

Fat Tails and Carbon Taxes: A Spectral Resolution of the Climate Economics Debate

Dr. Tamás Nagy

tnagyphd@gmail.com

Draft

Abstract

The central debate in climate economics — between Nordhaus-style moderate carbon pricing and Stern-style aggressive policy — is fundamentally a disagreement about the tail behavior of the damage distribution. We organize this through the Latent Number ρ : heuristically, when $\rho > 2$, thin-tail regimes admit well-posed second-moment reasoning and standard cost-benefit analysis (CBA) is not ruled out by tail divergence alone; when $\rho \leq 2$, fat-tail effects dominate in the Pareto sense (variance diverges as $\rho \downarrow 2$), connecting to Weitzman’s analysis of catastrophic climate risk. In a stylized proportional rescaling recorded in the companion script, policy charges scale like $\rho/(\rho-2)$ relative to a baseline marginal-damage anchor — a device that places a thin-tail calibration (e.g. $\rho \approx 5$, anchored at \$80/ton in the numerics) and a fatter-tail calibration (e.g. $\rho \approx 2.5$, anchored at \$150/ton) on one spectral continuum. Tipping points enter as discrete damage jumps; the machine-checked numerics include a +40% uplift on the \$80/ton anchor. Growth uncertainty is encoded as a strictly positive downward correction to the certainty-equivalent discount rate (without exporting a closed-form $\sigma_g^2/(2\rho_g)$ identity to the formal layer). Seventeen real-arithmetic theorems in `climate_economy_proof.py` are machine-verified in the Platonic kernel with no user-side axioms in that file.

1. Introduction

Climate change policy hinges on a single number: the social cost of carbon (SCC) — the present value of damages from emitting one additional ton of CO₂. Nordhaus’s DICE line of work (2017) is often reported in the tens of dollars per ton in policy discussions; Stern (2006) argues for much higher values (on the order of \$200/ton in influential summaries), demanding aggressive immediate action. The companion formalization fixes a thin-tail *numerical anchor* at \$80/ton for the quadratic-damage block (`dice_baseline_scc`) so that tipping-point and tail scenarios are comparable in one table — not as a claim that DICE uniquely equals \$80.

The factor-of-four disagreement stems from two sources: 1. **Discount rate:** Stern uses $\delta \approx 0.1\%$ (ethical, near-zero pure time preference); Nordhaus uses $\delta \approx 1.5\%$ (market-based). 2. **Tail risk:** Weitzman (2009) showed that if the damage distribution has fat tails, the expected cost of climate change can be *infinite* — the “dismal theorem.”

We show that both disagreements reduce to a single spectral parameter: the Latent Number ρ of the damage and growth distributions.

1.1 Contributions

- **Spectral damage characterization:** $\rho > 2$ as the thin-tail regime for second-moment reasoning; $\rho \leq 2$ as the variance-divergence wall (Pareto-type tail index), motivating precautionary emphasis (Weitzman).
- **Unified scaling device:** proportional rescaling $\propto \rho/(\rho-2)$ in the convergent regime, nesting high- ρ and lower- ρ calibrations on one continuum in the formal layer.
- **Discount rate resolution:** growth uncertainty lowers effective discount, with magnitude depending on ρ_g .
- **Tipping point premium:** +40% SCC when tipping points (ice sheets, AMOC) are included.
- Machine-verified: 17 theorems in `climate_economy_proof.py`, 0 user axioms in that file (bootstrap real arithmetic only).

Formalization scope. The Platonic theorems are deliberately low-dimensional *templates*: they prove ordered-field consequences from explicit hypotheses (e.g. proportional scaling laws), not the empirical content of integrated assessment models. IAM interpretation is carried in prose; when prose and types disagree, the types win.

2. The DICE Damage Function

2.1 Quadratic Damages

Nordhaus models climate damage as a quadratic function of temperature increase:

$$D(T) = \alpha T^2$$

where T is degrees Celsius above pre-industrial levels and D is damage as a fraction of GDP.

Theorem 1 (Superlinear Damage). *Damage grows faster than linearly in temperature: if $T_2 > T_1 > 0$, then $D_2/D_1 > T_2/T_1$.*

Machine-verified: `climate_economy_proof.py`, theorem `damage_superlinear`. \square

Theorem 2 (Increasing Marginal Damage). *The marginal damage from an additional degree of warming increases with the current temperature level: $D'(T_2) > D'(T_1)$ when $T_2 > T_1$.*

Machine-verified: `climate_economy_proof.py`, theorem `marginal_damage_increasing`. \square

2.2 Social Cost of Carbon

Theorem 3 (SCC Positive). *The social cost of carbon is strictly positive: $SCC = D'(T) \cdot \beta > 0$ where $\beta \in (0, 1)$ is the discount factor.*

Machine-verified: `climate_economy_proof.py`, theorem `scc_positive`. \square

3. Weitzman's Fat Tails

3.1 The Dismal Theorem

Weitzman (2009) observed that if climate sensitivity (the temperature response to doubling CO₂) has a fat-tailed distribution, the expected value of damage can be infinite — making standard CBA meaningless.

Theorem 4 (Fat-Tail Amplification — comparative statics). *Fix a representation in which a variance proxy V and the tail-shape proxy $r = (\rho - 2)/\rho$ satisfy $V_1 r_1 = V_2 r_2$ with $V_i > 0$. If $r_2 < r_1$ (looser tail / smaller ρ in this coordinate), then $V_2 > V_1$. Thus moving ρ toward the critical value from above raises the variance proxy; the Pareto identity $\text{Var}(X) \propto 1/(\rho - 2)$ for $\rho > 2$ motivates the rescaling but is not re-derived as a closed form in the script.*

Machine-verified: climate_economy_proof.py, theorem fat_tail_amplification. \square

Theorem 5 (Thin-Tail Regime — reduced form). *In the formal layer, a nonnegative “tail risk” quantity τ and index $\rho > 2$ are tied by $\rho \cdot \tau = C$ with $C > 0$. This is a bookkeeping normalization for the thin-tail regime, not a proof that all moments of a physical damage random variable are finite.*

Machine-verified: climate_economy_proof.py, theorem thin_tails_cba_valid. \square

3.2 The Precautionary Premium

Theorem 6 (Precautionary Premium). *Optimal policy includes a premium beyond expected marginal damage. The premium is inversely proportional to ρ : fatter tails demand more precaution.*

Machine-verified: climate_economy_proof.py, theorem precautionary_premium_inverse_rho. \square

4. The Discount Rate Debate

4.1 The Ramsey Equation

Theorem 7 (Ramsey Discount Rate). *The social discount rate satisfies $r = \delta + \eta g$, where δ is pure time preference, η is the elasticity of marginal utility, and g is the growth rate.*

Machine-verified: climate_economy_proof.py, theorem ramsey_discount_rate. \square

Theorem 8 (Lower Discount \rightarrow Higher SCC). *The SCC is inversely related to the discount rate. Stern’s low δ implies high SCC; Nordhaus’s high δ implies low SCC.*

Machine-verified: climate_economy_proof.py, theorem lower_discount_higher_scc. \square

4.2 Growth Uncertainty and the Effective Discount Rate

Theorem 9 (Uncertainty Lowers Effective Discount). *The formal theorem assumes only that $r_{\text{eff}} = r_{\text{cert}} - c$ for some correction c with $0 < c < r_{\text{cert}}$, hence $r_{\text{eff}} < r_{\text{cert}}$. In continuous-time portfolio models with stochastic growth, a correction of order $\sigma_g^2/(2\rho g)$ is standard folklore; that closed form is **not** exported to the Platonic statement.*

Machine-verified: climate_economy_proof.py, theorem growth_uncertainty_lowers_discount. \square

This block is therefore a *sign* result in the kernel: uncertainty generates a strictly positive wedge lowering the certainty-equivalent discount relative to the perfect-foresight Ramsey benchmark. Any quantitative mapping to Stern-Nordhaus numbers remains outside the verified fragment.

5. Tipping Points

5.1 Irreversible Regime Shifts

Theorem 10 (Tipping Point Damage Jump). *Crossing a tipping point (e.g., ice sheet collapse at $T_c \approx 3^\circ C$) causes a discrete damage jump: $D_{post} = D_{pre} + \Delta_{jump}$.*

Machine-verified: climate_economy_proof.py, theorem tipping_point_damage_jump. \square

Theorem 11 (Tipping Probability Monotone). *The probability of triggering a tipping point increases monotonically with temperature.*

Machine-verified: climate_economy_proof.py, theorem tipping_probability_monotone. \square

Theorem 12 (Tipping Premium). *The optimal carbon tax with tipping points exceeds the baseline: $Tax_{tip} = Tax_{base} + premium$, where the premium reflects the expected damage from irreversible regime shifts.*

Machine-verified: climate_economy_proof.py, theorem tipping_premium_positive. \square

6. The Latent Resolution

6.1 Unified Carbon Tax

Theorem 13 (Carbon Tax Decreasing in ρ — proportional class). *Under the hypothesis $\rho_1 Tax_1 = \rho_2 Tax_2$ with $0 < \rho_1 < \rho_2$, the formal proof yields $Tax_2 < Tax_1$. Interpreting Tax as a carbon-price instrument, fatter tails (smaller ρ in this coupling) correspond to higher prices within this one-parameter family — a comparative static, not a full optimal-tax characterization.*

Machine-verified: climate_economy_proof.py, theorem carbon_tax_decreases_with_rho. \square

6.2 Stern and Nordhaus on One Spectrum

Theorem 14 (Stern-Nordhaus as ρ -Spectrum). *The Stern-Nordhaus disagreement maps to different positions on the ρ spectrum:*

Position	ρ	SCC estimate	Implication
Nordhaus	~5	~\$80/ton	Moderate carbon tax, gradual transition
Intermediate	~3	~\$120/ton	Significant carbon tax, faster action
Stern	~2.5	~\$150/ton	High carbon tax, immediate aggressive action

The empirical question is: what is ρ for climate damages? This is estimable from paleoclimate data, climate model ensembles, and extreme event statistics.

Machine-verified: climate_economy_proof.py, theorem stern_nordhaus_as_rho_spectrum. \square

7. Numerical Calibration

7.1 Scenarios

Scenario	ρ	SCC (\$/ton)	Tipping premium	Total SCC
DICE baseline	~ 5	\$80	+\$32	\$112
Fat-tail	~ 2.5	\$150	+\$60	\$210
Near-critical (illustrative)	~ 2.1	\$400+	+\$160	\$560+

Machine-checked rows: `dice_baseline_scc`, `fat_tail_scc_higher`, and `tipping_scc_premium` fix the \$80 baseline, the strict ordering between the \$80 and \$150 anchors, and the \$80-to-\$112 tipping uplift with a 40% premium parameter. The near-critical row extends the same scaling story **outside** the verified numerics and should be read as qualitative extrapolation only. \square

7.2 Policy Implications

At $\rho = 5$ on the thin-tail anchor: the stylized carbon price including tipping points is \sim \$112/ton in the verified numerics — above many headline DICE-era SCC quotes but still in a moderate band. At $\rho = 2.5$ with the fat-tail anchor and the same tipping arithmetic: \sim \$210/ton, in the ballpark of high-Stern narratives.

Tail index and discounting jointly matter in full IAMs; within this note, the formal layer separates Ramsey identities, a sign result for growth uncertainty, and tail-driven proportional rescaling. The empirical question is how to estimate a damage-tail index alongside discounting — not reducible to ethics alone, but also not collapsed to a single ρ without econometric substance.

8. Discussion

8.1 Relationship to Existing Literature

- **Nordhaus (2017)**: DICE model — the thin-tail *calibration row* here uses $\rho \approx 5$ and the \$80 anchor by construction, not as a uniqueness theorem for DICE.
- **Stern (2006)**: Stern Review — corresponds in this narrative to lower ρ and/or lower effective discounting; the table’s $\rho \approx 2.5$ column is illustrative.
- **Weitzman (2009)**: Catastrophic/climate tail risk — connected here to the $\rho \rightarrow 2^+$ variance wall and precautionary scaling, not to a literal re-proof of the dismal-theorem setup inside Platonic.
- **Pindyck (2013)**: Critique of IAMs — the Latent framework provides the parametric structure Pindyck argued was missing.

8.2 Connection to Companion Papers

- **Optimal taxation** (companion): The carbon tax is a Pigouvian tax — the spectral framework applies to both income and environmental taxation.
- **Het agents** (companion): Climate damages are distributionally unequal; low- ρ wealth distributions amplify the welfare cost.
- **Bounded rationality** (companion): IAMs are computationally bounded; the Latent truncation makes the intractable DICE extensions tractable.

9. Conclusion

The climate economics debate is a spectral question. The Latent Number ρ of the damage distribution organizes when thin-tail CBA is tenable ($\rho > 2$ for finite-variance Pareto tails), how precautionary wedges scale near the $\rho \rightarrow 2^+$ wall, and how to place alternative calibrations on one continuum. Discounting and tail risk remain distinct mechanisms; the formal file separates Ramsey-style identities, a sign result for growth uncertainty, and tail-driven proportional rescaling. By foregrounding tail index alongside discounting, the Latent framing sharpens which parts of the disagreement are genuinely empirical.

During the preparation of this work the author used large language models in order to assist with manuscript drafting, literature search, and coding assistance. After using these tools, the author reviewed and edited the content as needed and takes full responsibility for the content of the published article.

References

- Nordhaus, W.D. (2017). Revisiting the Social Cost of Carbon. *Proceedings of the National Academy of Sciences*, 114(7), 1518–1523.
- Pindyck, R.S. (2013). Climate Change Policy: What Do the Models Tell Us? *Journal of Economic Literature*, 51(3), 860–872.
- Stern, N. (2006). *The Economics of Climate Change: The Stern Review*. Cambridge University Press.
- Weitzman, M.L. (2009). On Modeling and Interpreting the Economics of Catastrophic Climate Change. *Review of Economics and Statistics*, 91(1), 1–19.