we have immediately from (45) that $\omega^{\max} < 1$.

Analogously, the minimum probability of exceeding the constraint occurs when the abatement rate is one.

$$E_t^{\min} = B(t), \tag{47}$$

$$m_{t+1}^{\min} - 1 = (1 - \delta_m(t))(m_t - 1) + \frac{E_t^{\min}}{M_B}, \tag{48}$$

$$F_{t+1}^{\min} = \Omega \log_2\left(m_{t+1}^{\min}\right) + EF(t), \tag{49}$$

$$T_{t+1}^{\min} = \tilde{H}_{t+1} + \beta_2 F_{t+1}^{\min} + \beta_3 (O - T)(t), \tag{50}$$

$$\omega_{t+1}^{\min} = 1 - \text{NCDF}\left(\frac{T^* - P_{t+1}^{\min}}{\sigma_{H,t}}\right), \text{ where} \tag{51}$$

$$P_{t+1}^{\min} \equiv \mu_t T_t + \beta_2 F_{t+1}^{\min} + \beta_3 (O - T)(t). \tag{52}$$

From (45) and (51), current emissions have only a small effect on the probability of exceeding the target. First, current emissions are small relative to concentrations that have built up over centuries. Second, the forcing equation is logarithmic, further limiting the effect of current emissions on the current temperature. Indeed, near term temperatures are largely a function of inertia in the climate and total GHG concentrations.

The set of probabilities achievable with an interior abatement rate expands over time, since the planner can lower future temperatures via a sustained reduction in emissions. On the other hand, the uncertain climate feedback parameter has a multiplicative effect over time. Therefore, future temperatures are more uncertain and therefore are more difficult to control. Figure 1 plots the mean of 10,000 simulated values of $\omega_t^{\min}$ and $\omega_t^{\max}$ for the given parameter values.

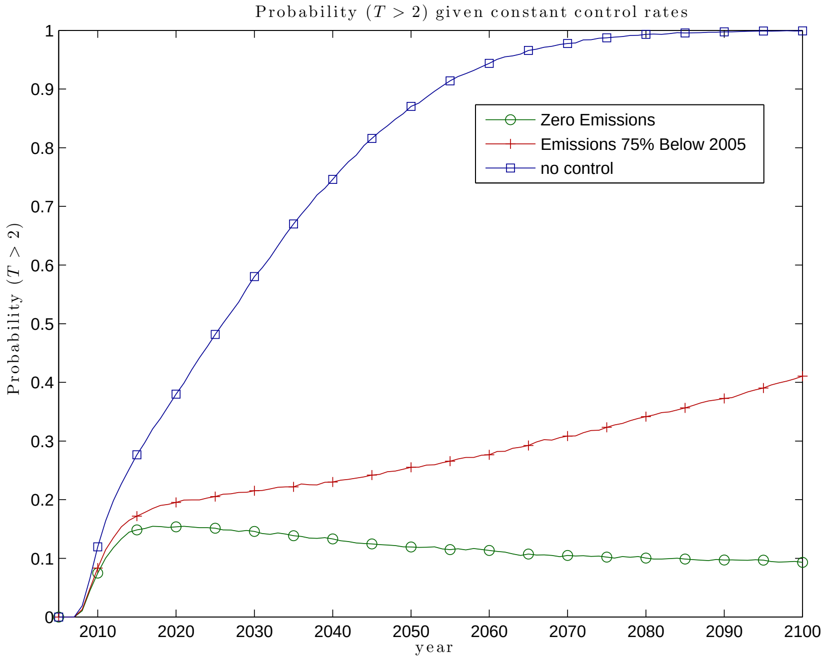

Figure 1: Probability of exceeding $T^* = 2°C$, given various emissions policies. The graph is the fraction of 10,000 simulations, each of which draw a random realization of β_1 from the prior distribution and a sequence ν_t, for which the temperature exceeds the target in the given year, for the given abatement policy. In all figures the data is annual, but markers are plotted only every five years for clarity.

For the first period, the probability of exceeding 2°C is nearly zero. Hence the current probabilistic constraint is non-binding for almost any ω. Given current information, however, there exists an approximately 15% chance that the 2°C target will eventually be exceeded, even with an immediate, permanent drop to zero emissions. Therefore, values of $\omega < 0.15$ are infeasible given today's information. An immediate, permanent, 75% drop in endogenous emissions below 2005 levels violates the constraint for any $\omega < 0.36$.[13]

Figure 1 also shows that a zero emissions policy can eventually overcome the inertia in the climate and reduce the probability of exceeding 2°C to near zero. A policy of zero emissions over time will slowly return the GHG concentrations to preindustrial levels (see equation 7).

The planner can choose an emissions policy over a period of decades such that most probabilistic constraints may be eventually satisfied with interior abatement rates, given current information. Nonetheless, the planner has little control over the climate on a year-

[13] An immediate 75% drop below 2005 emissions is a far stricter policy than, for example, a 350 ppm GHG target or the Kyoto agreement of 7% below 1990 levels.

to-year basis. Therefore, regardless of the emissions policy, in any given period the planner will likely find the probabilistic constraint to be either non-binding or infeasible.

4 Results

4.1 Optimal policy and uncertainty

Appendix A.3 details the solution method. First, we analyze how the optimal abatement policy varies with the probabilistic target and the uncertainty. Figure 2 plots the optimal abatement policy for two different true values of β_1: the prior, for which a doubling of GHGs causes a steady state temperature change of $\Delta T_{2\times} = 3.08°$C, and a relatively high value for which $\Delta T_{2\times} = 4°$C.

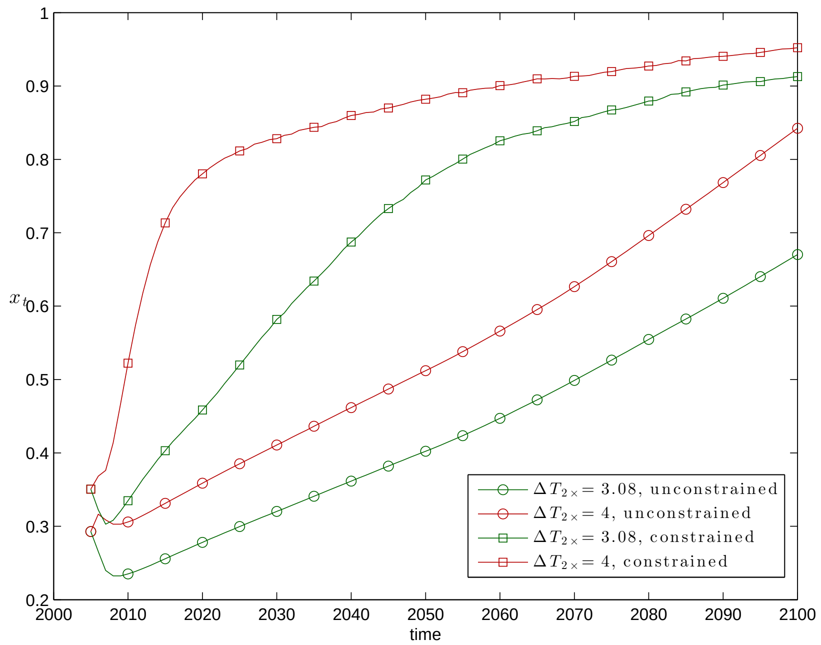

Figure 2: Emissions abatement rate for $\Delta T_{2\times} = \{3.08, 4\}$, unconstrained and constrained with $\omega = 0.5$. Mean of 10,000 simulations.

In both cases, the probabilistic constraint increases the abatement rate. The constraint increases the abatement rate initially as the planner must begin reducing emissions immediately to prevent the target from eventually being exceeded. The initial abatement rate with the probabilistic constraint is 35%, in contrast to the unconstrained initial abatement rate of 29%. As new information arrives, when $\Delta T_{2\times} = 4$, the planner updates the prior and therefore must increase the abatement rate to meet the target. When $\Delta T_{2\times} = 4$, the constrained planner must dramatically increase the abatement rate to 71% by 2015. In contrast, when $\Delta T_{2\times} = 3.08$, the constrained planner has more time to slow the climate inertia, and the abatement rate is only 33% in 2015.

The unconstrained abatement rate drops over the period 2005-2010. Kelly and Tan (2013) show that, while overall learning of the climate sensitivity is a slow process, the planner can rule out extreme values of $\Delta T_{2\times}$ relatively quickly if the true climate sensitivity is close to the prior. Therefore, learning reduces one motivation for early abatement- to insure against potentially extreme values of the climate sensitivity. The abatement rate then rises over time as abatement becomes less expensive, unabated emissions rise due to economic growth, and wealthier future households are more willing to purchase abatement.

The difference between the constrained and unconstrained abatement rates in 2010-2060 is much greater when $\Delta T_{2\times} = 4$. When $\Delta T_{2\times} = 4$, the optimal temperature rises. The planner learns the climate is more sensitive to GHG concentrations, and so the abatement expenditure required to keep the temperature at a given level increases, but the benefits are unchanged. In contrast, by definition, the probabilistic constraint employs a fixed target. Therefore, the difference between the unconstrained and constrained policies rises with $\Delta T_{2\times}$ because the constraint is inflexible: T^* and ω cannot adjust as new information arrives.

Figure 3 plots the average temperature changes for the above two cases.

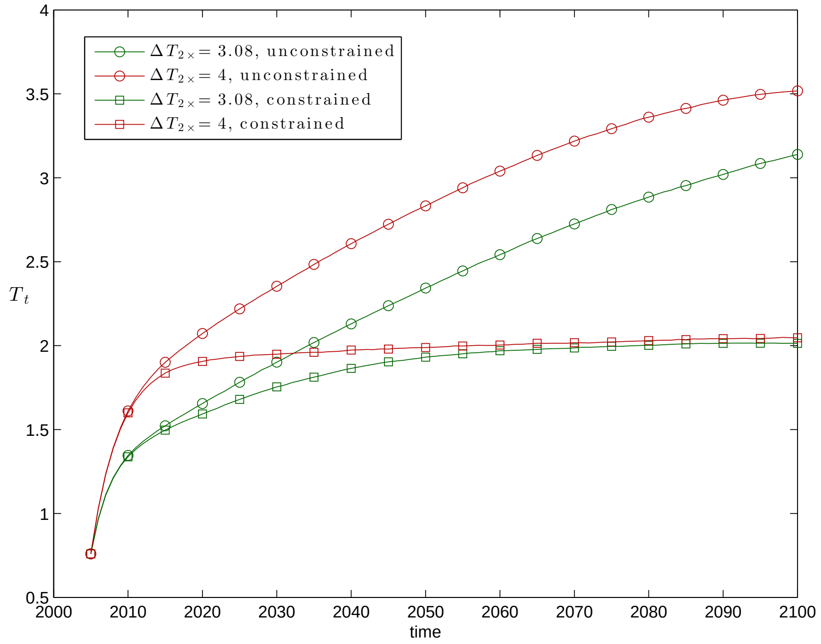

Figure 3: Temperature change for $\Delta T_{2\times} = \{3.08, 4\}$, unconstrained and constrained with $\omega = 0.5$. Mean of 10,000 simulations.

When $\Delta T_{2\times} = 3.08$, the unconstrained optimal temperatures cross the target in 31 years. Therefore, the planner has more time to slow climate change and can spread out the increase in the abatement rate. In contrast, when $\Delta T_{2\times} = 4$, the unconstrained temperature crosses the target in only 14 years. Therefore, in the constrained case, the abatement rate must rise more quickly to keep the temperature below the target. Notice the unconstrained optimal temperature rises when $\Delta T_{2\times}$ is higher as the planner responds to the higher required expenditure to keep the temperature at a given level by letting the temperature rise more. However, the target stays fixed at 2°C.

Next, we examine how abatement policy and the temperature respond to changes in the probability of exceeding the target, ω. We solve the model (38) for various values of ω, and simulate each solution 10,000 times. Figure 4 reports the mean optimal abatement rate for each ω, assuming the true climate sensitivity equals the prior.

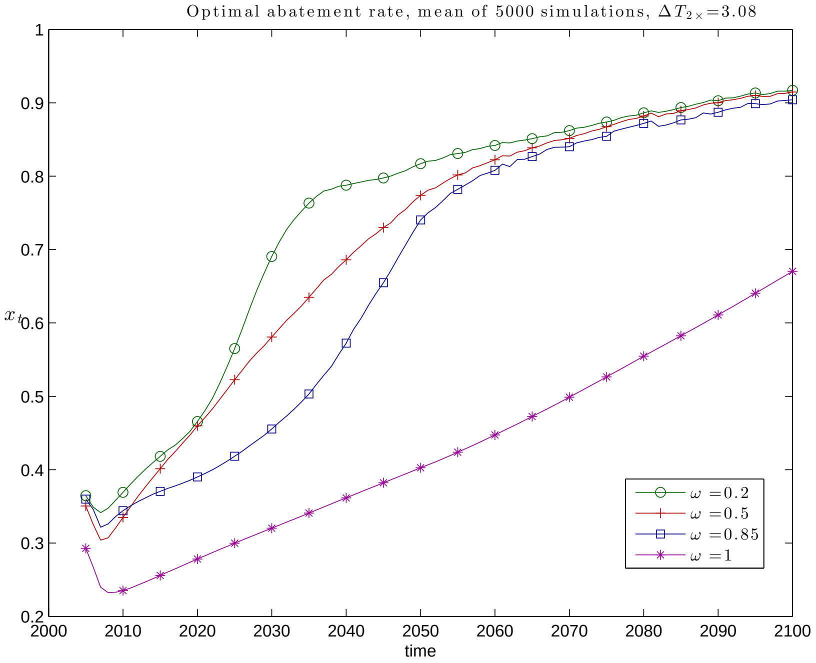

Figure 4: Optimal constrained abatement policy. Each curve is the mean of 10,000 simulations, with true $\Delta T_{2\times} = 3.08$, and the reported value of ω.

In Figure 4, $\omega = 1$ corresponds to the unconstrained optimum. As the maximum allowable probability of exceeding the target (ω) decreases, the constraint becomes more strict, causing the optimal abatement to rise. Non-monotonic regions arise in the first few periods, as the optimal abatement rate responds to a complex set of changing variables. Uncertainty is decreasing, causing the planner to become less concerned about the risk that the climate sensitivity and damages will be high. The cost of abatement and emissions intensity are falling exogenously, while unconstrained emissions and wealth are rising due to growth in productivity and capital.

The planner must ramp up abatement spending relatively early when $\omega = 0.2$, since in fact the temperature must converge to a level in which random weather shocks cause the temperature to exceed 2°C with probability of at most 0.2. This temperature is below 2°. In contrast, when $\omega = 0.85$ the planner can wait a few more years before ramping up abatement, since the constraint allows the temperature to converge to slightly above 2°C, requiring only that the weather shock causes the temperature to fall below 2°C with probability 0.15.

The unconstrained optimal abatement is far less than the constrained optimal abatement, even for $\omega = 0.85$. Regardless of ω, the ultimate temperature must stay relatively close to 2°, since the standard deviation of the weather shocks is only 0.11°. In contrast, the unconstrained optimal temperature rises to over 3°C, which requires much less abatement.

After 2050, abatement rates are similar for $\omega = 0.2$ and $\omega = 0.85$. At this point, learning has resolved much of the uncertainty and the climate no longer has an upward trend. The difference in abatement rates therefore reflects only that the mean steady state temperature differs with ω. However, the mean steady state temperature is only about 0.2° higher when $\omega = 0.85$ versus $\omega = 0.2$. Therefore, the abatement rates are relatively close in later years.

Figure 5 shows the temperature change for the same simulations as Figure 4.

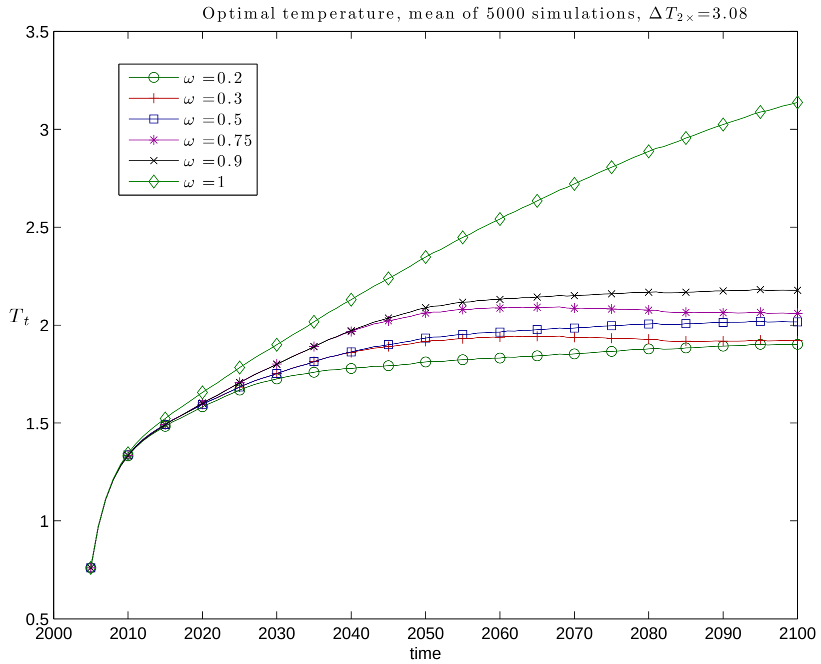

Figure 5: Optimal constrained temperature. Each curve is the mean of 1500 simulations, with true $\Delta T_{2\times} = 3.08$ and the reported value of ω.

The unconstrained optimal maximum mean temperature change is 3.26°C, so the 2° target is about 1.26 degrees too stringent, if the true climate sensitivity equals the prior.[14] Smaller values of ω imply a smaller maximum mean temperature change. For $\omega = 0.5$, the maximum mean temperature is very slightly below 2°C by 2100. The weather shock is normally distributed with mean zero. Hence, starting at 2°C, the temperature next period exceeds the constraint with probability 0.5. Similarly, since the standard deviation of the weather shock is 0.11°, with probability 0.2 the weather shock is greater than 0.09°. Hence,

[14]Our model is based on the Nordhaus DICE model, which has no tipping points, irreversibilities, etc. Other models with these features may feature smaller optimal maximum temperature changes.

the planner sets the maximum temperature change when $\omega = 0.2$ to slightly below 1.91 (actually 1.89). A high temperature shock implies costs rise next period, as the planner must increase the abatement level, whereas reducing the temperature below 1.91 increases abatement costs today, but reduces the risk of higher abatement costs later. Abatement costs are convex, however, which implies the planner will reduce the temperature below 1.91 as a precaution. This effect is small, however.

Overall, the maximum mean temperature is not very sensitive to ω. At $\omega = 0.9$, the maximum temperature change is 2.16°C (the level at which a weather shock causes the temperature to drop below 2°C with probability 0.1). The difference in maximum temperatures between $\omega = 0.2$ and $\omega = 0.9$ is only 0.25°C. The weather shock does not have enough variance to cause changes in ω to generate large differences in temperature. Although the maximum temperature must approach the unconstrained maximum temperature as $\omega \to 1$, the convergence is highly nonlinear. Suppose for example that $\omega = 0.999997$. Then, the planner can set the abatement rate so that the temperature converges to 2.5°C, since at 2.5° the probability of a large negative weather shock that causes the temperature to decrease to 2° is exactly $1 - \omega$. However, even with $\omega = 0.999997$, the mean temperature of 2.5° is still far from the unconstrained temperature of 3.26°. Therefore, the imposition of a probabilistic constraint causes a significant drop in temperature and requires significantly more abatement, for almost any value of ω.

The above analysis is for the case where the true climate sensitivity equals the prior. Sufficient time exists to slow the inertia in the climate and stabilize the temperature, both because the true climate sensitivity is moderate and because the planner is not very surprised by the true climate sensitivity. However, many realizations of the uncertain climate sensitivity cause the temperature to exceed the target with probability one or zero irrespective of ω. If the climate sensitivity is sufficiently high, the temperature exceeds the target with probability one, causing zero emissions regardless of ω. If the climate sensitivity is sufficiently low, the cost of reducing the temperature becomes inexpensive and the unconstrained optimum falls below 2°C. In this case, emissions equal the unconstrained optimum regardless of ω. Therefore, much of the distribution of uncertainty over the climate sensitivity results in emissions which are independent of ω. For example, Figure 6 plots average temperature change over time as a function of ω when $\Delta T_{2\times} = 5$.[15]

[15]Weitzman (2009) using IPCC data assigns prior probability that $\Delta T_{2\times} \geq 5 = 0.07$.

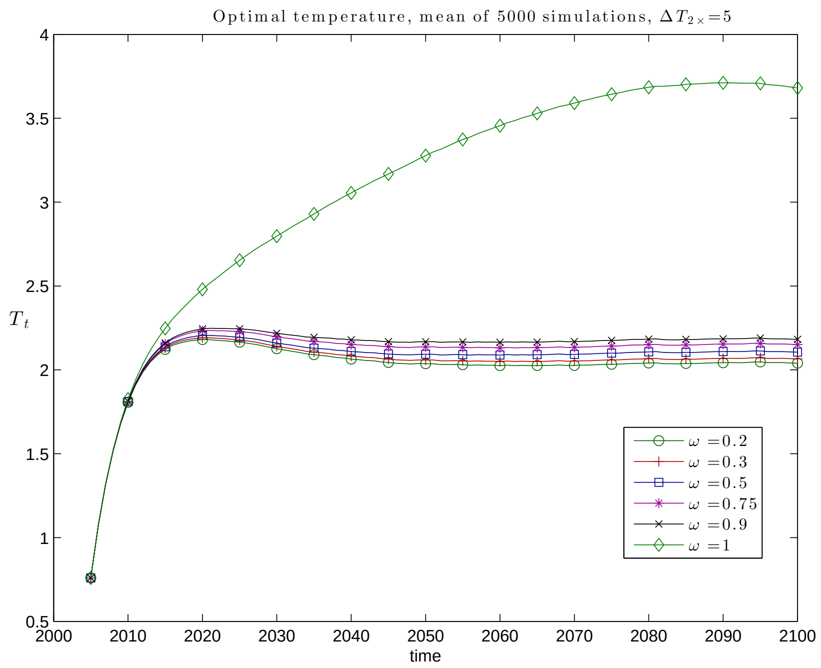

Figure 6: Optimal constrained temperature. Each curve is the mean of 5000 simulations, with true $\Delta T_{2\times} = 5$ and the reported value of ω.

The mean optimal unconstrained temperature change rises to 3.71°C, indicating the optimal unconstrained temperature is sensitive to climate sensitivity. The mean optimal unconstrained temperature change increases because when the climate is more sensitive to GHG concentrations, the abatement expenditures required to keep the temperature at a given level rises, but the benefits are unchanged. In contrast, the target is by definition inflexible and does not vary with the resolution of uncertainty. The mean maximum temperature change falls with ω as in Figure 5, but is not very sensitive to ω. When $\Delta T_{2\times} = 5$ and $\omega = 0.2$, on average the temperature exceeds the target for about 104 years. Since $\Delta T_{2\times} = 5$ is unexpectedly high, the climate has unexpected inertia, and crosses the target after just a few years. The problem is exacerbated by the fact that the planner believes $\Delta T_{2,\times} < 5$ for many years, and underestimates the effect of emissions on the climate. Once the temperature crosses the target, the planner sets emissions to zero, but the climate continues to rise on average because the inertia is unexpectedly strong. Note that Figure 5 is an average of 5000 simulations. In 18% of the simulations, the weather shocks are low enough such that the temperature falls below 2°C in less than 40 years. Conversely, 34% of the simulations require 150 years or more to reduce the temperature back to the target.

4.2 Welfare Loss

The probabilistic constraint at least weakly restricts the choice set for the planner, and therefore must result in a welfare loss.[16] Our interest is in how the welfare loss varies with ω and how uncertainty affects the welfare loss.

Figure 7 graphs the welfare loss as a percentage of the welfare of the unconstrained problem, $\omega = 1$.

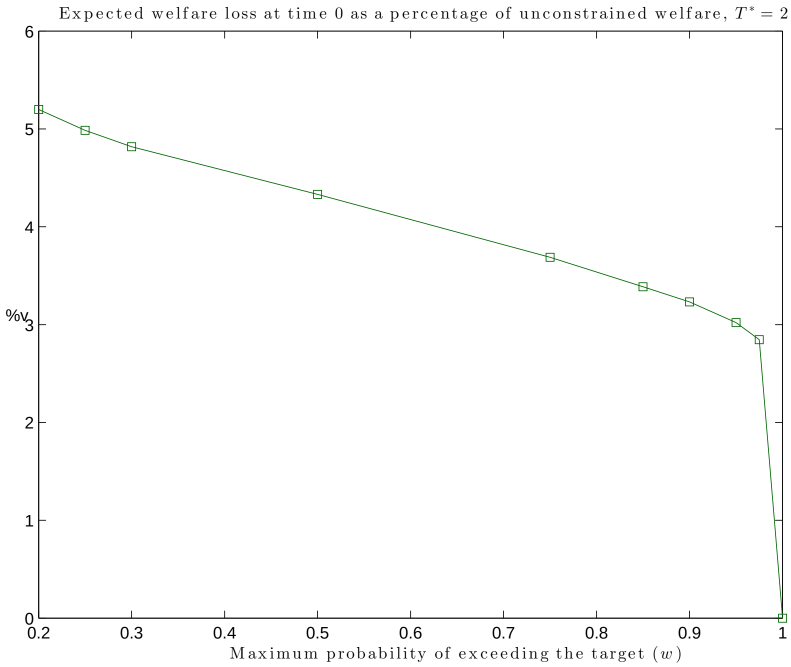

Figure 7: Welfare loss as a function of ω. The graph plots $1 - V(s_0; \omega) / V(s_0; 1)$, where V is the solution to (37), as a function of ω.

From the graph, the unconstrained problem has no welfare loss, and welfare loss is decreasing in ω since relaxing the constraint reduces the welfare loss. The slope is relatively flat for $\omega < 0.975$, since even relatively large values of ω force the maximum climate change to be close to two degrees, rather the optimum which is 3.26°C and rises with the true value of the climate sensitivity (see Figure 5). Welfare loss is about 3-5% of the unconstrained

[16] We are following, for example, Nordhaus (2007), who imposes a 2°C constraint a version of the model with no uncertainty, and then calculates the welfare loss. Other authors use cost effective analysis (CEA) or cost risk analysis (CRA), which replace the damage function with the probabilistic constraint (see Neubersch, Held, and Otto 2014, for an excellent discussion of CEA and CRA). A damage function, despite being uncertain, allows for a transparent interpretation of the welfare costs of temperature changes, which is the focus of this paper.

policy, depending on ω (Table 2 reports the exact welfare losses). For most of the state space, the probability of exceeding the target is either zero or one regardless of ω, and so the welfare loss does not vary significantly with ω over most of the state space.

Mastrandrea et al. (2010) and Neubersch, Held, and Otto (2014) calibrate a value of $\omega = 0.33$ based on an interpretation IPCC statements calling for policies for which achieving the 2° target is likely. We therefore let $\omega = 0.33$ be the baseline stabilization policy recommendation. Table 2 then indicates that the baseline stabilization policy results in a welfare loss of approximately 4.7%.

Figure 8 plots the welfare loss for various true values of $\Delta T_{2\times}$.

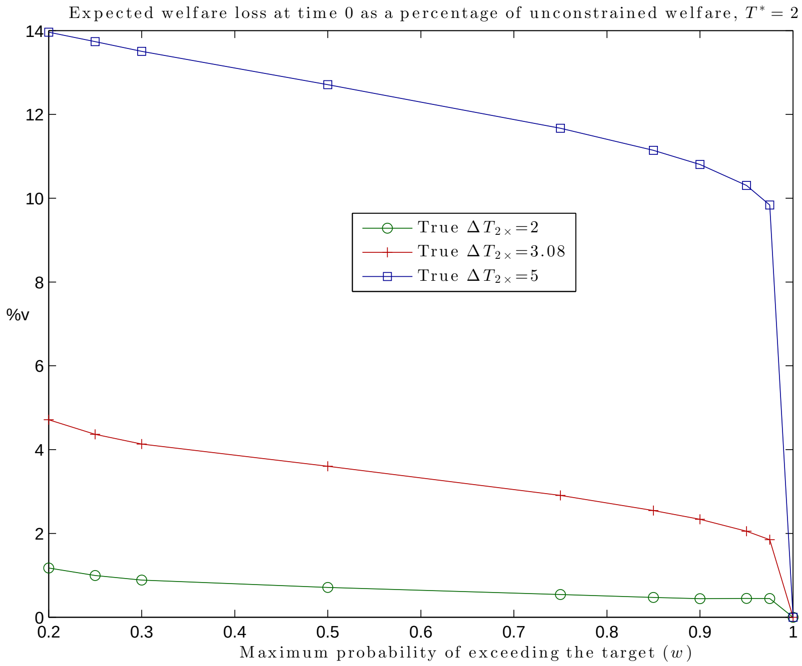

Figure 8: Welfare loss as a function of ω and the true $\Delta T_{2\times}$.

The welfare loss increases with the true value of $\Delta T_{2\times}$, rising to 14% when $\Delta T_{2\times} = 5$ and $\omega = 0.2$.

When the arrival of new information indicates the climate sensitivity is higher than expected, the unconstrained planner learns that the cost of keeping the temperature at a given level has risen. The planner increases abatement in response to the higher expected damages. Nonetheless, the high climate sensitivity implies temperature is on a higher trajectory. In contrast, the constrained planner must increase abatement still further, to keep the temperature on the same trajectory despite the high climate sensitivity, because the target does not change. Therefore, higher values of $\Delta T_{2\times}$ cause greater welfare losses, because the constraint is inflexible.

Traeger (2014) shows that, for the DICE model under certainty, the maximum temperature change is approximately 3.6°C above preindustrial.[17] Therefore, in the certainty version of the model, the target is set 1.6° too low. After imposing a pure ($\omega = 0$) stabilization target on the certainty version of the model,[18] we find the welfare loss of a 2°C target is 2.84%. This result depends on the specification of the damage function, assumptions about the discount rate, the lack of threshold effects, etc. In the certainty version of the model, more pessimistic parameter assumptions would lower the maximum optimal temperature change to 2° or lower, in which case imposing the 2° target would create no welfare loss.

In contrast, a welfare loss exists in the model with uncertainty, irrespective of parameter assumptions, because the target is inflexible and abatement policy overreacts to transient shocks. Table 2 shows that the welfare loss in the model with uncertainty ranges between 5.2% for $\omega = 0.2$ to 2.85% for $\omega = 0.975$. The welfare loss with uncertainty differs from the welfare loss with certainty for three reasons. First, the optimal maximum temperature changes with the resolution of uncertainty, but the target is inflexible, creating a welfare loss. In contrast, adjusting the target is not necessary under certainty. Second, with uncertainty the constrained planner must respond to transient weather shocks. These shocks are absent in the model with certainty. Third, the difference between the maximum unconstrained mean temperature and the maximum constrained mean temperature when the true climate sensitivity equals the prior is smaller under uncertainty than the difference between the maximum unconstrained temperature and the maximum constrained temperature under certainty. The risk of GHG emissions eventually causing large temperature changes makes the planner more cautious under uncertainty. Therefore, the 2° target is not as restrictive when the true climate sensitivity equals the prior with uncertainty relative to certainty.

The first two channels cause a larger welfare loss relative to certainty, while the third channel causes a smaller welfare loss relative to certainty. Overall the first two channels dominate for all values of ω. The welfare loss of the baseline model ($\omega = 0.33$) under uncertainty is 66% greater than the model with certainty. Table 2 indicates the welfare loss under uncertainty is 82.8% greater than the model with certainty for $\omega = 0.2$, although for

[17] In our model with uncertainty, the maximum of the mean temperature change is 3.26°C. The lower temperature occurs because the risk averse planner responds to the risk of damages from a potentially high climate sensitivity by increasing abatement.

[18] A probabilistic target cannot be imposed on the certainty version of the model, since the planner can always choose a path which satisfies the target with probability one, and any path which exceeds the target also does so with probability one.

$\omega = 0.975$ the third effect becomes more important, and the welfare loss is only slightly higher under uncertainty.

Unfortunately, the three channels by which uncertainty affects the welfare loss of a stabilization target are difficult to isolate. For example, eliminating the weather shocks from the model removes the need for the optimal constrained abatement policy to respond to transient shocks. However, without weather shocks, the planner learns the climate sensitivity after one observation, and the model reduces to the certainty case.

However, some channels can be approximately isolated. For example, consider an exercise designed to isolate the welfare loss caused by the response of abatement policy to transient shocks. We perform a simulation which keeps two sets of temperature data. The first set of temperature data includes the weather shocks and the second set of data sets all shocks equal to zero. We use the temperature data with shocks to compute the updates to the mean and variance of the prior uncertainty distribution, and use the temperature data without shocks to compute damages, the optimal decisions, and the future temperature. In this way, the speed of learning is nearly identical to the previous case, but abatement policy does not directly adjust to transient shocks, which are absent.[19] Table 2 shows that eliminating weather shocks (except for updating priors) reduces the welfare loss by 6.5-9.7%.

The third effect is that the probabilistic target used under uncertainty is less restrictive than the pure stabilization target under certainty. Under certainty, the maximum unconstrained temperature is 3.6°, whereas the constrained target is 2°C, a difference of 1.6°C. In contrast, Figure 5 shows that the unconstrained maximum mean temperature is 3.26°, and that the maximum constrained mean temperature is 1.89-2.24°, depending on ω. When $\omega = 0.2$, the temperature difference is 1.37° versus 1.6° under certainty. Therefore, for $\omega = 0.2$, the third effect, the difference between optimal and constrained temperature, is similar for uncertainty and certainty. The welfare loss for $\omega = 0.2$ is 82.8% higher under uncertainty, indicating that removing the third effect causes the welfare loss to increase.

Overall, then the majority of the added welfare loss in the model with uncertainty results from the inflexibility of the target. The climate sensitivity is uncertain, and the inability of the constrained planner to adjust policy as new information arrives reduces the value of learning in the model.

[19] That is, we compute $x\left(S\left(T\left(\nu=0\right),\mu\left(\nu\right),\eta\left(\nu\right)\right)\right)$, which differs from the abatement policy when weather shocks affect the temperature, eventually slightly affecting the speed of learning.

5 Conclusions

In this paper, we evaluated the common policy recommendation of a 2°C temperature limit in the context of an integrated assessment model with uncertainty and learning about the climate sensitivity. Because climatic processes are still highly uncertain and subject to inertia, it is difficult to envision controlling the climate to the precise degree implied by the 2°C target. Indeed, we show that even with an immediate reduction of all GHG emissions to zero, the temperature eventually exceeds 2°C with probability 0.15. Further, we show that as uncertain climatic processes evolve, the temperature randomly moves to a place where the 2° target is either impossible to satisfy or is satisfied even with no abatement. Our results therefore cast doubt on how workable a stabilization target is in an uncertain environment.

Further, we show that stabilization targets affect welfare in three ways: first, with uncertainty, as new information arrives the optimal temperature path adjusts. Because the stabilization target is by definition inflexible, adhering to a stabilization target causes welfare loss. Second, as the climate randomly evolves over time, the temperature will exceed the target for at least some periods. In this case, the planner must set an excessively high abatement rate to immediately return the temperature back to the target. Third, uncertainty makes the risk averse planner more cautious: the average unconstrained temperature change is lower under uncertainty. Thus, the welfare loss associated with setting the target lower than is optimal is less under uncertainty. We show that the total welfare loss for the baseline policy is 4.7%, most of which is due to inflexibility of the stabilization target.

Our model may be extended in a number of ways. We could consider other targets such as limit GHG concentrations to 350 ppm or limiting sea level rise or ocean acidification. Our results will also likely extend to regulations other than climate change. For example, some fisheries regulations try to achieve a particular stock of fish, even though fish stocks are affected by many uncertain processes other than the size of the catch. One may envision stabilization targets as providing some welfare benefits outside the current model. For example, they could serve as a commitment device to induce firms to invest in irreversible abatement capital.

A 2°C stabilization targets is easier to convey to the public than, for example, a particular carbon tax. Since damages are a function of temperature, it is also easier to understand the impacts of 2°C temperature limit. However, one must use caution in that a stabilization target may convey the false impression that we have precise control over an uncertain climate and that our understanding of the optimal temperature change will not change as new information arises.

A Appendix

A.1 Exogenous variables

The DICE model includes many variables which change exogenously over time. Further, unlike DICE, we assume the ocean temperature also changes exogenously to reduce the state space. Reducing the state space from seven to six variables significantly reduces the computation time. Traeger (2014) presents a deterministic DICE model with exogenous ocean temperature. Therefore, we take the evolution of the exogenous variables directly from that study. For completeness, they are listed below. We refer the reader to Traeger (2014) for details.

Population growth:

$$L(t) = L_0 + \left(\bar{L} - L_0\right)\left(1 - \exp\left[-g_L t\right]\right). \tag{53}$$

TFP growth:

$$A(t) = A_0 \exp\left[g_{A,0} \frac{1 - \exp\left[-\delta_A t\right]}{\delta_A}\right]. \tag{54}$$

Backstop cost:

$$\Psi(t) = \frac{\sigma(t)}{a_2} a_0 \left(1 - \frac{1 - \exp\left[g_\Psi t\right]}{a_1}\right). \tag{55}$$

Emissions intensity of output:

$$\sigma(t) = \sigma_0 \exp\left[g_{\sigma,0} \frac{1 - \exp\left[-\delta_\sigma t\right]}{\delta_\sigma}\right]. \tag{56}$$

Exogenous emissions:

$$B(t) = B_0 \exp\left[-\delta_B t\right]. \tag{57}$$

Decay rate of GHGs:

$$\delta_m(t) = \bar{\delta}_m + \left(\delta_{m,0} - \bar{\delta}_m\right) \exp\left[-\delta_m^* t\right]. \tag{58}$$

Exogenous forcing:

$$EF(t) = EF_0 + 0.01\left(EF(100) - EF_0\right) \cdot \min\left\{t, 100\right\}. \tag{59}$$

Heat uptake from ocean:

$$O\left(t\right) = \max\left\{\Delta_{T1} + \Delta_{T2}t + \Delta_{T3}t^2, 0\right\}. \tag{60}$$

Discount factor:

$$\beta\left(t\right) = \exp\left[-\delta_u + (1-\eta)\log\left(\frac{A\left(t+1\right)}{A\left(t\right)}\right) + \log\left(\frac{L\left(t+1\right)}{L\left(t\right)}\right)\right]. \tag{61}$$

A.2 Tables

Parameter	Description	Value
L_0	Initial population	6514
$\bar{L}$	Steady state population	8600
g_L	decline rate in population growth	0.035
A_0	Initial TFP	0.0058
g_{A0}	Initial TFP growth rate	0.0131
δ_A	Decline rate in TFP growth rate	0.001
γ	Capital share	0.3
δ_k	Depreciation rate of capital	0.1
ϕ	Coefficient of risk aversion	2
δ_u	Pure rate of time preference	0.015
g_Ψ	Backstop cost growth rate	-0.005
a_0	Initial backstop cost	1.17
a_1	Backstop cost parameter	2
a_2	Cost convexity	2
b_1	Damage function parameter	0.00284
b_2	Damage function convexity	2
σ_0	Initial emissions intensity	0.1342
$g_{\sigma,0}$	Initial growth rate in σ	-0.0073
δ_σ	Decline rate in emissions intensity growth	0.003
B_0	Initial exogenous emissions	1.1
δ_B	decay rate in exogenous emissions	0.0105
M_B	Pre-industrial GHG stock (gigatons)	590
δ_m^*	decay rate in GHG decay rate	0.0083
$\bar{\delta}_m$	steady state GHG decay rate	0.01
$\delta_{m,0}$	initial decay rate	0.014
EF_0	Initial exogenous forcing	-0.06
EF_{100}	Exogenous forcing at $t = 100$	0.3
Ω	Forcing coefficient	3.8
α	Ocean heat uptake	0.2837^{-1}
ξ	Heat transfer coefficient from the ocean	0.07
Δ_{T1}	Ocean Temperature Parameter	0.7
Δ_{T2}	Ocean Temperature Parameter	0.02
Δ_{T3}	Ocean Temperature Parameter	-0.00007
k_0	Initial capital per TFP adjusted person	3.6261
T_0	Current air warming above pre-industrial	0.76
m_0	Current GHG stock, relative to pre-industrial	1.371
μ_0	Mean of climate feedback prior distribution	0.65
η_0	Precision of climate feedback prior distribution	0.13^{-2}
ρ	Precision of weather shock	0.11^{-2}

Table 1: Description and values of model parameters.

	Total	No weather shocks	
ω	Loss (%)	Loss (%)	Difference (%)
0.20	5.20	4.86	6.49
0.25	4.99	4.57	8.29
0.30	4.82	4.42	8.16
0.33	4.73	4.34	8.23
0.50	4.33	3.95	8.84
0.75	3.69	3.39	8.12
0.85	3.39	3.07	9.47
0.90	3.23	2.92	9.73
0.95	3.02	2.73	9.49
0.97	2.85	2.60	8.78
1.00	0	0	NA
		Certainty	
0	2.84		

Table 2: Decomposition of the welfare loss. The baseline case is $\omega = 0.33$.

A.3 Solution

We solve the model by forming a grid of values for the state space, and then assume v takes the form of a cubic spline with continuous first and second derivatives. The model can then be solved by choosing an initial spline, optimizing at each grid point, and then using the solution to update the parameters of the spline.[20]

Once the model converges, we obtain the optimal decision rules, $x(s)$ and $k'(s)$. We then simulate the model using the decision rules and the transition equations (17)-(24). The algorithm is:

1. Draw a true value of the climate feedback parameter β_1^* from $N[\mu_0, \frac{1}{\eta_0}]$.

2. Given s_0 compute $x(s_0) = x_0$.

3. Given x_0, s_0, β_1^*, and a randomly drawn weather shock ν_0, compute s_1 from transition equations.

4. Repeat steps (2)-(3) for np years.

5. Repeat steps (1)-(4) ns times with different draws for ν and β_1.

[20]Kelly and Kolstad (1999b) give a detailed explanation of this solution method, except they use neural networks rather than splines.

6. Compute means over all simulations to get the expected value of each variable in each time period.

The above algorithm gives the expected value of each variable conditional on the prior distribution for β_1. In some cases, we fix a value for β_1 and vary only ν across simulations. In this case, we obtain the expected value of each variable conditional on β_1.

References

Bruckner, T., and K. Zickfield, 2009, "Emissions Corridors for Reducing the Risk of a Collapse of the Atlantic Thermohaline Circulation," *Mitigation and Adaptation Strategies for Global Change*, 14, 61–83.

Copenhagen Accord, 2009, "Report of the Conference of the Parties on its fifteenth session, held in Copenhagen from 7 to 19 December 2009," Discussion paper, United Nations Framework Convention on Climate Change, http://unfccc.int/resource/docs/2009/cop15/eng/11a01.pdf.

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.

den Elzen, M. G. J., M. Meinshausen, and D. van Vuuren, 2007, "Multi-gas Emissions Envelopes to Meet Greenhouse Gas Concentration Targets: Cost versus Certainty of Limiting Temperature Increase," *Global Environmental Change-Human and Policy Dimensions*, 17, 260–280.

European Commission, 2007, "Communication from the Commission to the Council, the European Parliament, the European Economic and Social Committee and the Committee of the Regions - Limiting global climate change to 2 degrees Celsius - The way ahead for 2020 and beyond," Discussion paper, European Commission, Brussels, Belgium, http://eurlex.europa.eu/LexUriServ/LexUriServ.do?uri=CELEX:52007DC0002:EN:NOT.

Hansen, J., 2005, "A Slippery Slope: How Much Global Warming Constitutes 'Dangerous Anthropogenic Interference'?," *Climatic Change*, 68, 269–79.

Hare, B., and M. Meinshausen, 2006, "How Much Warming Are We Committed to and How Much Can Be Avoided?," *Climatic Change*, 75, 111–149.

Harvey, L., 2007, "Allowable CO2 Concentrations Under the United Nations Framework Convention on Climate Change as a Function of the Climate Sensitivity Probability Distribution Function," *Environmental Research Letters*, 2, 10.

Held, H., E. Kriegler, K. Lessmann, and O. Edenhofer, 2009, "Efficient Climate Policies Under Technology and Climate Uncertainty," *Energy Economics*, 31, S50–S61.

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., et al., 2005, "Avoiding Dangerous Anthropogenic Interference with the Climate System," *Climatic Change*, 73, 227–238.

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., and C. D. Kolstad, 2001, "Solving Infinite Horizon Growth Models with an Environmental Sector," *Journal of Computational Economics*, 18, 217–31.

Kelly, D. L., and Z. Tan, 2013, "Learning and Climate Feedbacks: Optimal Climate Insurance and Fat Tails," University of Miami working paper.

Keppo, I., B. O'Neill, and K. Riahi, 2007, "Probabilistic Temperature Change Projections and Energy System Implications of Greenhouse Gas Emissions Scenarios," *Technological Forecasting and Social Change*, 74, 936–961.

Kvale, K., et al., 2012, "Carbon Dioxide Emission Pathways Avoiding Dangerous Ocean Impacts," University of Victoria Working Paper.

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

Lemoine, D., and I. Rudik, 2014, "Steering the Climate System: Using Inertia to Lower the Cost of Policy," Discussion Paper 14-03, University of Arizona, Department of Economics.

Lemoine, D. M., and C. P. Traeger, 2014, "Watch Your Step: Optimal Policy in a Tipping Climate," *American Economic Journal: Economic Policy*, 6, 1–31.

Lorenz, A., M. Schmidt, E. Kriegler, and H. Held, 2012, "Anticipating Climate Threshold Damages," *Environmental Modeling and Assessment*, 17, 163–175.

Mastrandrea, M., et al., 2010, "Guidance Note for Lead Authors of the IPCC Fifth Assessment Report on Consistent Treatment of Uncertainties Intergovernmental Panel on Climate Change IPCC," Discussion paper, IPCC.

Neubersch, D., H. Held, and A. Otto, 2014, "Operationalizing climate targets under learning: An application of cost-risk analysis," *Climatic Change*, 126, 305–18.

Nordhaus, W., 2007, "DICE Model Version as of May 22, 2007," `http://www.econ.yale.edu/nordhaus/DICEGRAMS/DICE2007.htm`.

O'Neill, B., and M. Oppenheimer, 2002, "Climate Change - Dangerous Climate Impacts and the Kyoto Protocol," *Science*, 296, 1971–2.

Portney, P., I. Parray, H. Gruenspecht, and W. Harrington, 2003, "Policy Watch: The Economics of Fuel Economy Standards," *Journal of Economic Perspectives*, 17, 203–17.

Richels, R., A. Manne, and T. Wigley, 2004, "Moving beyond concentrations: the challenge of limiting temperature change," Discussion Paper 04-11, AEI-Brookings Joint Center for Regulatory Studies, `http://ssrn.com/abstract=545742`.

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

Rudik, I., 2014, "Targets, Taxes, and Learning: Optimizing Climate Policy Under Knightian Damages," University of Arizona, Department of Economics, `http://dx.doi.org/10.2139/ssrn.2516632`.

Schmidt, M., A. Lorenz, H. Held, and E. Kriegler, 2011, "Climate Targets Under Uncertainty: Challenges and Remedies," *Climatic Change*, 104, 783–791.

Schubert, R., et al., 2006, "The Future Oceans - Warming Up, Rising High, Turning Sour," Discussion paper, German Advisory Council on Global Change, Special report, `http://www.wbgu.de`.

Stavins, R., 1993, "Experience with Market-Based Environmental Policy Instruments," in K. Maler, and J. Vincent (ed.), *Handbook of Environmental Economics*, vol. 1, . pp. 355–435, Elsevier Science, Amsterdam.

Stocker, T., Q. Dahe, and G. Plattner, 2013, "Technical Summary," in *Climate Change 2013: The Physical Science Basis* . pp. 1–79, IPCC Secretariat.

Traeger, C., 2014, "A 4-stated DICE: Quantitatively Addressing Uncertainty Effects in Climate Change," *Environmental and Resource Economics*, 59, 1–37.

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