Appendices

In the following, we provide the details and derivations for various results in the main text. As in the main text, throughout the appendix we let lower-case variables capture potential realizations of random vectors. Additional plots and figures can be found in our online notebook:

`https://lphansen.github.io/mitigation_book`.

Appendix A HJB Equations

As part of dynamic programming backward induction, we start by giving the post-jump HJB equations. The post-jump value functions are then inputs into the pre-jump HJB equations. Recall that there are L possible jump realizations from the pre-jump setting: $L-1$ potential damage curvature realizations and one technology jump realization.

A.1 A numerically convenient transformation

Prior to constructing post-damage curvature jump value functions, we describe two numerically convenient transformations. Note that the damage function in this setting depends upon the temperature at which the damage information event occurred, $\hat{y}$. This dependence has a special structure that we exploit in computations. At the time of the jump this new state variable will be initialized at $\overline{y}$. Integrating (6), we construct the log damage function:

$$\hat{n}(y) = \begin{cases} \lambda_1 y + \frac{\lambda_2}{2} y^2 & \text{pre damage event} \\ \lambda_1 y + \frac{\lambda_2}{2}\hat{y}^2 + \frac{\lambda_2}{2}(y-\hat{y}+\overline{y})^2 + \frac{\lambda_3(\ell)}{2}(y-\hat{y})^2 - \frac{\lambda_2}{2}(\overline{y})^2 & \text{post damage event.} \end{cases} \tag{19}$$

Note that this construction depends implicitly on both ℓ and $\hat{y}$, a dependence that we now make more explicit. Define a new state variable by transforming the temperature anomaly:

$$z \stackrel{\text{def}}{=} y - \hat{y} + \overline{y}.$$

We transform this log damage function for $y \geqslant \hat{y}$ by constructing:

$$\begin{aligned}\hat{m}(z;\ell) &\stackrel{\text{def}}{=} \hat{n}(y;\ell,\hat{y}) - (\lambda_1\hat{y} + \frac{1}{2}\lambda_2\hat{y}^2) \\ &= \lambda_1(y-\hat{y}) + \frac{1}{2}\lambda_2(y-\hat{y}+\overline{y})^2 + \frac{1}{2}\lambda_3(\ell)(y-\hat{y})^2 - \frac{1}{2}\lambda_2(\overline{y})^2 \\ &= \lambda_1(z-\overline{y}) + \frac{1}{2}\lambda_2 z^2 + \frac{1}{2}\lambda_3(\ell)(z-\overline{y})^2 - \frac{1}{2}\lambda_2(\overline{y})^2\end{aligned}$$

Importantly, the $\hat{m}$ is constructed so as to depend on z, but not separately on y and $\hat{y}$.

We construct some post technology jump value functions by first using z and $\hat{m}$ and then transforming these value functions by subtracting the adjustment:

$$\lambda_1\hat{y} + \frac{1}{2}\lambda_2\hat{y}^2.$$

We subtract this term because log damages enter the HJB equations with a negative sign since damages reduce the utility of the planner. This two-step approach gives us a numerically convenient way to capture dependence of value functions on the temperature anomaly realization, $\hat{y}$, when the damage information event takes place.

A.2 Post damage and technology jumps

We start with HJB equation after both a technology jump and one of the damage curvature jumps has taken place. Denote the resulting value function $V^{\ell,L,\hat{y}}$ for $1 \leqslant \ell \leqslant L-1$ and the value function counterpart constructed using z and $\hat{m}$ as $W^{\ell,L}$. Importantly, $W^{\ell,L}$ does not depend on $\hat{y}$.

Since we are conditioning on the technology jump having already occurred, $\phi_0 = 0$, and there is no incentive for a further R&D investment targeted for the discovery of clean technology. The resulting output technology is:

$$C_t + I_t^k = \alpha K_t.$$

For the relevant state space, control set, and distortion set we use

$$x = \{z, \hat{k}\}, \qquad \Phi = \{i^k\}, \qquad \Gamma = \{h\}.$$

Observe that temperature no longer varies over time. Since we use the transformed temperature variable

z, we also use $\hat{m}$ instead of $\hat{n}$ in the HJB equation for $W^{\ell,L}$.

To compute $W^{\ell,L}$, we solve the HJB equation

$$\begin{aligned}0 = &\max_{i^k}\min_{h^k}\ \delta\log(\alpha - i^k) + \delta\hat{k} - \delta\hat{m}(\cdot;\ell) - \delta W^{\ell,L} + \xi_m\frac{\left|h^k\right|^2}{2}\\ &+ \frac{\partial W^{\ell,L}}{\partial\hat{k}}\left[-\mu_k + i^k - \frac{\kappa}{2}\left(i^k\right)^2 - \frac{|\sigma_k|^2}{2} + \sigma_k h\right] + \frac{\partial^2 W^{\ell,L}}{\partial\hat{k}^2}\frac{|\sigma_k|^2}{2}\end{aligned}$$

Using guess and verify, the robustly optimal choice variables are given by:

$$\begin{aligned}h^* =& -\frac{1}{\xi_m}\sigma_k\\ i^{k*} =& \frac{(1+\alpha\kappa) - \sqrt{(1+\alpha\kappa)^2 - 4\kappa(\alpha-\delta)}}{2\kappa},\end{aligned}$$

and the value function by:

$$W^{\ell,L}\left(z,\hat{k}\right) \quad = \quad \log\left(\alpha - i^{k*}\right) - \frac{1}{2\delta\xi_m}|\sigma_k|^2 + \frac{1}{\delta}\left[-\mu_k + i^{k*} - \frac{\kappa}{2}\left(i^{k*}\right)^2 - \frac{|\sigma_k|^2}{2}\right] + \hat{k} - \hat{m}(z;\ell).$$

As a second step, we then construct:

$$V^{\ell,L,\hat{y}} = W^{\ell,L} - \left(\lambda_1\hat{y} + \frac{1}{2}\lambda_2\hat{y}^2\right),$$

which follows since the damage adjustment, $\hat{n}$, enters negatively into the discounted objective of the planner.

A.3 Post-damage but pre-technology

We next compute each of the post-damage curvature jump, pre-technology jump value functions $V^{\ell,\hat{y}}$ where only a damage curvature jump has been realized for $\ell = 1, ..., L-1$.

Again, we start with the HJB for the value functions W^ℓ, $1 \leqslant \ell \leqslant L-1$ using $\hat{m}$. The modified state vector x, control set Φ, and distortion set Γ are

$$x = \{\hat{k}, z, \hat{r}\}, \qquad \Phi = \{i^k, i^r, e\}, \qquad \Gamma = \{h, q(\theta), g^L\}.$$

with $0 \leqslant e \leqslant \beta\alpha k$, $q(\theta) \geqslant 0, \sum_{\theta\in\Theta} q(\theta) = 1$, and $g^L \geqslant 0$. The function W^ℓ solves:

$$\begin{aligned}
0 \;=\; & \max_{i^k, i^r, e} \min_{h, q(\theta), g^L} \delta \log\left[\alpha - i^k - i^r - \alpha\phi_0\left(1 - \frac{e}{\beta\alpha k}\right)^{\phi_1}\right] + \delta\hat{k} - \delta\hat{m}(z;\ell) - \delta W^\ell(\hat{k}, z, \hat{r}) \\
& + \frac{\partial W^\ell}{\partial \hat{k}}(\hat{k}, z, \hat{r})\left[-\mu_k + i^k - \frac{\kappa}{2}\left(i^k\right)^2 - \frac{|\sigma_k|^2}{2} + \sigma_k h\right] + \frac{\partial^2 W^\ell}{\partial \hat{k}\,\partial \hat{k}'}(\hat{k}, z, \hat{r})\frac{|\sigma_k|^2}{2} \\
& + \frac{\partial W^\ell}{\partial y}(\hat{k}, z, \hat{r}) e\left[\sum_{\theta\in\Theta} \theta q(\theta)\right] + \frac{\partial^2 W^\ell}{\partial y\,\partial y'}(\hat{k}, z, \hat{r})\frac{|\varsigma|^2}{2} e^2 \\
& + \frac{\partial W^\ell}{\partial \hat{r}}(\hat{k}, z, \hat{r})\left[-\zeta + \psi_0 (i^r)^{\psi_1} \exp\left(-\psi_1\left(\hat{r} - \hat{k}\right)\right) - \frac{|\sigma_r|^2}{2} + \sigma_r h\right] + \frac{\partial^2 W^\ell}{\partial \hat{r}\,\partial \hat{r}'}(\hat{k}, z, \hat{r})\frac{|\sigma_r|^2}{2} \\
& + \mathcal{J}^L g^L\left[W^{\ell,L}(\hat{k}, z, \hat{r}) - W^\ell(\hat{k}, z, \hat{r})\right] + \xi_m \mathcal{J}^L\left[1 - g^L + g^L \log g^L\right] + \frac{\xi_m}{2} h' h + \xi_a \sum_{\theta\in\Theta} q(\theta) \log\left(\frac{q(\theta)}{p(\theta)}\right).
\end{aligned}$$

where, as an input, we impose function $W^{\ell,L}$ as constructed in Section A.2.

As a second step we compute:

$$V^{\ell,\hat{y}}(\hat{k}, y, \hat{r}) = W^\ell(\hat{k}, y - \hat{y} + \bar{y}, \hat{r}) - \left(\lambda_1 \hat{y} + \frac{1}{2}\lambda_2 \hat{y}^2\right).$$

Importantly,

$$V^{\ell,L,\hat{y}} - V^{\ell,\hat{y}} = W^{\ell,L} - W^\ell,$$

which is needed to make the HJB equation for W^ℓ relevant for the construction of $V^{\ell,\hat{y}}$.

A.4 Post-technology but pre-damage

Let V^L denote the pre-damage information event, post-technology jump value function. Once the technology jump has happened the subsequent damage curvature jumps are inconsequential. This is true because temperature anomaly no longer increases in response to economic production since $\phi_0^L = 0$ and it remains at a level for which damage curvature uncertainty is inconsequential.

In this case, we directly construct V^L restricting the temperature anomaly to reside in $0 \leqslant y < \hat{y}$. The function, V^L can be computed in essentially the same way as $W^{\ell,L}$, described in Section A.2. The state vector x, control set Φ, and distortion set Γ are

$$x = \{\hat{k}\}, \qquad \Phi = \{i^k\}, \qquad \Gamma = \{h\}.$$

The HJB equation for V^L is

$$0 = \max_{i^k} \min_{h^k} \delta \log(\alpha - i^k) + \delta \hat{k} - \delta \hat{n}(\cdot;\ell) - \delta V^L + \xi_m \frac{|h^k|^2}{2} + \frac{\partial V^L}{\partial \hat{k}} \left[-\mu_k + i^k - \frac{\kappa}{2} \left(i^k \right)^2 - \frac{|\sigma_k|^2}{2} + \sigma_k h \right] + \frac{\partial^2 V^L}{\partial \hat{k}^2} \frac{|\sigma_k|^2}{2},$$

which again can be solved using guess and verify, applying a close counterpart to the solution in Section A.2. Note, from the construction of equation (19) that $\hat{n}$ does not depend on ℓ over the temperature anomaly domain $0 \leqslant y < \hat{y}$.

A.5 Pre-technology and pre-damage

Finally, we present the HJB equation for the value function V prior to any jump realization. We need as inputs to this equation the post jump value functions: $V^{\ell,\hat{y}}$ for $1 \leqslant \ell \leqslant L-1$ and V^L. Importantly, we only need the function $V^{\ell,\hat{y}}$ for $\hat{y} = y$ as prior to jumps being realized, the recursive equation compares the current continuation value, V, to what happens hypothetically at temperature anomaly y if a jump happens to occur at y. Thus we define:

$$V^\ell \stackrel{\text{def}}{=} V^{\ell,y}$$

for $\ell = 1, 2, ...L-1$. The state vector x, control set Φ, and distortion set Γ are

$$x = \{\hat{k}, y, \hat{r}\}, \qquad \Phi = \{i^k, i^r, e\}, \qquad \Gamma = \{h, q(\theta), (g^\ell)_{\ell=1}^L\}$$

with $0 \leqslant e \leqslant \beta\alpha k$, $q(\theta) \geqslant 0, \sum_{\theta\in\Theta} q(\theta) = 1$, $g^\ell \geqslant 0$, and $g^L \geqslant 0$. The complete HJB equation is:

$$\begin{aligned}
0 \;=\; & \max_{i^k, i^r, e} \min_{h, q(\theta), g^\ell, g^L} \delta \log \left[\alpha - i^k - i^r - \alpha\phi_0 \left(1 - \frac{e}{\beta\alpha k} \right)^{\phi_1} \right] + \delta \hat{k} - \delta \hat{n}(y) - \delta V(\hat{k}, y, \hat{r}) \\
& + \frac{\partial V}{\partial \hat{k}}(\hat{k}, y, \hat{r}) \left[-\mu_k + i^k - \frac{\kappa}{2} \left(i^k \right)^2 - \frac{|\sigma_k|^2}{2} + \sigma_k h \right] + \frac{\partial^2 V}{\partial \hat{k}\, \partial \hat{k}'}(\hat{k}, y, \hat{r}) \frac{|\sigma_k|^2}{2} \\
& + \frac{\partial V}{\partial y}(\hat{k}, y, \hat{r}) e \left[\sum_{\theta\in\Theta} \theta q(\theta) \right] + \frac{\partial^2 V}{\partial y\, \partial y'}(\hat{k}, y, \hat{r}) \frac{|\varsigma|^2}{2} e^2 \\
& + \frac{\partial V}{\partial \hat{r}}(\hat{k}, y, \hat{r}) \left[-\zeta + \psi_0 (i^r)^{\psi_1} \exp\left(-\psi_1 \left(\hat{r} - \hat{k} \right) \right) - \frac{|\sigma_r|^2}{2} + \sigma_r h \right] + \frac{\partial^2 V}{\partial \hat{r}\, \partial \hat{r}'}(\hat{k}, y, \hat{r}) \frac{|\sigma_r|^2}{2} \\
& + \sum_{\ell=1}^L \mathcal{J}^\ell g^\ell \left[V^\ell(\hat{k}, y, \hat{r}) - V(\hat{k}, y, \hat{r}) \right] + \xi_m \sum_{\ell=1}^L \mathcal{J}^\ell \left[1 - g^\ell + g^\ell \log g^\ell \right] + \frac{\xi_m}{2} h' h + \xi_a \sum_{\theta\in\Theta} q(\theta) \log \left(\frac{q(\theta)}{p(\theta)} \right).
\end{aligned}$$

Appendix B Social planner preferences: IES not equal to one

We adopt a recursive representation of preferences in continuous time for the planner extended to allow $\rho \neq 1$. We start by forming the continuation value for each calendar date as follows:

$$\exp(V_t) = \left(\delta \int_0^\infty \exp(-\delta\tau) \left(C_{t+\tau}/N_{t+\tau}\right)^{1-\rho} d\tau\right)^{\frac{1}{1-\rho}}$$

where $\exp(V)$ gives an ordinally equivalent representation of preferences since $\exp(\cdot)$ is an increasing function. These preferences are dynamically consistent with a recursive representation.[33] The particular representation is homogeneous of degree one in consumption and $\exp(V)$.

The following differential equation gives the local representation expressed in terms of V:

$$\lim_{\epsilon \downarrow 0} \frac{1}{\epsilon} \left(V_{t+\epsilon} - V_t\right) = -\frac{\delta}{1-\rho}\left[\left(\frac{(C_t/N_t)^{1-\rho}}{\exp\left[(1-\rho)V_t\right]}\right) - 1\right] \tag{20}$$

which is a backward recursion linking future continuation values and current consumption to the current continuation value.

We introduce stochasticity under a unitary risk specification, which we implement by using conditional expectation for the recursive representation of V :

$$\lim_{\epsilon \downarrow 0} \frac{1}{\epsilon} \left[\mathbb{E}\left(V_{t+\epsilon} \mid \mathfrak{F}_t\right) - V_t\right] = -\frac{\delta}{1-\rho}\left[\left(\frac{(C_t/N_t)^{1-\rho}}{\exp\left[(1-\rho)V_t\right]}\right) - 1\right].$$

As in the case where $\rho = 1$, we use the adjustments derived in Section 3 to replace the local means of the continuation value process with the robust counterparts.

Appendix C Production function interpretation of abatement

The type of mathematical formulation used to represent abatement in our framework has been applied in the literature on climate-economics many times since the work of Nordhaus. See, for example, Nordhaus (2017). One possible interpretation is as follows. Consider a fixed proportions technology for energy and capital. In such a setting, a reduction in dirty energy less than the required proportion of capital must be

[33] This consistency is evident by raising both sides of the equation for $\exp(V_t)$ to the power $1-\rho$ and scaling by $\frac{1}{1-\rho}$ to obtain an ordinally equivalent transformation.

replaced by a clean alternative subject to a convex cost.

Recall the output equation

$$\text{output} = \alpha k \left[1 - \phi_0 \, (b)^{\phi_1}\right] \tag{21}$$

for

$$b = \left(1 - \frac{e}{\beta \alpha k}\right) \mathbf{1}_{\{e \leqslant \beta \alpha k\}}. \tag{22}$$

For notational simplicity, we drop the dependence of ϕ_0 on the technology discovery jump. We assume that $\phi_1 > 1$. The term

$$\alpha k \phi_0 \, (b)^{\phi_1}$$

is often referred to as an abatement cost, perhaps a cost that is even external to the firm.

Although we find this representation to be substantively interesting, in our framework we may also view the right-hand side of (21) as a production relation without reference to an abatement cost. This, of course, is a matter of interpretation, but it also impacts how we think of output and of plausible calibrations of ϕ_0 and ϕ_1. Both have their attractive features and limitations. The production function interpretation allows for curvature that is missed in a fixed proportion technology but assumes that all of the energy inputs are dirty. Observe that when $\mathcal{E}_t = 0$ the output or output net of the abatement cost is $\alpha K_t(1 - \phi_0)$, depending on how one views the technology and cost structure.

We now verify the desired mathematical properties that support a production function interpretation. We show that the first derivatives are positive and that the second-derivative matrix is negative semi-definite.

Notice first that the candidate production function is homogeneous of degree one in (e, k). Then consider the partial derivatives with respect to e:

$$\begin{aligned}
\frac{\partial}{\partial e}\text{output} &= \alpha k \, \phi_0 \, \phi_1 \, (b)^{\phi_1 - 1} \, \frac{1}{\beta \alpha k} \\
&= \frac{\phi_0 \, \phi_1}{\beta} \, (b)^{\phi_1 - 1} \\
&> 0 \\
\frac{\partial^2}{\partial e^2}\text{output} &= -\frac{\phi_0 \phi_1 (\phi_1 - 1)}{\beta^2 \alpha k} \, (b)^{\phi_1 - 2} \\
&< 0.
\end{aligned}$$

If $\phi_1 > 2$, both derivatives are zero at $e = \beta\alpha k$. This remains true for $e > \beta\alpha k$. Notice that the marginal product of emissions increases in the parameter ϕ_0.

Next, we consider derivatives with respect to k:

$$\begin{aligned}
\frac{\partial}{\partial k}\text{output} &= \alpha\left(1-\phi_0\,(b)^{\phi_1}\right) - \phi_1\phi_0\alpha k\,(b)^{\phi_1-1}\left(\frac{e}{\beta\alpha k^2}\right)\\
&= \alpha\left(1-\phi_0\,(b)^{\phi_1}\right) - \phi_1\phi_0\,(b)^{\phi_1-1}\left(\frac{e}{\beta k}\right)\\
\frac{\partial^2}{\partial k^2}\text{output} &= -\frac{\phi_0\phi_1(\phi_1-1)}{\beta^2\alpha k}\,(b)^{\phi_1-2}\left(\frac{e}{k}\right)^2\\
&< 0.
\end{aligned}$$

The first derivative is $\alpha(1-\phi_0) \geqslant 0$ when $k \to \infty$ and $\alpha > 0$ when $\beta\alpha k \leqslant e$. Given the negative second derivative, the first derivative remains positive for $k > 0$.

The simple relationship between the second derivatives with respect to e and k is to be anticipated, since the first derivatives are homogeneous of degree zero. Consistent with this relationship, the cross partial is

$$\frac{\partial^2}{\partial e\partial k}\text{output} = \frac{\phi_0\,\phi_1(\phi_1-1)}{\beta^2\alpha k}\,(b)^{\phi_1-2}\left(\frac{e}{k}\right).$$

The negative semi-definite Hessian matrix follows since

$$\begin{aligned}
\begin{bmatrix}\mathsf{r}_1 & \mathsf{r}_2\end{bmatrix}\begin{bmatrix}\frac{\partial^2}{\partial e^2}\text{output} & \frac{\partial^2}{\partial e\partial k}\text{output}\\ \frac{\partial^2}{\partial e\partial k}\text{output} & \frac{\partial^2}{\partial k^2}\text{output}\end{bmatrix}\begin{bmatrix}\mathsf{r}_1\\ \mathsf{r}_2\end{bmatrix} &= \frac{\partial^2}{\partial e^2}\text{output}\begin{bmatrix}\mathsf{r}_1 & \mathsf{r}_2\end{bmatrix}\begin{bmatrix}1 & -\frac{e}{k}\\ -\frac{e}{k} & \left(\frac{e}{k}\right)^2\end{bmatrix}\begin{bmatrix}\mathsf{r}_1\\ \mathsf{r}_2\end{bmatrix}\\
&= \frac{\partial^2}{\partial e^2}\text{output}\left(\mathsf{r}_1 - \mathsf{r}_2\frac{e}{k}\right)^2\\
&\leqslant 0.
\end{aligned}$$

To compute the implied energy demand elasticity, the first-order conditions for a given price p expressed in units of damaged consumption imply

$$p = \left(\frac{1}{n}\right)\frac{\partial}{\partial e}\text{output} = \frac{\phi_0\,\phi_1}{n\beta}\,(b^*)^{\phi_1-1}$$

for $e^* \leqslant \beta\alpha k$. Solving this equation for b^* gives:

$$b^* = \left(\frac{n\beta p}{\phi_0\,\phi_1}\right)^{\frac{1}{\phi_1 - 1}}$$

Using the original construction of b given in (22)

$$e^* = \beta\alpha k\,(1 - b^*) = \beta\alpha k\left[1 - \left(\frac{n\beta p}{\phi_0\,\phi_1}\right)^{\frac{1}{\phi_1 - 1}}\right]$$

for $e^* \leqslant \beta\alpha k$. Thus

$$\frac{d}{dp}e^* = -p^{\frac{2-\phi_1}{\phi_1 - 1}}\left(\frac{\beta\alpha k}{\phi_1 - 1}\right)\left(\frac{n\beta}{\phi_0\phi_1}\right)^{\frac{1}{\phi_1 - 1}} = -\frac{1}{p}\left(\frac{\beta\alpha k}{\phi_1 - 1}\right)b^*.$$

We then compute the elasticity, which is given by the product

$$\left(\frac{d}{dp}e^*\right)\left(\frac{p}{e^*}\right) = -\left(\frac{1}{\phi_1 - 1}\right)\frac{b^*}{1 - b^*}$$

When computing the price elasticities we plug in e^* to deduce an implied shadow price, p. We computed elasticities at the initial states and find:

- baseline ($\xi_m = \infty$): -0.053
- less aversion ($\xi_m = .1$): -0.079
- more aversion ($\xi_m = .05$): -0.129 .

This production function specification allows for $e < 0$, which could be viewed as carbon capture. A weaker requirement is that the implied

$$1 - \phi_0(b)^{\phi_1} > 0,$$

otherwise production would cease. While carbon capture could be allowed under this technology, we imposed the restriction that $e \geqslant 0$ in our computations. This restriction binds in only a small fraction of our state space and only after the more extreme damage curves are revealed ($\lambda_3 \geqslant 13/57 \approx 0.2281$, among the twenty damage specifications considered). It was not encountered in our simulations.

Appendix D Jump densities

The contributions to the *ex ante* density (relative to a uniform density) for the first jump time are computed as follows. The initial jump could be a damage information event or a technology discovery jump. First, by conditioning on the state variable history, we find that the *ex post* probability that any of the jumps has occurred in the time period t to be:

$$F_t = 1 - \exp\left[-\int_0^t \sum_{\ell=1}^{L} G_\tau^{*\ell} \mathcal{J}^\ell(X_\tau) d\tau\right],$$

which increases in t. We then compute the probability for any of the jumps occurring by differentiating F_t with respect to t:

$$f_t = \sum_{\ell=1}^{L} G_t^{*\ell} \mathcal{J}^\ell(X_t) \exp\left[-\int_0^t \sum_{\hat{\ell}=1}^{L} G_\tau^{*\hat{\ell}} \mathcal{J}^{\hat{\ell}}(X_\tau) d\tau\right].$$

We unpack this density construction so as to apply to each of the separate jump possibilities,

$$f_t^\ell \stackrel{\text{def}}{=} G_t^{*\ell} \mathcal{J}^\ell(X_t) \exp\left[-\int_0^t \sum_{\hat{\ell}=1}^{L} G_\tau^{*\hat{\ell}} \mathcal{J}^{\hat{\ell}}(X_\tau) d\tau\right], \quad \ell = 1, ..., L \tag{23}$$

which can be viewed as contributions to a density with respect to a uniform measure over the possible time horizons and a uniform measure over discrete states. To obtain the corresponding contributions to the *ex ante* densities of interest, we take conditional expectations over the future values of the state variables.

Appendix E Parameter values for the example economy

We quantitatively discipline our analysis by choosing model parameters based on (i) external empirical estimation and measurement of relevant economic and asset pricing outcomes, (ii) direct calibration using simplified versions of the model to match steady-state implications to empirical macroeconomic moments, and (iii) indirect calibration comparing outcomes from the simplified model settings to observable empirical macroeconomic and asset pricing moments. We validate these calibrated parameter choices by examining how the calibrated model fits additional macroeconomic and asset pricing empirical moments not included in the calibration. In addition, we perform sensitivity analysis for key parameters and specifications related to preferences, abatement technology, and climate damage as an additional validation of our quantitative

results. In what follows, we provide the details for calibration, validation, and sensitivity analysis.

The model parameters chosen for preferences, production, and the dynamics of the capital stock, knowledge stock, climate and climate damages are given in Tables 8 – 13. Steady-state values that are adapted to empirical moments for capital dynamics calibration are given in Table 14. The initial values and ranges for the state variables are provided in Tables 15 – 16. The descriptions of these choices are given below.

Parameter	Value
δ	.01
ρ	1
ξ_m	$\{.05, .1, \infty\}$
ξ_a	$\{.00175, .00349, \infty\}$

Table 8: Preferences

Parameter	Value
α	.115
ϕ_0	.5
ϕ_1	3
β	.12

Table 9: Production

In constructing Table 8, we set the model ambiguity parameter ξ_a so that at the initial date, the implied increase in the model average matched a run with model ambiguity aversion replaced by model misspecification concerns for all of the evolution equations, including the temperature evolution equation. We performed the computation for each choice of ξ_m. The implied values of ξ_a are reported in the last row.

Parameter	Value
μ_k	.043
κ	6.67
σ_k	[.01, 0, 0]

Table 10: Capital dynamics

Parameter	Value
ζ	0
ψ_0	.1
ψ_1	.5
σ_r	[0, 0, .0078]
χ	.0009

Table 11: Knowledge dynamics

Parameter	Value
$\bar{\theta}$	1.86 / 1000
ς	$[0,\ 1.2 \times \bar{\theta},\ 0]$

Table 12: Climate dynamics

Parameter	Value
λ_1	.00017675
λ_2	$2 \times .0022$
$\{\lambda_3(\ell)\}_{\ell=1,\ldots,L-1}$	$\left\{\frac{1}{3}\frac{\ell-1}{L-2}\right\}_{\ell=1,\ldots,L-1}$
d_0	1.5
d_1	.36
d_2	.0001
$\underline{y}$	1.5
$\bar{y}$	2.5

Table 13: Climate damages

Variable	Value
Investment/capital[a]: i^k	.090
Growth rate of capital[a]: η	.020
Marginal value of capital[a]: π	2.50

Table 14: Imposed steady states for the model specification without climate impacts.

State variables	values
K_0	739
Y_0	1.1
R_0	13.1

Table 15: State Variable Initial Values

State variables	range
$\log(K)$	$[4, 9]$
Y	$[0, 5]$
$\log(R)$	$[1, 8]$

Table 16: State variable ranges

E.1 Preferences

The subjective discount rate is set to a value of $\delta = .01$, consistent with the value used by others in the macroeconomics and asset pricing literature, including Barnett et al. (2020, 2022) and Barrage and Nordhaus (2024). The baseline choice of the IES is unitary, imposed by setting $\rho = 1$, which is a convenient

benchmark configuration. As a sensitivity analysis, we examine outcomes for $\rho = 3/4$ and $\rho = 4/3$. As Hansen et al. (2024) emphasize, the specification of ρ has an important impact for the overall consumption-savings choice and for this reason could interact with calibration choices for the productivity of capital. In our baseline analysis we choose $\xi_m \in \{.05, .1\}$. These values provide probability distortions that we view as reasonable based on the outcomes reported in the results for our numerical example and useful in illustrating the effects of the model uncertainty mechanism in our framework. Since the numerical magnitude of ξ_m reflects social preferences, we are very receptive to considering implications for alternative values of ξ_m. Our computations and code are easily implemented for such explorations.

E.2 Capital dynamics

The choices of parameters for the productivity and evolution of productive capital follow Barnett et al. (2022), who use an undamaged version of the consumption capital model to calibrate the economic growth rate to a value of 2%, a marginal value of capital of 2.5, and an investment-capital ratio of .09, consistent with empirical values from the BEA and World Bank databases. Specifically, the calibration approach uses externally estimated model parameters, as well as the growth rate η and steady state values for (π, i_k), as inputs for the relevant non-climate-state evolution equations and FOCs. We then invert the system of equations to derive the remaining model parameters of interest, exploiting the tractable recursive structure of this problem as follows:

a) From the first-order investment conditions, solve for κ given (i^k, π):

$$\pi = \left(1 - \kappa i^k\right)^{-1}$$

b) From the growth equation, solve for $\overline{\mu}_K$ given η, κ, and i^k:

$$\eta = -\mu_K + i^k - \frac{\kappa}{2}\left(i^k\right)^2,$$

c) Given (c, i^k), we determine α from the output constraint:

$$\alpha = i^k + \delta\pi.$$

For our chosen set of inputs, the resulting parameter values are given by $\kappa = 6.667, \mu_k = .043$ and $\alpha = .115$. Finally, the capital volatility is set to $\sigma_k = [.01, 0, 0]$, matching annual percent changes in the time series of GDP from the World Bank database.

E.3 Knowledge stock dynamics

For the R&D investment parameters, we make the following choices. We assume the depreciation of R&D stock is given by $\zeta = 0$ and set the curvature value to be $\psi_1 = .5$, each for simplicity and computational tractability purposes. We then set the R&D investment productivity scaling to be $\psi_0 = .1$ and set the multiplicative factor that translates US R&D stock to a global value in units of the arrival rate for our Poisson jump process for the technology discovery to be $\chi = .0009$. These values ensure that the simplified model without a damage information event and without uncertainty aversion generates initial R&D investment values below 1% of GDP, which is the peak value for major U.S. R&D investment programs as estimated by Stine (2008). Moreover, these parameter values are chosen to be consistent with estimates of returns to R&D investment from Lucking et al. (2019) and Bloom et al. (2019), and to produce an expected arrival time for the technology discovery jump consistent with proposed policy timelines of achieving carbon neutrality between 2050 and 2080. Alhamdan et al. (2023) similarly cite major U.S. R&D investment programs such as the Manhattan project as an anticipated timeline for one type of major green energy innovation, nuclear fusion development. The volatility of knowledge stocks is set to $\sigma_r = [0, 0, .0078]$, matching the annual percent changes in the time series of U.S. R&D stocks from the BLS database.

E.4 Climate dynamics and damages

Our choices for the emissions component of the production technology (i.e., the abatement cost parameters following the interpretation of Nordhaus and others), reported in Table 9, are as follows. We set $\phi_1 = 3$, similar to the values of the estimated parameter of Cai and Lontzek (2019) and Barrage and Nordhaus (2024). Although the value of ϕ_0 is highly uncertain, we choose $\phi_0 = .5$ as a reasonable benchmark for the fraction of lost output to achieve zero emissions. We also consider $\phi_0 = .1$, consistent with Barrage and Nordhaus (2024), for a sensitivity analysis comparison. The value of the intensity of the output emissions $\beta = .12$ comes from the implied emissions intensity value for 2020 from Cai and Lontzek (2019), which together with the chosen subjective discount rate generates annual carbon emissions consistent

with estimates by Figueres et al. (2018) of about 10.5 GtC in a simplified version of the model without uncertainty aversion or jumps.

The climate dynamics parameter values are set as follows. Specifically, $\bar{\theta}$ is the average value across 144 climate model outcomes constructed from pulse experiments provided by Joos et al. (2013) and Geoffroy et al. (2013), and matches the values reported in Masson-Delmotte et al. (2021). The value for ς is chosen based on a parameter uncertainty interpretation of TCRE specification. Specifically, the chosen value is based on the implied standard deviation associated with the coefficient of the Matthew's approximation for a constant emissions path near the current value.

As noted previously, for climate damages, the functional form and values of λ_1, λ_2, and $\lambda_3(\ell)$ are roughly consistent with the range of climate damage specifications from the literature, including Nordhaus (2019), Weitzman (2012), and the very recent study by Waidelich et al. (2024). We choose baseline values for our intensity function of $\underline{y} = 1.5$ and $\overline{y} = 2.5$, motivated by the literature on climate thresholds and tipping points, which we elaborate on in Remark E.1. We set $\mathsf{d}_0 = 1.5$ and $\mathsf{d}_2 = .0001$, and then impose $\mathsf{d}_1 = .36$, such that the modal point of the density implied by the jump intensity is $y_m \approx 2.0$ based on a simplified model of the temperature pathway.[34]

Remark E.1. *The choices of $\underline{y}$ and $\overline{y}$, used for the range over which the damage curvature jump realizations occur, are motivated by the literature on "climate tipping points" and "climate thresholds." Specifically, climate scientists, economists, and others are concerned about potentially drastic shifts in climate-carbon dynamics and realized economic damages resulting from continued climate change. Drijfhout et al. (2015) and Armstrong McKay et al. (2022) enumerate and characterize various potential thresholds of potential consequence, with the latter stating that "[t]he Earth may have left a safe climate state beyond $1^{\circ}C$ global warming. A significant likelihood of passing multiple climate tipping points exists above $\sim 1.5^{\circ}C$, particularly in major ice sheets. The probability of a tipping point increases further in the Paris range of 1.5 to $< 2^{\circ}C$ warming." Rogelj et al. (2018) and Rogelj et al. (2019) suggest a $1.5^{\circ}C$ target to limit damages from such events, while noting $2.0^{\circ}C$ as a potentially more plausible target that also includes risks of more severe damage consequences.*

However, consensus on the plausibility and timeline for such thresholds is far from decided. Hansen et al. (2025) notes that while tipping point concerns are real, "[m]any tipping point processes are reversible if

[34] For an elaboration of the ramifications of our damage infrmation event jump intensity and some alternatives, see the Online Appendix P.

Earth cools, but the recovery time varies and may be long for some feedbacks" and that the most threatening and catastrophic thresholds are still likely to occur far into the future. Ritchie et al. (2021) conclude that "the point of no return" for climate thresholds is highly uncertain based on their analysis using recent developments in dynamical systems theory.

We note that our damage information event specification is not an immediate drop in output or a significant realization of climate damage. Instead, our damage curvature jumps represent information revelation about the severity of climate damage function curvature going forward, including the possibility of more severe consequences similar to tipping points or thresholds. Thus, our damage function specification allows for concerns about potential tipping points or thresholds in such a way that we believe is still consistent with the more moderate dynamics emphasized by Hansen et al. (2025) and Ritchie et al. (2021), rather than stark "falling off the cliff" specifications used in some other settings.

E.5 Climate model uncertainty

We construct 144 different TCRE's by using 100 GtC pulse experiment results of Joos et al. (2013) tracing out the resulting carbon in the atmosphere for 9 different models. We then use these as inputs into 16 model approximations for temperature responses using the approximation in Geoffroy et al. (2013) to build the collection of $\theta(\ell)$'s used in our analysis.

E.6 Initial values

The initial value of capital is set so that our initial GDP matches the 2020 World GDP value of \$85 trillion estimated by the World Bank National Accounts data. With our choice of $\alpha = .115$, we end up with $K_0 = 739.13$. The initial value of knowledge capital is set to $R_0 = 13.1$, which converts and scales the 2020 value for US R&D stocks in the BLS database to a global value. As noted above, this value scaled by χ leads to an expected arrival time of a breakthrough technology discovery jump of around 30-50 years in the simplified model with no damage information event and without uncertainty aversion. The initial value of atmospheric temperature anomaly is set to $Y_0 = 1.1$ degrees Celsius to match recent estimates from the IPCC AR6 (Masson-Delmotte et al. 2021).

E.7 Some additional asset-pricing implications

For additional context about the model's calibration and quantitative implications, we examine (shadow) local asset pricing moments in the form of the implied risk-free rate from the model solution. We examine the outcomes under the baseline ($\xi_m = \infty$), less aversion ($\xi_m = .1$), and more aversion ($\xi_m = .05$) specifications. We calculate numerically the one-month risk-free rate using simulations that account for all relevant uncertainty components. The simulation-based expression for the monthly, risk-free rate expressed in terms of annualized time units is given by

$$r_{f,t} = \delta - \frac{1}{\epsilon} \log \left(\tilde{\mathbb{E}} \left[\frac{C_t/N_t}{C_{t+\epsilon}/N_{t+\epsilon}} \mid \mathfrak{F}_t \right] \right)$$

where $\tilde{\mathbb{E}}[\cdot \mid \mathfrak{F}_t]$ denotes the expectation under the uncertainty-adjusted probability distribution, $N_t = \exp[\hat{n}(Y_t)]$, and $\epsilon = 1/12$ is the monthly specification. The results are provided in Table 17. Importantly, the values for each uncertainty aversion specification are within the range of reasonable values used in the asset pricing literature based on empirical estimates of the risk-free rate. Moreover, the decreasing values of the risk-free rate as the uncertainty aversion increases reflect increasing concerns of negative outcomes that will impact future consumption. Note that because we are examining equilibrium solutions under social optimality, these do not represent market-based prices from the decentralized solution. Nevertheless these values provide context for the quantitative plausibility of our model's calibration.

	baseline ($\xi_m = \infty$)	less aversion ($\xi_m = .1$)	more aversion ($\xi_m = .05$)
risk-free rate ($r_{f,0}$)	2.8%	2.6%	2.2%

Table 17: Annualized one-month risk-free rate for the full model computed under the baseline ($\xi_m = \infty$), less aversion ($\xi_m = .1$), and more aversion ($\xi_m = .05$) specifications. The values are computed at the initial time period using stochastic simulations under the corresponding (uncertainty-adjusted) probability distribution.

Appendix F Model sensitivity

We now discuss sensitivity analysis exercises that provide further insight into the economic mechanisms underlying our quantitative implications. We first outline the implications for changing the subjective rate of discount (δ). We then explore the effects of changing the intertemporal elasticity of substitution (IES)

(ρ). We also provide the detailed results for sensitivity to uncertainty aversion (ξ_m) in the technology jump only model. After that, we examine the consequences of changing the abatement technology scaling parameter (ϕ_0). Additionally, we compare the effect of changing the combination of the R&D knowledge stock scaling parameter (χ) and the R&D investment effectiveness parameter (ψ_0), while briefly noting implications for changing the knowledge stock volatility (σ_r) as well. To reduce the extensive computational cost of calculating the sensitivity analysis outcomes, we use a coarse grid resolution rather than the fine grid resolution applied for our main analysis. As shown in Appendix G, this grid choice has no impact on the qualitative implications of the model solutions, and only minor quantitative impacts that are most noticeable in the more risk aversion case. The exception to this grid choice are the results reported in Figure 11 and Table 21, which use the fine grid resolution. We provide a subset of results and intuition in what follows, with detailed analysis and computational results provided in the Online Appendix.

F.1 Subjective rate of discount

The prior environmental economics literature has explored the sensitivity of the SCC to changes in the "discount rate." Often these analyses feature the implied discount rate used in computing present values abstracting from stochastic discounting. In our setting with misspecification, stochastic discounting and endogenously determined changes in the probability measure are central ingredients in valuation, as is expected from our understanding of asset pricing.

	SVRD			SCGW		
ξ_m	$\delta = .02$	$\delta = .015$	$\delta = .01$	$\delta = .02$	$\delta = .015$	$\delta = .01$
∞	1.90	2.79	3.90	37.4	47.8	60.6
.1	3.03	4.63	6.48	62.4	88.5	119.7
.05	4.49	7.16	10.05	112.9	178.0	254.6

Table 18: Initial SVRD (columns 1 - 3) and SCGW (columns 4 - 6) for alternative values of the subjective discount rate: $\delta = .02$, $\delta = .015$, and $\delta = .01$. Outcomes are shown for the three different values of the aversion parameter ξ_m in the respective rows.

In Table 18 we report computations for both the social value of R&D and the social cost of global warming. Our results show that increasing δ leads to drops in valuations, confirming a sensitivity often noted in the literature on environmental economics. We see drops in the SVRD by about 28% to 29% by increasing δ from .01 to .015, and by an additional 32% to 37% when increasing δ to .02. For the SCGW, the reductions are slightly more modest, dropping by about 21% to 30% by increasing δ from .01 to .015,

and by an additional 22% to 37% when increasing δ to .02.

F.2 Intertemporal elasticity of substitution (IES)

We also consider alternative values of the IES using the preferences defined in Appendix B. Much of the asset pricing literature has studied the consequences of changing the IES on asset valuation within the setting of an endowment economy. However, our economy is a production economy and changing the IES has a big impact on production outcomes.

	R&D investment			Capital investment			Emissions		
ξ_m	$\rho = 3/4$	$\rho = 1$	$\rho = 4/3$	$\rho = 3/4$	$\rho = 1$	$\rho = 4/3$	$\rho = 3/4$	$\rho = 1$	$\rho = 4/3$
∞	.014	.007	.003	.878	.773	.702	9.104	9.265	9.377
.1	.034	.018	.009	.855	.760	.696	8.626	8.852	9.041
.05	.075	.044	.024	.809	.730	.679	7.829	8.169	8.471

Table 19: Initial investments and emissions for different specifications of the IES. R&D and capital investment are reported as ratios relative to total output. Outcomes are shown for the three different values of the aversion parameter ξ_m in the respective rows.

Table 19 shows the results for two alternative specifications of the IES that differ from unity: higher ($\rho = 3/4$) and lower ($\rho = 4/3$) IES scenarios. As expected from growth models, investments in both types of capital relative to output are higher when the elasticity is greater (ρ is smaller). Conversely, emissions decrease when the elasticity is greater (ρ is smaller). The sensitivities in emissions prior to any jumps are consistent with the investment differences since higher investment helps support more initial growth in output.[35]

Inverse IES	SVRD	SCGW
$\rho = 3/4$	8.83	161.16
$\rho = 1$	6.48	119.75
$\rho = 4/3$	4.67	89.66

Table 20: Social values at the initial time period for less aversion to misspecification uncertainty.

Table 20 shows the social value of the R&D stock and the social cost of global warming for $\rho = 3/4$,

[35] When discussing our paper, Eric Renault reminded us that there can be seemingly counterintuitive interactions between the IES and the risk aversion in recursive utility models, noting that the latter is not a "pure" risk aversion parameter. Indeed, in dynamic stochastic settings, the intertemporal composition of risk comes into play when exploring the preference implications. See Cai and Lontzek (2019) and Hambel et al. (2021) for related discussions when exploring the SCC. The changing implications for consumption and investment induced by changes in the IES make the risk aversion comparisons all the more tricky.

$\rho = 1$, and $\rho = 4/3$. The magnitude of the social valuations of the corresponding state variables increases when the elasticity is greater (ρ is smaller). These increases in the SVRD and SCGW are consistent with the investment and emissions responses shown in Table 19. Valuations for the baseline and more aversion cases across this range of IES values are given in Online Appendix M.

F.3 Uncertainty aversion sensitivity

To further detail the model sensitivity to uncertainty aversion, we provide the numerical values for the SVRD and R&D investment outcomes corresponding to Figure 11. The numerical values, provided in Table 21, show the responses to increasing uncertainty aversion (decreasing ξ_m) for the technology jump only model where the only uncertainty channel is the technology jump channel. As noted in the main text, initial increases in uncertainty aversion (lower values of ξ_m) lead to higher SVRDs that augment R&D investment. As the misspecification concerns continue to increase to very high values (very low values of ξ_m), the SVRDs drop and thus the R&D investment does as well. The SVRD increases to more than double the baseline case at its maximum, and then diminishes to about 30% higher than the baseline case for the smallest value of ξ_m that we examine. The R&D investment value rises to nearly 5.5 times higher than the baseline case at its maximum, and then decreases to a value just under double the baseline case for the smallest value of ξ_m that we examine. This non-monotonic response coincides with the net utility benefit from a potential discovery initially being the dominant effect from uncertainty aversion, and then the uncertainty-adjusted probabilities pushing more weight to the likelihood of a delayed technology success takes over for more severe uncertainty aversion concerns.

ξ_m	∞	.10	.05	.02	.015	.01	.009	.008	.007	.006	.005
SVRD	4.26	6.44	8.90	9.92	9.35	8.21	7.85	7.44	7.02	6.55	5.56
R&D inv.	.0079	.0180	.0344	.0426	.0379	.0292	.0268	.0240	.0214	.0186	.0134

Table 21: Social value of R&D (technology jump only) as a function of the robustness parameter, ξ_m. Row 1 shows the value of ξ_m, row 2 shows the knowledge stock valuations, and row 3 shows the R&D investment-to-output ratios.[36]

[36]The implied SVRD totals from this table differ from those reported in Table 6 by between 1.2 and 5.6 percent because of our use of simulation methods to compute the flow decompositions.

F.4 Abatement technology sensitivity

As we noted in Appendix C, economic models of climate change often make reference to an "abatement cost" with ϕ_0 representing the magnitude of the cost of full abatement and ϕ_1 capturing the convexity of the cost function. Under this alternative interpretation, previous work led by Nordhaus use an alternative calibration of $\phi_0 = .1$. Table 22 shows the R&D-to-output ratio and emissions for our baseline value of $\phi_0 = .5$, the Nordhaus-motivated alternative of $\phi_0 = .1$, and an intermediate value of $\phi_0 = .3$. We find that decreasing the value of ϕ_0 leads to i) substantially less R&D as a fraction of output, ii) a more modest but notable proportional reduction in emissions, and iii) the same qualitative response to model uncertainty implications. These responses reflect the reduced output loss induced by a reduction in emissions that would occur for lower values of ϕ_0.

	R&D investment			emissions		
ξ_m	$\phi_0 = .5$	$\phi_0 = .3$	$\phi_0 = .1$	$\phi_0 = .5$	$\phi_0 = .3$	$\phi_0 = .1$
∞	.0066	.0055	.0030	9.26	9.03	8.40
.1	.0183	.0125	.0047	8.85	8.63	8.11
.05	.0438	.0266	.0070	8.17	7.99	7.73

Table 22: Initial robust actions for two alternative initial specifications of ϕ_0. R&D investment is reported as a ratio relative to total output. Outcomes are shown for the three different values of the aversion parameter ξ_m in the respective rows.

Online Appendix N shows the numerical outcomes for changing the value of ϕ_0 for different choices of the R&D knowledge stock related parameters, namely χ and ψ_0. Importantly, the qualitative implications for changing ϕ_0 and for the impact of uncertainty aversion persist in these alternative parameter specifications.

F.5 R&D knowledge stock parameter sensitivity

Given the novel implications of model uncertainty emphasizing proactive R&D investment early prior to the revelation of damage severity and targeted responses after the severity is fully revealed, we also provide sensitivity analysis relative to the specification of R&D investment and knowledge stock evolution in our model. Specifically, we examine model outcomes for different combinations of the R&D knowledge stock scaling parameter (χ), the knowledge stock volatility (σ_r), and the R&D investment effectiveness parameter (ψ_0). First, in terms of quantitative implications, the most significant impacts come from variation in χ and ψ_0, whereas changes in σ_r have a trivial impact. As such, we do not report results for changes in

σ_r, but they are available upon request. Importantly, the trivial impact of shifting the knowledge stock volatility value further highlights the significance of the technology jump in our model, and the role of uncertainty about this jump.

Table 23 shows the 9 variations of the R&D related parameter values we explore: $\psi_0 \in \{.07, .11, .15\}$ and $\chi \in \{.0014, .0009, .0007\}$.

First, we can see that increasing ψ_0, which improves the efficiency of R&D investment, leads to a decrease in R&D investment, and an increase in the probability of a technology jump occurring first. The (conditional) median year of a technology jump occurring slightly increases for $\xi_m = \infty$, but decreases for $\xi_m = .1$ and $\xi_m = .05$. Next, decreasing χ, which decreases the likelihood arrival rate of a technology jump occurrence for a given level of the knowledge stock, leads to an increase in R&D investment and an increase in the (conditional) median year of a technology jump occurring for all cases except when $\psi_0 = .07$ and $\xi_m = .05$. The probability of a technology jump occurring first decreases in all cases. The final, and most important, takeaway when comparing outcomes across uncertainty aversion scenarios is that the qualitative results are essentially identical to the results reported in the main text across all R&D knowledge stock parameters combinations we explore. As mentioned before, Online Appendix N shows the numerical outcomes for changing the R&D knowledge stock related parameters χ and ψ_0 across different choices of ϕ_0. Again, the qualitative implications for the impact of uncertainty aversion persist in these alternative parameter specifications.

ψ_0	χ	emissions	R&D investment	technology-first probability	conditional tech-jump median year
.07	.0014	9.31	.0042	.6403	18.49
.07	.0009	9.15	.0072	.5206	21.14
.07	.0007	9.05	.0091	.4555	22.49
.11	.0014	9.41	.0040	.6861	18.82
.11	.0009	9.27	.0067	.5748	21.60
.11	.0007	9.19	.0084	.5122	23.01
.15	.0014	9.48	.0037	.7223	18.95
.15	.0009	9.35	.0061	.6184	21.77
.15	.0007	9.28	.0076	.5586	23.20

(A) Simulation outputs under baseline across combinations of R&D parameter values. R&D investment is reported as a ratio relative to total output. Technology-first probability is the probability that the technology jump is the first jump to occur. Conditional tech-jump median year is the corresponding conditional median year of a technology jump occurring.

ψ_0	χ	emissions	R&D investment	technology-first probability	conditional tech-jump median year
.07	.0014	8.79	.0175	.4170	22.49
.07	.0009	8.38	.0316	.1931	27.64
.07	.0007	8.15	.0388	.0980	30.34
.11	.0014	9.16	.0095	.5837	21.00
.11	.0009	8.87	.0178	.4050	25.13
.11	.0007	8.69	.0236	.3012	27.49
.15	.0014	9.33	.0067	.6625	20.40
.15	.0009	9.11	.0121	.5174	24.01
.15	.0007	8.97	.0161	.4289	26.02

(B) Simulation outputs under less aversion across combinations of R&D parameter values. R&D investment is reported as a ratio relative to total output. Technology-first probability is the probability that the technology jump is the first jump to occur. Conditional tech-jump median year is the corresponding conditional median year of a technology jump occurring.

ψ_0	χ	emissions	R&D investment	technology-first probability	conditional tech-jump median year
.07	.0014	7.95	.0495	.0455	33.28
.07	.0009	7.56	.0552	.0011	37.61
.07	.0007	7.41	.0521	.0001	34.03
.11	.0014	8.69	.0251	.3564	26.06
.11	.0009	8.21	.0425	.1099	33.05
.11	.0007	7.96	.0499	.0346	36.93
.15	.0014	9.07	.0137	.5513	23.02
.15	.0009	8.69	.0264	.3233	28.38
.15	.0007	8.45	.0349	.1954	31.64

(C) Simulation outputs under more aversion across combinations of R&D parameter values. R&D investment is reported as a ratio relative to total output. Technology-first probability is the probability that the technology jump is the first jump to occur. Conditional tech-jump median year is the corresponding conditional median year of a technology jump occurring.

Table 23: Simulation outputs for the full model for varying R&D knowledge stock parameter configurations. Column 1 shows the value of ψ_0 used. Column 2 shows the value of χ used. Column 3 shows the emissions value. Column 4 shows the R&D investment-to-output ratio value. Column 6 shows the median year of a technology jump. Rows 1 through 3 assume a low R&D investment efficiency parameter. Rows 4 through 6 assume a medium R&D investment efficiency parameter. Rows 7 through 9 assume a high R&D investment efficiency parameter. Rows 1, 4, and 7 assume a small knowledge stock (inverse) scaling. Rows 2, 5, and 8 assume a medium knowledge stock (inverse) scaling. Rows 3, 6, and 9 assume a large knowledge stock (inverse) scaling. Panel (A) shows outcomes for the baseline uncertainty neutral case. Panel (B) shows outcomes for the less uncertainty aversion case. Panel (C) shows outcomes for the more uncertainty aversion case.

Appendix G Computational method

We provide an outline of the computational algorithm used to solve the HJB equations in what follows, and direct the reader to our online appendix for full details.

G.1 Policy iteration

For simplicity, we denote the state space, control set, and distortion set by:

$$X = \{\log k, y, \log r\}, \qquad \Phi = \{i^k, i^r, \mathcal{E}\}, \qquad \Gamma = \{h, q(\theta), g^\ell\}.$$

State variables, controls, and distortions that are not applicable to an HJB equation at a particular jump state are omitted. The value function $\hat{v}(X)$ depends on the state X.

A pseudo-code of the numerical algorithm for solving the HJB equation via policy iteration is given in Algorithm 1. The algorithm involves an iterative process that alternates between policy improvement and policy evaluation steps. The subsequent sections then explain further details for each step in the algorithm.

Algorithm 1 Solving the HJB equation via policy iteration.

Input: Initial guesses $\hat{v}^0$, Φ^0, Γ^0; values for $\Delta\tau$, ε_{relax}, ϵ_{KSP}, ϵ_{outer}, N_{max}

Output: Optimal value function $\hat{v}^*$

$\hat{v} \leftarrow \hat{v}^0$

Repeat

Step 1: Compute optimal actions Φ^* from the first-order condition of the maximization problem using $\hat{v}$, Φ, Γ. Update using numerical relaxation scheme with weight ε_{relax}.

Step 2: Compute optimal probability distortions Γ^* from the first-order condition of the minimization problem using $\hat{v}$, Φ^*, Γ.

Step 3: Update the value function $\hat{v}^*$ by solving (conditionally) the HJB PDE with a false transient method. $\hat{v}$, Φ^*, Γ^* are used in the PDE. Convergence tolerance given by ϵ_{KSP}.

Step 4: $\hat{v}^{prior} \leftarrow \hat{v} \quad \hat{v} \leftarrow \hat{v}^* \quad \Phi \leftarrow \Phi^* \quad \Gamma \leftarrow \Gamma^*$

Until $\left\| \frac{\hat{v}^* - \hat{v}^{prior}}{\Delta\tau} \right\|_\infty \leqslant \epsilon_{outer}$ or $N \geqslant N_{max}$

G.2 Updating rules for Φ and Γ at policy improvement steps

In solving HJB equations, we often encounter complex, highly non-linear equations for our optimal control choices that do not admit analytical solutions. To address this challenge, we implement a numerical relaxation updating scheme. The relaxation weight ε determines the updating rate of the control vector

Φ. The control vector is updated according to

$$\Phi_{n+1} = \varepsilon_{relax}\Phi^* + (1 - \varepsilon_{relax})\Phi_n$$

where Φ_{n+1} is the next iteration's optimal control vector, Φ^* is the updating optimal control vector value using the current value function solution, and Φ_n is the previous iteration's optimal control vector. The control updating is embedded in the outer policy/false-transient iteration. For emissions, the implementation also compares feasible stationary and boundary candidates before applying the relaxation update.

We note that every probability distortion from misspecification aversion has an analytical expression as a function of the value function and its derivatives. These can be computed and plugged in directly to the HJB expressions, and the iterative process for the HJB equation solution updates these values as the value function is updated until convergence is achieved.

G.3 Solving the linear PDE equation at the policy evaluation step

Updating value functions, given the state variables, controls, and distortions, is done by solving the linear PDE implicitly. To mitigate the potential instability of the non-linear HJB, we add a false transient (time) dimension and iterate the solution with an implicit time stepping procedure until the false time derivative converges to zero. In addition, we solve the linear system using a stabilized Bi-Conjugate Gradient solver with the ILU preconditioner provided by PETSc, a widely used software library developed by researchers at Argonne National laboratory to efficiently solve sparse linear/nonlinear systems[37]. We outline how we construct the linear system from the HJB equation in what follows.

First, note that we can write the HJB equation from our model using the following form:

$$A\hat{v} + B_{\hat{k}}\frac{\partial \hat{v}}{\partial \hat{k}} + B_y\frac{\partial \hat{v}}{\partial y} + B_{\hat{r}}\frac{\partial \hat{v}}{\partial \hat{r}} + C_{\hat{k}}\frac{\partial^2 \hat{v}}{\partial \hat{k}\,\partial \hat{k}'} + C_y\frac{\partial^2 \hat{v}}{\partial y\,\partial y'} + C_{\hat{r}}\frac{\partial^2 \hat{v}}{\partial \hat{r}\,\partial \hat{r}'} + D = 0.$$

More generally, our system can be expressed in the conditionally linear form:

$$\begin{aligned} 0 \;=\; & V_t(x) + \mathbb{A}(x; V, V_x, V_{xx})V(x) + \mathbb{B}(x; V, V_x, V_{xx})V_x(x) \\ & + \frac{1}{2}tr[\mathbb{C}(x; V, V_x, V_{xx})'V_{xx}(x)\mathbb{C}(x; V, V_x, V_{xx})] + \mathbb{D}(x; V, V_x, V_{xx}) \end{aligned}$$

[37] We particularly thank the research professionals from UChicago's Macro-Finance Research Program - Bin Cheng, Pengyu Chen, and Zhaoyang Xu - and our coauthor Hong Zhang, who worked together to integrate the PETSc software (see Balay et al. (2025)) into our computational algorithm.

where x is our state variable vector, $V_x(x) = \frac{\partial V}{\partial x}(x)$ and $V_{xx}(x) = \frac{\partial^2 V}{\partial x \partial x'}(x)$ are used for notational simplicity. The solution is obtained by finding a $V(x)$ such that the HJB equality, first-order condition, and $V_t(x) = 0$ hold. Conditional on an initial guess, $V^0(x)$, we calculate the coefficients $\mathbb{A}, \mathbb{B}, \mathbb{C}$ and $\mathbb{D}$ based on the flow utility and the evolution of the state variables. We substitute the calculated coefficients, rearrange the terms, and apply a backward finite difference to $V_t^0(x)$ to obtain the final expression

$$\begin{aligned}\hat{V}(x) &= V^0(x) + [\mathbb{A}(x, V^0, V_x^0, V_{xx}^0)\hat{V}(x) + \mathbb{B}(x, V^0, V_x^0, V_{xx}^0)\hat{V}_x(x) \\ &\quad + \frac{1}{2} tr[\mathbb{C}(x, V^0, V_x^0, V_{xx}^0)'\hat{V}_{xx}(x)\mathbb{C}(x, V^0, V_x^0, V_{xx}^0)] + \mathbb{D}(x, V^0, V_x^0, V_{xx}^0)]\Delta t.\end{aligned}$$

We solve numerically for $\hat{V}(x)$ from this linear HJB system using the Preconditioned Bi-Conjugate Gradient Stabilized (BCGS) Krylov subspace (KSP) method from van der Vorst (1992), with ILU preconditioner.

The relative KSP residual tolerance statistic used to determine convergence of the solution for the linear system is given by

$$\frac{\left\|P^{-1}(AV_n - b)\right\|_2}{\left\|P^{-1}(AV_0 - b)\right\|_2} \leqslant \epsilon_{KSP},$$

where P^{-1} represents the ILU pre-conditioner and $AV = b$ represents the conditional linear system for a given value function V. The implementation of this step terminates when this statistic is less than ϵ_{KSP}.

The false-transient residual is the statistic used to determine convergence for the "outer loop" of the algorithm. This statistic is computed as

$$\left\| \frac{(V^{n+1} - V^n)}{\Delta\tau} \right\|_\infty \leqslant \epsilon_{outer}$$

where V^{n+1} is the next iteration's value function, V^n is the previous iteration's value function, and $\Delta\tau$ is the false transient step size. The implementation of this step terminates when this statistic is less than ϵ_{outer} or when the stated maximum iteration count is reached.

G.4 Finite-differences scheme

The finite difference scheme used in our computations is as follows. For calculating first- and second-order derivatives used to compute optimal controls, we use central differences at all interior points in the state space. At boundary points corresponding to the minimum state values we use forward differences, and

at the maximum boundary values we apply backward differences. The benefit of using central differences is that they provide greater accuracy in approximating derivatives than one-sided differences, while also leading to faster convergence in iterative numerical algorithms when the algorithm converges.

Central differences, however, can generate spurious oscillations. Therefore, for each of our drift terms $\mathbb{B}_n(x; V, V_x, V_{xx})V_{x_n}(x)$, where $x_n \in \{\hat{k}, y, \hat{r}\}$, we use an implicit upwind finite-difference scheme that selects forward or backward differences depending on the sign of $\mathbb{B}_n$. This is implemented in the matrix assembly step for the PetSc solver by switching which neighbor enters the stencil based on the sign of $\mathbb{B}_n$. This yields a monotone discretization of the first-order part of the HJB operator at interior grid points. The second derivatives are approximated with central differences at all interior points in the state space, forward differences at minimum boundary values, and backward differences at maximum boundary values. The controls used in policy evaluation are computed from the model's first-order conditions using the derivative approximations described above.

Finally, we note two additional features that improve the numerical stability of our algorithm. First, our model includes non-trivial volatility terms, which help disperse local truncation errors. Second, we use relatively small time steps for the false-transient terms in our algorithm, which help dampen spurious oscillations and feedback loops. Together with the implicit upwind finite-difference scheme, these features promote numerical stability. We assess numerical accuracy using the convergence diagnostics and grid-refinement results reported below.

G.5 Boundary Conditions

Boundary conditions in our algorithm come in two forms. First, the specification of the damage information event introduces an explicit "physical" or "economic" boundary value at $\bar{y}$ in the form of a value matching condition. The value matching condition is given by the robust (certainty equivalent) boundary aggregation of the $L-1$ post-damage curvature jump continuation values

$$V = -\xi_m \log\left[\frac{1}{L-1}\sum_{\ell=1}^{L-1}\exp\left(-\frac{V^\ell}{\xi_m}\right)\right],$$

which is achieved in the limit as the jump intensity goes to infinity. Numerically, this condition can be approximated by imposing a large (but finite) value of the damage information event arrival rate while maintaining the robust jump terms in the HJB equation. We use this approximation of the true value

matching condition for the pre-damage curvature jump settings based on the values of the post-jump value function at $\bar{y}$ across the $\hat{k}$ and $\hat{r}$ dimensions of the state space.

Along the remaining boundaries of the state space, there are no obvious "physical" or "economic" boundary conditions to impose. Therefore, with the PDE operator discretized up to the boundary, we use the "natural" one-sided finite difference closure implied by our upwind finite-difference scheme at all remaining boundary points of the state space. Specifically, we apply forward differences at boundary points corresponding to the minimum state values and backward differences while at the maximum boundary values for both first- and second-order derivatives. To test the potential impact of these boundary conditions on the achieved numerical solutions, we informally explore different grid sizes and bounds on the state space and find essentially no change in the outcomes for these sensitivity analyses.

G.6 Hyperparameters and numerical accuracy

We now provide additional details on the algorithm hyper-parameters and convergence criteria to enhance replicability. We also provide the outcomes for various numerical accuracy tests. Table 24 shows the state variable domain space, the number of grid points used within the state domain, and the discretization step size along the state and false transient domains. Table 25 provides the values used for the relaxation weight, the relative KSP residual tolerance, the false-transient residual tolerance, and the maximum iterations for each scenario. The "initial" and "warm" designations indicate the value-function initialization type. "Initial" denotes the first solve using the stage-specific initialization implemented by the solver: the default state-dependent guess for W^{ℓ} and the mean of the post-damage continuation values for V. "Warm" denotes initialization from an existing solution for the corresponding equation on the same grid. We note that the post-technology jump cases $W^{\ell,L}$ and V^{L} have analytic solutions which we use in our main results. Thus, we do not include these cases in Tables 24 and 25 as no false-transient step size, relaxation weight, KSP residual tolerance, outer stopping tolerance, or maximum iterations are needed in those settings. The state domain, grid points, and grid spacing used to compute the numerical solutions for those settings match the pre-technology jump counterparts ($W^{\ell,L}$ uses the same values as W^{ℓ} and V^{L} uses the same values as V). Importantly, the reported values reflect the choice of a broad state space domain with sufficient discretization to derive accurate numerical solutions. Our next set of results confirms that we do in fact achieve a high level of accuracy from our numerical solutions.

Value Function	State domain $[\hat{k}] \times [Y/z] \times [\hat{r}]$	Grid points $n_{\hat{k}} \times n_Y/n_z \times n_{\hat{r}}$	Grid spacing $(\Delta\hat{k}, \Delta Y/\Delta z, \Delta\hat{r})$	False-transient step ($\Delta\tau$)
W^ℓ (initial)	$[4,9]\times[0,5]\times[1,8]$	$51\times201\times71$	$(0.1, 0.025, 0.1)$	0.1
W^ℓ (warm start)	$[4,9]\times[0,5]\times[1,8]$	$51\times201\times71$	$(0.1, 0.025, 0.1)$	0.025
V (initial and warm)	$[4,9]\times[0,2.5]\times[1,8]$	$51\times101\times71$	$(0.1, 0.025, 0.1)$	0.05

Table 24: Numerical discretization criteria. Reported values are for the full-model baseline-parameter solutions. The middle entries for columns 2-4 are denoted by the labels Y/z, n_Y/n_z, and $\Delta Y/\Delta z$ to reflect that the post-damage state vector is $(\log K, z, \log R)$, while the pre-damage state vector is $(\log K, Y, \log R)$.

Value Function	Relaxation weight (ε)	KSP residual tolerance (ϵ_{KSP})	Outer stopping tolerance (ϵ_{outer})	Maximum iterations
W^ℓ (initial)	0.1	10^{-13}	10^{-7}	10,000
W^ℓ (warm start)	0.025	10^{-13}	10^{-7}	100,000
V (initial and warm)	0.05	10^{-13}	10^{-7}	50,000

Table 25: Numerical convergence criteria. Reported values are for the full-model baseline-parameter solutions. Post- and pre-damage state vectors are given by $(\log K, z, \log R)$ and $(\log K, Y, \log R)$, respectively.

Table 26 shows the number of equations corresponding to each type of value function solved, the maximum absolute KSP residual, the maximum false-transient residual, and the maximum reported outer iterations. Note that the maximum absolute KSP residual provided by PETSc that we report can exceed or be significantly smaller than the relative KSP residual tolerance that is used as the stopping criteria statistic. The central takeaway from these reported values is that each of these outcomes demonstrates a high level of accuracy achieved in deriving our solutions, emphasizing reliability of the numerical solutions we have computed.[38]

[38]For completeness, we note that the maximum analytic HJB residuals for $W^{\ell,L}$ and V^L are both 2.19×10^{-17}.

Value Function	Number of equations	Max absolute KSP residual norm	Max false-transient residual	Max outer iterations
W^{ℓ}	60	2.78×10^{-10}	9.9994×10^{-8}	16,027
V	3	2.75×10^{-10}	9.9922×10^{-8}	5,410

Table 26: Convergence diagnostics and residuals. Reported values are for the full-model baseline-parameter solutions, including the uncertainty-neutral, less-aversion, and more-aversion specifications. The post-damage curvature jump row contains 20 damage curvature specifications for each aversion specification. The pre-damage curvature jump row contains one equation per aversion specification and incorporates all 20 post-damage curvature jump continuation values. Post- and pre-damage state vectors are given by $(\log K, z, \log R)$ and $(\log K, Y, \log R)$, respectively.

Tables 27 - 28 provide details for additional accuracy results via a grid-refinement check. For these tests, we compare social valuation, optimal control, and value function outcomes across state space grid resolutions. Table 27 summarizes the two different grid specifications used in this comparison. The grid used in the main body of the paper is $(n_{\hat{k}}, n_z, n_{\hat{r}}) = (51, 201, 71)$ for the post-damage curvature jump state and $(n_{\hat{k}}, n_Y, n_{\hat{r}}) = (51, 101, 71)$ for the pre-damage curvature jump state. We compare these results, referred to as the fine grid resolution case, with those for a coarse grid resolution with corresponding grids of $(n_{\hat{k}}, n_z, n_{\hat{r}}) = (26, 101, 36)$ for the post-damage curvature jump state and $(n_{\hat{k}}, n_Y, n_{\hat{r}}) = (26, 51, 36)$ for the pre-damage curvature jump state.

	post-damage jump		pre-damage jump	
Grid specification	$n_{\hat{k}} \times n_Y / n_z \times n_{\hat{r}}$	grid scaling	$n_{\hat{k}} \times n_Y / n_z \times n_{\hat{r}}$	grid scaling
coarse resolution	$26 \times 101 \times 36$	$1.00\times$	$26 \times 51 \times 36$	$1.00\times$
fine resolution	$51 \times 201 \times 71$	$7.70\times$	$51 \times 101 \times 71$	$7.66\times$

Table 27: State variable grids used in the refinement exercise. Grid sizes are reported as the total number of nodes used along each dimension when solving the HJB equation. Grid scaling is the number of total nodes relative to the original grid size.

Table 28 provides the value function, social valuation, and optimal control results for the two grid specifications. The magnitudes for differences in the value function outcomes are very modest, ranging between .02% and .14%. The differences in social valuations and controls across grid resolutions are also relatively small, though somewhat larger than the value function outcomes. For the SVRD and SCGW, the differences range between .15% and 3.01% in absolute value. For R&D investment and emissions, the

differences range between .02% and 6.13% in absolute value. The slightly larger differences reflect the impact of needing numerical derivatives to compute the social valuation and optimal control outcomes from the value function solutions. These results show that the grid choice for our model results provides a high level of accuracy for all of the main outcomes derived from our numerical solutions.

ξ_m	$V(X_0)$	SVRD	R&D investment	SCGW	emissions
∞	4.87	3.90	.0066	60.6	9.26
.1	4.78	6.48	.0183	119.7	8.85
.05	4.64	10.05	.0438	254.6	8.17

(a) Coarse grid resolution.

ξ_m	$V(X_0)$	SVRD	R&D investment	SCGW	emissions
∞	4.87	3.91	.0066	60.5	9.26
0.1	4.79	6.52	.0185	118.9	8.85
0.05	4.65	10.36	.0465	258.9	8.15

(b) Fine grid resolution.

Table 28: Grid-refinement comparisons of key model outcomes. Panel 28a shows the results using the coarse grid resolution. Panel 28b shows the results using the fine grid resolution. Column 1 identifies the uncertainty aversion cases considered: baseline ($\xi_m = \infty$), less aversion ($\xi_m = .1$), and more aversion ($\xi_m = .05$). Columns 2-6 show the value function, the SVRD, R&D investment-to-output ratio, the SCGW, and emissions, respectively, evaluated at the initial state variable vector X_0.

Our final accuracy check comes from evaluating the two methods used to compute our marginal social valuations. The first method is a "direct" computation using the finite-difference derivative of the HJB value function at the initial state, computed on the HJB solution grid. The second method is an "indirect" calculation based on the Monte Carlo evaluation of the marginal valuation formula given by

$$\frac{\partial V}{\partial x}(X_0) \cdot \Lambda_0 \;=\; \widetilde{E}\Big[\int_0^\infty \mathrm{Dis}_t\left(\Lambda_t \cdot \mathrm{Scf}_t\right) dt \;\Big|\; X_0, \Lambda_0\Big].$$

Summing over the flow i, ii and iii terms computed in this manner, and normalizing by marginal utility, we get an expression in the same units as our direct method. While the two methods are identical mathematically, the second method carries additional numerical errors from

i) numerical derivative approximation (e.g., finite differences and grid discretization)

ii) simulation-based approximation (e.g., time discretization and finite time boundary)

iii) Monte Carlo sampling (e.g., sampling noise from finite stochastic pathways)

We first note that the Monte Carlo standard errors, computed from the cross-sectional dispersion of the simulation pathways, are sufficiently small (0.002–0.006) as to be inconsequential for the overall numerical error. To evaluate the "numerical derivatives" errors, we again use grid refinement to compare outcomes across specifications. If the gap between the indirect and direct methods is driven by derivative approximation, enhanced state variable grid resolution refinement should systematically shrink the difference between the results from the two specifications. We therefore compute the social valuations based on the indirect and direct methods using the two grid resolutions outlined in Table 27 for each specification. As the fine grid solution has 2.0× the number of nodes along each dimension as the coarse grid solution, the first-order discretization errors should roughly halve between the coarse- and fine-grid solutions.

ξ_m	grid	indirect SVRD	direct SVRD	absolute gap	relative gap
∞	coarse	4.39	4.27	0.12	2.87%
.1	coarse	6.82	6.41	0.40	6.28%
.05	coarse	9.67	8.70	0.97	11.13%
∞	fine	4.32	4.26	0.05	1.21%
.1	fine	6.64	6.44	0.19	2.99%
.05	fine	9.40	8.90	0.50	5.63%

(a) Technology jump only model.

ξ_m	grid	indirect SVRD	direct SVRD	absolute gap	relative gap
∞	coarse	4.01	3.90	0.11	2.79%
0.1	coarse	6.80	6.48	0.32	4.87%
0.05	coarse	10.83	10.05	0.77	7.70%
∞	fine	3.95	3.91	0.04	1.11%
0.1	fine	6.66	6.52	0.14	2.19%
0.05	fine	10.78	10.36	0.42	4.07%

(b) Full baseline model.

Table 29: Grid-refinement comparisons of indirect and direct SVRD calculations across various model specifications. Panel 29a shows the results for the technology jump only model. Panel 29b shows the results for the full baseline model. Column 1 identifies the uncertainty aversion cases considered: baseline ($\xi_m = \infty$), less aversion ($\xi_m = .1$), and more aversion ($\xi_m = .05$). Column 2 identifies the state space grid cases considered. Columns 3 and 4 give the indirect and direct SVRD values, respectively. Columns 5 and 6 give the absolute and relative gap in valuations, respectively.

Table 29 shows the results of our grid-refinement audit of the numerical discrepancies between the "direct" and "indirect" social valuation calculations for the technology jump only and full model specifications. The top panel, 29a, shows the outcomes for the technology jump only model. The bottom panel, 29b, shows the outcomes for the full model. Note that an accurate characterization of the grid refine-

ment effect across methods requires an additional adjustment for the full model setting. Because the full model includes the damage information event, simulations are terminated at the sooner of the temperature anomaly threshold $\bar{y}$ or year 65, whichever comes first. This differs from the technology only jump model, for which we terminate the simulations at 120 years. This introduces a small negative terminal value error for the indirect method in the full model. We adjust for this error in the full model calculation under the indirect method by adding an interpolated value gradient term at the final simulation point. While overall this additional error is quite small, without this adjustment the grid refinement can appear to make the error gap worse in the baseline uncertainty aversion case because of the offsetting signs of the error term components. The grid refinement comparison across methods highlights the following main outcomes:

- Halving the grid spacing roughly halves the gap, consistent with first-order discretization error.
- The valuation gap between methods grows as ξ_m falls, coinciding with evaluating enhanced nonlinearities from uncertainty aversion on the same grid generating larger approximation errors.
- The direct derivative valuations move by at most 3.0% across grids, highlighting that most of the improvement in valuation discrepancy is coming from the indirect method errors.

While grid-refinement generates clear improvement in the numerical discrepancies across methods, we investigate the remaining errors to better understand the sources of numerical inaccuracies. We calculate implied derivative approximation errors from the "indirect" method by computing and aggregating up the point-wise residual derivative errors along each step of the simulation-based marginal social valuation construction. Removing this residual derivative approximation error captures nearly all of the remaining error. Therefore, the remaining errors after accounting for numerical derivative approximation (-0.16% to -0.30%) must pertain to sources such as monthly Euler step for simulations, interpolation of outcomes to off-grid points, the evaluation of the terminal term at the last monthly node rather than at the exact boundary-crossing time, and right-endpoint accumulation of the jump-augmented discounting term.

We characterize the nature and magnitude of these different error components in Table 30. Altogether, our characterization of the numerical errors highlights two key results. First, our numerical results appear to have a high level of accuracy. Second, neither the ordering of the social valuations across ξ_m nor the relative importance of the three flows in our decompositions is affected by remaining numerical errors.

error source	nature	error size	behavior
numerical derivative approximation	discretization error in the FK representation	$\approx 1\%$ ($\xi_m = \infty$); $\approx 2\%$–5% (finite ξ_m)	$\sim$ halves when grid spacing halves; grows as ξ_m falls
time discretization, boundary approximation	residual numerical errors	$\lesssim 0.3\%$	small but statistically distinguishable from zero
Monte Carlo sampling	random error	≈ 0.05–0.07% (s.e.)	falls as $1/\sqrt{N}$; irrelevant

Table 30: Decomposition of the direct-vs-indirect discrepancy for the full model with fine grid resolution.

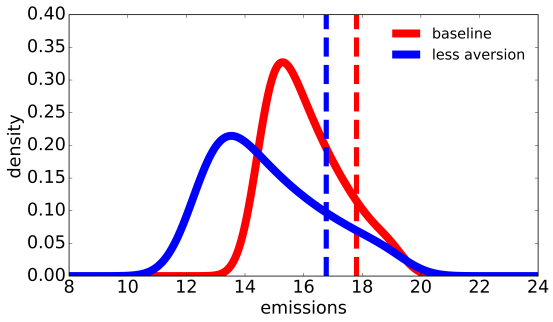
Appendix H Shapley value decomposition

Table 31 provides the full set of results from a Shapley value decomposition for each of different uncertainty channels in our model framework. This additive decomposition, which accounts for the cross-dependencies of multiple channels, validates the results found by our first-order index in Tables 1 and 2. Specifically, the technology channel is the dominant contributor to the uncertainty impact for the SVRD: 85% for $\xi_m = .1$ and 81% for $\xi_m = .05$. Moreover, the technology channel is a similarly relevant contributor to the uncertainty impact for the SCGW: 83% for $\xi_m = .1$ and 76% for $\xi_m = .05$.

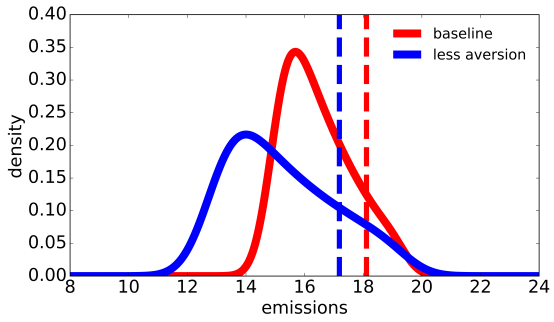
	SVRD		SCGW	
Uncertainty channel	$\xi_m = .1$	$\xi_m = .05$	$\xi_m = .1$	$\xi_m = .05$
climate	0.19 (7%)	0.52 (8%)	3.8 (6%)	15.7 (8%)
damage	0.31 (12%)	1.00 (15%)	7.9 (14%)	38.7 (20%)
productivity	-0.11 (-4%)	-0.27 (-4%)	-1.6 (-3%)	-6.5 (-3%)
technology	2.22 (85%)	5.19 (81%)	48.2 (83%)	150.4 (76%)

Table 31: Shapley decompositions of the changes in the two social valuations, SVRD and SCGW, between the baseline and all-channel specifications. Entries are channel-specific Shapley contributions, and the numbers in the parentheses report percentages of the corresponding all-minus-baseline change. Technology combines the R&D Brownian and technology-jump channels.

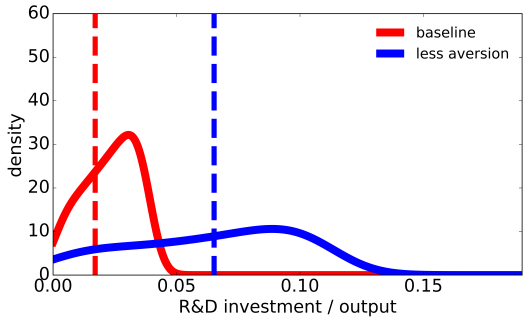
Appendix I Pre- and post-damage information event comparison



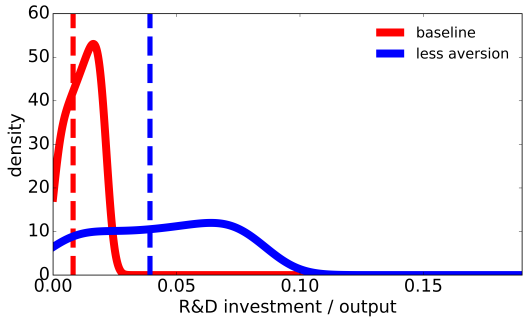
(a) Emissions, baseline median state



(b) Emissions, less aversion median state



(c) R&D investment, baseline median state



(d) R&D investment, less aversion median state

Figure 12: Robustly optimal R&D investment and emissions before and after a damage information event for two different state variable scenarios. Each case uses the median capital and knowledge stock values for no-jump simulations at 2°C. Panels (a) and (c) use the median state values under the baseline solution. Panels (b) and (d) use the median state values under the less-aversion solution. The median capital and knowledge stock values for the baseline case are $k = 1505$ and $r = 35$. The median capital and knowledge stock values for the less-aversion case are $k = 1508$ and $r = 56$. The post-damage information event policies are evaluated at $z = 2.5$. Red curves depict the baseline responses and blue curves depict the less-aversion responses. The solid curves are Gaussian kernel densities of the 20 post-damage curvature jump outcomes, with equal weights and Scott's bandwidth multiplied by 0.8. The dashed vertical lines give the outcomes before the damage curvature jump. Emissions are reported in gigatons of carbon (GtC). R&D investment is reported as a ratio relative to output abstracting from the energy input (αK). For the baseline median state initialization, pre-jump R&D investment is .017 and .065 and emissions are 17.80 GtC and 16.77 GtC under the baseline and less aversion policy responses, respectively. For the less aversion median state initialization, R&D investment is .0081 and .0392 and emissions are 18.11 GtC and 17.19 GtC under the baseline and less aversion policy responses, respectively.

	emissions (GtC)		R&D investment / output	
ξ_m	pre-event	post-event	pre-event	post-event
baseline median state ($K = 1505$, $R = 35$, $\hat{y} = 2^\circ$C)				
∞	17.8	15.8	.017	.024
		(14.6, 18.9)		(.002, .038)
.1	16.8	14.4	.065	.069
		(12.45, 18.9)		(.002, .110)
less aversion median state ($K = 1508, R = 56, \hat{y} = 2^\circ$C)				
∞	18.1	16.2	.008	.012
		(15.0, 19.0)		(.001, .021)
.1	17.2	14.9	.039	.045
		(12.9, 19.0)		(.001, .083)

Table 32: Robustly optimal emissions and R&D investment choices before and after a damage information event for two different state variable scenarios. Each case uses the median capital and knowledge stock values for no-jump simulations at 2°C. The top panel uses the median state values under the baseline solution, and the bottom panel state values under the less-aversion solution. The post-damage information event policies are evaluated at $z = 2.5$. Post-event reported values are medians across the 20 post-event, pre-technology policy values. Parentheses below each median report the minimum and maximum of the 20 equally weighted post-event, pre-technology policy values. Emissions are reported in gigatons of carbon (GtC). R&D investment is reported as a ratio relative to output abstracting from the energy input (αK).

Appendix J Two counteracting forces for the single technology jump model

In this appendix, we will produce characterizations of two counter-acting forces that impact the marginal valuation of the knowledge stock pertinent for a new discovery made more likely by R&D investments.

J.1 Basic input

Suppose there is a single jump indexed by L, in particular say the technology jump. Recall the date t flow contribution to the corresponding marginal valuation:

$$\Lambda_t \frac{\partial \mathcal{J}^L}{\partial x}(X_t)\xi_m \left[1 - \exp\left(-\frac{1}{\xi_m}\left[V^L(X_t) - V(X_t)\right]\right)\right],$$

and the jump probability adjustment to discounting:

$$\exp\left[-\int_0^t G_u^{L*}\mathcal{J}^L(X_u)du\right]. \tag{24}$$

Robustness concerns are captured by a reduction in the jump intensities, G_u^{L*}, delaying the prospective probability of a technology discovery. Notice that this reduction increases the discount contribution as given by (24).

For the purposes of this appendix, we do not incorporate this term in the discount adjustment and instead study:

$$\exp\left[-\int_0^t G_u^{L*}\mathcal{J}^L(X_u)du\right]\Lambda_t \frac{\partial \mathcal{J}^L}{\partial x}(X_t)\xi_m \left[1 - \exp\left(-\frac{1}{\xi_m}\left[V^L(X_t) - V(X_t)\right]\right)\right],$$

For pedagogical simplicity, we hold fixed the controls as functions of the states so both the state process, X, and the stochastic response process, Λ, are invariant with respect to changes in the penalty parameter, ξ_m. As with the use of endowment economies in asset pricing, this allows us to focus on the impacts on the marginal valuation. Given this invariance, in what follows we drop the contribution from Λ_t.

J.2 Force 1: a push for reducing the marginal valuation

To obtain an *ex post* measure of the jump time density, compute the negative of the time derivative of (24), which is given by:

$$G_t^{L*}\mathcal{J}^L(X_t)\exp\left[-\int_0^t G_u^{L*}\mathcal{J}^L(X_u)du\right]. \tag{25}$$

See Appendix D. Provided that

$$\int_0^t G_u^{L*}\mathcal{J}^L(X_u)du$$

diverges to ∞, the jump time density (25) will integrate to one over the time dimension reflecting when a jump might happen.

Since the G_t^{L*} adjustments to this density push the event times further into the future, this delay makes the R&D investments less attractive, capturing a force for *reducing* the R&D investments in the presence of subjective discounting (via δ).

J.3 Force 2: a push for enhancing the marginal valuation

Since the jump time *ex post* density given by (25) absorbs the intensity $G_t^{L*}\mathcal{J}^L(X_t)$, we need to offset these terms in what remains. In so doing, we use the following relations:

$$\frac{\partial\mathcal{J}^L}{\partial x}(X_t)=\mathcal{J}^L(X_t)\frac{\partial\log\mathcal{J}^L}{\partial x}(X_t), \tag{26}$$

$$\mathcal{J}^L(X_t)\xi_m\left[1-\exp\left(-\frac{1}{\xi_m}\left[V^L(X_t)-V(X_t)\right]\right)\right]\geqslant G_t^{L*}\mathcal{J}^L(X_t)\left[V^L(X_t)-V(X_t)\right]. \tag{27}$$

Inequality (27) follows by inspecting the objective of the Problem 3.3 used to make a robust adjustment to a jump intensity. The left-hand side of (27) includes the entropy penalty contribution, in contrast to the right-hand side.[39] Provided that

$$\frac{\partial\log\mathcal{J}^L}{\partial x}\geqslant 0,$$

[39] In fact this entropy term would be absent in a constrained counterpart to our robust adjustment.

the following relationship holds:

$$\begin{aligned}
&G_t^{L*}\mathcal{J}^L(X_t)\frac{\partial \log \mathcal{J}^L}{\partial x}(X_t)\left[V^L(X_t)-V(X_t)\right] \\
&\leqslant \frac{\partial \log \mathcal{J}^L}{\partial x}(X_t)\mathcal{J}^L(X_t)\xi_m\left[1-\exp\left(-\frac{1}{\xi_m}\left[V^L(X_t)-V(X_t)\right]\right)\right] \\
&= \frac{\partial \mathcal{J}^L}{\partial x}(X_t)\xi_m\left[1-\exp\left(-\frac{1}{\xi_m}\left[V^L(X_t)-V(X_t)\right]\right)\right],
\end{aligned}$$

where the last equality follows from relation (26). Finally, multiply the first and last terms by expression (24) to obtain:

$$\begin{aligned}
&\left(\exp\left[-\int_0^t G_u^{L*}\mathcal{J}^L(X_u)du\right]G_t^{L*}\mathcal{J}^L(X_t)\right)\left[\frac{\partial \log \mathcal{J}^L}{\partial x}(X_t)\left[V^L(X_t)-V(X_t)\right]\right] \\
&\leqslant \exp\left[-\int_0^t G_u^{L*}\mathcal{J}^L(X_u)du\right]\frac{\partial \mathcal{J}^L}{\partial x}(X_t)\xi_m\left[1-\exp\left(-\frac{1}{\xi_m}\left[V^L(X_t)-V(X_t)\right]\right)\right].
\end{aligned} \tag{28}$$

Consider the expression on the left-hand side of (28), which is given as the product of two terms. The first term is the *ex post* jump density given previously by (25). The second term:

$$\frac{\partial \log \mathcal{J}^L}{\partial x}(X_t)\left[V^L(X_t)-V(X_t)\right]$$

captures the second counteracting force of uncertainty aversion. A decrease in ξ_m increases $V^L(X_t)-V(X_t)$, working primarily through a reduction in $V(X_t)$.

A complementary derivation gives:

$$\mathcal{J}^L(X_t)\xi_m\left[1-\exp\left(-\frac{1}{\xi_m}\left[V^L(X_t)-V(X_t)\right]\right)\right]=\mathcal{J}^L(X_t)G_t^{L*}\xi_m\left[\exp\left(\frac{1}{\xi_m}\left[V^L(X_t)-V(X_t)\right]\right)-1\right]$$

since

$$G_t^{L*}=\exp\left(-\frac{1}{\xi_m}\left[V^L(X_t)-V(X_t)\right]\right).$$

Now multiply both sides by $\frac{\partial \log \mathcal{J}^L}{\partial x}(X_t)$ to obtain:

$$\begin{aligned}
=&\frac{\partial \mathcal{J}^L}{\partial x}(X_t)\xi_m\left[1-\exp\left(-\frac{1}{\xi_m}\left[V^L(X_t)-V(X_t)\right]\right)\right] \\
&=\mathcal{J}^L(X_t)G_t^{L*}\frac{\partial \log \mathcal{J}^L}{\partial x}(X_t)\xi_m\left[\exp\left(\frac{1}{\xi_m}\left[V^L(X_t)-V(X_t)\right]\right)-1\right]
\end{aligned}$$

where we used (26) in the first expression. Next multiply the second expression by (24) to obtain

$$\left(\exp\left[-\int_0^t G_u^{L*}\mathcal{J}^L(X_u)du\right]G_t^{L*}\mathcal{J}^L(X_t)\right)\left[\frac{\partial\log\mathcal{J}^L}{\partial x}(X_t)\xi_m\left[\exp\left(\frac{1}{\xi_m}\left[V^L(X_t)-V(X_t)\right]\right)-1\right]\right].$$

The expression

$$\frac{\partial\log\mathcal{J}^L}{\partial x}(X_t)\xi_m\left[\exp\left(\frac{1}{\xi_m}\left[V^L(X_t)-V(X_t)\right]\right)-1\right]$$

is increasing and convex in $V^L(X_t)-V(X_t)$. The convexity implies our previous inequality as follows by a gradient inequality for a convex function:

$$\xi_m\left[\exp\left(\frac{1}{\xi_m}\left[V^L(X_t)-V(X_t)\right]\right)-1\right]\geqslant V^L(X_t)-V(X_t).$$

Appendix K Two-step and partial success technology discovery

To complement the reported SVRD and R&D investment outcomes for the different two-step technology discovery setting reported in Table 7 of Section 9.2, we report quantiles of economic capital investment relative to output and R&D investment relative to output after the first jump for the three two-jump model specifications. The simulations used to construct these quantiles are done under the baseline probabilities with the respective control laws imposed.

The outcomes are shown in Table 33, which provides the .1, .5, and .9 quantiles of the two different investment types after the first jump for the three different two-jump model specifications and across the three model uncertainty aversion scenarios. We note that there is considerable heterogeneity, particularly for specification iii. However, R&D investment ratios even at the .9 quantile are smaller after the first jump for specifications i) and ii). In specification iii), where the initial advancement of the knowledge stock is only relevant for the first technology jump, we see R&D investment at the .5 and .9 quantiles is higher than the pre-jump value for $\xi_m=\infty$, whereas only the .9 quantile exceeds the pre-jump value for $\xi_m=.1$ and $\xi_m=.05$. For economic capital investment, the post-jump values decrease as uncertainty aversion increases, compensating for the increasing R&D investment. The heterogeneity patterns, however, are similar to those for R&D investment.

Panel A: Capital investment / total output

	.1 quantile			.5 quantile			.9 quantile		
ξ_m	(i)	(ii)	(iii)	(i)	(ii)	(iii)	(i)	(ii)	(iii)
∞	.7647	.7470	.7339	.7757	.7676	.7701	.7778	.7726	.7772
.1	.7632	.7450	.7292	.7730	.7601	.7646	.7755	.7656	.7742
.05	.7596	.7343	.7253	.7684	.7472	.7575	.7713	.7547	.7692

Panel B: R&D investment / total output

	.1 quantile			.5 quantile			.9 quantile		
ξ_m	(i)	(ii)	(iii)	(i)	(ii)	(iii)	(i)	(ii)	(iii)
∞	.0027	.0066	.0033	.0040	.0087	.0098	.0067	.0114	.0396
.1	.0046	.0122	.0059	.0062	.0151	.0150	.0094	.0194	.0473
.05	.0082	.0214	.0105	.0100	.0264	.0220	.0140	.0329	.0524

Table 33: Quantiles for investments after the first jump realization for different values of ξ_m for the three two-step specifications. R&D investment is reported as a ratio relative to total output.

To further evaluate the sensitivity of the competing forces underlying the proactive R&D response to uncertainty aversion, we examine another variant of the simplified model without a damage information event. Specifically, we look at the impact of the one-step technology advancement being only a partial success. In this setting, we assume that the terminal abatement technology parameter is given by $\phi_0^{(post)} = 0.5 \times \phi_0$, where ϕ_0 is the initial abatement technology parameter value used in the baseline model.

ξ_m	SVRD	SCGW	R&D investment	capital investment	emissions
∞	2.71	279.4	.0032	.7216	8.14
.1	3.29	286.6	.0047	.7193	8.11
.05	4.07	297.8	.0072	.7156	8.07

Table 34: Social valuations, initial investments, and emissions for the partial technology success specification. R&D investment is reported as a ratio relative to output. Outcomes are shown for the three different values of the aversion parameter ξ_m in the respective rows.

Table 34 shows the social valuations and the robustly optimal control outcomes for this partial technological change success setting. A few key outcomes emerge from this setting. First, the quantitative magnitude of the SVRD is notably reduced when the technology jump is only a partial improvement in the abatement technology parameter. The corresponding R&D investment values are lower as a result. Second, the quantitative magnitude of the SCGW is substantially increased without the technology discovery jump leading to an absolute breakthrough outcome. As a result, emissions are also initially lower in this setting. Third, we see that capital investment also declines, consistent with investment tradeoff in the allocation of output. Finally, the qualitative response to increased uncertainty aversion from the main model setting remains. As ξ_m decreases, R&D investment is amplified. Going from the baseline setting to the more aversion setting, the R&D investment relative to total output more than doubles. While emissions decrease as uncertainty aversion increases, the sensitivity to uncertainty remains quite modest, just as in our main specification. These results suggest that the proactive R&D response to uncertainty can persist even without an absolute-breakthrough in abatement technology.

Appendix L Comparison to a real options formulation

As we noted in Section 10, there are similarities, but also differences, in our R&D investment problem and a real options problem.

As a reminder, in standard option pricing, payoff volatility is valued because the investor can exercise the option when the payoff is sufficiently positive. Increases in volatility enhance the right tail of the payoff distribution. The experiment that we explore, however, changes uncertainty aversion and does not increase the right tail probability. Given the ability to hedge, there is no direct role for risk aversion in the familiar options pricing formula.

Since our social planner does not have the ability to replicate the uncertainty of the R&D investment, we find that the pedagogically revealing real options analysis of Miao and Wang (2007) provides interesting benchmarks for comparison. Two of the specifications they consider essentially solve single-agent decision problems, and standard replication arguments do not apply. In both cases, there is uncertainty in the payoff dynamics. While they explore risk aversion with exponential utility, we obtain analogous conclusions with very similar supporting computations starting with linear utility and incorporating robustness concerns. We use their functional forms for pedagogical simplicity.

The decision-maker decides when to exercise the option with a known cost of doing so. Other than access to this stochastic payoff, the decision-maker can invest in a riskless technology. In the first specification, exercising the option provides a discrete payoff. In the second specification, exercising the option entitles the decision maker with access to a stochastic flow. To support our investigation, we make the two payoffs comparable by scaling the flow so that the present values of the discrete payoff and flow payoff processes are the same in the absence of uncertainty about the payoff dynamics. In the case of the discrete payoff, the uncertainty concerns are resolved when the option is exercised, in contrast to the case with a flow payoff as the uncertainty concerns about the transient dynamic persist.

As we will now show, uncertainty has obverse impacts for the two payoffs. This aversion makes the investor more bullish on the investment for the discrete payoff specification in the sense that there is a lower investment threshold. The outcome gets reversed with the stochastic flow payoff and the investor is less bullish on the investment and exercises the option at a higher threshold with enhanced aversion.

Consider a robust counterpart to Model I in Miao and Wang (2007). We let the utility be linear in consumption and the income process be:

$$dY_t = \nu dt + \sigma dW_t.$$

For purposes of illustration, we set $\nu = .1$ and $\sigma = .3$.[40] When the option is exercised, we assume a payoff equal to the discounted present value (with a constant discount rate δ) of current and future Y_t minus a prespecified cost I. In our illustration, we set $I = 10$. The risk-free rate and subjective rate of discount are the same, which we set as $\delta = .02$. We refer to this as the *discrete-payoff model*.

The post exercise value function satisfies:

$$\max_c -\delta V^0 + \delta c + V_w^0(\delta w - c) = 0,$$

which can be solved by setting $V^0(w) = \delta w$. The value function, V, prior to exercising the option solves the HJB equation:

$$\max_c \min_h -\delta V + \delta c + V_w(\delta w - c) + V_y \nu + \sigma h V_y + \frac{\xi_m}{2} h^2 + \frac{\sigma^2}{2} V_{yy}.$$

[40] Miao and Wang (2007) use these numbers in the illustration they report in their Figure 4.

Guess that $V(w,y) = \delta[w + G(y)]$. Then

$$0 = -\delta G + G_y \nu + \frac{\sigma^2}{2} G_{yy} - \frac{\sigma^2 \delta}{2\xi_m} (G_y)^2 \tag{29}$$

where the minimizing h solves:

$$h^* = -\frac{\sigma}{\xi_m} V_y = -\frac{\sigma\delta}{\xi_m} G_y. \tag{30}$$

This equation for G essentially matches Miao and Wang (2007)'s G function in their Proposition 1. The discounted expected value of the option payoff with income realization, y, is

$$\frac{1}{\delta} y + \frac{\nu}{\delta^2}$$

obtained by solving the Feynman-Kac equation:

$$0 = -\delta F + y + \nu F_y + \frac{\sigma^2}{2} F_{yy}.$$

Thus, value matching gives:

$$\delta G(\bar{y}) = \bar{y} + \frac{\nu}{\delta} - \delta I,$$

and smooth pasting gives

$$\delta G_y(\bar{y}) = 1.$$

We next change the payoff on the investment by letting it be continuous, given by the Y process. We refer to this as the *flow-payoff model.* The value function, V^o, after exercising satisfies the HJB equation:

$$\max_c \min_h -\delta V^o + \delta c + V^o_w(\delta w + y - c) + \nu V^o_y + \sigma h V^o_y + \frac{1}{2}\sigma^2 V^o_{yy} + \frac{\xi_m}{2} h^2 = 0.$$

We guess that $V^o(w,y) = \delta[w + F(y)]$. With this guess, we solve

$$0 = -\delta F + y + F_y \nu + \frac{\sigma^2}{2} F_{yy} - \frac{\sigma^2 \delta}{2\xi_m} (F_y)^2$$

where

$$h^* = -\frac{\sigma}{\xi_m} V^o_y = -\frac{\delta\sigma}{\xi_m} F_y.$$

The solution for F is:

$$F(y) = \frac{1}{\delta}\left[y + \frac{\nu}{\delta} - \frac{\sigma^2}{2\xi_m \delta}\right],$$

implying that $h^* = -\sigma/\xi_m$.

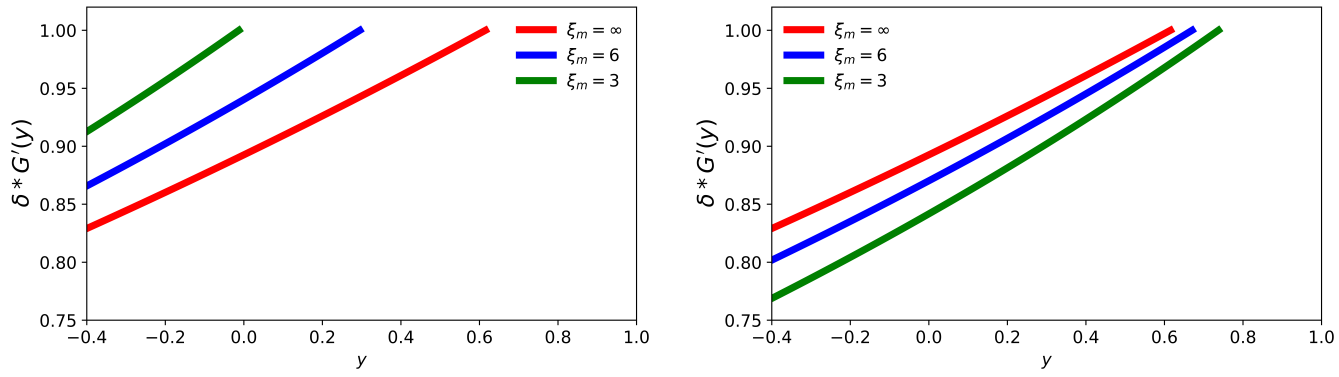
As with the discrete-payoff specification, the value function prior to the option being exercised solves (29) with $V(w, y) = \delta[w + G(y)]$. The minimizing h^* is given by formula (30). The value-matching condition is:

$$\delta G(\bar{y}) = \delta[F(\bar{y}) - I] = \bar{y} + \frac{\nu}{\delta} - \frac{\sigma^2}{2\xi_m \delta} - \delta I$$

and the smooth-pasting condition is:

$$\delta G_y(\bar{y}) = \delta F_y(\bar{y}) = 1.$$

Notice that the value matching condition is the same for both models in the limiting case in which $\xi_m \to \infty$. In both models, the exercise threshold $\bar{y}$ satisfies: $\delta G_y(\bar{y}) = 1$ even for $\xi_m < \infty$, but, of course, the G functions are different. Figure 13 plots the derivatives for each model specification and for alternative values of the robustness parameter, ξ_m. Recall that larger values of the penalty parameter imply less aversion. For the discrete-payoff model, the derivatives decrease with ξ_m, resulting in higher values of y when the option is exercised. For the stochastic flow-payoff model the derivatives increase with ξ_m, resulting in lower values of the threshold $\bar{y}$.



(A) discrete-payoff model

(B) flow-payoff model

Figure 13: Plots of δG_y for different values of ξ_m.

Consistent with this finding, we report in Figure 14 that uncertainty aversion has opposite implications for the threshold $\bar{y}$ for the two payoff specifications. This figure reports the investment thresholds for alternative values of the penalty parameter, ξ_m. The implied uncertainty aversion for low values of ξ_m

makes the investor more bullish on the investment for the discrete payoff specification in the sense that there is a lower investment threshold. The outcome gets reversed with the stochastic flow payoff: the investor is less bullish on the investment and exercises the option at a higher threshold with enhanced aversion. By design, the two curves converge when $\xi_m \to \infty$.[41]

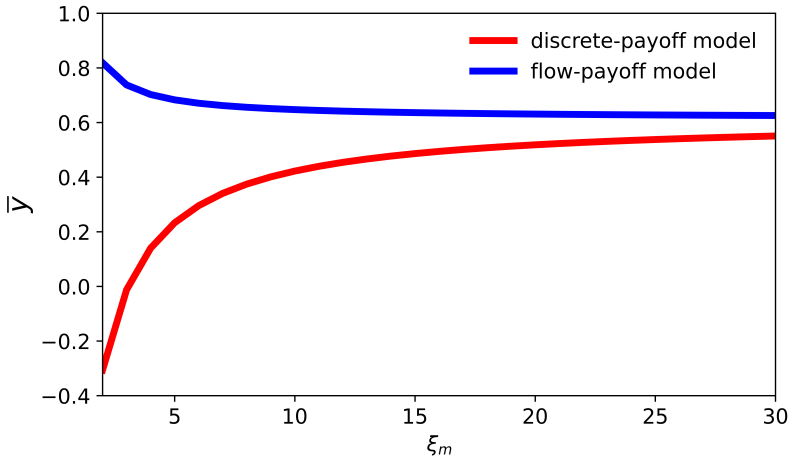


Figure 14: Solution of exercise threshold $\overline{y}$ for each model as functions of the robustness parameter, ξ_m.

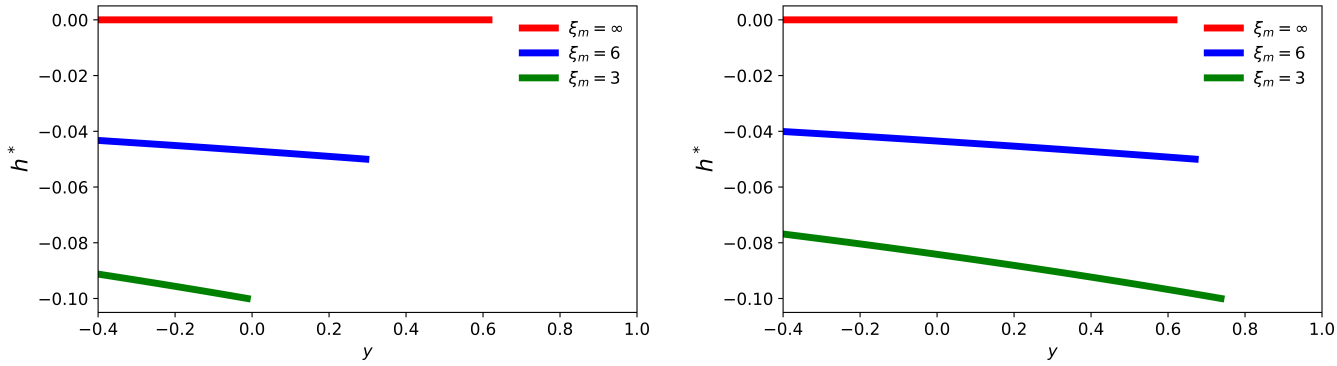
We plot the pre-exercise drift distortions h^* in Figure 15. The post-jump distortions are zero for the discrete-payoff specification, and, by the smooth pasting condition, equal to the distortion at the time the option is exercised for the flow-payoff specification. Recall that the h^*'s are expressed as drifts in a standard Brownian increment. They are not extreme for these illustrations, as the magnitude of the local mean distortion never exceeds ten percent of the standard deviation. The drift distortions prior to exercise are a bit larger in magnitude for the discrete-payoff specification.

Why the difference? In both cases, robustness concerns induce more cautious assessments of the payoff prospects. In the discrete-payoff specification, the uncertainty is resolved as soon as the payoff is exercised, increasing the valuation; while in the flow-payoff specification, the post exercise valuation is diminished under robustness concerns because of the continued uncertainty about future payoffs.

Although there are intriguing similarities between this comparison of real options in different environments and our findings, the investment opportunities are quite different. Rather than an exercise decision, our social planner can make investments that increase the likelihood of a discovery. Uncertainty aversion

[41] These results have counterparts to the risk aversion investigation reported in Miao and Wang (2007), albeit without the explicit present-value link between the two payoffs.

makes the planner skeptical about the payoff horizon in our analysis. These differences make our uncertainty aversion impact not monotone as displayed in Table 21 and in contrast to the real options findings reported in Figure 14.



(A) discrete-payoff model

(B) flow-payoff model

Figure 15: Plots of h^* for different values of ξ_m.

Online Appendices

Below we provide additional details and derivations for various results, which will be made available online in our online notebook:

`https://lphansen.github.io/mitigation_book`.

Appendix M Sensitivity analysis details: IES

ρ	SVRD	SCGW	R&D investment	capital investment	emissions
.75	5.62	81.95	.0137	.8781	9.1039
1	3.90	60.63	.0066	.7730	9.2645
1.33	2.76	47.55	.0033	.7022	9.3774

Table 35: Simulation results across IES values for $\xi_m = \infty$

ρ	SVRD	SCGW	R&D investment	capital investment	emissions
.75	8.83	161.16	.0338	.8550	8.6261
1	6.48	119.75	.0183	.7598	8.8517
1.33	4.67	89.66	.0095	.6958	9.0414

Table 36: Simulation results across IES values for $\xi_m = .1$

ρ	SVRD	SCGW	R&D investment	capital investment	emissions
.75	13.16	341.81	.0751	.8086	7.8291
1	10.05	254.64	.0438	.7300	8.1693
1.33	7.49	187.79	.0244	.6785	8.4711

Table 37: Simulation results across IES values for $\xi_m = .05$

Appendix N Sensitivity analysis details: (ϕ_0, ψ_0, χ)

ϕ_0	ψ_0	χ	emissions	R&D investment	technology-first probability	conditional tech-jump median year
.1	.07	.0014	8.51	.0017	.6411	19.30
.1	.07	.0009	8.26	.0027	.5134	22.19
.1	.07	.0007	8.12	.0032	.4434	23.62
.1	.11	.0014	8.63	.0020	.6847	19.71
.1	.11	.0009	8.40	.0031	.5657	22.82
.1	.11	.0007	8.27	.0037	.4984	24.40
.1	.15	.0014	8.72	.0021	.7206	19.88
.1	.15	.0009	8.51	.0032	.6097	23.09
.1	.15	.0007	8.39	.0038	.5452	24.73
.3	.07	.0014	9.09	.0034	.6413	18.72
.3	.07	.0009	8.89	.0058	.5201	21.46
.3	.07	.0007	8.78	.0072	.4541	22.85
.3	.11	.0014	9.20	.0034	.6864	19.06
.3	.11	.0009	9.03	.0056	.5737	21.94
.3	.11	.0007	8.93	.0070	.5103	23.40
.3	.15	.0014	9.28	.0032	.7224	19.19
.3	.15	.0009	9.13	.0052	.6173	22.12
.3	.15	.0007	9.04	.0065	.5566	23.61
.5	.07	.0014	9.31	.0042	.6403	18.49
.5	.07	.0009	9.15	.0072	.5206	21.14
.5	.07	.0007	9.05	.0091	.4555	22.49
.5	.11	.0014	9.41	.0040	.6861	18.82
.5	.11	.0009	9.27	.0067	.5748	21.60
.5	.11	.0007	9.19	.0084	.5122	23.01
.5	.15	.0014	9.48	.0037	.7223	18.95
.5	.15	.0009	9.35	.0061	.6184	21.77
.5	.15	.0007	9.28	.0076	.5586	23.20

Table 38: Simulation results under the full model for $\xi_m = \infty$, including the probability that the technology jump is the first jump to occur and the corresponding conditional median year of a technology jump occurring.

ϕ_0	ψ_0	χ	emissions	R&D investment	technology-first probability	conditional tech-jump median year
.1	.07	.0014	8.19	.0031	.5234	20.76
.1	.07	.0009	7.83	.0048	.3440	24.08
.1	.07	.0007	7.64	.0055	.2528	25.56
.1	.11	.0014	8.43	.0030	.6125	21.00
.1	.11	.0009	8.12	.0047	.4552	24.65
.1	.11	.0007	7.94	.0057	.3666	26.50
.1	.15	.0014	8.59	.0028	.6713	20.97
.1	.15	.0009	8.31	.0044	.5318	24.64
.1	.15	.0007	8.15	.0054	.4497	26.56
.3	.07	.0014	8.61	.0108	.4550	21.72
.3	.07	.0009	8.18	.0188	.2445	26.13
.3	.07	.0007	7.95	.0230	.1456	28.30
.3	.11	.0014	8.95	.0069	.5928	20.95
.3	.11	.0009	8.65	.0123	.4224	24.90
.3	.11	.0007	8.45	.0160	.3244	27.08
.3	.15	.0014	9.13	.0053	.6649	20.54
.3	.15	.0009	8.89	.0092	.5222	24.15
.3	.15	.0007	8.73	.0120	.4364	26.13
.5	.07	.0014	8.79	.0175	.4170	22.49
.5	.07	.0009	8.38	.0316	.1931	27.64
.5	.07	.0007	8.15	.0388	.0980	30.34
.5	.11	.0014	9.16	.0095	.5837	21.00
.5	.11	.0009	8.87	.0178	.4050	25.13
.5	.11	.0007	8.69	.0236	.3012	27.49
.5	.15	.0014	9.33	.0067	.6625	20.40
.5	.15	.0009	9.11	.0121	.5174	24.01
.5	.15	.0007	8.97	.0161	.4289	26.02

Table 39: Simulation results under the full model for $\xi_m = .1$, including the probability that the technology jump is the first jump to occur and the corresponding conditional median year of a technology jump occurring.

ϕ_0	ψ_0	χ	emissions	R&D investment	technology-first probability	conditional tech-jump median year
.1	.07	.0014	7.75	.0055	.3439	22.89
.1	.07	.0009	7.32	.0073	.1415	25.95
.1	.07	.0007	7.15	.0073	.0718	26.29
.1	.11	.0014	8.16	.0045	.5104	22.80
.1	.11	.0009	7.75	.0070	.3078	27.14
.1	.11	.0007	7.53	.0081	.2054	29.19
.1	.15	.0014	8.41	.0038	.6066	22.35
.1	.15	.0009	8.05	.0061	.4298	26.66
.1	.15	.0007	7.84	.0074	.3282	28.94
.3	.07	.0014	7.81	.0283	.1162	28.88
.3	.07	.0009	7.35	.0333	.0093	32.17
.3	.07	.0007	7.18	.0309	.0012	29.16
.3	.11	.0014	8.52	.0152	.4138	24.70
.3	.11	.0009	8.03	.0258	.1716	30.75
.3	.11	.0007	7.76	.0307	.0749	33.94
.3	.15	.0014	8.88	.0094	.5706	22.69
.3	.15	.0009	8.50	.0172	.3624	27.65
.3	.15	.0007	8.25	.0224	.2423	30.55
.5	.07	.0014	7.95	.0495	.0455	33.28
.5	.07	.0009	7.56	.0552	.0011	37.61
.5	.07	.0007	7.41	.0521	.0001	34.03
.5	.11	.0014	8.69	.0251	.3564	26.06
.5	.11	.0009	8.21	.0425	.1099	33.05
.5	.11	.0007	7.96	.0499	.0346	36.93
.5	.15	.0014	9.07	.0137	.5513	23.02
.5	.15	.0009	8.69	.0264	.3233	28.38
.5	.15	.0007	8.45	.0349	.1954	31.64

Table 40: Simulation results under the full model for $\xi_m = .05$, including the probability that the technology jump is the first jump to occur and the corresponding conditional median year of a technology jump occurring.

Appendix O HJB Equations with 2 Technology Steps

As we are using the simplified model setting with only technology jumps, there are only three possible jump states to consider in this setting: the final technology state, the intermediate technology state, and the pre-technology jump state. As in the baseline specification, we start by giving the final jump state HJB equation. This final technology value function is an input into the HJB equation for the intermediate technology state, which we give second. Lastly, we construct the HJB equation for the pre-technology jump state, which depends on the value function for the intermediate technology state. We denote when the intermediate technology jump has occurred by a superscript $\tilde{L}$. This intermediate-technology jump realization corresponds to the state where the abatement technology parameter is given by $\tilde{\phi}_0 = \frac{1}{2}\phi_0$. We denote when the final technology jump has occurred by a superscript L. This final-technology jump realization corresponds to the state where the abatement technology parameter is given by $\phi_0^L = 0$. As noted in Section 9, we close down all uncertainty channels except for the technology jump distortions in this simplified model setting with only technology jumps.

O.1 Final technology

The value function characterization for the final-technology jump state is essentially identical to those given in the baseline post-damage, post-technology jump setting. This is because $\phi_0 = 0$, so full abatement and no R&D investment are the Planner's optimal choices. Therefore, temperature remains constant and R&D knowledge stock is irrelevant to the Planner and so the problem simplifies to the same capital-based settings as before. We refer the readers to Appendix A for details related to this setting.

O.2 Intermediate technology and pre-technology jump

The remaining two cases are the HJB equation for the value function $V^{\tilde{L}}$ for after the intermediate technology jump realization and the value function V prior to any jump realization in this two-step technology setting. The characterization of these two HJB equations is very similar, but with critical differences in the model inputs used for each HJB equation. For the intermediate technology state HJB equation, the value of the abatement technology parameter is $\tilde{\phi}_0$, the value function realized when a jump occurs is V^L, and the relevant jump intensity and probability distortion used for the jump process ambiguity concerns are given by $\mathcal{J}^L$ and g^L. For the pre-technology jump state HJB equation, the value of the abatement

technology parameter is ϕ_0, the value function realized when a jump occurs is $V^{\tilde{L}}$, and the relevant jump intensity and probability distortion used for the jump process ambiguity concerns are given by $\mathcal{J}^{\tilde{L}}$ and $g^{\tilde{L}}$.

Therefore, the complete HJB equation for the intermediate technology state is:

$$\begin{aligned}
0 \;=\; & \max_{i^k, i^r, e} \min_{g^L} \delta \log \left[\alpha - i^k - i^r - \alpha \tilde{\phi}_0 \left(1 - \frac{e}{\beta \alpha k} \right)^{\phi_1} \right] + \delta \hat{k} - \delta \hat{n}(y) - \delta V^{\tilde{L}}(\hat{k}, y, \hat{r}) \\
& + \frac{\partial V^{\tilde{L}}}{\partial \hat{k}}(\hat{k}, y, \hat{r}) \left[-\mu_k + i^k - \frac{\kappa}{2} \left(i^k \right)^2 - \frac{|\sigma_k|^2}{2} \right] + \frac{\partial^2 V^{\tilde{L}}}{\partial \hat{k} \, \partial \hat{k}'}(\hat{k}, y, \hat{r}) \frac{|\sigma_k|^2}{2} \\
& + \frac{\partial V^{\tilde{L}}}{\partial y}(\hat{k}, y, \hat{r}) \bar{\theta} e + \frac{\partial^2 V^{\tilde{L}}}{\partial y \, \partial y'}(\hat{k}, y, \hat{r}) \frac{|\varsigma|^2}{2} e^2 \\
& + \frac{\partial V^{\tilde{L}}}{\partial \hat{r}}(\hat{k}, y, \hat{r}) \left[-\zeta + \psi_0 (i^r)^{\psi_1} \exp\left(-\psi_1 \left(\hat{r} - \hat{k} \right) \right) - \frac{|\sigma_r|^2}{2} \right] + \frac{\partial^2 V^{\tilde{L}}}{\partial \hat{r} \, \partial \hat{r}'}(\hat{k}, y, \hat{r}) \frac{|\sigma_r|^2}{2} \\
& + \mathcal{J}^L g^L \left[V^L(\hat{k}, y, \hat{r}) - V^{\tilde{L}}(\hat{k}, y, \hat{r}) \right] + \xi_m \mathcal{J}^L \left[1 - g^L + g^L \log g^L \right],
\end{aligned}$$

and the complete HJB equation for the pre-technology jump state is:

$$\begin{aligned}
0 \;=\; & \max_{i^k, i^r, e} \min_{g^{\tilde{L}}} \delta \log \left[\alpha - i^k - i^r - \alpha \phi_0 \left(1 - \frac{e}{\beta \alpha k} \right)^{\phi_1} \right] + \delta \hat{k} - \delta \hat{n}(y) - \delta V(\hat{k}, y, \hat{r}) \\
& + \frac{\partial V}{\partial \hat{k}}(\hat{k}, y, \hat{r}) \left[-\mu_k + i^k - \frac{\kappa}{2} \left(i^k \right)^2 - \frac{|\sigma_k|^2}{2} \right] + \frac{\partial^2 V}{\partial \hat{k} \, \partial \hat{k}'}(\hat{k}, y, \hat{r}) \frac{|\sigma_k|^2}{2} \\
& + \frac{\partial V}{\partial y}(\hat{k}, y, \hat{r}) \bar{\theta} e + \frac{\partial^2 V}{\partial y \, \partial y'}(\hat{k}, y, \hat{r}) \frac{|\varsigma|^2}{2} e^2 \\
& + \frac{\partial V}{\partial \hat{r}}(\hat{k}, y, \hat{r}) \left[-\zeta + \psi_0 (i^r)^{\psi_1} \exp\left(-\psi_1 \left(\hat{r} - \hat{k} \right) \right) - \frac{|\sigma_r|^2}{2} \right] + \frac{\partial^2 V}{\partial \hat{r} \, \partial \hat{r}'}(\hat{k}, y, \hat{r}) \frac{|\sigma_r|^2}{2} \\
& + \mathcal{J}^{\tilde{L}} g^{\tilde{L}} \left[V^{\tilde{L}}(\hat{k}, y, \hat{r}) - V(\hat{k}, y, \hat{r}) \right] + \xi_m \mathcal{J}^{\tilde{L}} \left[1 - g^{\tilde{L}} + g^{\tilde{L}} \log g^{\tilde{L}} \right].
\end{aligned}$$

Note that the constraints $0 \leqslant e \leqslant \beta \alpha k$, $g^{\tilde{L}} \geqslant 0$, and $g^L \geqslant 0$ still hold as in the main specification. In addition, the value of R_t is reinitialized at R_0 in the intermediate technology state for specification iii).

Appendix P Damage function jump intensity

Jump intensities imply densities over the possible success times. Here we illustrate implications for the damage information event. We have already computed implications for planner solutions. When the temperature anomaly is an increasing function of calendar time, a density over the potential time realizations of the event implies a corresponding density over the potential temperature anomalies that trigger the

event, which we now illustrate.

Denote increasing temperature anomaly by Y_t starting from an initial time $t = 0$. Then for a given intensity $\mathcal{J}$, the jump time distribution is:

$$1 - \exp\left[-\int_0^t \mathcal{J}(Y_s)ds\right]$$

with a density given by:

$$\mathcal{J}(Y_t)\exp\left[-\int_0^t \mathcal{J}(Y_s)ds\right].$$

With the familiar change of variables formula, we may compute a density over the alternative temperature anomalies.

For purposes of illustration, consider the special case in which[42]

$$Y_t = Y_0 + \alpha t.$$

Then the change-of-variables formula gives the following implied density over the temperature anomalies:

$$\frac{1}{\alpha}\mathcal{J}(y)\exp\left[-\frac{1}{\alpha}\int_{Y_0}^{y} \mathcal{J}(\tilde{y})d\tilde{y}\right]$$

In the computations that follow, we consider three values for the temperature response in degrees Celsius per year: $\alpha = .014, .020, .025$. As rough approximations, we took the .1, .5, and .9 quantiles of Figure 1 and combined them with an annual carbon emissions estimate of about 10.5 GtC per year based on Figueres et al. (2018). The choice of $\alpha = .020$ matches the approximation suggested by Berkeley Earth's Global Temperature Report 2025 prepared by Robert Rohde, which states: "Over the long-term man-made global warming has been responsible for gradually increasing temperatures at a rate of $\approx 0.20^{\circ}$/decade…".

Figure 16 gives the implied densities expressed in terms of the temperature anomaly for the three different values of α. Increasing (decreasing) α shifts the density to the right (left) and makes it less (more) peaked at the interior mode. Figure 17 reports the damage information event jump intensity function used in our calculations along with what happens when we double and halve d_1. Figure 18 shows the implied densities expressed in terms of the temperature anomaly. Increasing (decreasing) d_1 shifts the density to

[42]This is indeed an illustration because both the planner solutions and actual empirical evidence suggest alternative forms of nonlinearity

the left (right) and makes it more (less) peaked at the interior mode.

We note that for Figures 16 - 18 we impose a numerical cutoff of $y = 2.49$ because the intensity function diverges as y approaches $\overline{y} = 2.5$. While the density functions spike up near the upper boundary, we confirm that they do in fact tend to zero at $\overline{y} = 2.5$, as expected.

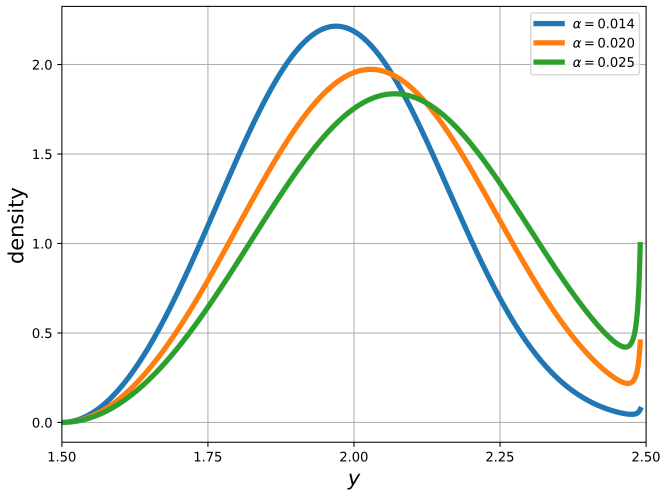


Figure 16: Damage information event jump densities for alternative α's.

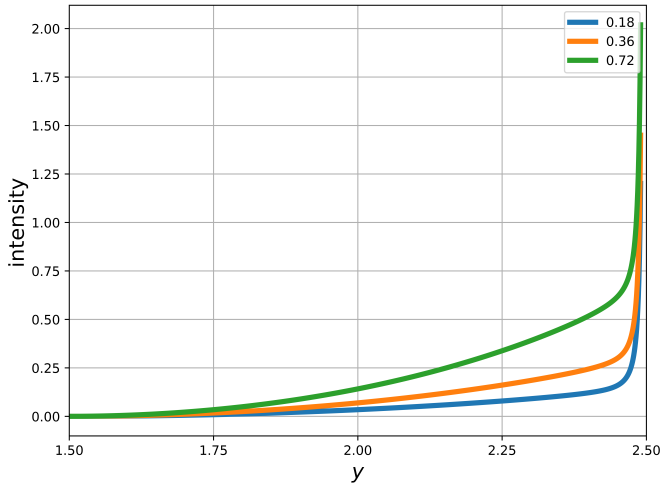


Figure 17: Jump intensity functions for the damage information event for alternative values of d_1. We used $\mathsf{d}_1 = .36$ in our calculations.

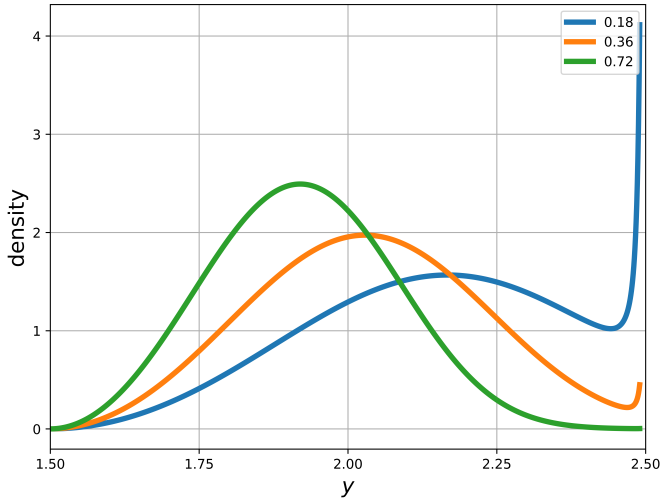


Figure 18: Damage information event jump densities for alternative d_1's. The parameter $\alpha = .020$ when generating these densities. We used $\mathsf{d}_1 = .36$ in our computations.