Economic Damages of Delayed Climate Action

By Kent D. Daniel, Robert B. Litterman, and Gernot Wagner*

Draft: November 17, 2025

Delayed climate mitigation imposes substantial economic costs by shifting the burden of adjustment onto future generations. We quantify these welfare losses within a climate-economy model that allows us to calculate the deadweight loss (DWL) of underpricing carbon pollution. We simulate policy delay by constraining initial mitigation years and comparing resulting welfare outcomes to an unconstrained baseline. We show analytically that delay raises the required expected entry carbon price: with interior solutions, the expected price maximizes utility and is increasing in the deteriorated climate state at re-entry. Across scenarios, expected re-entry prices are higher by roughly 0.4-0.9% per additional year of delay. The consumption-equivalent DWL even for short delays of 5 to 15 years ranges from 14-32% of first-period consumption, or roughly $8-19 trillion (2020 USD) in one-time compensation. DWLs rise steeply but concavely in the length of delay, reflecting catch-up pricing and abatement once the constraint lifts.

JEL: D62,G12,Q54
Keywords: climate risk, climate policy, cost of carbon

Climate economics is, at its core, the economics of delaying optimal choice. Consequences of delaying climate mitigation are profound and quantifiable, as every year without meaningful reductions in greenhouse gas emissions increases their

* Daniel: Columbia Business School, New York, NY 10027; Litterman: Kepos Capital, New York, NY 10036; Wagner: Columbia Business School, New York, NY 10027; gwagner@columbia.edu. We thank Theo Moers for excellent research assistance.

concentration and commits the world to higher temperature and greater climate damages. From an economic perspective, these delays are an implicit transfer of welfare from future generations to the present, an intertemporal reallocation that is driven not by efficiency but political and institutional frictions. Understanding the dynamics of this delay and quantifying the resulting deadweight loss (DWL) is essential in understanding the true cost of inaction.

Most climate-economic integrated assessment models (IAMs) seek to identify the optimal mitigation path that maximizes intertemporal social welfare under a set of assumed parameters. Yet governments rarely follow the paths that economists identify as socially optimal. Corporate lobbying (Oreskes and Conway, 2011) and other interest-group politics (Mildenberger, 2020), in part via public opinion (Dechezleprêtre et al., 2025; Mildenberger and Tingley, 2019), institutional constraints (Bertram et al., 2024), behavioral barriers (Wagner and Zeckhauser, 2012), and other political economy considerations (Meckling, Sterner and Wagner, 2017; Meckling, 2025) defer action, even when — or perhaps especially when — the social planner's problem is well understood.

The more explicit the attempt at pricing the negative climate externality, the louder are the voices of vested interests lobbying against climate policy. This delay in climate action moves the world further off the efficient frontier, which does not just lead to greater economic damages as reflected in higher social cost of carbon (SCC) calculations (Moore et al., 2024) but measurable DWLs.

Here we examine these costs explicitly. Building on a carbon asset pricing framework (Bauer, Proistosescu and Wagner, 2024), which extends the Epstein–Zin recursive preference structure of Daniel, Litterman and Wagner (2019), we quantify DWLs of delaying optimal policy by comparing carbon price paths under constrained and unconstrained conditions, analyzing sensitivities of various model parameters, including technological progress and learning. In contrast to Daniel, Litterman and Wagner (2019), we further show that the optimal expected carbon price in a delayed scenario will be higher under standard assumptions than in an unconstrained scenario.

We present simple heuristics about the high and quickly accumulating costs of delayed climate action, finding DWLs of delay of between 14-32% of first-period consumption, or \$8-19 trillion (2020 USD), even for relatively short delays of between 5 and 15 years. These numbers are significantly higher than typically calculated, using SCC-based measures. Our optimal carbon price in our base case here is roughly \$200, below, for example, that calculated by Moore et al. (2024)'s "synthetic distribution" with a mean SCC of around \$280. Meanwhile, Bilal and Känzig (2025) calculate an SCC above \$1500 and a welfare cost of (only) around 30%. A key difference to our analysis: we solve for the 'optimal' carbon price by considering marginal disutility of damages, instead of calculating the SCC, the discounted value of the stream of expected future damages.

I. Socio-economic modeling choices

To explore how postponing climate policy affects welfare and the socially optimal carbon price path, we endow a representative agent with recursive Epstein-Zin (EZ) preferences and place it within a binomial decision tree where utility is maximized at each step. Such preferences allow us to disentangle risk over time from risk across states of nature. This distinction follows Epstein and Zin (1989, 1991), with a long history in financial economics, and a more recent one in modeling the financial implications of climate risks (Ackerman, Stanton and Bueno, 2013; Traeger, 2014; Lemoine and Traeger, 2014).

The representative agent's preferences follow the recursive Epstein–Zin specification,

$$U_t = \Big((1-\beta)c_t^{\rho} + \beta\big[\mathbb{E}_t(U_{t+1}^{\alpha})\big]^{\rho/\alpha}\Big)^{1/\rho}, \tag{1}$$

where $\beta := (1+\delta)^{-1} > 0$ is the one-year discount factor, with $\delta > 0$ denoting the pure rate of time preference; $c_t > 0$ is consumption at time t; $\rho := 1 - 1/\sigma$, where $\sigma > 0$ is the elasticity of intertemporal substitution (EIS); and $\alpha := 1 - \gamma$, where $\gamma > 0$ is the coefficient of relative risk aversion (RA). The term $\mathbb{E}_t(U_{t+1}^{\alpha})$ represents the certainty equivalent of future utility.

When $\alpha = \rho$, that is, when risk aversion and intertemporal substitution coincide, the recursive formulation in Equation (1) collapses to the standard time-additive expected-utility form with constant relative risk aversion.

For the terminal period T, we assume exogenous consumption growth $g > 0$ and define terminal utility as

$$U_T = \left[\frac{1-\beta}{1-\beta(1+g)^{\rho}}\right]^{1/\rho} c_T. \tag{2}$$

This specification cleanly separates two central preference parameters: σ, which governs willingness to substitute consumption over time, and γ, which governs aversion to risk across uncertain future states.[1]

II. Optimization

Following Daniel, Litterman and Wagner (2019) and Bauer, Proistosescu and Wagner (2024), we embed the representative agent in a finite-horizon probability landscape. Our model has six decision times $T_0, \ldots, T_5$ (Figure 1). At every node (t, s) of the tree, the agent maximizes EZ utility in (1) and chooses a node-specific mitigation level $m_{t,s} \in [0, \bar{m}]$ with upper mitigation bound $\bar{m}$, subject to climate dynamics, resource constraints, abatement costs, climate damages, and the technological feasibility of mitigation. Each choice commits the agent to a continuation policy for all downstream nodes in a given branch.

The climate state evolves according to the impulse response function (IRF) of Joos et al. (2013) for atmospheric CO_2 concentration C and the TCRE mapping from cumulative emissions to the global mean surface *temperature anomaly* θ (in °C above preindustrial), see Eqs. (A2)–(A3) in Appendix A.A1. Our im-

[1] See Appendix A.A1 for parameter values used in our main specification.

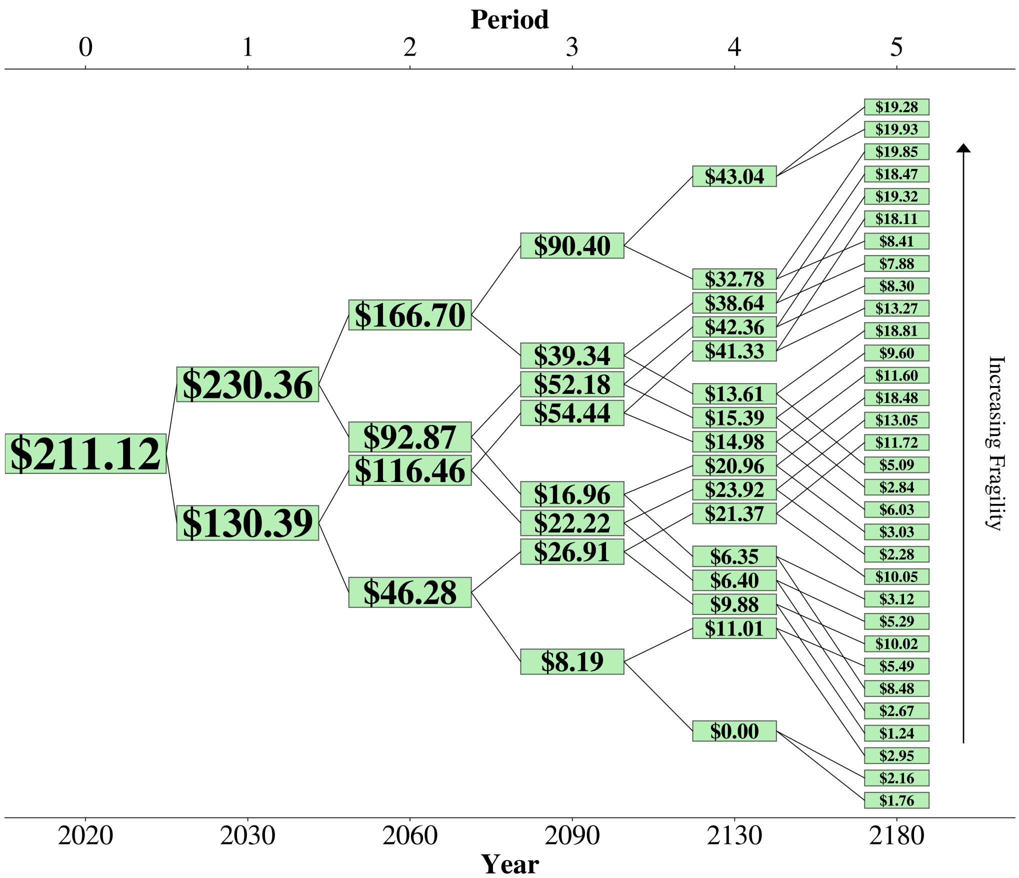

Figure 1. : Optimal price paths (unconstrained baseline)

Note: The binomial non-recombining tree shows the optimal node-level shadow-price trajectory on the stochastic decision tree used in the model. Each node represents a shadow price (in 2020 USD per ton CO_2) at the beginning of the indicated year, conditional on previous realizations of climate and economic uncertainty.

plementation uses a carbon cycle model with persistent parameters from Joos et al. (2013) and uncertain effective TCRE $\lambda_{\text{eff}} \sim \mathcal{N}(0.52, 0.21^2)$. For analytical results we employ a stylized finite-dimensional representation that preserves the key monotonicity properties.

Define at node (t, s) the marginal willingness to pay $\Phi_{t,s}(\theta_{t,s})$ to avoid one additional ton of CO_2 emitted in $[T_t, T_{t+1})$. Mitigation has node-specific marginal cost $\kappa'_{t,s}(m_{t,s}, L_{t,s}) = \partial \kappa_{t,s}(m_{t,s}, L_{t,s})/\partial m_{t,s}$.

At decision time t, let $\mathcal{S}_t$ denote the set of nodes (states) and $\{\pi_{t,s}\}_{s \in \mathcal{S}_t}$ the probabilities of those nodes, conditional on information at the start of period t. For any node-level variable $x_{t,s}$[2], write its cross-node expectation as

$$x_t := \mathbb{E}_t\left[x_{t,S}\right] = \sum_{s \in \mathcal{S}_t} \pi_{t,s} x_{t,s}. \tag{3}$$

We summarize period t by the expected objects $\Phi_t := \mathbb{E}_t[\Phi_{t,S}]$ and $\kappa'_t := \mathbb{E}_t[\kappa'_{t,S}]$ when needed (using (3)).

The climate state at node (t, s) is $(C_{t,s}, \theta_{t,s})$, with $C_{t,s}$ and $\theta_{t,s}$ generated by Eqs. (A2)–(A3). Period-t expectations are $C_t := \mathbb{E}_t[C_{t,S}]$ and $\theta_t := \mathbb{E}_t[\theta_{t,S}]$. On the cost side, mitigation has node-specific marginal cost $\kappa'_{t,s}(m_{t,s}, L_{t,s})$, where $L_{t,s}$ indexes both exogenous technological progress and endogenous learning-by-doing. Past mitigation lowers future costs by shifting down the marginal cost curve.

III. Carbon price paths under delay

To test the cost of delay, we impose a zero-mitigation constraint for the first decision node and vary the length L of that period by shifting the initial decision time between 5, 10, and 15 years: $L \in \{5, 10, 15\}$. Each constrained run is then evaluated against two baseline scenarios, depending on the figure—the optimal expected price at the same L, and one common $L = 10$ baseline. Figure 2 shows the resulting optimal carbon price paths in expectation over time. We here find

[2]Subscript (t, s) denotes a node-level object; subscript t alone denotes its period-t expectation across $s \in \mathcal{S}_t$ with weights $\pi_{t,s}$.

that the carbon price paths in expectation of each delayed scenario lie above the baseline scenario's levels.[3]

If the optimum at time T_t is interior (i.e., $m^\star_{t,s} \in (0, \bar{m})$ for all $s \in \mathcal{S}_t$) and baseline emissions $E_t > 0$, the first–order condition at each node equates the node–specific marginal abatement cost to the node–specific marginal damage:

$$\tau_{t,s} := \kappa'_{t,s}\Big(m^\star_{t,s}(\theta_{t,s}, L_{t,s}),\, L_{t,s}\Big) = \Phi_{t,s}(\theta_{t,s}), \qquad s \in \mathcal{S}_t. \tag{4}$$

where $\kappa'_{t,s}(m_{t,s}, L_{t,s}) = \partial\kappa_{t,s}/\partial m_{t,s}$. We summarize the decision period by the expected carbon price,

$$\tau_t := \mathbb{E}_t[\tau_{t,S}] = \mathbb{E}_t\big[\Phi_{t,S}(\theta_{t,S})\big] = \sum_{s \in \mathcal{S}_t} \pi_{t,s}\, \Phi_{t,s}(\theta_{t,s}), \tag{5}$$

which is the probability–weighted average across all nodes at time t. Learning–by–doing and exogenous technological progress enter through $L_{t,s}$, shifting $\kappa'_{t,s}$ and thereby altering both the node–level prices $\tau_{t,s}$ and the expected price τ_t.

For period-level values, define the expected total mitigation cost and expected marginal abatement cost as $\kappa_t := \mathbb{E}_t[\kappa_{t,S}]$ and $\kappa'_t := \mathbb{E}_t[\kappa'_{t,S}]$. Note that because $m_{t,s}$ is node-specific, $\kappa'_t \neq \partial\kappa_t/\partial m$ in general. At the node level, however, $\kappa'_{t,s}(m_{t,s}, L_{t,s}) = \partial\kappa_{t,s}(m_{t,s}, L_{t,s})/\partial m_{t,s}$.

Delaying mitigation creates a deviation from optimal choice. By not allowing for mitigation for the first L years the world reaches the first unconstrained decision date $T_1 = L$ with a worse climate state: higher cumulative emissions, higher atmospheric CO_2 across persistence reservoirs, and higher temperatures. In that state, marginal damages are higher than they would have been without delay, and the representative agent's marginal willingness to pay to avoid one ton of CO_2, $\Phi_L(\theta_L)$, is higher than it would have been without delay.

Worse still, delaying mitigation also means postponing learning-by-doing. Forc-

[3]See Figure A1 for other outputs, like emissions and economic damages.

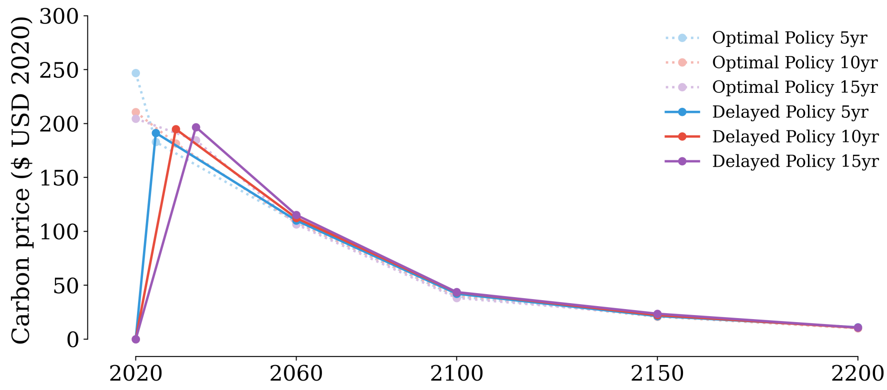

Figure 2. : Optimal CO_2-price paths under delayed policy implementation.

Note: Six decision times are used in all runs, but the length of the first decision period is varied between 5, 10, and 15 years, while all subsequent time steps remain fixed. Each delay scenario (5 yr, 10 yr, 15 yr) is solved as an independent run. We also show a canonical baseline scenario with a decision time at the 10-year step. The resulting carbon price paths show that postponing mitigation leads to a sharp upward adjustment in the first active period, followed by convergence toward the optimal no-delay trajectory.

ing $m_t = 0$ in the early window removes that source of endogenous cost decline. (See Proposition A.1.) As a result, when the world begins to act optimally at $T_1 = L$, it faces both a more fragile climate state and a less mature—more expensive—abatement cost curve. The representative agent's optimal response is therefore to start the policy period with a higher expected carbon price than in the no-delay baseline, and to immediately mitigate more aggressively. This jump in the required expected entry carbon price is structural: it comes from state dependence in climate damages and from foregone technological progress, not from a particular calibration of parameters.

Formally, under standard convexity of abatement costs and monotonicity of damages in the climate state (Appendix A.A1), the optimal expected carbon price at $T_1 = L$ in the delayed scenario, τ_L^{delay}, is weakly higher than the optimal expected carbon price at the same time in the no-delay baseline, τ_L^{base}, with strict inequality whenever the no-mitigation constraint was binding. The state-dependence logic is formalized in Proposition A.1 (Appendix A.A1).

Table 1—: Delay-to-price elasticity at re-entry

L (years)	Year	τ_L^{base}	τ_L^{delay}	$\Delta\log(\tau)$	$\eta(L)$ (%/yr)
5	2025	\$182.97	\$191.35	0.044810	0.896
10	2030	\$181.55	\$194.83	0.070607	0.706
15	2035	\$184.49	\$196.71	0.064150	0.428
				Average	0.677

We summarize the expected price impact with the delay-to-price elasticity

$$\eta(L) := \frac{\partial \log \tau_L^{\text{delay}}}{\partial L}, \tag{6}$$

which is the percent increase in the required expected entry carbon price at the first unconstrained decision time per additional year of forced delay. Because our model is solved at discrete delay lengths only, we estimate $\eta(L)$ by finite differences using $L \in \{5, 10, 15\}$, i.e., $\eta(L) \approx (\log \tau_L^{\text{delay}} - \log \tau_L^{\text{base}})/L$.

Table 1 reports the resulting elasticities. For a 5-year delay, the required re-entry expected price rises from \$182.97 to \$191.35, which corresponds to about an 0.896% higher expected price per year of delay. For 10 and 15 years of delay, the effect remains positive but declines to 0.706% and 0.428% per year, respectively. Averaging across the three scenarios yields an elasticity of about 0.677% per year. A simple log–linear fit of the delayed expected prices on the delay length gives a smaller, global elasticity of 0.276% per year (SE: 0.000487), with a 95% confidence interval of [0.181%, 0.372%]. This confirms Proposition A.1: In every scenario we consider, a binding delay raises the required expected entry carbon price. The elasticity is declining in L, which implies a concave delay-price relationship: early years of inaction are disproportionally costly, as they push the system into a higher-damage (and technically, low-learning) state, while additional years of delay add to the carbon debt at a diminishing marginal rate.

As a robustness check, we also estimate a pooled delay-to-price elasticity across the three scenarios by regressing the delayed re-entry expected price on the length of the delay,

$$\log(\tau_L^{\text{delay}}) = \alpha + \eta^{\text{OLS}} L + \varepsilon_L, \tag{7}$$

which yields $\widehat{\log(\tau_L^{\text{delay}})} = 5.2417 + 0.0028L$, with $R^2 = 0.97$. The slope coefficient $\eta^{\text{OLS}} = 0.0028$ implies that, on average across the 5-15 year range, each extra year of delay raises the required expected entry carbon price by about 0.276% per year. The 95% confidence interval corresponds to [0.181%, 0.372%] per year.

IV. Estimating deadweight losses (DWLs) of delay

To quantify the societal cost of delayed action, we compute the DWL associated with postponing mitigation. Specifically, we determine the additional consumption in the first period required to restore lifetime utility of the representative agent to the level of the unconstrained (baseline) case. Denoting baseline utility at the root as U_0^*, first-period consumption in the delayed scenario as c_0^D, and the expected (certainty-equivalent) future utility as $\text{CE}_1^D := (\mathbb{E}_0[U_1^\alpha])^{1/\alpha}$, we define the consumption-equivalent DWL $\phi \geq 0$ implicitly by

$$U_0^* = \Big((1-\beta)\big((1+\phi)c_0^D\big)^\rho + \beta(\text{CE}_1^D)^\rho\Big)^{1/\rho} \tag{8}$$

Solving for ϕ yields[4]

$$\phi = \left[\frac{(U_0^*)^\rho - \beta(\text{CE}_1^D)^\rho}{(1-\beta)(c_0^D)^\rho}\right]^{1/\rho} - 1, \quad (\rho \neq 0). \tag{9}$$

Applying this metric, we find that the DWL of delayed mitigation rises with the duration of inaction (Table 2). In our main specification, enforced bans on mitigation force higher entry expected carbon prices at $T_1 = L$, which we capture with the delay-to-price elasticity $\eta(L)$; that higher required expected starting price translates directly into a larger consumption-equivalent DWL ϕ. In our main

[4] In the case that $\rho = 0$, we apply the Cobb-Douglas limit as derived in Appendix A.A2

specification, banning mitigation for five years, ten years, and fifteen years produces a DWL of delay of roughly 14%, 22%, and 32% of first-period consumption, respectively. In monetary terms, these correspond to about \$8.3tn, \$12.8tn, and \$18.8tn in one-time global compensation at the start of the policy window. Each additional year of delay raises the DWL by about \$1.05 trillion per year over the 5–15 year range, i.e., roughly 1.8 percentage points of first-period consumption per year. A simple log-log fit implies $\phi(L) \propto L^{0.73}$, indicating sub-linear scaling and modestly declining marginal losses as the delay lengthens.

These DWLs increase with delay length, but not linearly. Longer bans on mitigation give the representative agent at the next decision time a more polluted atmosphere and therefore require a higher expected starting carbon price at $T_1 = L$ under our standard assumptions. The agent then responds by catching up: once mitigation is finally allowed, the optimal policy sharply raises the expected carbon price at $T_1 = L$ relative to the no-delay baseline and moves immediately to very aggressive abatement. This catch-up behavior is economically painful in the short run, which shows up in $\phi(L)$, but it stabilizes the long run by limiting further deterioration of the climate-economy state. Numerically, in our baseline calibration $\phi(L)$ rises quickly between 0 and 10 years of delay and continues to rise thereafter, though at a slower rate (Table 2). In our benchmark runs, then, the exogenously induced delay has a clearly measurable cost: for the 5-15 year range we study, every additional year without mitigation forces the social planner to start the policy period with an expected carbon price between about 0.4 and 0.9 percent higher than it otherwise would have been, with an average of 0.7 percent. The regression-based estimate is smaller because it smooths across the three scenarios, but it preserves the sign and the basic message: delay makes the first feasible expected carbon price higher.

Table 2—: Social cost of delaying climate action under alternative baselines

First period length L (years)	Canonical baseline		Aligned baseline		Difference (p.p.)
	DWL (%)	DWL (2020 USD tn)	DWL (%)	DWL (2020 USD tn)	
5	14.24	8.2	13.00	7.5	-1.24
10	21.82	12.6	21.87	12.7	$+0.05$
15	32.21	18.7	33.10	19.1	$+0.89$

Note: Deadweight loss (DWL) represents the consumption-equivalent compensation required for lifetime utility in the delayed-mitigation scenario to equal that in the corresponding baseline. The canonical baseline fixes the first decision period at 10 years across all runs to enable direct DWL comparisons. The aligned baseline matches each delay scenario to an unconstrained run with the same decision timeline (e.g., 5-year delay vs. 5-year baseline). The minor difference for the 10-year scenario reflects stochastic draws in the model's Monte Carlo simulations. Dollar values are in trillions of 2020 USD.

V. Evaluating parameter importance

Understanding why delay is so costly requires unpacking which structural primitives make the carbon-debt difference $\tau_L^{\text{delay}} - \tau_L^{\text{base}}$ and its elasticity as per Equation (6) large, and therefore drive the DWL penalty ϕ. Figure 3 offers a first look, plotting the DWL of delay against four structural drivers: EIS, PRTP, exogenous technological change, and endogenous learning. Lines show within-delay OLS fits of the expected DWL on the parameter value. The fitted Gaussian curves show the distribution of $\phi(L)$ for the different delay lengths given our parameter space. The central result is that impatience dominates. When societies heavily discount the future, no amount of technological progress or learning can offset the welfare lost from postponing mitigation.

Parameter sensitivities in Table 3 show which economic mechanisms drive the DWL of delaying climate action in a multivariate regression, i.e., they show partial OLS effects within our parameter grid. We estimate these effects over a broad random draw of the model's structural parameters, each sampled independently from its prior probability distribution. This approach ensures that coefficients capture partial effects across the full range of plausible economic and technological states. Each parameter thus maps a structural assumption into an economic

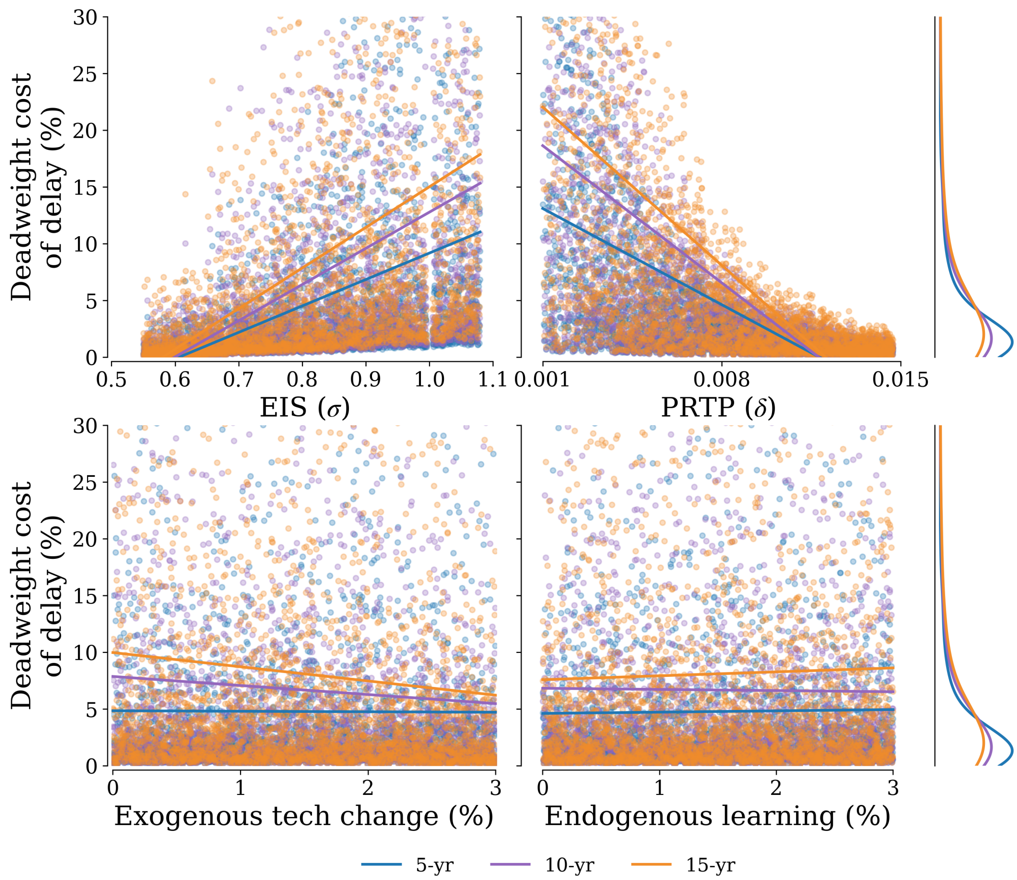

Figure 3. : Variance decomposition of deadweight losses (DWLs) by structural parameters.

Note: The figure truncates the vertical axis at 30% to improve visibility. This range contains approximately 95% of observations.

intuition about risk, time, and technology.

Higher risk aversion (RA) raises the shadow value of insurance against uncertain climate damages. A one-point increase in γ raises the DWL by roughly 0.25 percentage points, or about \$147 billion. This result is highly significant. Agents who are risk-averse value early mitigation more strongly as protection against catastrophic tail outcomes. Nonetheless, RA across states of nature matters less than aversion across time (see EIS): what dominates is not which climate future occurs but how long society waits to act.

The elasticity of intertemporal substitution (EIS) is the single strongest behavioral determinant of delay costs. A 0.1 increase in σ (within the sampled range of 0.55-1.1) raises the DWL by about 5%, or \$3 trillion ($p < 0.001$). In the unconstrained baseline, such a society is willing to sacrifice some near-term consumption (via costly early mitigation) in exchange for much lower climate damages later. A binding delay prevents that optimal intertemporal trade, so the DWL $\phi(L)$ from delay is larger when σ is high.

With a negative coefficient significant at the 10% level, faster exogenous technological change cushions the economy against delay. A one percentage-point increase in the exogenous rate of technological change reduces DWL by roughly 0.7%, or about \$440 billion. Independent innovation lowers future abatement costs and partially offsets delay. When technology improves independently of early action, postponement hurts less because future abatement is cheaper. Conversely, technological stagnation amplifies the cost of delay.

Endogenous technological learning has a positive coefficient that is significant at the 10% level with similar magnitude to exogenous change. A one percentage-point increase in learning intensity raises DWL by about 0.6%, or \$380 billion. The mechanism is path dependence: delaying mitigation slows learning-by-doing, delaying cost reductions and locking in higher future abatement costs. Inaction today undermines tomorrow's productivity gains.

The pure rate of time preference (PRTP, δ) has a large and statistically significant effect. The coefficient implies that raising δ by 1 percentage point lowers DWL by roughly 25.6 percentage points, or $ 15.3 trillion. More impatient societies (higher δ, lower β) care relatively less about distant future damages. In our regressions, this shows up as a lower measured DWL from delay. Conversely, patient societies (low δ) view delay as extremely expensive. This confirms that how we value time, not technology or static risk, is a first-order driver of the DWL variation across scenarios.

The backstop premium, that is, the long-run cost ceiling for zero-carbon technology, is statistically insignificant and economically negligible. A 1% change in the backstop price shifts DWL by less than 0.01%, or about $0.4 billion. In the model, this parameter adds a surcharge to the marginal cost of over-mitigation, that is, for mitigation levels above 100%, corresponding to net carbon removal. This captures the real-world cost gap between eliminating emissions and achieving net-negative emissions through technologies such as direct air capture. Because optimal policy paths in our delay experiments rarely enter the over-mitigation regime, the DWL effects of the premium remain small.

Consumption growth enters negatively: A one percentage-point faster consumption-growth rate reduces the DWL by about 0.6%, or $340 billion. Even though it is not significant, the sign aligns with the theoretical expectation that faster growth decreases the DWL of delay.

Both delay length indicator variables are positive and highly significant. Extending the first decision period from 5 to 10 years raises the DWL by about 3.5%, or $2.1 trillion; extending to 15 years increases them by roughly 8.5%, or $5.1 trillion. The rise is steep but concave, consistent with the model's adaptive catch-up dynamics: once mitigation begins, expected carbon prices jump sharply, partially—but never fully—recovering lost welfare.

Taken together, our analysis shows that the economics of delay is fundamentally about time preference and intertemporal trade-offs: Impatience and substitution, not technology or risk, explain most of the DWL of inaction, underscoring that the true price of delay is paid in lost time.

Table 3—: Regression results: determinants of deadweight loss (DWL) and mitigation delay outcomes

	(1)	(2)	(3)
Const	-0.0457**	-10.5381	-6301.7279
	(0.0199)	(8.3528)	(4994.9064)
RA (γ)	0.0046***	0.2454***	146.7690***
	(0.0002)	(0.0847)	(50.6514)
EIS (σ)	0.4767***	49.7372***	29742.5249***
	(0.0057)	(2.5267)	(1510.9441)
Tech. change (exog.)	-1.4735***	-73.5360*	-43974.0590*
	(0.0935)	(38.5723)	(23066.0244)
Tech. learning (endo.)	0.8882***	63.7416*	38117.0615*
	(0.0883)	(34.3644)	(20549.6858)
PRTP (δ)	-23.5410***	-2556.3668***	-1528691.37***
	(0.2509)	(111.3807)	(66604.948)
log(Backstop premium)	-0.0009	-0.6235	-372.8457
	(0.0020)	(0.8472)	(506.6172)
Cons. growth	-3.5662***	-57.4224	-34338.2212
	(0.1814)	(77.6601)	(46440.2569)
Delay 10	0.0469***	3.4608***	2069.5203***
	(0.0015)	(0.3973)	(237.5642)
Delay 15	0.1088***	8.5027***	5084.5728***
	(0.0021)	(0.7598)	(454.3852)
R^2	0.7639	0.2034	0.2034
Adj. R^2	0.7636	0.2026	0.2026
N	8830	8830	8830

Note: Results from an OLS regression with time fixed effects (Delay 10 and 15). Heteroskedasticity-robust standard errors in parentheses. All specifications include the same nine regressors. Delay 10 and Delay 15 are indicator variables for the different delay periods and reference delay 5. The dependent variable in column (1) is the utility loss (in %) from delaying optimal climate policy. Column (2) uses the consumption-equivalent DWL ($\phi(L)$ in %), and column (3) uses the absolute DWL in billions of 2020 USD.

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

REFERENCES

Ackerman, Frank, Elizabeth A Stanton, and Ramón Bueno. 2013. "Epstein–Zin utility in DICE: Is risk aversion irrelevant to climate policy?" *Environmental and Resource Economics*, 56(1): 73–84.

Bauer, Adam Michael, Cristian Proistosescu, and Gernot Wagner. 2024. "Carbon Dioxide as a Risky Asset." *Climatic Change*, 177(5): 72.

Bertram, Christoph, Elina Brutschin, Laurent Drouet, Gunnar Luderer, Bas van Ruijven, Lara Aleluia Reis, Luiz Bernardo Baptista, Harmen-Sytze de Boer, Ryna Cui, Vassilis Daioglou, et al. 2024. "Feasibility of peak temperature targets in light of institutional constraints." *Nature Climate Change*, 14(9): 954–960.

Bilal, Adrien, and Diego R Känzig. 2025. "The macroeconomic impact of climate change: Global vs. local temperature." National Bureau of Economic Research.

Burke, Marshall, Solomon M. Hsiang, and Edward Miguel. 2015. "Global non-linear effect of temperature on economic production." *Nature*, 527(7577): 235–239.

Burke, Marshall, W. Matthew Davis, and Noah S. Diffenbaugh. 2018. "Large potential reduction in economic damages under UN mitigation targets." *Nature*, 557(7706): 549–553.

Daniel, Kent D., Robert B. Litterman, and Gernot Wagner. 2019. "Declining CO_2 price paths." *Proceedings of the National Academy of Sciences*, 116(42): 20886–20891.

Dechezleprêtre, Antoine, Adrien Fabre, Tobias Kruse, Bluebery Planterose, Ana Sanchez Chico, and Stefanie Stantcheva. 2025. "Fighting climate change: International attitudes toward climate policies." *American Economic Review*, 115(4): 1258–1300.

Dietz, Simon, James Rising, Tobias Stoerk, and Gernot Wagner. 2021. "Economic impacts of tipping points in the climate system." *Proceedings of the National Academy of Sciences*, 118(34): e2103081118.

Epstein, Larry G, and Stanley E Zin. 1989. "Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework." *Econometrica: Journal of the Econometric Society*, 937–969.

Epstein, Larry G, and Stanley E Zin. 1991. "Substitution, risk aversion, and the temporal behavior of consumption and asset returns: An empirical analysis." *Journal of political Economy*, 99(2): 263–286.

Howard, Peter H., and Thomas Sterner. 2017. "Few and Not So Far Between: A Meta-analysis of Climate Damage Estimates." *Environmental and Resource Economics*, 68(1): 197–225.

Intergovernmental Panel on Climate Change (IPCC). 2023. "Climate Change 2023: Synthesis Report. Contribution of Working Groups I, II and III to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change." Intergovernmental Panel on Climate Change (IPCC), Geneva, Switzerland.

Joos, F., R. Roth, J. S. Fuglestvedt, G. P. Peters, I. G. Enting, W. von Bloh, V. Brovkin, E. J. Burke, M. Eby, N. R. Edwards, T. Friedrich, T. L. Frölicher, P. R. Halloran, P. B. Holden, C. Jones, T. Kleinen, F. T. Mackenzie, K. Matsumoto, M. Meinshausen, G.-K. Plattner, A. Reisinger, J. Segschneider, G. Shaffer, M. Steinacher, K. Strassmann, K. Tanaka, A. Timmermann, and A. J. Weaver. 2013. "Carbon dioxide and climate impulse response functions for the computation of greenhouse gas metrics: a multi-model analysis." *Atmospheric Chemistry and Physics*, 13(5): 2793–2825.

Lemoine, Derek, and Christian Traeger. 2014. "Watch your step: optimal policy in a tipping climate." *American Economic Journal: Economic Policy*, 6(1): 137–166.

Meckling, Jonas. 2025. "The geoeconomic turn in decarbonization." *Nature*, 645(8082): 869–876.

Meckling, Jonas, Thomas Sterner, and Gernot Wagner. 2017. "Policy sequencing toward decarbonization." *Nature Energy*, 2(12): 918–922.

Mildenberger, Matto. 2020. *Carbon captured: How business and labor control climate politics.* MiT Press.

Mildenberger, Matto, and Dustin Tingley. 2019. "Beliefs about climate beliefs: the importance of second-order opinions for climate politics." *British Journal of Political Science*, 49(4): 1279–1307.

Moore, Frances C, Moritz A Drupp, James Rising, Simon Dietz, Ivan Rudik, and Gernot Wagner. 2024. "Synthesis of evidence yields high social cost of carbon due to structural model variation and uncertainties." *Proceedings of the National Academy of Sciences*, 121(52): e2410733121.

Oreskes, Naomi, and Erik M Conway. 2011. *Merchants of doubt: How a handful of scientists obscured the truth on issues from tobacco smoke to global warming.* Bloomsbury Publishing USA.

Rose, Steven K., Delavane B. Diaz, and Geoffrey J. Blanford. 2017. "Understanding the Social Cost of Carbon: A Model Diagnostic and Inter-comparison Study." *Climate Change Economics*, 8(02): 1750009.

Traeger, Christian P. 2014. "Why uncertainty matters: discounting under intertemporal risk aversion and ambiguity." *Economic Theory*, 56(3): 627–664.

Wagner, Gernot, and Richard J Zeckhauser. 2012. "Climate policy: hard problem, soft thinking." *Climatic change*, 110(3): 507–521.

Appendix

A1. Model and dynamics

This section formally describes the model setup and proves Proposition A.1. Decisions occur on a finite set of times $T_0, T_1, \ldots, T_N$, measured in calendar years (e.g., $T_0 = 2020, T_1 \in \{2025, 2030, 2035\}$, etc.). At each decision time T_t and node, the social planner (our representative agent) chooses the mitigation rate $m_t \in [0, \bar{m}]$ that applies over the entire subsequent interval $[T_t, T_{t+1})$. Values $m_t > 1$ are net removal of CO_2 (direct air capture and related negative-emissions technologies) at a backstop premium. In our benchmark calibration, $\bar{m} = 1.5$ and the backstop premium is \$10,000.

Following Daniel, Litterman and Wagner (2019), we embed the planner in a non-recombining binomial tree in which each node inherits a current fragility state that indexes how severe climate damages have turned out so far. The high-fragility branches correspond to high realized damages (or bad climate/economic states), and the low-fragility branch to more benign outcomes. Uncertainty therefore resolves gradually along the tree rather than all at once. At each node, the agent re-optimizes given the currently realized state.

This structure matters for two reasons. First, it makes the problem explicitly stochastic as future consumption, temperature, and damages differ across branches. Second, it allows us to separate aversion to risk across time from aversion to risk across states of nature. The same structure also lets us impose politically relevant constraints on early climate policy: we can constrain the agent not to mitigate for an initial window and then ask how the system behaves once the constraint is lifted.

We analyze delays in mitigation by imposing an exogenous no-mitigation period of length L years. Formally, for a given $L \in \{5, 10, 15\}$, we impose $m_t = 0$ for all decision times $T_t < L$, and relax this constraint for $T_t \geq L$.[5] We compare each

[5] In our standard calibration, this corresponds to constraining T_0 only.

scenario to a common unconstrained baseline with a node at $T_t = 10$ in which the planner is free to choose m_t at all decision times.

Preferences follow a standard Epstein-Zin recursive specification as in Equation (1) with terminal utility given by Equation (2). For our main specification, parameter values are as follows: PRTP(δ) = 0.002; EIS(σ) = 0.833; RA(γ) = 10; consumption growth p.a. = 0.02; exog. tech change = 0.015; endo. tech learning = 0; baseline emissions = SSP2; backstop premium = 10 000.

Emissions, climate dynamics, costs, and damages

For expositional clarity, throughout this subsection we fix an arbitrary realization (i.e., path) of uncertainty and suppress state indices; all objects $(m_t, \theta_t, L_t, \kappa_t, \Phi_t, \tau_t)$ are thus defined along a single path in our binomial tree.

Let $E_t > 0$ denote the baseline (business-as-usual) CO_2 emissions over the interval $[T_t, T_{t+1})$, based on a reference socioeconomic pathway (e.g., SSP2). The planner can abate a fraction $m_t \in [0, \bar{m}]$ of those baseline emissions, so realized emissions over that interval are

$$e_t = (1 - m_t)E_t. \tag{A1}$$

The climate state at time T_t is characterized by two key variables: atmospheric CO_2 concentration C_t (in ppm) and temperature anomaly θ_t (in °C above preindustrial).

Atmospheric CO_2 concentration C_t evolves as the convolution of past emissions with the impulse response function (IRF) of the carbon cycle, following Joos et al. (2013):

$$C_t = C_0 + \chi \int_0^t \Psi(t - s) e_s \, ds, \qquad \text{where } \Psi(s) = a_0 + \sum_{i=1}^{3} a_i \exp(-s/b_i), \tag{A2}$$

with $\chi = 1/7.8 = 0.128$ ppm/GtCO_2 and coefficients $a_0 = 0.2173$, $a_1 = 0.2240$,

$a_2 = 0.2824$, $a_3 = 0.2763$, and time constants $b_1 = 394.4$, $b_2 = 36.54$, $b_3 = 4.304$ years. All $a_i, b_i > 0$, so concentrations are strictly increasing in cumulative emissions. These parameters capture the multiple carbon-cycle reservoirs (atmosphere, mixed-layer ocean, deep ocean, biosphere) and the long-lived airborne fraction a_0.

The global mean surface temperature anomaly θ_t is linked to cumulative emissions through the Transient Climate Response to Cumulative Emissions (TCRE) framework following AR6:

$$\theta_t = \lambda_{\text{eff}} \int_0^t e_u \, du, \qquad \lambda_{\text{eff}} := \frac{\lambda}{1 - f_{\text{nc}}}. \tag{A3}$$

Here e_u denotes CO_2 emissions at time u measured in thousand gigatonnes of CO_2 per year ($TtCO_2$/yr), so that $\int_0^t e_u \, du$ is cumulative emissions in thousands $GtCO_2$. The parameter $\lambda > 0$ (in K per 1000 $GtCO_2$) is the TCRE for CO_2-only warming, while $f_{\text{nc}} \in (0, 1)$ scales in the contribution from non-CO_2 forcing. Following Bauer, Proistosescu and Wagner (2024), we take the effective TCRE to be $\lambda_{\text{eff}} \sim \mathcal{N}(0.52,\, 0.21^2)$ K per $TtCO_2$.

This formulation implies three key properties for our analytical results:

(i) C_t and θ_t are strictly increasing in the emissions path $\{e_s\}_{s \le t}$ because $\Psi(\zeta) \ge 0$ and $\lambda_{\text{eff}} > 0$.

(ii) The multi-timescale IRF ensures that past emissions affect concentrations far into the future, with fraction a_0 remaining indefinitely.

(iii) High emissions in $[T_i, T_{i+1})$ permanently elevate both C_t and θ_t for all subsequent times.

Based on Burke, Davis and Diffenbaugh (2018); Rose, Diaz and Blanford (2017); Howard and Sterner (2017); Dietz et al. (2021), damages are represented as the sum of an aggregate temperature-based loss component and an additional component from climate tipping points:

$$d_t = D^{(k)}(\theta_t) + d_{\text{tp}}(\theta_t), \qquad D^{(k)}(\theta_t) = \delta_1^{(k)}\theta_t + \delta_2^{(k)}\theta_t^2 \tag{A4}$$

where $k \in \{\text{statistical, structural, meta}\}$ indexes the aggregate damage family, and $d_{\text{tp}}(\theta_t)$ captures the expected effect of climate tipping events. The coefficients $(\delta_1^{(k)}, \delta_2^{(k)})$ were calibrated by (Bauer, Proistosescu and Wagner, 2024) from the respective sources and may vary across periods[6], but for any fixed t, each $D^{(k)}$ is quadratic in θ_t and increasing on the temperature range we study (0-6°C).

The structural IAM function (Rose, Diaz and Blanford, 2017) and the meta-analytic function (Howard and Sterner, 2017) are convex ($\delta_{2,t}^{(\cdot)} > 0$) over our calibration. The statistical function (Burke, Hsiang and Miguel, 2015) is convex through mid-century and becomes mildly concave in late-century ($\delta_{2,t}^{(\text{stat})} < 0$ for t after 2100) but remains increasing on the relevant temperature range. The tipping component $d_{\text{tp}}(\theta_t)$ is also quadratic with positive curvature (Dietz et al., 2021), so it raises marginal damages at higher temperatures. To not rely too heavily on single estimates, our main specification averages across the three aggregate families with equal probability and adds the tipping component in every draw. Hence, we define the model-averaged damages at time t as

$$\bar{D}(\theta_t) := \mathbb{E}_k \left[D^{(k)}(\theta_t) + d_{\text{tp}}(\theta_t) \right] \quad \forall k, \tag{A5}$$

where the expectation is over the three aggregate families with equal weights. In our main specification, this means $\bar{D}(\theta_t)$ is twice continuously differentiable and satisfies

$$\frac{d\bar{D}(\theta_t)}{d\theta_t} > 0 \quad \text{and} \quad \frac{d^2\bar{D}(\theta_t)}{d\theta_t^2} \geq 0 \tag{A6}$$

[6]The statistical specification from (Burke, Hsiang and Miguel, 2015), for example, uses distinct mid- and end-century calibration.

for temperatures between 0 and 6 °C.

The period-t marginal abatement cost curve (MACC) follows the exponential form and is calibrated to IPCC AR6 Working Group III data (Intergovernmental Panel on Climate Change , IPCC), consistent with Bauer, Proistosescu and Wagner (2024). For a mitigation rate $m_t \in [0, \bar{m}]$ and technology and learning state L_t, which captures both exogenous and endogenous technological progress, we specify

$$\text{(A7)} \qquad \tau_t(m_t, L_t) = \begin{cases} L_t \tau_0 \big(e^{\xi m_t} - 1\big), & 0 \le m_t \le 1, \\ L_t(\tau_0 + \tau^{\text{prem}})\big(e^{\xi m_t} - 1\big), & m_t > 1, \end{cases}$$

where $\tau_0 > 0$ and $\xi > 0$ are level and curvature parameters, and $\tau^{\text{prem}} > 0$ is a backstop premium representing the additional cost of net-negative emissions (e.g., direct air capture). The corresponding total mitigation cost function is obtained by integrating the marginal cost curve for $m_t \le 1$,

$$\text{(A8)} \qquad \kappa_t(m_t, L_t) = L_t\, \tau_0 \left(\frac{e^{\xi m_t} - 1}{\xi} - m_t \right), \quad (0 \le m_t \le 1),$$

and analogously with $\tau_0 + \tau^{\text{prem}}$ for $m_t > 1$.[7] On each regime $m_t \in [0, 1]$ and $m_t > 1$, the function $\kappa_t(\cdot, \cdot)$ is twice continuously differentiable, strictly increasing and convex in m_t, and weakly increasing in the learning factor L_t. Higher L_t indicates less technological progress and therefore higher costs, while lower L_t reflects learning-by-doing and innovation that shift the MACC downward.

The learning factor L_t evolves according to cumulative mitigation experience and exogenous technological improvement. Hence,

$$\text{(A9)} \qquad L_t = \big(1 - \psi_0 - \psi_1 X_t\big)^{(Y_t - Y_{\text{ref}})},$$

where Y_t is the calendar year at decision time T_t, Y_{ref} is the reference year used

[7] Note that this creates a level jump at $m_t = 1$.

for calibration (2030 in our baseline), and parameters $\psi_0 \geq 0$ and $\psi_1 \geq 0$ capture exogenous and endogenous technological progress, respectively. The term X_t represents the weighted average mitigation up to time t,

$$X_t := \frac{\int_0^t m(\zeta)E(\zeta)d\zeta}{\int_0^t E(\zeta)d\zeta}, \tag{A10}$$

so that stronger cumulative mitigation or faster exogenous innovation lowers L_t and thereby reduces future abatement costs.

Let y_t be gross resources available for consumption at time T_t. Actual consumption is then determined by

$$c_t = y_t - \kappa_t(m_t, L_t) - d_t(\theta_t). \tag{A11}$$

Delay thus reduces consumption through a more deteriorated climate state causing higher damages, and slower cost decline raising mitigation costs.

Optimal expected carbon prices

To compare policies at a given decision time t, we now take expectations across nodes, as in the main text (see Equation (3)).

Proposition A.1. (Delay raises the expected entry carbon price) *Assume baseline emissions are strictly positive in all periods prior to* T_1*, i.e.,* $E_t > 0$ *for all* $T_t < T_1$*. Suppose:*

- *(i) The delay constraint is binding in the baseline period, i.e., there exists* $t < T_1$ *with* $m_t^{base} > 0$*, while in the delayed scenario* $m_t^{delay} = 0$ *for all* $T_t < T_1$*;*
- *(ii) The carbon-cycle and TCRE mappings in* (A2)–(A3) *satisfy* $\Psi(\zeta) \geq 0$ *and* $\lambda_{eff} > 0$*, so climate dynamics are monotone;*
- *(iii) Model-averaged damages are increasing and weakly convex:* $\bar{D}'(\theta) > 0$ *and* $\bar{D}''(\theta) \geq 0$ *(cf.* (A5)*);*

(iv) The optimal m_{T_1} is interior in both the base and delay scenarios.

Then, for every node $s \in \mathcal{S}_{T_1}$, $\tau^{delay}_{T_1,s} \geq \tau^{base}_{T_1,s}$, and therefore, taking expectations across nodes,

$$\tau_1^{delay} \geq \tau_1^{base}, \tag{A12}$$

with strict inequality whenever $\theta^{delay}_{T_1,s} > \theta^{base}_{T_1,s}$ on a set of nodes with positive probability (equivalently, when $\mathbb{E}_1[\theta^{delay}_{1,S}] > \mathbb{E}_1[\theta^{base}_{1,S}]$).

Proof. Fix an arbitrary node $s \in \mathcal{S}_{T_1}$ and consider the unique history (path) leading to s under the baseline and delayed scenarios. By construction, $m_t^{\text{delay}} = 0$ for all $T_t < T_1$, while by (i) there exists some $t < T_1$ with $m_t^{\text{base}} > 0$. From (A1), for at least one such t we have $e_t^{\text{delay}} = (1-0)E_t = E_t > (1-m_t^{\text{base}})E_t = e_t^{\text{base}}$. Hence cumulative emissions up to T_1 are (weakly) higher along the delayed path. By (ii) and the monotonicity of the carbon-cycle IRF and TCRE mappings (A2)–(A3), higher past emissions imply a (weakly) worse climate state at T_1: $C^{\text{delay}}_{T_1,s} \geq C^{\text{base}}_{T_1,s}$ and $\theta^{\text{delay}}_{T_1,s} \geq \theta^{\text{base}}_{T_1,s}$, with strict inequality if baseline mitigation was positive on a nonzero set. By (iii), marginal damages are increasing in temperature, so $\Phi_{T_1,s}(\theta^{\text{delay}}_{T_1,s}) \geq \Phi_{T_1,s}(\theta^{\text{base}}_{T_1,s})$, strictly when $\theta^{\text{delay}}_{T_1,s} > \theta^{\text{base}}_{T_1,s}$. Under the interior first-order condition at node (T_1, s), $\tau^{\text{delay}}_{T_1,s} = \Phi_{T_1,s}(\theta^{\text{delay}}_{T_1,s})$ and $\tau^{\text{base}}_{T_1,s} = \Phi_{T_1,s}(\theta^{\text{base}}_{T_1,s})$, which implies the path-wise inequality $\tau^{\text{delay}}_{T_1,s} \geq \tau^{\text{base}}_{T_1,s}$. Taking expectations across nodes yields $\tau_1^{\text{delay}} = \mathbb{E}_{T_1}[\tau^{\text{delay}}_{T_1,S}] \geq \mathbb{E}_{T_1}[\tau^{\text{base}}_{T_1,S}] = \tau_1^{\text{base}}$, with strict inequality if $\theta^{\text{delay}}_{T_1,s} > \theta^{\text{base}}_{T_1,s}$ on a set of nodes with positive probability. Finally, delay implies a smaller technology stock (higher $L_{T_1,s}$) because learning-by-doing accumulates more slowly when early mitigation is zero. By (A9), $L^{\text{delay}}_{T_1,s} \geq L^{\text{base}}_{T_1,s}$, and since mitigation costs are increasing in L_t, $\kappa_{T_1,s}(m_{T_1,s}, L^{\text{delay}}_{T_1,s}) \geq \kappa_{T_1,s}(m_{T_1,s}, L^{\text{base}}_{T_1,s})$. Hence the total resource cost under delay is weakly higher, reinforcing the conclusion that $\tau_1^{\text{delay}} \geq \tau_1^{\text{base}}$. □

Table A1—: OLS estimate of the delay-to-price elasticity

	(1)
L (years)	0.002761
	(0.000487)
Constant	5.241717
	(0.004929)
R^2	0.969818
Observations	3

Note: Dependent variable: $\log(\tau_L^{\text{delay}})$. Robust standard errors in parentheses. 95% confidence interval for the slope: [0.001806, 0.003715], i.e. [0.181%, 0.372%] per year.

Regression results delay-to-price elasticity

A2. Analytical solutions

Cobb-Douglas limit of the EZ aggregator and the first periodic compensation

Solving for the consumption–equivalent DWL of delay when $\rho \neq 0$ is done in the main text. Here we consider the Cobb–Douglas case $\rho = 0$ (i.e. $\sigma = 1$).

As $\rho \to 0$, the Epstein–Zin aggregator $U_0 = \left((1-\beta)c_0^\rho + \beta\,\mathrm{CE}_1^\rho\right)^{1/\rho}$ converges to the geometric (Cobb–Douglas) form $U_0 = c_0^{1-\beta}\mathrm{CE}_1^\beta$. Let the delayed path have first–period consumption c_0^D and continuation certainty equivalent CE_1^D. We scale the first–period consumption by $(1+\phi)$ and require the compensated delayed utility to equal the baseline utility $U_0^\star$: $U_0^\star = \left((1+\phi)c_0^D\right)^{1-\beta}(\mathrm{CE}_1^D)^\beta$. Solving for ϕ gives the closed-form consumption-equivalent transfer:

$$\phi_{\mathrm{CD}} = \left[\frac{U_0^\star}{(c_0^D)^{1-\beta}\,(\mathrm{CE}_1^D)^\beta}\right]^{\frac{1}{1-\beta}} - 1. \tag{A13}$$

A3. Extended outputs

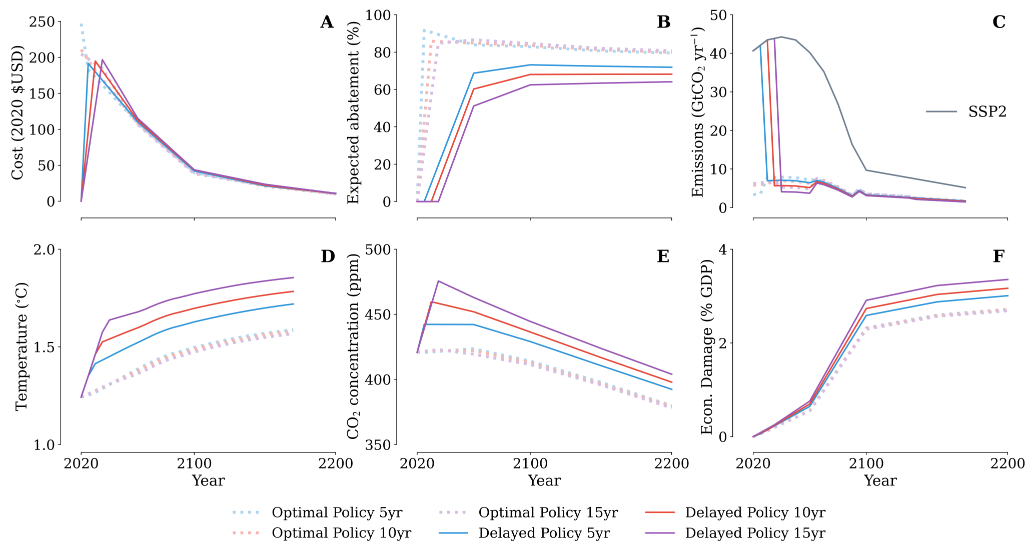

Figure A1. : Different outputs of the single-period optimal scenarios with delay periods $L \in \{5, 10, 15\}$.

Note: A: Expected carbon prices in 2020 \$/tCO_2. B: Expected abatement in %. C: Emissions in GtCO_2/year with SSP2 emissions baseline. D: Temperature in degrees Celsius. E: CO_2 concentration in ppm. F: Economic damages in % of Gross Domestic Product.

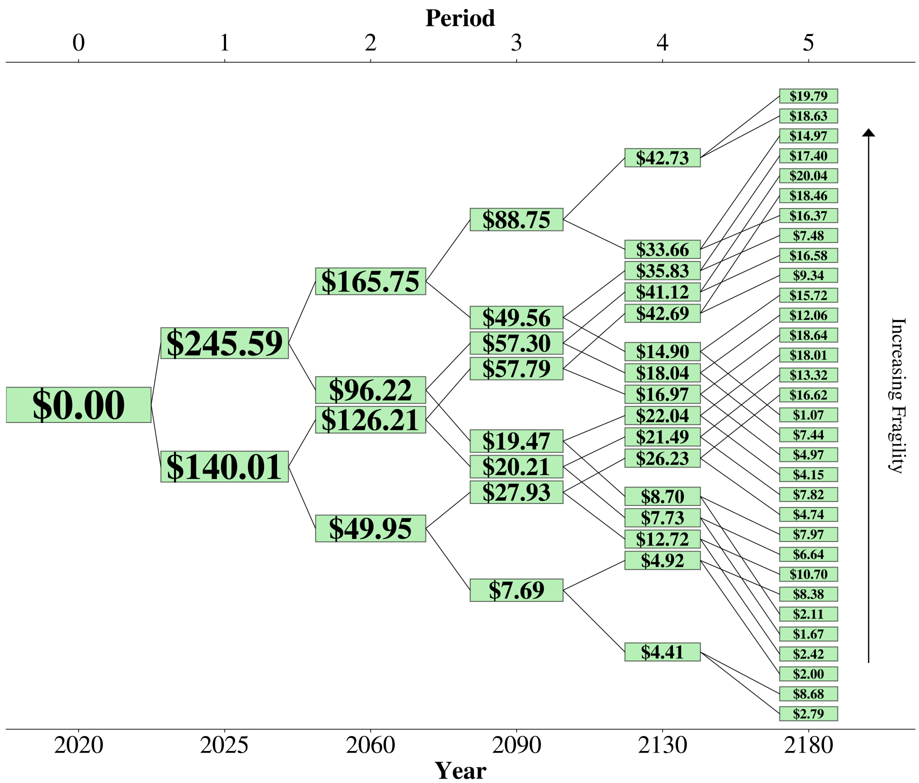

(a) Optimal price tree with 5-year delay.

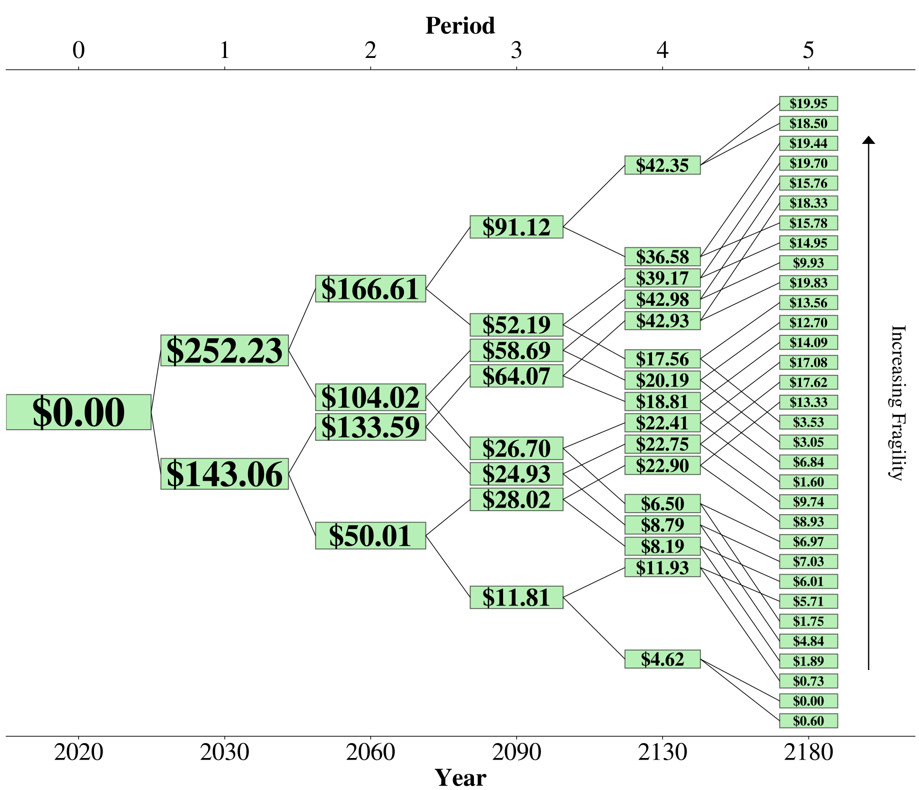

(b) Optimal price tree with 10-year delay.

Figure A2. : Decision trees for delayed runs.

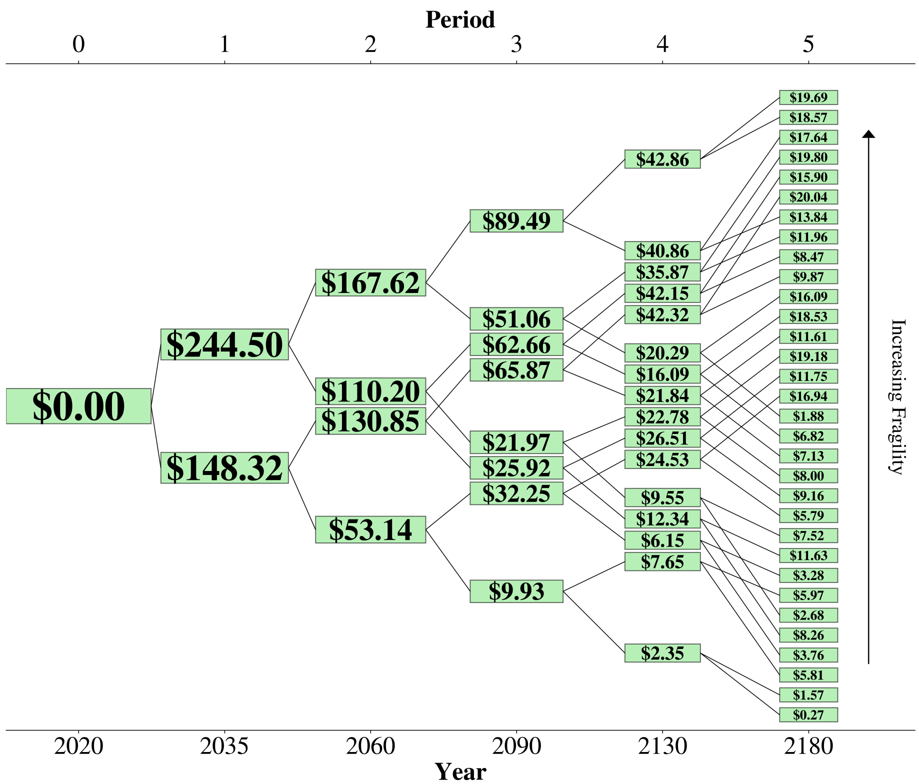

(c) Optimal price tree with 15-year delay.

Figure A2. : Decision trees for delayed runs (continued).