Optimal Climate Policy Mix (Extended Abstract)

1 Introduction

Climate change poses complex macro-financial risks that require a portfolio of policy responses. This paper develops a stochastic general-equilibrium model to study the optimal mix of climate policies – combining mitigation (emissions reduction via carbon pricing and clean technology investment) and adaptation (investment in resilience) – in the presence of uncertain climate dynamics and risks. We incorporate multiple interacting instruments (carbon taxes, clean R&D subsidies, adaptation spending, and negative emissions) into a unified social planner problem, extending standard single-policy Integrated Assessment Models. The model features a recursive utility preference structure and a rich production technology, alongside a detailed stochastic climate system with carbon-cycle feedbacks and low-probability tipping point shocks. We solve for the planner's optimal controls and characterize the Hamilton–Jacobi–Bellman (HJB) equations that govern the dynamic optimization. Quantitatively, our results demonstrate a strong complementarity between mitigation and adaptation: each policy enhances the effectiveness of the other. Higher carbon prices reduce future climate damages (lowering required adaptation costs), while early adaptation investments shield the economy and smooth consumption through climate shocks, facilitating more aggressive long-run mitigation. The optimal policy mix is time-varying – the relative emphasis on carbon pricing versus resilience-building evolves as climate risks and economic conditions change. We find that policy recommendations are sensitive to the assumed damage function (how damages scale with temperature) and climate sensitivity. When damages escalate sharply with warming, the planner optimally sets a high carbon tax and rapid emissions cuts; under milder damage scenarios, adaptation still plays a critical role in managing localized impacts. Overall, an integrated policy portfolio yields substantial welfare gains by reducing climate-induced volatility and improving macro-financial stability relative to single-instrument approaches.

2 Model Framework

Preferences: The representative agent has Epstein–Zin recursive utility, which disentangles risk aversion (γ) from intertemporal substitution elasticity (ψ). The continuation value V_t satisfies a stochastic differential utility representation. For consumption C_t and continuation value V_t, the planner's aggregator $f(C, V)$ is defined such that the value function evolves as:

$$V_t = E_t\left[\int_t^\infty f(C_s,\ V_s)\, ds\right],$$

with $f(C, V)$ normalized so that γ governs relative risk aversion and ψ the elasticity of intertemporal substitution (β is the time-preference/discount rate). This recursive formulation captures long-run risk and the planner's aversion to uncertain, delayed climate damages. Higher risk aversion (larger γ) increases the weight on worst-case climate outcomes, while lower ψ (EIS) makes the planner less willing to shift consumption over time, both affecting the optimal policy paths.

Production Technology: Output is produced using a nested CES-in-Cobb-Douglas production function. We consider three types of capital: clean capital K_c (e.g. renewable energy infrastructure), dirty capital K_d (fossil-fuel-based capital), and residual capital K_r (general infrastructure and other capital). Clean and dirty capital are imperfect substitutes with constant elasticity of substitution ϑ_{cd}, which we parameterize via $\rho_{cd} = (\vartheta_{cd} - 1)/\vartheta_{cd}$. This clean-dirty aggregate then combines with residual capital in a Cobb-Douglas fashion. Formally, letting $X_t = K_{c,t} + K_{d,t} + K_{r,t}$ be total productive capital and $\omega_{j,t} = K_{j,t}/X_t$ the share of type $j \in \{c, d, r\}$, the production function is:

$$Y_t \;=\; A\Big(\big(\omega_{c,t}^{\rho_{cd}} + \omega_{d,t}^{\rho_{cd}}\big)^{1/\rho_{cd}}\Big)^{\alpha_{cd}}\; K_{r,t}^{\,\alpha_r},$$

with $\alpha_{cd} + \alpha_r = 1$. Equivalently, output can be written as $Y_t = A_t^{\text{eff}} X_t$, linear in total capital X_t but with time-varying productivity $A_t^{\text{eff}} = A\,\omega_{r,t}^{\alpha_r}\big(\omega_{c,t}^{\rho_{cd}} + \omega_{d,t}^{\rho_{cd}}\big)^{\alpha_{cd}/\rho_{cd}}$ that depends on the composition of capital. This structure captures the idea that shifting the mix toward clean capital raises effective productivity by avoiding climate damages, while reliance on dirty capital can drag down productivity through environmental externalities.

Investment and Technological Change: The representative agent (or planner) allocates investment $I_{j,t}$ across the three capital types $j \in \{c, d, r\}$ and also into research & development (R&D) for clean technology. Investment in physical capital increases $K_{j,t}$ (subject to adjustment costs), while R&D spending $I_{\text{RD},t}$ advances an endogenous clean technology factor $\nu_{c,t}$. We model $\nu_{c,t}$ as the productivity or efficiency level of the clean sector, which evolves according to:

$$d\nu_{c,t} = \Big(\theta_{RD} + \epsilon_{IR}\frac{I_{\text{RD},t}}{X_t} - \delta_\nu\,\nu_{c,t}\Big)dt \;+\; \sigma_\nu\,dB_t^\nu,$$

where θ_{RD} is exogenous baseline improvement, ϵ_{IR} is the effectiveness of R&D investment in raising clean productivity, δ_ν is depreciation of knowledge, and σ_ν the volatility of tech progress. This formulation captures endogenous technological change: higher R&D investment accelerates growth in clean TFP $\nu_{c,t}$, reducing future clean energy costs. In the decentralized economy, there may be underinvestment in R&D due to knowledge spillovers, which justifies an optimal R&D subsidy in the planner's solution.

Mitigation and Adaptation Policy Instruments: In our centralized (social planner) scenario, the government internalizes environmental externalities by using a portfolio of tools:

- A carbon price (tax on emissions) that disincentivizes dirty capital use. In the decentralized interpretation, the carbon tax adjusts the private return on dirty vs. clean

investment, effectively shifting investment toward clean capital. The shadow price of carbon in the planner's problem corresponds to the social cost of carbon, reflecting the marginal damage of emissions.

- Clean technology subsidies (modeled as a temporary subsidy $G_{s,t}$ to R&D or clean investment) to overcome innovation market failures and jump-start the transition to a low-carbon economy.

- Adaptation investment $G_{a,t}$ which builds resilience (e.g. flood defenses, hardened infrastructure). Adaptation directly reduces the impact of climate change on the economy by lowering damage intensity. In the model, higher $G_{a,t}$ reduces the magnitude and/or frequency of climate-induced losses – for example, by lowering the probability of disaster or the fraction of capital destroyed when extreme events occur. This is captured via an adaptation effectiveness function $f_a(G_{a,t})$ that attenuates the damage from climate shocks.

- Negative emissions investment $G_{cc,t}$ (carbon capture and storage or other carbon dioxide removal). This is a direct mitigation tool that removes CO_2 from the atmosphere. We model carbon removal as an action that subtracts from the atmospheric carbon stock at a rate $\Psi(G_{cc,t})$, with diminishing returns (concave capture function) reflecting increasing marginal cost of removal.

All government expenditures $G_{i,t}$ (on adaptation, carbon removal, R&D subsidies, etc.) plus private consumption and investment must satisfy the resource constraint: $Y_t = C_t + \sum_j I_{j,t} + \sum_i G_{i,t}$ at each instant. The planner thus allocates output among consumption, various investments, and policy expenditures optimally.

3 Stochastic Climate System

The climate is represented by a stochastic model with feedbacks and shocks, driving economic damages and informing optimal policy. We track atmospheric CO_2 concentration and global temperature in a stylized but quantitatively calibrated system:

Carbon Stock Dynamics: Let $E_{S,t}$ denote the excess atmospheric CO_2 stock (or an equivalent measure of cumulative emissions). Emissions flow into E_S from economic activity, and outflows represent natural absorption and intentional removal. The evolution is given by:

$$\dot{E}_{S,t} \;=\; \zeta_{ET}\Big(F_{\text{emit}}(I_{d,t}) \;-\; \delta_{ES}E_{S,t} \;-\; \Psi(G_{cc,t})\Big) \;+\; \text{shock}_E(t).$$

Here $F_{\text{emit}}(I_{d,t})$ is the CO_2 emission rate from dirty capital usage (proportional to dirty investment or output; e.g. $F_{\text{emit}} = \eta_d I_{d,t}$ with η_d the emissions intensity), δ_{ES} is the natural carbon removal rate (ocean and biosphere uptake), and ζ_{ET} converts between emission units and concentration. Negative emissions enter as $-\Psi(G_{cc,t})$, reducing atmospheric CO_2. The term $\text{shock}_E(t)$ represents exogenous jumps in carbon – tipping-point events that release large CO_2 spurts. Specifically, we include a Poisson carbon-cycle tipping shock with some

arrival rate (hazard) λ_E that, when it occurs, instantaneously increases $E_{S,t}$ by a random amount (e.g. simulating permafrost melt releasing methane/CO_2, or forest dieback). These irreversible emission shocks capture low-probability high-impact events in the carbon cycle. Moreover, climate feedback effects are incorporated: higher temperatures can weaken natural CO_2 uptake or even turn sinks into sources. We introduce a state variable CF_t representing a climate feedback factor (e.g. net carbon released from feedback processes). Its dynamics (calibrated to represent phenomena like permafrost carbon feedback) adjust with temperature – for instance, $dCF_t = \kappa_{CF}\big(\theta_{CF} - \nu_{TS}T_{S,t} - CF_t\big)dt$ – meaning as surface temperature T_S rises, the equilibrium level of feedback emissions $\theta_{CF} - \nu_{TS}T_S$ increases, pushing CF_t upward. A higher CF_t then effectively adds to atmospheric CO_2 by reducing natural absorption (this appears as the $\xi_{CF}CF_t$ term reducing the sink in the E_S equation). Through CF_t, positive feedback loops can accelerate CO_2 accumulation if warming continues unchecked.

Temperature Dynamics: We adopt a two-layer global temperature model. Let $T_{S,t}$ be the surface temperature anomaly (degrees above pre-industrial) and $T_{L,t}$ the lower ocean temperature (deep ocean heat content). The temperature system is governed by:

$$\dot{T}_{S,t} = \alpha_{ES}\, E_{S,t} \;-\; \kappa_{TS}\, T_{S,t} \;+\; \alpha_T(T_{L,t} - T_{S,t}),$$
$$\dot{T}_{L,t} = \kappa_T(T_{S,t} - T_{L,t}) \;,$$

where α_{ES} is a climate sensitivity parameter (impact of CO_2 on warming), and κ_{TS}, κ_T govern heat exchange and dissipation between surface and deep ocean. This structure implies that rising atmospheric carbon drives surface warming, which gradually diffuses to the deep ocean. The two layers converge to equilibrium over decades, capturing thermal inertia. Higher emissions push E_S up, increasing the radiative forcing on T_S; if emissions are curbed, natural cooling ($-\kappa_{TS}T_S$) dominates and surface temperature can stabilize. Importantly, stochastic climate shocks also impact the economy: besides the carbon tipping events described above, we include climate disasters—sudden shocks directly to economic output or capital. In the model, these are represented as Poisson jumps in the aggregate capital/wealth X_t. For example, an extreme hurricane or wildfire might instantly destroy a fraction of X_t. We let such events occur with an intensity (hazard rate) that increases with temperature $T_{S,t}$, reflecting that higher global temperatures lead to more frequent and severe natural disasters. Each disaster reduces X_t by a random proportion (drawn from a distribution of damage severities). Adaptation investment $G_{a,t}$ enters by mitigating these shocks – effectively, G_a can lower the hazard (frequency) $\lambda_X(T_S)$ or the impact distribution of disasters (e.g. through better defenses). In our formulation, the expected drift of wealth X_t includes a term $-f_a(G_{a,t})\,\Lambda(T_{S,t})$ representing the reduced loss from disasters due to adaptation efforts (where $\Lambda(T_S)$ is the baseline temperature-dependent loss rate, and $f_a(\cdot)$ is increasing in G_a). Thus, mitigation (lower E_S, T_S) and adaptation (G_a) work in tandem to curb the probability and impact of catastrophic shocks.

4 Planner's Optimal Control Problem and HJB Structure

The social planner (benevolent government) chooses consumption and investment paths to maximize the representative agent's welfare, taking into account the coupled economic-climate system. Formally, the planner solves the optimal control problem:

$$V_0 \;=\; \sup_{\{C_t,\, I_{c,t}, I_{d,t}, I_{r,t}, I_{\mathrm{RD},t},\, G_{a,t}, G_{cc,t}, G_{s,t}\}_{t\geq 0}} E_0\left[\int_0^\infty f(C_t,\; V_t)\, dt\right],$$

subject to the entire set of state dynamics (capital accumulation equations, $\nu_{c,t}$ evolution, $E_{S,t}$ carbon cycle, $T_{S,t}, T_{L,t}$ climate system, etc.) and the resource constraint $Y_t = C_t + I_{c,t} + I_{d,t} + I_{r,t} + I_{\mathrm{RD},t} + G_{a,t} + G_{cc,t} + G_{s,t}$ for all t. The vector of state variables in $V_t = V(X_t, \nu_{c,t}, G_{s,t}, E_{S,t}, CF_t, T_{S,t}, T_{L,t})$ captures the economy's wealth and technology (X, ν_c), any temporary policy state like the R&D subsidy G_s (which could have its own decay, e.g. phasing out over time), and the climate states (E_S, CF, T_S, T_L). The planner internalizes that dirty investment I_d raises emissions (driving E_S up), while adaptation G_a and carbon removal G_{cc} can offset damages, etc.

We solve this infinite-horizon stochastic control problem via dynamic programming. The associated Hamilton–Jacobi–Bellman equation for the value function $V(\mathrm{state}_t)$ is:

$$0 \;=\; \max_{C_t,\, I_{j,t},\, G_{i,t}} \Big\{\, f(C_t,\; V_t) \;+\; \mathcal{A}V(X_t, \nu_{c,t}, G_{s,t}, E_{S,t}, CF_t, T_{S,t}, T_{L,t}) \,\Big\},$$

where $\mathcal{A}$ is the infinitesimal generator of the controlled state process (the differential operator capturing the drift of V given the chosen controls and stochastic shocks). This HJB equation characterizes the first-order conditions for optimal policies at each instant. In practice, it yields intuitively interpretable policy rules: for example, the optimal carbon tax equates the marginal welfare gain of emitting one less unit of CO_2 (in terms of reduced future climate damages) to the marginal utility cost of lower current consumption; similarly, the optimal adaptation investment equalizes the marginal benefit of risk reduction (lower expected damage from disasters) to the opportunity cost in terms of foregone output. The HJB embeds precautionary motives due to risk aversion – the planner weighs not just expected outcomes but also the distribution of shocks, assigning higher shadow costs to policies that expose the economy to significant risk. Solving the HJB (e.g. via numerical methods) gives the value function V and feedback control policies $C^*(\cdot), I_j^*(\cdot), G_i^*(\cdot)$ as functions of the state. These optimal policies define the climate policy mix: an intertemporal plan for carbon pricing, clean investment, adaptation spending, and carbon removal that maximizes social welfare.

5 Quantitative Results: Optimal Policy Paths and Interactions

Mitigation–Adaptation Complementarity: A central finding is that mitigation and adaptation are complementary strategies, not substitutes. In the optimal mix, each policy

enhances the effectiveness of the other. High carbon taxes (mitigation) reduce the accumulation of CO_2 and lower the probability of extreme warming, which in turn curbs the frequency of disaster shocks – this makes a given level of adaptation investment more effective (since there are fewer or less severe events to protect against). Conversely, early and substantial adaptation spending (resilience infrastructure) stabilizes the economy in the face of near-term climate shocks, preventing large losses of capital. This stability maintains economic and political capacity to implement stringent mitigation over time. In other words, adaptation "buys time" for mitigation by preserving growth and reducing volatility, thereby avoiding situations where climate damages spiral and render high carbon prices economically or politically untenable. Numerically, we find that a policy portfolio of both carbon pricing and adaptation yields higher welfare than either policy alone, confirming that a balanced approach can smooth consumption and investment paths despite climate risk. Each instrument addresses different temporal and spatial aspects of climate risk: carbon pricing deals with the long-run systemic risk of high global temperatures, while adaptation deals with short-run idiosyncratic shocks and localized damages. The two in tandem provide a more robust shield against uncertainty. Notably, the model indicates some trade-off at the margin: if adaptation technology becomes dramatically more effective (e.g. better flood defenses), the optimal carbon tax is slightly lower, as the economy can tolerate marginally higher emissions when it is very resilient. However, this effect is modest for realistic parameters – even with strong adaptation, aggressive mitigation remains necessary to limit long-term warming. Likewise, if mitigation becomes cheaper (e.g. via a breakthrough in clean tech), the planner can ease adaptation spending a bit. Overall, the optimal solution is characterized by positive co-deployment of both strategies, rather than an "all-in" on one or the other.

Dynamic Time Paths of Carbon Price and Adaptation: The optimal carbon price (social cost of carbon) in our model is forward-looking and state-contingent, rising or falling depending on climate trajectory and technological progress. In a baseline calibration, we find the optimal carbon tax is high from the outset to immediately start bending down the emissions curve. This high initial carbon price reflects the planner's desire to avoid irreversible carbon tipping risks and to capitalize on early abatement opportunities (when the atmosphere has less accumulated CO_2). Over time, as the economy decarbonizes and CO_2 concentrations stabilize, the carbon price path can adjust. If negative emission technologies scale up (due to G_{cc} investments), the effective carbon price may eventually level off or decline in the very long run because the option to pull CO_2 out of the atmosphere reduces the required tax pressure on remaining emissions. Correspondingly, adaptation investment follows a distinct trajectory: the planner ramps up G_a early on to build a resilient capital stock before the worst climate impacts hit. By investing in adaptation early (e.g. strengthening infrastructure), the economy experiences smaller losses from the initial waves of climate-related disasters, which pays off in higher capital accumulation and output in later periods. As global temperature stabilizes (due to mitigation efforts), the needed adaptation expenditures can plateau – essentially, once a certain level of resilience is achieved and as the risk of additional warming is curbed, adaptation spending stabilizes at a maintenance level (to account for the remaining but controlled risk). We also observe state-dependent adjustments: if a big climate shock occurs early (e.g. an unlucky disaster or tipping event),

the planner responds by temporarily increasing adaptation spending (to rebuild and further fortify defenses) and also raising the carbon tax (to further discourage emissions and hedge against worse outcomes). Thus, the policy paths are contingent on realized climate shocks, illustrating the model's capacity to adjust policies in response to new information or events – an important aspect of dynamic climate risk management.

Sensitivity to Damage Functions and Climate Parameters: The optimal policy mix is highly sensitive to assumptions about how damages scale with temperature and the probability of extreme events. We analyze scenarios with different damage functions – for instance, one where damages (or disaster intensity) rise convexly with T_S (steep acceleration of harm at higher temperatures) versus a more linear or concave damage response. In cases of steep damage-temperature relationships (or high climate sensitivity α_{ES}), the planner prioritizes aggressive mitigation: carbon taxes are set very high to avoid reaching those dangerous temperature thresholds, and clean investment is accelerated, as preventing an extra degree of warming yields enormous avoided damage. Adaptation is still used, but it cannot fully counteract the systemic risk of runaway damages, so mitigation takes the lead. By contrast, in scenarios with more moderate climate sensitivity or less convex damages, the relative emphasis shifts – the planner still prices carbon and invests in clean tech (to an extent consistent with long-run targets like the Paris Agreement), but recognizes that completely eliminating climate risk is unlikely or less urgent, so it leans more on adaptation to handle the moderate level of damages that do occur. For example, if each degree of warming only slowly increases disaster risk, it might be optimal to accept a bit more warming and deal with the consequences via resilience measures, rather than incur extremely high abatement costs. Our model quantifies these trade-offs, showing that mitigation and adaptation expenditures should be calibrated to damage curvature: when damages are highly nonlinear, it pays to invest more in mitigation upfront, whereas if damages are linear, a balanced or adaptation-heavy strategy can suffice without much welfare loss. We also examine the impact of climate-tipping probabilities: if the likelihood of a catastrophic carbon release (tipping event) is higher, the carbon tax acquires an additional precautionary premium – essentially, the mere risk of a huge jump in E_S pushes the optimal policy toward more stringent emission cuts at all times to keep temperatures as low as possible (reducing the chance of tipping). In summary, the model highlights that robust policy design must account for these deep uncertainties. The value of flexibility (option to adjust G_a and carbon taxes as we learn more about climate response) is significant, and the optimal policy mix under uncertainty often involves over-investment in precautionary measures (mitigation) in the near term, combined with adaptation to protect against shocks, especially in high-damage scenarios.

Financial Stability and Economic Resilience: A notable quantitative result is that the integrated policy mix greatly improves the economy's macro-financial resilience. By dampening climate-related shocks to production and wealth, adaptation policies reduce the volatility of consumption and investment. Mitigation lowers the long-run variance of the climate state (CO_2 and temperature), thereby shrinking the tail risk of extremely adverse outcomes (such as an unchecked warming causing a collapse in economic output). Together, these effects mean that the risk-adjusted growth path under the optimal mix is much

smoother and more stable. In simulations, we see that without adaptation, the economy experiences larger abrupt drops in capital (when disasters strike), causing persistent recessions that also discourage investment (including clean investment), potentially leading to financial stress (e.g. credit defaults or asset fire-sales in a market context). Without mitigation, on the other hand, the gradual accumulation of climate risk leads to growing probabilities of catastrophic shocks, which raises risk premia and depresses asset values as investors anticipate future losses. The planner's strategy of coordinated mitigation and adaptation avoids these pitfalls: adaptation insures the economy in the short run, while mitigation insures the long run. The outcome is lower risk in both the real economy and the financial system – fewer sudden asset write-downs from disasters and a more predictable environment for long-term investment. This has an analog in a decentralized economy: optimal climate policy contributes to financial stability by reducing the physical climate risks that can destabilize banks and markets, and by providing a clearer long-term price signal (carbon price path) that guides investment smoothly rather than forcing abrupt transitions later. Our findings underscore that climate policy is not only about environmental targets but also about managing transition risks and physical risks in the macroeconomy.

6 Implications for Climate Risk Pricing and Policy

Climate Risk Pricing: The model's optimal carbon tax can be interpreted as the price of climate risk in the economy. In our recursive utility setting, this price reflects both expected damages and the uncertainty/risk around those damages. The presence of stochastic shocks and risk aversion means the planner's carbon price is higher than it would be under risk-neutral expected value calculations – essentially there is a climate risk premium. This premium compensates for the possibility of extreme events and tipping points. Thus, one implication is that asset markets and policymakers should price carbon at a level that accounts for risk, not just average damages. If translated to a market setting, assets exposed to carbon or climate risk (e.g. fossil fuel stocks, coastal real estate) should carry a discount reflecting this risk, and conversely, investments in resilience or clean technology have additional value by lowering the economy's risk profile. Our framework provides a way to calculate the social cost of carbon (SCC) under uncertainty: it is essentially the shadow price in the HJB – the Lagrange multiplier on the carbon stock equation – which we find is elevated by the possibility of disasters. This has practical policy relevance: standard IAM estimates of the SCC might be too low if they ignore tail risks, whereas our results suggest incorporating risk and ambiguity could justify much higher carbon prices or more urgent mitigation (consistent with emerging climate finance literature on the pricing of rare disasters and uncertainty).

Macroeconomic Resilience: A key takeaway is the importance of adaptation for macroeconomic and financial resilience. Even with an optimal carbon tax, some amount of warming and climate shocks are unavoidable (due to already accumulated greenhouse gases and inertia in mitigation). Our results show that investing in adaptive capacity – e.g. climate-resilient infrastructure, disaster preparedness, drought-resistant agriculture – has outsized benefits in terms of maintaining economic stability. For policymakers, this means that fiscal spending on adaptation is an investment in economic resilience that can pay for itself by averting

large future losses. Furthermore, adaptation helps maintain productive capital in the long run, which supports growth and also means more resources are available for further mitigation or consumption. In a finance context, greater resilience lowers the default risk of loans, stabilizes insurance markets, and reduces the chance that climate events trigger systemic financial crises. Thus, climate adaptation policy should be seen as part of a broader strategy for financial stability and sustainable development, not merely as a humanitarian or environmental expenditure.

Capital Allocation and Transition Dynamics: The optimal policy mix will significantly redirect capital flows in the economy. With credible carbon pricing and R&D support, we expect to see accelerated investment in clean capital and gradual phasing out of dirty capital. Our model suggests an initial surge in clean energy investment (amplified by the R&D subsidy that temporarily boosts returns to clean innovation) leading to a clean technology boom. Dirty industries, meanwhile, face divestment as their effective cost of capital rises with the carbon tax. However, because we include adaptation and negative emissions, the transition is smoother: even sectors that cannot decarbonize immediately (e.g. certain industrial processes or aviation) can continue operating for some time because adaptation measures blunt the impact of their emissions and eventual carbon removal can offset residual emissions. In essence, capital is optimally reallocated across four fronts: (1) scaling up clean productive capital, (2) maintaining essential residual capital, (3) fortifying assets via adaptation capital (which doesn't produce new output but preserves existing output), and (4) building carbon removal capacity for the future. Policymakers can facilitate this reallocation by providing clear long-term price signals (a rising carbon price trajectory) and by directly investing in public goods (resilience and innovation). One implication is that stranded asset risks (the risk that dirty capital becomes unusable) are mitigated under the optimal mix, because the policy is announced and implemented gradually and paired with adaptation – giving firms time to adjust and protecting the economy from sudden shocks. This contrasts with delayed action scenarios, where a sudden late introduction of a high carbon tax could cause abrupt asset stranding and financial turmoil. Our findings reinforce the view that a predictable, well-communicated policy path is crucial: it allows the private sector to reallocate capital in advance, minimizing transitional disruptions and supporting a robust economy throughout the decarbonization process.

Fiscal Policy and Public Finance: Implementing the optimal climate policy mix has important fiscal implications. The carbon tax generates substantial public revenue (especially in early decades when emissions are still high). In the planner's setup, these revenues are implicitly redistributed lump-sum or used to fund the adaptation and clean tech subsidies. One could imagine a real-world policy using carbon tax revenues to finance adaptation projects and R&D subsidies, as our model would prescribe. This recycling of carbon revenue creates a win-win: it corrects the externality and finances further risk reduction measures. Our results suggest that an efficient fiscal policy package would include: (i) a robust carbon tax whose proceeds are earmarked in part for climate investment, (ii) temporary support for nascent clean technologies (declining as those technologies become competitive), and (iii) sustained funding for adaptation and negative emissions R&D. Over time, as private clean

investments become self-sustaining and carbon emissions fall (shrinking the tax base), the government may need to find other revenue or reallocate budget to continue funding adaptation and any remaining carbon removal efforts. However, by that point, the economy is larger and more resilient, so its capacity to shoulder such expenditures is improved. We also note that coordination between monetary, fiscal, and regulatory policy is implied: for instance, financial regulators might incorporate the presence of optimal adaptation in their stress tests (since the government is addressing physical risk proactively), and central banks might support the transition by maintaining stable financial conditions as capital reallocates. While our model is a social planner analysis, the insights translate into guidance for real-world policymakers designing comprehensive climate policy: a mix of pricing carbon, subsidizing innovation, investing in public resilience, and enabling carbon removal is required to achieve climate goals in a cost-effective and financially stable manner.

Conclusion: This research provides a theoretical and quantitative case for an integrated climate policy portfolio. The optimal climate policy mix simultaneously targets the root cause of climate change (GHG emissions) and its symptoms (climate damages), using all available levers – carbon pricing, technological innovation, adaptation infrastructure, and carbon dioxide removal. Such a coordinated approach not only improves welfare by lowering the long-run temperature trajectory and smoothing consumption, but also ensures a more resilient macro-financial system in the face of climate risks. These findings inform policymakers that meeting ambitious climate targets (such as the Paris Agreement) necessitates multi-faceted action: pricing carbon alone is not enough; it must be complemented by proactive adaptation and innovation policy. The framework developed here can be used to evaluate policy trade-offs under uncertainty, helping to design climate strategies that are robust against the unknowns of future climate sensitivity, damage intensities, and tipping-point probabilities. The clear message is that an optimal mix of mitigation and adaptation is mutually reinforcing – by embracing both, we can achieve a sustainable and stable economic future even as we confront the profound challenge of climate change.