The Geoengineering Put: Ecological Debt, Termination Risk, and Asset Prices

Samuel Hartzmark* Kelly Shue†

August 27, 2026

Abstract

Geoengineering offers an inexpensive, technically proven and fast-acting way to lower global temperatures without removing carbon from the atmosphere. This suggests climate-vulnerable countries will be increasingly likely to implement it. The asset pricing consequences of this have been largely overlooked. We develop a macro-finance framework in which geoengineering acts as a put option on climate damages, funded by the "ecological debt" of masked warming. If the program stops, temperatures snap back in a catastrophic termination shock. We show that the anticipation of geoengineering initially subsidizes fossil fuel firms, but ultimately transforms them into toxic assets as ecological debt deepens. Geoengineering and carbon dioxide removal switch from substitutes to complements once termination risk is present. Geoengineering reshapes innovation: deployment lowers temperatures while the world waits for a green breakthrough, but it weakens innovation incentives, because innovators are paid through carbon prices that deployment suppresses. Termination risk partially revives research incentives and directs it toward carbon removal, the technology that can retire ecological debt. The potential for geoengineering implementation can influence current dynamics even if its implementation is not imminent. Our model underscores the value of research today to better understand this technology and its impact.

*Boston College, Carroll School of Management and NBER.
†Yale University, School of Management and NBER.

1 Introduction

The economic framework to understand global warming typically revolves around the gargantuan costs to humanity of climate change (around $10 trillion per year)[1] and the immense cost required to address the problem (around $1 trillion per year)[2]. Further complicating the issue is the free-rider problem whereby sovereign actors face an incentive to under-invest in the transition to a low-carbon economy, as the benefits of one country's costly mitigation are globally shared. The potential deployment of climate geoengineering fundamentally alters this paradigm.

Geoengineering offers a physically plausible and scientifically proven way to cool the planet that is orders of magnitude less expensive (around $1 to $20 billion per year) than commonly discussed remedies to climate change (McClellan et al., 2012; Smith and Wagner, 2018). While geoengineering decreases the earth's temperature, it does not reduce the underlying atmospheric carbon stock. By simply masking the symptoms of accumulating carbon, it therefore introduces the risk of a termination shock: cessation of geoengineering would induce a rapid spike in global temperatures dictated by the high underlying carbon stock, with potentially massive economic costs.

In spite of this and other risks, as carbon accumulates it becomes increasingly likely that geoengineering will be implemented. This could occur with international consensus, but the low cost and ease of implementation of geoengineering means that it is plausible that a single country could deploy it unilaterally (McClellan et al., 2012; Smith and Wagner, 2018). This shifts the climate dilemma from a free-rider problem to what Weitzman (2015) termed the free-driver problem, where one actor can lower the global thermostat.[3] Given the heterogeneous impact of global warming, countries facing relatively high costs of climate change may want to implement the technology, even in the face of international opposition.

The asset pricing implications of geoengineering risk have been largely overlooked in the climate finance literature, which has instead focused on transition risk. Transition risk refers to the pricing of asset values that can occur as policy or technological change moves the economy away from fossil fuels. Aggressive carbon pricing and environmental regulation can

[1] For example, Bilal and Känzig (2026) estimates an impact of about 20% of GDP, around $20 trillion in today's dollars.

[2] For example, the IPCC estimates costs of about 1% to 4% of GDP, roughly $1 to $4 trillion `https://www.ipcc.ch/report/ar6/wg3/chapter/chapter-3/`

[3] Hartzmark and Shue (2026) argue that if the deploying country is held internationally responsible for adverse side effects, the strategic logic of geoengineering reverts from a free-driver to a free-rider problem. Every country prefers another to bear the blame for deployment, generating a war of attrition. Rather than deployment being brought forward in time by a unilateral actor, deployment can be delayed relative to the globally optimal timing. The asset pricing analysis developed in this paper applies under either strategic interpretation, since what matters is that unilateral deployment is technologically feasible and plausible.

erode the cash flows of carbon-intensive firms and impair the fiscal capacity of sovereigns whose tax bases depend on fossil fuel extraction. A growing literature explores the asset pricing implications of transition risk across equities, corporate debt, and sovereign markets (Bolton and Kacperczyk, 2021, 2023; van der Ploeg and Rezai, 2020; Engle et al., 2020; Acharya et al., 2025; Barnett et al., 2020). A more recent strand examines biodiversity risk, the financial exposure of firms and sovereigns to the loss of natural capital and the ecosystem services on which production depends (Dasgupta, 2021; Giglio et al., 2024; Garel et al., 2024). Geoengineering constitutes another important source of climate-related risk, with asset pricing implications that are sharply different from those of transition risk.

In this paper, we develop a macro-finance framework to study the economic implications of climate geoengineering in a general equilibrium setting. We define geoengineering broadly as any technological intervention that lowers realized global temperatures without altering the underlying atmospheric carbon stock. We focus on solar radiation management (SRM), a class of technologies that cool the planet by reflecting a small fraction of incoming sunlight. The most prominent example is stratospheric aerosol injection (SAI), which disperses reflective particles such as sulfur dioxide into the stratosphere.[4] The physics of SAI are illustrated by the eruption of Mount Pinatubo in 1991 that injected about 20 million tons of sulfur dioxide into the stratosphere, cooling the planet by about 0.5 degrees Celsius over the next two years. Engineering estimates put direct delivery costs at a few billion dollars per year for an early program (Smith and Wagner, 2018; McClellan et al., 2012), plausibly $10 to $20 billion per year at full scale — a fraction of a percent of the annual capital expenditure required for global decarbonization.[5] Given this asymmetric cost structure, a sovereign nation facing severe heat damages may possess both the means and the incentive to act unilaterally.

We formalize the macroeconomic consequences of geoengineering deployment by introducing two organizing concepts. The first is the "geoengineering put." Deployment caps global temperatures and truncates the right tail of climate damages, acting as a put option on the climate. However, because geoengineering only masks the warming effects of greenhouse gases without removing them from the atmosphere, exercise of this put is accompanied by the continuous accumulation of "ecological debt": the gap between the realized temperature and the temperature the world would experience given the carbon stock in the absence of geoengineering. Thus, deployment carries an accumulating catastrophic tail risk.

[4]Our definition of geoengineering also applies to other SRM approaches including marine cloud brightening, surface or cloud albedo enhancement, and space-based reflectors, as well as radiative forcing technologies, such as cirrus cloud thinning.

[5]Smith and Wagner (2018) price a program that halves the rate of warming at roughly $2.25 billion per year averaged over its first fifteen years. A program masking on the order of 1–2°C decades later requires roughly an order of magnitude more injection, which is the basis of the $10–20 billion at-scale figure.

This tail risk is recognized in the geophysical literature as a termination shock (Parker and Irvine, 2018; Trisos et al., 2018). A termination shock occurs because of the difference in lifespans between carbon dioxide and stratospheric aerosols. Carbon dioxide lingers in the atmosphere for centuries whereas stratospheric aerosols dissipate in one to two years. If geoengineering were to be suddenly halted, temperatures would snap back to the level determined by the underlying carbon stock, and the speed of the warming would greatly exceed that of baseline climate change. This sudden jump in global temperatures could lead to catastrophic destruction of global capital and productivity. In our framework, a termination shock represents a default on ecological debt.

We organize our analysis into three nested regimes, which successively activate the main economic mechanisms. In the baseline (Regime 1), geoengineering does not exist and the economy faces standard climate risks. In Regime 2, geoengineering is deployed once the carbon stock passes a known threshold. Continuation of geoengineering is certain, allowing us to isolate the moral hazard channel. In the full model (Regime 3), we introduce a termination hazard. Our key insight is that the same physical model of geoengineering has opposite effects on asset prices depending on whether termination risk is present.

We numerically solve the planner's Hamilton-Jacobi-Bellman (HJB) equation to obtain the forward-looking social cost of carbon, equilibrium dynamics, and asset prices for each regime. The model generates four main insights regarding how termination risk affects climate policy and asset prices and a second set of results regarding how geoengineering and the termination risk it creates affect innovation.

First, the geoengineering put acts as a moral hazard device. It reduces realized temperatures, which decreases the optimal carbon tax, raises the value of fossil fuels, and slows the green transition. If geoengineering were to mask all warming perfectly, the price of temperature damages would disappear from the social cost of carbon, and climate risk would no longer be priced through temperature damages.

Second, termination risk fundamentally alters the cross-section of expected returns. The positive termination risk premium on fossil fuels means that they will crash when the representative investor's marginal utility is highest. We show that the value of fossil fuel firms is non-monotonic in the carbon stock. At first, geoengineering helps fossil fuels by lowering the carbon tax. But eventually, it makes them toxic assets as ecological debt builds up.

Third, the relationship between geoengineering and carbon dioxide removal (CDR) flips. Without termination risk, geoengineering crowds out CDR. Since the carbon tax is suppressed, the price of carbon is too low to fund CDR operations. With termination risk, geoengineering and CDR are complements because CDR can pay down ecological debt, reducing the damages from termination risk.

Fourth, the social cost of carbon can be higher under geoengineering than it would be without geoengineering. At high enough concentrations, the next unit of carbon becomes more socially costly in Regime 3 than in Regime 1. That is, not only does the next unit cause more flow damages, but it also increases the probability of a latent catastrophe.[6]

The paper's second major contribution is to show that geoengineering and the termination risk it creates influences incentives for green innovation. The strongest argument for solar geoengineering is peak shaving: deploy SRM to buy time while the world waits for a technological breakthrough (Keith and MacMartin, 2015). We model two breakthroughs: a green energy cost collapse and the emergence of carbon removal at scale. The arrival rates of these breakthroughs depend on the market value of the innovation prize. We show that, while geoengineering can serve as peak shaving, it can also reduce incentives to innovate. Innovators will be paid through carbon prices, which deployment can suppress.

Termination risk partly revives innovation incentives and redirects them. Without termination risk, geoengineering suppresses both green energy and carbon removal research. With termination risk, carbon removal research revives disproportionately, because removal is the only technology that can retire ecological debt. Revival can be late, however. Under our baseline calibration, when geoengineering is deployed, it holds down the amount of removal research (relative to no geoengineering) until the carbon stock rises to 663 ppm. From today's concentration levels, the availability of geoengineering delays the arrival of a carbon removal breakthrough by approximately 6.2 years.

Our analysis shows that carbon removal has a built-in incentive problem: it is paid through the price of carbon, and that price falls as removal succeeds. Every ton removed makes the remaining carbon less dangerous, so the market price of carbon drops below the cost of removal before removal is complete. Potential developers of carbon removal breakthroughs look ahead to these declining prices and under-invest today. We show that an advance commitment contract (Kremer et al., 2020) resolves this problem: a promise to buy removal at a set price in the future restores the incentive to innovate.

Finally, we consider whether welfare is higher in a regime with geoengineering technology and termination risk than in a regime in which geoengineering technology does not exist. To do so, we evaluate four factors: (1) how long deployment postpones the need for a breakthrough; (2) how likely it is that a breakthrough occurs before deployment starts; (3) how much termination risk it incurs once deployed; and (4) how dampened incentives become after deployment reduces the need for further breakthroughs. We find that the world with geoengineering has modestly higher welfare. Under our baseline calibration, a

[6]Above 663 ppm under the baseline calibration, the marginal unit of carbon is more socially costly with intervention than without it.

green breakthrough often arrives before deployment starts, and once deployment does begin, the capped temperatures buy decades of avoided damages before much termination risk has built up. However, the margin is small and sensitive to assumptions. It becomes negative if deployment begins after substantial carbon accumulation, or if the program is likely to be interrupted while breakthroughs are slow to arrive.

Our paper contributes to several strands of the literature. First, we extend the asset pricing literature on climate-related risks (Giglio et al., 2021; Acharya et al., 2025; Giglio et al., 2024; Barnett et al., 2020) to a setting with geoengineering and termination risk. Following research on rare disasters (Barro, 2006; Wachter, 2013; Gabaix, 2012), our model features a stochastic discount factor with endogenous jump risk. We also follow Gabaix (2012) in deriving large risk premia from the risk-neutral magnification of the termination hazard, despite the low physical probabilities of termination.

This paper is closely related to the economics literature on geoengineering (Barrett, 2008; Weitzman, 2015; Moreno-Cruz and Keith, 2013; Heutel et al., 2018; Belaia et al., 2021). Our contribution is to formalize the asset pricing consequences, showing how termination risk creates novel risk premia, and to solve for the forward-looking equilibrium, yielding quantitative magnitudes. Keith and MacMartin (2015) and MacMartin et al. (2014) develop a peak-shaving scenario in which the world chooses early and moderate deployment under cooperative governance. Our baseline calibration describes a complementary scenario in which a vulnerable nation follows a deterministic deployment rule regardless of international consensus. In supplementary analysis, we explore intermediate deployment rules that span the range between these two extremes.

Our paper builds on the integrated assessment literature (Nordhaus, 2017; Golosov et al., 2014; Barrage, 2020), which solves for the social cost of carbon (SCC). We show that geoengineering transforms the SCC by creating a regime switch between suppressed and tipping-enhanced damages as a function of the geoengineering state. This links to the tipping points literature of Lemoine and Traeger (2014) and Cai and Lontzek (2019). To our knowledge, we are the first to establish that the SCC may be higher with geoengineering than without.

Our paper underscores the importance of considering the implications of geoengineering today, even if deployment remains years away. Although it is unclear whether markets currently price geoengineering in a meaningful way, our model implies that failing to do so would be a mistake. Once markets account for geoengineering, our baseline calibration predicts a 10% increase in brown firm value and a roughly 5% decrease in the social cost of carbon relative to a world in which it is not considered. Perhaps more importantly, the uncertainty surrounding many of our calibration parameters points to the need for further research. Research conducted today will help determine whether and how geoengineering

should be deployed.

The remainder of the paper proceeds as follows. Section 2 develops the model. Section 3 presents the theoretical results under three nested regimes. Section 4 presents quantitative results from the solved model. Section 5 concludes.

2 Model

This section presents the model along with motivation for parameter estimates which will be used in the calibration. In Section 2.1, we first describe the physical climate system. Then, in Section 2.2, we introduce geoengineering, ecological debt and termination risk. In Section 2.3, we define the dividends of firms as reduced-form functions of the carbon tax. Section 2.4 introduces the planner's problem in three steps. In Sections 2.5-2.7, we complete the planner's problem by adding the two breakthrough technologies, endogenous innovation and the post-arrival regime tree. In Section 2.8, we develop the asset pricing framework.

2.1 Carbon Stock and Climate

The main slow-moving state variable is the stock of carbon in the atmosphere, S_t, measured in parts per million (ppm):

$$dS_t = e(\tau_t)\,dt + \sigma_S\,dZ_t \tag{1}$$

where $e(\tau_t)$ is aggregate net emissions as a function of the prevailing carbon tax τ_t, $\sigma_S > 0$ captures natural carbon cycle volatility, and Z_t is a standard Brownian motion. Emissions are strictly decreasing in the carbon tax up to the finite shutdown tax $\bar{e}/\psi$, and zero beyond. Higher carbon taxes induce substitution away from fossil fuels and toward cleaner alternatives. We calibrate $\sigma_S = 1.0$ ppm/$\sqrt{\text{yr}}$ to the observed interannual variability of the net airborne flux, roughly 0.5 - 1 ppm/yr in the Global Carbon Budget accounting (Friedlingstein et al., 2023).

Note that carbon removal does not appear in the above equation. Carbon removal breakthrough will be modeled in the next section, but carbon removal prior to the breakthrough is negligible. Cumulative carbon removal as of 2025 total only 1 $MtCO_2$ (based on CDR.fyi market data), which is four orders of magnitude smaller than the annual fossil and land-use emissions of roughly 40 $GtCO_2$ (Friedlingstein et al., 2023). Thus, we round current removal capacity to zero. The shadow value of removal capacity, q (S), is still a well-defined valuation object and plays a central role in our asset pricing results, because it requires no installed capacity to price.

Also note that Equation (1) is a net-emissions equation and does not contain a natural

decay term. Our calibration is already net of natural sinks (ocean and land which absorb carbon) during the emissions era. $\bar{e} = 3.0$ ppm/yr in our calibration is set to the observed *net* accumulation rate of around 5 ppm/yr. After the emissions era (when emissions cease), true net drift would become slightly negative. However, carbon decay is slow. Natural drawdown would retire less than a third of the accumulated stock over the 100-year timescale of our results (Joos et al., 2013). For simplicity, we round it to zero.

Absent geoengineering, global temperature anomaly is a strictly increasing function of the carbon stock:

$$T_t^{un} = \Psi(S_t), \qquad \Psi' > 0 \tag{2}$$

Throughout the paper, temperature refers to warming relative to the preindustrial level, so today's temperature corresponds to $T \approx 1.4°\text{C}$, and the 280 ppm appearing in Ψ below is the preindustrial carbon stock. This increasing function captures the logarithmic relationship between CO_2 concentration and radiative forcing, combined with climate feedbacks that amplify warming.

We calibrate Ψ as the *transient* concentration-temperature mapping: realized temperature as a function of the current stock. In particular, we choose $\Psi(S) = 2.3 \log_2(S/280)$, an effective transient slope of 2.3 °C per doubling, such that today's carbon stock yields today's temperature: $\Psi(427) = 1.4°\text{C}$ at the current concentration of $S_0 = 427$ ppm.

We decompose climate damages to aggregate productivity into a temperature component and a carbon-stock component,

$$\Omega(T) + \Omega_{acid}(S), \qquad \Omega(T) = \omega(T) + \zeta \cdot h(T - \bar{T}) \tag{3}$$

where $\omega(T)$ represents standard climate damages (e.g., heat stress or agricultural yield losses) and is smooth, increasing, and convex. $\bar{T}$ is a temperature threshold at which the climate crosses an irreversible tipping point, such as disintegration of ice sheets or large-scale release of permafrost carbon (Armstrong McKay et al., 2022). $h(\cdot)$ is a smooth, increasing, and convex function, and $h(x) \approx 0$ for $x \ll 0$, and $\zeta > 0$ is a coefficient scaling the magnitude of damages from tipping points. The formulation follows the stochastic tipping specification in Lemoine and Traeger (2014) and Cai and Lontzek (2019).

$\Omega_{acid}(S)$ represents the cost of ocean acidification and other *non-temperature* costs of atmospheric carbon. It cannot be masked by geoengineering and depends directly on the stock S. Estimates of damages from ocean acidification are unfortunately very limited. We calibrate $\Omega_{acid}(S)$ to add \$20/tCO_2 to the carbon price at current concentrations, and also explore sensitivity to assumptions in the \$10–30/tCO_2 range. The existence of $\Omega_{acid}(S) > 0$ implies that the social cost of carbon remains positive even if geoengineering reduces standard

climate damages to zero. We denote $\tau_{acid}(S)$ as the carbon tax that prevails if temperature damages are completely masked.

We calibrate $\omega(T)$ as quadratic in warming, $\omega(T) = \omega_2 T^2$, following DICE, the standard integrated assessment model of climate economics (Nordhaus, 2017). We set the coefficient so that the social cost of carbon in the no-geoengineering baseline equals \$125/tCO$_2$ at the current carbon stock, which places the model's carbon price in the middle of published estimates. The implied coefficient is $\omega_2 = 0.00439$, 1.86 times the corresponding DICE value.[7]

Assumption 1 (Damage Regularity)**.** *Ω is continuously differentiable, increasing, and convex, and twice continuously differentiable away from the tipping threshold $\bar{T}$. The calibrated tipping term is a quadratic hinge, so its second derivative steps at $\bar{T}$.*

2.2 Geoengineering and Ecological Debt

We assume the existence of a climate-vulnerable sovereign actor (Region V) that deploys geoengineering at intensity $g_t \geq 0$. Geoengineering artificially lowers the realized global temperature without altering the atmospheric carbon stock:

$$T_t = \Psi(S_t) - \theta\, g_t \tag{4}$$

where $\theta > 0$ captures the cooling efficiency of the intervention.

The vulnerable region V in our model matches potential deployers in the real world: countries where climate damages may soon become so great that it becomes privately optimal to deploy SRM unilaterally. One motivating example involves wet-bulb temperature, a combined measure of heat and humidity. At wet-bulb temperatures around 35°C, sustained exposure is considered an approximate limit beyond which even healthy people in shaded, ventilated environments suffer from deadly heat stress (Sherwood and Huber, 2010). If a country faces repeated exposures to these conditions, it might choose to implement SRM rather than accept large numbers of heat-related fatalities. This possibility is empirically relevant because wet bulb temperatures have come close to this threshold in localized areas of South Asia and the Persian Gulf (Raymond et al., 2020), and projections identify the Indus and Ganges basins, the Persian Gulf, and eastern China as potential hotspots in the

[7]Published SCC estimates vary mainly with the discount rate: Rennert et al. (2022) report roughly \$185/tCO$_2$ at a 2% near-term rate and \$80/tCO$_2$ at 3%, and our target sits between them. The planner discounts welfare at $\rho = 2.0\%$; the rate $r = 4.0\%$ prices assets. Our coefficient exceeds the DICE value for a mechanical reason: we calibrate temperature as a transient function of the current carbon stock, which produces less warming per unit of carbon than the equilibrium sensitivity DICE uses, so matching the same damages per ton of carbon requires more damage per degree of warming.

future (Pal and Eltahir, 2016; Im et al., 2017). A climate-vulnerable country might also be pressured to deploy before reaching the wet-bulb survivability threshold. For example, governments could react early if they experience significant losses in agriculture or labor productivity due to heat and other climate changes.

In our model, V should be interpreted as a reduced-form representation of countries for which the private value of geoengineering is unusually high, even if the global costs and risks of SRM are substantial.

Assumption 2 (Geoengineering Policy)**.** *The geoengineering intensity $g(S)$ as a known, exogenous policy function that is weakly increasing in the carbon stock: $g'(S) \geq 0$.*

Our propositions require only the general form in Assumption 2. Our baseline calibration further assumes a specific rule for $g(S)$, hereafter referred to as the *thermostat*:

$$g(S) = \max\left(0, \frac{\Psi(S) - T^*}{\theta}\right), \qquad T^* = 2.0°\text{C} \tag{5}$$

Deployment activates when the unmitigated temperature reaches T^*, at $S = \Psi^{-1}(T^*) \approx 512$ ppm, and holds realized temperature constant at T^* everywhere above the activation threshold. This rule is motivated by three considerations:

First, we set the activation threshold for deployment ($g(S) > 0$) at the point when lethal wet-bulb events shift from episodic to recurrent. Today's wet-bulb events are isolated and short (Raymond et al., 2020). Recurrent 35°C wet-bulb events in the Indus-Ganges basins and the Persian Gulf are predicted to occur once warming exceeds 2°C (Im et al., 2017; Pal and Eltahir, 2016). Warming reaches that level at ≈ 512 ppm, which exceeds the current 427 ppm by 85 ppm, consistent with the fact that nothing has been deployed yet. Our assumption imply that, with the carbon stock rising at roughly 3 ppm per year and accelerating, deployment is about two to three decades away.

Second, the intensity. Any geoengineering program that masks less than 100% of incremental warming above T^* is not credible, because realized temperature would drift above T^* and Region V would unilaterally top the program up. The same argument disciplines whoever operates the program, including successor coalitions that may take operations after V launches. The thermostat is the unique deployment rule consistent with V's threat point. This is consistent with Weitzman (2015)'s free-driver logic: the global thermostat is set by the most vulnerable actor with deployment capacity.

Third, exogeneity. We assume a global planner sets the carbon tax but takes the deployment rule $g(S)$ as given. This asymmetry is deliberate. Carbon taxation requires global coordination, while geoengineering is cheap enough that a vulnerable country can deploy it

on its own. We maintain this asymmetry throughout the paper.[8]

In robustness analysis, we also study geoengineering deployment functioning as partial masking: $g(S) = m(\Psi(S) - T^*)/\theta$ with masking share $m \in \{0.25, 0.5, 1\}$. The $m = 0.5$ case represents a governed, moderate program in the spirit of Keith and MacMartin (2015), who recommend early, limited, responsive deployment under international governance. Our baseline thermostat represents scenario in which that governance does not arrive: warming reaches recurrent lethal levels and deployment is determined by the most vulnerable actor.

We define the *ecological debt* as:

$$F(S) \equiv \theta\, g(S) \tag{6}$$

Ecological debt represents the gap between realized temperatures and the temperature the world would experience given the carbon stock in the absence of geoengineering. When $g = 0$, ecological debt is zero, so there is no masked warming. When $g > 0$, ecological debt tracks the carbon stock; it increases as long as carbon is emitted, and decreases if carbon is removed from the atmosphere. If deployment ever stops, the world experiences a termination shock as temperatures jump to the unmasked level. The shock is fast and severe because of a mismatch in timescales: carbon dioxide remains in the atmosphere for centuries, but stratospheric aerosols dissipate within one to two years, so temperature snaps back rather than rising gradually as in baseline climate change.

Assumption 3 (Termination Hazard)**.** *The termination hazard rate is:*

$$\lambda(S) = \bar{\lambda} \cdot \mathbf{1}\{g(S) > 0\} \tag{7}$$

constant at $\bar{\lambda} > 0$ *while a program is deployed and zero otherwise: with no program, there is nothing to terminate.*

For simplicity, we assume the hazard rate is constant, so the probability of termination does not depend on the size of the program or the state of the climate. However, the damage from termination increases with the accumulated ecological debt. These are all our results need: a positive probability of termination and losses conditional on termination that increase with ecological debt. We will show that even though the physical hazard is flat, the market price of termination risk increases with the carbon stock. We calibrate $\bar{\lambda} = 0.5\%$ per year, which implies an expected program lifetime of 200 years, and we report sensitivity analysis

[8]The globally coordinated carbon tax is itself an optimistic assumption that does not match current reality. Solving the model as if one did can still revealing, because it shows that the problems we identify arise even when the cost of carbon is fully internalized.

in lifetime units. In robustness checks, we also consider a rising hazard $\lambda_0 + \lambda_1 g(S)$, which could mirror a world in which more aggressive deployment provokes stronger geopolitical opposition.

Termination in our model is absorbing: after a termination, geoengineering is never redeployed. We make this assumption mainly for tractability, but it may also be realistic. A program can end for many reasons, including political opposition, war, or technical accident. If opposition caused the halt, the same forces would stand in the way of redeployment. And whatever the cause, a world that has just suffered catastrophic losses may lack the capacity or the political will to redeploy.

The permanent productivity loss upon termination is:

$$\Gamma(F) = 1 - \exp\left(-\Gamma_0 F^{\beta}\right) \tag{8}$$

where $\Gamma_0 > 0$ scales the severity of termination damage and $\beta > 1$ determines its convexity. The functional form has three natural properties. If there is no ecological debt, termination causes no damage: $\Gamma(0) = 0$. Damage increases with the debt: $\Gamma' > 0$. And damage cannot exceed total productivity: $\Gamma(F) \to 1$ as $F \to \infty$. $\beta > 1$ implies that damage is convex at low levels of debt, so terminating a young program with little masked warming would do limited harm. However, terminating a mature program with a large accumulated debt would be catastrophic.

Aggregate productivity A_t is constant except at termination:

$$\frac{dA_t}{A_t} = -\Gamma(F_t)\, dN_t \tag{9}$$

where dN_t is a Poisson jump process with hazard rate $\lambda(S_t)$. When termination occurs ($dN_t = 1$), a fraction $\Gamma(F_t)$ of productivity is permanently destroyed. After termination, geoengineering ceases ($g = 0$, $\lambda = 0$) and the economy continues as a standard climate economy at the inherited carbon stock, with lower productivity. Climate damages, abatement costs, and the other costs in the model reduce consumption directly as proportional wedges on output (Section 2.4). We abstract from trend growth and other sources of productivity risk for simplicity.

Termination in our model is meant to represent a broad range of potential failure modes for geoengineering. A program may terminate because deployment causes (or is about to cause) negative side-effects, such as disruption of regional monsoons, damage to the ozone layer, or crop losses (Robock et al., 2009; Proctor et al., 2018), and the deployer or a global coalition chooses to stop. A geoengineering program may also end through misattribution

and backlash. If a major weather disaster is blamed on geoengineering, political support can collapse whether or not the blame is accurate (Parker and Irvine, 2018). Other reasons for termination includes wars, sanctions, or losing the ability to fund a program that needs constant operation. Regardless of the exact cause, termination means a catastrophic snapback in temperatures, as modeled in our paper.[9]

2.3 Firms and Dividends

We consider three types of firms: Brown (fossil fuel), Green (renewable energy), and CDR (carbon dioxide removal). Instead of modeling the detailed industry structure of energy markets, we work with reduced-form dividend functions that depend on the prevailing carbon tax.

Assumption 4 (Brown Dividends)**.** *Brown firm dividends $D^B(\tau)$ are strictly decreasing in the carbon tax τ up to a finite shutdown price $\bar{\tau}_{shut}$, with $D^B(0) > 0$ and $D^B(\tau) = 0$ for $\tau \geq \bar{\tau}_{shut}$.*

A higher carbon tax reduces fossil fuel profitability by raising production costs and shifting energy demand toward cleaner substitutes.

Assumption 5 (Green Dividends)**.** *Green firm dividends $D^G(\tau)$ are strictly increasing in the carbon tax τ up to a saturation price τ_{cap}, with $D^G(0) \geq 0$ and D^G constant beyond τ_{cap}.*

A higher carbon tax increases the demand for renewable energy and therefore the profit of Green firms, until they have fully replaced fossil energy and there is no longer any demand left to capture.

We impose these bounds as restrictions but note that our results rely only on two assumptions: dividends are monotone in the tax, and the Brown shutdown price is finite. We can interpret these reduced form dividends as describing competitive energy markets where Brown and Green energy are substitute inputs into production.

CDR firms operate physical carbon removal infrastructure, such as direct air capture plants, and earn revenue at the prevailing carbon price. Because the scale of current carbon removal is extremely small (see Section 2.1), we value the sector through the shadow value of a marginal unit of removal capacity, denoted $q(S)$: the present value of the carbon-price revenue that one unit of capacity would earn, net of depreciation, if it existed. This is well

[9] A limitation of our model is that it does not capture the variation potential losses associated with different types termination. If termination is caused by real side effects, a program harmful enough to be shut down would also plausibly impose costs while operating, and our consumption measure omits them. Adding those costs would make geoengineering less attractive in every state. If termination follows from misattributed blame, no such costs exist, and the model is complete as it stands.

defined even if there is no installed capacity. Section 2.6 extends the valuation to the case where the removal technology arrives stochastically.

After termination, all firms operate in an economy with permanently lower productivity: dividends of any firm i are scaled by $(1 - \Gamma(F))$ relative to what the firm would earn in a baseline economy at the same carbon stock.

2.4 The Planner's Problem

A representative global investor has CRRA preferences with risk aversion $\gamma > 1$ and discount rate $\rho > 0$:

$$\mathbb{E}_0 \int_0^\infty e^{-\rho t} \frac{C_t^{1-\gamma}}{1-\gamma} \, dt \tag{10}$$

Aggregate consumption equals output net of the costs the economy incurs:

$$C_t = A_t \, c_t, \qquad c_t = (1 - \Omega(T_t) - \Omega_{acid}(S_t)) \left(1 - \frac{a_t^2}{2\psi}\right) - x_t \tag{11}$$

Output is reduced by climate damages and by the resource cost of abatement. Abatement $a_t = \bar{e} - e_t$ is the reduction of emissions below the unabated level $\bar{e}$, with quadratic cost $a^2/(2\psi)$. ψ determines the cost of cutting emissions. The marginal abatement cost a/ψ equals the carbon tax in equilibrium. x_t is the cost of spending on carbon removal after a breakthrough, and $x_t = 0$ before. We omit research costs from the resource constraint because R&D is a small share of output.

A global social planner sets the carbon tax τ_t equal to the Social Cost of Carbon (SCC), taking Region V's geoengineering policy $g(S)$ as given. The coordinated tax is an optimistic assumption: decentralized national carbon taxes would be lower than the global SCC, and carbon would accumulate faster.

Homogeneity

Because flow utility is CRRA and every cost enters as a proportional wedge on output, the value function is homogeneous of degree $1 - \gamma$ in productivity:

$$V(S, A) = A^{1-\gamma} \, v(S) \tag{12}$$

where $v(S)$ solves the reduced HJB equation below. Under homogeneity, the planner's problem can be simplified to a one-dimensional ODE in S, and a permanent productivity loss of Γ at termination enters the continuation value as the factor $(1 - \Gamma)^{1-\gamma}$.

The HJB Equation: Incremental Construction

We construct the HJB equation in three nested regimes, each adding one mechanism. $u(c) = c^{1-\gamma}/(1-\gamma)$ represents flow utility per unit of $A^{1-\gamma}$ and $c(S,\tau)$ represents normalized consumption in state S under tax τ, from Equation (11).

Regime 1: No Geoengineering. Set $g \equiv 0$, so $F = 0$, $\lambda = 0$, and $T = \Psi(S)$. This is the standard climate economy without geoengineering. The reduced HJB for $v(S)$ is:

$$\rho\, v = \max_{\tau} \left\{ u(c(S,\tau)) + e(\tau)\, v' + \tfrac{1}{2}\sigma_S^2\, v'' \right\} \tag{13}$$

where primes denote derivatives with respect to S. The first-order condition shows that the planner chooses τ such that the marginal abatement cost is equal to the marginal value of a lower carbon stock (converted to consumption units):

$$\tau = SCC^{R1}(S) = \frac{-v'(S)}{u'(c(S,\tau))\,(1-\Omega-\Omega_{acid})} \tag{14}$$

In this regime, a higher carbon stock raises raises raises standard temperature damages and acidification damages, but does not affect termination damages.

Regime 2: Geoengineering Without Termination Risk. Geoengineering is activated $g(S) > 0$, but termination risk is zero ($\lambda \equiv 0$). Realized temperature is $T = \Psi(S) - \theta g(S)$. The HJB becomes:

$$\rho\, v = \max_{\tau} \left\{ u(c(S,\tau)) + e(\tau)\, v' + \tfrac{1}{2}\sigma_S^2\, v'' \right\}, \qquad T = \Psi(S) - \theta g(S) \tag{15}$$

The difference between Regimes 1 and 2 is that temperature damages Ω are evaluated at the artificially low realized temperature rather than at $\Psi(S)$. Acidification damages do not differ from Regime 1 because they depend directly on the carbon stock. This decline in temperature damages reduces the planner's incentive to tax carbon (the moral hazard channel).

Regime 3: The Full Model. Activate $\lambda(S) > 0$ and $\Gamma(F)\, dN_t$. The HJB acquires a jump term reflecting the expected value loss from termination:

$$\begin{aligned} \rho\, v = \max_{\tau} \Big\{ & u(c(S,\tau)) + e(\tau)\, v' + \tfrac{1}{2}\sigma_S^2\, v'' \\ & + \lambda(S) \left[(1-\Gamma(F(S)))^{1-\gamma}\, v^{R1}(S) - v \right] \Big\} \end{aligned} \tag{16}$$

The new term is the expected value change from a termination event. After termination, the economy moves immediately to Regime 1 with the same carbon stock and permanently lower productivity. The continuation value is the Regime 1 value $v^{R1}(S)$, scaled by the factor $(1-\Gamma(F))^{1-\gamma}$. Under this setup, the world after deployment faces a small chance each instant of a large permanent loss, and this possibility lowers value throughout deployment.

The SCC Under the Full Model

The first-order condition of the Regime 3 HJB yields the full Social Cost of Carbon, which can be decomposed into two channels:

$$SCC^{R3}(S) = \underbrace{SCC^{flow}(S)}_{\text{Flow Damage Channel}} + \underbrace{SCC^{term}(S)}_{\text{Termination Fragility Channel}} \tag{17}$$

The flow damage channel represents the standard marginal cost of carbon. After geoengineering deployment via a thermostat policy, realized temperature will be held constant at T^*, so an additional unit of carbon causes no additional temperature damage. The only flow cost remaining at the margin is acidification. The termination fragility channel is new. An additional unit of carbon raises $g(S)$, increasing ecological debt $F(S)$ and the damage $\Gamma(F)$ if termination occurs. The full expression takes the same form as Equation (14):

$$SCC^{R3}(S) = \frac{-v'(S)}{u'(c(S,\tau))\,(1-\Omega-\Omega_{acid})} \tag{18}$$

where $v(S)$ now solves the Regime 3 HJB, so the slope of the value function includes both flow damages and termination risk.

We solve the HJB numerically by value function iteration on a grid over S (Section 4). The resulting SCC is forward looking: it accounts for the future path of carbon accumulation, geoengineering deployment, and termination risk. Note that, once the carbon price is high enough to drive emissions to zero, raising it further can do no more to reduce emissions, so we report the carbon price in such states as the shutdown level ($\bar{e}/\psi$). The marginal harm from emissions can continue to rise beyond this point, but no tax above the shutdown level changes behavior.

2.5 Breakthroughs and Removal Capacity

Two technologies can arrive, each as a Poisson event, independent of each other and of the termination process:

- A green cost collapse $N_{G,t}$, at hazard $\eta_G(S)$. The abatement cost parameter falls from ψ to ψ_L, such that a modest carbon tax now eliminates most emissions. Under the calibrated ψ_L, a \$75/tCO$_2$ tax cuts emissions by 90%. Nuclear fusion is the example of such a discrete collapse. Notes that a green cost collapse does not describe smooth cost declines associated with developments in solar and wind, as these are incremental, already underway, and already part of the baseline emissions function $e(\tau)$, not of N_G.

- A carbon removal breakthrough $N_{C,t}$, at hazard $\eta_C(S)$. A meaningful carbon removal sector comes into existence. Our calibrations assume this sector has capacity $\bar{R} = 15$ GtCO$_2$/yr (1.92 ppm/yr) at marginal cost c_R = \$100/tCO$_2$. The values of these calibration parameters are motivated by portfolio assessments (National Academies of Sciences, Engineering, and Medicine, 2019) that cite roughly 10 GtCO$_2$/yr as a successful mid-century target and 20 GtCO$_2$/yr as an aggressive upper bound. Our chosen parameters lie in the middle, and we show our sensitivity to each.[10]

We model removal as a fixed capacity at a fixed cost per ton, instead of an upward-sloping supply curve. This setup reflects the engineering literature on removal technologies, and the capacity cap prevents us from generating unrealistically large amounts of removal in response to a sudden jump in carbon prices after a termination. The cap also determines a realistic speed of the exit. Even at the maximum capacity of the carbon removal sector, removing 1.92 ppm of carbon per year means it will take decades to unwind a century of buildup. For simplicity, we assume that the removal capacity is immediately at full strength when the breakthrough occurs. This assumption favors the case for geoengineering. Real ramp up the removal industry from today's negligible megatons to 15 gigatons per year could take decades.

Operation and procurement. We assume that after the CDR breakthrough, the planner has bought all the removal capacity and decides when to run it. The planner runs capacity at rate $r \in [0, \bar{R}]$ if doing so increases value, paying c_R per ton from consumption. We make this assumption because leaving removal in private hands would have an incentive problem. A private operator runs capacity only if the carbon price covers the removal cost. But removing carbon reduces the ecological debt, which reduces the social cost of carbon, reducing the carbon price as removal occurs. The carbon price may fail to cover the removal

[10]The removal cost assumption of \$100 per ton in generous in the sense that direct air capture today costs several times \$100 per ton, and \$100 per ton is a common long-run goal for a mature carbon removal industry, at the low end of assessed ranges (National Academies of Sciences, Engineering, and Medicine, 2019). We report sensitivity to \$150 per ton in robustness analysis. Potential carbon removal breakthroughs include direct air capture (5–10 Gt/yr at scale), enhanced rock weathering (2–4), ocean alkalinity enhancement (1–10), and BECCS (2–5).

cost before the carbon stock becomes safe, despite each removed ton reducing the severity of a possible termination. An advance commitment by the planner to buy removal at cost c_R solves this problem by keeping capacity running according to the planner's rule regardless of the carbon price. The commitment does not pay any rents to the innovator, and we report its expenditures as a share of output. Section 4.6 indicates the price path that motivates the commitment.[11]

Effective safety. Termination damage depends on the size of the ecological debt, and at small debt the damage is negligible, so the planner may stop removing carbon while some debt remains: the small residual risk is cheaper to live with than the cost of removing more carbon. We therefore need a benchmark for when the danger has effectively passed, and we set it at a termination loss of 1%. We define *effective safety* as the carbon stock S_{eff} at which termination, if it occurred, would destroy only 1% of productivity. Formally, $S_{\text{eff}} = \Psi^{-1}(T^* + F_1)$, where $F_1 = (-\ln(0.99)/\Gamma_0)^{1/\beta} = 0.53°\text{C}$ is the debt at which the conditional loss equals 1%; at the baseline severity this gives $S_{\text{eff}} = 600$ ppm. Throughout the paper, *rescue* means bringing the carbon stock down to effective safety before termination occurs.[12]

2.6 Innovation

We model innovation so that it responds to financial incentives. We do not model researchers' decisions directly. Instead, we assume a reduced-form in which research responds to incentives. The rate at which each breakthrough arrives is increasing in the value of the patent its innovator would receive, linearly up to a cap:

$$\eta_i(S) = \min\left(\kappa_i \max\left(\hat{\pi}_i(S), 0\right),\ \eta_{max}\right), \qquad i \in \{G, C\} \tag{19}$$

A free-entry model with quadratic research costs can delivers this linear reduced form. We impose a cap of $\eta_{max} = 25\%$ per year, which translates to a minimum expected arrival time of 4 years. This reflects the time it takes to build a new energy or removal system.

The patent value $\hat{\pi}_i(S)$ is what the innovator's claim would be worth if breakthrough i occurred when the carbon stock is S, measured per unit of productivity A. Like all assets

[11] Advance commitment is needed for the green breakthrough, because the green sector sells energy at market prices and does not depend on the carbon price for its revenue.

[12] Removing carbon all the way down to the activation threshold would remove all ecological debt, end deployment, and eliminate termination risk entirely. We call that absolute safety. The planner stops at effective safety instead of absolute safety, because removing the remaining carbon costs more than it is worth once the conditional loss is down to 1%.

in the model, we compute it from the economy that prevails after the breakthrough and price it using the model's stochastic discount factor (SDF). $\hat{\pi}_C$ is the market value of a CDR developer, which we return to when we take the model to the cross-section of firms. Measuring the patent per unit of A keeps research incentives stable as the economy grows.[13]

Geoengineering suppresses innovation through lower expected future carbon taxes. The reward to the innovator depends on the subsequent path of carbon taxes after the breakthrough. After a breakthrough, the economy will continue to deploy, which depresses the carbon tax further in a world with geoengineering. This means that the innovator gets less in each future state, the patent is worth less, and there is less research as a result. The advance commitment for CDR described in the previous section does not affect this. Advance commitment ensures that carbon removal activity continues to operate, but it buys removal services at cost, so it contributes nothing to the incentive to develop the breakthrough in the first place.

The green patent earns most of its value from flow profits, while the removal patent earns most of its value from termination.

- $\hat{\pi}_G(S)$ is the value of the green patent. After a green breakthrough, the innovator earns the green sector's profits, which flow at rate $D^G(\tau)$. If termination occurs, the carbon tax rises sharply, which shifts energy demand away from fossil fuels, leading to an increase in green profits. But Assumption 5 caps green dividends at $D^G(\tau_{cap})$, so the value of a permanently high tax is at most $D^G(\tau_{cap})/r$.

- $\hat{\pi}_C(S)$ is the value of the removal patent. After a CDR breakthrough, the removal sector earns the difference between the carbon price over cost, $(\tau - c_R)^+\bar{R}$ per unit time. Profits for removal are close to zero before termination because the advance commitment pays exactly removal costs when the tax is below c_R, and geoengineering suppresses the carbon tax toward or below c_R. The main value of the removal patent comes from termination. After termination, the carbon tax grows beyond c_R and removal earns $\tau - c_R$ for every ton removed (with the flow of removal bounded by the capacity $\bar{R}$). Thus, the value of the removal patent grows with ecological debt.

Both types of patents gain from termination, but the green patent's gain is capped at $D^G(\tau_{cap})/r$, while the gain of the removal patent grows with the ecological debt. As debt grows, research incentives shift toward removal.

[13] We make an approximation below the activation threshold. If we want to compute a patent value there, we must know how the carbon stock evolves as we wait for the breakthrough. But that depends on the patent values we are trying to compute. We break the circle by using the carbon dynamics of the model in which the breakthroughs cannot arrive. In this region, the two versions indicate nearly the same carbon path, differing by roughly 0.1 to 0.2 ppm per year.

2.7 Solving the Full Model

The full model consists of eight economies, one for each combination of technology state (neither breakthrough, green only, removal only, or both) and termination status. Each random event moves the world from one economy to another. A breakthrough adds its technology, and a termination moves the economy to its post-termination counterpart. Each economy is an optimizing problem of the same form as Section 2.4, with the economies it can jump to as continuation values. We solve by backward induction, starting from the economies that no further event can leave. No circularity arises, because nothing that happens after a breakthrough depends on the taxes, values, or research intensities that prevailed before it.

1. *Post-termination economies*, one for each technology state at the moment of termination (neither breakthrough, green only, removal only, or both). If there are breakthroughs in hand, they will be used. In particular, a post-termination economy with removal available uses it to reduce the carbon stock. We assume research ceases at termination because a breakdown severe enough to require termination also disrupts the innovation system. This assumption impacts how costly termination is, and we present the variant in which innovation does not cease after termination as a robustness check.

2. *The economy with both breakthroughs*, v_B. Removal operates under the advance commitment agreement, and the termination hazard remains as long as $g > 0$.

3. *The economy with removal only*, v_C. Removal lowers the carbon stock, g declines along the thermostat as the stock of carbon falls, and the termination hazard remains as long as $g > 0$. A green breakthrough can still occur in this economy and moves it to v_B.

4. *The economy with the green collapse only*, v_G. Cheap green energy reduces the flow of new emissions, but it does not remove the carbon already in the atmosphere. The stock, the masked warming, and the termination hazard therefore all remain. A CDR breakthrough can still occur in this economy and moves it to v_B.

5. *The pre-breakthrough regimes* 1, 2, and 3, solved last, with the research intensities $\eta_G(S)$ and $\eta_C(S)$ built from the economies in steps 1–4.

The last step of the backward induction is the pre-breakthrough problem in item 5. The Regime 3 problem adds two jump terms to Equation (16), one for each breakthrough:

$$\begin{aligned}\rho\, v = \max_{\tau} \Big\{ & u(c(S,\tau)) + e(\tau)\, v' + \tfrac{1}{2}\sigma_S^2\, v'' + \lambda(S)\left[(1-\Gamma(F(S)))^{1-\gamma}\, v^{post}(S) - v\right] \\ & + \eta_G(S)\left[v_G(S) - v\right] + \eta_C(S)\left[v_C(S) - v\right] \Big\}\end{aligned} \tag{20}$$

If a breakthrough occurs, the economy jumps to the corresponding continuation value defined above, v_G or v_C. The planner takes the research intensities $\eta_i(S)$ as given, exactly as she takes $g(S)$ as given. The termination destination $v^{post}(S)$ is the post-termination economy with research shut down; in the version of the model without breakthroughs it coincides with v^{R1}, which is why Equation (16) wrote it that way. Regimes 1 and 2 are the same equation with the deployment and termination terms shut off, each with research intensities computed from its own continuation economies.

Setting $\kappa_G = \kappa_C = 0$ shuts down research entirely, leaving a world in which no breakthrough will ever occur and today's technologies are permanent. We call this special case the *no-innovation benchmark*. Comparing the full model against the no-innovation benchmark allows us to see what endogenous innovation adds.

2.8 Asset Pricing

The Stochastic Discount Factor

Under CRRA utility, the stochastic discount factor evolves as:

$$\frac{d\Lambda_t}{\Lambda_t} = -r_t\,dt - \gamma\eta_t\,dZ_t + \sum_{j\in\{term,G,C\}} \left[(J^{\Lambda}_{j,t} - 1)\,dN_{j,t} - h_{j,t}(J^{\Lambda}_{j,t} - 1)\,dt\right] \tag{21}$$

where the sum has three terms, one for each event that can jump the economy: termination (hazard $h_{term} = \lambda$), a green breakthrough ($h_G = \eta_G$), and a CDR breakthrough ($h_C = \eta_C$). The coefficient η_t is the compensation investors require for bearing the random year-to-year fluctuations of the carbon stock in Equation (1); these fluctuations are small, and we set $\eta_t = 0$ below. In the no-innovation benchmark, only the termination term is present. For each event, the jump in the pricing kernel is the marginal utility ratio across the event:

$$J^{\Lambda}_{j,t} = \left(\frac{C_{t+}}{C_{t-}}\right)^{-\gamma} \tag{22}$$

At termination, consumption falls for two reasons. Productivity is destroyed by the factor $1 - \Gamma(F_t)$, and the post-termination tax forces a further contraction as the economy pays for a burst of abatement. Writing $J^c_t \equiv c^{R1}(S_t)/c(S_t)$ for the second effect, the consumption jump is $C_{t+}/C_{t-} = (1 - \Gamma(F_t)) \cdot J^c_t$, and both factors push the pricing-kernel jump above one.

Risk-Neutral Pricing

The change of measure to the risk-neutral probability $\mathbb{Q}$ scales each physical hazard by its pricing-kernel jump. For termination:

$$\lambda_t^{\mathbb{Q}} = \lambda_t \cdot J_t^{\Lambda} = \lambda_t \cdot [(1 - \Gamma(F_t)) \cdot J_t^c]^{-\gamma} \tag{23}$$

This is the rare-disaster mechanism of Barro (2006) applied to the termination shock. Termination arrives when marginal utility is high, so termination is priced as if it were more likely than it actually is. The physical hazard is constant while a program is deployed, so when $\lambda^{\mathbb{Q}}(S)$ rises with the carbon stock, it is because the damage conditional on termination has grown.

The same change of measure applies to the two potential breakthroughs. A breakthrough always raises welfare, but it can lower consumption at the moment it occurs, because the economy immediately puts resources into the new opportunity. After a green breakthrough, firms facing the carbon tax spend on newly cheap abatement. After a CDR breakthrough, the advance commitment begins paying for removal. Depending on the state, the event is therefore good or bad news for consumption, and its risk-neutral hazard is adjusted down or up accordingly; we discuss the signs with the results (Section 4.7).

Valuation Equations

As noted above, we set the compensation for carbon-stock fluctuations to zero, $\eta_t = 0$. The fluctuations are small ($\sigma_S = 1.0$ ppm$/\sqrt{\text{yr}}$), so consumption's exposure to them is second order next to the jumps, and jump risk carries all of the risk adjustment.[14] Each firm's value then solves an ODE on the carbon-stock grid, discounting dividends at r, carrying the physical drift and diffusion of S, and loading each jump at its risk-neutral hazard.

$$rV^i(S) = D^i(\tau(S)) + \mu_S(S)V^{i\prime}(S) + \tfrac{1}{2}\sigma_S^2 V^{i\prime\prime}(S) + \sum_j h_j^{\mathbb{Q}}(S)\left[V_j^i(S) - V^i(S)\right] \tag{24}$$

Here $\mu_S(S)$ is the equilibrium drift of the carbon stock and $V_j^i(S)$ is firm i's value in the economy that event j leads to. At termination, this value is scaled by the productivity loss, $V_{post}^i(S) = (1 - \Gamma(F(S))) \cdot V^{i,R1}(S)$.

The value of one unit of carbon removal capacity, $q(S)$ from Section 2.3, solves the same equation with two changes. Its dividend is the carbon-price revenue the unit would earn, $\tau(S)\phi$, and it is discounted at $r + \delta$ because removal capacity depreciates. Note that we

[14] A full treatment would add the covariance term $-\gamma\eta_t\sigma_S V^{i\prime}(S)$ to Equation (24).

compute q from the unit's revenue without subtracting the operating cost c_R, while the patent value $\hat{\pi}_C$ in Section 2.6 is computed from profits net of that cost, $(\tau - c_R)^+ \bar{R}$. The reason for this difference is that subtracting the cost would reduce q whenever the carbon price is less than c_R.[15]

In Regimes 1 and 2 the termination jump vanishes ($\lambda = 0$). In the no-innovation benchmark the breakthrough jumps vanish as well, and asset values reduce to perpetuity-like solutions of their dividend flows.

Termination Risk Premia

The expected excess return of any firm i satisfies:

$$\mu^i - r = \underbrace{\gamma\eta\,\sigma_{i,Z}}_{\text{Diffusion premium (=0 here)}} + \underbrace{TRP^i}_{\text{Jump risk premia}} \tag{25}$$

where the jump risk premium sums the three events:

$$TRP^i = -\sum_j h_j \left(J_j^\Lambda - 1\right)\left(J_j^i - 1\right) \tag{26}$$

and $J_j^i \equiv V_j^i / V^i$ is the proportional jump in firm i's valuation at event j. The termination component represents compensation investors demand for bearing termination risk. Its sign depends on J^i. Firms that lose value at termination ($J^i < 1$) covary negatively with the jump in the stochastic discount factor, so they have a positive risk premium. Firms that gain value at termination ($J^i > 1$) have a negative risk premium. The two breakthrough components price exposure to breakthrough news, and Section 4.7 reports the full decomposition.

Three Nested Regimes

We analyze the economy under three nested parameter restrictions, each isolating a distinct economic mechanism:

	Regime 1 (No geo.)	Regime 2 (Geo, no term.)	Regime 3 (Full)
$g(S)$	$\equiv 0$	> 0 (Assumption 2)	> 0 (Assumption 2)
$\lambda(S)$	$\equiv 0$	$\equiv 0$	> 0 (Assumption 3)
$\Gamma(F)\,dN_t$	absent	absent	active
Arrivals N_G, N_C	live	live	live

[15] Intuitively, with termination as the only jump and ignoring the drift and diffusion terms, the solution is $V^i(S) = [D^i(\tau(S)) + \lambda^{\mathbb{Q}}(S) V^i_{post}(S)]/(r + \lambda^{\mathbb{Q}}(S))$, which is the dividend flow discounted at the jump-adjusted rate plus the probability-weighted post-termination value.

In all three regimes, both breakthroughs can occur. The rates at which they occur, $\eta_G(S)$ and $\eta_C(S)$ depend each regime's patent values (Section 2.6), such that a regime that suppresses the carbon tax also generates slower research. Setting $\kappa_G = \kappa_C = 0$ turns research off, and represents our no-innovation benchmark. This completes the model. We now derive the theoretical results under each regime.

3 Theoretical Results Under Three Nested Regimes

We present our theoretical results in two parts. First, we state and prove Propositions 1 - 3 in the no-innovation benchmark (research off). Then we turn research on and derive the innovation results in the full model in Section 3.4. We present calibration numbers in the later quantitative sections using the full model, but we use the no-innovation benchmark to show how results change with innovation.

3.1 Regime 1: Baseline Climate Risk

Let $g \equiv 0$, so $F = 0$, $\lambda = 0$, and $T = \Psi(S)$. This regime describes the standard climate economy without geoengineering. The SCC is strictly positive and increasing in the carbon stock, generating a gradual transition from Brown to Green energy. Because dividends are monotone in the tax (Assumptions 4–5), as the stock of carbon grows, Brown valuations fall, Green valuations rise, and the value of a unit of removal capacity rises.

3.2 Regime 2: The Geoengineering Put

Let $g(S) > 0$ (Assumption 2) but keep $\lambda \equiv 0$. In this regime, geoengineering operates with no risk of termination. The HJB is Equation (15). The only change from Regime 1 is that temperature damages are a function of the low realized temperature $\Psi(S) - \theta g(S)$.

Proposition 1 (The Geoengineering Put: Moral Hazard)**.**

(a) The SCC in Regime 2 is strictly lower than in Regime 1, $SCC^{R2}(S) < SCC^{R1}(S)$, for all S where $g(S) > 0$. Geoengineering lowers realized temperature, which lowers damages of each additional unit of carbon.

(b) Under the thermostat policy (Equation 5), realized temperature is constant at T^* above the activation threshold, so an additional unit of carbon causes no additional temperature damage. The Regime 2 carbon tax only contains the acidification component, $SCC^{R2}(S) = \tau_{acid}(S)$.

(c) The lower carbon tax raises Brown valuations relative to Regime 1, $V^{B,R2}(S) > V^{B,R1}(S)$.

(d) The lower carbon tax reduces Green valuations relative to Regime 1, $V^{G,R2}(S) < V^{G,R1}(S)$.

(e) Above the activation threshold, carbon-stock fluctuations no longer move realized temperature, so temperature risk does not impact asset prices. The stock of carbon affects valuations only through acidification damages, which cannot be masked.

Proof sketch. We compare the Regime 1 and Regime 2 HJBs (Equations 13–15). Because $\Omega(\Psi(S) - \theta g) < \Omega(\Psi(S))$, the value function declines less steeply in S under Regime 2, so the planner is less concerned about carbon accumulation and the SCC is lower. Part (b) follows because $T = T^*$ above activation, so only the acidification term contributes to $v'(S)$. Parts (c) and (d) follow from Assumptions 4–5. □

Part (e) demonstrates the key moral hazard concern. If the geoengineering put exists with no termination risk, temperature risk does not matter for asset prices. The only risks in Regime 2 are the timing of the two breakthrough technologies (in the no-innovation benchmark, there is no priced risk at all).

Without termination risk, geoengineering also lowers the value of carbon removal. A unit of removal capacity earns the carbon price and geoengineering suppresses that price, so $q^{R2}(S) < q^{R1}(S)$ everywhere.

3.3 Regime 3: Termination Risk and the Toxic Put

Let termination risk $\lambda(S) > 0$ (Assumption 3), with productivity destruction $\Gamma(F)\, dN_t$. The HJB is Equation (16).

Proposition 2 (The Social Cost of Carbon Under Termination Risk)**.**

(a) For sufficiently small λ, $SCC^{R2}(S) < SCC^{R3}(S) < SCC^{R1}(S)$. *Termination risk raises the carbon tax relative to Regime 2, but the tax remains below the no-geoengineering level.*

(b) Under the thermostat policy, the SCC above the activation threshold equals an acidification component plus a termination component, $SCC^{R3}(S) = \tau_{acid}(S) + SCC^{term}(S)$. *Carbon is priced not for flow temperature damages, which are masked, but because each additional unit of carbon raises* $g(S)$, *which increases ecological debt and the damage if termination occurs.*

Proof sketch. The Regime 3 HJB contains an additional negative term $\lambda[(1-\Gamma)^{1-\gamma}v^{R1}-v]$, which steepens $v'(S)$ relative to Regime 2 and raises the SCC. The SCC remains below Regime 1 because flow temperature damages are lower. Part (b) follows from the same argument as Proposition 1(b). □

The carbon price in Regime 3 depends on two channels. Within the deployment region the price rises with the carbon stock, because deeper debt implies greater damage conditional on termination. At the activation threshold, the price drops discretely, because temperature flow damages change to zero, and just below activation it declines as the deployment era approaches (Figure 2). Under the baseline calibration, the Regime 3 price also crosses the Regime 1 price at a stock $\hat{S}$, quantified in Section 4.2. Above $\hat{S}$, the marginal unit of carbon is more socially costly under geoengineering than in a world without it, because the catastrophe embedded in the debt outweighs the flow damages that geoengineering suppresses.

The results below use two facts about the termination event. Let $\bar{\tau}(S) \equiv SCC^{R1}(S)$ be the post-termination carbon tax. This is the Regime 1 tax, evaluated at the current carbon stock. In the moderate-debt region, $\bar{\tau}(S)$ exceeds the tax that prevails while geoengineering operates, so termination brings about a sharp tax increase, and this drives the risk premium signs in Proposition 3. The ordering reverses in the high-debt region. In our baseline calibration, the Regime 3 tax crosses above the Regime 1 tax at a stock $\hat{S}$, and above $\hat{S}$ termination lowers the tax.

Proposition 3 (The Toxic Put: The Cross-Section of Termination Risk Premia)**.** *For S such that Brown firms earn positive pre-termination dividends ($\tau^{R3}(S) < \bar{\tau}_{shut}$):*

(a) Brown firms lose value at termination, because the post-termination tax sharply reduces their dividends ($J^B < 1$). They carry a strictly positive termination risk premium, which increases with the carbon stock.

(b) At moderate carbon stocks, Green firms gain value at termination, because the tax increase shifts energy demand toward them ($J^G > 1$). Their termination risk premium is negative. Green firms are disaster insurance.

(c) At carbon stocks just above the activation threshold, removal capacity gains value at termination, because the post-termination price of carbon far exceeds the price that prevails while geoengineering operates ($J^{CDR} > 1$). Removal capacity is a termination hedge.

Proof sketch. The sign of each firm's valuation jump at termination determines the sign of its premium through Equation (26). The post-termination tax exceeds the pre-termination tax in this region, and dividends are monotone in the tax by Assumptions 4–5. □

Parts (b) and (c) hold in the stated regions. In Section 4.7, we report regions where the signs reverse.

Geoengineering first raises and then lowers the value of Brown firms. At low carbon stocks, near the onset of geoengineering, the ecological debt is small, $\Gamma(F) \approx 0$, and the termination risk premium is negligible, so the dividend boost from the low carbon tax dominates and geoengineering raises Brown value. In this phase, geoengineering acts as a put option for Brown firms, protecting their cash flows from carbon taxation. But as carbon accumulates, the termination risk premium grows with the ecological debt while the dividend benefit is bounded, since D^B can never exceed $D^B(0)$ no matter how low the tax falls. In our baseline calibration the two Brown valuation curves therefore cross at a threshold S^*, quantified in Section 4.3, and the threshold falls as the hazard or the severity rises. Above S^*, geoengineering lowers Brown value below its Regime 1 level, $V^{B,R3}(S) < V^{B,R1}(S)$. The put has become toxic. Investors holding Brown equity above S^* collect elevated dividends in exchange for bearing large losses at termination.

Counterintuitively, the premium changes sign if the carbon stock reaches very high levels. If the Regime 3 tax reaches the shutdown price, then Brown dividends are zero before termination. Thus, termination, which shifts taxes to the lower Regime 1 tax, helps Brown firms rather than harming them. In this region, Brown firms have value because of the option that termination might occur. Whether the deployment era ever reaches this extreme is a quantitative question, taken up in Section 4.7.

In Regime 2, geoengineering and carbon removal are substitutes, meaning that the availability of geoengineering makes carbon removal less valuable ($q^{R2} < q^{R1}$). Termination risk reverses this substitutability. The Regime 3 value of removal capacity contains a term that Regime 2 lacks, the value of removal capacity in the termination state. After termination, the warming from the accumulated carbon stock is fully felt, so the carbon price jumps to $\bar{\tau}$. Each unit of removal capacity then earns that higher price of carbon. Thus, $q^{R3}(S) > q^{R2}(S)$ at every carbon stock, and in our baseline calibration q^{R3} crosses above the no-geoengineering value q^{R1} at 603 ppm. Beyond that crossing, geoengineering and carbon removal are complements, meaning that the availability of geoengineering makes carbon removal more valuable. Without termination risk, geoengineering eliminates the motive for removal. With termination risk, geoengineering creates the motive, because the ecological debt it accumulates can be retired safely only by removing carbon. Green energy can slow the growth of new ecological debt but cannot remove the carbon already in the atmosphere. Geoengineering can therefore raise the value of carbon removal.

Summary of Cross-Regime Comparisons

Table 1 summarizes the qualitative predictions. Some comparisons are global sign results, and others hold in the regions described above. We distinguish three ranges of the carbon stock. At low carbon stocks, geoengineering raises Brown value and the premiums are small. At moderate carbon stocks, geoengineering lowers Brown firms value below its Regime 1 value (the toxic put) and raises the value of removal capacity above its Regime 1 level (geoengineering and carbon removal become complements). At high carbon stocks, the Regime 3 tax exceeds the Regime 1 tax, Brown dividends fall before termination rather than after it, and the premium signs invert. Green valuations fall in every range because the green gain at termination is bounded, while removal capacity is the one asset whose value at termination grows with the ecological debt. The two rightmost columns preview the innovation results, which are described in the next section.

Table 1: Cross-Regime Comparative Statics

	τ	V^B	V^G	q^{CDR}	η_G	η_C
Regime 2 vs. 1	$\downarrow^*$	$\uparrow^*$	$\downarrow^*$	$\downarrow^*$ (substitutes)	$\downarrow^*$	$\downarrow^*$
Regime 3 vs. 2	$\uparrow^*$	$\downarrow$	$\uparrow$	$\uparrow^*$	$\uparrow$	$\uparrow\uparrow$
Regime 3 vs. 1 (low S)	$\downarrow$	$\uparrow$	$\downarrow$	ambiguous	$\downarrow$	$\downarrow$
Regime 3 vs. 1 (mod. S)	$\downarrow$	$\downarrow$ (toxic)	$\downarrow$	$\uparrow$ (complements)	$\downarrow$	$\downarrow$
Regime 3 vs. 1 (high S)	$\uparrow$	$\downarrow$	$\downarrow$	$\uparrow$	$\downarrow$	$\uparrow$ (above S_η)

Notes: Entries marked * are global sign results established in the propositions and hold at every carbon stock where deployment is active. Unmarked entries hold in the indicated range of the carbon stock under our baseline calibration, with the ranges located in Section 4. The $\uparrow^*$ for τ in the second row is established for termination hazard rates $\bar{\lambda}$ up to the calibrated value.

3.4 Geoengineering, Termination Risk, and the Direction of Innovation

The results so far do not allow for innovation. This subsection turns research on and presents the paper's second set of theoretical results. Geoengineering lowers realized temperatures, buying time for a green energy or carbon removal breakthrough. However, geoengineering also weakens research incentives, because innovators are paid through carbon prices that deployment suppresses. We show that termination risk partially revives research incentives and directs it toward carbon removal.

Proposition 4 (Innovation Crowding-Out and Redirection)**.**

(a) Without termination risk, geoengineering suppresses both research directions, $\eta_i^{R2}(S) < \eta_i^{R1}(S)$ for $i \in \{G, C\}$, because the tax path that pays either innovator is lower in every state after the breakthrough. If the acidification tax remains below c_R along the entire path after a CDR breakthrough, removal capacity earns nothing in Regime 2, and the removal patent's value in deployment states is zero.

(b) Termination risk shifts research toward removal. The ratio η_C/η_G rises from Regime 2 to Regime 3 wherever the green patent's gain at termination is at its bound $D^G(\tau_{cap})/r$ while the removal patent's gain grows with the ecological debt.

Proof sketch. Each patent pays profits that rise with the carbon tax after its breakthrough. In Regime 2, geoengineering still follows $g(S)$ after either breakthrough, so the carbon tax remains suppressed in every state relative to Regime 1, and both patent values fall. For part (b), green profits at termination are bounded by the dividend cap regardless of the ecological debt. Carbon removal profits grow with ecological debt, because termination conditional on a higher carbon stock implies the Regime 1 carbon price starts further above the removal cost c_R and more carbon stock must be removed before the price falls back to c_R. The formal proof is presented in Appendix A. □

Part (a) is the innovation analogue of the moral hazard result. The removal patent is a claim on a carbon price whose temperature component geoengineering has eliminated. Part (b) rests on the asymmetry from Section 2.6. Both patents gain if termination occurs, but the two gains scale differently. The green gain is limited because green dividends are capped and because termination destroys a fraction Γ of the green sector's profits along with the rest of the economy. The carbon removal gain faces no such cap and grows with the ecological debt.

The proposition compares Regime 3 with Regime 2. Adding termination risk to a world with geoengineering shifts research toward removal. A different question is whether removal research under geoengineering is higher or lower than in Regime 1, with no geoengineering at all. That answer depends on the calibration which we present in Section 4.4. There we show that geoengineering depresses removal research below its no-geoengineering level throughout the phase in which it subsidizes fossil fuel firms.

Proposition 5 (Only Removal Opens an Exit)**.** *A CDR breakthrough is the only technology that lowers the carbon stock, thus making it the only thing that could allow the world to pay off ecological debt. A green breakthrough allows for cheap energy to stop our ecological debt from increasing, but this does not address the carbon stock, the masked warming, or the termination threat. Absent carbon-cycle noise, termination occurs with probability one if there is no CDR breakthrough.*

Proof sketch. The proof is immediate from the state dynamics. Only removal changes the sign of dS. In the economy with only the green breakthrough, the termination hazard is bounded below by $\bar{\lambda} > 0$ for all time, so absent a later CDR breakthrough, termination is an almost-sure event. □

This simple proposition inverts a common intuition. In most discussions of climate policy, a green energy breakthrough such as fusion solves the climate problem, because emissions go to zero and warming stops. In our model, this only works if the green breakthrough arrives before deployment begins. If geoengineering is already underway, stopping emissions is not enough, because the carbon already in the atmosphere keeps the geoengineering program running and the termination hazard alive. How likely rescue by a CDR breakthrough is once geoengineering is underway is a quantitative question that we take up in Section 4.5.

Welfare With and Without Geoengineering

The final question is whether the world is better off with geoengineering available, i.e., whether welfare is higher in Regime 3 than in Regime 1. In our model, deployment is not a policy choice by the planner. Instead, Region V deploys via $g(S)$ if geoengineering technology is available and the carbon stock ever reaches the activation threshold. The comparison between Regime 3 and Regime 1 therefore speaks to the geoengineering moratorium debate (Biermann et al., 2022). Leaving aside the question of whether a moratorium could actually be enforced, this welfare comparison tells us whether we would want it. Regime 1 is the world where the moratorium succeeds and Regime 3 is the world without one.

Before comparing welfare across the two regimes, we first quantify the risk that deployment creates. The *landing probability* $p(S)$ is the probability that the world never suffers a termination shock, computed from a starting point where the carbon stock is S and neither breakthrough has yet occurred. The world avoids a termination shock if a breakthrough stops the carbon stock from ever reaching the activation threshold, so that deployment never begins, or if deployment begins and the economy reaches effective safety before termination. Note that a CDR breakthrough during deployment only starts the removal era. The termination hazard remains live until the economy reaches effective safety, so a breakthrough raises the landing probability only in proportion to the chance that its removal era finishes safely. The landing probability solves a linear ODE system stated in Appendix A. This structure is why a simpler calculation, the ratio $\eta_C/(\eta_C + \bar{\lambda})$, which is the probability that a CDR breakthrough arrives before termination, overstates the odds of rescue.

We now turn to the welfare comparison between Regime 3 and Regime 1, which determines whether the world would want geoengineering to be available. Availability improves welfare from carbon stock S_0 only if $v^{R3}(S_0) \geq v^{R1}(S_0)$, where each value function embeds its

own equilibrium research intensities. We decompose the welfare difference between Regime 3 and Regime 1 into two terms:

$$\underbrace{v^{R3}(\eta^{R3}) - v^{R1}}_{\text{net effect}} = \underbrace{\left[v^{R3}(\eta^{R1}) - v^{R1}\right]}_{\text{value of buying time}} + \underbrace{\left[v^{R3}(\eta^{R3}) - v^{R3}(\eta^{R1})\right]}_{\text{deterrence effect}} \tag{27}$$

The value of buying time is the welfare gain from deployment when innovation incentives are undistorted relative to those in Regime 1 (the strongest case that can be made for availability). The deterrence effect is the welfare lost because deployment weakens innovation incentives. The region where availability lowers welfare increases with the termination risk and the intensity of the termination shock, and decreases with removal capacity and the masking share. The latter follows because a partial masking program carries the full termination risk while delivering only part of the cooling. Section 4.5 presents the welfare comparison as a function of the expected breakthrough time and the expected geoenginnering program lifetime.

Discussion

Advance market commitments. Without a purchase commitment, carbon removal capacity runs only when the carbon price covers the removal cost, and the price stops covering the cost while the carbon stock is still dangerous, because removing carbon lowers the very price that pays for removal (Section 2.5). An advance market commitment is a binding agreement, signed before a technology exists, to purchase from its developer at specified terms once it arrives (Kremer et al., 2020). Agreements of this form already exist for carbon removal, such as those organized by Frontier, a coalition of firms that commits to future removal purchases. The simple advance commitment in our model sets the purchase price at cost. This keeps capacity running when the carbon price stops covering removal, but it pays the innovator no profit, so it does not affect research incentives. A commitment with a purchase price above cost would give the innovator a profit and therefore raise removal research before the breakthrough arrives. For simplicity, we leave this above-cost version outside the model.

Peak shaving. Part of the tax reduction in Regimes 2 and 3 is efficient waiting for a breakthrough rather than pure moral hazard. Keith and MacMartin (2015) advocate this use of geoengineering: an early, moderate, responsive program that shaves the temperature peak while decarbonization and removal scale up. The value of buying time in Equation (27) measures the benefit of peak shaving, and the deterrence effect measures the accompanying cost of weakened innovation incentives. The masking share $m = 0.5$ in our robustness

analysis approximates a moderate program.[16]

The green paradox. Sinn (2008) argues that, with forward-looking owners of fossil reserves, the anticipation of a green breakthrough accelerates extraction before the breakthrough arrives. Our model omits this channel. If the green paradox channel operated, carbon and therefore ecological debt would accumulate faster than our calibration implies.

Ocean acidification. Acidification damages grow with the carbon stock and cannot be masked, so even in Regime 2, where deployment continues forever with no termination risk, the value of buying time is finite.

4 Quantitative Results

This section presents quantitative results from model. We solve the planner's HJB equations by value function iteration on a grid of 600 points over $S \in [280, 1200]$ ppm, and we solve the asset pricing ODEs on the same grid by finite differences. All three regimes are computed from a single set of parameter values, varying only the geoengineering and termination risk channels.

4.1 Calibration

Table 2 presents all parameter values. The climate, damage, and acidification parameters were discussed in Section 2. We set risk aversion $\gamma = 8$, within the range of the rare disaster literature (Barro, 2006; Wachter, 2013), the abatement cost parameter ψ, set so that emissions reach zero at a marginal abatement cost of \$320/tCO_2, consistent with cost estimates for hard-to-abate sectors (IPCC, 2022; Gillingham and Stock, 2018), and the expected program lifetime $1/\bar{\lambda} = 200$ years, with robustness over lifetimes of 100 to 300 years.

We set the scale of the termination damage function $\Gamma(F)$ by specifying the loss at one reference point: a termination that occurs when the carbon stock is 800 ppm, and ecological debt of $F = 1.48$°C of masked warming destroys 12.5% of productivity. This value is the midpoint of the 10–15% range built from the per-degree damage estimates of Burke et al. (2015), scaled up assuming termination delivers its warming an order of magnitude faster than climate change (Trisos et al., 2018). With $\beta = 2.5$ held fixed, the target implies $\Gamma_0 = 0.0498$, and the severity curve reaches 1.0% at 600 ppm and 5.4% at 700 ppm. We check plausibility at carbon stocks the model actually reaches. In this case, the 99.5th percentile of simulated carbon stocks over a 200-year horizon is 968 ppm, where the implied termination

[16] A full treatment of peak shaving would also begin before warming reaches T^* and would phase out deliberately, while our deployment rule has a fixed activation threshold and declines only if the carbon stock falls.

consumption drop is 33.5%. This is within the 30–40% range of historical macroeconomic disasters (Barro, 2006). Our robustness analysis also reports every result for the milder alternative with 3.8% destruction set at 800 ppm.

Table 2: Baseline Calibration

Parameter	Description	Value	Source
Climate			
T_{cs}	Transient climate slope	2.3 °C/doubling	$\Psi(427) = 1.4$°C
S_0	Current CO_2 stock	427 ppm	NOAA GML mean, start of 2026
$\bar{T}$	Tipping threshold	1.5 °C	Armstrong McKay et al. (2022)
ω_2	Damage coefficient	0.00439	R1 SCC = \$125 at S_0; 1.86×DICE
ζ_s	Tipping damage scale	0.005	IPCC tipping estimates
Ω_{acid}	Acidification term	\$20/t R2 tax at S_0	Acidification literature
σ_S	Carbon cycle volatility	1.0 ppm/yr$^{1/2}$	Global Carbon Budget flux
Geoengineering and Termination			
T^*	Thermostat target	2.0 °C (activation 512 ppm)	Wet-bulb recurrence; Paris
θ	Cooling efficiency	0.5 °C/unit	Robock et al. (2009)
m	Masking share	1	Section 2.2
$\bar{\lambda}$	Termination hazard	0.005/yr (lifetime 200 yrs)	Sensitivity 100–300 yrs
Γ_0	Termination damage scale	0.0498	12.5% destruction at 800 ppm
β	Severity convexity	2.5	Held fixed
Emissions and Abatement			
$\bar{e}$	Baseline net emissions	3.0 ppm/yr	Observed net accumulation rate
ψ	Abatement cost parameter	120	MAC ≈ \$320/t at full abatement
Breakthroughs and Innovation			
ψ_L	Post-collapse green cost	90% cut at \$75/t	Specification choice (Section 2.5)
$\bar{R}$	Removal capacity	15 $GtCO_2$/yr (1.9 ppm/yr)	NASEM (2019)
c_R	Removal marginal cost	\$100/tCO_2	Mature-DAC range (low end)
κ_G	Green research productivity	0.00260	expected 50 yrs (achieved 52.8)
κ_C	CDR research productivity	0.887	expected 50 yrs from activation
η_{max}	Cap on breakthrough rate	0.25/yr	Build-out time bound
S_{eff}	Effective safety stock	600 ppm	1% residual conditional loss
Preferences and Pricing			
γ	Risk aversion	8	Barro (2006); Wachter (2013)
ρ	Time preference	0.02	Standard
r	Risk-free rate	0.04	Standard
δ	CDR depreciation	0.05	IEA
ϕ	CDR removal efficiency	0.5 ppm/yr per unit	Normalization

We set the research productivities κ_G and κ_C so that the expected time to each breakthrough (measured from the activation threshold along the Regime 3 path) is 50 years. Setting these two symmetrically ensures that any asymmetry in outcomes comes from the difference in prizes. The expected time to breakthrough holds exactly for CDR. For the green breakthrough, it holds only approximately, because the expected breakthrough time depends on the green patent value, which is not known until the model is solved, and solving the model requires κ_G. When we solve the model, the expected time to the green breakthrough

is 52.8 years rather than 50.

Note that the research productivities and the removal capacity are the parameters with the least direct evidence behind them. Our calibration is designed so that a reader can substitute other values. Figure 7 shows welfare for various values of both parameters.

Figure 1 shows our calibrated geoengineering policy primitives. Our baseline deployment rule is the thermostat policy with full masking, i.e., $m = 1$. This means that anywhere deployment occurs, the realized temperature never rises above T^*. We choose $m = 1$ as the baseline, because a program that masked only part of the warming would leave realized temperature drifting above T^*, and Region V would increase deployment. In our robustness analysis, we explore partial deployment, $m \in \{0.25, 0.5\}$. We also hold the deployment rule $g(S)$ fixed across Regimes 2 and 3, so that every difference between the two regimes can be attributed to termination risk.

The thermostat rule g(S) is shown in panel (a) of Figure 1. It triggers at 512 ppm. The ecological debt $F(S) = \Psi(S) - T^*$ is shown in panel (b). At 600 ppm, it reaches 0.53 °C of masked warming, and at 800 ppm, it reaches 1.48 °C. The two termination risks are shown in panel (c). The physical termination hazard is constant at $\bar{\lambda} = 0.5\%$ per year while a program is deployed. The risk-neutral termination hazard $\lambda^{\mathbb{Q}}(S)$ increases with the carbon stock, since the jump in the pricing kernel at termination scales with the ecological debt.

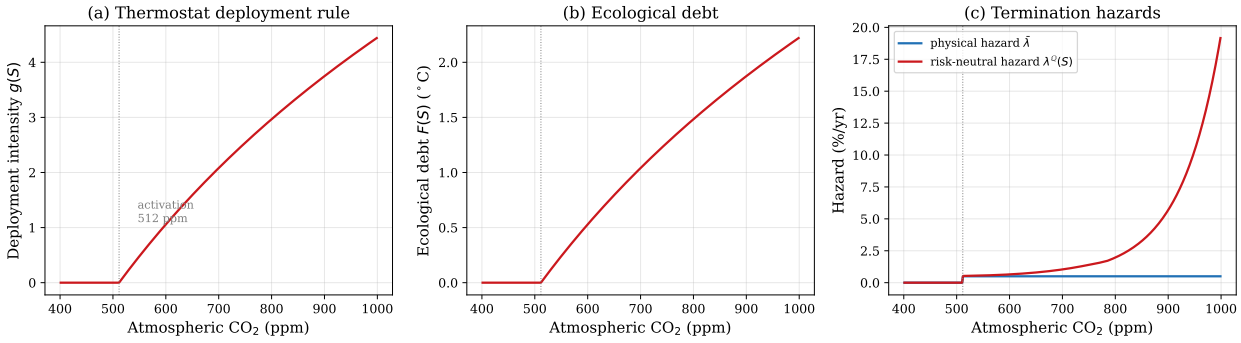


Figure 1: **Calibrated Model Primitives.** Panel (a) shows the thermostat deployment rule $g(S)$, which activates at 512 ppm. Panel (b) presents the ecological debt $F(S) = \Psi(S) - T^*$. Panel (c) reports the termination hazards.

4.2 The Carbon Price and Emissions

Figure 2 plots the Social Cost of Carbon under all three regimes, with research active. In Regime 1 the carbon price rises with the carbon stock, from \$99/tCO$_2$ at current concentrations to \$148/tCO$_2$ at 600 ppm and \$186/tCO$_2$ at 800 ppm. These values sit below the

no-innovation benchmark, whose price at current concentrations is the \$125/tCO$_2$ used to set the damage coefficient (Section 2.1), because the possibility of a breakthrough lowers the cost of carbon today.

In Regime 2, geoengineering removes the temperature component of the carbon price (Proposition 1b), and the price falls to \$38/tCO$_2$ at 600 ppm. This sits slightly below the acidification component of \$40/tCO$_2$, because, as in Regime 1, the possibility of a breakthrough lowers the cost of carbon today.

In Regime 3, termination risk partially restores the carbon price (Proposition 2a). As the carbon stock gets close to the activation threshold, an additional unit of carbon causes temperature damage up until the activation threshold, but less and less time remains until the activation threshold, so the price declines. At the activation threshold, the price drops discretely, because an additional unit of carbon stops raising realized temperature at all, which leaves the acidification and termination components. The termination component then grows with the ecological debt, and in our baseline calibration the Regime 3 price crosses above the Regime 1 price at $\hat{S} \approx 663$ ppm. This means that, starting at the crossing point, an additional unit of carbon is more socially costly under geoengineering than without it, since the extra carbon added to a possible termination outweighs the flow damages suppressed by geoengineering. The slight flattening of the curve around 600 ppm comes from the fact that this is the effective safety boundary at $S_{\text{eff}} = 600$ ppm, where the value of a CDR breakthrough switches slope because that is where a future removal era would stop. The flat segment at \$320/tCO$_2$ beginning just below 800 ppm reflects the reporting convention described in Section 2.4: the tax there has already driven emissions to zero, so we report the price at the shutdown level.

Figure 3 shows simulated paths of the carbon stock and ecological debt. We simulate 600 paths over 250 years at annual steps, starting from $S_0 = 427$ ppm in 2026, with breakthrough and termination switched off. Each path follows its regime's equilibrium net emissions plus the carbon-cycle noise from Equation (1), which is the only source of variation across simulated paths. The generated paths show how each regime's policies move the carbon stock before any event occurs. In the full model, most paths are interrupted, and 47.3% receive a green breakthrough before deployment ever begins.

Panel (a) shows that Regime 2 produces the highest carbon path. The suppressed price keeps emissions at roughly 2.6 ppm per year at 600 ppm, against 1.6 under Regime 1 and 1.9 under Regime 3, and at 800 ppm Regime 2 emissions are still 2.5 ppm per year while the Regime 3 price has driven emissions to zero. The lower carbon tax therefore does not just raise fossil fuel dividends. It sustains higher emissions, and the higher emissions accumulate into a permanently higher carbon path.

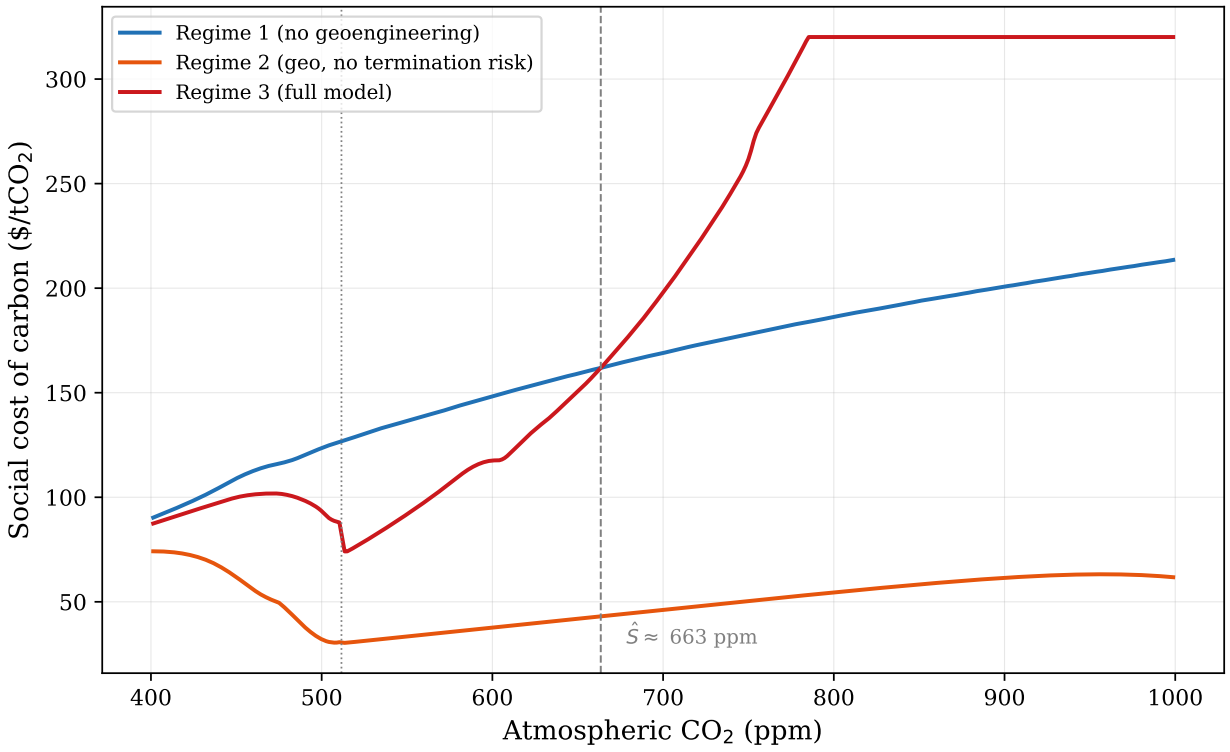


Figure 2: **Social Cost of Carbon.** The SCC under the three regimes. The Regime 3 price drops discretely at the activation threshold, where the temperature component reaches zero, and is thereafter composed of the acidification and termination components. In our baseline calibration, the Regime 3 price crosses above the Regime 1 price at $\hat{S} \approx 663$ ppm.

Panel (b) tracks the ecological debt under Regime 3. On the median of the 600 paths, deployment activates in 2067, and the carbon stock reaches the toxic put threshold S^* in 2128, roughly a century from today.[17]

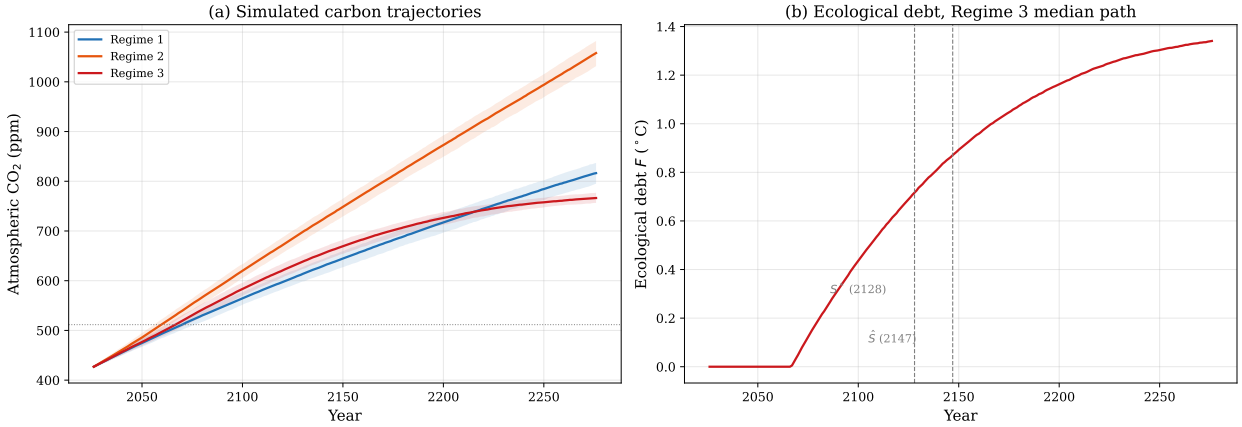


Figure 3: **Simulated Dynamics.** (a) Simulated paths of the carbon stock (median and 90% bands) under the three regimes, with breakthroughs and termination switched off. (b) The ecological debt under Regime 3, with the years in which the median path reaches the activation threshold and the toxic put threshold S^* marked.

4.3 The Toxic Put

Figure 4 compares Brown firm valuations in Regime 3 with those in Regime 1. In Regime 1, Brown firm valuations fall as the carbon stock grows, because the rising carbon tax reduces Brown firm dividends. In Regime 3, Brown firm valuations rise as the carbon stock approaches the activation threshold, because the era of suppressed carbon taxes is imminent and the termination risk premium is still small. It peaks around the threshold.

Above the activation threshold, the two forces in Proposition 3 compete. At low carbon stocks, the Regime 3 curve lies above the Regime 1 curve, because the dividend gain from the suppressed tax outweighs the small termination risk premium. In our baseline calibration the two curves intersect at $S^* \approx 634$ ppm, where the ecological debt of 0.71 °C causes a termination shock that destroys more value than the suppressed tax creates. Beyond S^*, Brown firms have lower value than in Regime 1 at the same carbon stock. The put has become toxic.

[17]The simulated paths assume the social planner implements a globally coordinated carbon tax, which slows carbon accumulation along the way. If actual policy is weaker, as is currently the case, carbon accumulates faster and deployment activates earlier.

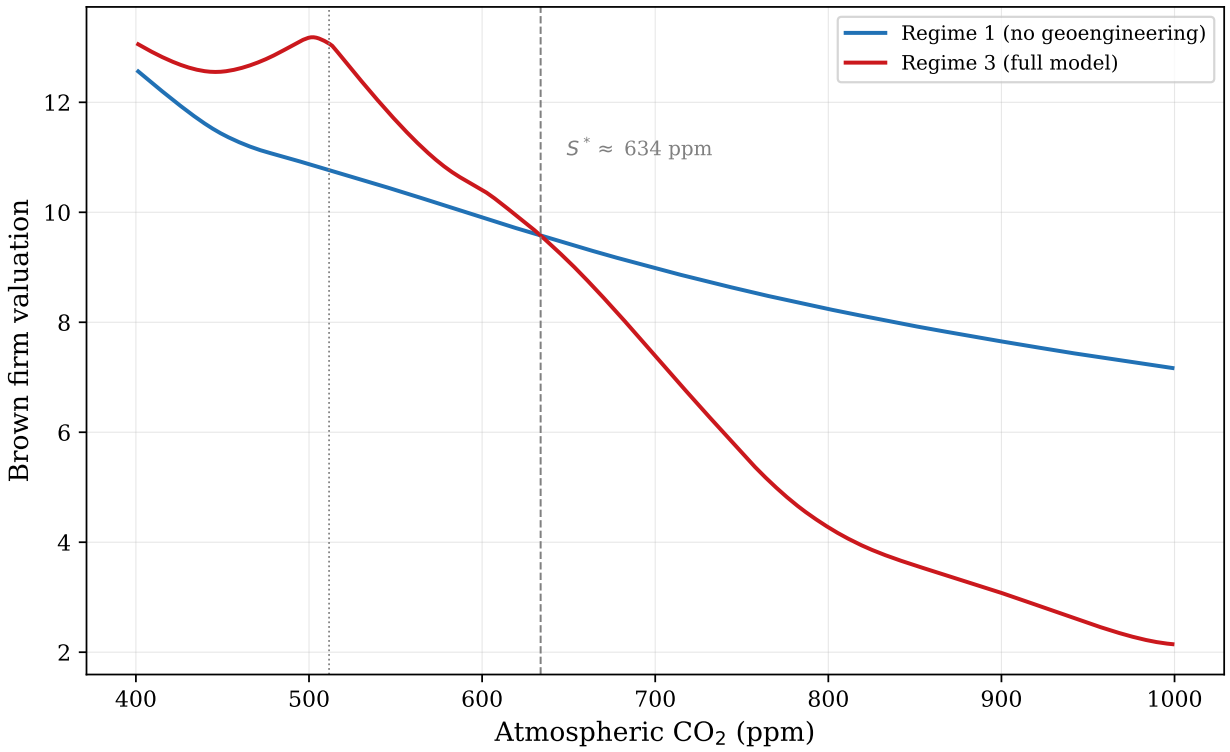


Figure 4: **Brown Firm Valuations and the Toxic Put.** Brown valuations under Regime 1 and Regime 3. In our baseline calibration, the Regime 3 curve crosses below the Regime 1 curve at $S^* \approx 634$ ppm. Below the crossing, geoengineering raises Brown value. Above it, Brown firms are worth less than in a world with no geoengineering at the same carbon stock.

4.4 Innovation Across Regimes

Figure 5 displays the equilibrium research intensities $\eta_G(S)$ and $\eta_C(S)$ under the three regimes. The figure offers the quantitative counterpart of Proposition 4.

Panel (a) shows green research. The Regime 1 research intensity is flat at 2.00% per year because, after a green breakthrough, green energy displaces brown energy entirely, so green profits equal their cap $D^G(\tau_{cap})$ and the green patent value barely varies with the carbon stock. The Regime 2 research intensity lies strictly below the Regime 1 level everywhere due to the crowding out channel described in Proposition 4a. The Regime 3 research intensity sits between the former two intensities. The termination risk increases the green patent value relative to Regime 2 because the post-termination tax raises green profits. However, the gain is capped, so green research intensity never rises above its Regime 1 level. At high carbon stocks beyond 936 ppm, even the gain relative to Regime 2 reverses, because the productivity destruction at termination cuts the post-termination green sector profits.

Panel (b) shows carbon removal research intensity under the three regimes. In Regime 2, removal research is essentially zero above the activation threshold. In this region, geoengineering has eliminated the temperature component of the price of carbon, and the remaining acidification component of \$20/tCO_2 is far below the removal cost of \$100/tCO_2. In Regime 1, removal research grows steadily with the carbon stock as the carbon price climbs past the removal cost c_R. In Regime 3, the carbon price while geoengineering operates stays below the removal cost, and the advance commitment contract pays exactly the removal cost, leaving the carbon removal sector with no profit margin. The removal patent's value is derived purely from termination risk.

Regime 1 offers a steady profit to the carbon patent innovator whereas Regime 3 offers a rare prize tied to termination risk. The prize is modest while the ecological debt is modest, and the steady Regime 1 profit margin is worth more. As carbon accumulates, the prize grows, and removal research under geoengineering overtakes its Regime 1 level at $S_\eta \approx 663$ ppm. This crossing lies above the toxic put threshold $S^* \approx 634$ ppm, where Brown value under geoengineering falls below its Regime 1 level. Because $S^* < S_\eta$, throughout the entire phase in which geoengineering subsidizes fossil fuel firms, it also depresses removal research below its Regime 1 level. Below the activation threshold the ordering flips. Regime 3 removal research intensity lies slightly above Regime 1 (0.57 against 0.44% per year at current concentrations), because the prospect of future deployment gives the removal patent value from the risk of termination.[18]

[18]The Regime 3 intensity reaches the cap η_{max} at 753 ppm and is flat beyond. The local peak at 586 ppm and trough at 603 ppm have the same source as the flattening of the carbon price in Figure 2: the value of a future removal era changes slope at the effective safety boundary, where removal would stop.

Comparing Panels (a) and (b) reveals an important insight about how geoengineering and termination risk directs research toward carbon removal. Moving from Regime 2 to Regime 3 multiplies removal research by orders of magnitude but increases green research by less than a percentage point. Termination risk revives research in the only technology (carbon removal) that can retire the ecological debt. The revival is nonetheless costly at the start. Measured from current concentrations, the expected time to the CDR breakthrough is 64.5 years without geoengineering (Regime 1) and 70.8 years with it (Regime 3), so deployment delays the expected CDR breakthrough by roughly 6.2 years.

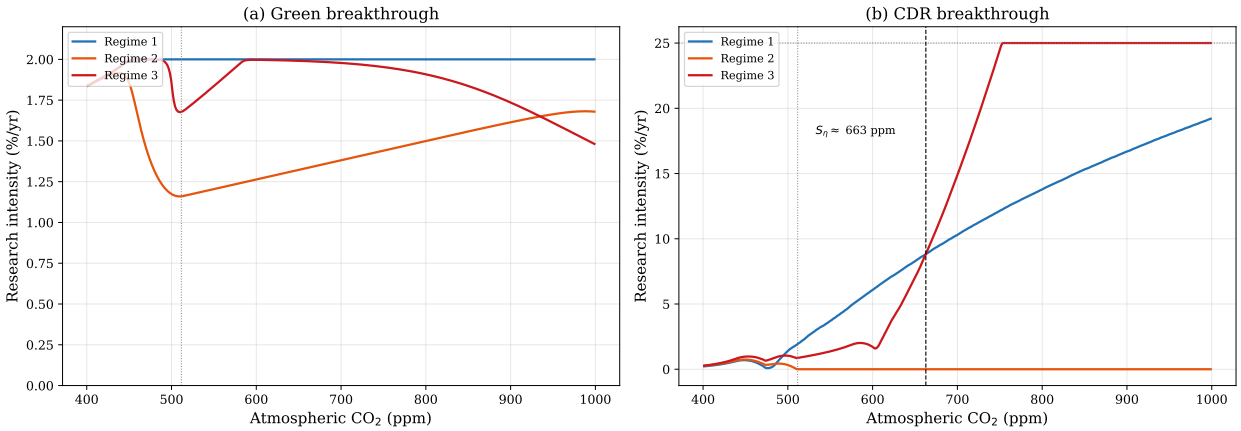


Figure 5: **Innovation Across Regimes.** Equilibrium research intensities by regime. (a) The green breakthrough. (b) The CDR breakthrough, with the crossing $S_\eta \approx 663$ ppm marked with a vertical dotted line.

4.5 Rescue and Welfare

Figure 6 reports the landing probability, the probability that the world never suffers a termination shock. Starting from the current carbon stock, we compute the landing probability as 70.8%. Two-thirds of the overall landing probability, or 47.3 percentage points, comes from paths on which a green breakthrough arrives before the carbon stock reaches the activation threshold, so geoengineering deployment never begins.[19]

Conditional on geoengineering deployment, the landing probability is much lower. Evaluated just above the effective safety boundary, the conditional landing probability is 30.3%,

[19]In our baseline calibration, a CDR breakthrough before activation does not prevent deployment, because the removal cost exceeds the carbon price, so no one would choose to operate carbon removal. Once deployment begins and the growing ecological debt raises the carbon price, carbon removal begins operating under the advance commitment contract.

of which 16.1 percentage points come through the breakthroughs and 14.2 percentage points through carbon-cycle noise drifting the carbon stock back below the activation threshold[20]

Our model delivers a surprising and underappreciated implication of green innovation. A green energy breakthrough is widely viewed as the best possible climate news. However, once geoengineering is deployed, a green breakthrough lowers the probability of landing. It stops emissions, but it also caps the carbon price below the removal cost, so the removal patent earns nothing, removal research stops, and the ecological debt stays in place with no funded exit. Termination then arrives before rescue with probability close to one. The energy breakthrough still raises welfare at most carbon stocks, because cheap energy is valuable and the ecological debt stops deepening. In our baseline calibration, however, a green breakthrough that arrives at a sufficiently high carbon stock lowers welfare outright, because losing the funded chance of a CDR breakthrough outweighs those gains.

The ratio $\eta_C/(\eta_C+\bar{\lambda})$ represents the probability that a CDR breakthrough arrives before termination. This probability is 76.9% evaluated just above the effective safety boundary, where the conditional landing probability is only 30.3%. This ratio dramatically overstates the conditional landing probability because a CDR breakthrough does not immediately bring about landing. The removal era that follows (the period in which removal operates under the advance commitment) lasts decades under a live termination hazard, and below 632 ppm residual emissions exceed removal capacity, so reaching effective safety also depends on a green breakthrough arriving during the removal era. Panel (b) plots the two components of the conditional landing probability, the exit through breakthroughs and the exit through carbon-cycle noise, together with the probabilities that a removal era ends in effective safety, q_C for the economy with removal only and q_B for the economy with both breakthroughs.

We now turn to welfare. Figure 7 displays the welfare difference between Regime 3 and Regime 1 at current concentrations over the expected breakthrough time and the expected program lifetime. This welfare difference is the sum of the value of buying time and the deterrence effect in Equation (27).

Blue regions are combinations of beliefs under which welfare is higher in Regime 3 than in Regime 1 at current concentrations, $v^{R3}(S_0) > v^{R1}(S_0)$, so the world is better off with geoengineering. The red region shows the reverse. It occupies short program lifetimes combined with slow breakthroughs, deployment that is likely to be interrupted while waiting for a tech-

[20]The landing probability line represents the world's full future path, so a pre-breakthrough economy that reaches the effective safety boundary is not yet safe, because emissions push the carbon stock back up. The conditional-on-deployment line takes on a different value at the activation threshold compared to the landing probability line, because it starts from a deployed state and reports whether the carbon stock reaches effective safety before termination, so carbon-cycle noise that drifts the stock across the nearby boundary counts as success. The carbon noise channel fades quickly away from the boundary, and the two lines coincide above roughly 650 ppm.

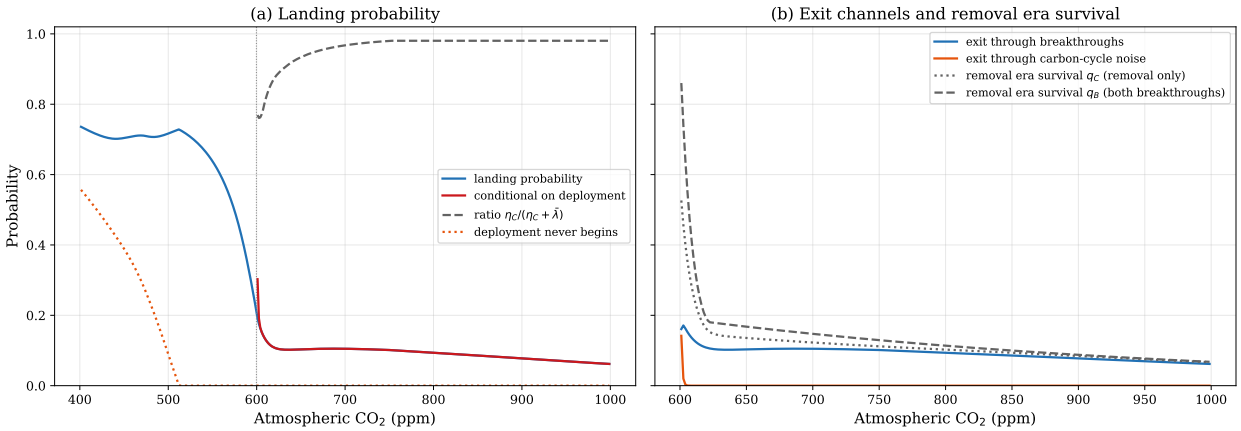


Figure 6: **The Landing Probability.** (a) The landing probability, its conditional-on-deployment version, the ratio $\eta_C/(\eta_C+\bar{\lambda})$, and the contribution of paths on which deployment never begins. (b) The probabilities that each removal era ends in effective safety.

nology that is unlikely to come in time. The comparison is not monotone in breakthrough speed. At short lifetimes, faster breakthroughs improve the case for geoengineering, rescuing a program that will not survive long. At long lifetimes the gradient reverses, because when the program is durable, slower breakthroughs make the temporary cooling worth more in the long interim, and the largest welfare gain on the grid sits at the longest lifetime and the slowest breakthrough.

At our baseline beliefs, welfare is higher in Regime 3 than in Regime 1. However, the different is small at +0.037% of consumption, with a value of buying time of +0.047% and a deterrence effect near zero (−0.009%).

If we shut off innovation (the no-innovation benchmark), welfare is higher in Regime 3 than in Regime 1 for every belief combination shown. I.e., without the possibility of green or CDR breakthroughs, the world is better off with geoengineering technology available under every belief combination shown. Allowing for the possibility of breakthroughs creates the red region in which the world is better off without geoengineering technology being available. The intuition is as follows. Without innovation, the only cost of geoengineering is termination risk itself, and the value of masking temperature damages for as long as the program survives outweighs the cost of termination risk. With innovation, deployment also delays the breakthroughs the world is waiting for, and that additional cost is what turns the welfare comparison negative for some beliefs.

Figure 8 turns to a timing question: What is the availability of geoengineering technology worth to a world that has reached a given carbon stock without geoengineering technology? Making the technology available at carbon stock S means the world switches to Regime 3

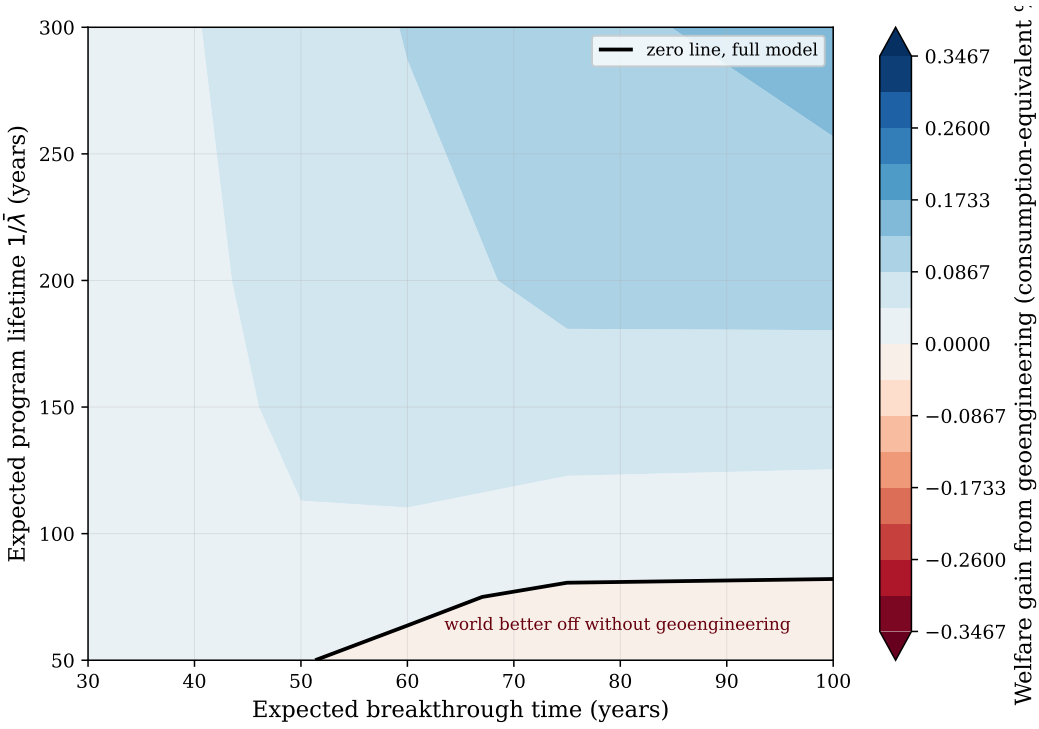


Figure 7: **Welfare With and Without Geoengineering, by Beliefs.** The welfare gain from geoengineering at current concentrations, in consumption equivalents, over the expected breakthrough time and the expected program lifetime. The black contour marks the belief combinations at which welfare is exactly equal in Regime 3 and Regime 1. In the no-innovation benchmark, the gain is positive everywhere on this domain.

at that stock. We assume the deployment rule is the thermostat policy. If the carbon stock exceeds the activation threshold, Region V deploys immediately, realized temperature falls to T^*, and the termination hazard switches on with ecological debt of $\Psi(S) - T^*$.

The figure plots the welfare gain from this switch, decomposed into the value of buying time and the deterrence effect of Equation (27).[21] The value of buying time increases with the carbon stock above the activation threshold, peaks at around 675 ppm, and then drops quickly because making the technology available at a deep carbon stock means an immediate large ecological debt, and a termination from that debt destroys more than the cooling is worth. The deterrence effect becomes positive above 652 ppm, where deployment causes removal research to exceed its Regime 1 level. So its effect on innovation is no longer a cost. The net effect becomes negative above 818 ppm. A world that reaches so high a carbon stock without geoengineering is better off never gaining the geoengineering technology.

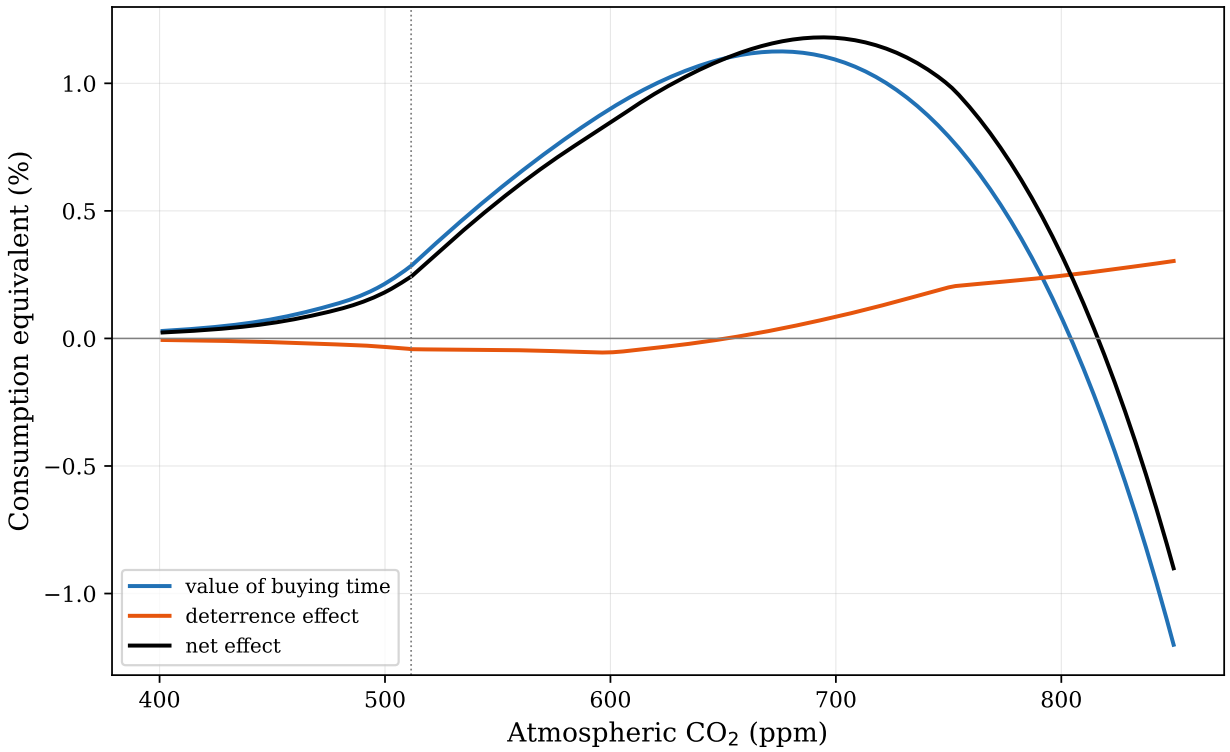


Figure 8: **The Welfare Decomposition.** The welfare gain from making geoengineering available at each carbon stock, decomposed into the value of buying time and the deterrence effect, in consumption equivalents.

[21] Equation (27) is exact in units of the value function. Each term is plotted here as a consumption equivalent, which is multiplicative, $1 + CE_{total} = (1 + CE_{buy})(1 + CE_{det})$. This implies the two plotted terms do not add up exactly to the plotted net effect (the difference is less than a thousandth of a percentage point at current concentrations).

4.6 The Advance Commitment

Figure 9 shows why the advance commitment contract is needed. Panel (a) plots the carbon price in the economy where the CDR breakthrough has occurred but the green breakthrough has not yet occurred, computed without the advance commitment, as a function of the carbon stock. The removal era moves the world leftward along this curve, as removal reduces the carbon stock and the corresponding carbon price.

The leftward motion stops at the dash-dotted line, 632 ppm, where residual emissions have risen to equal the removal capacity. At 632 ppm, net removal is zero despite the carbon price covering the removal cost. The curve falls below 632 ppm because the price remains well defined at these carbon stocks. But they are reached only if the residual emissions fall to zero due to a green breakthrough and the leftward motion resumes. The dotted line at 609 ppm marks where the price drops below the removal cost, the point at which a market on its own would stop paying for removal. In our baseline calibration this cutoff is never reached, because the carbon stock stalls at 632 ppm first. The advance commitment matters instead in the economies after a green breakthrough. There the carbon price is capped below the removal cost permanently, so without committed purchases removal would never run at all (Section 2.5).

Panel (b) reports the cost of the commitment. Removal operates only above the effective safety boundary, so procurement outlays are zero below 600 ppm and jump to roughly 1.5% of output where removal runs, peaking at 1.58%.[22]

4.7 Risk Premia Across Firms

Figure 10 decomposes the expected excess returns of Brown and Green firms into the three jump components of Equation (26): termination, the green breakthrough, and the CDR breakthrough. All values are computed in the Regime 3 economy in which none of these events has yet occurred, at each carbon stock, in basis points per year. The sign of each risk premium component depends on how the asset's value moves with consumption at each event. An asset that gains value at an event that lowers consumption is a hedge and earns a negative premium. An asset that gains value at an event that raises consumption earns a positive premium.

With respect to carbon removal, the model contains three objects that should not be confused. The value of a unit of removal capacity, q, prices installed capacity and is the quantity compared across regimes in Section 3. The patent value $\hat{\pi}_C$ prices the developer

[22]Removing at full capacity means purchasing 15 $GtCO_2$ per year at \$100 per ton, about \$1.5 trillion per year.

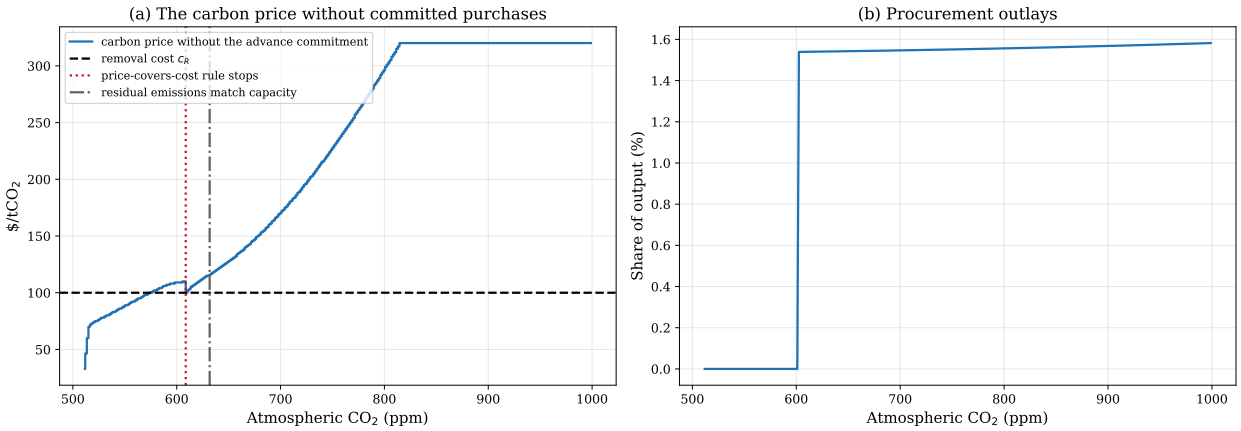


Figure 9: **The Carbon Price During the Removal Era.** (a) The carbon price along the removal era, computed without the advance commitment, against the removal cost c_R. The dotted line marks where a price-covers-cost rule would stop removal (609 ppm). The dash-dotted line marks the carbon stock at which residual emissions match removal capacity and the carbon stock stops falling (632 ppm). (b) Procurement outlays during the removal era as a share of output.

of the technology and determines the rate of the CDR breakthrough. Before the CDR breakthrough, investors cannot hold a firm that operates capacity and earns profits from removal. They can only hold the potential patent developer, whose sole asset is the prospect of the patent.

For Brown firms, the termination component dominates and rises steeply with the carbon stock. This is the toxic put expressed in expected returns. The green-breakthrough component turns positive at high carbon stocks. Brown firms gain value when the green breakthrough arrives there, because the planner sets a lower carbon tax while emissions collapse anyway, but the breakthrough also raises consumption on impact, so the gain commands a positive premium rather than serving as a hedge. The CDR-breakthrough component is a modest hedge throughout.[23]

At moderate carbon stocks, Green firms pair a negative termination component, the disaster insurance of Proposition 3b, with the chance of a green breakthrough. The insurance sign reverses near 778 ppm, visible in panel (b), where the productivity destruction and the shrinking gap between the pre-termination and post-termination taxes turn the termination component positive.

[23]The sign change at very high carbon stocks described in Section 3.3 never occurs on the reporting grid, because the possibility of breakthroughs supports Brown values at high carbon stocks.

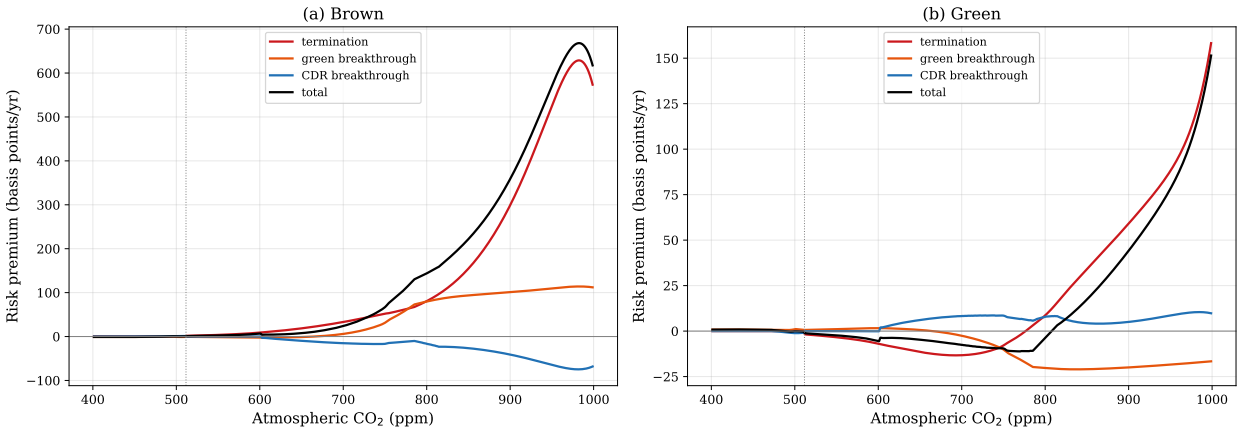


Figure 10: **Risk Premia by Jump Component.** Expected excess returns in the Regime 3 economy before any of the three events has occurred, decomposed into the termination, green-breakthrough, and CDR-breakthrough components of Equation (26), in basis points per year. (a) Brown firms. (b) Green firms.

4.8 Robustness

Table 3 explores the robustness of the baseline calibration results. In each row, we resolve the full model with a single parameter changed, and report: the toxic put threshold S^*, the tax crossing $\hat{S}$, the value of buying time and the deterrence effect at current concentrations, and the landing probability.

The expected program lifetime (which is decreasing in the termination hazard) moves the results most. At the baseline expected lifetime of 200 years, the subsidy phase, the range of carbon stocks in which geoengineering raises Brown value, runs from the activation threshold at 512 ppm up to S^*, a span of 122 ppm. At an expected lifetime of 150 years, S^* falls to 595 ppm and the subsidy phase narrows accordingly. This shows that the more likely termination is, the earlier the market prices the toxic put.

The severity of the termination shock determines the location of the crossings. Assuming a milder severity alternative, 3.8% destruction at 800 ppm, the toxic put threshold moves to 801 ppm and the tax crossing to 847 ppm.

Our results are strongly sensitive to the masking share. Because the thermostat policy adjusts the deployment to keep the actual temperature at T^*, the cooling efficiency θ does not affect the outcome. We can vary the masking share m. Both crossings move deep into the carbon stock range for $m = 0.25$ (see 844 and 918 ppm). For $m = 0.25$, the situation is very different. There is no crossing at all on the reporting grid, the sign of the value of buying time flips, and the deterrence effect becomes positive. A policy that masks a quarter of the

Table 3: Robustness: One-at-a-Time Deviations from the Baseline Calibration

	S^* (ppm)	$\hat{S}$ (ppm)	Buying time (CE %)	Deterrence (CE %)	Landing prob. (%)
Baseline (lifetime 200, $m = 1$)	634	664	0.050	-0.002	69.8
Program lifetime 100 years	545	600	0.021	0.002	75.3
Program lifetime 150 years	595	635	0.038	-0.001	72.7
Program lifetime 300 years	685	706	0.064	-0.005	65.1
Geopolitical amplification ($\lambda_0 + \lambda_1 g$)	589	625	0.051	0.003	73.5
Milder termination shock (3.8% at 800 ppm)	801	847	0.062	-0.012	71.2
Masking share $m = 0.5$	844	918	0.014	-0.002	73.9
Masking share $m = 0.25$	–	–	-0.002	0.012	60.9
Expected breakthrough time 30 years	641	667	0.020	-0.002	68.4
Expected breakthrough time 75 years	621	656	0.087	-0.006	65.0
Expected breakthrough time 100 years	609	648	0.120	-0.022	57.8
Removal capacity $\bar{R} = 5$ Gt/yr	607	639	0.044	-0.004	71.1
Removal capacity $\bar{R} = 10$ Gt/yr (consensus scale)	620	651	0.046	0.000	71.4
Removal capacity $\bar{R} = 20$ Gt/yr	638	677	0.048	0.001	66.2
Removal cost $c_R = \$150$/t	610	634	0.060	0.006	80.6
Cap on breakthrough rate $\eta_{max} = 0.5$/yr	634	664	0.050	-0.003	69.8
Cap on breakthrough rate $\eta_{max} = 1.0$/yr	634	664	0.050	-0.003	69.8
Acidification floor \$10/t	629	661	0.070	-0.009	69.3
Acidification floor \$30/t	636	665	0.033	0.004	68.6
Green collapse: 95% reduction at \$75/t	632	663	0.048	-0.001	69.9

Notes: Each row solves the full model with one change from the baseline calibration. S^* is the toxic put threshold. $\hat{S}$ is the tax crossing. Buying time and deterrence are the two welfare terms of Equation (27), evaluated at current concentrations in consumption-equivalent percent. Landing is the landing probability starting from current concentrations. The robustness rows are solved on a 320-point grid, whereas the numbers in the main text are based off of a calibration using a 600-point grid. This is why the baseline row is slightly different from the text. Dashes mark a crossing that does not exist on the reporting grid under that deviation.

warming carries the full risk of termination but delivers only a quarter of the cooling. The geopolitical amplification variant where the termination hazard increases with the intensity of deployment moves all crossings closer to the activation threshold.

5 Conclusion

By providing a cheap and fast-acting way to cool the planet without reducing the carbon stock, geoengineering introduces risks that are sharply different from the transition and biodiversity risks studied in existing climate-finance literature. In this paper, we develop a macro-finance framework in which which geoengineering acts as a put option on climate damages, funded by a growing "ecological debt" of masked warming. If the program ever stops, temperatures snap back in a catastrophic termination shock. Our model allows us to evaluate the argument that geoengineering improves welfare by buying time while the world waits for a climate technology breakthrough.

Our analysis yields four main results. First, geoengineering without termination risk is pure moral hazard. It removes the temperature component of the carbon price, raises fossil fuel valuations, and reduces the value of carbon removal.

Second, termination risk changes the cross-section of asset prices. The prospect of geoengineering initially subsidizes fossil fuel incumbents by suppressing the carbon tax, but as ecological debt builds, those same assets become toxic. Investors who view geoengineering as a reprieve from carbon taxation are missing the termination exposure that comes with it.

Third, geoengineering and carbon removal switch from substitutes to complements, because the ecological debt that geoengineering accumulates can be retired safely only by removing carbon. Policymakers should therefore evaluate geoengineering jointly with carbon removal, rather than treating one as a substitute for the other.

Fourth, geoengineering and the associated termination risk direct innovation. Deployment lowers realized temperatures while the world waits for a breakthrough. But it weakens the incentive to produce one, because innovators are paid through carbon prices that deployment suppresses. We consider two breakthroughs the world might wait for. A green breakthrough stops emissions but leaves the carbon stock in place. Carbon removal is the only technology that can retire ecological debt, and the termination risk created by deployment pushes research strongly in that direction. Our analysis also shows that carbon removal has a built-in incentive problem: it is paid through the price of carbon, and that price falls as removal succeeds. An advance commitment contract solves this problem by promising to buy removal at a set price in the future.

On net, the world with geoengineering has slightly higher welfare in our baseline calibra-

tion, but the margin is small, and it can flip depending on how fast breakthroughs arrive and how long geoengineering programs survive before a termination shock.

More broadly, our model offers a framework for thinking through the consequences of geoengineering before they arrive. Our calibration results depend on parameters that remain highly uncertain, including the termination hazard, the severity of the termination shock, and the cost and productivity of breakthrough technologies. That uncertainty points directly to the need for further research. Research done today will help determine whether geoengineering should be deployed, and if so, how. As the physical conditions for deployment strengthen, the asset pricing and innovation predictions of the model should become testable in market data. We leave that empirical work for future research.

References

Viral V. Acharya, Stefano Giglio, Stefano Pastore, Johannes Stroebel, Zhenhao Tan, and Tiffany Yong. Climate transition risks and the energy sector. Working Paper 33413, National Bureau of Economic Research, 2025.

David I. Armstrong McKay, Arie Staal, Jesse F. Abrams, Ricarda Winkelmann, Boris Sakschewski, Sina Loriani, Ingo Fetzer, Sarah E. Cornell, Johan Rockström, and Timothy M. Lenton. Exceeding 1.5°C global warming could trigger multiple climate tipping points. *Science*, 377(6611):eabn7950, 2022. doi: 10.1126/science.abn7950.

Michael Barnett, William Brock, and Lars Peter Hansen. Pricing uncertainty induced by climate change. *Review of Financial Studies*, 33(3):1024–1066, 2020.

Lint Barrage. Optimal dynamic carbon taxes in a climate–economy model with distortionary fiscal policy. *Review of Economic Studies*, 87(1):1–39, 2020.

Scott Barrett. The incredible economics of geoengineering. *Environmental and Resource Economics*, 39(1):45–54, 2008.

Robert J. Barro. Rare disasters and asset markets in the twentieth century. *Quarterly Journal of Economics*, 121(3):823–866, 2006.

Mariia Belaia, Juan B. Moreno-Cruz, and David W. Keith. Optimal climate policy in 3D: Mitigation, carbon removal, and solar geoengineering. *Climate Change Economics*, 12(3): 2150008, 2021. doi: 10.1142/S2010007821500081.

Frank Biermann, Jeroen Oomen, Aarti Gupta, et al. Solar geoengineering: The case for an international non-use agreement. *WIREs Climate Change*, 13(3):e754, 2022.

Adrien Bilal and Diego R Känzig. The macroeconomic impact of climate change: Global versus local temperature*. *The Quarterly Journal of Economics*, 141(2):889–944, 05 2026. ISSN 0033-5533. doi: 10.1093/qje/qjag011. URL `https://doi.org/10.1093/qje/qjag011`.

Patrick Bolton and Marcin Kacperczyk. Do investors care about carbon risk? *Journal of Financial Economics*, 142(2):517–549, 2021.

Patrick Bolton and Marcin Kacperczyk. Global pricing of carbon-transition risk. *Journal of Finance*, 78(6):3677–3754, 2023. doi: 10.1111/jofi.13272.

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

Yongyang Cai and Thomas S. Lontzek. The social cost of carbon with economic and climate risks. *Journal of Political Economy*, 127(6):2684–2734, 2019.

Partha Dasgupta. The economics of biodiversity: The Dasgupta review. Technical report, HM Treasury, London, 2021.

Robert F. Engle, Stefano Giglio, Bryan Kelly, Heebum Lee, and Johannes Stroebel. Hedging climate change news. *Review of Financial Studies*, 33(3):1184–1216, 2020.

Pierre Friedlingstein, Michael O'Sullivan, Matthew W. Jones, et al. Global carbon budget 2023. *Earth System Science Data*, 15(12):5301–5369, 2023.

Xavier Gabaix. Variable rare disasters: An exactly solved framework for ten puzzles in macro-finance. *Quarterly Journal of Economics*, 127(2):645–700, 2012.

Alexandre Garel, Arthur Romec, Zacharias Sautner, and Alexander F. Wagner. Do investors care about biodiversity? *Review of Finance*, 28(4):1151–1186, 2024.

Stefano Giglio, Bryan Kelly, and Johannes Stroebel. Climate finance. *Annual Review of Financial Economics*, 13:15–36, 2021.

Stefano Giglio, Theresa Kuchler, Johannes Stroebel, and Olivier Wang. The economics of biodiversity loss. Working Paper 32678, National Bureau of Economic Research, 2024.

Kenneth Gillingham and James H. Stock. The cost of reducing greenhouse gas emissions. *Journal of Economic Perspectives*, 32(4):53–72, 2018.

Mikhail Golosov, John Hassler, Per Krusell, and Aleh Tsyvinski. Optimal taxes on fossil fuel in general equilibrium. *Econometrica*, 82(1):41–88, 2014.

Samuel Hartzmark and Kelly Shue. Geoengineering: Free rider or free driver? Working paper, 2026.

Garth Heutel, Juan Moreno-Cruz, and Soheil Shayegh. Solar geoengineering, uncertainty, and the price of carbon. *Journal of Environmental Economics and Management*, 87:24–41, 2018.

Eun-Soon Im, Jeremy S. Pal, and Elfatih A. B. Eltahir. Deadly heat waves projected in the densely populated agricultural regions of south asia. *Science Advances*, 3(8):e1603322, 2017. doi: 10.1126/sciadv.1603322.

IPCC. Climate change 2022: Mitigation of climate change. In *Contribution of Working Group III to the Sixth Assessment Report of the IPCC*. Cambridge University Press, 2022.

Fortunat Joos, Raphael Roth, Jan S. Fuglestvedt, et al. 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, 2013.

David W. Keith and Douglas G. MacMartin. A temporary, moderate and responsive scenario for solar geoengineering. *Nature Climate Change*, 5:201–206, 2015.

Michael Kremer, Jonathan Levin, and Christopher M. Snyder. Advance market commitments: Insights from theory and experience. *AEA Papers and Proceedings*, 110:269–273, 2020.

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

Douglas G. MacMartin, Ken Caldeira, and David W. Keith. Solar geoengineering to limit the rate of temperature change. *Philosophical Transactions of the Royal Society A*, 372: 20140134, 2014.

Justin McClellan, David W. Keith, and Jay Apt. Cost analysis of stratospheric albedo modification delivery systems. *Environmental Research Letters*, 7(3):034019, 2012. doi: 10.1088/1748-9326/7/3/034019.

Juan B. Moreno-Cruz and David W. Keith. Climate policy under uncertainty: A case for solar geoengineering. *Climatic Change*, 121(3):431–444, 2013.

National Academies of Sciences, Engineering, and Medicine. Negative emissions technologies and reliable sequestration: A research agenda. Technical report, The National Academies Press, Washington, DC, 2019.

William D. Nordhaus. Revisiting the social cost of carbon. *Proceedings of the National Academy of Sciences*, 114(7):1518–1523, 2017.

Jeremy S. Pal and Elfatih A. B. Eltahir. Future temperature in southwest asia projected to exceed a threshold for human adaptability. *Nature Climate Change*, 6:197–200, 2016. doi: 10.1038/nclimate2833.

Andy Parker and Peter J. Irvine. The risk of termination shock from solar geoengineering. *Earth's Future*, 6(3):456–467, 2018.

Jonathan Proctor, Solomon Hsiang, Jennifer Burney, Marshall Burke, and Wolfram Schlenker. Estimating global agricultural effects of geoengineering using volcanic eruptions. *Nature*, 560(7719):480–483, 2018.

Colin Raymond, Tom Matthews, and Radley M. Horton. The emergence of heat and humidity too severe for human tolerance. *Science Advances*, 6(19):eaaw1838, 2020. doi: 10.1126/sciadv.aaw1838.

Kevin Rennert, Frank Errickson, Brian C. Prest, et al. Comprehensive evidence implies a higher social cost of co_2. *Nature*, 610:687–692, 2022.

Alan Robock, Allison Marquardt, Ben Kravitz, and Georgiy Stenchikov. Benefits, risks, and costs of stratospheric geoengineering. *Geophysical Research Letters*, 36(19):L19703, 2009.

Steven C. Sherwood and Matthew Huber. An adaptability limit to climate change due to heat stress. *Proceedings of the National Academy of Sciences*, 107(21):9552–9555, 2010. doi: 10.1073/pnas.0913352107.

Hans-Werner Sinn. Public policies against global warming: a supply side approach. *International Tax and Public Finance*, 15(4):360–394, 2008.

Wake Smith and Gernot Wagner. Stratospheric aerosol injection tactics and costs in the first 15 years of deployment. *Environmental Research Letters*, 13(12):124001, 2018.

Christopher H Trisos, Giuseppe Amatulli, Jessica Gurevitch, Alan Robock, Lili Xia, and Brian Zambri. Potentially dangerous consequences for biodiversity of solar geoengineering implementation and termination. *Nature Ecology & Evolution*, 2(3):475–482, 2018. doi: 10.1038/s41559-017-0431-0.

Frederick van der Ploeg and Armon Rezai. The risk of policy tipping and stranded carbon assets. *Journal of Environmental Economics and Management*, 100:102258, 2020.

Jessica A. Wachter. Can time-varying risk of rare disasters explain aggregate stock market volatility? *Journal of Finance*, 68(3):987–1035, 2013.

Martin L. Weitzman. A voting architecture for the governance of free-driver externalities, with application to geoengineering. *Scandinavian Journal of Economics*, 117(4):1049–1068, 2015.

A Derivations and Supporting Results

In this appendix, we use numerical verification based on the slope of the value function in the production grid (a level ordering of value functions does not determine the ordering of their derivatives). We note these claims as a result of our baseline calibration.

A.1 Derivation of the SCC Decomposition

The planner's value function $V(S, A)$ solves the HJB equation. An additional unit of carbon lowers welfare for two reasons. First, it increases flow damages. The temperature changes according to the transient slope net of the deployment response, $dT/dS = \Psi'(S) - \theta g'(S)$, and acidification damages increase with the stock of carbon. Under the thermostat policy, the temperature component equals zero above the activation threshold. Second, a marginal unit of carbon increases the ecological debt, increasing termination damages. Integrating the two channels using the stochastic discount factor over all future dates yields the decomposition of the carbon price in Equation (17). □

A.2 Derivation of the Stochastic Discount Factor

The representative investor has CRRA utility over consumption $C_t = A_t c(S_t)$, meaning that marginal utility is proportional to $C_t^{-\gamma}$. Applying Ito's lemma for jump-diffusion processes, with the carbon-cycle Brownian shock and the three Poisson events for termination and the two breakthroughs, yields Equation (21). At termination, consumption is scaled by the productivity loss and by the abatement contraction induced by the higher post-termination tax,

$$\frac{C_{t+}}{C_{t-}} = (1 - \Gamma(F_t)) \cdot J_t^c \tag{28}$$

and the pricing-kernel jump is this consumption ratio raised to the power $-\gamma$. □

A.3 Derivation of the Expected Return Equation

A no arbitrage condition delivers the expected return equation. The product of the pricing kernel and any asset value, plus accumulated dividends, must be a martingale:

$$d(\Lambda_t V_t^i) = V_{t-}^i d\Lambda_t + \Lambda_{t-} dV_t^i + d\langle \Lambda^c, V^{i,c} \rangle_t + (\Delta\Lambda_t)(\Delta V_t^i) \tag{29}$$

Using Ito's lemma for jump-diffusions to expand this product, dividing by its current value, taking conditional expectations, and using $\mathbb{E}_t[d\Lambda_t/\Lambda_t] = -r_t\, dt$ delivers Equation (25), the

expected excess return as the sum of the three jump risk premia. □

A.4 Proof of Proposition 4

Fix a carbon stock S and a research direction i. The patent value $\hat{\pi}_i(S)$ is the value of the profits earned after the breakthrough, priced under the model's stochastic discount factor along the path started at S. Compare the Regime 1 and Regime 2 economies after the breakthrough. They solve the same optimizing problem, except that Regime 2 values temperature damages at the masked temperature. Its carbon tax is therefore weakly lower at all states, and strictly lower above the activation threshold (Proposition 1). Green profits $D^G(\tau)$ increase with the tax, and removal profits $(\tau - c_R)^+\bar{R}$ never decrease with it. Profits are therefore weakly lower at all points of the Regime 2 path, and strictly lower on states reached with positive probability. Since the weights of the pricing kernel are strictly positive, $\hat{\pi}_i^{R2}(S) < \hat{\pi}_i^{R1}(S)$, the research intensities inherit this ranking through Equation (19). To prove the exact zero result, we note that the Regime 2 tax under the thermostat equals the acidification component $\tau_{acid}(S)$ in all states where it is deployed. If this component remains below the removal cost c_R along the whole path after a CDR breakthrough, then removal profits are zero pointwise. This means that the removal patent earns nothing in deployment states. The condition holds throughout the reporting range in our calibration, and the small remaining patent value comes from paths that diffuse below the activation threshold.

For part (b), write each patent value as a flow part plus a gain at termination. The green gain is bounded, uniformly in the ecological debt, by the capped green dividend flow discounted at the risk-free rate, $D^G(\tau_{cap})/r$. The removal gain grows with the ecological debt, because termination at a higher carbon stock means the Regime 1 price starts further above the removal cost and more carbon must be removed before the price falls back to it. Wherever the green bound binds and the removal gain is still growing, the ratio η_C/η_G rises from Regime 2 to Regime 3. Our calibration satisfies this condition over the reporting range below the cap on the breakthrough rate. □

A.5 The Landing-Probability System

The landing probability is the solution to the following system of four linear differential equations on the same grid as the value functions. These are solved backward in the same pass. The two removal-era survival probabilities q_C and q_B are the solutions to the following Feynman-Kac equations on their economies' equilibrium dynamics (with termination hazard live while deployment is active and absorption at the effective safety boundary). The equation for q_C includes the source term $\eta_G^C(S), q_B(S)$, because a green breakthrough during the

removal era moves the economy to the faster removal era with both technologies. The last two equations control the landing probability p and the corresponding probability p_G in the economy with only the green breakthrough:

$$0 = \tfrac{1}{2}\sigma_S^2\, p'' + e(\tau(S))\, p' - [\lambda(S) + \eta_C(S) + \eta_G(S)]\, p + \eta_C(S)\, q_C(S)\, \mathbf{1}\{g > 0\} + \eta_G(S)\, p_G(S) \tag{30}$$

$$0 = \tfrac{1}{2}\sigma_S^2\, p_G'' + e_G(S)\, p_G' - \left[\bar{\lambda} + \eta_C^G(S)\right] p_G + \eta_C^G(S)\, q_B(S) \tag{31}$$

Here $e_G(S)$ and $\eta_C^G(S)$ are the equilibrium net emissions and removal research intensity of the economy with only the green breakthrough. Each source term is a breakthrough rate multiplied by a survival probability. A breakthrough that is later followed by termination does not count as a landing. □

B Post-Termination Dynamics

After termination, geoengineering stops ($g = 0$, $\lambda = 0$) and the world continues as in Regime 1 with the inherited carbon stock. Productivity is reduced by the termination loss, $A_{t+} = A_{t-}(1 - \Gamma(F_t))$. Dividends and removal profits are scaled down by the same factor relative to a Regime 1 economy.

The post-termination value of a marginal unit of removal capacity is the damaged Regime 1 value at the inherited carbon stock,

$$q_{t+}(S) = (1 - \Gamma(F_t)) \cdot q^{R1}(S) \tag{32}$$

A marginal unit does not change the aggregate drift, so this value is priced on the ordinary Regime 1 path, with its carbon-price revenue falling as the equilibrium evolves.

The value of removal capacity at termination ($J^{CDR} = q_{t+}/q_{t-}$) depends on the state. At carbon stocks just above the activation threshold, capacity gains value at termination since the post-termination carbon price is much higher than the price during operation of geoengineering. This gain reverses near 599 ppm in our baseline calibration, beyond which capacity loses value at termination. At higher carbon stocks, the pre-termination value of removal capacity already includes the high carbon price and the value of a possible termination, while the post-termination value is scaled down by the productivity loss. Thus, capacity loses value at termination and its premium turns positive.