Stochastic Responses and Marginal Valuation: Supporting Information*

Lars Peter Hansen† and Panagiotis Souganidis‡

October 17, 2025

1 An illustrative comparison to the co-state equation

It is pedagogically revealing to consider a linear-quadratic stochastic control problem. For such a problem, we exploit the quasi-analytical solutions to illustrate the distinction and relation between our results and the analysis of co-state processes.

Consider the state equation:

$$dX_t = (\mathbb{A}X_t + \mathbb{B}F_t)\, dt + \mathbb{C}dW_t.$$

where F is a control or decision process. Suppose the utility function (inclusive of a scaling by δ) is

$$-\frac{1}{2}{F_t}'\mathbb{P}F_t - \frac{1}{2}{X_t}'\mathbb{Q}X_t$$

where we take both $\mathbb{P}$ and $\mathbb{Q}$ to be symmetric.[1] Let Ψ be the co-state process, which satisfies the backward stochastic differential equation:

$$-d\Psi_t = -\mathbb{Q}X_t dt + (\mathbb{A}' - \delta\mathbb{I})\, \Psi_t dt - Z_t dW_t.$$

*This Supplemental Appendix benefited from comments from Jessie Liao and particularly from Xiaoyang Xu.

†University of Chicago.

‡University of Chicago.

[1]As is well known in linear-quadratic control theory, we may always transform the control to achieve the separability in the quadratic utility function.

Finally,

$$0 = -\mathbb{P}F_t + \mathbb{B}'\Psi_t.$$

Suppose that the $\mathbb{A}$ matrix has eigenvalues whose real parts are less than or equal to zero.[2] Then we may express:

$$\Psi_0 = -\int_0^\infty \exp(-t\delta)\exp(t\mathbb{A})' E(\mathbb{Q}X_t \mid X_0)\, dt \tag{1}$$

by in effect solving the co-state equation forward. While valid, this is not the representation that interests us. In particular, the dependence of the control on the state is omitted in the marginal utility contribution.[3]

Next, we consider an alternative. We initially suppose that the optimal control is

$$F_t^* = -\mathbb{F}^* X_t,$$

and that we entertain alternatives

$$G_t = F_t - F_t^*,$$

and we note that, by construction, the optimal $G_t^* = 0$.

Then the utility function can be rewritten as

$$\begin{aligned} &-\frac{1}{2}(G_t + F_t^*)' \mathbb{P}(G_t + F_t^*) - \frac{1}{2}X_t{}'\mathbb{Q}X_t \\ &= -\frac{1}{2}G_t{}'\mathbb{P}G_t - \frac{1}{2}X_t{}'\mathbb{Q}^* X_t - X_t{}'\mathbb{N}^* G_t, \end{aligned}$$

where

$$\mathbb{Q}^* \stackrel{\text{def}}{=} \mathbb{Q} + \mathbb{F}^{*\prime}\mathbb{P}\mathbb{F}^* \quad \text{and} \quad \mathbb{N}^* \stackrel{\text{def}}{=} -\mathbb{F}^{*\prime}\mathbb{P}.$$

Finally, we modify $\mathbb{A}$ to be:

$$\mathbb{A}^* \stackrel{\text{def}}{=} \mathbb{A} - \mathbb{B}\mathbb{F}^*$$

[2] This eigenvalue restriction need not be satisfied, as Hamiltonian methods lead to solving the joint state-costate system and the co-state will appear in the state evolution equation when we substitute for the control F_t.

[3] This omission is intimately connected to the so-called Envelope Theorem and the relation between co-states and partial derivatives of value functions.

The backward stochastic differential equation satisfied by the co-state variable Ψ_t is

$$-d\Psi_t = -\mathbb{Q}^* X_t dt - \mathbb{N}^{*\prime} G_t dt + \left(\mathbb{A}^{*\prime} - \delta \mathbb{I}\right) \Psi_t dt - Z_t dW_t,$$

complemented with

$$0 = \mathbb{B}' \Psi_t - \mathbb{P} G_t - \mathbb{N}^{*\prime} X_t.$$

Solving the co-state equation forward yields

$$\Psi_0 = -\int_0^\infty \exp(-t\delta) \exp\left(t\mathbb{A}^*\right)' E\left(\mathbb{Q}^* X_t \mid X_0\right) dt. \tag{2}$$

Of course, we would obtain the same formula for Ψ_t by simply imposing the control $F_t^* = -\mathbb{F}^* X_t$ and not entertain deviations away from this.

Formulas(1) and (2) give two different representations for the initial co-state in terms of forward-looking expectations. As we verify next, representation (2) is the one that is easiest to connect to the asset-pricing type representation that interests us. We note that it is not our aim to revisit solutions to control problems, but instead to obtain substantively interesting marginal value representations that enhance interpretability and are more generally applicable.

Next we interpret formula (2) using our incorporation of marginal responses. The state dynamics (with the optimal control imposed) are given by

$$dX_t = \mathbb{A}^* X_t dt + \mathbb{C} dW_t,$$

and, given the linearity of the state dynamics, the marginal response process is deterministic and solves

$$d\Lambda_t = \mathbb{A}^* \Lambda_t dt,$$

giving the solution

$$\Lambda_t = \exp\left(\mathbb{A}^* t\right) \Lambda_0.$$

Thus we may rewrite each coordinate of formula (2) as:

$$\Lambda_0 \cdot \Psi_0 = -\int_0^\infty \exp(-\delta t) E\left(\Lambda_t \cdot \mathbb{Q}^* X_t \mid X_0, \Lambda_0\right) dt,$$

where Λ_0 is the corresponding coordinate vector.

As we emphasized in the body of the paper, we need not start from an optimal solution. The same formula applies if simply impose a control, say $F_t = -\mathbb{F} X_t$.

2 Recursive utility extension

With an abuse of notation, let $\{V_t; t \geqslant 0\}$ be the continuation utility associated with consumption process $\{C_t : t \geqslant 0\}$. In what follows, let $\widehat{C} \stackrel{\text{def}}{=} \log C$. Under a unitary risk aversion, the local evolution of the continuation value is:

$$\lim_{\epsilon \downarrow 0} \frac{1}{\epsilon} \left[E\left(V_{t+\epsilon} \mid \mathfrak{F}_t\right) - V_t \right] = \frac{\delta}{\rho - 1} \left(\exp\left[(1-\rho) \left(\widehat{C}_t - V_t \right) \right] - 1 \right)$$

where $1/\rho$ is the elasticity of intertemporal substitution and δ is a subjective rate of discount. Now write: $\widehat{C}_t = \hat{c}(X_t)$ and $V_t = V(X_t)$. Compute

$$\begin{aligned} &\frac{\partial}{\partial x} \left(\frac{\delta}{\rho - 1} \right) \left[\exp\left[(1-\rho) \left[\hat{c}(x) - V(x) \right] \right] - 1 \right] \\ &= -\delta \exp\left[(1-\rho) \left[\hat{c}(x) - V(x) \right] \right] \left[\frac{\partial \hat{c}}{\partial x}(x) - \frac{\partial V}{\partial x}(x) \right] \end{aligned}$$

Thus δ is replaced by:

$$\delta \exp\left[(1-\rho) \left[\hat{c}(x) - V(x) \right] \right]$$

which impacts both the flow term from the utility gradient and discount rate term in the FK equation for DV unless $\rho = 1$. The change in probability measure used for robust adjustments applies to

$$\lim_{\epsilon \downarrow 0} \frac{1}{\epsilon} \left[E\left(V_{t+\epsilon} \mid \mathfrak{F}_t\right) - V_t \right]$$

as in Section 3.1.

3 Derivation of results from Section 3.1.2

We provide more details supporting the analysis in Section 3.1.2. Recall that artificial process $\{\overline{X}_t : t \geqslant 0\}$ evolves according to

$$d\overline{X}_t = \left[\mu\left(\overline{X}_t\right) + \sigma\left(\overline{X}_t\right) h^*\left(\overline{X}_t\right) \right] dt + \sigma\left(\overline{X}_t\right) dW_t^{\overline{H}^*} \quad \overline{X}_0 = \overline{x}. \tag{3}$$

The original process $\{X_t : t \geqslant 0\}$ is replaced by

$$dX_t = \left[\mu(X_t) + \sigma(X_t)h^* \left(\overline{X}_t\right)\right] dt + \sigma(X_t) dW_t^{\overline{H}^*} \quad X_0 = x. \tag{4}$$

If $X_0 = \overline{X}_0 = x$, then (3) and (4) yield that $X_t = \overline{X}_t$ for all $t > 0$. Therefore, for all x,

$$V(x) = \overline{V}(x, x),$$

Since V is the value function obtained by minimizing over H, initializing $\overline{X}$ at some value different from the initial condition for X results in a suboptimal choice of H and hence a larger initial value function. It follows that, for all $x, \bar{x}$,

$$\overline{V}(x, \bar{x}) \geqslant V(x).$$

Moreover,

$$\frac{\partial \overline{V}}{\partial \overline{x}}(x, x) = 0 \text{ and hence } \frac{\partial \overline{V}}{\partial x}(x, x) = DV(x).$$

4 Derivation of Results from Section 3.2

Given a pre-specified $\hat{h}(x)$, a particular discounted relative entropy K can be computed by solving the Feynman-Kac equation:

$$0 = \frac{\delta}{2}\hat{h}'\hat{h} - \delta K + \mathbb{A}K + DK \cdot \left(\sigma \hat{h}\right). \tag{5}$$

For this same K, other $h's$ will solve this same equation. We consider an entire family of alternative probabilities through a convex set constructed with the inequality

$$\mathcal{S}(x) = \left\{ s : \frac{1}{2}s's + \frac{1}{\delta}DK(x) \cdot [\sigma(x)s] \leqslant K(x) - \frac{1}{\delta}\mathbb{A}K(x) \right\},$$

including ones with smaller discounted relative entropies state-by-state. For convenience, we rewrite the constraint as:

$$\frac{1}{2} \mid s - \hat{s} \mid^2 \leqslant \frac{1}{2}R^2,$$

where

$$\hat{s} \stackrel{\text{def}}{=} -\frac{1}{\delta}\sigma' DK \qquad \text{and} \qquad R^2 \stackrel{\text{def}}{=} K - \frac{1}{\delta}\mathbb{A}K + \frac{1}{2} \mid \hat{s} \mid^2.$$

With this extra structure, the robust valuation is reflected in the HJB equation

$$0 = \delta U - \delta V + \mathbb{A} V + \mathcal{H}(DV, \cdot), \tag{6}$$

where

$$\begin{aligned} \mathcal{H}(p, \cdot) \stackrel{\text{def}}{=} \min_{s \in S(x)} (\sigma' p) \cdot s &= (\sigma' p) \cdot s^* \\ &= -\frac{1}{\delta}(\sigma' p) \cdot (\sigma' DK) - R \mid \sigma' p \mid \end{aligned} \tag{7}$$

and R is the positive square root of R^2. In (7),

$$s^*(p, \cdot) = -\frac{1}{\delta}\sigma' DK - \frac{R}{|\sigma' p|}\sigma' p.$$

Consider the analog to Approach 1 in Section 3.1, where, when performing our marginal value calculations, we take the drift to be:

$$\mu(x) + \sigma(x) s^*[DV(x), x].$$

To elaborate, first compute

$$D_p s^*(p, \cdot)' = -\frac{R}{\mid \sigma' p \mid}\sigma \left[I - \left(\frac{1}{p'\sigma\sigma' p} \right) \sigma' p p' \sigma \right],$$

and note that

$$D_p s^*(p, \cdot)' \sigma' p = 0.$$

This result is just a special case of the Envelope Theorem since s^* solves the corresponding minimization problem. Thus,

$$D_p \mathcal{H}(p, \cdot) = \sigma s^*(p, \cdot).$$

Moreover,

$$D_x \mathcal{H}(p, x) = D_x[\sigma(x) s^*(p, x)' \sigma(x)'] p.$$

Next impose that $p = DV$, and note that

$$D_p \mathcal{H}(p, x) \cdot D_x(DV \cdot \lambda) = \sigma s^* \cdot D_x(DV \cdot \lambda)$$

contributes the change in drift for x in the robust-adjusted dynamics for the function $DV \cdot \lambda$.

The term:

$$D_x \mathcal{H}(DV, x) \cdot \lambda$$

captures the altered contribution to the robust-adjusted λ dynamics. Thus differentiating (6) and forming a dot product with λ gives:

$$\begin{aligned}\delta DV \cdot \lambda =& \delta U \cdot \lambda + \mathbb{B}(DV \cdot \lambda) \\ &+ D_p \mathcal{H}(DV, x) \cdot D_x(DV \cdot \lambda) + D_x \mathcal{H}(DV, x) \cdot \lambda \\ =& \delta U \cdot \lambda + \mathbb{B}^*(DV \cdot \lambda)\end{aligned} \tag{8}$$

where $\mathbb{B}^*$ is the robust-adjusted generator under Approach 1.

For Approach 2, we start by refining the structure and partition the state vector into two components

$$X_t = \begin{bmatrix} X_{1,t} \\ X_{2,t}. \end{bmatrix}$$

The first contains the endogenous states and the second the additional exogenous states that impact the decision environment but cannot be influenced by the decision maker.

Suppose that the more structured uncertainty pertains only to the state dynamics for $\{X_{2,t}\}$. Thus:

$$\mu(x) = \begin{bmatrix} \mu_1(x) \\ \mu_2(x_2) \end{bmatrix} \quad \text{and} \quad \sigma(x) = \begin{bmatrix} \sigma_1(x) \\ \sigma_2(x_2) \end{bmatrix}.$$

The Approach 1 construction has s^* depending on the entire state vector, including the endogenous states. Since we only target the misspecification of the $\{X_{2,t}\}$ dynamics, the entropy function, K, depends only on x_2 as does $\hat{s}$ and R^2.

We now form an artificial process:

$$d\overline{X}_{1,t} = \mu_1 \left(\overline{X}_{1,t}, X_{2,t}\right) dt + \sigma_1 \left(\overline{X}_{1,t}, X_{2,t}\right) dW_t,$$

and set

$$\bar{s}\left(\bar{x}_1, x_2\right) = -\frac{1}{\delta}\sigma_2(x_2)' D_2 K(x_2) - \frac{R(x_2)}{\mid \sigma(\bar{x}_1, x_2)' DV\left(\bar{x}_1, x_2\right) \mid} \sigma\left(\bar{x}_1, x_2\right)' DV\left(\bar{x}_1, x_2\right),$$

noting that $\bar{s}$ satisfies the flow entropy constraint.

We use $\bar{s}\left(\overline{X}_{1,t}, X_2, t\right)$ to construct our altered specification for dW_t, and we now let

$\overline{V}(x_1, x_2, \bar{x}_1)$ denote the resulting value function. Then

$$V(x) = \overline{V}(x_1, x_2, x_1),$$

$$\overline{V}(x_1, x_2, \bar{x}_1) \geqslant V(x)$$

and, analogously to our derivation for Approach 2 in Section 3.1.2,

$$\frac{\partial \overline{V}}{\partial \bar{x}_1}(x_1, x_2, x_1) = 0.$$

Hence

$$\frac{\partial \overline{V}}{\partial x}(x_1, x_2, \bar{x}_1) = DV(x).$$

5 Derivation of Results from Section 3.3

5.1 Robust learning solution for the discrete parameter example

The minimization problem on the right side of equation (22) has the exponential tilting solution

$$\pi_i^* = \frac{\exp\left(-\frac{1}{\xi}\eta_i\right)\hat{z}_i}{\sum_j \exp\left(-\frac{1}{\xi}\eta_j\right)\hat{z}_j},$$

where the η_i's are entries of

$$\eta(p, x) \stackrel{\text{def}}{=} \kappa(y)' \left[\varsigma(y)\varsigma(y)'\right]^{-1/2} \sigma(y)' p.$$

The minimized objective for $p = DV$ is

$$\begin{aligned}\mathbb{M}(DV, x) &= \eta(DV, x) \cdot (\pi^* - \hat{z}) + \xi \sum_i (\log \pi_i^* - \log \hat{z}_i)\pi_i^* \\ &= -\xi \log\left(\sum_i \exp\left[-\frac{1}{\xi}\eta_i(DV, x)\right]\hat{z}_i\right) - \eta(DV, x) \cdot \hat{z}\end{aligned}$$

5.2 Wonham filtering

We suppose that Z is a discrete state Markov jump process with I states and time invariant transitions. We continue to assume that its realizations are coordinate vectors and

$$\nu(y,z) = \kappa(y)z,$$

where the z realization captures alternative "regimes" for the state dynamics and each regime is governed by a column of $\kappa(y)$. We denote by Θ the intensity matrix which has nonnegative off diagonal entries and row sums equal to zero.

An object of interest is the vector $\widehat{Z}_t = \mathbb{E}(Z_t \mid \mathfrak{F}_t)$ whose coordinates are the filtered probability for the discrete states. Under the reduced information, we find that the state dynamics are given

$$\begin{aligned} dY_t &= \kappa(Y_t)\widehat{Z}_t dt + \left[\varsigma(Y_t)\varsigma(Y_t)'\right]^{1/2} d\widehat{W}_t \\ d\widehat{W}_t &= \left[\varsigma(Y_t)\varsigma(Y_t)\right]^{-\frac{1}{2}} \left[\kappa(Y_t)dZ_t - \kappa(Y_t)d\widehat{Z}_t + \varsigma(Y_t)dW_t\right], \end{aligned} \tag{9}$$

with $\widehat{W}$ a multivariate Brownian motion under the $\mathfrak{F}$ filtration. The recursive filtering equation for $\widehat{Z}$ is

$$\begin{aligned} d\widehat{Z}_t &= \Theta'\widehat{Z}_t dt + \left[\operatorname{diag}\left(\widehat{Z}_t\right) - \widehat{Z}_t\widehat{Z}_t'\right]\kappa(Y_t)'\left[\varsigma(Y_t)\varsigma(Y_t)'\right]^{-1}\left[dY_t - \kappa(Y_t)\widehat{Z}_t dt\right] \\ &= \Theta'\widehat{Z}_t dt + \left[\operatorname{diag}\left(\widehat{Z}_t\right) - \widehat{Z}_t\widehat{Z}_t'\right]\kappa(Y_t)'\left[\varsigma(Y_t)\varsigma(Y_t)'\right]^{-1/2} d\widehat{W}_t, \end{aligned} \tag{10}$$

where $\operatorname{diag}(\hat{z})$ is the diagonal matrix with the entries of $\hat{z}$ on the diagonal.

We rewrite (9) and (10) in terms of the composite state vector

$$X_t = \begin{bmatrix} Y_t \\ \widehat{Z}_t \end{bmatrix},$$

which evolves according to

$$dX_t = \mu(X_t)dt + \sigma(Y_t)d\widehat{W}_t$$

where

$$\mu(x) = \begin{bmatrix} \kappa(y) \\ \Theta' \end{bmatrix} \hat{z} \quad \text{and} \quad \sigma(y) = \begin{bmatrix} \left[\varsigma(y)\varsigma(y)'\right]^{1/2} \\ \left[\operatorname{diag}(\hat{z}) - \hat{z}\hat{z}'\right]\kappa(y)'\left[\varsigma(y)\varsigma(y)'\right]^{-1/2} \end{bmatrix}.$$

While the discrete hidden state Z now evolves stochastic over the time, the mathematical form of the smooth ambiguity-robustness adjustment remains virtually the same as the formulas in Section 3.3 of the main paper.

6 A substantive example

Consider a simplified model of climate-economic dynamics. This model is very close to that used by Barnett et al. (2025). Output is using an AK type production function with adjustment costs and climate damages. Output is split between consumption and two different types of investment with distinct intertemporal contributions to production: a conventional capital investment, I_t^k, and an investment in R&D, I_t^r (the superscripts denote the investment type):

$$C_t + I_t^k + I_t^r = \left(\frac{\alpha}{N_t}\right)\Upsilon(K_t, \mathcal{E}_t) \tag{11}$$

where Υ is a constant returns to scale production function given by:

$$\Upsilon(K_t, \mathcal{E}_t) \stackrel{\text{def}}{=} \left(1 - \phi_0 \left(A_t\right)^{\phi_1}\right),$$

for $\phi_1 \geqslant 2$ and $0 < \phi_0 \leqslant 1$, where A_t is a measure of abatement. Abatement is defined to be

$$A_t \stackrel{\text{def}}{=} \left(1 - \frac{\mathcal{E}_t}{\beta\alpha K_t}\right)\mathbf{1}_{\{\mathcal{E}_t < \beta\alpha K_t\}},$$

where $\mathbf{1}$ is an indicator function that assigns one to the event in the $\{\cdot\}$ brackets. Abatement only plays a role in the production when $\mathcal{E}_t \leqslant \beta\alpha K_t$ analogous to specifications used in so-called DICE models created by Nordhaus and his collaborators. See, for instance, Nordhaus (2017). The term N_t is included as a proportional reduction in output induced by global warming, which we view as an externality.

Carbon emissions $\mathcal{E}_t$ are a proxy for a "dirty" energy input into production. When emissions fall short of the threshold $\beta\alpha K_t$, there is a corresponding convex adjustment in the output given by the right-hand side of (11).

Our specification of N_t uses a piece-wise log quadratic specification as a function of the temperature anomaly. By temperature anomaly we mean temperature net of its preindustrial counterpart. Let the date t anomaly be denoted by Y_t and a potential realization by

y. We suppose that

$$\hat{n}(y) = \lambda_1 y + \frac{\lambda_2}{2} y^2 \qquad\qquad y \leqslant \tilde{y}$$
$$\hat{n}(y) = \lambda_1 y + \frac{\lambda_2}{2}\tilde{y}^2 + \frac{\lambda_2}{2}\left[(y - \tilde{y} + \bar{y})^2 - \bar{y}^2\right] + \frac{\lambda_3(\ell)}{2}(y - \tilde{y})^2 \quad y > \tilde{y} \tag{12}$$

for $\ell = 1, ..., L-1$. The variable $\tilde{y}$ serves a threshold triggered by a Poisson jump prior to an upper limit $\bar{y}$. The intensity, $\mathcal{J}^n$, depends on the temperature anomaly and is specified so that this jump takes place in the interval $[\underline{y}, \bar{y}]$. See Barnett et al. (2025) for the specific details. We shift the derivative of the logarithms of damages with respect to temperature changes from $\lambda_1 + \lambda_2\tilde{y}$ to $\lambda_1 + \lambda_2\bar{y}$ when the jump happens. We also increase the slope by including a term $\lambda_3(\ell)(y - \tilde{y})$, where the coefficient $\lambda_3(\ell)$ is *ex ante* uncertain.

To capture damage uncertainty, we allow for several possible values of λ_3 indexed by ℓ in formula (12). The λ_3's are equally spaced over a pre-specified range. As a baseline probability specification, we take each possible value to be equally likely. Mathematically, we treat each such possibility as a distinct jump with the jump intensity given by $[1/(L-1)]\mathcal{J}^n$.

For pedagogical simplicity, we consider the case of a single technology jump to a fully productive green technology. Initially $\phi_0 > 0$ and then at some point in the future a fully green technology becomes economically viable, in which case $\phi_0^L = 0$. Once this happens dirty energy is no longer needed to produce output. We use the superscript L to denote a realization of technology advance in the future.

There are three state variables in this model. The first is the stock of productive capital, K_t, evolves as

$$dK_t = K_t\left[-\mu_k + \frac{I_t^k}{K_t} - \frac{\kappa}{2}\left(\frac{I_t^k}{K_t}\right)^2\right]dt + K_t\sigma_k dW_t, \tag{13}$$

where σ_k is a row vector with the same dimension as the underlying Brownian motion. New investment, I_t^k, augments the capital stock, K_t, subject to an adjustment cost captured by the curvature parameter κ. Capital is broadly conceived to include human capital and intangible capital.

A state vector R_t captures the stock of R&D-induced knowledge capital and evolves as

$$dR_t = -\zeta R_t dt + \psi_0 \left(I_t^r\right)^{\psi_1} \left(R_t\right)^{1-\psi_1} dt + R_t\sigma_r dW_t, \tag{14}$$

where $0 < \psi_1 < 1$ and I_t^r is an investment in research and development. The date t

knowledge stock, R_t, gives the jump intensity for the R&D discovery at that date, $\mathcal{J}_t^L = R_t$. The term $\sigma_r dW_t$ reflects an exogenous stochastic inflow of information about the future likelihood of a technological advance.

The third state variable is the temperature anomaly, Y_t. We follow the simplified climate dynamics based on an approximation from the geoscience literature to capture its dynamics. Specifically, Matthews et al. (2009) and others have purposefully constructed an approximation for climate model impacts:

$$\text{temperature anomaly } (Y_t) \approx \text{ TCRE}(\theta) \times \text{ cumulative emissions },$$

where TCRE is an acronym for the Transient Climate Response to cumulative emissions.

Let Y_t denote the date t temperature anomaly (temperature net of a pre-industrial baseline). Our specific form is given by

$$dY_t = \mathcal{E}_t \theta dt + \sigma_y(Y_t) dW_t \tag{15}$$

where θ is the equally weighted average of the elements of set of TCRE's implied by alternative climate models. The term $\sigma_y(Y_t)dW_t$ captures short time scale fluctuations in how emissions impact temperature change. We allow for this local variance to increase as the earth warms.[4]

The state vector X_t consists of K_t, R_t, and Y_t. We construct the μ vector and σ matrix for X by stacking the equations: (13), (14), and (15) and substituting time invariant functions of the state vector, X_t, for the controls, I_t^k, I_t^r, and $\mathcal{E}_t$

While Barnett et al. (2025) feature the design of robustly optimal policies, here we consider potentially suboptimal policies; but we restrict them to be time-invariant functions of the state vector. Suboptimal baselines are commonly used for marginal valuations.

We take instantaneous utility function be logarithmic in consumption. We substitute consumption from constraint (11) and take instantaneous utility contribution to be:

$$\delta \log \left[\left(\frac{\alpha}{N_t} \right) \Upsilon(K_t, \mathcal{E}_t) - I_t^k - I_t^r \right]$$

[4]The local variance specification differs from that used by Barnett et al. (2025). This previous work omitted the state dependence, but it included emissions as a scale factor to capture randomness in the temperature response to emissions. Their approach gives rise to an additional second derivative term in value function that is not present here. While they show that its impact on marginal valuation is small, we feature only forward-looking formulas for the first-derivative in this paper.

We use equation (12) to construct N_t as a function of the temperature state variable Y_t.

In addition to the Brownian motion inputs, the model specification has two types of Poisson jumps, a damage curve realization jump for which there are $L-1$ possibilities, and technology jump eliminating the need for dirty energy. The $L-1$ damage curves are treated as equally likely under baseline probabilities, given that a damage curve realization jump takes place. All L jumps have intensities that depend on endogenous state variables: either temperature or the stock of knowledge.

Given the diffusion dynamics, the jump intensities, and the instantaneous utility contribution, we can derive the Feynman-Kac formula for the value function V.

Two marginal valuations are of particular interest: the marginal cost of climate change (the partial derivative of the value function with respect to temperature) and the marginal benefit of knowledge pertinent to the discovery of a new clean production technology (the partial derivative of the value function with respect to knowledge stock). We use the approach described in Remark 2.4 to deduce these costs and benefits. Consistent with formulation in this paper, we adopt a pre-jump perspective. The jump impacts are captured by continuation values conditioned on the respective L jumps.

Consider first, the marginal social cost of emissions. This is given by intertemporal contribution:

$$-\frac{\partial V}{\partial y}(X_t)\theta.$$

We take the negative because temperature increases have an adverse impact on the value function. The minus sign converts "negative benefits" into positive costs. The marginal benefit comes from the instantaneous utility contribution:

$$\frac{\delta\alpha}{C_t N_t}\Upsilon_e(K_t, \mathcal{E}_t)$$

where from the output equation, (11), C_t is a pre-specified, time invariant function of the state vector, X_t, since the controls and $\log N_t$ are time invariant functions of the state vector, X_t. Our representation and decompositions are aimed at the dynamic cost specification.

Consider second, the marginal benefit of investment in research and development as given by the dynamic contribution

$$\frac{\partial V}{\partial r}(X_t)\psi_0\left(\frac{I_t^r}{R_t}\right)^{\psi_1-1}.$$

The marginal cost is given by static contribution:

$$\frac{\delta}{C_t}.$$

In this case, our analysis is directed at the marginal benefit term.

These marginal valuations given so far, are expressed in utility units. Dividing these expressions by current period marginal (multiplying them by $\frac{C_t}{\delta}$) converts the costs and benefits into consumption units.

Finally, while this model has three state variables, for computational purposes there is some simplification because of the homogeneity properties of the technology. For instance, see Barnett et al. (2025).

References

Barnett, Michael, William A. Brock, Lars Peter Hansen, and Hong Zhang. 2025. Uncertainty, Social Valuation, and Climate Change Policy. SSRN working paper.

Nordhaus, William D. 2017. Revisiting the Social Cost of Carbon. *Proceedings of the National Academy of Sciences* 114:1518–1523.