

WORKING PAPER

The Optimal Carbon-Tax Trajectory: A Finite-Horizon HJB Approach

Closed Form, Numerical Algorithm, and a DICE-Style 25-Year Path

Supervised by Dr. T. Zamrik

13-MAY-2026

Gradient flow of the Gaussian mixture: $\dot{\mathbf{x}} = \nabla f(\mathbf{x})$, $f(x, y) = \sum_{i=1}^n w_i \mathcal{N}(x, y; \mu_i, \Sigma_i)$



WP-2026-74091446
zamrik.com



Zamrik Quant Research

The research arm of Zamrik Quant Lab

Copyright

© 2026 T. Zamrik. All rights reserved. Published by Zamrik Quant Research. No part of this document may be reproduced, stored, distributed or adapted, in whole or in part, by any means, without the prior written permission of the publisher, except for brief quotation with full citation. Requests for permission: info@zamrik.com.

Suggested Citation

Optimal Control & Decision Desk, Zamrik Quant Research (2026). 'The Optimal Carbon-Tax Trajectory: A Finite-Horizon HJB Approach: Closed Form, Numerical Algorithm, and a DICE-Style 25-Year Path'. *Zamrik Quant Research Working Paper WP-2026-74091446*. Available at <https://zamrik.com/research-items/wp-2026-74091446/>.

About Zamrik Quant Research

Zamrik Quant Research is the research arm of Zamrik Quant Lab. It works through five research desks, from market microstructure and derivatives to optimal control, stochastic analysis and applied mathematics, and publishes everything it finds in the open: research papers, notebooks that run on Kaggle, free tools that answer one practical question, and essays.

The five desks are Market Microstructure & Trading; Derivatives & Volatility; Optimal Control & Decision; Stochastic Analysis; and Applied Mathematics. Each paper, notebook, tool and essay belongs to exactly one desk. Zamrik Quant Research is run by a small, independent team, headed by Dr. T. Zamrik. We are self-funded: no sponsors, no advertisers, no outside capital.

Every publication passes through the same five functions before anyone outside sees it: research, data, software, verification and production. We never retrofit the mathematics to a conclusion chosen in advance. Every number is checked two independent ways. Every figure is re-run from its code. When we find an error, or replace a piece with a better one, we record it in the public corrections log at <https://zamrik.com/corrections/>.

Legal Notice

This paper is research and education. Nothing in it is investment, financial, legal or engineering advice; nothing is a recommendation to buy or sell shares of stock, derivative contracts, bonds, or any other instrument or property; and nothing is a forecast. You are responsible for your own decisions. Figures and code are provided as they are. Care and independent checks reduce errors but cannot rule them out, so no warranty is given. The views expressed are those of the Optimal Control & Decision Desk.

WORKING PAPER



The Optimal Carbon-Tax Trajectory: A Finite-Horizon HJB Approach

*Closed Form, Numerical Algorithm, and a DICE-Style 25-Year
Path*

Supervised by Dr. T. Zamrik¹

13-MAY-2026

WP-2026-74091446

Abstract. We derive the optimal time-varying carbon-tax trajectory as the solution of a finite-horizon Hamilton–Jacobi–Bellman problem for a planner who balances the output cost of the tax against quadratic climate damage on the cumulative-emissions state. For the linear–quadratic specification the value function and the optimal tax admit a closed form: a stationary, Pigouvian term plus a finite-horizon correction that decays at an explicit rate. With physical anchors from the IPCC Sixth Assessment Report (cumulative emissions of 2390 GtCO₂ by 2019, 0.45 °C of warming per 1000 GtCO₂) and an illustrative, DICE-style damage and cost calibration, the 25-year planner’s initial tax is \$2.61/tCO₂ against a stationary benchmark of \$5.80, 55 % lower. The dollar levels follow the illustrative anchors; the size of the gap is set mainly by how the terminal condition values the stock left at the horizon. Every headline number is computed twice, by the closed form and by an independent method: an ODE solve, a direct optimisation that does not use the HJB equation, Monte Carlo, or finite differences.

WP-2026-74091446 · Optimal Control & Decision Desk · Working Paper · Published
13-MAY-2026 · Reading time: about 32 min

1. Introduction

How should a planner who knows the climate problem will be resolved — one way or another — within a generation set the carbon tax today? The infinite-horizon Pigouvian answer, in which the tax is permanently calibrated to the steady-state shadow price of the climate stock, sits at the centre of integrated-assessment practice [1], [2], [3]. Yet most policy timetables are short: nationally determined contributions are revised every five years, central banks plan transition stress tests over a decade, and the major emitters’ net-zero commitments cluster around 2050. Whether the finite-horizon optimal tax should differ from its infinite-horizon counterpart — and by how much — is a quantitative question, not a philosophical one. This paper answers it for a linear–quadratic model with a DICE-style quadratic damage function [4].

We study a stochastic planner who balances the convex output cost of a carbon tax against

¹Correspondence: Optimal Control & Decision Desk, Zamrik Quant Research, info@zamrik.com.



quadratic damage on the cumulative-emissions state, over a deterministic horizon T . For the linear-quadratic specification we characterise the value function and the optimal feedback law in closed form through a Riccati ODE system, prove the verification theorem, and verify the closed form against an explicit finite-difference scheme and against a direct optimisation that does not use the HJB equation. With the physical anchors and the illustrative cost calibration of Section 7 and $T = 25$ years, the finite-horizon optimal initial tax is \$2.61/tCO₂, against an infinite-horizon Pigouvian benchmark of \$5.80/tCO₂: a horizon correction of 55 %. The dollar levels inherit the illustrative anchors; the size of the correction is governed mainly by the terminal condition (Section 5). Those levels sit far below observed carbon prices and are not a forecast of them; the robust finding is the relative gap between the finite-horizon tax and the Pigouvian benchmark, not its dollar level.

The paper makes three contributions.

1. **Closed-form Riccati reduction with explicit horizon-decay rate.** The value function admits a quadratic ansatz $V(t, E) = c_2(t) E^2 + c_1(t) E + c_0(t)$, reducing the HJB partial differential equation to three coupled ODEs. The leading coefficient c_2 has a fully analytic solution governed by a single rate $\lambda = \sqrt{\rho^2 + 2\kappa\beta_E}$, where $\kappa = 2b^2/\alpha$. At the base calibration the half-life $\log 2/\lambda \approx 22$ years, comparable to the horizon T itself, which is why the horizon correction is large.
2. **The horizon correction and what sets it.** We decompose the optimal initial tax into a stationary component and a horizon correction that decays at rate λ , and show that its size is set mainly by the terminal valuation: with a terminal coefficient equal to the stationary one the correction falls from 55 % to 16 %. Under one-at-a-time sensitivity over the discount rate, the damage coefficient, the terminal coefficient, the emissions volatility, the horizon and the reference stock, the damage coefficient and the horizon move the initial tax most over the ranges of Table 4; the volatility does not move it at all (certainty equivalence).
3. **Every number checked twice.** The closed form is checked against an ODE solve of the full Riccati system, against a direct optimisation of the discretised problem that uses neither the HJB equation nor the Riccati system, against the Monte Carlo cost of the feedback, and against an explicit finite-difference scheme for the HJB equation. The scheme converges at first order, because the advection CFL bound ties its inner time step to ΔE ; its relative error at the reference state is 2.0×10^{-4} at $(\Delta t, \Delta E) = (0.25, 6.25)$. It is the natural numerical engine for extensions in which the closed form is unavailable.

The paper sits in a literature on integrated assessment that we read in three strands. The first, exemplified by DICE [4], [2] and the climate-and-risk framework of [5], [6], focuses on infinite-horizon calibration of the social cost of carbon. The second, descending from [3] and [7], derives optimal taxes in general-equilibrium settings with rich structure but typically without analytic feedback laws. The third, anchored by [8], [9], emphasises uncertainty in climate sensitivity. Our contribution is methodological and quantitative: a



small, sharp model that gives the horizon correction in closed form, with an openly stated calibration, and checked by independent numerical methods.

The remainder of the paper is organised as follows. Section 2 fixes the planner's problem. Section 3 derives the HJB equation and reduces it to a Riccati system. Section 4 states and proves the verification theorem, with the proof relegated to Appendix A. Section 5 studies the properties of the optimal feedback law. Section 6 presents the numerical algorithms: closed-form integration of the Riccati system, an explicit finite-difference scheme, and forward Monte Carlo simulation under the optimal policy. Section 7 states the calibration: what is sourced and what is illustrative. Section 8 reports the numerical results and their second checks. Section 9 conducts the sensitivity analysis. Sections 10 and 11 contain the discussion and conclusion. Appendices give the verification proof, the Riccati derivation, and full algorithm pseudocode.

2. The planner's problem

This section fixes notation and states the planner's problem. The state is cumulative emissions; the control is the carbon tax; the planner minimises the expected discounted sum of output cost and damage flow over a finite horizon T , with a terminal damage on the residual stock.

2.1 State and control

Fix a complete filtered probability space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}_{t \in [0, T]}, \mathbb{P})$ carrying a one-dimensional standard Brownian motion $\{W_t\}$. Let $E_t \in \mathbb{R}$ denote the cumulative anthropogenic carbon emissions at time t , measured in gigatonnes of CO_2 (GtCO_2) cumulated since some reference epoch. The control is the carbon tax $\tau_t \geq 0$, in $\$/\text{tCO}_2$, levied uniformly on the marginal tonne of emissions.

Baseline emissions in the absence of any tax are denoted $a > 0$ and expressed in GtCO_2/yr . Tax-induced abatement is linear in the tax at elasticity $b > 0$, so that $b\tau_t$ is the abatement flow induced by the tax. Residual stochastic emissions enter through a diffusion of intensity $\sigma_E > 0$. The emissions stock therefore evolves as

$$dE_t = (a - b\tau_t) dt + \sigma_E dW_t, \quad E_0 \in \mathbb{R} \text{ given.} \quad (2.1)$$

The linearity in τ_t and the Gaussian diffusion are the two ingredients that, combined with the quadratic cost structure below, produce the LQ structure exploited in Section 3. The state is allowed to take negative values formally, which models net carbon removal at the aggregate; physical non-negativity is not imposed because the policy-relevant region of state space lies well above zero.

Definition 1 (Admissible control). A process $\tau = \{\tau_t\}_{t \in [0, T]}$ is *admissible* if it is $\{\mathcal{F}_t\}$ -progressively measurable, satisfies $\tau_t \geq 0$ almost surely, and obeys the integrability condition $\mathbb{E} \int_0^T \tau_t^2 dt < \infty$. We denote the set of admissible controls by $\mathcal{A}$.



2.2 Damage-augmented cost flow

The flow cost combines the output cost of the tax and the damage cost of the stock. Following the convex-tax tradition in [1], [7], the per-instant output cost of a tax τ is quadratic with coefficient $\alpha > 0$; the quadratic form captures both administrative dead-weight loss and the microeconomic distortion from the marginal abatement-cost schedule rising with depth. The stock damage flow is quadratic in E with coefficient $\beta_E > 0$. The terminal damage on the stock at T is quadratic with coefficient $\gamma_T > 0$, absorbing the planner's valuation of post-horizon climate harms in a single closed-form penalty. Both cost terms are measured in per cent of world output per year.

Definition 2 (Cost functional). For an admissible control $\tau \in \mathcal{A}$ and initial state $E_0 = E$, the planner's *cost functional* is

$$J(t, E; \tau) = \mathbb{E} \left[\int_t^T e^{-\rho(s-t)} \left(\frac{1}{2} \alpha \tau_s^2 + \frac{1}{2} \beta_E E_s^2 \right) ds + e^{-\rho(T-t)} \frac{1}{2} \gamma_T E_T^2 \mid E_t = E \right],$$

where $\rho > 0$ is the planner's pure rate of time preference.

The quadratic structure should not be read as a claim that the true damage function is quadratic; rather, it is a tractable approximation whose coefficient is set in Section 7 so that the damage flow is 0.95 % of world output at 2 °C of CO₂-induced warming, an illustrative, DICE-style anchor. Alternative damage specifications are discussed in Section 10.

2.3 Admissibility and integrability

We collect the assumptions that ensure the value function below is well defined and finite.

Assumption 1 (Parameter regularity). $\alpha, b, \beta_E, \gamma_T, \rho, \sigma_E > 0$; $a \geq 0$; $T \in (0, \infty)$.

Assumption 2 (Square integrability of state). For every admissible τ , $\mathbb{E} \sup_{t \in [0, T]} E_t^2 < \infty$.

Assumption 2 is implied by Assumption 1, the integrability built into Definition 1, and standard moment estimates for linear SDEs with bounded drift; see [10], Chapter 1. We record it as a separate assumption because the verification theorem of Section 4 cites it directly.

2.4 The value function

The planner solves

$$V(t, E) = \inf_{\tau \in \mathcal{A}} J(t, E; \tau), \quad (t, E) \in [0, T] \times \mathbb{R}. \quad (2.2)$$

The value function is the central object of the rest of the paper. Section 3 characterises it as the unique classical solution of an HJB equation with quadratic ansatz; Section 4 verifies that this characterisation implies the candidate optimal feedback law is in fact optimal.



Proposition 1 (Well-posedness). Under Assumptions 1 and 2, for every $\tau \in \mathcal{A}$ and every $(t, E) \in [0, T] \times \mathbb{R}$, the SDE for $\{E_s\}_{s \in [t, T]}$ admits a unique strong solution with $\mathbb{E} \sup_{s \in [t, T]} E_s^2 < \infty$, the cost functional $J(t, E; \tau)$ is finite, and the value function $V(t, E)$ is finite.

Proof. Existence and uniqueness of the strong solution follow from the standard Lipschitz-and-linear-growth criterion: the drift $(a - b\tau_s)$ is progressively measurable with $\mathbb{E} \int_0^T |a - b\tau_s|^2 ds < \infty$ by Definition 1, and the diffusion is constant σ_E . See [10], Chapter 1. The second-moment bound follows from Gronwall applied to the squared SDE; J is finite because each of its three integrands is non-negative and dominated by a polynomial in the state with finite second moment; V is finite as the infimum of a non-empty set of finite quantities. $\square$

Remark 1. The non-negativity constraint $\tau \geq 0$ in Definition 1 is the natural physical constraint — the planner cannot subsidise emissions. The closed-form solution below ignores this constraint in the unconstrained quadratic minimisation, then verifies *ex post* that the optimal feedback is non-negative on the policy-relevant region of state space (Section 5). Tax negativity would arise only at deep-negative E , which we exclude as economically uninteresting.

3. HJB equation and Riccati reduction

We derive the HJB partial differential equation, eliminate the control by pointwise minimisation, and reduce the resulting nonlinear PDE to a system of three coupled ordinary differential equations via a quadratic ansatz.

3.1 The HJB equation

Applying the dynamic programming principle to the value function defined in Section 2 yields the standard HJB equation; see [11], Chapter III.

Proposition 2 (HJB equation). If $V \in C^{1,2}([0, T] \times \mathbb{R}) \cap C([0, T] \times \mathbb{R})$ is the value function, then it satisfies

$$-\partial_t V + \rho V = \inf_{\tau \geq 0} \left\{ \frac{1}{2} \alpha \tau^2 + (a - b\tau) \partial_E V + \frac{1}{2} \sigma_E^2 \partial_{EE}^2 V + \frac{1}{2} \beta_E E^2 \right\}, \quad (t, E) \in [0, T] \times \mathbb{R},$$

with terminal condition $V(T, E) = \frac{1}{2} \gamma_T E^2$.

The minimiser inside the brackets is found by setting the first-order condition to zero. Differentiating with respect to τ ,

$$\alpha \tau - b \partial_E V = 0 \implies \tau^*(t, E) = \frac{b}{\alpha} \partial_E V(t, E), \quad (3.1)$$

provided the resulting expression is non-negative (we return to this in Section 5). Substituting τ^* back into the bracketed expression yields the *reduced* HJB equation



$$-\partial_t V + \rho V = a \partial_E V - \frac{b^2}{2\alpha} (\partial_E V)^2 + \frac{1}{2} \sigma_E^2 \partial_{EE}^2 V + \frac{1}{2} \beta_E E^2, \quad (3.2)$$

which is nonlinear in $\partial_E V$ but quadratic in E and $\partial_E V$ jointly. This quadratic structure is what allows the closed form.

3.2 Quadratic ansatz

The terminal condition is quadratic in E ; the source term $\frac{1}{2} \beta_E E^2$ is quadratic in E ; the nonlinearity is quadratic in $\partial_E V$. We therefore seek a solution of the form

$$V(t, E) = c_2(t) E^2 + c_1(t) E + c_0(t), \quad (3.3)$$

with coefficients $c_2, c_1, c_0 : [0, T] \rightarrow \mathbb{R}$ to be determined. Computing the partial derivatives gives $\partial_E V = 2c_2 E + c_1$ and $\partial_{EE}^2 V = 2c_2$, and $\partial_t V = \dot{c}_2 E^2 + \dot{c}_1 E + \dot{c}_0$. Substituting into the reduced HJB and matching powers of E , we obtain a system of three coupled ODEs in time.

Theorem 1 (Riccati system). Let $\kappa := 2b^2/\alpha$. In time-to-horizon $s = T - t$ the coefficients (c_2, c_1, c_0) satisfy

$$\begin{aligned} \frac{dc_2}{ds} &= \frac{1}{2} \beta_E - \rho c_2 - \kappa c_2^2, & c_2(0) &= \frac{1}{2} \gamma_T, \\ \frac{dc_1}{ds} &= -(\rho + \kappa c_2) c_1 + 2a c_2, & c_1(0) &= 0, \\ \frac{dc_0}{ds} &= -\rho c_0 - \frac{\kappa}{4} c_1^2 + a c_1 + \sigma_E^2 c_2, & c_0(0) &= 0. \end{aligned}$$

The system is well-posed on $[0, T]$: the right-hand side is locally Lipschitz, and the explicit closed form for c_2 below confirms existence and uniqueness globally.

Proof. Substitute the quadratic ansatz into the reduced HJB equation and match the coefficients of E^2 , E and E^0 ; the term-by-term computation is in Appendix B. $\square$

3.3 Closed form for c_2

The equation for c_2 is a scalar Riccati ODE with constant coefficients. Its right-hand side is a quadratic in c_2 with discriminant $\rho^2 + 2\kappa\beta_E > 0$, so it has two real roots.

Proposition 3 (Riccati roots and decay rate). The two real roots of the Riccati right-hand side are

$$c_2^\infty = \frac{-\rho + \sqrt{\rho^2 + 2\kappa\beta_E}}{2\kappa} > 0, \quad c_2^- = \frac{-\rho - \sqrt{\rho^2 + 2\kappa\beta_E}}{2\kappa} < 0,$$

and the horizon-decay rate is

$$\lambda := \sqrt{\rho^2 + 2\kappa\beta_E}.$$

The closed-form solution with terminal value $c_2(T) = \frac{1}{2} \gamma_T$ is

$$c_2(t) = \frac{c_2^\infty (c_2^T - c_2^-) - c_2^- (c_2^T - c_2^\infty) e^{-\lambda(T-t)}}{(c_2^T - c_2^-) - (c_2^T - c_2^\infty) e^{-\lambda(T-t)}}, \quad c_2^T := \frac{1}{2} \gamma_T.$$



Proof. The right-hand side of the c_2 equation factors as $-\kappa(c_2 - c_2^\infty)(c_2 - c_2^-)$, and $\kappa(c_2^\infty - c_2^-) = \lambda$. The Möbius variable $u := (c_2 - c_2^\infty)/(c_2 - c_2^-)$ therefore satisfies the linear equation $du/ds = -\lambda u$, so $u(s) = u(0)e^{-\lambda s}$; solving for c_2 with $u(0) = (c_2^T - c_2^\infty)/(c_2^T - c_2^-)$ gives the formula. Linear-quadratic Riccati equations of this kind are treated in [10], Chapter 6. $\square$

Remark 2 (Steady-state limit). As $T \rightarrow \infty$ with the terminal value held fixed, $e^{-\lambda(T-t)} \rightarrow 0$ for every fixed t , and $c_2(t) \rightarrow c_2^\infty$ — the infinite-horizon steady state. The half-life of the horizon correction is $\log 2/\lambda$; at the base calibration of Section 7 it is 22.2 years, comparable to the horizon $T = 25$ used here. The horizon correction is therefore not negligible at this time scale.

The remaining coefficients c_1 and c_0 do not admit closed-form solutions in elementary functions for general parameters, but they are linear ODEs with known forcing and are computed by quadrature. We defer the numerical recipe to Section 6.

4. Verification theorem

The closed form of Section 3 produces a candidate value function and a candidate optimal feedback law. The verification theorem confirms that the candidate value function coincides with the true value function and that the candidate feedback law is in fact optimal in the class of admissible controls. Without verification, all that has been shown is that a smooth solution of the HJB equation exists; the link to the original control problem requires an independent argument.

4.1 Statement

Theorem 2 (Verification). Let $\widehat{V} \in C^{1,2}([0, T] \times \mathbb{R}) \cap C([0, T] \times \mathbb{R})$ satisfy the reduced HJB equation of Section 3 with terminal condition $\widehat{V}(T, E) = \frac{1}{2}\gamma_T E^2$. Define the candidate feedback law

$$\hat{\tau}(t, E) := \frac{b}{\alpha} \partial_E \widehat{V}(t, E),$$

and assume the SDE

$$d\widehat{E}_t = (a - b\hat{\tau}(t, \widehat{E}_t)) dt + \sigma_E dW_t, \quad \widehat{E}_0 = E,$$

has a unique strong solution that satisfies Assumption 2. Assume further the polynomial-growth bound

$$|\widehat{V}(t, E)| + |\partial_E \widehat{V}(t, E)| \leq K(1 + E^2) \quad \text{for some } K > 0, \text{ uniformly in } t.$$

Then

$$\widehat{V}(t, E) = V(t, E) \quad \text{for all } (t, E) \in [0, T] \times \mathbb{R},$$

and $\hat{\tau}$ is an optimal feedback control.

The proof, given in full in Appendix A, proceeds by Itô's formula applied to $e^{-\rho s} \widehat{V}(s, E_s)$ under an arbitrary admissible control, followed by a martingale argument and a compar-



ison of the resulting integral inequality against the cost functional of Definition 2. The polynomial-growth bound and the integrability of Definition 1 together ensure that the stochastic integral is a true martingale rather than a local martingale.

4.2 Application to the quadratic ansatz

The candidate $\widehat{V}$ constructed in Section 3 satisfies the hypotheses of Theorem 2. We record this as a corollary.

Corollary 1 (Optimal value function and policy). Let (c_2, c_1, c_0) solve the Riccati system of Theorem 1. Then

$$V(t, E) = c_2(t) E^2 + c_1(t) E + c_0(t),$$

and the optimal carbon tax is the affine feedback law

$$\tau^*(t, E) = \frac{b}{\alpha} (2 c_2(t) E + c_1(t)).$$

Proof. The polynomial-growth bound holds because c_2, c_1, c_0 are bounded on $[0, T]$ (Theorem 1 plus Proposition 3). The closed-loop SDE under $\hat{\tau}$ becomes $dE_t = (a - b\hat{\tau}(t, E_t)) dt + \sigma_E dW_t$; the drift is affine in E with bounded, time-continuous coefficient, so a unique strong solution exists with $\mathbb{E} \sup_{t \leq T} E_t^2 < \infty$ by standard SDE theory [10], Chapter 1. The hypotheses of Theorem 2 are satisfied, and the conclusion follows. $\square$

Remark 3 (Why verification matters). One could in principle attempt to compute the value function directly by simulating many candidate controls and taking the infimum of their cost. For the LQ problem this is unnecessary because the closed form is available, but the verification theorem is what licenses the substitution: it certifies that the smooth solution of the HJB is the value function, rather than merely a solution of the equation. This certification is independent of the closed-form structure and would carry over to model extensions in which only a numerical solution is available, provided the regularity hypotheses are satisfied.

Remark 4 (Non-negativity of the optimal tax). The candidate feedback $\hat{\tau}(t, E) = (b/\alpha)(2c_2E + c_1)$ may in principle take negative values. With $a \geq 0$ and $E \geq 0$ we have $c_2(t), c_1(t) \geq 0$ for all $t \in [0, T]$ (Proposition 4), so non-negativity holds on the policy-relevant region; Section 8 confirms it on the plotted grid. The closed form should be read as the *interior* optimal policy; the case in which the unconstrained optimum dips below zero is not reached at the base calibration.

5. Properties of the optimal policy

Corollary 1 yields the optimal feedback in closed form, $\tau^*(t, E) = (b/\alpha)(2c_2(t) E + c_1(t))$. This section extracts the properties of τ^* that are policy-relevant: the affine structure, the infinite-horizon limit, the sign and magnitude of the horizon correction, and the comparative statics in the structural parameters.



5.1 Affine feedback and steady-state limit

Proposition 4 (Affine structure). The optimal carbon tax is affine in the cumulative-emissions state E , with time-varying slope and intercept:

$$\tau^*(t, E) = \frac{b}{\alpha} c_1(t) + \frac{2b}{\alpha} c_2(t) E.$$

Both the slope $(2b/\alpha)c_2(t)$ and the intercept $(b/\alpha)c_1(t)$ are positive for all $t \in [0, T]$ under Assumption 1 with $a \geq 0$.

Proof. Affinity is immediate from Corollary 1 and the linearity of the Riccati ansatz in E . Positivity of c_2 on $[0, T]$ follows from the closed form of Proposition 3: $c_2^T = \frac{1}{2}\gamma_T > 0$ and $c_2^\infty > 0$, and the Möbius interpolation between the two stays strictly positive. Positivity of c_1 follows from the linear ODE of Theorem 1: with $a \geq 0$, $c_2 > 0$, and $c_1(T) = 0$, backward integration produces $c_1(t) \geq 0$ for all t . $\square$

Proposition 5 (Steady-state Pigouvian limit). Let $c_1^\infty = 2a c_2^\infty / (\rho + \kappa c_2^\infty)$ be the steady-state solution of the c_1 ODE. Then

$$\lim_{T \rightarrow \infty} \tau^*(0, E) = \tau_\infty^{\text{Pig}}(E) := \frac{b}{\alpha} (2c_2^\infty E + c_1^\infty),$$

uniformly on compact subsets of $\mathbb{R}$.

Proof. As $T \rightarrow \infty$ with t fixed, $e^{-\lambda(T-t)} \rightarrow 0$, so the explicit formula in Proposition 3 yields $c_2(t) \rightarrow c_2^\infty$. The c_1 ODE then becomes a linear, time-invariant ODE with constant forcing $2a c_2^\infty$, whose stable fixed point is exactly c_1^∞ . Substituting into the formula of Corollary 1 gives the claim. $\square$

5.2 Sign and magnitude of the horizon correction

The horizon correction at state E and time $t = 0$ is

$$\Delta(E) := \tau_\infty^{\text{Pig}}(E) - \tau^*(0, E) = \frac{b}{\alpha} [2(c_2^\infty - c_2(0)) E + (c_1^\infty - c_1(0))]. \quad (5.1)$$

We show that Δ is strictly positive on $E > 0$ whenever the terminal coefficient lies below an explicit bound, which the base calibration satisfies.

Lemma 1 (Positive horizon correction). Under Assumption 1 with $a \geq 0$ and $\gamma_T < 2\beta_E/(\rho + \lambda)$, we have $c_2(0) < c_2^\infty$ and $c_1(0) < c_1^\infty$, so $\Delta(E) > 0$ for all $E > 0$.

Proof. Since $\lambda^2 - \rho^2 = 2\kappa\beta_E$, the positive root is $c_2^\infty = (\lambda - \rho)/(2\kappa) = \beta_E/(\rho + \lambda)$, so the hypothesis reads $c_2^T = \frac{1}{2}\gamma_T < c_2^\infty$. Then $u(0) < 0$ in the proof of Proposition 3, and $u(s) = u(0)e^{-\lambda s} < 0$ gives $c_2 < c_2^\infty$ on $[0, T]$. In the c_1 equation the moving target $f(c_2) := 2ac_2/(\rho + \kappa c_2)$ is increasing in c_2 , so $f(c_2(s)) < f(c_2^\infty) = c_1^\infty$. Starting from $c_1 = 0 < c_1^\infty$, c_1 cannot reach c_1^∞ : at $c_1 = c_1^\infty$ its derivative $(\rho + \kappa c_2)(f(c_2) - c_1^\infty)$ is negative. Hence $c_1(0) < c_1^\infty$, and both terms of $\Delta(E)$ are positive for $E > 0$. At the base calibration $\gamma_T = 2.89 \times 10^{-7}$, well below the bound 3.14×10^{-6} . $\square$



Remark 5. The positivity of the horizon correction is the central qualitative fact of the paper: a finite-horizon planner whose terminal valuation lies below the stationary one sets a *lower* carbon tax than the infinite-horizon planner. The terminal penalty $\frac{1}{2}\gamma_T E^2 = 3 \cdot \frac{1}{2}\beta_E E^2$ prices the stock left at T as three years of undiscounted flow damage, whereas the stationary coefficient prices it as $2c_2^\infty/\beta_E \approx 33$ years. At the base calibration the correction is 55 % of the stationary benchmark at E_{ref} ; with $\gamma_T = 2c_2^\infty$ (the stationary curvature at the horizon, with $c_1(T) = 0$ still) it is 16 %, and with $\gamma_T = 0$ it is 59 %.

5.3 Comparative statics

We collect the signed responses of $\tau^*(0, E_{\text{ref}})$ to the structural parameters; sensitivity magnitudes are reported in Section 9.

Proposition 6 (Comparative statics). Under Assumption 1 with $a \geq 0$ and $E_{\text{ref}} > 0$,

$$\frac{\partial \tau^*(0, E_{\text{ref}})}{\partial \beta_E} > 0, \quad \frac{\partial \tau^*(0, E_{\text{ref}})}{\partial \gamma_T} > 0, \quad \frac{\partial \tau^*(0, E_{\text{ref}})}{\partial \rho} < 0, \quad \frac{\partial \tau^*(0, E_{\text{ref}})}{\partial \sigma_E} = 0,$$

and, when $\gamma_T < 2\beta_E/(\rho + \lambda)$, $\partial \tau^*(0, E_{\text{ref}})/\partial T > 0$.

Proof. The right-hand side of the c_2 equation increases in β_E and decreases in ρ , and $c_2(0) = \frac{1}{2}\gamma_T$ increases in γ_T ; by the comparison theorem for scalar ODEs, $c_2(s)$ increases in β_E and γ_T and decreases in ρ . The right-hand side of the c_1 equation decreases in ρ and increases in c_2 , since its c_2 -derivative is $2a - \kappa c_1$ and c_1 stays below $2a/\kappa$: it starts at 0, and at $c_1 = 2a/\kappa$ its derivative is $-2a\rho/\kappa < 0$. Comparison passes the signs to c_1 . For the horizon, $\tau^*(0, E)$ is the closed form evaluated at $s = T$; under the hypothesis c_2 increases in s (proof of Lemma 1), and c_1 increases in s because it starts below its rising target $f(c_2(s))$ and cannot cross it. The volatility enters only the c_0 equation, which does not affect $\partial_E V$: the tax is certainty-equivalent and $\partial \tau^*/\partial \sigma_E = 0$ exactly. Section 9 checks every sign numerically over the ranges of Table 4. $\square$

Remark 6 (Horizon versus discount rate). Over the ranges of Table 4 the horizon moves the initial tax by \$2.66 (T from 10 to 50 years) and the discount rate by \$1.37 (ρ from 0.01 to 0.05); the damage coefficient, over the likely range of the TCRE, moves it by \$3.80. These are statements about the ranges chosen, not about the parameters in general.

6. Numerical algorithms

This section gives the three algorithms used in the paper: the closed-form Riccati integrator that produces the analytic value function and feedback law, an explicit backward-time finite-difference scheme for the reduced HJB PDE that provides an independent verification, and a forward Monte Carlo simulator under the optimal feedback that produces representative state trajectories.



6.1 Closed-form Riccati integration

The Riccati coefficient c_2 is computed analytically from Proposition 3. The coefficients c_1 and c_0 are computed by one classical fourth-order Runge–Kutta (RK4) march of their equations in time-to-horizon $s = T - t$, from $s = 0$ to $s = T$, with c_2 taken from the closed form at every stage; their values at any t are read from the march. Algorithm 1 gives the recipe; the inputs are the parameter tuple $(\alpha, b, \beta_E, \gamma_T, \rho, a, \sigma_E, T)$, the query times $\{t_k\}_{k=0}^N$, and the number of steps N_q .

Algorithm 1 — Riccati closed-form integration

Closed-form evaluation of c_2 at every query time, and one RK4 march of the c_1 and c_0 equations over $s \in [0, T]$.

```

1 kappa    <- 2*b^2/alpha
2 lambda   <- sqrt(rho^2 + 2*kappa*beta_E)
3 c2_inf   <- (-rho + lambda) / (2*kappa)
4 c2_minus <- (-rho - lambda) / (2*kappa)
5 c2_T     <- 0.5*gamma_T
6 c2s(s)   <- (c2_inf*(c2_T - c2_minus) - c2_minus*(c2_T -
    c2_inf)*exp(-lambda*s))
    / ((c2_T - c2_minus) - (c2_T - c2_inf)*exp(-lambda*s))
7
8 f(s, y)  <- ( -(rho + kappa*c2s(s))*y1 + 2*a*c2s(s),
    -rho*y0 - 0.25*kappa*y1^2 + a*y1 + sigma_E^2*c2s(s) ) # y
    = (y1, y0) = (c1, c0)
9
10 h <- T / N_q; y[0] <- (0, 0)
11 for i = 0..N_q - 1:
12     s <- i*h
13     k1 <- f(s, y[i]);          k2 <- f(s + h/2, y[i] + (h/2)*k1)
14     k3 <- f(s + h/2, y[i] + (h/2)*k2); k4 <- f(s + h, y[i] + h*k3)
15     y[i+1] <- y[i] + (h/6)*(k1 + 2*k2 + 2*k3 + k4)
16 for each t in {t_k}:
17     c2(t) <- c2s(T - t); (c1(t), c0(t)) <- y interpolated at s = T - t
18 return {(c2(t_k), c1(t_k), c0(t_k))}
```

Proposition 7 (Convergence of Algorithm 1). The values of c_2 are exact; at the nodes of the march the errors in c_1 and c_0 are $O(N_q^{-4})$ under Assumption 1.

Proof. The right-hand side of the (c_1, c_0) system is smooth in s (through the closed form of c_2) and in (c_1, c_0) , and Lipschitz on bounded sets, uniformly on $[0, T]$ by the boundedness of c_2 (Proposition 3). The classical fourth-order Runge–Kutta method then has global error $O(h^4)$ with $h = T/N_q$. $\square$

6.2 Explicit finite-difference scheme

The reduced HJB PDE is parabolic with quadratic nonlinearity in $\partial_E V$. An explicit backward-time scheme is direct to implement and serves as an independent check on the closed form; it also generalises directly to non-quadratic damage specifications in which the closed form is unavailable.



The scheme discretises V on a tensor grid $\{(t_k, E_j)\}$, with $t_k = k \Delta t$, $k = 0, 1, \dots, K = T/\Delta t$, and $E_j = E_{\min} + j \Delta E$, $j = 0, 1, \dots, J$. The first and second spatial derivatives are approximated by centred differences in the interior and one-sided differences at the boundary. The time march is backward in t , starting from the terminal condition.

Two CFL conditions must be respected simultaneously. The parabolic CFL $\Delta t \leq \Delta E^2 / \sigma_E^2$ controls the explicit diffusion step; the advection CFL $\Delta t \leq \Delta E / |a|$ controls the drift step. When a user-supplied outer time step violates either, the scheme substeps internally to the smaller of the two CFL bounds.

Algorithm 2 — Explicit backward-time FDM with CFL substepping

Backward-time march from the terminal condition $V(T, E) = \frac{1}{2} \gamma_T E^2$ under an inner step satisfying both the parabolic and the advection CFL bounds.

```

1 K      <- round(T / Deltat)
2 J      <- round((E_max - E_min) / DeltaE)
3 E[j]   <- E_min + j*DeltaE for j = 0..J
4 cfl_diff <- DeltaE^2 / sigma_E^2
5 cfl_adv  <- DeltaE / |a|
6 Deltat_in <- 0.5*min(cfl_diff, cfl_adv)
7 n_sub   <- ceil(Deltat / Deltat_in)
8 Deltat_in <- Deltat / n_sub
9 V[K][j] <- 0.5*gamma_T*E[j]^2 for j = 0..J
10 for k = K, K-1, ..., 1:
11   Vk <- V[k]
12   repeat n_sub times:
13     for j = 1..J - 1:
14       dV <- (Vk[j+1] - Vk[j-1]) / (2*DeltaE)
15       ddV <- (Vk[j+1] - 2*Vk[j] + Vk[j-1]) / DeltaE^2
16       F[j] <- -rho*Vk[j] + 0.5*beta_E*E[j]^2 + a*dV -
17         0.5*(b^2/alpha)*dV^2 + 0.5*sigma_E^2*ddV
18       # one-sided extrapolation at boundaries j=0 and j=J
19       dV[0], dV[J] <- (Vk[hyperlink{zzref-1}{1}] - Vk[0])/DeltaE,
20         (Vk[J] - Vk[J-1])/DeltaE
21       ddV[0], ddV[J] <- ddV[hyperlink{zzref-1}{1}], ddV[J-1]
22       F[0], F[J] <- (same RHS with one-sided dV, ddV)
23       Vk <- Vk + Deltat_in*F
24   V[k-1] <- Vk
25 return V

```

Proposition 8 (Consistency of Algorithm 2). At interior nodes the local truncation error of Algorithm 2 is $O(\Delta t_{\text{in}}) + O(\Delta E^2)$. Because the advection bound forces $\Delta t_{\text{in}} \leq \frac{1}{2} \Delta E / a$, refining ΔE refines the inner step in proportion, and the leading term is first order in ΔE .

Proof. Taylor expansion of the centred first and second differences gives $O(\Delta E^2)$ at interior nodes; the explicit Euler step gives $O(\Delta t_{\text{in}})$. The inner step is $\Delta t / n_{\text{sub}}$ with $n_{\text{sub}} = \lceil \Delta t / \Delta t_{\text{stable}} \rceil$, so it never exceeds $\frac{1}{2} \min(\Delta E^2 / \sigma_E^2, \Delta E / a)$. $\square$



Convergence of the explicit scheme for the nonlinear equation is not proved here; the natural framework for it is that of viscosity solutions [12]. Table 2 measures it instead: the error halves with each halving of the grid, and moving the boundaries further out does not change it at the reference state.

6.3 Forward Monte Carlo simulation

Trajectories of the optimal closed-loop state $\{E_t^*\}$ are produced by the Euler–Maruyama scheme applied to the closed-loop SDE.

Algorithm 3 — Forward Monte Carlo under optimal feedback

Euler–Maruyama simulation of the closed-loop SDE with the optimal feedback evaluated from the Riccati coefficients at each step.

```

1 K_sim <- round(T / Deltat_sim)
2 rng    <- seed
3 for m = 1..N_paths:
4   E[m][0] <- E_0
5   for k = 0..K_sim - 1:
6     t      <- k*Deltat_sim
7     tau     <- max((b/alpha)*(2*c2(t)*E[m][k] + c1(t)), 0)
8     xi      <- rng.normal(0, 1)
9     E[m][k+1] <- E[m][k] + (a - b*tau)*Deltat_sim +
       sigma_E*sqrt(Deltat_sim)*xi
10 return E, tau paths

```

The path-averaged tax $\bar{\tau}(t) = N_{\text{paths}}^{-1} \sum_m \tau^*(t, E_t^{(m)})$ is the *expected* optimal tax path under the model’s emissions noise; Figure 3 shows its 95 % band.

7. Calibration

This section states the base calibration and, for each entry, whether it is taken from a source or set as an illustrative anchor. Table 1 lists the values. The state is cumulative CO₂, whose near-proportional link to warming is the finding of [13]. Costs are measured in per cent of world output per year, so V is a discounted sum of such percentages; no dollar value of world output enters the model.



Table 1: Base calibration. Three entries are taken from sources; the others are illustrative, DICE-style choices, stated as such and varied in Section 9.

Parameter	Symbol	Value	Units	Source
Cumulative CO ₂ , 1850–2019	E_{ref}	2 390	GtCO ₂	IPCC AR6 [14], SPM D.1.2 (± 240)
Warming per 1 000 GtCO ₂ (TCRE)	$\zeta \times 10^3$	0.45	°C	IPCC AR6 [14], SPM D.1.1
Baseline emissions	a	40	GtCO ₂ /yr	Global Carbon Budget 2023 [15]
Damage coefficient	β_E	9.62×10^{-8}	% of output per (GtCO ₂) ²	illustrative anchor
Output-cost coefficient	α	8.0×10^{-4}	% of output per (\$/tCO ₂) ²	illustrative anchor
Tax elasticity of emissions	b	0.40	GtCO ₂ /yr per \$/tCO ₂	illustrative
Terminal coefficient	$\gamma_T = 3\beta_E$	2.89×10^{-7}	% of output per (GtCO ₂) ²	illustrative
Discount rate	ρ	0.03	yr ⁻¹	illustrative
Emissions volatility	σ_E	5	GtCO ₂ yr ^{-1/2}	illustrative
Horizon	T	25	yr	base case

Cumulative emissions and warming. The state E is cumulative anthropogenic CO₂ since 1850. The IPCC Sixth Assessment Report puts the 1850–2019 total at $2\,390 \pm 240$ GtCO₂ and the transient climate response to cumulative emissions (TCRE) at 0.45 °C per 1 000 GtCO₂, likely range 0.27 – 0.63 °C [14]. We take $E_{\text{ref}} = 2\,390$ GtCO₂ and map the stock to CO₂-induced warming by $\Delta T = \zeta E$, which gives 1.08 °C at E_{ref} ; non-CO₂ forcing is outside the model.

Baseline emissions. Total anthropogenic CO₂ emissions, fossil and land use, were 40.7 ± 3.2 GtCO₂ in 2022 [15]; we take $a = 40$ GtCO₂/yr.

Damage. The flow damage $\frac{1}{2}\beta_E E^2 = d_2(\zeta E)^2$, with $d_2 = \beta_E/(2\zeta^2)$, is quadratic in CO₂-induced warming. We set β_E so that 2 °C of warming, reached at $E = 2/\zeta \approx 4\,444$ GtCO₂, costs 0.95 % of output per year: an illustrative, DICE-style anchor, not a value read from a DICE release. This gives the β_E of Table 1 and a damage flow of 0.27 % of output at E_{ref} . Because $\beta_E \propto \zeta^2$, the TCRE likely range scales β_E by 0.36 to 1.96 ; Section 9 uses that range.

Output cost and abatement. The coefficient α is set so that a $\$50/\text{tCO}_2$ tax costs 1 % of output per year: $\frac{1}{2}\alpha \cdot 50^2 = 1.00$. The elasticity $b = 0.40$ means that a tax of $a/b = \$100/\text{tCO}_2$ would abate the whole baseline flow. Both are illustrative. The tax is costed directly rather than through an abatement-cost curve, so the dollar levels of τ^* are quantities of this model, not estimates of a social cost of carbon.



Terminal condition, discount rate, volatility, horizon. The choice $\gamma_T = 3\beta_E$ values the stock left at T at three years of flow damage; it drives the size of the horizon correction (Remark 5) and is varied in Section 9. The discount rate $\rho = 0.03$ is illustrative and is swept over $[0.01, 0.05]$. The volatility $\sigma_E = 5 \text{ GtCO}_2 \text{ yr}^{-1/2}$ is illustrative, large enough to make the stochastic terms visible in the finite-difference check; it enters only c_0 , never the tax. The horizon $T = 25$ years places the end of the planning window near the middle of the century.

Derived quantities. The Riccati coefficient is $\kappa = 2b^2/\alpha = 400$; the horizon-decay rate is $\lambda = \sqrt{\rho^2 + 2\kappa\beta_E} = 0.0313 \text{ yr}^{-1}$, with half-life $\log 2/\lambda = 22.2 \text{ yr}$. The Riccati roots are $c_2^\infty = 1.57 \times 10^{-6}$ and $c_2^- = -7.66 \times 10^{-5}$; the stationary intercept is $c_1^\infty = 4.10 \times 10^{-3}$. At $t = 0$: $c_2(0) = 9.11 \times 10^{-7}$, $c_1(0) = 8.73 \times 10^{-4}$, $c_0(0) = 0.2980$. The terminal value $c_2(T) = \frac{1}{2}\gamma_T = 1.44 \times 10^{-7}$ is 0.092 times c_2^∞ , inside the condition of Lemma 1. Section 8 checks these numbers with an independent ODE solve.

8. Numerical results

This section reports the results at the base calibration of Section 7 and checks each headline number a second way. Figures 1 and 2 show the value function and the optimal tax over the (t, E) plane, Figure 3 a closed-loop trajectory, Table 2 the finite-difference check, and Table 3 the second checks.

8.1 Value function and optimal policy contours

The value function $V(t, E) = c_2(t) E^2 + c_1(t) E + c_0(t)$ at the base case is plotted in Figure 1 over $(t, E) \in [0, 25] \times [1\,500, 4\,000]$. At $t = T$ it is the terminal parabola $\frac{1}{2}\gamma_T E^2$; at earlier times it is larger at every E , because more flow damage remains to be incurred, and its curvature $c_2(t)$ rises from $c_2(T) = \frac{1}{2}\gamma_T$ toward c_2^∞ as the remaining horizon lengthens ($c_2(0) = 6.3 c_2(T)$). The right panel compares $V(0, E)$ with the finite-difference solution of Algorithm 2: the largest relative gap on $[1\,800, 3\,200] \text{ GtCO}_2$ is 2.7×10^{-4} . This checks Theorem 1, Proposition 2 and Corollary 1 by a method that does not use the quadratic ansatz.




Figure 1: Value function and its finite-difference check.

Figure 1. Left: value function $V(t, E)$ at the base calibration, in per cent of world output (a discounted sum over years); the dashed line marks E_{ref} . Right: $V(0, E)$ in closed form (solid), the terminal parabola (dashed), and the finite-difference solution of Algorithm 2 at $(\Delta t, \Delta E) = (0.25, 6.25)$ (points).

The optimal tax $\tau^*(t, E) = (b/\alpha)(2c_2(t)E + c_1(t))$ is plotted in Figure 2. It rises in E , because the marginal damage grows with the stock, and it falls in t at fixed E , because $c_2(t)$ and $c_1(t)$ fall toward their terminal values $c_2(T) = \frac{1}{2}\gamma_T < c_2^\infty$ and $c_1(T) = 0$ as the horizon approaches. Both properties hold at every node of the plotted grid, where the smallest tax is \$0.22/tCO₂, so the positivity of Proposition 4 and Remark 4 holds on the policy-relevant region.




Figure 2: Optimal tax and its direct-optimisation check.

Figure 2. Left: optimal tax $\tau^*(t, E)$ in \$/tCO₂, with the deterministic optimal path



from E_{ref} . Right: $\tau^*(0, E)$ (solid) and the stationary benchmark $\tau_{\infty}^{\text{Pig}}(E)$ (dashed); circles: the initial tax from direct optimisation of the discretised problem, with no HJB equation; square: the same at a 600-year horizon, the check of Proposition 5.

At the reference state $E_{\text{ref}} = 2390 \text{ GtCO}_2$ and $t = 0$ the optimal tax is

$$\tau^*(0, E_{\text{ref}}) = \$2.61 / \text{tCO}_2, \quad (8.1)$$

against the infinite-horizon Pigouvian benchmark

$$\tau_{\infty}^{\text{Pig}}(E_{\text{ref}}) = \$5.80 / \text{tCO}_2. \quad (8.2)$$

The horizon correction is 55.0 % of the benchmark: the finite-horizon planner's initial tax is less than half the stationary one. This is the quantitative form of Lemma 1. The levels follow from the illustrative anchors of Section 7; the share of the correction is governed mainly by the terminal coefficient (Remark 5).

Two second checks of these numbers use no part of the closed form. First, solving the full system of Theorem 1, c_2 included, with an adaptive eighth-order Runge–Kutta method at relative tolerance 10^{-12} reproduces c_2 and c_1 on $[0, T]$ to 3×10^{-15} and 7×10^{-14} relative, and $\tau^*(0, E_{\text{ref}})$ to 1×10^{-15} ; this checks Proposition 3 and the march of Algorithm 1 (Proposition 7). A straight-line fit of $\log |u(s)|$, with $u = (c_2 - c_2^{\infty}) / (c_2 - c_2^-)$ from the ODE solution, gives $\lambda = 0.03126$ and a half-life of 22.18 years. Second, minimising the discretised cost directly over piecewise-constant tax paths, as one quadratic programme with neither the HJB equation nor the Riccati system, gives an initial tax of \$2.6135 and \$2.6137 at $N = 1000$ and 2000 steps and \$2.613962 after Richardson extrapolation, against the closed form \$2.613962 (relative gap 4.8×10^{-8}): the feedback of Corollary 1 is the minimiser, as Theorem 2 and Theorem 3 assert. With the horizon stretched to 600 years the same programme gives \$5.80364, against $\tau_{\infty}^{\text{Pig}}(E_{\text{ref}}) = \$5.80363$ (Proposition 5).

8.2 Representative closed-loop trajectory

Figure 3 shows the optimal closed-loop path from $E_0 = E_{\text{ref}}$, deterministically ($\sigma_E = 0$) and as a 95 % band over 20000 simulated paths with $\sigma_E = 5$ (Algorithm 3). The stock rises to 3372 GtCO₂ at T (band 3323 to 3421). Along the path:

- the finite-horizon tax $\tau^*(t, E_t)$ *falls* from

\$2.61/tCO₂ at $t = 0$ to \$0.49/tCO₂ at $t = T$: the boundary condition $c_1(T) = 0$ and the small terminal slope $c_2(T) = \frac{1}{2}\gamma_T$ pull it down as the horizon approaches, and this outweighs the rising stock; - the stationary benchmark $\tau_{\infty}^{\text{Pig}}(E_t)$ *rises* from \$5.80 to \$7.35/tCO₂, tracking the stock alone.

The two curves move in opposite directions over $[0, T]$. The quadratic programme of the previous subsection checks the path: its stock path agrees with the closed-loop path to 1×10^{-6} GtCO₂ at every node, and its last tax, \$0.4875, matches the closed-form tax at the same mid-step time, \$0.4875.




Figure 3: Closed-loop trajectory under optimal feedback.

Figure 3. Left: the optimal tax along the deterministic optimal path (solid) with the 95 % band of the stochastic paths (shaded), the stationary benchmark on the same path (dashed), and the tax path from direct optimisation (circles, at mid-step). Right: cumulative emissions on the deterministic path (solid), the 95 % band (shaded), and the stock path from direct optimisation (circles).

8.3 Finite-difference verification

Table 2 compares $V(0, E_{\text{ref}})$ from Algorithm 2 with the closed form $V(0, E_{\text{ref}}) = 7.58917$ on two domains: the base grid $[E_{\text{ref}} - 800, E_{\text{ref}} + 1600]$ and one twice as wide.

Table 2: Algorithm 2 against the closed form at $(0, E_{\text{ref}})$ (Proposition 8), on the base domain and on one twice as wide; Δt and the inner step in years, ΔE in GtCO₂, $V_{\text{FDM}}(0, E_{\text{ref}})$ in per cent of world output. The closed form is exact up to the RK4 march of Proposition 7, so the error is the scheme's own.

Δt	ΔE	inner step	V_{FDM}	rel. error	rel. error, wide
1	25	0.25	7.59519	7.9×10^{-4}	7.9×10^{-4}
0.5	12.5	0.125	7.59217	4.0×10^{-4}	4.0×10^{-4}
0.25	6.25	0.0625	7.59067	2.0×10^{-4}	2.0×10^{-4}
0.125	3.125	0.03125	7.58992	9.9×10^{-5}	9.9×10^{-5}

The error halves with each halving of the grid (successive ratios 2.005, 1.999, 1.999): first order, as Proposition 8 predicts once the advection bound ties the inner step to ΔE . Doubling the domain does not change the value at E_{ref} , so the boundaries are not a source of error there, and Richardson extrapolation of the two finest rows leaves 5.3×10^{-8} . Near



the edges of the grid the one-sided closure does matter: at $E = 3\,400$ GtCO₂, within one horizon's drift of the upper edge, the relative error is 5.7×10^{-3} , which is why Figure 1 compares the two on $[1\,800, 3\,200]$.

8.4 Shadow value of the stock

The marginal value of the stock at (t, E) is $\partial_E V(t, E) = 2c_2(t)E + c_1(t)$, in per cent of world output (summed over years) per GtCO₂. At $(0, E_{\text{ref}})$ it is 5.23×10^{-3} , of which the slope term $2c_2E$ carries 83 %; the stationary value is 1.16×10^{-2} . The tax is $\tau^* = (b/\alpha) \partial_E V$ exactly, the first-order condition of Section 3, so within the model the planner prices the whole of its own shadow value; the horizon correction is entirely the gap between the two shadow values. Converting either into dollars per tonne would need a dollar value of world output and an abatement-cost curve, which the model does not have.

8.5 Checked twice

Table 3 lists each headline number by two methods. Theorem 2 also says that no admissible policy does better than the feedback. Costing two mis-scaled feedbacks, $\tau = m \tau^*$ with $m = 0.5$ and $m = 1.5$, by the same Monte Carlo gives 7.59618 and 7.59614, above $V(0, E_{\text{ref}})$ by at least 5 072 standard errors, while the optimal feedback's cost sits on $V(0, E_{\text{ref}})$ up to the time-step bias (Theorem 3, Corollary 1, Proposition 1).



Table 3: Each headline number by two methods (Theorem 1, Theorem 2, Proposition 3, Proposition 5). The Monte Carlo row uses 200 000 antithetic paths at $\Delta t = 0.01$; its residual is time-step bias, not noise: at $\Delta t = 0.005$ it is -4.8×10^{-7} , and the extrapolation to $\Delta t \rightarrow 0$ leaves 3.1×10^{-7} .

Quantity	Closed form	Second method	Method	Relative gap
$\tau^*(0, E_{\text{ref}})$ (\$/t)	2.613962	2.613962	ODE solve	1×10^{-15}
$\tau^*(0, E_{\text{ref}})$ (\$/t)	2.613962	2.613962	direct optimisation	4.8×10^{-8}
$\tau_{\infty}^{\text{Pig}}(E_{\text{ref}})$ (\$/t)	5.803626	5.803638	direct optimisation, $T = 600$	2.1×10^{-6}
horizon correction (%)	54.960	54.960	direct optimisation	1.7×10^{-6}
half-life (yr)	22.1763	22.1763	fit to the ODE solution	1×10^{-15}
$V(0, E_{\text{ref}})$	7.589170	7.589170	ODE solve	6×10^{-16}
$V(0, E_{\text{ref}})$	7.589170	7.589160 ± 0.000001	Monte Carlo	-1.3×10^{-6}
$V(0, E_{\text{ref}})$	7.589170	7.589170	finite differences, Richardson	5.3×10^{-8}
E_T on the optimal path (GtCO ₂)	3372.206	3372.206	direct optimisation	9.3×10^{-11}

9. Sensitivity analysis

We report the sensitivity of $\tau^*(0, E_{\text{ref}})$ to one-at-a-time changes of six inputs, each over the range of Table 4, holding the others at the base calibration. Every endpoint is re-solved independently by the ODE method of Section 8; the largest relative difference from the closed form is 1×10^{-15} .



Table 4: Sensitivity ranges (Proposition 6). For comparison, the Stern Review uses a pure rate of time preference of 0.1 %/yr [16], below the bracket for ρ .

Input	Range	Base	Basis of the range
Discount rate ρ	[0.01, 0.05]	0.03	illustrative bracket
Damage coefficient β_E	$[3.46 \times 10^{-8}, 1.89 \times 10^{-7}]$	9.62×10^{-8}	TCRE likely range [14]
Terminal coefficient γ_T	$[1.44 \times 10^{-7}, 5.77 \times 10^{-7}]$	2.89×10^{-7}	half to double
Emissions volatility σ_E	[2.5, 10]	5	half to double
Horizon T	[10, 50] yr	25 yr	illustrative bracket
Reference stock E_{ref}	[2 150, 2 630]	2 390	AR6 likely range [14]

9.1 Sweeps

Figure 4 reports $\tau^*(0, E_{\text{ref}})$ across the ranges of ρ , β_E , γ_T and T . The volatility is omitted because it does not move the tax (Proposition 6), and E_{ref} enters linearly through the affine feedback of Proposition 4.




Figure 4: One-at-a-time sensitivity sweeps for the initial optimal tax.

Figure 4. Sensitivity of $\tau^*(0, E_{\text{ref}})$ to the discount rate, the damage coefficient (over the likely range of the TCRE), the terminal coefficient and the horizon, holding the others at the base case (diamond); circles: the endpoints re-solved by the ODE method.

The signs are those of Proposition 6. Doubling β_E raises $\tau^*(0, E_{\text{ref}})$ by 91 %; over the likely range of the TCRE the tax runs from \$1.09 to \$4.89. Extending the horizon from 25 to 50 years raises it by 54 %, to \$4.03, still 31 % below the stationary benchmark. Raising ρ from 0.01 to 0.05 lowers it by 40 %; at $\rho = 0.015$ it is \$3.19. The terminal coefficient, from half to double, moves it by \$0.34; the volatility does not move it at all.



9.2 Tornado diagram

The tornado diagram in Figure 5 ranks the six inputs by the span of $\tau^*(0, E_{\text{ref}})$ over their ranges.




Figure 5: Tornado diagram of the initial optimal tax over the ranges of Table 4.

Figure 5. Tornado diagram over the ranges of Table 4, sorted by span: the damage coefficient, the horizon, the discount rate, the reference stock, the terminal coefficient, and the volatility (zero span).

Over these ranges the damage coefficient, whose range is the physical uncertainty in the TCRE, and the horizon move the initial tax most; the discount rate comes third. The ranking describes the ranges chosen, not the parameters in general. Pindyck [17] argues that in integrated assessment models the damage function and the discount rate are chosen with little empirical guidance, and the same caution applies to the illustrative entries of Table 1.

10. Discussion

The closed form and its sensitivity analysis support three statements about the model. They are implications of the model at an illustrative calibration, not policy recommendations. We discuss them in turn, then the limitations of the model and the extensions to which it is most naturally suited.

10.1 Implications

The tax path falls toward the horizon. At fixed E the optimal tax falls in t , and along the optimal path the fall outweighs the rising stock (Figure 3). This is a property of the terminal valuation: $c_2(T) = \frac{1}{2}\gamma_T$ is 0.09 of c_2^∞ and $c_1(T) = 0$, so the planner prices less and less of the stock as T approaches. A finite-horizon model whose terminal condition undervalues the residual stock produces a declining tax path for that reason alone.



The horizon correction is large, and the terminal condition sets it. At the base calibration the initial tax is 55 % below the stationary benchmark. The correction depends more on γ_T relative to $2c_2^\infty$ than on the discount rate: with $\gamma_T = 2c_2^\infty$ it falls to 16 %, and with $\gamma_T = 0$ it is 59 %. Taxes from models with different horizons can therefore be compared only together with their terminal conditions.

How the terminal condition enters. The terminal condition $\frac{1}{2}\gamma_T E^2$ does not enter the tax through the discount factor; it enters through the Riccati slope $c_2(t)$ and through the boundary condition of the c_1 equation. At our calibration $c_2(T) < c_2^\infty$, so the slope at the horizon is below its stationary value, and the Möbius solution carries this information backward to $t = 0$ at rate λ . The size of the correction is a function of the *gap* between the terminal proxy and the stationary value, not of the level of either.

10.2 Limitations

The model is intentionally minimal. Five limitations are worth highlighting.

Costed tax, constant output. The tax is costed directly by $\frac{1}{2}\alpha\tau^2$, not through an abatement-cost curve, and world output does not grow: both cost terms are shares of a constant output. The dollar levels of τ^* are therefore model quantities, not estimates of the social cost of carbon.

Scalar state. The state is cumulative emissions; capital, technology, and population are absent. A two-state generalisation to (K, E) , in which capital accumulation interacts with the abatement decision, would follow the same Riccati structure with a 2×2 coefficient matrix; the Möbius solution generalises to the matrix-Riccati case.

Deterministic damage law. β_E and γ_T are constants; damage uncertainty enters only through the sensitivity analysis. A stochastic damage model, for example a Cox–Ingersoll–Ross process on the damage coefficient, would generate a precautionary component in the tax that is absent here; [18] and [6] make this case at length.

Single region. The model has one decision-maker who internalises all emissions. Multi-region versions, in which strategic interactions between blocs add a non-cooperative wedge to each region’s tax, require the methods of differential games rather than HJB optimal control.

Risk-neutral planner. The expected-cost criterion gives equal weight to all states of the world. A risk-sensitive or robust-control criterion would in general change the tax; we do not pursue it here.

10.3 Extensions

The framework is most naturally extended in three directions. (i) Two-region strategic interaction with cooperative and Nash equilibria; the matrix Riccati system generalises, and the gap between cooperative and Nash taxes quantifies the value of coordination. (ii) Climate-sensitivity uncertainty, for instance a random TCRE, which makes β_E a state variable; the value function becomes two-state and the closed-form structure survives if the specification preserves quadratic-in- E costs. (iii) Capital accumulation and growing



output, giving a multi-state model with a matrix Riccati system. Each of these uses the machinery of Sections 3–6 on a larger state vector.

11. Conclusion

We have solved a finite-horizon HJB problem for the optimal carbon tax under LQ-Gaussian primitives, with physical anchors from the IPCC Sixth Assessment Report and an illustrative, DICE-style cost calibration, over a 25-year horizon. The value function and the optimal feedback law admit a closed form through a three-equation Riccati system in which the leading coefficient is fully analytic, with horizon-decay rate $\lambda = \sqrt{\rho^2 + 2\kappa\beta_E} \approx 0.031 \text{ yr}^{-1}$ (half-life $\approx 22 \text{ yr}$).

The principal finding is qualitative and does not depend on the calibration: when the terminal condition values the residual stock below the stationary value, the finite-horizon tax lies below the infinite-horizon one (Lemma 1). At the base calibration the initial tax of \$2.61/tCO₂ is 55 % below the benchmark of \$5.80/tCO₂; the levels follow the illustrative anchors, and the share follows mainly the terminal coefficient. Over the ranges studied, the damage coefficient and the horizon move the initial tax most, and the volatility does not move it at all.

Every headline number was computed twice. The closed form agrees with an independent ODE solve to rounding, with direct optimisation of the discretised problem to 5×10^{-8} , with Monte Carlo up to the time-step bias, and with an explicit finite-difference scheme that converges at first order (relative error 2.0×10^{-4} on the third grid).

The framework extends naturally to multi-state models incorporating capital accumulation, multi-region strategic interaction, and stochastic damage; the Riccati machinery survives in each direction, with the scalar Möbius solution replaced by a matrix-Riccati system. We leave these extensions to subsequent work.

A. Proof of the verification theorem

This appendix gives the proof of Theorem 2, established in three steps: (i) application of Itô's formula to the discounted candidate value function; (ii) showing the resulting stochastic integral is a true martingale via the polynomial-growth bound and Assumption 2; (iii) closing the integral inequality against the cost functional and identifying the equality case with the candidate feedback law.

Assumption 3 (Hypotheses recalled). $\widehat{V} \in C^{1,2}([0, T] \times \mathbb{R}) \cap C([0, T] \times \mathbb{R})$ solves the reduced HJB equation of Section 3 with terminal condition $\widehat{V}(T, E) = \frac{1}{2}\gamma_T E^2$; the candidate feedback $\hat{\tau}(t, E) := (b/\alpha) \partial_E \widehat{V}(t, E)$ generates a unique strong solution $\{\widehat{E}_t\}$ of the SDE $d\widehat{E}_t = (a - b\hat{\tau}) dt + \sigma_E dW_t$ with $\mathbb{E} \sup_{t \leq T} \widehat{E}_t^2 < \infty$; and $|\widehat{V}(t, E)| + |\partial_E \widehat{V}(t, E)| \leq K(1 + E^2)$ uniformly in t .



A.1 Step 1. Itô expansion

Lemma 2 (Itô expansion of discounted candidate). For any admissible control $\tau \in \mathcal{A}$ with associated state $\{E_t\}$ starting at $E_0 = E$, and for any $t \in [0, T]$,

$$e^{-\rho t} \widehat{V}(t, E_t) = \widehat{V}(0, E) + \int_0^t e^{-\rho s} \mathcal{L}^{\tau_s} \widehat{V}(s, E_s) ds + M_t,$$

where

$$\mathcal{L}^{\tau} \widehat{V} := -\rho \widehat{V} + \partial_t \widehat{V} + (a - b\tau) \partial_E \widehat{V} + \frac{1}{2} \sigma_E^2 \partial_{EE}^2 \widehat{V},$$

and $\{M_t\}_{t \in [0, T]}$ is the stochastic integral

$$M_t := \int_0^t e^{-\rho s} \sigma_E \partial_E \widehat{V}(s, E_s) dW_s.$$

Proof. Apply Itô's formula to $f(s, x) = e^{-\rho s} \widehat{V}(s, x)$, then group deterministic and stochastic integrals. The function f is in $C^{1,2}$ by the smoothness hypothesis on $\widehat{V}$; the SDE coefficients are progressively measurable; standard Itô applies. $\square$

A.2 Step 2. Martingale property

Lemma 3. $\{M_t\}$ is a true martingale. Under Assumption 3 and Assumption 2, the process $\{M_t\}_{t \in [0, T]}$ is a square-integrable martingale.

Proof. It suffices to show $\mathbb{E} \int_0^T e^{-2\rho s} \sigma_E^2 (\partial_E \widehat{V}(s, E_s))^2 ds < \infty$. By the polynomial-growth hypothesis, $(\partial_E \widehat{V}(s, E_s))^2 \leq K^2(1 + E_s^2)^2 \leq 2K^2(1 + E_s^4)$. Assumption 2 gives $\mathbb{E} \sup_{s \leq T} E_s^2 < \infty$; the fourth-moment bound $\mathbb{E} \sup_{s \leq T} E_s^4 < \infty$ follows from the same standard SDE estimates [10], Chapter 1, since the drift is bounded affine in E and the diffusion is constant. The integral is then dominated by $2K^2 T(1 + \mathbb{E} \sup_{s \leq T} E_s^4) < \infty$. Square integrability and the Doob martingale property follow. $\square$

A.3 Step 3. Inequality and equality

Theorem 3 (Verification, restated). Under Assumption 3,

$$\widehat{V}(t, E) = V(t, E), \quad (t, E) \in [0, T] \times \mathbb{R},$$

and $\hat{\tau}$ is an optimal feedback control.

Proof. Fix (t, E) and an arbitrary admissible $\tau \in \mathcal{A}$ generating state $\{E_s\}_{s \in [t, T]}$ with $E_t = E$. By Lemma 2 applied on $[t, T]$ with shift of origin,

$$e^{-\rho T} \widehat{V}(T, E_T) = \widehat{V}(t, E) + \int_t^T e^{-\rho(s-t)} \mathcal{L}^{\tau_s} \widehat{V}(s, E_s) ds + M_T - M_t.$$

Adding the deterministic running cost $\int_t^T e^{-\rho(s-t)} L(\tau_s, E_s) ds$ with $L(\tau, E) = \frac{1}{2} \alpha \tau^2 + \frac{1}{2} \beta_E E^2$ to both sides, using the terminal condition $\widehat{V}(T, \cdot) = \Psi(\cdot) = \frac{1}{2} \gamma_T (\cdot)^2$, and rear-



ranging,

$$\widehat{V}(t, E) = \mathbb{E} \left[\int_t^T e^{-\rho(s-t)} L(\tau_s, E_s) ds + e^{-\rho(T-t)} \Psi(E_T) \right] - \mathbb{E} \int_t^T e^{-\rho(s-t)} [\mathcal{L}^{\tau_s} \widehat{V} + L] ds,$$

where the expectation of $M_T - M_t$ vanishes by Lemma 3. The reduced HJB equation gives $\mathcal{L}^{\tau} \widehat{V} + L \geq 0$ pointwise (Section 3, by the infimum definition of the reduced HJB), with equality if and only if $\tau = \hat{\tau}$. Therefore

$$\widehat{V}(t, E) \leq J(t, E; \tau) \quad \forall \tau \in \mathcal{A},$$

with equality at $\tau = \hat{\tau}$. The left-hand side is the candidate; the right-hand side ranges over the cost functional; taking infimum gives $\widehat{V}(t, E) = V(t, E)$, and the candidate feedback $\hat{\tau}$ attains the infimum, so it is optimal. $\square$

Remark 7 (Where the polynomial-growth bound enters). The bound is used twice. In Lemma 3 it ensures $(\partial_E \widehat{V})^2$ is integrable along the state process. In the rearrangement of Step 3, it ensures the boundary term $e^{-\rho T} \widehat{V}(T, E_T) = \frac{1}{2} \gamma_T \mathbb{E}[E_T^2]$ is finite, which is needed to subtract it from both sides. Without the bound, the argument requires localisation by stopping times, with no change in conclusion.

B. Riccati derivation: term-by-term

This appendix derives the system of Theorem 1 from the reduced HJB equation of Section 3 by substituting the quadratic ansatz and matching powers of E . Since $b^2/(2\alpha) = \kappa/4$ with $\kappa = 2b^2/\alpha$, the reduced equation reads

$$-\partial_t V + \rho V = a \partial_E V - \frac{\kappa}{4} (\partial_E V)^2 + \frac{1}{2} \sigma_E^2 \partial_{EE}^2 V + \frac{1}{2} \beta_E E^2, \quad V(T, E) = \frac{1}{2} \gamma_T E^2. \quad (\text{B.1})$$

For $V = c_2 E^2 + c_1 E + c_0$ we have $\partial_E V = 2c_2 E + c_1$, $\partial_{EE}^2 V = 2c_2$, $(\partial_E V)^2 = 4c_2^2 E^2 + 4c_2 c_1 E + c_1^2$ and $\partial_t V = \dot{c}_2 E^2 + \dot{c}_1 E + \dot{c}_0$. Substituting,

$$-\dot{c}_2 E^2 - \dot{c}_1 E - \dot{c}_0 + \rho(c_2 E^2 + c_1 E + c_0) = a(2c_2 E + c_1) - \frac{\kappa}{4}(4c_2^2 E^2 + 4c_2 c_1 E + c_1^2) + \sigma_E^2 c_2 + \frac{1}{2} \beta_E E^2. \quad (\text{B.2})$$

Matching the coefficients of E^2 , E and E^0 :

$$-\dot{c}_2 + \rho c_2 = \frac{1}{2} \beta_E - \kappa c_2^2, \quad -\dot{c}_1 + \rho c_1 = 2ac_2 - \kappa c_2 c_1, \quad -\dot{c}_0 + \rho c_0 = ac_1 - \frac{\kappa}{4} c_1^2 + \sigma_E^2 c_2. \quad (\text{B.3})$$

With $s = T - t$, so that $d/ds = -d/dt$, these are the three equations of Theorem 1; the terminal condition gives $c_2(T) = \frac{1}{2} \gamma_T$ and $c_1(T) = c_0(T) = 0$.

Remark 8. The cross term $4c_2 c_1 E$ of $(\partial_E V)^2$ is the source of the $-\kappa c_2 c_1$ term in the c_1 equation; the constant c_1^2 enters only the c_0 equation, with coefficient $\kappa/4$. The indepen-



dent ODE solve of Section 8 integrates exactly these three equations, c_2 included, and reproduces the closed form.

C. Algorithm pseudocode in full

This appendix collects the full pseudocode for the three algorithms introduced in Section 6, suitable for transcription into any imperative language. The reference implementation in Python is `hjb_carbon.py`; `checks.py` computes every number in the paper and `figures.py` draws the figures.

Algorithm 4 — Riccati closed-form integration — full pseudocode

Closed-form evaluation of c_2 , followed by one RK4 march of the c_1 and c_0 equations in time-to-horizon.

```

1 function riccati_solve(alpha, b, beta_E, gamma_T, rho, a, sigma_E, T,
  {t_k}, N_q):
2     kappa    <- 2*b^2/alpha
3     lambda   <- sqrt(rho^2 + 2*kappa*beta_E)
4     c2_inf   <- (-rho + lambda) / (2*kappa)
5     c2_minus <- (-rho - lambda) / (2*kappa)
6     c2_T     <- 0.5*gamma_T
7
8     # closed form for c2 in time-to-horizon s
9     function c2s(s):
10         e <- exp(-lambda*s)
11         return (c2_inf*(c2_T - c2_minus) - c2_minus*(c2_T - c2_inf)*e)
12             / ((c2_T - c2_minus) - (c2_T - c2_inf)*e)
13
14     # right-hand side of the (c1, c0) equations
15     function f(s, c1, c0):
16         q <- c2s(s)
17         return (-rho + kappa*q)*c1 + 2*a*q,
18             -rho*c0 - 0.25*kappa*c1^2 + a*c1 + sigma_E^2*q)
19
20     # one RK4 march from s = 0 (t = T) to s = T (t = 0)
21     h <- T / N_q
22     y[0] <- (0, 0)
23     for i = 0..N_q - 1:
24         s <- i*h
25         k1 <- f(s, y[i])
26         k2 <- f(s + h/2, y[i] + (h/2)*k1)
27         k3 <- f(s + h/2, y[i] + (h/2)*k2)
28         k4 <- f(s + h, y[i] + h*k3)
29         y[i+1] <- y[i] + (h/6)*(k1 + 2*k2 + 2*k3 + k4)
30
31     for each t in {t_k}:
32         c2[t] <- c2s(T - t)
33         (c1[t], c0[t]) <- y interpolated linearly at s = T - t

```



```
34 return {(c2[t_k], c1[t_k], c0[t_k])}
```

Algorithm 5 — Explicit backward-time FDM with CFL substepping — full pseudocode

Backward Euler in time, centred differences in space, internal substepping to satisfy both the parabolic and the advection CFL.

```
1 function hjb_fdm(alpha, b, beta_E, gamma_T, rho, a, sigma_E, T, E_min,
  E_max, DeltaE, Deltat, eta):
2   K      <- round(T / Deltat)
3   J      <- round((E_max - E_min) / DeltaE)
4   E[j]    <- E_min + j*DeltaE   for j = 0..J
5
6   cfl_diff <- DeltaE^2 / sigma_E^2
7   cfl_adv  <- DeltaE / |a|
8   Deltat_in <- eta * min(cfl_diff, cfl_adv)
9   n_sub    <- ceil(Deltat / Deltat_in)
10  Deltat_in <- Deltat / n_sub
11
12  V[K][j] <- 0.5*gamma_T*E[j]^2   for j = 0..J
13
14  for k = K, K - 1, ..., 1:
15    Vk <- V[k]
16    repeat n_sub times:
17      for j = 1..J - 1:
18        dV[j] <- (Vk[j+1] - Vk[j-1]) / (2*DeltaE)
19        ddV[j] <- (Vk[j+1] - 2*Vk[j] + Vk[j-1]) / DeltaE^2
20        dV[0] <- (Vk[1] - Vk[0]) / DeltaE
21        dV[J] <- (Vk[J] - Vk[J-1]) / DeltaE
22        ddV[0] <- ddV[1]
23        ddV[J] <- ddV[J - 1]
24      for j = 0..J:
25        F[j] <- -rho*Vk[j] + 0.5*beta_E*E[j]^2 + a*dV[j]
26              - 0.5*(b^2/alpha)*dV[j]^2 + 0.5*sigma_E^2*ddV[j]
27      Vk <- Vk + Deltat_in*F
28    V[k - 1] <- Vk
29
30  return V
```

Algorithm 6 — Forward Monte Carlo under optimal feedback — full pseudocode

Euler-Maruyama on the closed-loop SDE with the optimal tax evaluated from the Riccati coefficients at each step.

```
1 function forward_mc(alpha, b, beta_E, gamma_T, rho, a, sigma_E, T, E_0,
  Deltat_sim, N, seed):
2   K_sim <- round(T / Deltat_sim)
3   rng    <- seed
```



```

4   (c2, c1, c0) <- riccati_solve(... , {0, Deltat_sim, 2*Deltat_sim,
... , T}, N_q)
5
6   for m = 1..N:
7     E[m][0] <- E_0
8     for k = 0..K_sim - 1:
9       t <- k*Deltat_sim
10      tau <- max((b/alpha)*(2*c2[t]*E[m][k] + c1[t]), 0)
11      tau_paths[m][k] <- tau
12      xi <- rng.normal(0, 1)
13      E[m][k+1] <- E[m][k] + (a - b*tau)*Deltat_sim +
sigma_E*sqrt(Deltat_sim)*xi
14
15  return E, tau_paths

```

Remark 9. The three algorithms are pipelined: Algorithm 4 supplies the Riccati coefficients $(c_2(t), c_1(t), c_0(t))$ on the time grid; Algorithm 5 is independent and provides the finite-difference check; Algorithm 6 reads the Riccati coefficients to evaluate the optimal feedback law at each Monte Carlo step.

Remark 10 (Reproducibility). The reference implementation fixes every random seed. All numbers reported in Sections 7–9 are written by `checks.py` to one results file, from which the text, the tables and the figures are produced.



References

- [1] Nordhaus, W. D. (2008). *A Question of Balance: Weighing the Options on Global Warming Policies*. Yale University Press.
- [2] Nordhaus, W. D. (2017). Revisiting the social cost of carbon. *Proceedings of the National Academy of Sciences*, 114(7), 1518–1523.
<https://doi.org/10.1073/pnas.1609244114>
- [3] Golosov, M., Hassler, J., Krusell, P., & Tsyvinski, A. (2014). Optimal Taxes on Fossil Fuel in General Equilibrium. *Econometrica*, 82(1), 41–88.
- [4] Nordhaus, W. D. (2018). Projections and Uncertainties about Climate Change in an Era of Minimal Climate Policies. <https://doi.org/10.1257/pol.20170046>
- [5] Cai, Y., & Lontzek, T. S. (2019). The Social Cost of Carbon with Economic and Climate Risks. *Journal of Political Economy*, 127(6), 2684–2734.
<https://doi.org/10.1086/701890>
- [6] van den Bremer, T. S., & van der Ploeg, F. (2021). The Risk-Adjusted Carbon Price. *American Economic Review*, 111(9), 2782–2810.
<https://doi.org/10.1257/aer.20180517>
- [7] Barrage, L. (2020). Optimal Dynamic Carbon Taxes in a Climate-Economy Model with Distortionary Fiscal Policy. *Review of Economic Studies*, 87(1), 1–39.
<https://doi.org/10.1093/restud/rdz055>
- [8] Hassler, J., Krusell, P., & Olovsson, C. (2018). The consequences of uncertainty: climate sensitivity and economic sensitivity to the climate. *Annual Review of Economics*, 10, 189–205. <https://doi.org/10.1146/annurev-economics-080217-053229>
- [9] Hambel, C., Kraft, H., & van der Ploeg, F. (2021). Asset Pricing and Decarbonization: Diversification versus Climate Action. *Working Paper*.
<https://doi.org/10.2139/ssrn.3528239>
- [10] Yong, J., & Zhou, X. Y. (1999). *Stochastic Controls: Hamiltonian Systems and HJB Equations*. Springer.
- [11] Fleming, W. H., & Soner, H. M. (2006). *Controlled Markov Processes and Viscosity Solutions*. Springer.
- [12] Crandall, M. G., Ishii, H., & Lions, P.-L. (1992). User’s guide to viscosity solutions of second-order partial differential equations. *Bulletin of the American Mathematical Society*, 27(1), 1–67. <https://doi.org/10.1090/s0273-0979-1992-00266-5>
- [13] Matthews, H. D., Gillett, N. P., Stott, P. A., & Zickfeld, K. (2009). The proportionality of global warming to cumulative carbon emissions. *Nature*, 459, 829–832. <https://doi.org/10.1038/nature08047>
- [14] IPCC (2021). Summary for Policymakers. *Climate Change 2021: The Physical Science Basis. Contribution of Working Group I to the Sixth Assessment Report of*



the Intergovernmental Panel on Climate Change, 3–32.

<https://doi.org/10.1017/9781009157896.001>

- [15] Friedlingstein, P., O’Sullivan, M., Jones, M. W., Andrew, R. M., & et al. (2023). Global Carbon Budget 2023. *Earth System Science Data*, 15(12), 5301–5369.
<https://doi.org/10.5194/essd-15-5301-2023>
- [16] Stern, N. (2007). *The Economics of Climate Change: The Stern Review*. Cambridge University Press.
- [17] Pindyck, R. S. (2013). Climate Change Policy: What Do the Models Tell Us?. *Journal of Economic Literature*, 51(3), 860–872.
<https://doi.org/10.1257/jel.51.3.860>
- [18] Weitzman, M. L. (2012). GHG Targets as Insurance Against Catastrophic Climate Damages. *Journal of Public Economic Theory*, 14(2), 221–244.



Zamrik Quant Research

The research arm of Zamrik Quant Lab

About the Desk

Optimal Control & Decision Desk

Zamrik Quant Research


What to do under uncertainty: optimal control and HJB, mean-field games, agent-based models, irreversible decisions, betting, lotteries and insurance.

Most of this desk's questions have the same shape: someone must act now, without knowing what comes next, and live with the result. The Optimal Control & Decision desk turns such questions into Hamilton–Jacobi–Bellman equations, optimal stopping and singular control problems, and solves them.

When several parties optimise at once, the problem becomes a game: Dynkin games and wars of attrition between two or three players, and mean field games, in which a large population and its representative member shape each other through coupled Hamilton–Jacobi–Bellman and Fokker–Planck equations. Where the equations give out, agent-based models take over, and the desk says plainly which tool fits which case.


The applications are concrete: when to build an irreversible plant, how a company should pay dividends, how to run a home battery against a time-of-use tariff, when a wind turbine should turn into the wind, how a central bank should answer stagflation, and how much to bet when an edge is small. Its notebooks and essays carry the same reasoning to retirement, market timing, lotteries, insurance and prize competitions, each worked through to a decision a reader could actually make.

The desk is supervised by **Dr. T. Zamrik**, Head of Zamrik Quant Research, PhD, Imperial College London (2011).





Hamiltonian flow of the Lévy-stable density: $\dot{\mathbf{x}} = J \nabla f(\mathbf{x})$, $J = \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix}$, $\varphi(\xi) = \exp[i\mu\xi - |\sigma\xi|^\alpha(1 - i\beta \operatorname{sgn}(\xi)\Phi)]$



ZAMRIK QUANT RESEARCH

The research arm of Zamrik Quant Lab



WP-2026-74091446
zamrik.com