Probabilistic Stabilization Targets*

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

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

Current Version: October 1, 2016

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 the climate is difficult to control in the short run, although as learning resolves the planner eventually achieves the target with a sustained reduction in emissions over time.

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. Introduction of a stabilization target into the baseline model with uncertainty results in a welfare loss of 4.7%, which is 66% higher than the cost of introducing the target in the certainty version of the model.

JEL Codes: Q54, Q58, O44

*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

A common feature of international climate agreements are stabilization targets: the maximum allowable change in an environmental variable tied to greenhouse gas (GHG) emissions. The most well-known stabilization target, and a central component of the 2015 Paris Climate Accord, calls for limiting the increase in global mean surface temperature to 2°C above its pre-industrial level. Staying below the 2°C target is widely advocated to prevent irreversible and catastrophic climate damages. With a target in hand, researchers assume certainty equivalence and compute the least cost emissions path that ensures remaining below the stipulated threshold.

However, uncertain parameters and random weather shocks affect climate dynamics. A sequence of large shocks, or an unexpectedly sensitive relationship between temperature and GHG concentrations (the climate sensitivity), can cause the temperature to exceed the target, irrespective of the emissions policy. *Pure* stabilization targets are thus not feasible in a dynamic setting with uncertainty. An alternative is *probabilistic* stabilization targets (see for example, Held, Kriegler, Lessmann, and Edenhofer 2009) which instead require that the environmental variable stay beneath the target with a maximum allowable probability.[1] A probabilistic target is feasible if the maximum allowable probability of exceeding the target is sufficiently high. We analyze stabilization targets in an integrated assessment model of the climate and economy with uncertainty, Bayesian learning, and random weather shocks.[2] We show that a probabilistic target in an environment with uncertainty induces two welfare losses previously undocumented in the literature.

First, the arrival of new information changes beliefs about the climate sensitivity over time, changing the optimal temperature. For example, an unexpectedly high climate sensitivity raises the optimal temperature, since achieving a given temperature requires more abatement expenditures, while the damages from that temperature are unchanged. By definition, however, the target temperature remains unchanged. Therefore, a welfare cost ensues because the target is *inflexible*. Second, we show that stabilization targets force overly stringent policy responses to transient weather shocks. Weather shocks cause the temperature to exceed the target with positive probability. If so, abatement expenditures must rise to bring the temperature back to the target. The abatement cost is excessive, since the temperature naturally returns to the target as the shock dies out regardless. 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. In models which assume certainty, a welfare cost only occurs if the target is set below the temperature resulting from the optimal

[1] A pure stabilization target is a special case of a probabilistic target where the probability of exceeding the target is zero. Therefore, without loss of generality, we consider only probabilistic targets.

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

emissions policy. 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°C. This result is sensitive to the damage function, the discount rate, the lack of tipping points, and other assumptions in the 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, which reduces the welfare loss associated with the 2°C target.[3] In contrast, welfare losses associated with inflexibility and overreaction depend only on the interaction of uncertainty and the target and are therefore robust to alternative modeling assumptions.

We show that introducing a stabilization target into the baseline model with uncertainty causes a welfare loss of 4.7%, which is 66% higher than the welfare cost of introducing a stabilization target into the model with certainty.[4] 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 welfare loss is relatively insensitive to the policy parameter which specifies the maximum probability of exceeding the target. The welfare loss increases, however, with the unknown true value of the climate sensitivity. If the climate sensitivity is 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 the true value of the climate sensitivity.

Our results yield several implications for climate change policy. First, large and immediate reductions in GHG emissions are necessary in order to keep the temperature below the 2°C target with reasonable confidence: 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. Second, to mitigate the welfare costs of inflexibility, international climate agreements should feature provisions that call for adjusting the target temperature in response to better information about the climate sensitivity. Third, targeting a different environmental variable which is less variable then the temperature would reduce the costs of overreaction. For instance, a target on GHG concentrations, which are more stable on year-to-year basis, would lead to fewer overreactive abatement responses.

We also study the dynamics of learning in the presence of stabilization targets. Since the target reduces emissions, the planner learns the climate sensitivity more slowly. We find the standard deviation of the prior distribution of the climate sensitivity is approximately 10% higher in 2050 than without the stabilization target, if the true climate sensitivity is equal to the mean of the prior. If the true climate sensitivity is unexpectedly high, then learning

[3]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.

[4]We follow Neubersch, Held, and Otto (2014) and calibrate a maximum probability of exceeding the target of 0.33, based on an interpretation of IPCC statements calling for policies for which achieving the 2°C target is likely.

slows further. In turn, unexpectedly slow learning makes achieving the temperature target more difficult, because if the climate sensitivity is sufficiently high relative to the mean of the prior distribution, the planner does not abate emissions enough.

2 Stabilization Targets

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 warming targets. Other stabilization targets differ according to the choice of target. 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. Some policy groups such as 350.org favor a GHG concentration target.[5] Replacing a temperature target with an alternative target variable would not change our qualitative conclusions. However, the overreaction welfare loss would change quantitatively, as alternative target variables may be more or less variable than temperature.

A large literature computes the least-cost emissions path which stabilizes the climate at 2°C under certainty (Nordhaus 2007, Richels, Manne, and Wigley 2004, Lemoine and Rudik 2014). 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, result in warming that exceeds the target by a considerable margin.[6] Indeed, a branch of the literature focuses on the likelihood of meeting stabilization targets for various emissions scenarios proposed by policy makers, or what emissions paths satisfy the target for various values of the climate sensitivity. 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. This research provides an important first step in estimating the probability of exceeding the target. Here we take two important next steps by introducing learning and fat-tailed uncertainty. Learning allows the emissions policy to adjust as new information arrives, while fat tailed uncertainty implies a relatively large value of the climate sensitivity is expected to occur more often than if the

[5] den Elzen, Meinshausen, and van Vuuren (2007) and Lemoine and Rudik (2014) study GHG targets.

[6] Stabilization targets evolved as method of dealing with damage uncertainty in integrated assessment models. If expected damages are sufficiently high when the temperature exceeds a threshold, a limit on temperature changes is optimal. Somewhat paradoxically, climate sensitivity uncertainty is more difficult to address within this framework.

uncertainty was normally distributed. Learning and fat-tailed uncertainty have competing effects on the ability of the planner to stabilize the climate.

Learning allows the planner to more easily meet a target by quickly adjusting emissions if new information indicates the climate sensitivity is greater than previously thought. Further, learning allows the planner to move closer to the target, while still remaining 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 through feedback effects (Kelly and Tan 2015). Therefore, controlling the temperature requires a precise estimate of the climate sensitivity, which takes time. We find the optimal near term policy with Bayesian learning is similar to the case without learning. We find that learning eventually allows the temperature to move closer to the target, but the effect is marginal since learning is slow.

It is well known that the prior distribution of the climate sensitivity has a fat tail, since small, uncertain feedback effects eventually have large effects on the temperature (Kelly and Tan 2015, Weitzman 2009). We show that stabilizing the climate is more difficult when the climate sensitivity is high, as the climate has more inertia and learning slows. The welfare loss from inflexibility of the target also increases with the climate sensitivity.

We analyze 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 are 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.

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

Bayesian 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.

3 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 only two reservoirs exist for carbon (atmosphere and ocean/biosphere). 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 3.1-3.2 describe the economic and climate models briefly (refer to Traeger 2014, for a detailed discussion).

3.1 Economic system

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

$$Q_t = A(t) K_t^{\gamma} L(t)^{1-\gamma}. \tag{3.1}$$

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:

$$Y_t = \frac{1 - \Psi(t) x_t^{a_2}}{1 + b_1 T_t^{b_2}} A(t) K_t^{\gamma} L(t)^{1-\gamma}. \tag{3.2}$$

The presence of the damage function makes the model a DICE-type model of optimal policy, rather than a cost-effective analysis (CEA) framework typically used to study targets under certainty. Under certainty, CEA models do not include a damage function, but instead constrain temperatures to remain below the target. This is equivalent to a damage function which is zero for temperatures below the target and arbitrarily high for temperatures above the target. However, given climate sensitivity uncertainty, the target is exceeded with a strictly positive probability that depends on emissions. Therefore, the optimal policy with such a damage function is to emit zero for all time periods.[7] Since we study targets with climate sensitivity uncertainty, the model has a DICE-type damage function. For computational reasons we do not include damage uncertainty, but we note in Section 5.4 how the results would change if damages were also uncertain.

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

$$Y_t = C_t + K_{t+1} - (1 - \delta_k) K_t. \tag{3.3}$$

Period utility is constant relative risk aversion:

$$u = \frac{(C_t/L_t)^{1-\phi} - 1}{1 - \phi}. \tag{3.4}$$

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.

3.2 Climate System

Current period GHG emissions, E_t, 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:

$$E_t = (1 - x_t) \sigma(t) Q_t + B(t). \tag{3.5}$$

[7] Some CEA studies include a penalty function for exceeding the target. Such an approach is equivalent to using a damage function in which damages are zero below the target, and increasing above the target.

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:

$$M_{t+1} - M_B = (1 - \delta_m(t))(M_t - M_B) + E_t. \tag{3.6}$$

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

$$m_{t+1} - 1 = (1 - \delta_m(t))(m_t - 1) + \frac{E_t}{M_B}. \tag{3.7}$$

Radiative forcing of GHGs, F_t, increases the temperature:

$$F_{t+1} = \Omega \log_2(m_{t+1}) + EF(t). \tag{3.8}$$

Here $EF(t)$ is forcing from other sources, which grows exogenously and Ω is the increase in radiative forcing from a doubling of GHGs.

The global mean temperature evolves according to:

$$\hat{T}_{t+1} = \hat{T}_t + \frac{1}{\alpha}\left(F_{t+1} - \frac{\hat{T}_t - \Gamma}{\tilde{\lambda}} + \xi\left(\hat{O} - \hat{T}\right)(t)\right) + \tilde{\nu}_{t+1} \tag{3.9}$$

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:

$$\Delta T_{2\times} = \Omega\tilde{\lambda}. \tag{3.10}$$

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).[8]

[8]More precisely, the initial mean of the climate feedback parameter, defined below, is identical to Roe

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:

$$T_{t+1} = \tilde{\beta}_1 T_t + \beta_2 F_{t+1} + \beta_3 \left(O - T\right)(t) + \tilde{\nu}_{t+1}. \tag{3.11}$$

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 $\tilde{\lambda}$).[9] Uncertain climate feedbacks cause uncertainty in the climate sensitivity (Stocker, Dahe, and Plattner 2013).

3.3 Learning

Following the literature (Roe and Baker 2007, Kelly and Tan 2015), we 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. However, values of the feedback parameter greater than or equal to one imply infinite steady state temperature changes for any positive concentration of GHGs. Therefore, the computational solution requires truncating the prior distribution at a value less than one.[10]

When prior beliefs about the feedback parameter are normal, the prior distributions of the climate sensitivity and $\Delta T_{2\times}$ have fat upper tails (Kelly and Tan 2015). Thus, the probability of relatively large temperature changes from a doubling of GHGs is large relative to the normal distribution.

The weather shock $\tilde{\nu}_t$ occurs at the beginning of each period, before the abatement rate is chosen. We combine the two uncertain terms in equation (3.11) and denote the sum $\tilde{H}$:

$$\tilde{H}_{t+1} = \tilde{\beta}_1 T_t + \tilde{\nu}_{t+1}. \tag{3.12}$$

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:

$$\sigma^2_{\tilde{H}} = T_t^2/\eta_t + 1/\rho. \tag{3.13}$$

The planner observes $H_{t+1} = T_{t+1} - \beta_2 F_{t+1} - \beta_3 \left(O - T\right)(t)$ 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,

and Baker (2007). Our implied $\Delta T_{2\times}$ is slightly lower because of small differences in the other parameters of the model.

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

[10] See Costello, Neubert, Polasky, and Solow (2010) for a justification of truncation. We assume numerically that beliefs evolve ignoring the truncation. See appendix A.6.2 for more details of the truncation.

with mean and precision:

$$\mu_{t+1} = \frac{\eta_t \mu_t + \rho H_{t+1} T_t}{\eta_t + \rho T_t^2}, \tag{3.14}$$

$$\eta_{t+1} = \eta_t + \rho T_t^2. \tag{3.15}$$

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 and η_t.

3.4 Recursive Problem

The planner chooses emission abatement rate, x_t, and capital investment each period to maximize social welfare.

$$W = \max_{k_{t+1}, x_t} \mathrm{E}\left[\sum_{t=0}^{\infty} \exp\left(-\delta_u t\right) L_t u\left(\frac{C_t}{L_t}\right)\right]. \tag{3.16}$$

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, f denote the normal density function, and $s = [k, T, m, t, \mu, \eta]$. The recursive version of the social planning problem is:

$$V(s) = \max_{k', x \in [0,1]} \left\{ u\left(c\right) + \beta\left(t\right) \int_{-\infty}^{\infty} V\left[s'\right] f\left(H|\mu, \frac{1}{\eta}\right) d\tilde{H} \right\}, \tag{3.17}$$

subject to:

$$c = y - \exp\left(\frac{g_A\left(t\right)}{1-\gamma} + g_L\left(t\right)\right) k' + \left(1 - \delta_k\right) k. \tag{3.18}$$

$$T' = \beta_2 F' + \beta_3 \left(O - T\right)\left(t\right) + \tilde{H}', \tag{3.19}$$

$$F' = \Omega \log_2\left(m'\right) + EF\left(t\right), \tag{3.20}$$

$$m' = 1 + \left(1 - \delta_m\left(t\right)\right)\left(m - 1\right) + \frac{E}{M_B}, \tag{3.21}$$

$$E = \left(1 - x\right)\sigma\left(t\right) A\left(t\right)^{\frac{1}{1-\gamma}} L\left(t\right) k^{\gamma} + B\left(t\right). \tag{3.22}$$

$$\mu' = \frac{\eta\mu + \rho\tilde{H}'T}{\eta + \rho T^2}, \tag{3.23}$$

$$\eta' = \eta + \rho T^2. \tag{3.24}$$

$$t' = t + 1. \tag{3.25}$$

Equation (3.17) 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. Time, t, is a state variable which determines all exogenous variables. 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 unbounded. 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.

3.5 Stabilization Targets

As discussed in Section 2, stabilization targets differ primarily by the choice of target variable. Here we analyze a temperature target, which is the most common. The qualitative channels we find apply regardless of the target variable: all targets are by definition inflexible and force policy variables to respond to transient shocks. However, the magnitude of the effects differ depending on, for example, the volatility of any shocks to the target variable.

Other differences in stabilization targets include the probability of exceeding the target and possibly the timing (when the constraint becomes active). The model treats the probability of exceeding the constraint, ω, as a given policy (not optimized). The probabilistic stabilization target is a sequence of constraints:

$$Pr\left(T_{t+1} \geq T^*\right) \leq \omega, \ \forall t = 0, \ldots. \tag{3.26}$$

A pure stabilization target is a special case of equation (3.26) with $\omega = 0$. The constraint is always satisfied if $\omega = 1$, so $\omega = 1$ corresponds to the unconstrained case. The planner must choose abatement to prevent the temperature from exceeding the target in $t + 1$, and chooses abatement anticipating binding targets in future periods. A constraint exists for each period, $t = 1, \ldots$, which restricts the probability that $T_{t+1} \geq T^*$. The constraints limit the temperature change in every period. Alternative choices for the timing, such as compelling the planner to choose abatement today to restrict the probability of exceeding the target in all future periods (as opposed to anticipating today of being compelled in the future), are discussed in Appendix A.2.

The stabilization target is effectively a sequence of constraints on the abatement rate x, since x_t affects T_{t+1}. Appendix A.4 shows that condition (3.26) may be written as:

$$x_t \geq PC\left(s_t, T^*, \omega\right). \tag{3.27}$$

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

Mathematically, PC equals the abatement rate which reduces emissions and therefore GHG concentrations and radiative forcing by just enough to cause the probability of T_{t+1} exceeding T^* to be exactly ω. A stabilization target is therefore equivalent to a minimum abatement rate. For t close to the initial period, $PC \approx 0$. The temperature is very unlikely to exceed 2°C regardless of the emissions level, so the constraint is non-binding. Instead, in the early periods the planner chooses x_t anticipating that constraints in the future will bind with high probability.

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

$$V(s) = \max_{k', x \in [0,1]} \left\{ u(c) + \theta \left[x - PC(s, T^*, \omega) \right] + \beta(t) \int_{-\infty}^{\infty} V[s'] f(H|\mu, \sigma_H) \, d\tilde{H} \right\}, \quad (3.28)$$

subject to equations (3.18)-(3.25) and the Kuhn-Tucker conditions. 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 $i = 1, 2, \ldots$. Constraints on the probability that $T_{t+i+1} \geq T^*$ are implicit in the derivatives of V (shown formally in Appendix A.3), since the planner anticipates in period t facing a probabilistic constraint in $t+i$.

4 Feasibility

4.1 Tightness of the Constraint

Given current information and an emissions path $X_{t-1+i} \equiv \{x_j\}_{j=t \ldots t-1+i}$, an expected probability $\omega(s_t; X_{t-1+i})$, exists in period t of exceeding the target in $t+i$. Appendix A.4 calculates the probabilities $\omega_{t+i}^{\min} = \omega(s_t; X_{t-1+i} = 1)$ and $\omega_{t+i}^{\max} = \omega(s_t; X_{t-1+i} = 0)$, which are the expected probabilities in t of exceeding the target in $t+i$ when the abatement rate is set to 1 and 0, respectively, for all $t, \ldots, t-1+i$, conditional on current information. For $i = 1$, the probabilities satisfy:

$$1 = PC\left(s_t, T^*, \omega_{t+1}^{\min}\right) \quad (4.1)$$

$$0 = PC\left(s_t, T^*, \omega_{t+1}^{\max}\right) \quad (4.2)$$

Values of ω greater than $\omega_{t+i}^{\min}$ are expected to be feasible policies in $t+i$, conditional on current information. For $\omega > \omega_{t+i}^{\max}$, the constraint is expected not to bind in 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.[12] Conversely, any value of $\omega < \omega_{t+i}^{\min}$ implies in expectation $PC(s_{t+i}) > 1$, even

[12]Note that the planner may optimally set $x_t > 0$, even if the constraint is not expected to bind in $t+i$, anticipating that the constraint will bind at some point after $t+i$. Here a low value of $\omega_{t+i}^{\max}$ indicates

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 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 estimated values of $\omega_{t+i}^{\min}$ and $\omega_{t+i}^{\max}$ for $t = 2005$ and $i = 1, \ldots, 95$ from a sample of 10,000 simulations using the calibrated parameter values.

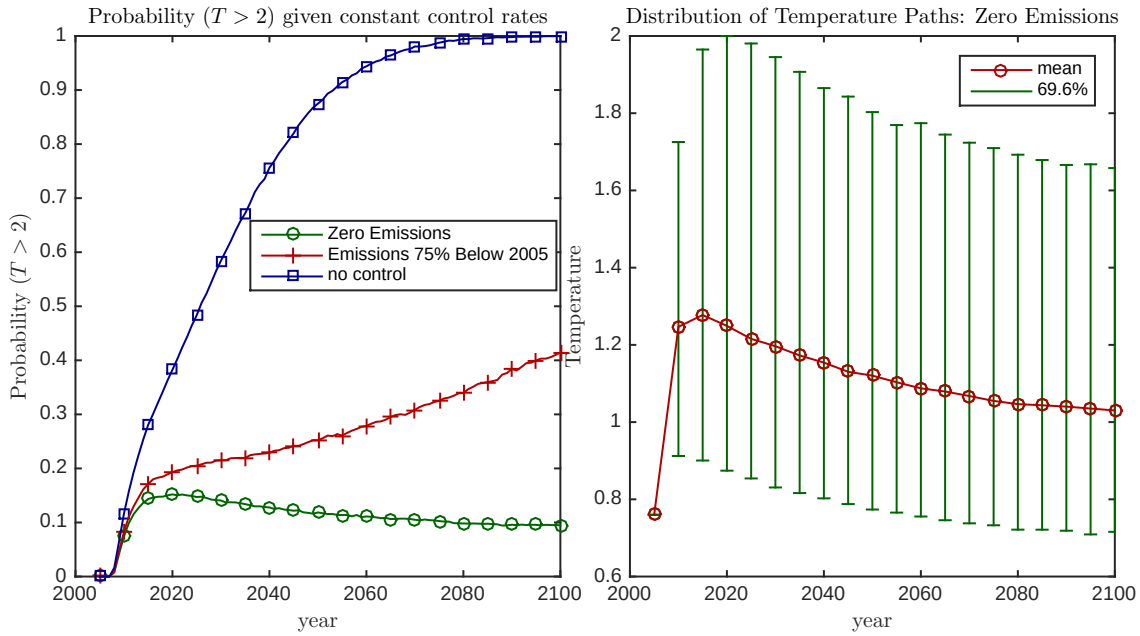

Figure 1: Probability of exceeding $T^* = 2$°C, given various emissions policies. The left panel gives the fraction of 10,000 simulations, each of which draw a random realization of β_1 from the prior distribution and a sequence ν_t, for which the temperature exceeds the target in the given year, for the given abatement policy. The right panel gives the 69.6% error bars, so in 15.2% of the simulations the temperature was greater than or equal to the top error bar. 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 15.2% chance that the 2°C target will eventually be exceeded, even with an immediate, permanent drop to zero emissions. Therefore, values of $\omega < 0.152$ are expected to be infeasible given today's information.

The fat upper tailed uncertainty, calibrated from Roe and Baker (2007), is visible in the right panel of Figure 1, where the error bar extends further above the mean than below. Fat

the constraint is relatively slack in $t+i$, since even a zero abatement policy will in expectation satisfy the constraint in $t+i$ for most values of ω.

tailed uncertainty implies very high values of the climate sensitivity occur relatively more often than under the normal distribution. Therefore, the temperature change exceeds the 2°C target with surprisingly high probability, regardless of the emissions policy.[13]

An immediate, permanent, 75% drop in endogenous emissions below 2005 levels is expected to violate the constraints for any $\omega < 0.41$ by 2100.[14]

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 3.7).

The planner chooses a current emissions policy, anticipating future emissions policies over a period of decades, which is expected to satisfy future constraints, given current information. However, learning and inertia in the climate may cause the temperature to enter the infeasible range, even along a path where all long run probabilistic constraints were expected to be satisfied.

4.2 A Feasible Constraint

Assuming emissions cannot be negative, constraint (3.27) may not be feasible; the set of $x \in [0, 1]$ which also satisfy (3.27) may be empty, violating a necessary condition for the existence of a maximum. The problem is not feasible when $PC(s, T^*, \omega) > 1$, 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 $x = 1$ whenever $PC > 1$. That is, whenever even zero emissions results in the target being exceeded with probability greater than ω, the planner must return to the feasible range as quickly as possible by setting $x = 1$.[15] This is appealing in the sense that

[13] Relatively high values of the feedback parameter cause both GHG concentrations and weather shocks to have more persistent effects on temperature, both of which amplify the variance of temperature over time.

[14] 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.

[15] 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 real world phenomena to calibrate the penalty to. 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. Finally, note that our solution is not equivalent to a sufficiently strong penalty function. Both a very strong penalty function and our assumption imply $x = 1$ when the target is exceeded.

it provides calibrated incentives to reduce emissions in the near term to avoid the jump in welfare losses when the target is crossed. Indeed, catastrophic damages (which we do not model) induce a similar jump in welfare loss beyond T^*. Both catastrophic damages and our assumption imply $x = 1$ when the target is exceeded. Of course, damages and costs interact differently with the other state variables, so the two approaches are not equivalent. Sufficiently convex damages would lead to even lower current emissions as the planner has a stronger incentive to reduce the risk of exceeding the target.

The feasible constrained problem is then:

$$\mathbf{V}(\mathbf{s}) = \max_{\mathbf{k}',\mathbf{x}\in[\mathbf{0},\mathbf{1}]} \left\{ u(c) + \theta \left[x - \min\left\{ PC(s, T^*, \omega), 1 \right\} \right] + \beta(t) \int_{-\infty}^{\infty} V[s'] f(H|\mu, \sigma_H) \, d\tilde{H} \right\}, \tag{4.3}$$

subject to equations (3.18)-(3.25) and the Kuhn-Tucker conditions.

Problem (4.3) 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 (4.3), through the derivatives of $V[s']$ (see Appendix A.3).

The requirement that $x = 1$ when exceeding the target affects near term abatement, even though the temperature is initially well below the target. Abatement costs are convex, so the planner raises current abatement to smooth abatement costs over time. In addition, the model now has an endogenous probability of crossing a "threshold" (the target) into a region with higher abatement costs. This creates two extra incentives to increase near term abatement, which are related to incentives in Lemoine and Traeger (2014) for the case of threshold damages. First, increasing abatement reduces the likelihood of crossing the target and incurring higher costs (the marginal hazard effect). Second, increasing current abatement reduces the climate inertia, which reduces the number of periods for which $x = 1$ upon crossing the threshold (similar to Lemoine and Traeger's differential welfare impact).

The next section analyzes the how the constraints affect optimal abatement and temperature paths, determines how the planner adjusts abatement today in response to future constraints, and determines how learning affects the constrained optimal policy.

But they have different effects on current emissions: if the penalty function is sufficiently strong, current emissions are less as the planner has a stronger incentive to reduce the risk of exceeding the target.

5 Results

5.1 Optimal policy and uncertainty

Appendix A.6 details the solution method. First, we analyze how the optimal abatement policy varies with the probabilistic target and the uncertainty. The left panel of Figure 2 plots the optimal abatement policy for two different true values of β_1: the mean of 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.

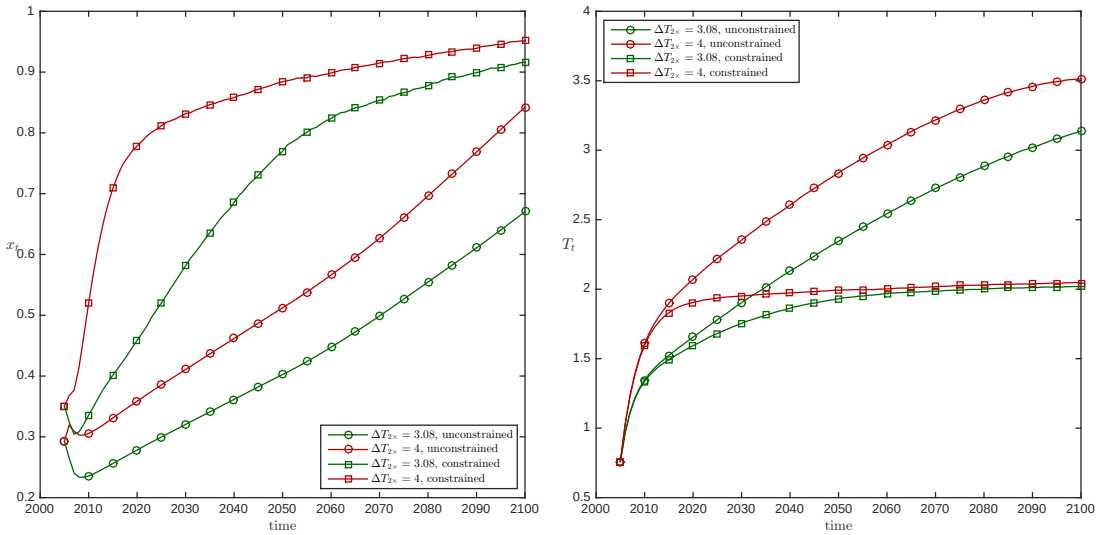

Figure 2: Emissions abatement rate and temperature for $\Delta T_{2\times} = \{3.08, 4\}$, unconstrained and constrained with $\omega = 0.5$. Mean of 10,000 simulations.

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 when the true climate sensitivity equals the mean of the prior. This result is found in other computational models with uncertainty that use the DICE framework (see for example Rudik 2016, Kelly and Tan 2015). Kelly and Tan (2015) 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 mean of the prior. Therefore, when $\Delta T_{2\times} = 3.08$, learning

reduces one motivation for early abatement- to insure against potentially extreme values of the climate sensitivity (note that abatement is monotonic when $\Delta T_{2\times} = 4$). 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.

The right panel of Figure 2 plots the average temperature changes for the above two cases. When $\Delta T_{2\times} = 3.08$, the unconstrained optimal temperature crosses 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 (4.3) for various values of ω, and simulate each solution 10,000 times. The left panel of Figure 3 reports the mean optimal abatement rate for various values of ω, assuming the true climate sensitivity equals the mean of the prior.

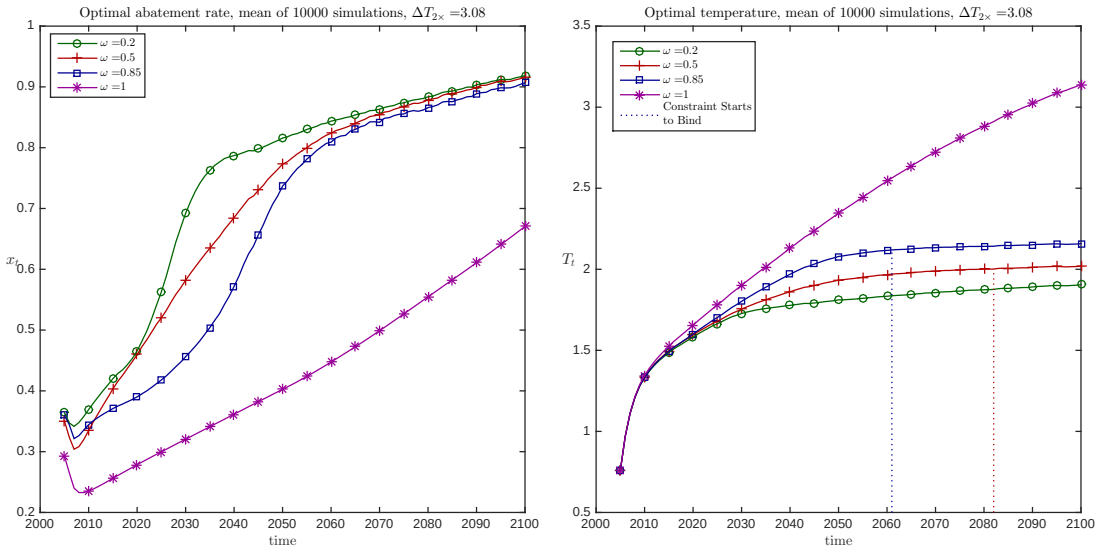

Figure 3: Optimal constrained abatement policy and temperature. Each curve is the mean of 10,000 simulations, with true $\Delta T_{2\times} = 3.08$, and the reported value of ω.

In Figure 3, $\omega = 1$ 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 in Figure 2, because learning reduces the risk that the climate sensitivity is high.

The planner must ramp up abatement spending relatively early when $\omega = 0.2$, since 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.

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.[16] The difference in abatement rates therefore reflects only that the mean maximum temperature differs with ω. The mean maximum temperature is about 0.25° higher when $\omega = 0.85$ versus $\omega = 0.2$.

The unconstrained optimal abatement is far less than the constrained optimal abatement, even for $\omega = 0.85$. Regardless of ω, the maximum 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 is 3.26°C, which requires much less abatement.

[16]Our results are sensitive to the endogenous rate of learning. Kelly and Tan (2015) perform a detailed analysis of the rate of learning in a model with fat tailed uncertainty about the climate sensitivity, and consider a number of extensions, including adding additional uncertainties. Adding additional uncertainties would amplify the effect of targets on emissions and welfare. Learning would slow and the variance of the temperature would increase, both causing the target to be more difficult to adhere to.

The right panel of Figure 3 shows the mean temperature change for various values of ω, assuming the true climate sensitivity equals the mean of the prior. 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 mean of the prior.[17] Smaller values of ω imply a smaller maximum mean temperature change. For $\omega = 0.5$, the mean temperature is very slightly below 2°C after 2120. 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, when $\omega = 0.2$ the temperature should converge to slightly below 1.91 (actually 1.89 by 2120). The planner stays below the target to reduce the likelihood of crossing the target and incurring the cost of $x = 1$ (the marginal hazard effect) and to reduce the number of periods for which $x = 1$ if the target is exceeded.

Overall, the maximum mean temperature is not very sensitive to ω. The difference in maximum temperatures between $\omega = 0.2$ and $\omega = 0.85$ 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, because of the large difference between the target and the unconstrained optimal temperatures.

The right panel of Figure 3 also shows the year in which the constraint just binds, on average. Holding emissions fixed, a smaller value of ω implies the constraint binds earlier. However, a smaller value of ω induces the planner to reduce emissions earlier, which tends to increase the time until the constraint just binds. The right panel of Figure 3 shows that the emissions effect dominates when the true climate sensitivity equals the mean of the prior. Some of the early abatement when the climate sensitivity was uncertain was not necessary *ex post* and so the constraint binds relatively late for low values of ω (the constraint binds after 2100 when $\omega = 0.2$).

The above analysis is for the case where the true climate sensitivity equals the mean of 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, if the climate sensitivity is sufficiently

[17] 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.

high, the temperature exceeds the target, regardless of ω. For example, Figure 4 plots abatement and mean temperature change over time as a function of ω when $\Delta T_{2\times} = 5$.[18]

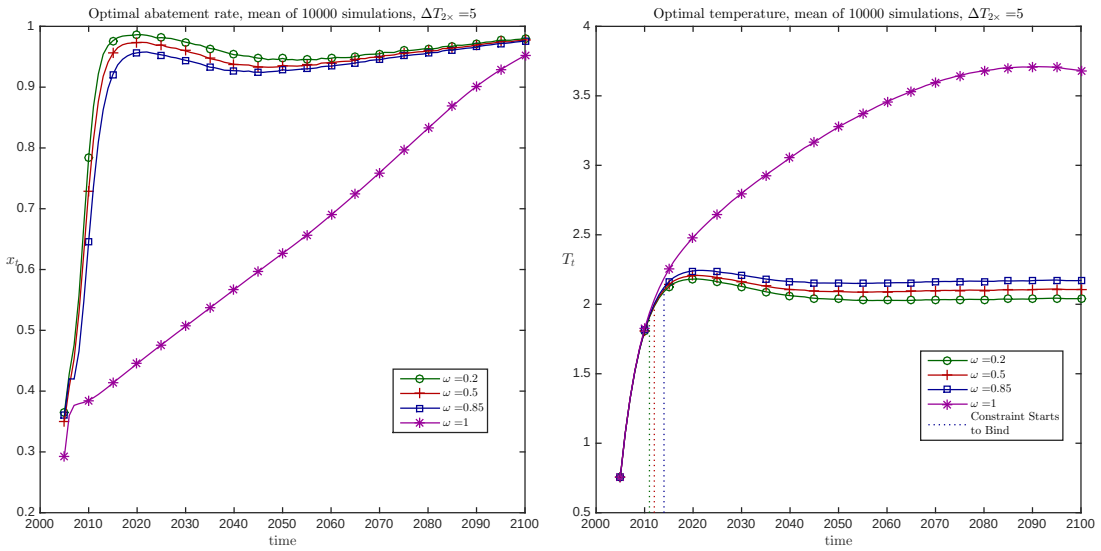

Figure 4: Optimal constrained abatement and temperature. Each curve is the mean of 10,000 simulations, with true $\Delta T_{2\times} = 5$ and the reported value of ω.

In the right panel of Figure 4, 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 3, 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 3 is an average of 10,000 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.

The right panel of Figure 4 also shows that the constraint binds after just a few years. As with the case where the true climate sensitivity equals the prior, as ω decreases holding emissions fixed the constraint binds earlier. However, as ω declines, the planner reduces

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

emissions, which tends to make the constraint bind later. When $\Delta T_{2\times} = 5$, the planner has little time to adjust emissions downward, so the first effect dominates and the constraint binds slightly earlier as ω decreases.

The left panel of Figure 4 shows abatement when $\Delta T_{2\times} = 5$. Because expectations are being revised upward, abatement does not decline initially as in Figure 3.

5.2 Optimal reaction to future constraints

Given a sequence of emissions levels, the planner knows that the temperature may exceed the target in the future. If so, the planner endures a cost of setting emissions equal to zero (and high damages). Ultimately, these costs motivate the planner to reduce emissions today (to smooth abatement costs over time, reduce the likelihood of exceeding the target, and reduce the time during which the target is exceeded). However, abatement today is also costly, and so how much near term abatement, and how often the target is crossed, are computational questions. If abatement is inexpensive, the planner may choose a high level of abatement, paying today to reduce the risk of additional costs in the future. The planner may even elect to increase abatement to the point where the probability of exceeding the target in the future is less than ω, if the cost of abatement is inexpensive relative to the cost of exceeding the target in the future. Conversely, if abatement is expensive, the planner may take on more risk, perhaps even to the point where today the planner expects the target to be exceeded in the future with probability greater than ω.

The probability today of exceeding the target in the future is therefore a constrained optimal policy, trading off these costs. In contrast, ω acts in the future to lower the threshold temperature beyond which the planner incurs the extra costs. Thus reducing ω creates an extra incentive to reduce emissions today, but is different from optimal probability of exceeding the target.

Table 2 gives the percentage of 10,000 simulations for which the temperature exceeds the target. In the unconstrained case the temperature exceeds the target for 72.1% of the simulations in 2050. As ω falls below one, the planner increases current abatement. Since abatement is low and abatement cost is convex, small increases in abatement are not very costly and by increasing abatement the planner reduces the probability of exceeding the target in the future and incurring the cost of setting emissions equal to zero. Thus for $\omega > 0.3$, the planner sets abatement high enough so that the probability of exceeding the target in 2050 is strictly less than ω.

However, given the convexity of abatement costs, eventually the cost of increasing abatement today rises to the point where further reductions in the probability of exceeding the constraint in the future are no longer optimal. For $\omega = 0.2$, the target is still exceeded approximately 35% of the time in 2050 (from Figure 1, an immediate 75% reduction in

emissions cause the target to be exceeded with probability 0.25 in 2050). The planner is unwilling to pay the cost of more severe emissions reductions today and accepts the penalty of occasionally exceeding 2° and having to set future emissions equal to zero. The probability of exceeding the target is therefore constrained optimal in the sense that the planner chooses a future probability of exceeding the target today given the costs and benefits of reducing that probability.

Finally, the last column of Table 2 shows the autocorrelation of periods in which the target is exceeded. The correlation is very positive: once the target is exceeded, inertia typically causes the the temperature to exceed the target in subsequent periods. In addition, only one true value of the climate sensitivity exists. Given an emissions path, if the climate sensitivity is surprisingly high, then probability of exceeding the target rises in every period. So an increase in the probability of exceeding the target in one period is correlated with exceeding the target in subsequent periods.

5.3 Learning

Learning plays an important role in the presence of a stabilization target. If learning reduces uncertainty relatively quickly, the planner can adjust abatement and the temperature is less likely to exceed the target. Further, learning directly affects the stringency of the probabilistic constraint.

With regard to the speed of learning, from equation (3.24), the precision of the feedback parameter grows approximately linearly, which implies the variance shrinks according to a power law (t^{-1}). Nonetheless, the rate of learning about the climate sensitivity differs from the rate of learning about the feedback parameter. From equation (3.10):

$$\Delta T_{2\times} = \Omega\tilde{\lambda} = \frac{\Omega}{\alpha\left(1-\tilde{\beta}_1\right)}, \tag{5.1}$$

so small differences in β_1 can imply relatively large differences in λ and $\Delta T_{2\times}$. Since the sensitivity of the climate to GHG concentrations is ultimately more important for determining abatement policy, especially with a stabilization target, we focus on the rate of learning about $\Delta T_{2\times}$ rather than the feedback parameter.

Kelly and Tan (2015) show that overall learning about the climate sensitivity is a slow process, although the planner can effectively rule out values of the climate sensitivity much higher than the true value relatively quickly. Figures 5 and 6 confirm these results. The left panel of Figure 5 plots the evolution of the mean of the prior for 200 simulations where the true $\Delta T_{2\times}$ equals the prior. The percentage of simulations for which the estimate of $\Delta T_{2\times}$ is above 3.5° falls from 28.5% initially to 2% in only 14 years. However, even in 2050 significant uncertainty remains. The difference between the largest and smallest value of the estimate

of $\Delta T_{2\times}$ is 0.42° in 2050. Further, learning slows as the true value increases. The right panel shows that if the $\Delta T_{2\times} = 5°$, in 2050 the difference between the largest and smallest estimate of $\Delta T_{2\times}$ is 0.89°.

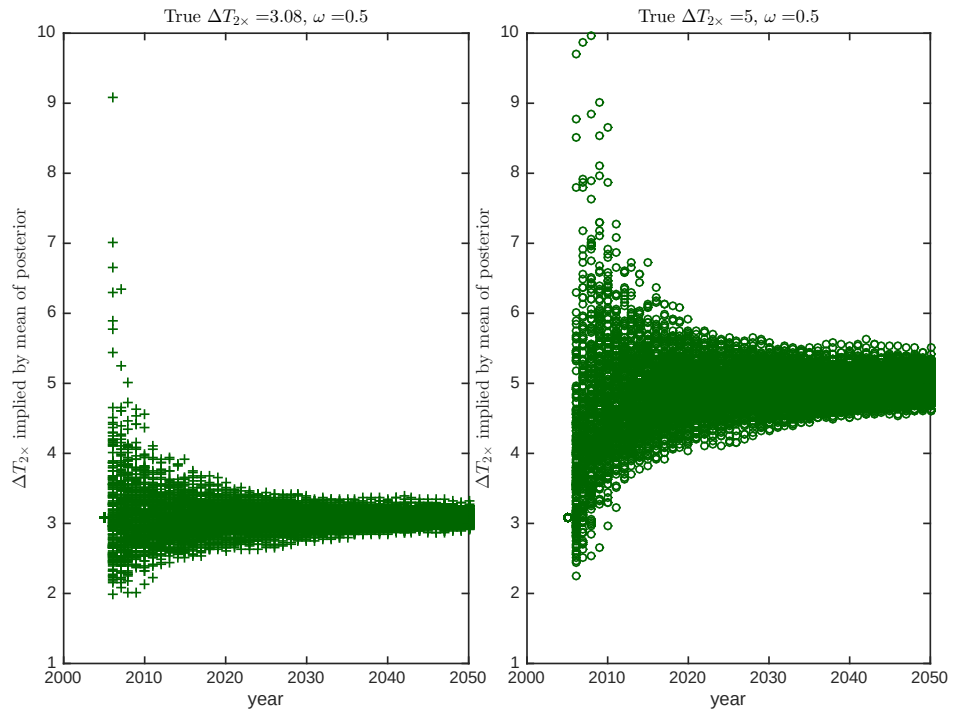

Figure 5: Evolution of the estimate of $\Delta T_{2\times} = 5°$ for 200 simulations when $\omega = 0.5$ and $\Delta T_{2\times}$ equals 3.08° (the prior) and 5°.

Alternatively, Figure 6 shows that the standard deviation of the estimate of $\Delta T_{2\times}$ falls initially as very high values are ruled out,[19] but then falls at a much slower rate after.

[19]However, if the true $\Delta T_{2\times}$ is greater than the prior, the standard deviation of $\Delta T_{2\times}$ does not decrease monotonically. The standard deviation of the feedback parameter does decrease monotonically (equation 3.24), but increases in the mean estimate of the feedback parameter implies small changes in the feedback parameter have larger changes in $\Delta T_{2\times}$. Therefore, increases in the mean of the feedback parameter estimate, which occur when the true value exceeds the prior, increase the standard deviation of the estimate of $\Delta T_{2\times}$ (equation 5.1).

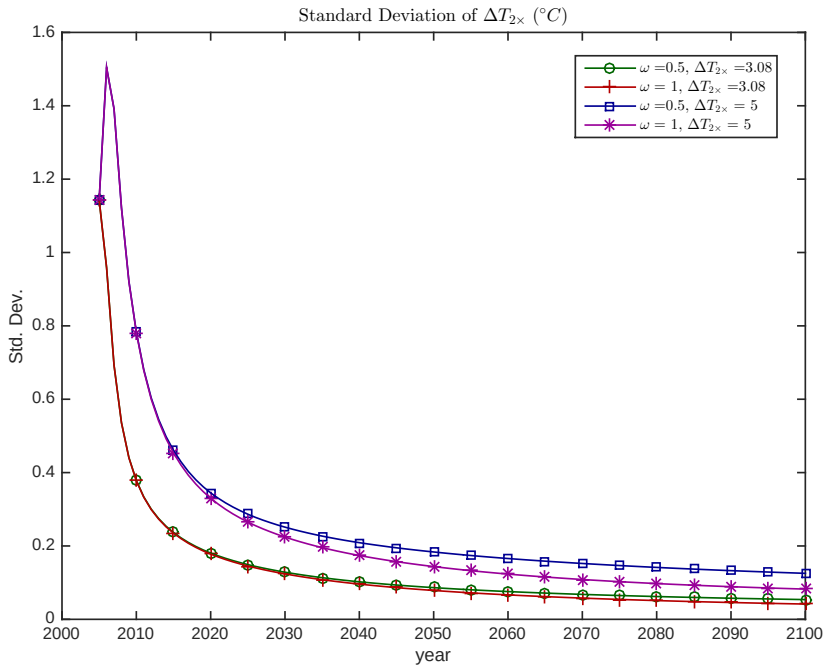

Figure 6: Evolution of the standard deviation of the prior belief distribution of $\Delta T_{2\times}$. Mean of 10,000 simulations.

Figure 6 also shows that the stabilization target slows the rate of learning. Because the probabilistic constraint reduces emissions, the climate change signal is less visible in the noisy weather shocks, which slows learning. When the true $\Delta T_{2\times} = 3.08$°C, the standard deviation of the estimate of $\Delta T_{2\times}$ is 9.8% higher in 2050 when the stabilization target is present. When the true $\Delta T_{2\times} = 5$°C, the standard deviation of the estimate of $\Delta T_{2\times}$ is 27.7% higher in 2050 when the stabilization target is present. Thus the stabilization target slows learning most when learning is most helpful for keeping the temperature below the target.

In turn, slow learning has two offsetting effects on achieving the temperature target, when the true climate sensitivity is greater than the mean of the prior. First, because the mean of the prior is below the true climate sensitivity, the planner underestimates the amount of abatement needed to meet the temperature target. Second, when learning is slow the planner keeps abatement high to reduce the risk of exceeding the target. The former effect must dominate when the climate sensitivity is sufficiently high. The true climate sensitivity is unobserved, and so a higher climate sensitivity will not change the second effect, but will worsen the first effect. Thus, slow learning and the inertia in the climate explain why the temperature exceeds the target when the climate sensitivity is relatively high (Figure 4).[20]

[20] Although consistent with the previous literature, learning on a scale of decades may still appear relatively rapid given that many climatic processes evolve over a period of centuries. Indeed, our two equation climate model with a single uncertainty likely overstates the rate of learning in complex climate models with many

The primary channel through which learning affects abatement policy is the effect of reducing the variance of the feedback parameter on the probabilistic constraint. Since the constraint binds frequently, changes in the constraint induced by a reduction in the variance of uncertainty directly affect the abatement rate. It is straightforward to verify that:

$$\frac{\partial PC\left(s_t, T^*, \omega\right)}{\partial 1/\eta_t} > (<, =)0 \;\Leftrightarrow\; \omega < (>, =)0.5. \tag{5.2}$$

Therefore, how learning (reducing the variance, $1/\eta$, of the feedback parameter) affects abatement depends on ω. In fact, learning changes abatement so that the temperature moves closer to the target, T^*.

First, suppose that $\omega < 0.5$, so that the constraint limits the temperature to a level below T^* (see the green line in the right panel of Figure 3). Recall from equation (3.19) that T_{t+1} is normally distributed. Since the normal distribution is symmetric and the temperature is below T^*, we have $\mathrm{E}\left[T_{t+1}\right] < T^*$. Next, learning reduces the variance of the estimate of the feedback parameter, which means extreme realizations of the climate sensitivity, which imply $T_{t+1} > T^*$, are less likely. Therefore, learning decreases $Pr\left(T_{t+1} \geq T^*\right)$. The constraint relaxes (PC decreases), and (if the constraint is binding) x_t falls and T_t rises closer to T^*. Therefore, for $\omega < 0.5$, learning increases the temperature by decreasing abatement.

Suppose now that $\omega > 0.5$, so that the constraint limits the temperature to a level above T^* (blue line in the right panel of Figure 3). Since T_{t+1} is normally distributed and the temperature exceeds T^*, we have $\mathrm{E}\left[T_{t+1}\right] > T^*$. Now learning makes extremely high realizations of the climate sensitivity less likely, but any realization above the mean still implies $T_{t+1} > T^*$, so $Pr\left(T_{t+1} \geq T^*\right)$ is unaffected by reducing the probability of extremely high realizations. However, realizations slightly less than the mean are more likely. Since $T^* < \mathrm{E}\left[T_{t+1}\right]$, $Pr\left(T_{t+1} \geq T^*\right)$ increases. The constraint tightens, so (if the constraint binds) x_t increases and T_t decreases. Therefore, for $\omega > 0.5$ learning decreases the temperature by increasing abatement.

In either case, learning gives the planner better control of the temperature, allowing the planner to move closer to the target. For the case of $\omega = 0.5$, reducing the variance has no effect on PC, since the normal distribution is symmetric.

Figure 7 confirms the above intuition. For $\omega = 0.25 < 0.5$, the abatement falls and the temperature rises, moving closer to the target relative to no learning.[21] For $\omega = 0.75 > 0.5$, abatement rises and the temperature falls, but is still moving closer to the target relative to no learning. For $\omega = 0.5$ the constraint is unaffected by learning and, since the constraint binds both with and without learning, the temperatures are almost identical. Note the true

uncertainties. Introducing additional uncertainties which slow learning would magnify the results here, since additional uncertainties would make a stabilization target more difficult to adhere to.

[21] Note that for the case without learning, the model is re-solved so that the planner anticipates no further information about the climate sensitivity is forthcoming.

climate sensitivity equals the prior in Figure 7, which allows us to focus on the effect of learning on the variance of uncertainty in isolation of the effect of learning on the mean of the prior.

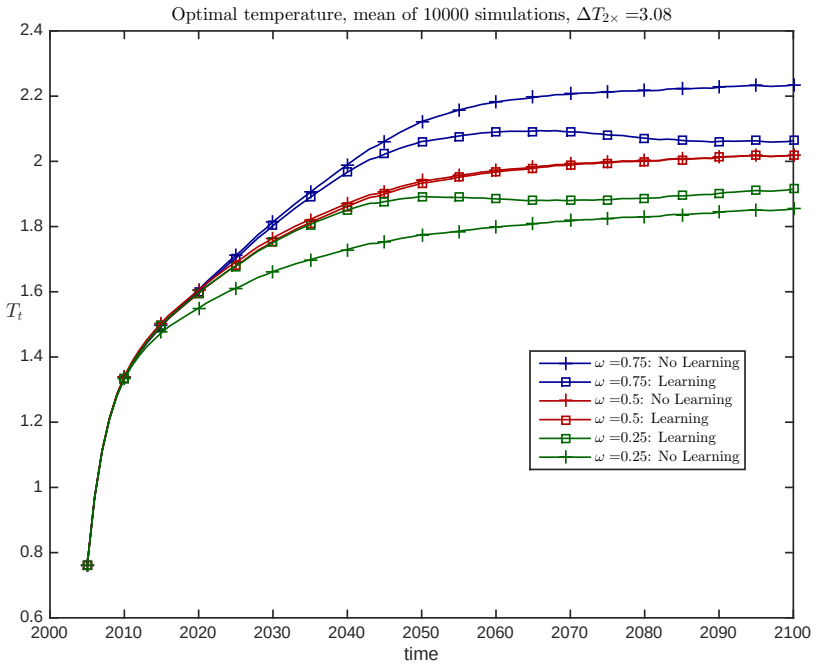

Figure 7: Optimal constrained temperature. Each curve is the mean of 10,000 simulations, with true $\Delta T_{2\times} = 3.08$ and the reported value of ω.

5.4 Welfare loss

The probabilistic constraint at least weakly restricts the choice set for the planner, and therefore must result in a welfare loss, relative to the unconstrained problem.[22] Our interest is in how the welfare loss varies with ω and how uncertainty affects the welfare loss.

The left panel of Figure 8 graphs the welfare loss as a percentage of the welfare of the unconstrained problem, $\omega = 1$.

[22]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.

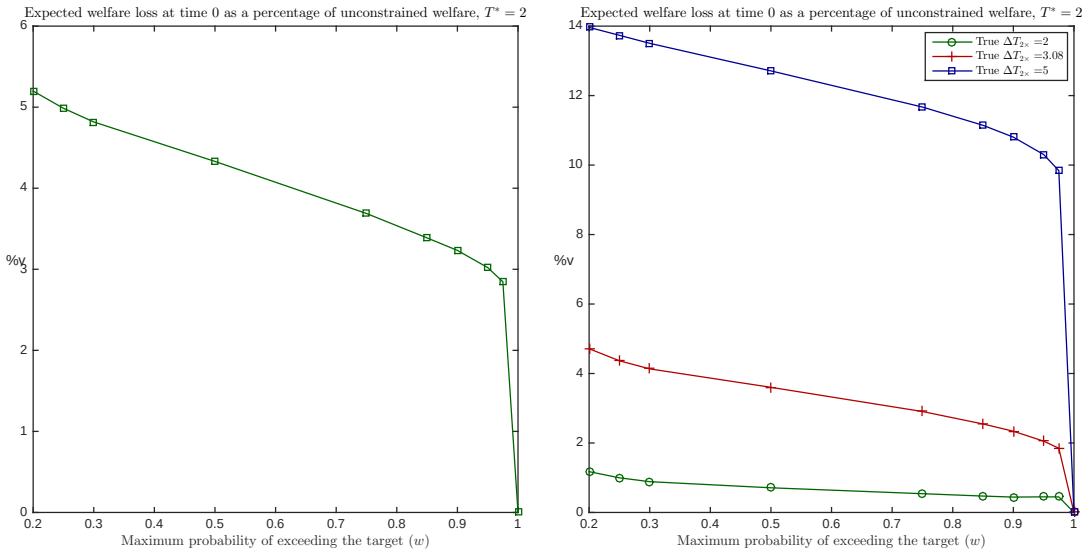

Figure 8: Welfare loss as a function of ω (left panel) and as a function of ω and the true $\Delta T_{2\times}$ (right panel). The graph plots $1 - V(s_0; \omega)/V(s_0; 1)$, where V is the solution to (3.28), as a function of ω. Left panel: mean of 10,000 simulations, each with a draw of the feedback parameter from the prior distribution and a draw of the weather shock for each period in the simulation. Right panel: mean of 10,000 simulations, each with the feedback parameter implied from the given value of the climate sensitivity and a draw of the weather shock for each period in the simulation.

From the left panel of Figure 8, 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 the right panels of Figures 3 and 4). Welfare loss is about 3-5% of the unconstrained policy, depending on ω (Table 3 reports the exact welfare losses).

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 target recommendation.[23] Table 3 then indicates that the baseline stabilization target results in a welfare loss of approximately 4.7%.[24]

The right panel of Figure 8 plots the welfare loss for various true values of $\Delta T_{2\times}$. The welfare loss increases with the true value of $\Delta T_{2\times}$, rising to 14% when $\Delta T_{2\times} = 5$ and $\omega = 0.2$.[25]

[23]One may alternatively interpret the IPCC statement as the optimal probability of exceeding the target at some point in the future is 0.33. If so, from Table 2, the appropriate calibration is $\omega \approx 0.2$.

[24]The welfare loss in consumption equivalents is approximately 1%. That is, the household would require an annual consumption increase of 1% in perpetuity to be willing to accept the stabilization constraint (see Table 3).

[25]Note that as the true climate sensitivity decreases, the full information optimal maximum temperature

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.[26] 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,[27] 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 tipping points, the lack of intrinsic preferences for low temperatures, 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 3 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 mean of the prior is smaller under uncertainty than the difference between the maximum unconstrained temperature and the maximum constrained temperature under certainty. Therefore, the 2° target is not as restrictive when the true

decreases, eventually to below two degrees. Yet, a small welfare loss still results. If the climate sensitivity is low, with full information the planner reduces abatement (left panel of Figure 2). The temperature still approaches a low maximum temperature because the climate sensitivity is low. In contrast, the constraint causes the planner to *increase* near term abatement relative to the unconstrained case, creating a welfare loss as the planner undershoots the optimal temperature.

[26] 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.

[27] 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.

climate sensitivity equals the mean of 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 3 indicates the welfare loss under uncertainty is 82.8% greater than the model with certainty for $\omega = 0.2$, although for $\omega = 0.975$ 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.[28] Table 3 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, the right panel of Figure 3 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

[28] 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.

learning in the model.

6 Conclusions

In this paper, we evaluated the common policy recommendation of a 2°C temperature target in the context of an integrated assessment model with uncertainty and learning about the climate sensitivity. Uncertainty fundamentally changes the feasibility and welfare implications of a stabilization target. We show that even with an immediate reduction of all GHG emissions to zero, the temperature eventually exceeds 2°C with probability 0.15. Given that the temperature is already 0.8°C above pre-industrial and climatic processes are highly uncertain and subject to inertia, it is difficult to envision controlling the climate to the precise degree implied by the 2°C target in the short run. Nonetheless, we show that as learning resolves the planner can eventually achieve the target with sustained emissions reductions over time. An immediate policy implication is that implementing a stringent emissions reduction policy is imperative if the temperature change is to be kept under 2°C: even with an immediate 75% reduction in emissions, the 2°C target is exceeded with probability 0.41 by 2100.[29]

Further, we show that adding uncertainty changes the welfare effect of stabilization targets 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, the temperature randomly evolves over time, occasionally exceeding the target. In this case, the planner must set an excessively high abatement rate to immediately return the temperature back to the target. Third, 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.

An important policy implication of our results is that introducing some flexibility into the target can reduce the welfare loss. For example, many regulations and international agreements have provisions whereby the regulation comes up for renewal after a prespecified number of years, at which point the regulation may be changed to reflect new information. Another policy implication is that a less variable target, such as GHG concentrations, is more attractive than temperature. The welfare loss of the target increases with the variance of the shock, through the overreaction channel. Since GHGs are less variable than temperature, less overreaction would occur.

Our model may be extended in a number of ways. We could consider other targets such

[29] Indeed, some scientists now argue that exceeding the 2°C target is inevitable (Sanford, Frumhoff, Luers, and Gulledge 2014).

as limit GHG concentrations to 350 ppm or limiting sea level rise or ocean acidification. Our approach would 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:

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

TFP growth:

$$A(t) = A_0 \exp\left[g_{A,0}\frac{1 - \exp\left[-\delta_A t\right]}{\delta_A}\right]. \tag{A.2}$$

Backstop cost:

$$\Psi(t) = \frac{\sigma(t)}{a_2} a_0 \left(1 - \frac{1 - \exp\left[g_\Psi t\right]}{a_1}\right). \tag{A.3}$$

Emissions intensity of output:

$$\sigma(t) = \sigma_0 \exp\left[g_{\sigma,0}\frac{1-\exp\left[-\delta_\sigma t\right]}{\delta_\sigma}\right]. \tag{A.4}$$

Exogenous emissions:

$$B(t) = B_0 \exp\left[-\delta_B t\right]. \tag{A.5}$$

Decay rate of GHGs:

$$\delta_m(t) = \bar{\delta}_m + \left(\delta_{m,0} - \bar{\delta}_m\right)\exp\left[-\delta_m^* t\right]. \tag{A.6}$$

Exogenous forcing:

$$EF(t) = EF_0 + 0.01\left(EF(100) - EF_0\right)\cdot\min\left\{t, 100\right\}. \tag{A.7}$$

Heat uptake from ocean:

$$O(t) = \max\left\{\Delta_{T1} + \Delta_{T2}t + \Delta_{T3}t^2, 0\right\}. \tag{A.8}$$

Discount factor:

$$\beta(t) = \exp\left[-\delta_u + \frac{1-\phi}{1-\gamma}\log\left(\frac{A(t+1)}{A(t)}\right) + \log\left(\frac{L(t+1)}{L(t)}\right)\right]. \tag{A.9}$$

A.2 Alternative Timing of the Target

Here we consider alternatives for the timing of the target. One alternative requires that the temperature limit holds only in the steady state. However, in the DICE the backstop cost $\Psi(t)$ eventually decreases to the point where zero emissions becomes optimal. Thus, under the optimal solution, the climate eventually returns back to prehistoric temperature levels. Therefore, a steady state temperature limit of 2°C is not binding and the unconstrained solution results.[30]

A second alternative is, at time t, to restrict the probability of exceeding the target in all future periods:

$$Pr\left(\text{ANY}\left\{T_{t+1} \geq T^*, T_{t+2} \geq T^*, \ldots\right\}\right) \leq \psi \;\forall\; t = 1, \ldots. \tag{A.10}$$

In contrast (3.26) restricts the probability of exceeding the target in $t+2$, but not until

[30]Intuitively, a steady state constraint is non-binding more generally. Since the transversality condition only requires T_t remain bounded, the temperature converges to the unconstrained solution, with a plan to satisfy the constraint at $t = \infty$, which never arrives.

$t+1$, and so on. Of course, the planner given (3.26) anticipates future constraints and optimally increases abatement today in response to future constraints, because smoothing abatement is optimal. Condition (A.10) is at least as stringent, because the planner *must* address anticipated future temperature changes earlier.

A natural question is how much more strict (A.10) is relative to (3.26). Condition (A.10) holds if and only if:

$$Pr\left(\text{ALL}\left\{T_{t+1} < T^*, T_{t+2} < T^*, \ldots\right\}\right) \geq 1-\psi, \tag{A.11}$$

$$Pr\left(T_{t+1} < T^*\right) Pr\left(\text{ALL}\left\{T_{t+2} < T^*, T_{t+3} < T^*, \ldots\right\} | T_{t+1} < T^*\right) \geq 1-\psi, \tag{A.12}$$

$$\left(1 - Pr\left(T_{t+1} \geq T^*\right)\right) Pr\left(\text{ALL}\left\{T_{t+2} < T^*, T_{t+3} < T^*, \ldots\right\} | T_{t+1} < T^*\right) \geq 1-\psi, \tag{A.13}$$

$$Pr\left(T_{t+1} \geq T^*\right) \leq \frac{A\left(s\right) - \left(1-\psi\right)}{A\left(s\right)} \tag{A.14}$$

$$A\left(s\right) \equiv Pr\left(T_{t+2} < T^* \text{ AND } T_{t+3} < T^* \ldots | T_{t+1} < T^*\right). \tag{A.15}$$

Thus, if the economy is at state s, constraints (3.26) and (A.10) are equivalent if:

$$\omega = \frac{A\left(s\right) - \left(1-\psi\right)}{A\left(s\right)} \tag{A.16}$$

So constraint (3.26) is equivalent at state s to a version of (A.10) with a (weaker) value of $\psi > \omega$. Note that the comparison is not exact because the value of ψ which equates the constraints varies over the state space.

Unfortunately, (A.10) is computationally infeasible. Nonetheless, equation (A.16) implies that, as $\omega \to 1$, $\psi \to \omega$. Further, the difference widens as ω decreases. Thus one can get a rough idea about constraint (A.10) with probability ψ by looking at the results for smaller values of ω. While no policy discussion exists regarding (A.10) versus (3.26), simulations indicate that constraint (A.10) at the calibrated value of $\omega = 0.33$ will likely imply far more abatement than is indicated in most policy discussions, favoring our choice of (3.26) over (A.10).

A.3 Presence of the Constraint in the Value Function Derivatives

The envelope theorem implies that:

$$V_s\left(s_t\right) = u_c\left(c\left(s_t\right)\right) c_s\left(s_t\right) - \theta_t PC_s\left(s_t, T^*, \omega\right) + \beta\left(t\right) \mathrm{E}\left[V_s\left(G\left(s_t\right)\right) G_s\left(s_t\right)\right]. \tag{A.17}$$

Here subscripts on functions denote derivatives, $s_{t+1} = G\left(s_t\right)$ is the law of motion for the state variables (given by 3.19 - 3.25), and decision variables are evaluated at the optimum.

Note that (A.17) also implies:

$$V_s(s_{t+1}) = u_c(c(s_{t+1}))\, c_s(s_{t+1}) - \theta_{t+1} PC_s(s_{t+1}, T^*, \omega) + \beta(t+1)\, \mathrm{E}\left[V_s(G(s_{t+1}))\, G_s(s_{t+1})\right]. \tag{A.18}$$

Substituting (A.18) into (A.17) implies:

$$\begin{aligned} V_s(s_t) =& u_c(c(s_t))\, c_s(s_t) - \theta_t PC_s(s_t, T^*, \omega) + \beta(t)\, \mathrm{E}\left[u_c(c(s_{t+1}))\, c_s(s_{t+1})\, G_s(s_t)\right] \\ & - \beta(t)\, \mathrm{E}\left[\theta_{t+1} PC_s(s_{t+1}, T^*, \omega)\, G_s(s_t)\right] \\ & + \beta(t)\, \beta(t+1)\, \mathrm{E}\left[V_s(G(s_{t+1}))\, G_s(s_{t+1})\, G_s(s_t)\right]. \end{aligned} \tag{A.19}$$

Thus, the derivative of the constraint the planner faces in period $t+1$ is part of the derivative of the value function in period t. Continuing to iterate forward implies that all future constraints are present in the derivative of the value function at time t. Therefore, the planner considers the impact of today's decisions on future constraints.

Equation (A.19) holds for the derivative with respect to each of the state variables, but is most useful when thinking about the envelope equation for temperature. Since PC is increasing in temperature, if the constraint binds ($\theta_t > 0$), the constraint reduces (makes more negative) the marginal value of temperature. In turn, the first order condition implies that reducing the marginal value of temperature increases optimal abatement today.

A.4 Derivation of Equation (3.27)

Rewriting the left hand side of (3.26) for $i = 1$ gives:

$$Pr(T_{t+1} \geq T^*) = Pr\left(\beta_2 F_{t+1} + \beta_3 (O - T)(t) + \tilde{H}_{t+1} \geq T^*\right), \tag{A.20}$$

$$= Pr\left(\tilde{H}_{t+1} \geq T^* - \beta_2 F_{t+1} - \beta_3 (O - T)(t)\right), \tag{A.21}$$

$$= 1 - \mathrm{NCDF}\left[\frac{T^* - \beta_2 F_{t+1} - \beta_3 (O - T)(t) - \mu_t T_t}{\sqrt{\frac{T_t^2}{\eta_t} + \frac{1}{\rho}}}\right]. \tag{A.22}$$

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

$$P_{t+1} \equiv \beta_2 F_{t+1} + \beta_3 (O - T)(t) + \mu_t T_t \tag{A.23}$$

be the expected (predicted) value of T_{t+1}, then using equation (3.13):

$$Pr(T_{t+1} \geq T^*) = 1 - \mathrm{NCDF}\left[\frac{T^* - P_{t+1}}{\sigma_{H,t}}\right]. \tag{A.24}$$

Constraint (3.26) is therefore equivalent to:

$$\text{NCDF}\left[\frac{T^{*}-P_{t+1}}{\sigma_{H,t}}\right]\geq 1-\omega, \tag{A.25}$$

$$\frac{T^{*}-P_{t+1}}{\sigma_{H,t}}\geq \text{NCDF}^{-1}\left[1-\omega\right], \tag{A.26}$$

$$F_{t+1}\leq \frac{1}{\beta_{2}}\left(T^{*}-\beta_{3}\left(O-T\right)\left(t\right)-\mu_{t}T_{t}-\sigma_{H,t}\cdot\text{NCDF}^{-1}\left[1-\omega\right]\right). \tag{A.27}$$

Equations (3.20)-(3.22) then imply (A.27) reduces to:

$$\begin{aligned}x_{t}\geq & 1+\frac{M_{B}}{\sigma\left(t\right)L\left(t\right)A\left(t\right)^{\frac{1}{1-\gamma}}k_{t}^{\gamma}}\left(1+\left(1-\delta_{m}\left(t\right)\right)\left(m_{t}-1\right)-\right.\\ & \left.\exp\left\{\frac{\log\left(2\right)}{\Omega\beta_{2}}\left[T^{*}-\beta_{3}\left(O-T\right)\left(t\right)-\mu_{t}T_{t}-\sigma_{H,t}\cdot\text{NCDF}^{-1}\left[1-\omega\right]\right]\right\}\right),\end{aligned} \tag{A.28}$$

$$x_{t}\geq PC\left(s_{t},T^{*},\omega\right). \tag{A.29}$$

Here PC is the right hand side of (A.28).

We now calculate the highest possible probability for which $T_{t+1}\geq T^{*}$, which occurs with a zero abatement rate. In this case, we have from (3.21):

$$E_{t}^{\max}=\sigma\left(t\right)A\left(t\right)^{\frac{1}{1-\gamma}}L\left(t\right)k_{t}^{\gamma}+B\left(t\right), \tag{A.30}$$

$$m_{t+1}^{\max}-1=\left(1-\delta_{m}\left(t\right)\right)\left(m_{t}-1\right)+\frac{E_{t}^{\max}}{M_{B}}, \tag{A.31}$$

$$F_{t+1}^{\max}=\Omega\log_{2}\left(m_{t+1}^{\max}\right)+EF\left(t\right). \tag{A.32}$$

Next, from (3.19):

$$T_{t+1}^{\max}=\tilde{H}_{t+1}+\beta_{2}F_{t+1}^{\max}+\beta_{3}\left(O-T\right)\left(t\right), \tag{A.33}$$

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

$$\omega_{t+1}^{\max}=Pr\left(T_{t+1}^{\max}\geq T^{*}\right), \tag{A.34}$$

$$\omega_{t+1}^{\max}=Pr\left\{\tilde{H}_{t+1}\geq T^{*}-\beta_{2}F_{t+1}^{\max}-\beta_{3}\left(O-T\right)\left(t\right)\right\}, \tag{A.35}$$

$$\omega_{t+1}^{\max}=1-\text{NCDF}\left(\frac{T^{*}-P_{t+1}^{\max}}{\sigma_{H,t}}\right),\text{ where} \tag{A.36}$$

$$P_{t+1}^{\max}\equiv\mu_{t}T_{t}+\beta_{2}F_{t+1}^{\max}+\beta_{3}\left(O-T\right)\left(t\right). \tag{A.37}$$

Since the difference between the target and the expected temperature in period $t+1$ under

maximum emissions is not infinite, we have immediately from (A.36) that $\omega^{\max} < 1$.

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

$$E_t^{\min} = B(t), \tag{A.38}$$

$$m_{t+1}^{\min} - 1 = (1 - \delta_m(t))(m_t - 1) + \frac{E_t^{\min}}{M_B}, \tag{A.39}$$

$$F_{t+1}^{\min} = \Omega \log_2\left(m_{t+1}^{\min}\right) + EF(t), \tag{A.40}$$

$$T_{t+1}^{\min} = \tilde{H}_{t+1} + \beta_2 F_{t+1}^{\min} + \beta_3 (O - T)(t), \tag{A.41}$$

$$\omega_{t+1}^{\min} = 1 - \text{NCDF}\left(\frac{T^* - P_{t+1}^{\min}}{\sigma_{H,t}}\right), \text{ where} \tag{A.42}$$

$$P_{t+1}^{\min} \equiv \mu_t T_t + \beta_2 F_{t+1}^{\min} + \beta_3 (O - T)(t). \tag{A.43}$$

A.5 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.

		Year			Auto-
		2005	2050	2075	correlation
	0.2	0	25.3	34.8	0.67
	0.25	0	30.4	37.3	0.67
	0.3	0	32.6	41.4	0.66
	0.5	0	35.5	50.6	0.67
ω	0.75	0	49.5	62.7	0.70
	0.85	0	51.8	69.1	0.73
	0.9	0	52.7	71.0	0.74
	0.95	0	53.8	72.8	0.76
	0.975	0	54.9	74.5	0.77
	1	0	72.1	93.6	0.88

Table 2: Columns 2-4: percent simulations for which $T_t > T^*$, 10,000 simulations, each with a draw of the feedback parameter from the prior distribution and a weather shock draw for each period in the simulation. Let $q_{it} = 1$ if and only if $T_{it} > T^*$ for time t and simulation i. Last column: correlation between q_{it} and $q_{i,t-1}$.

		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
(baseline case)	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 3: Decomposition of the welfare loss, mean of 10,000 simulations, each with a draw of the feedback parameter from the prior distribution and a weather shock draw for each period in the simulation. The second and third columns report $1 - V(s_0;\omega)/V(s_0;1)$. In the third column, only the evolution of the mean and variance of the prior are affected by the weather shocks. That is, $s' = [k', m', T'(\nu = 0), t', \mu'(\nu), \eta'(\nu)]$.

A.6 Solution

A.6.1 Solution Outline

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. Kelly and Tan (2015) give a detailed explanation of this solution method. The solution method requires specifying a maximum climate sensitivity, which is discussed in Section A.6.2. The constraint, $x \geq PC(s)$, is linear in the controls and is therefore straightforward to add to the optimization routine in the solution method (e.g. Matlab's optimization routines automatically generate the Kuhn-Tucker conditions). Potential points of non-differentiability in the value function caused by the constraint are discussed in Section A.6.3.

Table 4 gives the grid points used in the solution algorithm. Grid points are selected by a trial and error procedure starting with a relatively sparse grid and then adding grid points where the value function has significant curvature. Other important considerations include choosing grid points near the initial condition and including a zero-variance grid point for analysis of the certainty case.

State	Grid Points
k	0.53, 0.75, 1.18, 1.81, 2.62, 3.59, 4.69, 5.87, 7.12, 8.38, 9.63, 10.81, 12.30, 13.69, 14.97
T	0.07, 0.59, 1.48, 2.52, 3.41, 3.93, 6.00, 7.00
m	0.97, 1.29, 1.84, 2.48, 3.03, 3.35
$\hat{t}$	0.01, 0.05, 0.15, 0.27, 0.45, 0.55, 0.65, 0.75, 0.85, 0.90, 0.95, 0.98, 0.99
μ	0.40, 0.55, 0.65, 0.75, 0.80, 0.96
$\frac{1}{\eta}$	0, 0.0050, 0.0100, 0.0169
Total grid points/spline basis functions: 224,640	

Table 4: Grid points. k: normalized capital stock per productivity adjusted person, in thousands of dollars per productivity adjusted person. T: temperature in °C above preindustrial, m: greenhouse gas concentrations as a fraction of preindustrial levels. $\hat{t}$: equal to $\exp(0.02t)$, where t is years after 2005. μ: mean of the prior distribution of the feedback parameter (unit free). $1/\eta$: variance of the prior distribution of the feedback parameter (unit free).

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 (3.18)-(3.25). 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 n_p years.

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

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 information 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.

A.6.2 Maximum climate sensitivity

The solution algorithm requires a maximum value of the climate sensitivity. As discussed in Kelly and Tan (2015) appendix B.2, if climate sensitivity is unbounded, then the temperature can always exceed the largest temperature grid point. Interpolation of the value function outside the grid is likely to be inaccurate. We therefore set the largest value of the climate sensitivity equal to 15°C. Values greater than 15°C are collected into a mass point at 15°C.[31]

The results, however, are not sensitive to the maximum allowable climate sensitivity. First, the simulations aggregate probability mass greater than the maximum climate sensitivity into a mass point at the maximum climate sensitivity. Therefore, reducing the maximum $\Delta T_{2\times}$ essentially makes a small range of low probability events less harmful.

Further, unconstrained abatement is increasing in the climate sensitivity. Therefore, beyond a certain critical climate sensitivity, the 2°C target is almost certainly exceeded, so constrained emissions quickly fall to zero and remain at zero for a relatively long period of time. But unconstrained emissions also drop to near zero. The peak temperature increases for both the constrained and unconstrained models. Therefore, beyond a certain climate sensitivity, the only welfare difference between high and very high values of the climate sensitivity is in the far future, where the constrained model still has zero emissions in an

[31] The literature provides some expert opinion the upper bound of the climate sensitivity. Weitzman (2009) aggregates the densities of 22 studies and estimates the probability that $\Delta T_{2\times} > 10$°C is approximately 1%. However, he also suggests that because of feedback uncertainty, the probability that $\Delta T_{2\times} > 20$ is 1%. Roe and Baker (2007) plot the densities of several studies, none of which have any mass above 10°C. Similarly, Intergovernmental Panel on Climate Change (2007) normalize a number of papers so that the probability that $\Delta T_{2\times} > 10^\circ$ is zero. Our calibration matches the fitted distribution of Roe and Baker (2007), for which the probability of $\Delta T_{2\times} > 15$ is 1.8%. Therefore, a mass point exists at $\Delta T_{2\times} = 15$ of 1.8%.

attempt to return to 2°, while the unconstrained model remains at an optimal temperature which is increasing in the climate sensitivity, with more emissions.

For example, if all weather shocks are zero and emissions are zero, the target is exceeded for any $\Delta T_{2\times} \geq 5.36$. Therefore, the distribution can be truncated anywhere between 8-15°C with little change in the near term emissions path. Differences in the far future exist, but the low probability of such events and discounting combine to limit the effect on welfare.

Table 5 shows how the welfare loss changes with the maximum value of $\Delta T_{2\times}$. The maximum difference in the table is about 12%. Note that welfare losses are lower as a percentage of the unconstrained welfare when the maximum $\Delta T_{2\times}$ is higher. This occurs because both constrained and unconstrained welfare fall as the maximum $\Delta T_{2\times}$ increases, but the unconstrained welfare falls more. Total welfare (not as a percentage of unconstrained) is monotonically decreasing as a function of the maximum $\Delta T_{2\times}$ for all cases.

		Maximum $\Delta T_{2\times}$ (°C)		
		8	10	15
	0.20	5.54	5.35	5.20
	0.25	5.32	5.14	4.99
	0.30	5.16	4.97	4.82
	0.50	4.67	4.49	4.33
ω	0.75	4.04	3.85	3.69
	0.85	3.73	3.55	3.39
	0.90	3.58	3.39	3.23
	0.95	3.37	3.19	3.02
	0.98	3.19	3.01	2.85
	1	0	0	0
Minimum Feasible ω		0.152	0.152	0.152

Table 5: Welfare loss as a function of the maximum $\Delta T_{2\times}$, mean of 10,000 simulations, each with a draw of the feedback parameter from the prior distribution and a weather shock draw for each period in the simulation. The baseline maximum is 15°C above preindustrial (column 3). The last row gives the probability that the temperature exceeds the target at some time in the future with zero emissions (the peak in Figure 1).

A.6.3 Non-differentiability of the value function

The presence of the min operator in the constraint implies the value function is continuous, but not differentiable (kinked), along a particular hyperplane. Therefore, error will exist at each kink point as the spline has continuous first and second derivatives. Here we examine the error in the kink point.

The first step is to calculate the kink hyperplane. Below is a heuristic argument, for a more formal treatment, see Rincón-Zapatero and Santos (2009). Problem (4.3) has potentially three regions: (1) the constraint does not bind ($\theta = 0$), (2) the constraint binds and $PC(s) < 1$, and (3) the constraint binds with $PC(s) \geq 1$.

If the value function is differentiable at s_0, the envelope theorem implies the derivative is:

$$V_s(s) = u_c\left[c\left(x^*\left(s_0\right), s_0\right)\right] c_s\left(x^*\left(s_0\right), s_0\right) - \theta\left(s_0\right) \frac{\partial}{\partial s} \min\left\{PC\left(s_0\right), 1\right\} + \beta(t)\, \mathrm{E}\left[V_s\left[s'\left(s_0, x^*\left(s_0\right)\right)\right] s'_s\left(s_0, x^*\left(s_0\right)\right)\right]. \tag{A.44}$$

In the non-binding region (1), $\theta = 0$ and so the value function is differentiable via standard arguments. In region (2),

$$\theta\left(s_0\right) \frac{\partial}{\partial s} \min\left\{PC\left(s_0\right), 1\right\} = \theta\left(s_0\right) PC_s\left(s_0\right). \tag{A.45}$$

Therefore, one can again apply standard arguments since the constraint is continuously differentiable in this region. At any s_0 such that the constraint just binds, $\theta = 0$ and $x^*(s_0) = PC(s_0)$. Further, V is differentiable at s_0 if the left and right derivatives are equal. Since from the binding region $\theta \to 0$ as $s \to s_0$ and from the non-binding region $\theta = 0$, the left and right derivatives are equal.

Next, in the binding region where $PC(s_0) > 1$:

$$\theta\left(s_0\right) \frac{\partial}{\partial s} \min\left\{PC\left(s_0\right), 1\right\} = 0. \tag{A.46}$$

Therefore, the value function is differentiable in region (3). Finally, consider any point s_0 such that $PC(s_0) = 1$ and the constraint binds. In this case, the left and right derivatives differ. Let V_s^L denote the derivative in the region where $PC < 1$, then:

$$V_s^L(s) = u_c\left[c\left(x^*(s), s\right)\right] c_s\left(x^*(s), s\right) - \theta(s)\, PC_s(s) + \beta(t)\, \mathrm{E}\left[V_s\left[s'\left(s, x^*(s)\right)\right] s'_s\left(s, x^*(s)\right)\right]. \tag{A.47}$$

The limit as $s \to s_0$ from the left is thus:

$$V_s^L(s_0) = u_c\left(c\left(1, s_0\right)\right) c_s\left(1, s_0\right) - \theta\left(s_0\right) PC_s\left(s_0\right) + \beta(t)\, \mathrm{E}\left[V_s\left[s'\left(s_0, 1\right)\right] s'_s\left(s_0\right)\right]. \tag{A.48}$$

Let V_s^R denote the derivative where $PC > 1$. As $s \to s_0$ from the right:

$$V_s^R(s_0) = u_c\left(c\left(1, s_0\right)\right) c_s\left(1, s_0\right) + \beta\left(t\right) \mathrm{E}\left[V_s\left[s'\left(s_0, 1\right)\right] s'_s\left(s_0\right)\right]. \tag{A.49}$$

Hence any kink point s_0 in the value function is on a hyperplane which satisfies $PC\left(s_0\right) = 1$. Further, the left derivative is smaller than the right derivative. For states such that the derivative of V is negative (increases in T, m, and μ reduces lifetime utility), then the left slope is steeper than the right slope.

Figure 9 graphs a sample kink point.[32] The value function switches from concave to almost linear at the kink point, and the right slope is flatter than the left slope at the kink point, as expected. The green line is the (differentiable) spline approximation, which smooths the kink point. The magenta line is composed of two spline approximations, one on either side of the kink point. Hence, the magenta line need not be differentiable at the kink point. From the graph, the error associated with assuming the value function is differentiable is small.[33]

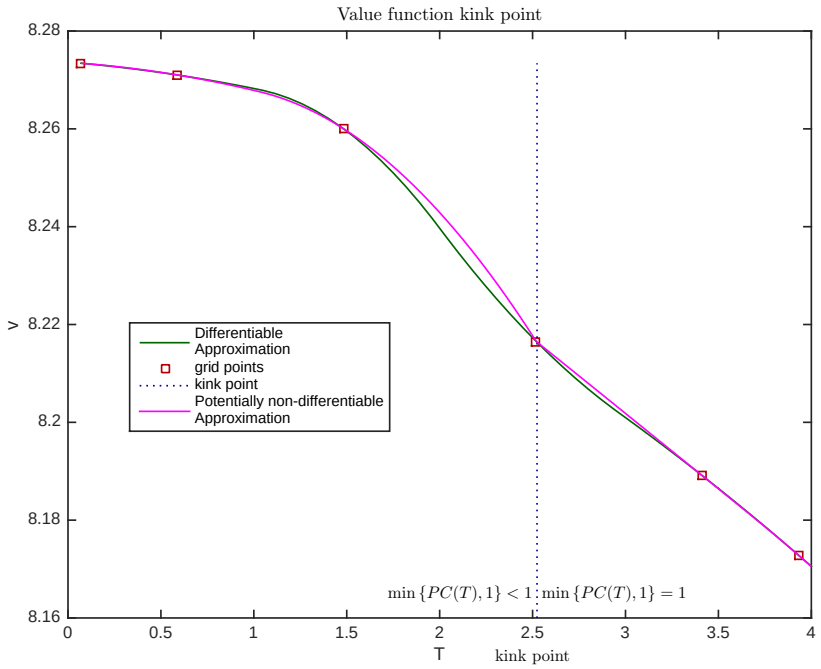

Figure 9: Sample kink point in the temperature dimension. At state $s_0 = [10.810, 2.524, 1.292, 0.146, 0.550, 0.0169]$, $PC\left(s_0\right) = 1$.

[32] A kink point was chosen which almost exactly equals a grid point. Between the grid points, it is not possible to separate interpolation error from error associated with using a continuous approximation of a function with a kink.

[33] We do not have a proof which bounds the approximation error. Nonetheless, a check of kink points in many areas of the state space revealed very small errors.

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, 2015, "Learning and Climate Feedbacks: Optimal Climate Insurance and Fat Tails," *Journal of Environmental Economics and Management*, 72, 98–122.

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.

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.

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.

Rincón-Zapatero, J. P., and M. S. Santos, 2009, "Differentiability of the Value Function Without Interiority Assumptions," *Journal of Economic Theory*, 144, 19481964.

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.

Rudik, I., 2016, "Optimal Climate Policy When Damages are Unknown," Iowa State University Working Paper.

Sanford, T., P. Frumhoff, A. Luers, and J. Gulledge, 2014, "The climate policy narrative for a dangerously warming world," *Nature Climate Change*, 4, 164–166.

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.

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.