Online Appendix: Optimal Climate Policy when Damages are Unknown - Ivan Rudik

The online appendix gives full details on the dynamic stochastic climate-economy model, describes how I use a linear climate model consistent with recent developments in climate science, describes my computational algorithms, provides an error analysis of the results, shows the full derivation of the carbon tax, and plots the certainty tax.

A The Full Stochastic Climate-Economy Model

The model is a Ramsey-Cass-Koopmans growth model coupled to a climate system. The model is governed by a representative policymaker whose objective is to maximize her expected discounted welfare. Each period lasts one year and the model begins in 2005.

In each period t, the policymaker has an endowment of capital K_t, labor L_t, and technology A_t. To improve numerical accuracy I express capital in effective labor terms, $k_t = \frac{K_t}{A_t L_t}$ (Traeger, 2014). Capital, labor and technology are combined in a Cobb-Douglas production function to produce gross output

$$Y_t^g = k_t^{\kappa_t}.$$

Warming of the Earth's surface causes damage to output, resulting in net output after damages

$$Y_t^n = Y_t^g \, \mathcal{L}_t,$$

where $\mathcal{L}_t$ is the fraction of output remaining after damages as defined in the main text (fractional net output)

$$\mathcal{L}_t = \frac{1}{1 + d_1 \, [T_t^s]^{d_2} \, w_t}.$$

The policymaker has three ways to use her remaining output after damages. First, she can use it for consumption C_t, also expressed in effective labor terms $c_t = \frac{C_t}{A_t L_t}$ to increase flow utility

$$U(c_t) = \frac{c_t^{1-\eta}}{(1-\eta)}, \quad \eta \neq 1.$$

Second, she can use it to abate some fraction $\alpha_t \in [0,1]$ of emissions from factor production with a cost given by

$$C(\alpha_t) = \Psi_t \alpha_t^{a_2}.$$

The residual output is left for investment into increasing the future capital stock, which depreciates at an annual rate of δ_k. Net emissions after abatement e_t is

$$e_t = \sigma_t(1-\alpha_t)Y_t^g + B_t.$$

B_t is emissions from exogenous land use change, and σ_t is the emissions intensity of output. Following the recent climate science literature (e.g. Matthews et al., 2009; IPCC, 2014; Knutti et al., 2017), I use a 1 state climate system that is a function of cumulative CO_2 emissions. Contemporaneous surface temperature (relative to preindustrial) is proportional to cumulative CO_2 emissions since preindustrial times

$$T_t^s = \zeta \sum_{s=1880}^{t} e_s,$$

where ζ is the transient climate response to emissions (TCRE). Cumulative emissions has a Markov transition

$$E_t = E_{t-1} + e_t$$

so that I can write temperature as

$$T_t^s = \zeta E_t.$$

The model's exogenously evolving processes are

$$L_t = L_0 + (L_\infty - L_0)\left(1 - \exp(-\delta_L t)\right) \qquad \text{(Labor population)}$$

$$g_{L,t} = \delta_L \left(\frac{L_\infty}{L_\infty - L_0} \exp(\delta_L t) - 1 \right)^{-1} \qquad \text{(Labor growth rate)}$$

$$A_t = A_0 \exp\left((1 - \exp(-\delta_A t)) \frac{g_{A,0}}{\delta_A} \right) \qquad \text{(Production technology)}$$

$$g_{A,t} = g_{A,0} \exp(-\delta_A \, t) \qquad \text{(Production technology growth rate)}$$

$$\beta_t = \exp\left(-\rho + (1-\eta) g_{A,t} + g_{L,t}\right) \qquad \text{(Growth adjusted discount factor)}$$

$$\sigma_t = \sigma_0 \exp\left[\frac{g_{\sigma,0}}{\delta_\sigma} \left(1 - \exp(-\delta_\sigma t)\right)\right] \qquad \text{(Gross emissions per unit of output)}$$

$$\Psi_t = \frac{a_0 \, \sigma_t}{a_2} \left(1 - \frac{1 - \exp(g_\Psi t)}{a_1}\right) \qquad \text{(Abatement cost coefficient)}$$

$$B_t = B_0 \, g_B^t \qquad \text{(Non-industrial CO}_2\text{ emissions)}$$

Table A1 reports the values of the model parameters. The calibration of the distribution over d_1, d_2, and ω_{t+1} are described in the main text.

Table A1: The parameters of the model.

Parameter	Value	Description
A_0	0.027	Initial production technology
$g_{A,0}$	0.009	Initial growth rate of production technology
δ_A	0.001	Change in growth rate of production technology
L_0	6514	Year 2005 population (millions)
L_∞	8600	Asymptotic population (millions)
δ_L	0.035	Rate of approach to asymptotic population level
σ_0	0.13	Initial emission intensity of output (Gigatons of carbon per unit output
$g_{\sigma,0}$	-0.0073	Initial growth rate of decarbonization
δ_σ	0.003	Change in growth rate of emissions intensity
a_0	1.17	Cost of backstop technology in 2005 ($1000 per ton of carbon)
a_1	2	Initial backstop technology cost / Final backstop technology cost
a_2	2.8	Abatement cost function exponent
g_Ψ	-0.005	Growth rate of backstop technology cost
B_0	1.1	Initial non-industrial CO_2 emissions (Gigatons of carbon)
g_B	-0.01	Growth rate of non-industrial emissions
κ	0.3	Capital elasticity in production
δ_k	0.1	Capital depreciation rate
E_0	454.90	Cumulative emissions in initial model year (Gigatons of carbon in 2005
ρ	0.015	Pure rate of time preference
η	2	1/EIS, and RRA
ζ	0.0016	TCRE (°C per 1,000 gigatons of carbon)
k_0	$137/(A_0 L_0)$	Year 2005 effective capital
μ_c	-5.38	d_1 location parameter
σ_c^2	0.38	d_1 scale parameter
$\exp\left(\mu_c + \frac{\sigma_c^2}{2}\right)$	0.00556	d_1 mean
$\exp\left(2\mu_c + \sigma_c^2\right)\left[\exp\left(\sigma_c^2\right) - 1\right]$	0.00379^2	d_1 variance
μ_0	1.88	Year 2005 d_2 mean
Σ_0	0.20	Year 2005 d_2 variance
μ_ω	-0.59	ω_t location parameter
σ_ω^2	1.18	ω_t scale parameter
$\exp\left(\mu_\omega + \frac{\sigma_\omega^2}{2}\right)$	1	ω_t mean
$\exp\left(2\mu_\omega + \sigma_\omega^2\right)\left[\exp\left(\sigma_\omega^2\right) - 1\right]$	1.22^2	ω_t variance

Without loss of generality, the robust control policymaker's problem is then

$$V_t(k_t, E_t, \mathcal{L}_t, \mu_t, \Sigma_t) =$$
$$\max_{c_t, \alpha_t} \left\{ u(c_t) + \mathbb{E}_{d_2} \left[-\theta log \left(\mathbb{E}_{d_1, \omega_{t+1}} \left[exp \left(-\frac{\beta_t V_{t+1}(k_{t+1}, E_{t+1}, \mathcal{L}_{t+1}, \mu_{t+1}, \Sigma_{t+1})}{\theta} \right) \right] \right) \right] \right\}$$

subject to transitions:

$$\begin{aligned}
k_{t+1} &= \exp\left(-(g_{L,t} + g_{A,t})\right) \left[(1-\delta_k)k_t - c_t + (1 - \Psi_t \alpha_t^{a_2}) Y_t^g \mathcal{L}_t\right], \\
E_{t+1} &= E_t + \sigma_t (1 - \alpha_t) Y_t^g + B_t, \\
\mathcal{L}_{t+1} &= \frac{1}{1 + d_1 \left[\zeta E_{t+1}\right]^{d_2} \omega_{t+1}}, \\
\mu_{t+1} &= \frac{(\sigma_w^2 + \sigma_c^2)\,\mu_t + \log(\zeta E_{t+1})\,\Sigma_t \left[\log\left(\frac{1}{\mathcal{L}_{t+1}} - 1\right) - (\mu_c + \mu_w)\right]}{(\sigma_w^2 + \sigma_c^2) + \left[\log(\zeta E_{t+1})\right]^2 \Sigma_t}, \\
\Sigma_{t+1} &= \frac{\Sigma_t\,(\sigma_w^2 + \sigma_c^2)}{(\sigma_w^2 + \sigma_c^2) + \left[\log(\zeta E_{t+1})\right]^2 \Sigma_t}.
\end{aligned}$$

Finally I constrain abatement to be less than 100 percent and I impose the resource constraint

$$\begin{aligned}
&\alpha_t \leq 1, \\
&c_t + \Psi_t \alpha_t^{a_2} Y_t^g \mathcal{L}_t \leq Y_t^g \mathcal{L}_t.
\end{aligned}$$

If the policymaker does not learn, then $\mu_{t+1} = \mu_t$ and $\Sigma_{t+1} = \Sigma_t$. Note that the transition for μ_{t+1} is stochastic and depends on the realization of ω_{t+1} and the true values of d_1 and d_2.

A.1 Climate System

To make the model tractable to solve numerically, I take advantage of recent findings in climate science that find warming is proportional to cumulative emissions of carbon (Matthews et al., 2009; IPCC, 2014; Knutti et al., 2017). This relationship comes about because a pulse of carbon into the atmosphere increases equilibrium temperature quickly by a constant amount (Matthews and Caldeira, 2008). This has lead scientists to conclude that contemporaneous temperature is almost entirely determined by the level of cumulative past emissions (Allen et al., 2009; Matthews et al., 2009). Past cumulative emissions is translated into temperature by a factor of proportionality called the *transient climate response to emissions* (TCRE).

Define the TCRE as (Williams et al., 2016; Dietz and Venmans, 2019)

$$\zeta \equiv \frac{\Delta T}{\Delta E} = \frac{\Delta T}{\Delta M}\frac{\Delta M}{\Delta E},$$

where T is temperature, M is atmospheric carbon, and E is cumulative emissions. Let Δ denote changes since preindustrial levels. $\frac{\Delta T}{\Delta M}$ gives the change in warming from additional carbon in the atmosphere and $\frac{\Delta M}{\Delta E}$ gives the change in atmospheric carbon concentrations given a change in emissions. Each of these terms is individually non-linear but the product of this ends up being a simple proportional relationship for two key reasons: (1) warming from a pulse of emissions realizes quickly, and then is constant over time and (2) the marginal effect of emissions on temperature is constant and independent of the state of the climate.

How does this come about?[1] First, $\frac{\Delta T}{\Delta M}$ is increasing and concave. Over time, greater CO_2 concentrations lead to more warming as additional heat gets trapped. The relationship is concave because both the relationship between CO_2 concentrations and temperature is approximately logarithmic, and the oceans' significant heat capacity leads to thermal inertia: a lag between changes in carbon concentrations and equilibrium warming. The second term, $\frac{\Delta M}{\Delta E}$, is increasing and convex. Additional emissions initially go into the atmosphere and increase atmospheric CO_2, while some of these emissions eventually get absorbed in land and ocean carbon sinks. The convex relationship arises because the strength of these sinks is decreasing in cumulative emissions. If there has been greater past cumulative emissions, a larger fraction of the next emission will remain in the atmosphere instead of being absorbed by one of the sinks. The curvature in $\frac{\Delta T}{\Delta M}$ and $\frac{\Delta M}{\Delta E}$ effectively cancel out so that ζ is constant and positive.

Using the TCRE climate model allows me to reduce the climate system to 1 state as opposed to other options such as the 5 state DICE climate system. This makes using standard tensor product approximation methods tractable and updates the climate system to match the most recent developments in climate science. I use a TCRE value of $\zeta = 0.0016$ (°C/1,000 gigatons of carbon). This value and the proportional translation of cumulative emissions to temperature follow closely to parallel work on climate model uncertainty (Berger and Marinacci, 2017) and recent

[1] See Figure 1 of Millar et al. (2016) for plots of these relationships.

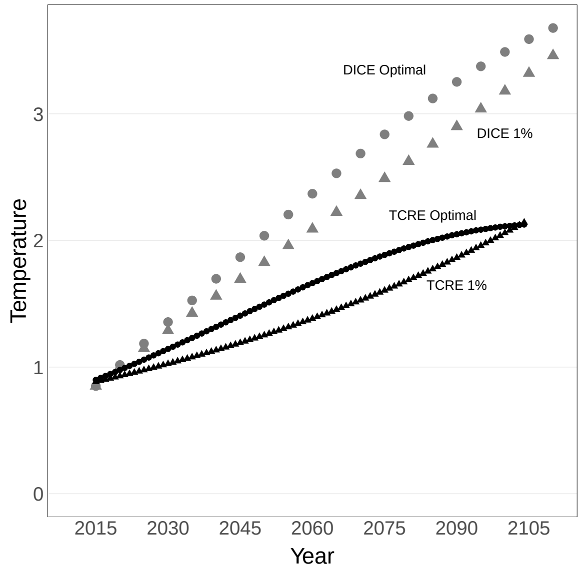

Figure A1: The temperature trajectories for DICE-2016 along the DICE-optimal emissions path (gray circles), DICE-2016 along an emissions path that grows cumulative emissions by 1% per year (gray triangles), the TCRE-based model used in this paper along the DICE-optimal emissions path (black circles) and the TCRE-based model used in the paper along an emissions path that grows cumulative emissions by 1% per year (black triangles).

climate economics papers (e.g. Dietz and Venmans, 2019; van der Ploeg, 2018). This value for the TCRE is in the middle of the most likely range estimated by the IPCC (IPCC, 2014; Knutti et al., 2017). ζ has been found to be constant up to various levels of cumulative emissions, including up to 2,000 GtC (Matthews et al., 2009), or 3,000 GtC (Leduc et al., 2015), and the TCRE has also been found to be approximately constant across all RCP scenarios (MacDougall and Friedlingstein, 2015).

Figure A1 plots temperature trajectories comparing the DICE climate system to the TCRE system. The DICE temperature trajectories are shown in gray and the TCRE trajectories are shown in black. Circle markers denote emissions trajectories that correspond to the optimal emissions trajectory from DICE-2016 while triangle markers correspond to an emissions trajectory that grows cumulative emissions by 1% per year. Both trajectories result in cumulative emissions of 1,300-1,400 GtC by 2105. The TCRE system predicts approximately 2.1°C of warming along these two trajectories which is well within the range of predicted temperatures of global climate models at this cumulative emissions level (IPCC, 2013). The DICE model predicts about 1°C more warming than the TCRE model and tends to be on the high end of temperature predictions from the suite of global climate models used in IPCC (2013).

A.2 Model Solution Method and Error Analysis

General approximation scheme The model is solved using value function iteration on a finite horizon. The collocation grid and polynomial interpolant are built using standard tensor product methods. The grid is constructed from a tensor product of one dimensional vectors of zeros of Chebyshev polynomials. For the learning frameworks, the grid has 2,100 grid points for 2005–2505 and 1,575 grid points for 2506–2605.[2] For the non-learning frameworks, the grid has 729 grid points.[3] The polynomial interpolant is constructed from a tensor product of Chebyshev polynomials. Below I test the sensitivity of my results to the number of collocation grid points, quadrature points, and the bounds for the state space.

Terminal value function The terminal year is 2605. The terminal continuation value function corresponding to 2606 has the policymaker not learning while holding her initial beliefs about $\{d_2, d_1, \omega_{t+1}\}$, and has all exogenous processes held constant at their 2606 levels. Changing the terminal value function to one where the policymaker does not expect damage stochasticity and believes the damage function to be exactly equal to that in the conventional DICE model does not significantly alter the results since the continuation value 600 years in the future is effectively discounted to zero.

Expectations Expectations over future states are taken using Gauss-Hermite quadrature with 11 unique points over the d_2 distribution, and 11 unique points over the joint $\{d_1, \omega_{t+1}\}$ distribution for a total of 121 quadrature points. Note that the support of the location parameter for a lognormal distribution is the entire real line. Depending on the draws of the random variables, it can take on any value in $(-\infty, \infty)$. Below I test the sensitivity of my results to the number of quadrature points to determine whether evaluating the random variables at additional points further into the tail of the distribution significantly affects my results.

[2] These correspond to 5 unique grid points for cumulative emissions, 3 for the location parameter, 7 for the scale parameter, 5 for fractional net output, and either 3 or 4 for effective capital.

[3] This corresponds to 9 on each state.

Table A2: Upper and lower bounds for each state. In the case of learning and RC+L frameworks the bounds are only for the first iteration of the adaptive grid algorithm.

	k	E	μ	Σ	$\mathcal{L}$
State Upper Bound: Uncertainty and RC	4.30	3500.0	–	–	1.0
State Lower Bound: Uncertainty and RC	0.65	454.9	–	–	0.3
Initial State Upper Bound: Learning and RC+L	6.0	5000.0	3.13	$.45^2$	1.0
Initial State Lower Bound: Learning and RC+L	0.0	454.9	0.63	0.5	0.4

Adaptive grid I use a time-varying set of bounds to approximate the learning and RC+L value functions.[4] Earlier years have tighter bounds since the set of possible realized states is smaller, while later years have wider bounds.[5] This allows me to use significantly fewer collocation points than on a standard hyperrectangle grid. I generate this adaptive grid using the following algorithm:

1. Solve the model on the time-invariant set of bounds given by Table A2. Use 3 unique grid points for effective capital, 5 for cumulative emissions, 3 for the location parameter, 7 for the scale parameter, and 5 for fractional net output.
2. Perform 1,000 simulations that are 301 years in length.
3. For each state
   (a) Recover the maximum and minimum simulated values in each year over the 1,000 simulations.
   (b) These time series will be noisy in later years even with a very large number of simulations, so fit a spline to the time series of maxima and another spline to the time series of minima.[6]
   (c) Set the new time t upper/lower bounds to a be weighted average of the fitted maxima/minima splines evaluated at time t, and the previous time

[4]This scheme is similar to that used in a number of recent stochastic implementations of DICE with even more states (Cai et al., 2015; Lontzek et al., 2015; Cai et al., 2016, Forthcoming). Cai (2018) provides more details on why this approach is particularly efficient for climate-economy models related to DICE and how it can allow researchers to solve extremely complicated models in minutes.

[5]For example, cumulative emissions can only grow so much in the first five years and the trajectory is non-decreasing since there are no negative emissions technologies.

[6]Smoothing out a time series of maxima or minima takes many more simulation runs than smoothing out the mean. The process can be greatly accelerated by fitting a spline.

t bounds, where weights on the bounds given by the splines are shown in Table A3.[7]

4. Solve the model on the new time-varying set of bounds obtained in Step 3. Use 3 unique grid points for effective capital, 5 for cumulative emissions, 3 for the location parameter, 7 for the scale parameter, and 5 for fractional net output.

5. Repeat steps 2–4 16 times.

6. Repeat steps 2–4 4 times but use 5 grid points for effective capital starting at year 2505.

7. Set the bounds in 2005 to effectively be an epsilon ball around the initial state since it is the only state reached in that year.

8. Solve the model on this final set of adapted bounds using 4 unique grid points for effective capital, 5 for cumulative emissions, 3 for the location parameter, 7 for the scale parameter, 5 for fractional net output.

I do not adapt the bounds for the scale parameter (i.e. the 0.0 weights in Table A3) because the misspecified damage function simulations tend to result in it jumping outside the bounds if its bounds are adapted, while the other states stay within the bounds even when the damage function is misspecified. I do not adapt the upper bound for fractional net output because the maximum simulated values are generally very close to 1.0, the initial upper bound. The scale parameter bounds are not adapted except for the initial period. The algorithm allows for tight bounds on the remaining states compared to the non-learning frameworks' hyperrectangle grid.

Figure A2 depicts the bounds and 1,000 simulated state trajectories. The final set of adapted grid bounds are the thick black lines, the simulated state trajectories are the thin gray lines, and the expected state trajectories are the dotted black lines. For the states where the bounds were adapted, the bounds are tight around the set of states actually reached in simulations. For about 20 total simulation years each, fractional net output and the location parameter are no longer within the state bounds where the value function approximant is reliable. Although this introduces error, 20 years outside the bounds is a negligible fraction of the 101,000 total years displayed in

[7] This is similar to damping procedures used in other iterative algorithms like infinite-horizon value function iteration.

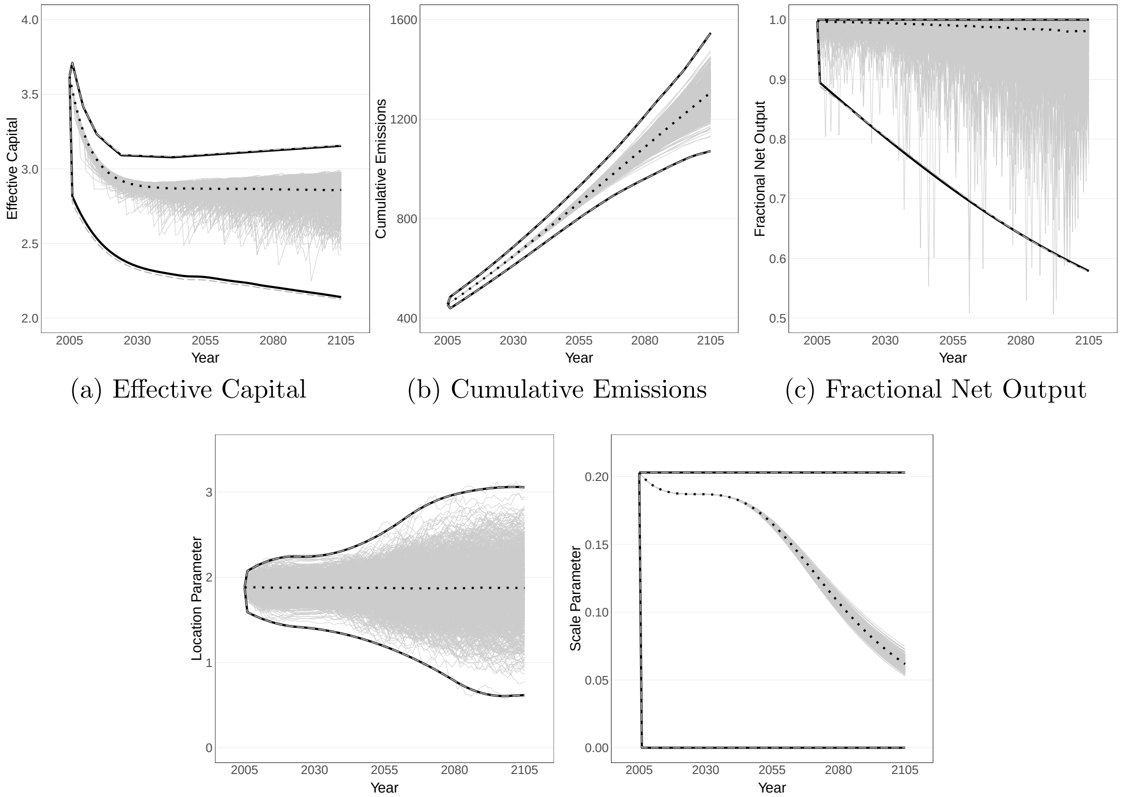

(a) Effective Capital (b) Cumulative Emissions (c) Fractional Net Output

(d) Location Parameter (e) Scale Parameter

Figure A2: The bounds for each state for the learning framework and the robust control and learning framework (thick black line). The bounds for each state for the learning framework and the robust control and learning framework in the second-to-last iteration of the adaptive grid algorithm (thick dashed gray line). Simulated state trajectories from the first 1,000 simulations (thin gray lines). Expected state trajectory (dotted black line).

the simulations (101 year simulation $\times$ 1,000 simulations). The figure also shows the bounds from the second-to-last iteration of the adaptive grid algorithm as a thick gray dashed line. These are almost identical to the actual state bounds in black, except for the capital lower bound. The average relative difference in the capital lower bounds between the final two algorithm iterations is less than 1%.

The non-learning frameworks do not use the adaptive grid. This is for two reasons. First, the non-learning frameworks have low dimensionality and can be solved on a standard hyperrectangle grid. Second, without learning, the states often drift outside adapted approximation bounds when facing a misspecified damage function. If the policymaker's beliefs about the damage function are correct, the adaptive grid can

Table A3: Weights on fitted spline bounds obtained from Step 3c in the adaptive grid algorithm. Weights on the previous iteration's bounds are 1 minus the weights in the table.

	k	E	μ	Σ	$\mathcal{L}$
Upper Bound Weight	0.3	0.4	0.5	0.0	0.0
Lower Bound Weight	0.1	0.3	0.5	0.0	0.5

easily be used for a non-learning policymaker. The state bounds for the non-learning frameworks are contained in Table A2.

Error analysis I investigate the accuracy of the model by calculating the relative differences in simulated policy trajectories when (i) increasing the number of unique collocation points by 1 by over each dimension (e.g. capital goes from $4 \rightarrow 5$, cumulative emissions goes from $5 \rightarrow 6$, etc), (ii) increasing the number of quadrature points by 2 over the d_2 distribution and the joint $\{d_1, \omega_{t+1}\}$ distribution, (iii) increasing the collocation upper bound by 10% and decreasing the collocation lower bound by 10% for the non-adapted grids,[8] and (iv) only doing the first 10 adaptive grid iterations instead of the full 20. For the frameworks with learning, (i)-(iii) use the same adaptive grid bounds as the baseline results. Tables A4 displays the maximum and average relative differences of the simulated trajectories' carbon taxes and consumption levels over the first 100 years.

Average differences in taxes and consumption for all frameworks are on the order of 10^{-3} or smaller, while the maximum error is larger but still relatively small.[9]

B Carbon Tax Derivation

The conventional definition of the time t carbon tax used in the dynamic stochastic integrated assessment literature is that it is the shadow cost of time t emissions. It is converted into dollar terms using time t marginal utility evaluated at the optimal consumption level. Without loss of generality I derive the carbon tax using the value

[8]For fractional net output I reduce the lower bound by 20% since increasing the upper bound beyond 1.0 corresponds to generating more than 100% of possible output without climate change.

[9]Changing the number of collocation points on capital by 2 instead of 1 generates similarly sized errors of about 10^{-3} for averages and 10^{-2} for maximums.

Table A4: Average and maximum differences in carbon taxes and consumption from increasing the number of collocation grid points or quadrature points, expanding the size of the approximation domain for the non-learning frameworks, and performing 10 fewer iterations of the adaptive grid algorithm for the frameworks with learning. The average and maximum carbon tax trajectory differences are taken from 50,000 simulations over 2005–2105 where each simulation has a random draw of $\{d_1, d_2, \omega_{2005}, ..., \omega_{2105}\}$.

Parameters Changed	Relative Differences	Uncertainty	Learning	RC	RC+L
# Collocation Points: Add 1 unique point to each state	Tax: Max	1.1%	1.4%	1.2%	0.7%
	Tax: Avg	0.2%	0.5%	0.2%	0.4%
	Consumption: Max	0.2%	0.08%	0.2%	0.06%
	Consumption: Avg	0.02%	9×10^{-3}%	0.02%	8×10^{-3}%
# Quadrature Points: Add 2 unique points to both distributions: $\{d_2\}$ and $\{d_1, \omega_{t+1}\}$	Tax: Max	0.06%	0.6%	0.04%	0.5%
	Tax: Avg	0.02%	0.3%	0.02%	0.2%
	Consumption: Max	0.03%	0.08%	0.03%	0.03%
	Consumption: Avg	7×10^{-3}%	7×10^{-3}%	8×10^{-3}%	8×10^{-3}%
Domain Bounds: Increase UB and decrease LB by 10%	Tax: Max	0.8%	–	0.9%	–
	Tax: Avg	0.3%	–	0.3%	–
	Consumption: Max	0.4%	–	0.2%	–
	Consumption: Avg	0.1%	–	0.09%	–
# Adaptive Iterations: Decrease from 20 to 10	Tax: Max	–	0.8%	–	0.8%
	Tax: Avg	–	0.3%	–	0.3%
	Consumption: Max	–	0.2%	–	0.2%
	Consumption: Avg	–	0.06%	–	0.06%

function for the robust control and learning framework. Recall the right hand side of the Bellman equation is

$$V_t(\boldsymbol{S_t}) = \max_{c_t,\alpha_t}\left\{U(c_t) + \mathbb{E}_{d_2}\left[-\theta\log\left(\mathbb{E}_{d_1,\omega_{t+1}}\left[\exp\left(\frac{-\beta_t\, V_{t+1}(\boldsymbol{S_{t+1}})}{\theta}\right)\right]\right)\right]\right\}$$

where $\mathbf{S_{t+1}} \equiv \{k_{t+1}, E_{t+1}, \boldsymbol{\mathcal{L}}_{t+1}, \mu_{t+1}, \Sigma_{t+1}\}$ and expectations are taken using the time t information set. To economize on space I omit the value function arguments for now. The shadow cost of emissions is the negative partial derivative of the right hand side of the Bellman with respect to time t emissions e_t, which yields

$$\theta\mathbb{E}_{d_2}\left\{\frac{\mathbb{E}_{d_1,\omega_{t+1}}\left[-\frac{\beta_t}{\theta}\exp\left(\frac{-\beta_t V_{t+1}}{\theta}\right)\frac{\partial V_{t+1}}{\partial \mathbf{S_{t+1}}}\frac{\partial \mathbf{S_{t+1}}}{\partial e_t}\right]}{\mathbb{E}_{d_1,\omega_{t+1}}\left[\exp\left(\frac{-\beta_t V_{t+1}}{\theta}\right)\right]}\right\}.$$

The θ terms cancel. Expand the expectation term over d_1, ω_{t+1} in the numerator of the fraction using the covariance identity to obtain

$$\beta_t\mathbb{E}_{d_2}\left\{\frac{\mathbb{E}_{d_1,\omega_{t+1}}\left[\exp\left(\frac{-\beta_t V_{t+1}}{\theta}\right)\right]\mathbb{E}_{d_1,\omega_{t+1}}\left[\frac{-\partial V_{t+1}}{\partial \mathbf{S_{t+1}}}\frac{\partial \mathbf{S_{t+1}}}{\partial e_t}\right]}{\mathbb{E}_{d_1,\omega_{t+1}}\left[\exp\left(\frac{-\beta_t V_{t+1}}{\theta}\right)\right]} + \frac{\mathrm{cov}_{d_1,\omega_{t+1}}\left(\exp\left(\frac{-\beta_t V_{t+1}}{\theta}\right), \frac{-\partial V_{t+1}}{\partial \mathbf{S_{t+1}}}\frac{\partial \mathbf{S_{t+1}}}{\partial e_t}\right)}{\mathbb{E}_{d_1,\omega_{t+1}}\left[\exp\left(\frac{-\beta_t V_{t+1}}{\theta}\right)\right]}\right\}.$$

We can cancel the expectations over the exponentials in the first term, and bring the denominator of the second term inside of the covariance operator such that it multiplies the first argument. The last step to obtain the optimal carbon tax is to divide by $u'(c_t)$ to translate into dollar terms

$$\frac{\beta_t}{u'(c_t)}\mathbb{E}_{d_2}\left\{\mathbb{E}_{d_1,\omega_{t+1}}\left[\frac{-\partial V_{t+1}}{\partial \mathbf{S_{t+1}}}\frac{\partial \mathbf{S_{t+1}}}{\partial e_t}\right] + \mathrm{cov}_{d_1,\omega_{t+1}}\left(\frac{\exp\left(\frac{-\beta_t V_{t+1}}{\theta}\right)}{\mathbb{E}_{d_1,\omega_{t+1}}\left[\exp\left(\frac{-\beta_t V_{t+1}}{\theta}\right)\right]}, \frac{-\partial V_{t+1}}{\partial \mathbf{S_{t+1}}}\frac{\partial \mathbf{S_{t+1}}}{\partial e_t}\right)\right\}. \tag{1}$$

The first term is the standard carbon tax expression without robust control. The second term is the additively separable adjustment for robust control.

Now focus on the first term (the standard carbon tax expression term) and ignore the leading $\frac{\beta_t}{u'(c_t)}$ for clarity. Expanding out the state vector, and recognizing that time $t+1$ capital is not a function of time t emissions, we have that

$$\begin{aligned}\mathbb{E}_{d_2,d_1,\omega_{t+1}}\left[\frac{-\partial V_{t+1}}{\partial \mathbf{S_{t+1}}}\frac{\partial \mathbf{S_{t+1}}}{\partial e_t}\right] =&\mathbb{E}_{d_2,d_1,\omega_{t+1}}\left[\frac{-\partial V_{t+1}}{\partial E_{t+1}}\frac{\partial E_{t+1}}{\partial e_t}\right]+\mathbb{E}_{d_2,d_1,\omega_{t+1}}\left[\frac{-\partial V_{t+1}}{\partial \mathcal{L}_{t+1}}\frac{\partial \mathcal{L}_t}{\partial e_t}\right]\\ &+\mathbb{E}_{d_2,d_1,\omega_{t+1}}\left[\frac{-\partial V_{t+1}}{\partial \mu_{t+1}}\frac{\partial \mu_{t+1}}{\partial e_t}\right]+\mathbb{E}_{d_2,d_1,\omega_{t+1}}\left[\frac{-\partial V_{t+1}}{\partial \Sigma_{t+1}}\frac{\partial \Sigma_{t+1}}{\partial e_t}\right].\end{aligned}$$

Next pass the expectations through. Recognize $\frac{\partial \mathcal{L}_{t+1}}{\partial e_t}$ and $\frac{\partial \mu_{t+1}}{\partial e_t}$ are random variables so we can apply the covariance identity to the expectations over those terms. And finally note that $\mathbb{E}_{d_2,d_1,\omega_{t+1}}\left[\frac{\partial \mu_{t+1}}{\partial e_t}\right]=0$ for a Bayesian, and $\frac{\partial E_{t+1}}{\partial e_t}=1$ by definition. Then we arrive at the final standard carbon tax expression before performing Taylor expansions

$$\begin{aligned}\mathbb{E}_{d_2,d_1,\omega_{t+1}}\left[\frac{-\partial V_{t+1}}{\partial \mathbf{S_{t+1}}}\frac{\partial \mathbf{S_{t+1}}}{\partial e_t}\right] =&\mathbb{E}_{d_2,d_1,\omega_{t+1}}\left[\frac{-\partial V_{t+1}}{\partial E_{t+1}}\right]\\ &+\mathbb{E}_{d_2,d_1,\omega_{t+1}}\left[\frac{-\partial V_{t+1}}{\partial \mathcal{L}_{t+1}}\right]\mathbb{E}_{d_2,d_1,\omega_{t+1}}\left[\frac{\partial \mathcal{L}_t}{\partial e_t}\right]+\mathrm{cov}_{d_2,d_1,\omega_{t+1}}\left(\frac{-\partial V_{t+1}}{\partial \mathcal{L}_{t+1}},\frac{\partial \mathcal{L}_{t+1}}{\partial e_t}\right)\\ &+\mathrm{cov}_{d_2,d_1,\omega_{t+1}}\left(\frac{-\partial V_{t+1}}{\partial \mu_{t+1}},\frac{\partial \mu_{t+1}}{\partial e_t}\right)+\mathbb{E}_{d_2,d_1,\omega_{t+1}}\left[\frac{-\partial V_{t+1}}{\partial \Sigma_{t+1}}\right]\frac{\partial \Sigma_{t+1}}{\partial e_t}.\end{aligned} \tag{2}$$

Recall that for a simple bivariate case, a second-order Taylor expansion of $E_{x,y}[f(x,y)]$ about (c,d) is

$$\begin{aligned}E_{x,y}[f(x,y)] \approx &E_{x,y}\Bigg[f(c,d)\\ &+f_x(x,y)(x-c)+f_y(c,d)(y-d)\\ &+\frac{1}{2}\big[f_{xy}(c,d)(x-c)^2+f_{xy}(c,d)(x-c)(y-d)\\ &+f_{yx}(c,d)(y-d)(x-c)+f_{yy}(c,d)(y-d)^2\big]\Bigg]\end{aligned}$$

where the term on the first line and on right hand side of the equality is the zeroth-order term, the second line contains the first-order terms, and the third and fourth lines contain the second-order terms. Now perform a second-order Taylor expansion of the expectations of the value function partial derivatives around

$$\upsilon := \left\{k_{t+1}, E_{t+1}, \mathbb{E}_{d_2,d_1,\omega_{t+1}}[\mathcal{L}_{t+1}], \mathbb{E}_{d_2,d_1,\omega_{t+1}}[\mu_{t+1}], \Sigma_{t+1}\right\}$$

where $\mathbb{E}_{d_2,d_1,\omega_{t+1}}[\mu_{t+1}] = \mu_t$.

Here I show the second-order Taylor expansion for $\mathbb{E}_{d_2,d_1,\omega_{t+1}}\left[\frac{-\partial V_{t+1}}{\partial E_{t+1}}\right]$ evaluated at υ. It consists of 1 zeroth-order Taylor term, five first-order Taylor terms, and twenty five second-order Taylor terms

$$
\begin{aligned}
\mathbb{E}_{d_2,d_1,\omega_{t+1}}\left[\frac{-\partial V_{t+1}}{\partial E_{t+1}}\right] \approx & \mathbb{E}_{d_2,d_1,\omega_{t+1}}\left[\frac{-\partial V_{t+1}}{\partial E_{t+1}}\right|_{\upsilon} && \text{(Zeroth-order)} \\
& + \frac{-\partial^2 V_{t+1}}{\partial E_{t+1}\partial k_{t+1}}\Big|_{\upsilon}(k_{t+1} - \mathbb{E}_{d_2,d_1,\omega_{t+1}}[k_{t+1}]) && \text{(First-order)} \\
& + \frac{-\partial^2 V_{t+1}}{\partial E_{t+1}^2}\Big|_{\upsilon}(E_{t+1} - \mathbb{E}_{d_2,d_1,\omega_{t+1}}[E_{t+1}]) \\
& + \frac{-\partial^2 V_{t+1}}{\partial E_{t+1}\partial \mathcal{L}_{t+1}}\Big|_{\upsilon}(\mathcal{L}_{t+1} - \mathbb{E}_{d_2,d_1,\omega_{t+1}}[\mathcal{L}_{t+1}]) \\
& + \frac{-\partial^2 V_{t+1}}{\partial E_{t+1}\partial \mu_{t+1}}\Big|_{\upsilon}(\mu_{t+1} - \mathbb{E}_{d_2,d_1,\omega_{t+1}}[\mu_{t+1}]) \\
& + \frac{-\partial^2 V_{t+1}}{\partial E_{t+1}\partial \Sigma_{t+1}}\Big|_{\upsilon}(\Sigma_{t+1} - \mathbb{E}_{d_2,d_1,\omega_{t+1}}[\Sigma_{t+1}]) \\
& + \frac{1}{2}\frac{-\partial^3 V_{t+1}}{\partial E_{t+1}\partial k_{t+1}^2}\Big|_{\upsilon}(k_{t+1} - \mathbb{E}_{d_2,d_1,\omega_{t+1}}[k_{t+1}])^2 && \text{(Second-order)} \\
& + \frac{1}{2}\frac{-\partial^3 V_{t+1}}{\partial E_{t+1}\partial k_{t+1}\partial E_{t+1}}\Big|_{\upsilon}(k_{t+1} - \mathbb{E}_{d_2,d_1,\omega_{t+1}}[k_{t+1}])(E_{t+1} - \mathbb{E}_{d_2,d_1,\omega_{t+1}}[E_{t+1}]) \\
& + \frac{1}{2}\frac{-\partial^3 V_{t+1}}{\partial E_{t+1}\partial k_{t+1}\partial \mathcal{L}_{t+1}}\Big|_{\upsilon}(k_{t+1} - \mathbb{E}_{d_2,d_1,\omega_{t+1}}[k_{t+1}])(\mathcal{L}_{t+1} - \mathbb{E}_{d_2,d_1,\omega_{t+1}}[\mathcal{L}_{t+1}]) \\
& + \frac{1}{2}\frac{-\partial^3 V_{t+1}}{\partial E_{t+1}\partial k_{t+1}\partial \mu_{t+1}}\Big|_{\upsilon}(k_{t+1} - \mathbb{E}_{d_2,d_1,\omega_{t+1}}[k_{t+1}])(\mu_{t+1} - \mathbb{E}_{d_2,d_1,\omega_{t+1}}[\mu_{t+1}]) \\
& + \frac{1}{2}\frac{-\partial^3 V_{t+1}}{\partial E_{t+1}\partial k_{t+1}\partial \Sigma_{t+1}}\Big|_{\upsilon}(k_{t+1} - \mathbb{E}_{d_2,d_1,\omega_{t+1}}[k_{t+1}])(\Sigma_{t+1} - \mathbb{E}_{d_2,d_1,\omega_{t+1}}[\Sigma_{t+1}]) \\
& + \frac{1}{2}\frac{-\partial^3 V_{t+1}}{\partial E_{t+1}^2\partial k_{t+1}}\Big|_{\upsilon}(E_{t+1} - \mathbb{E}_{d_2,d_1,\omega_{t+1}}[E_{t+1}])(k_{t+1} - \mathbb{E}_{d_2,d_1,\omega_{t+1}}[k_{t+1}]) \\
& + \frac{1}{2}\frac{-\partial^3 V_{t+1}}{\partial E_{t+1}^3}\Big|_{\upsilon}(E_{t+1} - \mathbb{E}_{d_2,d_1,\omega_{t+1}}[E_{t+1}])^2 \\
& + \frac{1}{2}\frac{-\partial^3 V_{t+1}}{\partial E_{t+1}^2\partial \mathcal{L}_{t+1}}\Big|_{\upsilon}(E_{t+1} - \mathbb{E}_{d_2,d_1,\omega_{t+1}}[E_{t+1}])(\mathcal{L}_{t+1} - \mathbb{E}_{d_2,d_1,\omega_{t+1}}[\mathcal{L}_{t+1}]) \\
& + \frac{1}{2}\frac{-\partial^3 V_{t+1}}{\partial E_{t+1}^2\partial \mu_{t+1}}\Big|_{\upsilon}(E_{t+1} - \mathbb{E}_{d_2,d_1,\omega_{t+1}}[E_{t+1}])(\mu_{t+1} - \mathbb{E}_{d_2,d_1,\omega_{t+1}}[\mu_{t+1}])
\end{aligned}
$$

$$
\begin{aligned}
&+\frac{1}{2}\left.\frac{-\partial^3 V_{t+1}}{\partial E_{t+1}^2 \partial \Sigma_{t+1}}\right|_{\upsilon}(E_{t+1}-\mathbb{E}_{d_2,d_1,\omega_{t+1}}[E_{t+1}])(\Sigma_{t+1}-\mathbb{E}_{d_2,d_1,\omega_{t+1}}[\Sigma_{t+1}]) \\
&+\frac{1}{2}\left.\frac{-\partial^3 V_{t+1}}{\partial E_{t+1}\partial \mathcal{L}_{t+1}\partial k_{t+1}}\right|_{\upsilon}(\mathcal{L}_{t+1}-\mathbb{E}_{d_2,d_1,\omega_{t+1}}[\mathcal{L}_{t+1}])(k_{t+1}-\mathbb{E}_{d_2,d_1,\omega_{t+1}}[k_{t+1}]) \\
&+\frac{1}{2}\left.\frac{-\partial^3 V_{t+1}}{\partial E_{t+1}\partial \mathcal{L}_{t+1}\partial E_{t+1}}\right|_{\upsilon}(\mathcal{L}_{t+1}-\mathbb{E}_{d_2,d_1,\omega_{t+1}}[\mathcal{L}_{t+1}])(E_{t+1}-\mathbb{E}_{d_2,d_1,\omega_{t+1}}[E_{t+1}]) \\
&+\frac{1}{2}\left.\frac{-\partial^3 V_{t+1}}{\partial E_{t+1}\partial \mathcal{L}_{t+1}^2}\right|_{\upsilon}(\mathcal{L}_{t+1}-\mathbb{E}_{d_2,d_1,\omega_{t+1}}[\mathcal{L}_{t+1}])^2 \\
&+\frac{1}{2}\left.\frac{-\partial^3 V_{t+1}}{\partial E_{t+1}\partial \mathcal{L}_{t+1}\partial \mu_{t+1}}\right|_{\upsilon}(\mathcal{L}_{t+1}-\mathbb{E}_{d_2,d_1,\omega_{t+1}}[\mathcal{L}_{t+1}])(\mu_{t+1}-\mathbb{E}_{d_2,d_1,\omega_{t+1}}[\mu_{t+1}]) \\
&+\frac{1}{2}\left.\frac{-\partial^3 V_{t+1}}{\partial E_{t+1}\partial \mathcal{L}_{t+1}\partial \Sigma_{t+1}}\right|_{\upsilon}(\mathcal{L}_{t+1}-\mathbb{E}_{d_2,d_1,\omega_{t+1}}[\mathcal{L}_{t+1}])(\Sigma_{t+1}-\mathbb{E}_{d_2,d_1,\omega_{t+1}}[\Sigma_{t+1}]) \\
&+\frac{1}{2}\left.\frac{-\partial^3 V_{t+1}}{\partial E_{t+1}\partial \mu_{t+1}\partial k_{t+1}}\right|_{\upsilon}(\mu_{t+1}-\mathbb{E}_{d_2,d_1,\omega_{t+1}}[\mu_{t+1}])(k_{t+1}-\mathbb{E}_{d_2,d_1,\omega_{t+1}}[k_{t+1}]) \\
&+\frac{1}{2}\left.\frac{-\partial^3 V_{t+1}}{\partial E_{t+1}\partial \mu_{t+1}\partial E_{t+1}}\right|_{\upsilon}(\mu_{t+1}-\mathbb{E}_{d_2,d_1,\omega_{t+1}}[\mu_{t+1}])(E_{t+1}-\mathbb{E}_{d_2,d_1,\omega_{t+1}}[E_{t+1}]) \\
&+\frac{1}{2}\left.\frac{-\partial^3 V_{t+1}}{\partial E_{t+1}\partial \mu_{t+1}\partial \mathcal{L}_{t+1}}\right|_{\upsilon}(\mu_{t+1}-\mathbb{E}_{d_2,d_1,\omega_{t+1}}[\mu_{t+1}])(\mathcal{L}_{t+1}-\mathbb{E}_{d_2,d_1,\omega_{t+1}}[\mathcal{L}_{t+1}]) \\
&+\frac{1}{2}\left.\frac{-\partial^3 V_{t+1}}{\partial E_{t+1}\partial \mu_{t+1}^2}\right|_{\upsilon}(\mu_{t+1}-\mathbb{E}_{d_2,d_1,\omega_{t+1}}[\mu_{t+1}])^2 \\
&+\frac{1}{2}\left.\frac{-\partial^3 V_{t+1}}{\partial E_{t+1}\partial \mu_{t+1}\partial \Sigma_{t+1}}\right|_{\upsilon}(\mu_{t+1}-\mathbb{E}_{d_2,d_1,\omega_{t+1}}[\mu_{t+1}])(\Sigma_{t+1}-\mathbb{E}_{d_2,d_1,\omega_{t+1}}[\Sigma_{t+1}]) \\
&+\frac{1}{2}\left.\frac{-\partial^3 V_{t+1}}{\partial E_{t+1}\partial \Sigma_{t+1}\partial k_{t+1}}\right|_{\upsilon}(\Sigma_{t+1}-\mathbb{E}_{d_2,d_1,\omega_{t+1}}[\Sigma_{t+1}])(k_{t+1}-\mathbb{E}_{d_2,d_1,\omega_{t+1}}[k_{t+1}]) \\
&+\frac{1}{2}\left.\frac{-\partial^3 V_{t+1}}{\partial E_{t+1}\partial \Sigma_{t+1}\partial E_{t+1}}\right|_{\upsilon}(\Sigma_{t+1}-\mathbb{E}_{d_2,d_1,\omega_{t+1}}[\Sigma_{t+1}])(E_{t+1}-\mathbb{E}_{d_2,d_1,\omega_{t+1}}[E_{t+1}]) \\
&+\frac{1}{2}\left.\frac{-\partial^3 V_{t+1}}{\partial E_{t+1}\partial \Sigma_{t+1}\partial \mathcal{L}_{t+1}}\right|_{\upsilon}(\Sigma_{t+1}-\mathbb{E}_{d_2,d_1,\omega_{t+1}}[\Sigma_{t+1}])(\mathcal{L}_{t+1}-\mathbb{E}_{d_2,d_1,\omega_{t+1}}[\mathcal{L}_{t+1}]) \\
&+\frac{1}{2}\left.\frac{-\partial^3 V_{t+1}}{\partial E_{t+1}\partial \Sigma_{t+1}\partial \mu_{t+1}}\right|_{\upsilon}(\Sigma_{t+1}-\mathbb{E}_{d_2,d_1,\omega_{t+1}}[\Sigma_{t+1}])(\mu_{t+1}-\mathbb{E}_{d_2,d_1,\omega_{t+1}}[\mu_{t+1}]) \\
&\left.+\frac{1}{2}\left.\frac{-\partial^3 V_{t+1}}{\partial E_{t+1}\partial \Sigma_{t+1}^2}\right|_{\upsilon}(\Sigma_{t+1}-\mathbb{E}_{d_2,d_1,\omega_{t+1}}[\Sigma_{t+1}])^2\right]. \qquad (3)
\end{aligned}
$$

Note that the expectation is taken over all terms in the expansion. The term on the right hand side of the equality on the first line is the zeroth-order term, the value function partial evaluated at υ, which is the state vector about which I do the Taylor expansion. This term is not random since it is evaluated at a specific state, so the expectation passes through.

The next five lines are the first-order Taylor terms. Here, another partial derivative of the value function is taken with respect to each state, then it is multiplied by the difference between the state and that state's value at υ, which is just the expectation of the state. The value function terms here have two partial derivatives, and are evaluated at υ. They also are not random so the expectation passes through those terms onto the terms capturing the difference between the state and its value at υ. The expectation of the difference between a variable and its expectation is zero so all first-order Taylor terms of the expansion are zero.

The second-order Taylor terms take another partial derivative, so these are the terms that include all of the triple value function partial derivatives. The triple partials are multiplied by the product of two differences between a state and that state's value at υ. The state in each difference corresponds to the state with respect to which the second and third partial derivatives are taken. As in zeroth-order and first-order Taylor terms, the triple partial derivatives are not random and the expectation passes through all of them onto the two difference terms. In the cases where at least one of the differences in the second-order Taylor term consists of a deterministic state $(k_{t+1}, E_{t+1}, \Sigma_{t+1})$, the expectation of the product of the two differences is again zero and the Taylor terms are zero. The remaining case is when both states in the product of differences are uncertain $(\mathcal{L}_{t+1},\ \mu_{t+1})$. After passing the expectation through the triple partial derivative terms, the difference terms are variances (e.g. $\mathbb{E}_{d_2,d_1,\omega_{t+1}}[(\mu_{t+1} - \mathbb{E}_{d_2,d_1,\omega_{t+1}}[\mu_{t+1}])^2] = \text{var}_{d_2,d_1,\omega_{t+1}}(\mu_{t+1})$ on the seventh to last line) or covariances (e.g. $\mathbb{E}_{d_2,d_1,\omega_{t+1}}[(\mu_{t+1} - \mathbb{E}_{d_2,d_1,\omega_{t+1}}[\mu_{t+1}])(\mathcal{L}_{t+1} - \mathbb{E}_{d_2,d_1,\omega_{t+1}}[\mathcal{L}_{t+1}])] = \text{cov}_{d_2,d_1,\omega_{t+1}}(\mu_{t+1}, \mathcal{L}_{t+1})$ on the eighth to last line). Similar logic follows for Taylor expansions for the other value function partials in equation (2).

Now define

$$ce := \left\{ k_{t+1}, E_{t+1}, \mathcal{L}_{t+1}\Big|_{d_1=\exp(\mu_c+\frac{1}{2}\sigma_c^2),\ \omega_{t+1}=\exp(\mu_\omega+\frac{1}{2}\sigma_\omega^2)}, \mu_t, 0 \right\}.$$

ce is a certainty state. Terms evaluated at ce are the terms that would arise if the policymaker happened to arrive at the same state at time t but states transitioned deterministically (with d_1, d_2, ω_t fixed to their expectations). I will add and subtract value function partial derivatives evaluated at ce, effectively adding 0, to obtain the certainty tax and the state uncertainty adjustment. After omitting the Taylor terms that are zero, the expectation of the cumulative emissions partial derivative is approximately

$$\begin{aligned}\mathbb{E}_{d_2,d_1,\omega_{t+1}}\left[\frac{-\partial V_{t+1}}{\partial E_{t+1}}\right] \approx \left.\frac{-\partial V_{t+1}}{\partial E_{t+1}}\right|_{ce} &+ \left[\left.\frac{-\partial V_{t+1}}{\partial E_{t+1}}\right|_{\upsilon} - \left.\frac{-\partial V_{t+1}}{\partial E_{t+1}}\right|_{ce}\right] \\ &+ \frac{1}{2}\left.\frac{-\partial^3 V_{t+1}}{\partial E_{t+1}\partial \mathcal{L}_{t+1}^2}\right|_{\upsilon} \mathrm{var}_{d_2,d_1,\omega_{t+1}}(\mathcal{L}_{t+1}) + \left.\frac{-\partial^3 V_{t+1}}{\partial E_{t+1}\partial \mathcal{L}_{t+1}\partial \mu_{t+1}}\right|_{\upsilon} \mathrm{cov}_{d_2,d_1,\omega_{t+1}}(\mathcal{L}_{t+1},\mu_{t+1}) \\ &+ \frac{1}{2}\left.\frac{-\partial^3 V_{t+1}}{\partial E_{t+1}\partial \mu_{t+1}^2}\right|_{\upsilon} \mathrm{var}_{d_2,d_1,\omega_{t+1}}(\mu_{t+1})\end{aligned}$$

The expectation of the fractional net output partial derivative is approximately

$$\begin{aligned}\mathbb{E}_{d_2,d_1,\omega_{t+1}}\left[\frac{-\partial V_{t+1}}{\partial \mathcal{L}_{t+1}}\right] \approx \left.\frac{-\partial V_{t+1}}{\partial \mathcal{L}_{t+1}}\right|_{ce} &+ \left[\left.\frac{-\partial V_{t+1}}{\partial \mathcal{L}_{t+1}}\right|_{\upsilon} - \left.\frac{-\partial V_{t+1}}{\partial \mathcal{L}_{t+1}}\right|_{ce}\right] \\ &+ \frac{1}{2}\left.\frac{-\partial^3 V_{t+1}}{\partial \mathcal{L}_{t+1}^3}\right|_{\upsilon} \mathrm{var}_{d_2,d_1,\omega_{t+1}}(\mathcal{L}_{t+1}) + \left.\frac{-\partial^3 V_{t+1}}{\partial \mathcal{L}_{t+1}^2\partial \mu_{t+1}}\right|_{\upsilon} \mathrm{cov}_{d_2,d_1,\omega_{t+1}}(\mathcal{L}_{t+1},\mu_{t+1}) \\ &+ \frac{1}{2}\left.\frac{-\partial^3 V_{t+1}}{\partial \mathcal{L}_{t+1}\partial \mu_{t+1}^2}\right|_{\upsilon} \mathrm{var}_{d_2,d_1,\omega_{t+1}}(\mu_{t+1})\end{aligned}$$

The expectation of the scale parameter partial derivative is approximately

$$\begin{aligned}\mathbb{E}_{d_2,d_1,\omega_{t+1}}\left[\frac{-\partial V_{t+1}}{\partial \Sigma_{t+1}}\right] \approx \left.\frac{-\partial V_{t+1}}{\partial \Sigma_{t+1}}\right|_{ce} &+ \left[\left.\frac{-\partial V_{t+1}}{\partial \Sigma_{t+1}}\right|_{\upsilon} - \left.\frac{-\partial V_{t+1}}{\partial \Sigma_{t+1}}\right|_{ce}\right] \\ &+ \frac{1}{2}\left.\frac{-\partial^3 V_{t+1}}{\partial \Sigma_{t+1}\partial \mathcal{L}_{t+1}^2}\right|_{\upsilon} \mathrm{var}_{d_2,d_1,\omega_{t+1}}(\mathcal{L}_{t+1}) + \left.\frac{-\partial^3 V_{t+1}}{\partial \Sigma_{t+1}\partial \mathcal{L}_{t+1}\partial \mu_{t+1}}\right|_{\upsilon} \mathrm{cov}_{d_2,d_1,\omega_{t+1}}(\mathcal{L}_{t+1},\mu_{t+1}) \\ &+ \frac{1}{2}\left.\frac{-\partial^3 V_{t+1}}{\partial \Sigma_{t+1}\partial \mu_{t+1}^2}\right|_{\upsilon} \mathrm{var}_{d_2,d_1,\omega_{t+1}}(\mu_{t+1})\end{aligned}$$

Substituting the Taylor expansion terms into equation (2) and then substituting equation (2) into equation (1) yields the full carbon tax expression.

C Certainty Tax

Figure A3 displays the certainty tax channel for each framework. The certainty tax is the remaining difference between the carbon tax in Figure 5 and the uncertainty channels in Figure 7 in the main text.

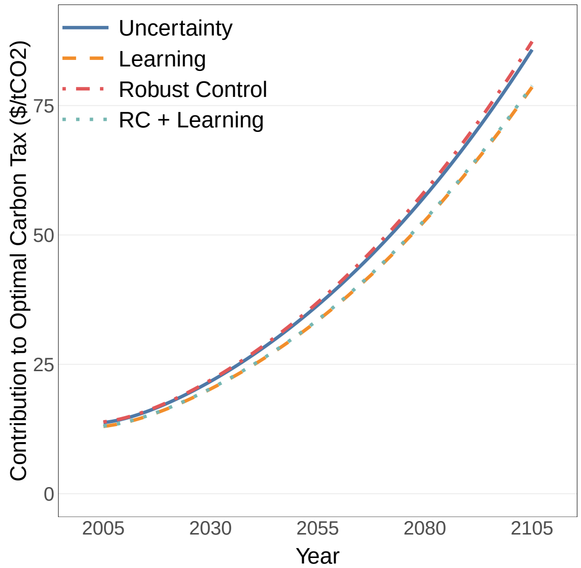

Figure A3: The mean certainty tax channel over 50,000 simulations for each framework. Each simulation samples a different set of $\{d_1, d_2, \omega_{2005}, ..., \omega_{2105}\}$.

References

Allen, Myles R, David J Frame, Chris Huntingford, Chris D Jones, Jason A Lowe, Malte Meinshausen, and Nicolai Meinshausen (2009) "Warming caused by cumulative carbon emissions towards the trillionth tonne," *Nature*, Vol. 458, No. 7242, pp. 1163–1166.

Berger, Loic and Massimo Marinacci (2017) "Model Uncertainty in Climate Change Economics," No. Cmcc, pp. 1–33.

Cai, Yongyang (2018) "Computational Methods in Environmental and Resource Economics," *Annual Review of Resource Economics*, Vol. 11.

Cai, Yongyang, Kenneth L Judd, Timothy M Lenton, Thomas S Lontzek, and Daiju Narita (2015) "Environmental tipping points significantly affect the cost- benefit assessment of climate policies," *Proceedings of the National Academy of Sciences*, p. 201503890.

Cai, Yongyang, Kenneth L. Judd, and Thomas S. Lontzek (Forthcoming) "The Social Cost of Carbon with Economic and Climate Risks," *Journal of Political Economy*.

Cai, Yongyang, Timothy M Lenton, and Thomas S Lontzek (2016) "Risk of multiple interacting tipping points should encourage rapid CO 2 emission reduction," *Nature Climate Change*, Vol. 6, No. 5, p. 520.

Dietz, Simon and Frank Venmans (2019) "Cumulative carbon emissions and economic policy: in search of general principles," *Journal of Environmental Economics and Management*, Vol. 96, pp. 108–129.

IPCC (2013) "Climate Change 2013: The Physical Science Basis. Contribution of Working Group I to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change," Cambridge University Press, Cambridge, United Kingdom and New York, NY, USA.

——— (2014) *Climate Change 2013: The Physical Science Basis: Working Group I Contribution to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change*: Cambridge University Press.

Knutti, Reto, Maria A. A. Rugenstein, and Gabriele C. Hegerl (2017) "Beyond equilibrium climate sensitivity," *Nature Geoscience*, Vol. 10, No. 10, p. 727.

Leduc, Martin, H Damon Matthews, and Ramón de Elía (2015) "Quantifying the limits of a linear temperature response to cumulative CO2 emissions," *Journal of Climate*, Vol. 28, No. 24, pp. 9955–9968.

Lontzek, Thomas S., Yongyang Cai, Kenneth L. Judd, and Timothy M. Lenton (2015) "Stochastic integrated assessment of climate tipping points indicates the need for strict climate policy," *Nature Climate Change*, Vol. 5, No. 5, pp. 441–444.

MacDougall, Andrew H and Pierre Friedlingstein (2015) "The Origin and Limits of the Near Proportionality Between Climate Warming and Cumulative CO2 Emissions," *Journal of Climate*, Vol. 28, No. 10, pp. 4217–4230.

Matthews, H Damon and Ken Caldeira (2008) "Stabilizing Climate Requires Near-Zero Emissions," *Geophysical Research Letters*, Vol. 35, No. 4.

Matthews, H Damon, Nathan P Gillett, Peter a Stott, and Kirsten Zickfeld (2009) "The Proportionality of Global Warming to Cumulative Carbon Emissions," *Nature*, Vol. 459, No. 7248, pp. 829–32.

Millar, Richard, Myles Allen, Joeri Rogelj, and Pierre Friedlingstein (2016) "The Cumulative Carbon Budget and its Implications," *Oxford Review of Economic Policy*, Vol. 32, No. 2, pp. 323–342.

van der Ploeg, Frederick (2018) "The Safe Carbon Budget," *Climatic Change*, Vol. 147, No. 1, pp. 47–59.

Traeger, Christian (2014) "A 4-Stated DICE: Quantitatively Addressing Uncertainty Effects in Climate Change," *Environmental and Resource Economics*, Vol. 59, No. 1, pp. 1–37.

Williams, Richard G, Philip Goodwin, Vassil M Roussenov, and Laurent Bopp (2016) "A framework to understand the transient climate response to emissions," *Environmental Research Letters*, Vol. 11, No. 1, p. 015003.