Mitigating Climate Risk in a Present-Biased World*

Yuan Li† Tak-Yuen Wong‡ Siqi Zhao§

June 2, 2026

Abstract

Present bias distorts intertemporal decision-making by overweighting upfront abatement costs relative to long-term climate resilience. We incorporate household present bias into a model where carbon accumulation accelerates capital depreciation and mitigation is a public good. The first-best allocation is decentralized via a carbon tax and an investment subsidy. Quantitatively, optimal policy levels are remarkably robust to present bias, but welfare gains from these policies depend critically on the severity of the bias. Decomposing these gains reveals a dynamic policy prioritization: carbon taxes are the primary welfare driver under low climate risk, whereas investment subsidies become paramount to build capital buffers as risks intensify. Finally, the value of commitment hinges on the elasticity of intertemporal substitution, governing the tension between splurging and smoothing motives.

Keywords: Climate Policy, Present Bias, Social Cost of Carbon, Recursive Utility, Value of Commitment
JEL Classification: Q54, H23, E22, E71

*We are grateful to seminar and conference participants at the Graduate Institute of Industrial Economics of National Central University, Shanghai University of Finance and Economics, and Taiwan Finance Association for helpful comments. Tak-Yuen Wong acknowledges financial support from the National Science and Technology Council of Taiwan (112-2410-H-007-058-MY2).

†School of Finance, Shanghai University of Finance and Economics; Shanghai Institute of International Finance and Economics. Email: yuanli4911@163.com.

‡Department of Quantitative Finance, National Tsing Hua University. Email: etywong110@gmail.com.

§School of Economics, Fudan University; Shanghai Institute of International Finance and Economics. Email: zhaosiqi.shufe@gmail.com.

1 Introduction

The fight against climate change is fundamentally an intertemporal problem. Because the costs of emission abatement are immediate while its benefits accrue over centuries, policy prescriptions depend critically on time preferences—a point central to the well-known debate between Nordhaus (2008) and Stern (2007). Indeed, even minor variations in discount rates can drastically alter optimal carbon taxes (Gerlagh and Liski, 2017). Yet, most climate-economy models typically assume time-consistent exponential discounting, ignoring substantial empirical evidence that individuals exhibit present bias (Laibson, 1997). This behavioral bias creates a profound distortion in climate economics: by systematically overweighting upfront costs relative to delayed rewards, present-biased decision-makers underinvest in mitigation. Consequently, standard policy prescriptions based on time consistency may prove inadequate, necessitating a reassessment of optimal climate interventions.

This paper aims to investigate how present bias reshapes the design of optimal climate policy. Given that mitigating climate risk demands not only emission abatement but also building a physical capital buffer against future damages, we examine policy interventions across both fronts. We show that the first-best allocation can be decentralized through a policy mix: a carbon tax and an investment subsidy. A striking insight of our analysis is the asymmetric impact of present bias: while the optimal levels of these policy instruments remain remarkably robust to the bias, the welfare gains they generate are highly sensitive to it. Consequently, present bias fundamentally reshapes dynamic policy prioritization. Furthermore, we reveal that a commitment device—designed to curb present-biased overconsumption—is not universally welfare-improving, but instead hinges critically on the household's elasticity of intertemporal substitution (EIS).

To derive these results, we develop a continuous-time stochastic general equilibrium model that integrates climate externalities, capital accumulation, and present-biased households. Consistent with the climate economics literature, production generates carbon emissions that accumulate in the atmosphere (e.g., Nordhaus, 2008; Golosov et al., 2014; Van den Bremer and Van der Ploeg, 2021). However, departing from standard models that focus on the level effects of climate damages, we emphasize the growth effects of climate change: atmospheric carbon accumulation accelerates the depreciation of physical capital (Dell et al., 2012). This implies that climate risk exerts a continuous influence on economic growth, equilibrium asset prices, and the marginal value of capital (Tobin's average q). To capture households' time inconsistency, we adopt a quasi-hyperbolic recur-

sive utility framework developed by Shigeta (2022), which integrates recursive preferences (Duffie and Epstein, 1992) with the present-future framework of Harris and Laibson (2013). This specification disentangles risk aversion from the EIS, allowing us to accurately evaluate how households trade off long-run climate uncertainties against intertemporal consumption smoothing. Crucially, the homogeneity property of our framework implies that the equilibrium dynamics are entirely governed by a single aggregate state variable: the ratio of the atmospheric carbon stock to the physical capital stock. Because new carbon emissions dominate natural decay at low carbon-capital ratios, but are offset by decay as the ratio rises, the economy converges to a unique steady state.

Governed by these aggregate dynamics, the decentralized economy is subject to fundamental externalities. First, climate risk mitigation is a public good. While individual firms can invest in emission abatement technologies to slow carbon accumulation, they fail to internalize the social benefits of reduced climate damages. Consequently, decentralized firms free-ride on aggregate mitigation efforts, resulting in zero private abatement. Second, aggregate capital accumulation generates a positive externality. Expanding the physical capital stock acts as a buffer against climate damages by lowering the carbon-capital ratio, thereby diluting the economy's overall exposure to climate risk. Individual firms ignore this social value and thus systematically underinvest relative to the social optimum. To quantify these distortions, we characterize the decentralized market equilibrium under the assumption that households are sophisticated—meaning they rationally anticipate the short-sighted behavior of their future selves—and contrast it with the first-best allocation chosen by a social planner who inherits the household's present bias.

To correct these dual externalities, we demonstrate that the first-best allocation can be decentralized through a policy package centered on a Pigouvian carbon tax and an investment subsidy. Consistent with the literature (e.g., Golosov et al., 2014; Van den Bremer and Van der Ploeg, 2021), the optimal carbon tax rate equals the social cost of carbon derived from the planner's problem, directly addressing the emissions externality. Complementing this, the investment subsidy aligns private investment with the social optimum, incentivizing firms to build sufficient physical resilience against future climate damages. While the equivalence between the decentralized market equilibrium and the first-best allocation is a standard result under time consistency, establishing this equivalence in a continuous-time framework with present bias is highly non-trivial. Because present-biased preferences induce time inconsistency, the proof requires matching both the current and continuation value functions of the households and the social planner, while accounting for strategic interactions among future selves. By overcoming this technical challenge, we provide a

tractable framework for incorporating behavioral frictions into dynamic general equilibrium models with externalities.

Building on this theoretical setup, we proceed to a quantitative analysis to derive policy implications. Central to our quantitative exercise is the household's horizon, a parameter that governs the severity of present bias. Intuitively, a shorter horizon implies that the future self arrives sooner, severely restricting the current self's ability to commit to long-run plans. Given that the fight against climate change demands sustained intertemporal effort, this inability to commit emerges as a critical friction. By varying this horizon, we explore how present bias reshapes three key dimensions of the economy: optimal policy design, the associated welfare gains, and the value of commitment devices.

Our first set of results examines how present bias influences the design of optimal climate policies. Strikingly, we find that both the optimal carbon tax and the optimal investment subsidy are largely insensitive to the household's horizon. This near-irrelevance result suggests that optimal policy levels derived from standard time-consistent models are remarkably robust to the degree of present bias. To ensure this finding is not driven by a general increase in impatience, we contrast the effects of present bias with variations in the long-run (exponential) discount rate. In stark contrast, even minor changes in the long-run discount rate drastically alter optimal policy levels. Consequently, our analysis yields a crucial policy implication: while implementing optimal climate policies requires a precise calibration of the long-run discount rate, policymakers need not precisely estimate the degree of households' present bias.

Second, despite the insensitivity of optimal policy levels to present bias, the household's horizon significantly affects the welfare gains derived from these policy interventions. Our quantitative analysis shows that jointly implementing the optimal carbon tax and investment subsidy increases welfare—measured by the certainty equivalent of the capital stock—by approximately 1.2% near the steady state. To understand the relative contribution of each instrument, we decompose these aggregate gains by implementing one policy at a time. When the carbon-capital ratio is low or households are relatively forward-looking, carbon pricing remains the primary instrument and accounts for a large share of the total welfare gains. Conversely, as atmospheric carbon accumulates and climate risks intensify, the relative importance of the investment subsidy rises substantially. These findings offer crucial guidance for dynamic policy prioritization: when facing implementation constraints in an economy transitioning toward its steady state, policymakers should prioritize carbon pricing; however, in high-risk or severely short-sighted environments, subsidizing capital

accumulation becomes paramount.

Third, we explore the value of commitment devices, which extend the household's horizon to restrict present-biased behavior. Formally, we quantify this value using the equivalent variation in aggregate capital required to make households indifferent between a baseline economy with a short horizon and a counterfactual with an extended horizon. We demonstrate that the desirability of such commitment hinges crucially on the EIS, which governs the tension between a splurging motive and a smoothing motive. In a high-EIS regime, households are highly willing to substitute consumption intertemporally; here, this commitment device increases welfare by disciplining the splurging motive and preventing underinvestment. Conversely, in a low-EIS regime where households are highly averse to consumption fluctuations, external commitment reduces welfare by constraining the flexibility needed to smooth consumption against climate shocks. Furthermore, we reveal that asset illiquidity reshapes these welfare effects along two distinct dimensions. On the one hand, physical capital illiquidity—induced by capital adjustment costs—acts as an endogenous substitute for external commitment. On the other hand, environmental illiquidity—governed by the slow decay of atmospheric carbon—operates as a friction that depresses the marginal return on investment, thereby reducing the value of commitment across all preference regimes. Finally, we confirm the robustness of our qualitative insights to alternative climate damage specifications, and extend our framework to explore how present bias interacts with the temporal resolution of long-run uncertainty.

Related literature. Our paper contributes to the literature on optimal climate policy in integrated assessment models. From the seminal work of Nordhaus (2008) to recent dynamic general equilibrium models (e.g., Golosov et al., 2014; Barrage, 2020; Van den Bremer and Van der Ploeg, 2021), a central feature of this literature is the assumption of time-consistent exponential discounting. Even recent studies that incorporate stochastic tipping points, model uncertainty, and disaster risks (Cai and Lontzek, 2019; Barnett et al., 2020; Hong et al., 2023) continue to rely on a time-consistent representative agent. This standard assumption, however, stands in stark contrast to substantial empirical evidence that individuals exhibit present bias (Laibson, 1997). We advance this literature by characterizing the optimal policy mix when the social planner inherits the households' present-biased preferences. Crucially, we uncover a striking asymmetric impact: while standard time-consistent models yield robust prescriptions for optimal policy levels, they miscalculate the associated welfare gains, thereby distorting dynamic policy prioritization between carbon taxes and investment subsidies.

This paper also builds on foundational studies of climate policy under non-constant time preferences (Karp, 2005; Karp and Tsur, 2011), and extends the recent quantitative literature on climate commitment (Gerlagh and Liski, 2017; Iverson and Karp, 2021). To ensure analytical tractability, existing studies typically rely on the discrete-time frameworks of Krusell et al. (2002, 2003) with simplifying assumptions, which limit the richness of capital dynamics and utility specifications. While yielding tractable closed-form solutions, these settings inevitably abstract from key macroeconomic frictions and climate uncertainties. We overcome these limitations by developing a continuous-time stochastic model featuring capital adjustment costs and quasi-hyperbolic recursive utility (Shigeta, 2022). Crucially, this preference specification disentangles risk aversion from the EIS. This methodological departure allows us not only to characterize how climate risk and time inconsistency jointly shape the stochastic discount factor and equilibrium asset prices, but also to uncover how illiquidity frictions and the EIS determine the value of commitment. Specifically, we demonstrate that physical capital illiquidity acts as an endogenous commitment device that restrains present-biased splurging. This insight bridges our framework with the recent macroeconomic literature on illiquidity and self-control (Maxted, 2025; Acharya et al., 2026; Beshears et al., 2025), extending their mechanisms to climate risk and general equilibrium policy design.

Section 2–4 contain the model setup, equilibrium characterization, and optimal policy design; derivations and proofs are in Appendices A–B. Sections 5–6 present the quantitative results and further discussions. Section 7 concludes. Online Appendix covers the model extensions under naive beliefs.

2 The Setting

Time is continuous with an infinite horizon, $t \in [0, \infty)$. The economy is populated by a continuum of identical households and firms, both with a unit measure. For notational clarity, we use **boldfaced** letters to denote aggregate variables, distinguishing them from their firm-level counterparts.

2.1 Technology and the Climate Block

Firm production and capital accumulation. A firm produces output proportional to its capital stock K_t. That is, its output is AK_t, where $A > 0$ is a constant productivity parameter.

Let I_t denote the firm's investment, then the evolution of the firm's capital stock is given by:

$$dK_t = (I_t - (\delta_K + \mathcal{D}(\mathbf{s}_t))K_t)\, dt + \sigma_K K_t d\mathcal{B}_t^K, \tag{1}$$

where $\mathcal{B}_t^K$ is a standard Brownian motion capturing continuous diffusion shocks to capital, δ_K is the depreciation rate of capital, and σ_K denotes the diffusion volatility.

In (1), $\mathcal{D}(\mathbf{s}_t)$ denotes the damage function, which is increasing in the carbon-capital ratio $\mathbf{s}_t \equiv \mathbf{S}_t/\mathbf{K}_t$. Here, $\mathbf{S}_t$ represents the stock of atmospheric carbon dioxide that exceeds the pre-industrial level attributable to human activities. This specification captures the negative effect of climate change on capital accumulation: a higher atmospheric carbon stock increases climate damages and accelerates capital destruction.[1] Our modeling approach is consistent with Dell et al. (2009; 2012) and Hambel et al. (2021), who emphasize the impact of climate change on growth rates rather than output levels. Accordingly, our framework differs from standard Integrated Assessment Models (see, e.g., Nordhaus, 2008; Golosov et al., 2014; Hassler et al., 2021), which focus on the level effects of carbon accumulation on economic losses.[2]

Carbon emission and mitigation technology. Let $E_t = E(M_t, K_t)$ denote the individual firm's net carbon emissions. A firm's baseline carbon emissions are increasing in output. Hence, more capital leads to higher carbon emissions. Firms are endowed with a mitigation technology: more mitigation spending M_t implies lower carbon emissions. We assume that this abatement technology is homogeneous of degree one in capital, that is $E(M_t, K_t) = e(m_t)K_t$ with $e'(m_t) < 0$, indicating that the emission reduction depends on the mitigation spending intensity relative to the capital stock, $m_t \equiv M_t/K_t$. Additionally, the marginal return to mitigation efforts is decreasing in m_t, which implies $e''(m_t) > 0$. Crucially, the mitigation of climate risk operates at the aggregate level, as the evolution of the climate state depends on the collective contributions of all firms. Consequently, aggregate carbon reduction is a public good.

The dynamics of the atmospheric carbon stock $\mathbf{S}_t$ are given by:

$$d\mathbf{S}_t = (\mathbf{E}_t - \delta_S \mathbf{S}_t)\, dt + \sigma_S \mathbf{S}_t d\mathcal{B}_t^S, \tag{2}$$

[1] Because our model economy is populated with a continuum of identical households and firms, the average of a micro-level variable equals the corresponding variable in the aggregate. For instance, the average of K_t equals the aggregate $\mathbf{K}_t$. Similarly, the average of I_t equals the aggregate $\mathbf{I}_t$. Our aggregation result is based on the exact law of large numbers (Duffie and Sun, 2012).

[2] Implicitly, the total damage $\mathcal{D}(\mathbf{s}_t)K_t$ is homogeneous of degree one in both the carbon stock and aggregate capital. An alternative approach is to model climate damages using a Poisson jump process as in the disaster-risk literature (see, e.g., Barro, 2006; Pindyck and Wang, 2013; Hong et al., 2023). Given an adaptation policy, disaster risk results in increased capital depreciation on average. However, our approach is numerically simpler.

where $\mathbf{E}_t$ denotes aggregate anthropogenic carbon emissions.[3] The parameter δ_S is the rate of natural carbon decay, capturing carbon absorption by natural sinks such as oceans, while σ_S denote the volatility of the carbon stock driven by the standard Brownian motion $\mathcal{B}_t^S$. This Brownian motion captures unexpected shocks to atmospheric carbon concentration arising from environmental events such as volcanic eruptions, earthquakes, or other man-made greenhouse gas emissions unrelated to production. The correlation coefficient between $\mathcal{B}_t^S$ and $\mathcal{B}_t^K$ is ϑ.[4]

In our economy, the carbon-capital ratio $\mathbf{s}_t$ serves as the key state variable. Applying Ito's Lemma to $\mathbf{s}_t$, it evolves according to:

$$d\mathbf{s}_t/\mathbf{s}_t = \mu_{\mathbf{s}}(\mathbf{s}_t)dt + \sigma_S d\mathcal{B}_t^S - \sigma_K d\mathcal{B}_t^K, \tag{3}$$

with the drift $\mu_{\mathbf{s}}(\mathbf{s}_t)$ given by

$$\mu_{\mathbf{s}}(\mathbf{s}_t; \mathbf{i}_t, \mathbf{m}_t) = \frac{\mathbf{e}(\mathbf{m}_t)}{\mathbf{s}_t} - \delta_S - \mathbf{i}_t + \delta_K + \mathcal{D}(\mathbf{s}_t) + \sigma_K^2 - \vartheta\sigma_S\sigma_K. \tag{4}$$

Here, $\mathbf{e}(\mathbf{m}_t)$, $\mathbf{i}_t$, and $\mathbf{m}_t$ denote the aggregate carbon emissions, investment, and mitigation scaled by capital $\mathbf{K}_t$, respectively. The squared volatility of $\mathbf{s}_t$ is a constant given by $\Sigma_{\mathbf{s}} \equiv \sigma_S^2 - 2\vartheta\sigma_S\sigma_K + \sigma_K^2$. Intuitively, the carbon-capital ratio increases with carbon emissions but decreases with capital accumulation and the natural decay of atmospheric carbon.

Firm's objective. The firm's net cash flows (dividends), Y_t, are given by:

$$Y_t = AK_t - (I_t + \Phi_t) - M_t,$$

which represent output net of investment costs and mitigation spending. Following the q theory of investment (Hayashi, 1982; Jermann, 1998), the total investment cost includes the capital adjustment cost, $\Phi_t = \Phi(I_t, K_t)$. As in Hayashi (1982), the adjustment cost is homogeneous of degree one in the capital stock K_t, that is, $\Phi(I_t, K_t) = \phi(i_t)K_t$, where $i_t \equiv I_t/K_t$ denotes the firm's investment-capital ratio and $\phi(i)$ is increasing and convex.

Given the representative household's equilibrium stochastic discount factor (SDF) $\mathbb{M}_t$, the firm

[3] We measure the stock of atmospheric carbon dioxide $\mathbf{S}_t$ in parts per million (ppm); accordingly, aggregate carbon emissions $\mathbf{E}_t$ are also expressed in ppm under (2). In practice, emissions are measured in gigatons of carbon (GtC), and the units can be readily converted by appropriately scaling $\mathbf{E}_t$. See, e.g., Hambel et al. (2021).

[4] Studies that employ Integrated Assessment Models typically model the evolution of carbon stock or temperature changes based on climate science (see, e.g., Nordhaus, 2008; Golosov et al., 2014; Gerlagh and Liski, 2017; Cai and Lontzek, 2019; Barrage, 2020). The specification (2) using a geometric Brownian motion has proved useful in continuous-time climate-economy models (Van den Bremer and Van der Ploeg, 2021; Hong et al., 2023).

chooses investment I_t and mitigation spending M_t to maximize its value Q_0 at $t = 0$:

$$Q_0 = \max_{\{I_t, M_t\}_{t\geq 0}} \mathbb{E}\left[\int_0^\infty \frac{\mathbb{M}_t}{\mathbb{M}_0} Y_t dt\right]. \tag{5}$$

The equilibrium SDF reflects both the time value of money and the risk premium and represents the representative household's marginal rate of substitution (MRS).

2.2 Households and Present-biased Preferences

Present-biased preferences. Households are short-term-focused and their rates of time preference decline over time. Formally, we endow households with quasi-hyperbolic preferences (Phelps and Pollak, 1968; Krusell et al., 2003) to capture present bias. Decision making with quasi-hyperbolic preferences is formalized as an intrapersonal game among multiple selves (Laibson, 1997).

Specifically, consider self 1 is born at time t_1. She regards $[t_1, t_2)$ as the current period and treats $[t_2, \infty)$ as the future periods, and in the current period, she controls equilibrium decisions. At time t_2, self 2 replaces self 1. She takes control in the new current period $[t_2, t_3)$ and treats $[t_3, \infty)$ as the new future periods. Thus, $T_n = t_{n+1} - t_n$ is the lifespan of self n. Assume that the lifespan of each self is exponentially distributed with the constant parameter $\xi > 0$. That is, the transition between selves is regulated by a Poisson process with a hazard rate of ξ. Consequently, the expected lifespan of each self, defined as $1/\xi$, represents the household's effective horizon.

Repeating the process, we can obtain a sequence of selves $n \in \{1, 2, ...\}$. They live in a sequence of time intervals $\{[t_1, t_2), [t_2, t_3), ...\}$, which are independent and identically distributed (i.i.d.). All selves discount exponentially with a discount rate $\rho > 0$, and each self further discounts her future by an additional quasi-hyperbolic discount factor $\beta \in (0, 1]$. This implies that, at time t, each self n applies the quasi-hyperbolic discount function $D_n(t, u)$ to future consumption at time u, where

$$D_n(t,u) = \begin{cases} e^{-\rho(u-t)}, & \text{if } u \in [t_n, t_{n+1}) \\ \beta e^{-\rho(u-t)}, & \text{if } u \in [t_{n+1}, \infty) \end{cases} \tag{6}$$

The magnitude of present bias is inversely related to the household's horizon $1/\xi$. A higher ξ implies a shorter horizon, which intensifies the present bias as the discrete discount jump β is expected to occur sooner. If $\xi = 0$, the horizon becomes infinite, and Equation (6) reduces to the traditional exponential discount function.

Models related to present-biased preferences require an assumption about how the current self forms expectations about the behavior of future selves, which influences their continuation value (O'donoghue and Rabin, 1999; O'donoghue and Rabin, 2001). In this paper, we focus on *sophisticated* households: the current self correctly perceives her future selves' present bias (Laibson, 1997). This framework for analyzing the planner's problem and market equilibrium under present bias follows Krusell et al. (2002) and Gerlagh and Liski (2017).[5]

Recursive utility. We also endow the representative household with recursive utility (Epstein and Zin, 1989), allowing us to disentangle the coefficient of relative risk aversion (RRA), γ, from the elasticity of intertemporal substitution (EIS), ψ. This separation is critical because, as we will show in Section 5, the welfare implications of present bias depend fundamentally on the magnitude of the EIS. We adopt the continuous-time formulation of Shigeta (2022), who integrates the stochastic differential utility model of Duffie and Epstein (1992) with the present-future framework of Harris and Laibson (2013). The lifetime utility of the current self is given by

$$J_0 = \mathbb{E}\left[\int_0^{\tau} f(C_t, J_t)dt + \beta^{\theta}\tilde{J}_{\tau}\right], \tag{7}$$

where τ denotes the arrival time of the future self. The Duffie-Epstein-Zin normalized aggregator for consumption C and utility J is given by

$$f(C, J) = \rho\theta J\left(\frac{C^{1-\psi^{-1}}}{((1-\gamma)J)^{1/\theta}} - 1\right), \tag{8}$$

where $\theta \equiv \frac{1-\gamma}{1-\psi^{-1}}$. Given the future consumption stream $\left\{\tilde{C}_u : t \leq u < \infty\right\}$ perceived by the current self, the continuation utility $\tilde{J}_t$ is given by:

$$\tilde{J}_t = \mathbb{E}_t\left[\int_t^{\infty} f(\tilde{C}_u, \tilde{J}_u)du\right]. \tag{9}$$

Equations (7)-(9) show that in the recursive utility framework with quasi-hyperbolic discounting, the current utility is captured by the aggregator (8), while continuation utility from future periods is additionally discounted by the quasi-hyperbolic discount factor β.[6] Note that in (7), the future

[5]The alternative assumption is that households are *naive*, meaning the current self incorrectly anticipates that future selves will act in a time-consistent manner. The Online Appendix provides the solution under naive beliefs.

[6]It is well-known that recursive preferences disentangle the RRA from the EIS. If the EIS is the inverse of the RRA, i.e., $\gamma = \psi^{-1}$, Equations (7)-(9) reduce to CRRA preferences with quasi-hyperbolic discounting as in Harris and Laibson (2013). However, unlike Harris and Laibson (2013) and Maxted (2025), who consider instantaneous gratification ($\xi \to \infty$), we focus on $\xi < \infty$ and let the length of the present time vary.

continuation utility $\tilde{J}_t$ is weighted by β^θ. This is because the inverse of the certainty equivalent in the discrete-time recursive utility formulation is nonlinear; in the continuous-time limit, the power θ on β reflects such a nonlinearity.[7]

Finally, we assume that the arrival of the future self is an aggregate event affecting all households simultaneously. This assumption eliminates *ex post* heterogeneity, allowing households to remain identical and facilitating aggregation into a representative household with preferences given by (7)-(9). In this environment, the discount shock, i.e., the downward jump in the discount factor upon the arrival of future selves, induces systematic risk in the market economy.

2.3 Financial Markets

The financial market is dynamically complete, allowing households to trade the following financial securities: (i) a risk-free asset that pays interest at the equilibrium rate of r_t, and (ii) the aggregate equity market. Let $\mathbf{Q}_t$ denote the ex-dividend value of the aggregate stock market and $\mathbf{D}_t$ be the aggregate dividends. The cum-dividend return from holding the risky equity is given by:

$$\frac{d\mathbf{Q}_t + \mathbf{D}_{t-}dt}{\mathbf{Q}_{t-}} = \mu_{\mathbf{Q},t-}dt + \sigma^K_{\mathbf{Q},t-}d\mathcal{B}^K_t + \sigma^S_{\mathbf{Q},t-}d\mathcal{B}^S_t + \left(\frac{\tilde{\mathbf{Q}}_t}{\mathbf{Q}_{t-}} - 1\right)d\mathcal{P}_t, \tag{10}$$

where $\mu_{\mathbf{Q},t-}$ is the endogenous expected return, and $\sigma^K_{\mathbf{Q},t-}$ and $\sigma^S_{\mathbf{Q},t-}$ capture the endogenous exposures to capital and climate diffusion risks, respectively.

The final term in (10) captures the valuation risk arising from stochastic shifts in time preferences. Specifically, the arrival of a future self is governed by a Poisson process $\mathcal{P}_t$ with the constant intensity ξ. If no arrival occurs at time t (i.e., $d\mathcal{P}_t = 0$), the household's discount function remains unchanged and the equity value satisfies $\mathbf{Q}_t = \mathbf{Q}_{t-}$. By contrast, when a future self arrives ($d\mathcal{P}_t = 1$), the household's discount factor drops discretely, triggering an instantaneous revaluation of the aggregate equity from its pre-arrival value $\mathbf{Q}_{t-}$ to the post-arrival value $\tilde{\mathbf{Q}}_t$. Unlike standard jump-diffusion models in which asset price jumps are driven by fundamental shocks (e.g., climate disasters or productivity changes), this jump risk is endogenous: it originates from the time inconsistency of the household's preferences and reflects sudden changes in valuation rather than shocks to economic fundamentals.

[7] See Appendix B in Shigeta (2022) for details.

3 Competitive Markets Solution

This section characterizes the competitive equilibrium of the laissez-faire economy. We adopt the following equilibrium concept.

Definition. (**Competitive Equilibrium**) A *recursive competitive equilibrium* consists of a set of allocations (consumption C_t, stock market portfolio Γ_t, investment I_t, and mitigation spending M_t) and a set of prices (risk-free rate r_t, aggregate equity value $\mathbf{Q}_t$, and the SDF $\mathbb{M}_t$) such that:

1. *Household Optimization*: Taking the equilibrium risk-free rate r_t and aggregate equity value $\mathbf{Q}_t$ as given, the representative household chooses her consumption C_t and allocation to the aggregate stock market Γ_t to maximize the lifetime utility defined by (7)-(9).

2. *Firm Optimization*: Taking the equilibrium SDF $\mathbb{M}_t$ as given, the representative firm chooses investment I_t and mitigation expenditure M_t to maximize its market value given by (5).

3. *Market Clearing*: The interest rate r_t, the aggregate stock price $\mathbf{Q}_t$, and the SDF $\{\mathbb{M}_t; t \geq 0\}$ are consistent with the household's and firm's optimal decisions and all markets clear.

While the definition of a competitive equilibrium is standard, time-inconsistent preferences imply intrapersonal games within both the household's and the firm's optimization problems. Following Grenadier and Wang (2007) and Harris and Laibson (2013), optimal decisions are therefore characterized by Markov perfect equilibrium (MPE) strategies. Specifically, we focus on stationary MPE, in which the current self and all future selves use the same Markov strategies.[8] We provide the details of the household's and the firm's optimization problems in Sections 3.1 and 3.2, respectively. Finally, we present the characterization of the market equilibrium in Section 3.3.[9]

3.1 Household's Optimization

We solve for a stationary MPE of the household's optimization problem. Under the assumption of sophistication, the current household is fully aware of her time-inconsistent preferences. When choosing consumption C and portfolio Γ, she correctly anticipates that all of her future selves will follow strategies $\tilde{C}$ and $\tilde{\Gamma}$ based on their own preferences. These beliefs determine the continuation value, which the current household takes as given when choosing C and Γ. In a stationary MPE, all

[8] Unlike the instantaneous gratification model of Harris and Laibson (2013), which admits a unique equilibrium, other equilibria may exist with $\xi < \infty$. However, the focus on stationary MPE is the most natural, helps with the characterization of the competitive equilibrium, and is informative for policy implications.

[9] The time-consistent benchmark can simply be obtained by setting $\xi = 0$ or $\beta = 1$.

selves adopt the same consumption strategy, so that $C = \tilde{C}$. Note that the MPE does not require $\Gamma = \tilde{\Gamma}$, as portfolio choices must satisfy financial market clearing.

Dynamic programming. The representative household's payoff-relevant states are wealth, the atmospheric carbon stock, and aggregate capital. Due to homogeneity, wealth W and the carbon-capital ratio $\mathbf{s}$ are sufficient state variables. Let $J(W, \mathbf{s})$ denote the current value function of the representative household. Given the return process (10), the representative firm's policies $(\mathbf{i}, \mathbf{m})$, and the strategies of future selves $(\tilde{C}, \tilde{\Gamma})$, the current value function satisfies the Hamilton-Jacobi-Bellman (HJB) equation:

$$\begin{aligned} 0 = &\max_{C,\Gamma} f(C, J) + (rW + (\mu_{\mathbf{Q}} - r)\Gamma - C) J_W(W, \mathbf{s}) + \frac{\Sigma_W(\mathbf{s})}{2}\Gamma^2 J_{WW}(W, \mathbf{s}) \\ &+ \mu_{\mathbf{s}}(\mathbf{s}; \mathbf{i}, \mathbf{m})\mathbf{s} J_{\mathbf{s}}(W, \mathbf{s}) + \frac{\Sigma_{\mathbf{s}}}{2}\mathbf{s}^2 J_{W\mathbf{s}}(W, \mathbf{s}) + \Sigma_{W\mathbf{s}}(\mathbf{s})\mathbf{s}\Gamma J_{W\mathbf{s}}(W, \mathbf{s}) + \xi \left(\beta^\theta \tilde{J}(\tilde{W}, \mathbf{s}) - J(W, \mathbf{s})\right), \end{aligned} \tag{11}$$

where Σ_W is the squared volatility of the household's wealth, and $\Sigma_{W\mathbf{s}}$ is the instantaneous covariance between wealth and carbon-capital ratio. Both terms are driven by the endogenous price process (detailed derivations for these volatility components are provided in Appendix A.1, Equations (A.11) and (A.12)).

The first line of the HJB equation (11) is standard, comprising the normalized aggregator $f(C, J)$ and the expected value change induced by wealth dynamics. In the second line of (11), the first two terms capture the change in the value driven by the stochastic evolution of the carbon-capital ratio $\mathbf{s}$, while the third term represents the intertemporal hedging motive against diffusion risks arising from the covariation between wealth and the carbon-capital ratio.

The last term of (11) captures the expected value change induced by present bias. Upon the arrival of a future self, the current household loses control over consumption and portfolio choices and evaluates future utility using long-run exponential discounting. This transition creates a discontinuity in both wealth and valuation. First, financial wealth jumps from W_{t-} to the post-arrival level $\tilde{W}_t$ due to the instantaneous revaluation of equity holdings:

$$\tilde{W}_t = W_{t-} + \left(\frac{\tilde{\mathbf{Q}}_t}{\mathbf{Q}_{t-}} - 1\right)\Gamma_{t-}.$$

After this discount shock, there are no further preference shocks ($d\mathcal{P}_t \equiv 0$). Accordingly, $\tilde{W}_t$ evolves according to (A.3) and $\tilde{\mathbf{Q}}_t$ follows the continuation return process (A.2). Second, the representative

household's current value $J(W, \mathbf{s})$ switches to her continuation value $\tilde{J}(\tilde{W}, \mathbf{s})$, which is additionally discounted by the present-bias factor β^θ. The current household's continuation value $\tilde{J}(\tilde{W}, \mathbf{s})$ satisfies the following valuation equation:

$$\begin{aligned} 0 =& f(\tilde{C}, \tilde{J}) + \left(\tilde{r}\tilde{W} + (\tilde{\mu}_{\mathbf{Q}} - \tilde{r})\tilde{\Gamma} - \tilde{C}\right) \tilde{J}_W(\tilde{W}, \mathbf{s}) + \frac{\Sigma_{\tilde{W}}(\mathbf{s})}{2}\tilde{\Gamma}^2 \tilde{J}_{WW}(\tilde{W}, \mathbf{s}), \\ &+ \mu_{\mathbf{s}}(\mathbf{s}; \tilde{\mathbf{i}}, \tilde{\mathbf{m}})\mathbf{s}\tilde{J}_{\mathbf{s}}(\tilde{W}, \mathbf{s}) + \frac{\Sigma_{\mathbf{s}}}{2}\mathbf{s}^2 \tilde{J}_{\mathbf{ss}}(\tilde{W}, \mathbf{s}) + \Sigma_{\tilde{W}\mathbf{s}}(\mathbf{s})\mathbf{s}\tilde{\Gamma}\tilde{J}_{W\mathbf{s}}(\tilde{W}, \mathbf{s}). \end{aligned} \tag{12}$$

Equation (12) differs from (11) in three aspects. First, the household no longer faces discount shocks ($\xi = 0$). Second, the current self has no control over future choices ($\tilde{C}$, $\tilde{\Gamma}$). Third, the dynamics of the continuation wealth $\tilde{W}$ are determined by the current self's beliefs about future market prices ($\tilde{r}$, $\tilde{\mathbf{Q}}$), while the anticipated evolution of the aggregate carbon-capital ratio $\mathbf{s}$ is driven by the beliefs about future firm policies ($\tilde{\mathbf{i}}$, $\tilde{\mathbf{m}}$).

Homogeneity property. The value functions $J_t = J(W_t, \mathbf{s}_t)$ and $\tilde{J}_t = \tilde{J}(\tilde{W}_t, \mathbf{s}_t)$ are homogeneous of degree $1 - \gamma$ in wealth. Specifically, $J(W, \mathbf{s})$ and $\tilde{J}(\tilde{W}, \mathbf{s})$ take the following form:

$$J(W, \mathbf{s}) = \frac{(u(\mathbf{s})W)^{1-\gamma}}{1-\gamma} \text{ and } \tilde{J}(\tilde{W}, \mathbf{s}) = \frac{(\tilde{u}(\mathbf{s})\tilde{W})^{1-\gamma}}{1-\gamma}, \tag{13}$$

where $u(\mathbf{s})$ and $\tilde{u}(\mathbf{s})$ are the values of a dollar to the representative household to be determined in the market equilibrium. In this formulation, $u(\mathbf{s})W$ and $\tilde{u}(\mathbf{s})\tilde{W}$ represent the certainty equivalent wealth (CEW).

First-order conditions. The current self's HJB equation (11) yields the optimal rules for consumption and portfolio allocation. First, the FOC for consumption, $f_C(C, J) = J_W(W, \mathbf{s})$, equates the marginal utility of consumption to the marginal utility of wealth. Substituting (13) into this FOC, we obtain a linear consumption rule:

$$C(W, \mathbf{s}) = \rho^\psi u(\mathbf{s})^{1-\psi} W, \tag{14}$$

which shows that present bias affects optimal consumption decision through the marginal propensity to consume $u(\mathbf{s})$. Second, simplifying the FOC for the market portfolio allocation Γ yields the optimal demand for the risky asset:

$$\Gamma = -\frac{\mu_{\mathbf{Q}} - r}{\Sigma_W(\mathbf{s})}\frac{J_W(W, \mathbf{s})}{J_{WW}(W, \mathbf{s})} - \frac{\Sigma_{W\mathbf{s}}(\mathbf{s})}{\Sigma_W(\mathbf{s})}\frac{\mathbf{s}J_{W\mathbf{s}}(W, \mathbf{s})}{J_{WW}(W, \mathbf{s})} + \frac{\xi\beta^\theta}{\Sigma_W(\mathbf{s})}\frac{\tilde{J}_W(\tilde{W}, \mathbf{s})}{J_{WW}(W, \mathbf{s})}\left(1 - \frac{\tilde{\mathbf{Q}}}{\mathbf{Q}}\right). \tag{15}$$

Portfolio demand is driven by three components: (i) the standard Merton trade-off between risk and expected excess return; (ii) the hedging demand against intertemporal diffusion risk; and (iii) a novel term unique to our model: the hedging demand against discount shocks, which induces jump risk in both wealth W and the aggregate equity value $\mathbf{Q}$.

3.2 Firm Investment and Mitigation

Since households are the firm's shareholders, the representative firm inherits their present bias and behaves as if it has time-inconsistent preferences. In particular, the current firm forms rational beliefs about the investment and mitigation strategies $(\tilde{i}, \tilde{m})$ of future selves and, taking the continuation firm value $\tilde{Q}_t$ as given, chooses corporate policies $(i,\ m)$ optimally to maximize the current firm value Q_t. In a stationary MPE, $i = \tilde{i}$ and $m = \tilde{m}$.

The firm maximizes its value given by (5) taking the equilibrium SDF $\mathbb{M}_t$ as given. In Appendix A.3, we show that the dynamics of the SDF follow:

$$\frac{d\mathbb{M}_t}{\mathbb{M}_{t-}} = -r_{t-}dt - \eta^K_{\mathbb{M},t-}d\mathcal{B}^K_t - \eta^S_{\mathbb{M},t-}d\mathcal{B}^S_t + (\chi_t - 1)\left(d\mathcal{P}_t - \xi dt\right). \tag{16}$$

On the right-hand side of (16), the first three terms are standard: $r_{t-} = -\mathbb{E}_t\left[\frac{d\mathbb{M}_t}{\mathbb{M}_{t-}}\right]\frac{1}{dt}$ is the equilibrium risk-free rate (Duffie, 2010), while $\eta^K_{\mathbb{M},t-}$ and $\eta^S_{\mathbb{M},t-}$ represent the market prices of risk associated with capital and carbon diffusion shocks, respectively.

The novel effect of time-inconsistent preferences in equilibrium asset pricing is that it induces the jump term in (16).[10] Specifically, upon the arrival of a new self at time t, the SDF adjusts discretely from $\mathbb{M}_{t-}$ to $\tilde{\mathbb{M}}_t$ by the endogenously determined market price of the discount shock, χ_t:

$$\frac{\tilde{\mathbb{M}}_t}{\mathbb{M}_{t-}} = \chi_t.$$

More precisely, $\tilde{\mathbb{M}}_t$ is the continuation SDF of the current household in the recursive competitive equilibrium; it follows the same form as in (16) but without the jump term ($d\mathcal{P}_t \equiv 0$) and is driven by the continuation asset price characteristics $\tilde{r}_t$, $\tilde{\eta}^K_{\mathbb{M},t}$, and $\tilde{\eta}^S_{\mathbb{M},t}$ (see Equation (A.32)).

In the intrapersonal game, the firm's payoff-relevant states are its own capital stock, the atmospheric carbon stock, and aggregate capital. Due to homogeneity, the firm's own capital stock K and the carbon-capital ratio $\mathbf{s}$ are sufficient state variables. Applying Ito's Lemma to the firm

[10] Since $\mathbb{E}_t[d\mathcal{P}_t] = \xi dt$, this term is a compensated jump martingale under the physical measure.

value $Q_t = Q(\mathbf{s}_t, K_t) = q(\mathbf{s}_t)K_t$ and using (16), we obtain the following HJB equation for Tobin's q, $q(\mathbf{s})$:

$$\begin{aligned} r(\mathbf{s})q(\mathbf{s}) = &\max_{i,m} A - i - \phi(i) - m + (i - \delta_K - \mathcal{D}(\mathbf{s}))\, q(\mathbf{s}) - \sigma_K(\eta_{\mathbb{M}}^K + \vartheta\eta_{\mathbb{M}}^S)\left(q(\mathbf{s}) - \mathbf{s}q'(\mathbf{s})\right) \\ &+ \left(\frac{\mathbf{e}(\mathbf{m})}{\mathbf{s}} - \delta_S - \mathbf{i} + \delta_K + \mathcal{D}(\mathbf{s}) - \sigma_S(\eta_{\mathbb{M}}^S + \vartheta\eta_{\mathbb{M}}^K)\right)\mathbf{s}q'(\mathbf{s}) \\ &+ \frac{\Sigma_{\mathbf{s}}\mathbf{s}^2}{2}q''(\mathbf{s}) + \xi\chi\left(\tilde{q}(\mathbf{s}) - q(\mathbf{s})\right). \end{aligned} \tag{17}$$

where $\tilde{q}(\mathbf{s})$ is the continuation firm value that satisfies $\tilde{Q}_t = \tilde{q}(\mathbf{s}_t)K_t$.[11]

The HJB equation (17) decomposes the required return on the firm into dividends, expected capital growth (net of climate damages), the stochastic evolution of the carbon-capital ratio, and the effect of the household's present bias. When a future self arrives, the firm's valuation jumps from $q(\mathbf{s})$ to $\tilde{q}(\mathbf{s})$, and the current self evaluates this jump using the endogenous market price of the discount shock, $\chi(\mathbf{s})$. Consequently, present-biased preferences directly influence the equilibrium equity value through the SDF. These returns are further adjusted for the market prices of capital risk, $\eta_{\mathbb{M}}^K(\mathbf{s})$, and carbon risk, $\eta_{\mathbb{M}}^S(\mathbf{s})$.

Optimal investment and mitigation. The HJB equation (17) yields two optimality conditions. First, the FOC with respect to investment i is

$$q(\mathbf{s}) = 1 + \phi'\left(i(\mathbf{s})\right), \tag{18}$$

which aligns with the standard q-theory of investment: the firm invests until the marginal benefit of an additional unit of capital, $q(\mathbf{s})$, equals its marginal cost, $1 + \phi'\left(i(\mathbf{s})\right)$.

Second, the FOC with respect to the firm's mitigation spending m is[12]

$$q'(\mathbf{s})\frac{\partial \mathbf{e}(\mathbf{m})}{\partial m} = 1.$$

The right-hand side is the marginal cost of mitigation, while the left-hand side is the marginal private benefit: the increase in firm value from a marginal reduction in aggregate emissions. However, as an individual firm is infinitesimal, its mitigation choice has a negligible effect on aggregate

[11] Details regarding the continuation SDF and continuation firm value are provided in Step 5 of Appendix A.3.

[12] Note that mitigation spending m enters the HJB equation (17) implicitly through the drift of the carbon-capital ratio, $\mu_{\mathbf{s}}(\mathbf{s})$. Recall from the state dynamics (4) that $\mu_{\mathbf{s}}(\mathbf{s}_t) = \mathbf{e}_t/\mathbf{s}_t - \mathbf{i}_t + \delta_K + \mathcal{D}(\mathbf{s}_t) - \delta_S + \sigma_K^2 - \vartheta\sigma_K\sigma_S$. Therefore, applying the chain rule to the drift term $q'(\mathbf{s})\mu_{\mathbf{s}}(\mathbf{s})$ yields the marginal benefit term $q'(\mathbf{s})\frac{\partial \mathbf{e}(\mathbf{m})}{\partial m}$.

mitigation $\mathbf{m}$ and aggregate emissions $\mathbf{e}(\mathbf{m})$, implying $\frac{\partial \mathbf{e}(\mathbf{m})}{\partial m} = 0$. As a result, the marginal private benefit of mitigation is zero, and firms optimally choose $m = 0$ in the decentralized equilibrium. Our setting thus features a classic free-rider problem: while mitigation raises capital value by reducing climate damages, its benefits are non-rival and non-excludable. Each firm then provides no mitigation effort, since it generates positive social benefits at the expense of private costs.

3.3 The Competitive Equilibrium

We now impose market clearing conditions to characterize the competitive equilibrium. First, the market for the risky asset clears: $W_t = \Gamma_t = \mathbf{Q}_t$. That is, the current household invests all wealth in the aggregate equity market and correctly anticipates that future selves will do the same ($\tilde{W}_t = \tilde{\Gamma}_t = \tilde{\mathbf{Q}}_t$).[13] Second, the goods market clears: $\mathbf{C}_t = A\mathbf{K}_t - \mathbf{I}_t - \Phi(\mathbf{I}_t, \mathbf{K}_t)$. Aggregate consumption equals production net of investment costs, with zero aggregate mitigation ($\mathbf{M}_t = 0$) due to its public-good nature.

Existence and uniqueness of the steady state. A central property of our model is that the economy converges to a unique steady state, $\mathbf{s}^{\text{steady}}$, where the expected drift of the carbon-capital ratio is zero, i.e., $\mu_{\mathbf{s}}(\mathbf{s}^{\text{steady}}) = 0$. While endogenous policies ($\mathbf{m}(\mathbf{s})$ and $\mathbf{i}(\mathbf{s})$) rule out analytical proofs, the underlying economic mechanism is straightforward. Recall from Equation (4) that the drift is given by $\mu_{\mathbf{s}}(\mathbf{s}) = \mathbf{e}(\mathbf{m})/\mathbf{s} - \delta_S - \mathbf{i}(\mathbf{s}) + \delta_K + \mathcal{D}(\mathbf{s}) + \text{const}$. When the carbon-capital ratio $\mathbf{s}$ is close to zero, the term $\mathbf{e}(\mathbf{m})/\mathbf{s}$ approaches infinity, driving the drift strictly positive ($\lim_{\mathbf{s}\to 0} \mu_{\mathbf{s}}(\mathbf{s}) > 0$). Intuitively, when the atmospheric carbon stock is small relative to the economy, additional emissions dominate the natural decay, causing $\mathbf{s}$ to grow. Conversely, as $\mathbf{s}$ increases, the rapid decline in $\mathbf{e}(\mathbf{m})/\mathbf{s}$ and the natural carbon decay δ_S outweighs the opposing effect of climate damages $\mathcal{D}(\mathbf{s})$. As a result, $\mu_{\mathbf{s}}(\mathbf{s})$ is monotonically decreasing and crosses zero exactly once. We verify this single-crossing property numerically in Figure 1.

Equilibrium characterization. To simplify notation, for any generic function $g(\mathbf{s})$, we define its elasticity with respect to $\mathbf{s}$, and the product of the elasticities of $g'(\mathbf{s})$ and $g(\mathbf{s})$, as follows:

$$\epsilon_g(\mathbf{s}) \equiv \frac{g'(\mathbf{s})\mathbf{s}}{g(\mathbf{s})} \quad \text{and} \quad \epsilon'_g(\mathbf{s}) \equiv \frac{g''(\mathbf{s})\mathbf{s}^2}{g(\mathbf{s})} = \epsilon_{g'}(\mathbf{s})\epsilon_g(\mathbf{s}). \tag{19}$$

Applying Ito's Lemma to the equilibrium return and SDF processes, and imposing the stationary

[13] As discussed in Section 3.1, the condition $\tilde{W}_t = \tilde{\Gamma}_t = \tilde{\mathbf{Q}}_t$ does not imply $\Gamma_t = \tilde{\Gamma}_t$ because the discrete jump in discounting implies $\mathbf{Q}_t \neq \tilde{\mathbf{Q}}_t$ and $W_t \neq \tilde{W}_t$. In words, the condition simply states that the current household expects to hold the entire equity supply at the perceived price $\tilde{\mathbf{Q}}_t$.

MPE and market clearing conditions, we reduce the household's HJB equations and the firm's valuation equations to a system of four ordinary differential equations (ODEs) (A.17), (A.20), (17), and (A.33). This system jointly determines the current values, $u(\mathbf{s})$ and $\mathbf{q}(\mathbf{s})$, and the continuation values, $\tilde{u}(\mathbf{s})$ and $\tilde{\mathbf{q}}(\mathbf{s})$. Importantly, all equilibrium objects are functions of the carbon-capital ratio $\mathbf{s}$, which summarizes the aggregate state of the economy. We relegate the full derivation of the asset returns, SDF dynamics, and the complete ODE system to Appendix A.3.

Key features of the competitive equilibrium. First, the equilibrium consumption-to-q ratio (which corresponds to the marginal propensity to consume out of wealth, or the dividend yield) is given by:[14]

$$\begin{aligned}\frac{\mathbf{c}(\mathbf{s})}{\mathbf{q}(\mathbf{s})} =& \rho + (\psi^{-1}-1)\left[\mathbf{i}(\mathbf{s}) - \delta_K - \mathcal{D}(\mathbf{s}) - \frac{\gamma}{2}\sigma_K^2\right] \\ &+ (\psi^{-1}-1)\left[\alpha_{\mathbf{s}}(\mathbf{s};\mathbf{i},\mathbf{m})\epsilon_{u\mathbf{q}}(\mathbf{s}) + \frac{\Sigma_{\mathbf{s}}}{2}\left(\epsilon'_{u\mathbf{q}}(\mathbf{s}) - \gamma\epsilon_{u\mathbf{q}}(\mathbf{s})^2\right)\right] \\ &- \frac{\xi(\psi^{-1}-1)}{1-\gamma}\left[1 - \beta^{\theta}\left(\frac{\tilde{u}(\mathbf{s})}{u(\mathbf{s})}\frac{\tilde{\mathbf{q}}(\mathbf{s})}{\mathbf{q}(\mathbf{s})}\right)^{1-\gamma}\right],\end{aligned} \tag{20}$$

where $\alpha_{\mathbf{s}}(\mathbf{s};\mathbf{i},\mathbf{m}) = \mu_{\mathbf{s}}(\mathbf{s};\mathbf{i},\mathbf{m}) - (1-\gamma)(\sigma_K^2 - \vartheta\sigma_K\sigma_S)$. The first line of (20) is standard in AK model (e.g., Pindyck and Wang, 2013): the long-term time preferences ρ, the economic growth rate $\mathbf{i}(\mathbf{s}) - \delta_K - \mathcal{D}(\mathbf{s})$, and the precautionary saving motive summarized by RRA γ and growth volatility σ_K. Specifically, when the EIS $\psi < 1$, the wealth effect dominates the substitution effect, such that higher growth net of precautionary savings increases the MPC.

The novel effects of atmospheric carbon accumulation and carbon risk on the MPC are reflected in the second line of (20). Again, suppose EIS $\psi < 1$. Firstly, a higher carbon drift $\alpha_{\mathbf{s}}$ reduces future wealth because carbon accumulation has a destructive effect, $u'(\mathbf{s}) < 0$ and $\mathbf{q}'(\mathbf{s}) < 0$. The dominant wealth effect then implies a reduction in the current consumption. Secondly, a higher carbon volatility $\Sigma_{\mathbf{s}}$ implies greater uncertainty regarding future damages. Risk-averse households then increases savings to maintain future consumption.[15] Finally, the last term of (20) captures the impact of present bias, which disappears when $\psi = 1$ as the wealth and substitution effects exactly offset each other.

[14] Because total wealth equals the aggregate equity value, $W = \mathbf{q}(\mathbf{s})\mathbf{K}$, and consumption is fully financed by dividends in the competitive equilibrium, $\mathbf{c}(\mathbf{s})/\mathbf{q}(\mathbf{s})$ also corresponds to the dividend yield.

[15] More precisely, $u(\mathbf{s})\mathbf{q}(\mathbf{s})$ is convex, in that $\epsilon'_{u\mathbf{q}}(\mathbf{s}) > 0$. Thus, households may prefer carbon volatility. But, when households are sufficiently risk averse, then $\epsilon'_{u\mathbf{q}}(\mathbf{s}) < \gamma\epsilon_{u\mathbf{q}}(\mathbf{s})^2$, implying that $\mathbf{c}(\mathbf{s})/\mathbf{q}(\mathbf{s})$ decreases with $\Sigma_{\mathbf{s}}$ as long as $\psi < 1$.

Second, the equilibrium risk-free rate is given by

$$\begin{aligned} r(\mathbf{s}) =& \rho + \psi^{-1}\left(\mathbf{i}(\mathbf{s}) - \delta_K - \mathcal{D}(\mathbf{s})\right) - \frac{\gamma(\psi^{-1}+1)}{2}\sigma_K^2 \\ & - \mu_{\mathbf{s}}(\mathbf{s};\mathbf{i},\mathbf{m})\mathcal{E}(u,\mathbf{q}) - \frac{\Sigma_{\mathbf{s}}}{2}\bar{\mathcal{E}}(u,\mathbf{q}) \\ & + \xi\left[\chi(\mathbf{s})\left(\frac{\tilde{\mathbf{q}}(\mathbf{s})}{\mathbf{q}(\mathbf{s})} - 1\right) - \frac{\psi^{-1}-1}{1-\gamma}\left(1 - \beta^{\theta}\left(\frac{\tilde{u}(\mathbf{s})}{u(\mathbf{s})}\frac{\tilde{\mathbf{q}}(\mathbf{s})}{\mathbf{q}(\mathbf{s})}\right)^{1-\gamma}\right)\right], \end{aligned} \tag{21}$$

where $\mathcal{E}(u,\mathbf{q})$ and $\bar{\mathcal{E}}(u,\mathbf{q})$, defined in (A.31), denote the coefficients on the carbon drift $\mu_{\mathbf{s}}(\mathbf{s};\mathbf{i},\mathbf{m})$ and squared volatility $\Sigma_{\mathbf{s}}$, respectively. The first line on the right-hand side of (21) is standard in recursive utility models, capturing the effects of impatience, expected economic growth, and precautionary savings. The second line demonstrates that carbon dynamics introduce additional saving motives: expected future damages and climate diffusion risks both encourage investment in the risk-free bond. The final line captures the effects of present bias.

Finally, the equilibrium risk premium is given by:

$$\begin{aligned} rp(\mathbf{s}) =& \eta_{\mathbb{M}}^{K}(\mathbf{s})\Big(\sigma_K\left(1-\epsilon_{\mathbf{q}}(\mathbf{s})\right) + \vartheta\sigma_S\epsilon_{\mathbf{q}}(\mathbf{s})\Big) + \eta_{\mathbb{M}}^{S}(\mathbf{s})\Big(\sigma_S\epsilon_{\mathbf{q}}(\mathbf{s}) + \vartheta\sigma_K\left(1-\epsilon_{\mathbf{q}}(\mathbf{s})\right)\Big) \\ & - \xi(\chi(\mathbf{s})-1)\left(\frac{\tilde{\mathbf{q}}(\mathbf{s})}{\mathbf{q}(\mathbf{s})} - 1\right). \end{aligned} \tag{22}$$

Equation (22) decomposes the equity risk premium into three components. The first two terms represent the compensation for exposures to capital and carbon diffusion risks, priced at $\eta_{\mathbb{M}}^{K}(\mathbf{s})$ and $\eta_{\mathbb{M}}^{S}(\mathbf{s})$. The final term captures a novel jump risk premium induced by present bias. Specifically, it compensates investors for the discrete revaluation of aggregate equity, $\tilde{\mathbf{q}}(\mathbf{s})/\mathbf{q}(\mathbf{s}) - 1$, triggered by the stochastic arrival of a future self, evaluated at the endogenous market price of the discount shock, $\chi(\mathbf{s}) - 1$.

4 Social Planner's Problem

We now analyze the planner's economy and characterize the first-best allocation. Since the social planner's objective is to maximize the household's welfare, we assume that the planner shares the same present-biased preferences with the representative household. In other words, the planner is time-inconsistent. To highlight the first-best allocation, we use an asterisk (*) as a superscript whenever emphasis is needed.

4.1 Planner's Allocation

The current planner chooses the aggregate consumption $\mathbf{C}_t$, investment $\mathbf{I}_t$, and mitigation spending $\mathbf{M}_t$ to maximize the current household's welfare subject to the aggregate resource constraint: $A\mathbf{K}_t = \mathbf{C}_t + \mathbf{I}_t + \Phi(\mathbf{I}_t, \mathbf{K}_t) + \mathbf{M}_t$, given the rational anticipation that all future planners choose $(\tilde{\mathbf{C}}, \tilde{\mathbf{I}}, \tilde{\mathbf{M}})$. In a stationary MPE, $(\mathbf{C}, \mathbf{I}, \mathbf{M}) = (\tilde{\mathbf{C}}, \tilde{\mathbf{I}}, \tilde{\mathbf{M}})$.

Dynamic programming. Let $V(\mathbf{K}, \mathbf{S})$ be the planner's current value and $\tilde{V}(\mathbf{K}, \mathbf{S})$ be the continuation value. The current value function satisfies the following HJB equation:

$$\begin{aligned} 0 = \max_{\mathbf{C},\mathbf{I},\mathbf{M}} & f(\mathbf{C}, V) + (\mathbf{I} - \delta_K \mathbf{K} - \mathcal{D}(\mathbf{s})\mathbf{K})\, V_{\mathbf{K}}(\mathbf{K}, \mathbf{S}) + (\mathbf{E}(\mathbf{M}, \mathbf{K}) - \delta_S \mathbf{S})\, V_{\mathbf{S}}(\mathbf{K}, \mathbf{S}) \\ & + \frac{\sigma_K^2 \mathbf{K}^2}{2} V_{\mathbf{KK}}(\mathbf{K}, \mathbf{S}) + \frac{\sigma_S^2 \mathbf{S}^2}{2} V_{\mathbf{SS}}(\mathbf{K}, \mathbf{S}) + \vartheta \sigma_K \sigma_S \mathbf{K}\mathbf{S} V_{\mathbf{KS}}(\mathbf{K}, \mathbf{S}) \\ & + \xi \left(\beta^\theta \tilde{V}(\mathbf{K}, \mathbf{S}) - V(\mathbf{K}, \mathbf{S}) \right). \end{aligned} \tag{23}$$

On the right-hand side of (23), the first line captures the normalized aggregator and the drift effects of investment and mitigation. The second line captures the effects of capital and carbon diffusion risks on $V(\mathbf{K}, \mathbf{S})$. The last line incorporates the expected value jump driven by the arrival of the next future planner. Given the present-bias discount factor β and the curvature of the recursive utility captured by θ, the continuation value $\tilde{V}(\mathbf{K}, \mathbf{S})$ is weighted by β^θ. The current planner's continuation value then satisfies

$$\begin{aligned} 0 = & f(\tilde{\mathbf{C}}, \tilde{V}) + \left(\tilde{\mathbf{I}} - \delta_K \mathbf{K} - \mathcal{D}(\mathbf{s})\mathbf{K} \right) \tilde{V}_{\mathbf{K}}(\mathbf{K}, \mathbf{S}) + \left(\mathbf{E}(\tilde{\mathbf{M}}, \mathbf{K}) - \delta_S \mathbf{S} \right) \tilde{V}_{\mathbf{S}}(\mathbf{K}, \mathbf{S}) \\ & + \frac{\sigma_K^2 \mathbf{K}^2}{2} \tilde{V}_{\mathbf{KK}}(\mathbf{K}, \mathbf{S}) + \frac{\sigma_S^2 \mathbf{S}^2}{2} \tilde{V}_{\mathbf{SS}}(\mathbf{K}, \mathbf{S}) + \vartheta \sigma_K \sigma_S \mathbf{K}\mathbf{S} \tilde{V}_{\mathbf{KS}}(\mathbf{K}, \mathbf{S}), \end{aligned} \tag{24}$$

given the anticipation that future planners choose $(\tilde{\mathbf{C}}, \tilde{\mathbf{I}}, \tilde{\mathbf{M}})$.

FOCs for investment and mitigation. The FOC for the aggregate investment $\mathbf{I}$ is

$$(1 + \Phi_I(\mathbf{I}, \mathbf{K}))\, f_{\mathbf{C}}(\mathbf{C}, V) = V_{\mathbf{K}}(\mathbf{K}, \mathbf{S}), \tag{25}$$

which equates the marginal cost of accumulating capital (left-hand side) with its marginal benefit (right-hand side). The marginal cost is the product of the marginal utility of consumption $f_{\mathbf{C}}(\mathbf{C}, V)$ and the unit price of capital (normalized to one) plus the marginal adjustment cost $(1 + \Phi_I(\mathbf{I}, \mathbf{K}))$. The nonseparability of recursive preferences indicates that $f_{\mathbf{C}}(\mathbf{C}, V)$ depends on not only consumption $\mathbf{C}$ but also the utility V. The marginal benefit is the increase in the current planner's value

function, $V_{\mathbf{K}}(\mathbf{K}, \mathbf{S})$, from having an additional unit of capital.

The FOC for the aggregate mitigation spending $\mathbf{M}$ is

$$f_{\mathbf{C}}(\mathbf{C}, V) = -\mathbf{E}_{\mathbf{M}}(\mathbf{M}, \mathbf{K}) \left(-V_{\mathbf{S}}(\mathbf{K}, \mathbf{S})\right). \tag{26}$$

The current planner optimally equates the marginal cost of mitigation, measured by the forgone current consumption utility $f_{\mathbf{C}}(\mathbf{C}, V)$, with its marginal social benefit. This benefit is the utility gain from reducing carbon emissions, captured by the marginal reduction in emissions, $-\mathbf{E}_{\mathbf{M}}(\mathbf{M}, \mathbf{K})$, multiplied by the marginal disutility of the carbon stock, $-V_{\mathbf{S}}(\mathbf{K}, \mathbf{S})$.

Simplifying the model solution. Because the planner's problem is homogeneous of degree one in capital, doubling the capital stock $\mathbf{K}$ doubles optimal consumption $\mathbf{C}$, investment $\mathbf{I}$, and mitigation $\mathbf{M}$. The value functions are thus homogeneous of degree $1-\gamma$ in $\mathbf{K}$ and take the following form:

$$V(\mathbf{K}, \mathbf{S}) = \frac{(b(\mathbf{s})\mathbf{K})^{1-\gamma}}{1-\gamma} \quad \text{and} \quad \tilde{V}(\mathbf{K}, \mathbf{S}) = \frac{(\tilde{b}(\mathbf{s})\mathbf{K})^{1-\gamma}}{1-\gamma}, \tag{27}$$

where $b(\mathbf{s})$ and $\tilde{b}(\mathbf{s})$ represent the certainty equivalent wealth (CEW) per unit of capital. In Appendix B.1, we substitute the conjectured value functions (27) into the HJB equation (23), the continuation value equation (24), and the FOCs (25) and (26) to establish a system of four ODEs. This system jointly determines the first-best solution for $b(\mathbf{s})$, $\tilde{b}(\mathbf{s})$, $\mathbf{i}^*(\mathbf{s})$, and $\mathbf{m}^*(\mathbf{s})$. Optimal scaled consumption $\mathbf{c}^*$ is then given by the resource constraint: $\mathbf{c}^*(\mathbf{s}) = A - \mathbf{i}^*(\mathbf{s}) - \phi(\mathbf{i}^*(\mathbf{s})) - \mathbf{m}^*(\mathbf{s})$.

First-best allocation and asset pricing implications. Based on the FOC for investment (25), we define the first-best Tobin's q as

$$\mathbf{q}^*(\mathbf{s}) = \frac{1 + \phi'(\mathbf{i}^*(\mathbf{s}))}{1 - \epsilon_b(\mathbf{s})}. \tag{28}$$

Crucially, unlike Tobin's q in the decentralized market (Equation (18)), the first-best $\mathbf{q}^*(\mathbf{s})$ incorporates the elasticity term $\epsilon_b(\mathbf{s})$. This term captures the positive externality of capital accumulation: aggregate investment dilutes the economy's exposure to climate risk by lowering the carbon-capital ratio. Because private firms ignore this social benefit, a wedge emerges between the private and social marginal value of capital.

Analogous to the decentralized market, we can also derive the first-best consumption-to-q ratio $\mathbf{c}^*(\mathbf{s})/\mathbf{q}^*(\mathbf{s})$, the risk-free rate $r^*(\mathbf{s})$, and the risk premium $rp^*(\mathbf{s})$. These formulations share similar intuitions with their market counterparts, but are evaluated at the first-best allocations. We relegate

the explicit expressions for these asset pricing variables to Appendix B.2.

4.2 Market Economy With Optimal Taxes and Subsidies Attains First-Best

The laissez-faire market economy fails to achieve the first-best allocation due to two externalities. First, the public-good nature of aggregate mitigation induces a classic free-rider problem, resulting in zero private abatement. Second, individual firms fail to internalize the benefits of capital accumulation on reducing the economy's exposure to climate risk. Consequently, the decentralized market features both severe under-mitigation and suboptimal capital accumulation.

A natural question then arises: can the planner implement the first-best allocation in a decentralized market economy? In what follows, we show that a policy mix—a Pigouvian carbon tax, a capital rebate, an investment subsidy, and a lump-sum levy—aligns private incentives with the social optimum. Taking the policies of future selves as given, the current planner implements a carbon tax τ_t^c, a capital rebate τ_t^{reb}, and an investment subsidy τ_t^i. In equilibrium, the policy rules chosen by the current self coincide with those adopted by future selves. Let $\mathbf{L}_t$ be a lump-sum levy determined by aggregate states. A firm's net cash flow Y_t is:

$$Y_t = AK_t - (I_t + \Phi_t + M_t) - \tau_t^c E_t + \tau_t^{reb} K_t + \tau_t^i (I_t + \Phi_t) - \mathbf{L}_t. \tag{29}$$

Given the equilibrium SDF $\mathbb{M}_t$ and these policies, each firm chooses investment I_t and mitigation M_t to maximize its value.

Carbon tax. To correct the externality associated with climate mitigation, the planner levies a Pigouvian carbon tax, τ_t^c. This tax rate is precisely equated to the marginal social damage of emissions, quantified by the social cost of carbon (SCC). Following Cai and Lontzek (2019) and Van den Bremer and Van der Ploeg (2021), the SCC is the marginal disutility of emitting an additional unit of carbon, normalized by the marginal utility of consumption:

$$SCC(\mathbf{s}) = -\frac{V_{\mathbf{S}}(\mathbf{K}, \mathbf{S})}{f_{\mathbf{C}}(\mathbf{C}^*, V)} = -\frac{b'(\mathbf{s})}{\rho} \left(\frac{\mathbf{c}^*(\mathbf{s})}{b(\mathbf{s})} \right)^{\psi^{-1}}.$$

The planner's mitigation FOC (26) equates the marginal cost of abatement to its marginal social benefit. This implies the SCC equals the inverse of the absolute marginal mitigation effect on

emissions. Thus, the optimal carbon tax is:

$$\tau_t^c = \tau^c(\mathbf{s}_t) = SCC_t = \frac{1}{-\mathbf{E}_M(\mathbf{K}, \mathbf{M}^*)}.$$

Since aggregate mitigation is a public good, individual firms treat the aggregate mitigation intensity $\mathbf{m}^*(\mathbf{s}_t)$ and emission intensity $\mathbf{e}(\mathbf{m}_t^*)$ as given. Thus, a firm's carbon tax liability is $\tau_t^c E_t = \tau^c(\mathbf{s}_t)\mathbf{e}(\mathbf{m}_t^*)K_t$.

Nevertheless, since emission abatement has diminishing marginal returns, its marginal cost (the SCC) exceeds its average cost ($\mathbf{m}^*$). If the planner retains all carbon tax revenue, the total tax burden on firms would exceed actual mitigation spending. This excess burden would depress the after-tax return on capital and distort private investment incentives. To prevent this, the planner rebates the excess revenue to firms proportionally to their capital. The rebate rate τ_t^{reb} is the difference between carbon tax revenue and mitigation spending per unit of capital:

$$\tau_t^{reb} = \tau^{reb}(\mathbf{s}_t) = \tau^c(\mathbf{s}_t)\mathbf{e}(\mathbf{m}_t^*) - \mathbf{m}^*(\mathbf{s}_t).$$

With this rebate, a firm's net tax burden per unit of capital is exactly $\mathbf{m}^*(\mathbf{s}_t)$. Firms thus fully finance the aggregate mitigation cost, $\mathbf{M}_t = \mathbf{m}^*(\mathbf{s}_t)\mathbf{K}_t$. Crucially, this eliminates the distortion on capital accumulation while preserving the Pigouvian price signal. Notably, this policy mix—a carbon tax and a capital rebate—is equivalent to a direct capital tax set at the optimal mitigation rate, $\tau_t^k = \mathbf{m}^*(\mathbf{s}_t)$. This capital tax exactly finances aggregate abatement.[16]

Investment subsidy. Resolving the emission externality alone is insufficient because firms fail to internalize the benefits of aggregate capital investment. To align private incentives with the social optimum, the planner provides an investment subsidy $\tau_t^i = \tau^i(\mathbf{s}_t)$:

$$\tau^i(\mathbf{s}_t) = \frac{\epsilon_b(\mathbf{s}_t)}{\epsilon_b(\mathbf{s}_t) - 1}. \tag{30}$$

By anchoring the subsidy to the elasticity of the first-best welfare $b(\mathbf{s})$, the planner ensures that decentralized capital accumulation tracks the social optimum. Under this subsidy, the firm's investment FOC becomes:

$$q(\mathbf{s}) = \left(1 + \phi'(i)\right)\left(1 - \tau^i(\mathbf{s})\right). \tag{31}$$

[16] The firm's net payment to the planner is $\tau_t^c E_t - \tau_t^{reb} K_t$. Substituting $E_t = \mathbf{e}(\mathbf{m}_t^*)K_t$ and $\tau_t^{reb} = \tau_t^c \mathbf{e}(\mathbf{m}_t^*) - \mathbf{m}^*(\mathbf{s}_t)$, this simplifies to $\mathbf{m}^*(\mathbf{s}_t)K_t$. Thus, a simple capital tax $\tau_t^k = \mathbf{m}^*(\mathbf{s}_t)$ achieves the same allocation.

Here, $q(\mathbf{s})$ is the private marginal benefit of capital (Tobin's q), and the right-hand side of (31) is the subsidy-adjusted marginal cost of investment. By lowering the private cost of capital, the subsidy induces firms to internalize the social benefits of their investment, perfectly replicating the socially optimal allocation.

To understand the role of this subsidy, it is useful to compare our model with standard climate-macro frameworks (e.g., Golosov et al., 2014; Gerlagh and Liski, 2017). In canonical models, physical capital and fossil fuels are complements in production. Higher capital leads to more emissions and damages, making carbon pricing the primary policy tool. In contrast, capital in our framework acts as a buffer against climate risk. Because physical capital and emission reduction are substitutes in mitigating aggregate risk, carbon pricing alone is insufficient. Importantly, because the planner in our model is also present-biased, this subsidy is not designed to correct a behavioral tendency to under-save. Instead, it is a necessary climate policy instrument that corrects the capital accumulation externality, aligning decentralized investment with the planner's optimal path.

Finally, the lump-sum levy associated with the optimal investment subsidy is given by $\mathbf{L}_t = \tau^i(\mathbf{s}_t)\left(\mathbf{I}_t^* + \mathbf{\Phi}_t^*\right)$. Substituting $\mathbf{L}_t$ into the firm's dividend (29) yields:

$$Y_t = AK_t - (I_t + \Phi_t + M_t) - \mathbf{m}^*(\mathbf{s}_t)K_t + \tau^i(\mathbf{s}_t)\left[(I_t + \Phi_t) - (\mathbf{I}_t^* + \mathbf{\Phi}_t^*)\right]. \tag{32}$$

Since decentralized firms choose zero mitigation ($M_t = 0$), Equation (32) shows that a firm's net subsidy depends on how much its investment deviates from the first-best level. Because $\tau^i(\mathbf{s}_t) > 0$, this marginal scheme penalizes under-investment. It drives firms to accelerate capital accumulation, ensuring they fully internalize the benefits of aggregate capital.

In summary, the planner successfully decentralizes the first-best allocation through a policy mix. First, to manage emissions, the planner pairs a Pigouvian carbon tax ($\tau_t^c = SCC_t$) with a capital rebate (τ_t^{reb}). Together, they act as an effective capital tax ($\tau_t^k = \mathbf{m}^*(\mathbf{s}_t)$) that finances the socially optimal mitigation and solves the free-rider problem. Second, to manage capital accumulation, the planner provides an investment subsidy (τ_t^i) governed by (30). This subsidy internalizes the positive externality of climate risk dilution, thereby eliminating the wedge between private and social incentives to invest.

5 Quantitative Analysis

This section quantitatively assesses the impact of present-biased preferences on equilibrium dynamics and the design of optimal climate policy. We first calibrate the model's preference parameters to match the empirical term structure of discount rates, while production and climate parameters follow the macro-finance literature and integrated assessment models. We then compare the decentralized market equilibrium with the social planner's first-best allocation to highlight the distortions induced by externalities. Next, we quantify the optimal carbon tax and investment subsidy, examining their sensitivity to the degree of present bias. Finally, we evaluate the welfare value of commitment and show how it depends on the EIS and asset illiquidity.

5.1 Calibration and Parameter Choices

Table 1: Baseline Parameter Values. This table reports the parameter values used in the benchmark calibration. Reported percentages are annualized rates unless otherwise noted. The three panels correspond to household preferences, production and investment, and carbon dynamics and mitigation.

Parameters	**Symbol**	**Value**
Panel A: Household preferences		
Discount shock arrival rate	ξ	0.1
Long-term discount rate	ρ	2.8%
Present-bias factor	β	0.78
Relative risk aversion	γ	8
Elasticity of intertemporal substitution	ψ	1.4
Panel B: Production and investment		
Productivity	A	30%
Adjustment cost parameter	ϕ_0	15
Capital diffusion volatility	σ_K	8%
Depreciation rate of capital	δ_K	6.5%
Panel C: Carbon dynamics and mitigation		
Max. carbon-induced depreciation	δ_0	0.46
Damage curvature parameter	υ	50
Emission intensity	λ	0.1%
Mitigation efficiency	α_1	17.8
Mitigation return elasticity	α_2	0.357
Volatility of carbon stock growth	σ_S	7.5%
Carbon decaying rate	δ_S	3%
Correlation between capital and carbon shocks	ϑ	0

The model features 17 parameters, classified into three groups: household preferences, carbon dynamics, and production technology. Table 1 reports the baseline parameterization.

Parameters for household preferences. We calibrate the time-preference parameters to match the empirical term structure of discount rates. Weitzman (2001) shows that discount rates drop from 4% in the short term to 1% in the long term. To replicate this declining term structure, we set the present-bias factor β to 0.78, the hazard rate ξ to 0.1 (implying a 10-year expected horizon), and the long-term discount rate ρ to 2.8%. Finally, consistent with Bansal and Yaron (2004), we set the RRA γ to 8 and the EIS ψ to 1.4.

Parameters for carbon dynamics. Following Golosov et al. (2014), we specify the damage function $\mathcal{D}(\mathbf{s})$ as:

$$\mathcal{D}(\mathbf{s}) = \delta_0(1 - e^{-\upsilon \mathbf{s}}), \tag{33}$$

which implies that damages increase with the carbon-capital ratio at a diminishing rate ($\mathcal{D}'(\mathbf{s}) > 0$, $\mathcal{D}''(\mathbf{s}) < 0$). We set $\delta_0 = 0.46$, capping the maximum carbon-induced increase in capital depreciation at about 3%. We set the curvature parameter $\upsilon = 50$ to match empirical carbon price estimates. For the mitigation technology, we specify the aggregate carbon emissions $\mathbf{E}_t$ as:

$$\mathbf{E}_t = \lambda A \mathbf{K}_t \left[1 - \left(\alpha_1 \frac{\mathbf{M}_t}{\mathbf{K}_t}\right)^{\alpha_2}\right].$$

We set the emission intensity λ to 0.1%, mitigation efficiency $\alpha_1 = 17.8$, and return elasticity $\alpha_2 = 0.357$. Following Hong et al. (2023), we set the carbon growth volatility $\sigma_S = 7.5\%$, the decay rate $\delta_S = 3\%$, and the correlation between capital and carbon shocks $\vartheta = 0$.

Parameters for production technology. Following Hayashi (1982), we use a quadratic capital adjustment cost, $\phi(i) = \frac{\phi_0}{2} i^2$. We set productivity $A = 30\%$ and the adjustment cost parameter $\phi_0 = 15$, standard in the asset pricing literature. Following Hong et al. (2023), we set capital diffusion volatility $\sigma_K = 8\%$ and the depreciation rate $\delta_K = 6.5\%$.

5.2 First-Best vs. Laissez-Faire Equilibrium

We first compare the first-best solution with the competitive equilibrium. As shown in Panel A of Figure 1, the first-best mitigation spending (solid line) follows an inverted U-shape. In contrast, due to the public-good nature of mitigation, private firms free-ride, resulting in zero aggregate abatement in the competitive market (dashed line). Consequently, the carbon growth

rate, $g_{\mathbf{S}}(\mathbf{s}) \equiv \mathbf{e}(\mathbf{m})/\mathbf{s} - \delta_S$, is consistently higher in the competitive equilibrium than in the first-best economy (Panel B of Figure 1). Furthermore, since the damage to the economic growth (33) depends solely on the carbon-capital ratio $\mathbf{s}$, the damage function is identical across both economies, as illustrated in Panel C of Figure 1. However, faster carbon accumulation in the laissez-faire market leads to a higher steady-state $\mathbf{s}^{\text{steady}}$ and, consequently, greater economic damage.[17] Specifically, steady-state damage to output growth (dots) reaches 0.59% in the market economy, compared to only 0.37% under the planner.

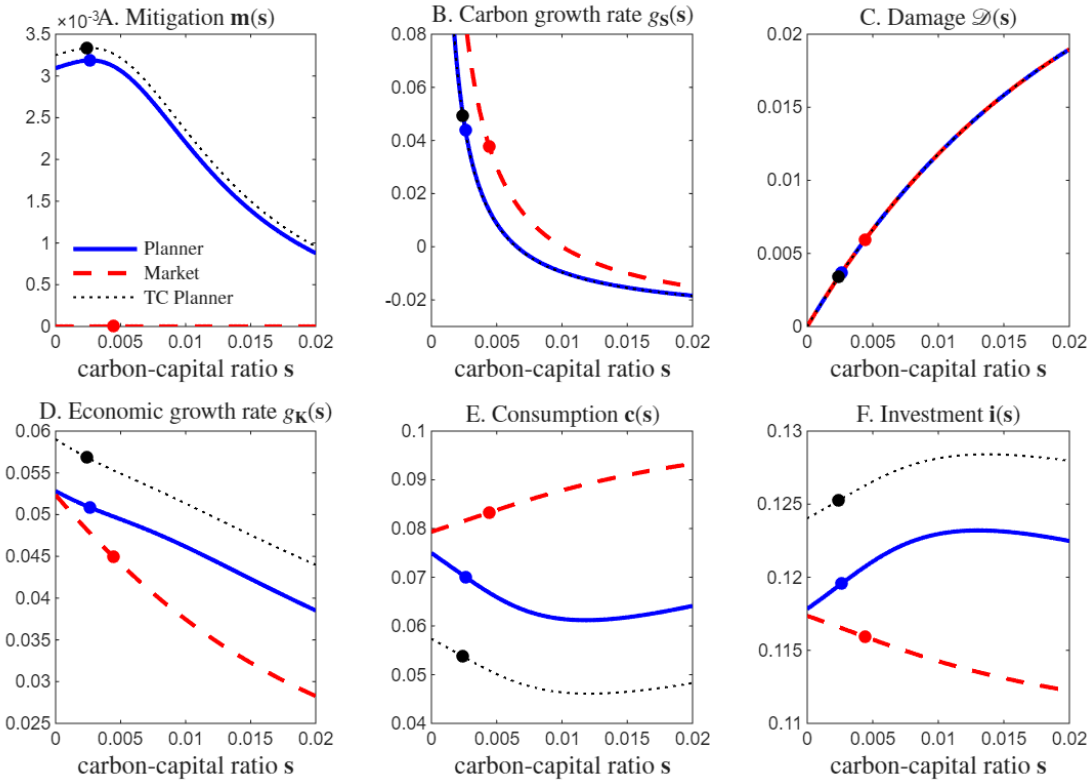


Figure 1: Resource allocation. Optimal allocations across three economies: (i) the economy with a present-biased planner (solid lines), (ii) the decentralized market with present-biased households (dashed lines), and (iii) the economy with a time-consistent planner (dotted lines). Dots denote steady states where the carbon-capital drift satisfies $\mu_{\mathbf{s}}(\mathbf{s}^{\text{steady}}) = 0$.

Panel D of Figure 1 plots the economic growth rate, $g_{\mathbf{K}}(\mathbf{s}) \equiv \mathbf{i}(\mathbf{s}) - \delta_K - \mathcal{D}(\mathbf{s})$. Given the same damage function, the disparity in the growth rates between the planner's economy and the market economy is driven entirely by the differences in capital investment $\mathbf{i}(\mathbf{s})$. As shown in Panel E of Figure 1, the first-best investment always exceeds market investment. Two reinforcing forces drive

[17] Accounting for the volatility adjustments in (4), the steady state features balanced carbon and capital growth. Thus, steady-state carbon growth rates (Panel B) are positive because physical capital is growing (Panel D).

this gap: (i) the benefits of capital accumulation on climate risk dilution discussed in (31), and (ii) the positive feedback from mitigation: active mitigation slows carbon growth and reduces future damages, which in turn raises the return on capital. Consequently, the planner invests more and consumes less than the laissez-faire market.[18]

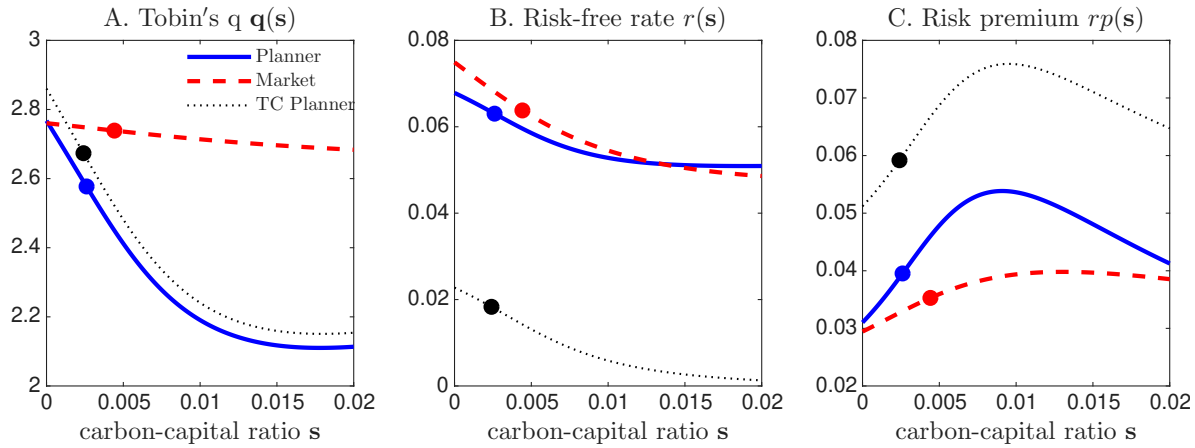


Figure 2: Asset prices. Tobin's q, the risk-free rate, and the risk premium across the three economies described in Figure 1. Dots indicate steady states.

Figure 2 shows that Tobin's q and the risk-free rate decrease with the carbon-capital ratio $\mathbf{s}$ in both economies, whereas the risk premium is non-monotonic. Intuitively, a higher $\mathbf{s}$ increases climate damages and reduces investment returns, causing the marginal value of capital as measured by Tobin's average q to decline. To hedge this risk, households reallocate wealth from risky equity to risk-free bonds, driving down the equilibrium interest rate.

Figures 1 and 2 further highlight the impact of present bias by comparing the baseline planner with a time-consistent counterfactual ($\xi = 0$ or $\beta = 1$, dotted lines). Panels A, E, and F of Figure 1 show that under time-inconsistent preferences, the planner focuses in the short run, and thus, consumes more and invests less in both physical capital and mitigation compared to a counterfactual planner. Consequently, the economy with present-biased preferences suffers from slower economic growth, $g_{\mathbf{K}}(\mathbf{s})$, and a higher steady-state carbon-capital ratio (Panel D of Figure 1). Specifically, steady-state damage drops to 0.34% under time consistency (Panel C of Figure 1), representing an 8% reduction compared to the present-biased baseline.[19]

[18] Since $\mu_{\mathbf{s}}(\mathbf{s}) = g_{\mathbf{s}}(\mathbf{s}) - g_{\mathbf{K}}(\mathbf{s}) + \sigma_K^2 - \vartheta\sigma_S\sigma_K$, the growth rate of the carbon-capital ratio is always lower in the planner's economy. As a result, for $\mathbf{s}$ lower than the steady-state level, $\mathbf{s}$ increases slower in the first-best solution; for $\mathbf{s}$ higher than the steady-state level, $\mathbf{s}$ declines faster in the planner's economy.

[19] Faster expected capital growth in the time-consistent economy leads to a higher Tobin's q (Panel A of Figure 2). Furthermore, because the time-consistent planner prioritizes long-term investment over current consumption, the economy borrows less, which lowers the risk-free rate. The same logic applies to the market economy; to avoid cluttered figures, we omit the time-consistent market solution.

5.3 Optimal Climate Policies

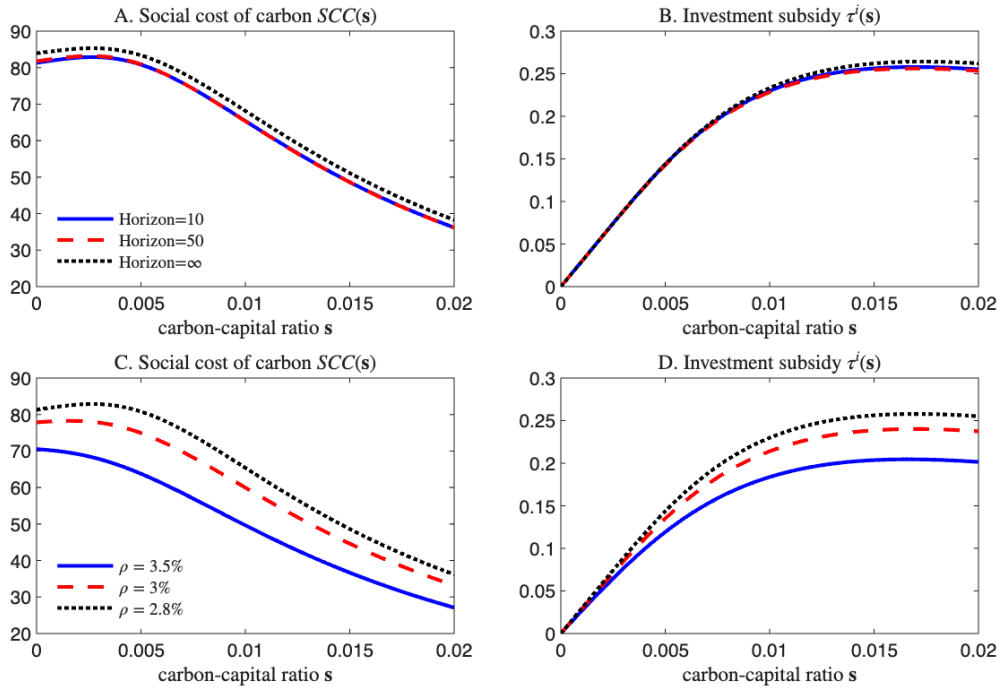


Figure 3: Optimal climate policies. The optimal carbon tax (left panels) and investment subsidy (right panels) across different household horizons $1/\xi$ (top panels) and long-term discount rates ρ (bottom panels).

We now turn to the optimal climate policies. As analyzed in Section 4.2, implementing the first-best allocation in a market economy requires both a carbon tax and an investment subsidy. Figure 3 plots the calibrated optimal policies. Panel A shows the optimal carbon tax, which equals the social cost of carbon (SCC). Near the steady state ($\mathbf{s} \approx 0.003$), the SCC is approximately 82 dollars per GtC. Beyond this point, the optimal carbon tax generally declines as the carbon-capital ratio rises. In our baseline, this decline (for $\mathbf{s} > 0.005$) is primarily driven by the diminishing marginal impact of carbon on capital depreciation under the concave damage function.[20]

Present bias and policy design. Perhaps surprisingly, the optimal carbon tax is largely insensitive to the household's planning horizon, $1/\xi$. To understand why, recall that $SCC(\mathbf{s}) = \frac{-V_{\mathbf{S}}(\mathbf{K},\mathbf{S})}{f_{\mathbf{C}}(\mathbf{C}^*,V)}$. Shortening the planning horizon, that is, strengthening present bias, generates two countervailing effects. First, it reduces the marginal disutility of carbon accumulation, $-V_{\mathbf{S}}(\mathbf{K},\mathbf{S})$, as future cli-

[20] It is worth noting that this hump-shaped pattern of the optimal carbon tax is robust to alternative damage specifications. As we will show in Section 6.3, the carbon tax also eventually declines under a convex damage function. In that setting, extreme climate damages severely depress consumption, causing the marginal utility of consumption to rise sharply, which ultimately lowers the SCC.

mate damages are discounted more heavily. Second, it lowers the marginal utility of consumption, $f_{\mathbf{C}}(\mathbf{C}^*, V)$, since households with shorter horizons tend to consume more today. These two forces push the SCC in opposite directions and quantitatively offset each other. This insensitivity aligns with Gerlagh and Liski (2017). Both studies confirm that quasi-hyperbolic discounting primarily distorts near-term capital accumulation, leaving the long-run social cost of carbon relatively unaffected.

This insensitivity extends to the optimal investment subsidy. Panel B of Figure 3 shows that the subsidy generally increases with the carbon-capital ratio, except at very high levels of $\mathbf{s}$. According to (30), the optimal subsidy is solely determined by the elasticity of the welfare with respect to the carbon-capital ratio $\epsilon_b(\mathbf{s}) = \frac{b'(\mathbf{s})\mathbf{s}}{b(\mathbf{s})}$. The two countervailing economic forces discussed above cancel each other out here as well: the effect of the planning horizon on the marginal value of carbon accumulation operates through $b'(\mathbf{s})$, while the current consumption channel goes through $b(\mathbf{s})$.[21]

This near-irrelevance result yields a crucial policy implication. Because the optimal SCC is robust to the exact profile of time preferences, policymakers can design carbon taxes without precisely calibrating the degree of present bias. In practice, estimating time-varying subjective discount rates is empirically challenging, and the strength of present bias varies across contexts. Our findings suggest that policymakers can economize on informational and implementation costs by computing the SCC as if households were time-consistent. Importantly, however, carbon taxation alone is insufficient. It must be paired with an investment subsidy to correct the capital accumulation externality, ensuring that private firms internalize the social value of capital in buffering climate risk.

Welfare value of climate policies. While present bias has a negligible impact on optimal policy levels, it significantly influences the welfare gains derived from these policies. Panel A of Figure 4 illustrates the welfare gains from jointly implementing the carbon tax and investment subsidy.[22] At low carbon-to-capital ratios, this optimal policy mix increases welfare by approximately 1.2%. However, present bias attenuates these gains relative to the time-consistent benchmark. This is because short-term-focused households heavily discount future climate damages, they perceive a lower benefit from long-term risk mitigation.

To decompose these aggregate gains, we evaluate the standalone welfare value of each policy.

[21] Varying the quasi-hyperbolic discount factor β similarly yields a minimal impact on optimal climate policies; the results are nearly identical to Figure 3. See Appendix C.1 for the analysis.

[22] The welfare gain is defined as $\frac{b(\mathbf{s})-u(\mathbf{s})\mathbf{q}(\mathbf{s})}{u(\mathbf{s})\mathbf{q}(\mathbf{s})}$, where $u(\mathbf{s})$ and $\mathbf{q}(\mathbf{s})$ are the certainty equivalent per dollar of wealth and the Tobin's q in the laissez-faire economy, respectively.

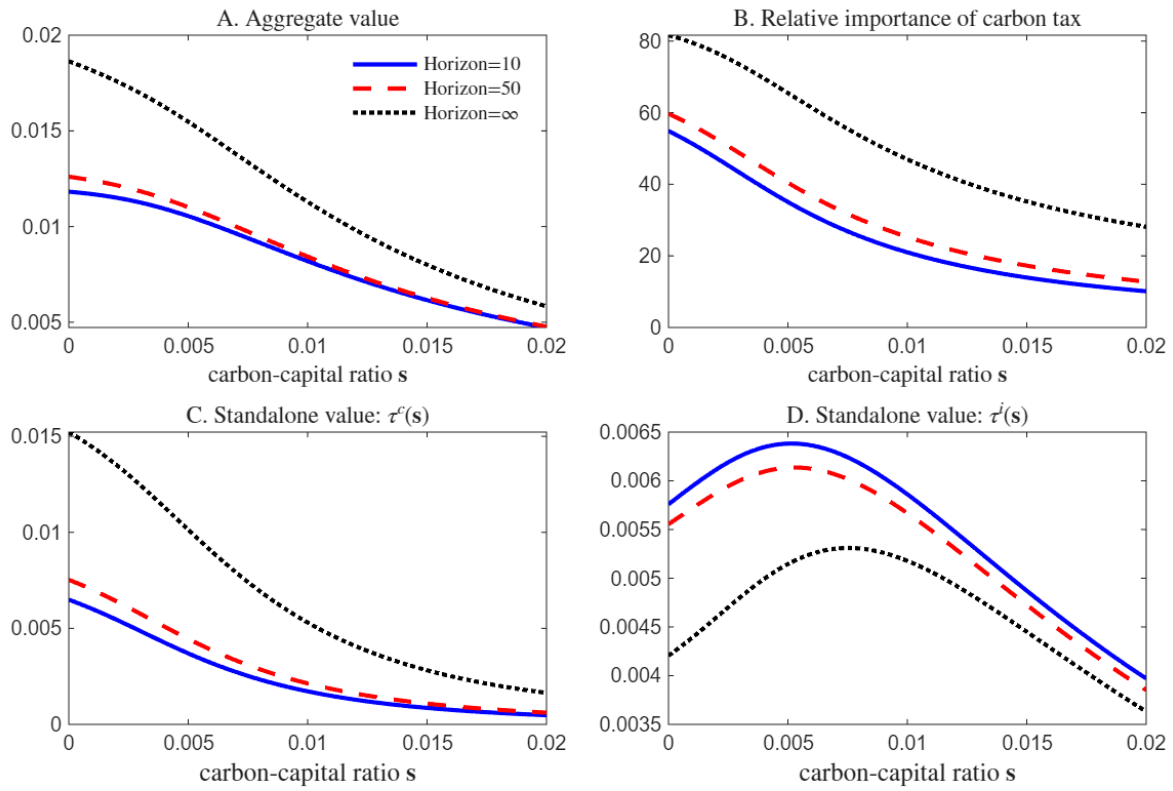


Figure 4: Welfare gains and policy decomposition. Welfare gains across different planning horizons $1/\xi$: aggregate gain from joint implementation (Panel A), relative importance of the carbon tax (Panel B), and standalone gains from the carbon tax (Panel C) and investment subsidy (Panel D).

Figure 4 demonstrates that both the carbon tax and the investment subsidy are crucial for improving welfare. For instance, near the steady state ($\mathbf{s} \approx 0.003$) under our baseline 10-year horizon, Panel B shows that the carbon tax accounts for roughly half of the total welfare gain. Specifically, the standalone gain from the optimal carbon tax is approximately 0.5% (Panel C), while the gain from the investment subsidy is about 0.6% (Panel D). However, the relative importance of these two policies varies significantly with the carbon-capital ratio and the degree of present bias. As Panel B illustrates, the relative contribution of the carbon tax declines as the carbon stock accumulates (higher $\mathbf{s}$) or as present bias becomes more severe (a shorter horizon).

These findings offer clear implications for the design and prioritization of climate policies. First, the carbon tax and the investment subsidy are indispensable complements. While the carbon tax directly addresses the emissions externality, the investment subsidy is essential to correct the capital accumulation externality. Relying on a single instrument leaves substantial welfare gains on the table. Second, our decomposition guides dynamic policy prioritization when policymakers face

institutional or political constraints that limit them to a single instrument. The optimal choice depends crucially on the state of the economy and the behavioral traits of households. When the carbon-capital ratio is low or households are relatively long-term focused, implementing a carbon tax should be the priority, as it delivers the majority of the standalone welfare gains. Conversely, as atmospheric carbon accumulates and climate risks intensify, or in economies characterized by severe present bias, the investment subsidy becomes the paramount instrument. In these high-risk or highly short-sighted environments, subsidizing capital accumulation to build a physical buffer against climate damages yields greater welfare gains than carbon pricing alone.

5.4 The Value of Commitment: The Role of EIS and Illiquidity

In this section, we explore what drives the value of commitment. We focus on how the household's EIS interacts with asset illiquidity. We proceed in two steps. First, we isolate the role of the EIS in determining whether extending the household's planning horizon improves welfare. Second, we examine how two distinct sources of illiquidity—capital illiquidity from adjustment costs (ϕ_0) and environmental illiquidity from persistent atmospheric carbon (δ_S)—reshape this commitment value.

The role of the EIS. We first analyze how the EIS shapes the value of commitment. Following Iverson and Karp (2021), we define the value of commitment, $\mathcal{V}(\mathbf{s})$, as the equivalent variation in aggregate capital that makes the household indifferent between the baseline economy—characterized by a short horizon of 10 years—and a counterfactual economy with an extended horizon.[23] Figure 5 shows that extending the horizon generates welfare gains in a high-EIS economy (Panel A) but reduces welfare in a low-EIS economy (Panel B). This divergence indicates that commitment is not universally beneficial; rather, its value depends on the fundamental tension between a *splurging motive* and a *smoothing motive*.

Consider first the case in which the EIS exceeds one. In this regime, households are highly willing to substitute consumption intertemporally, so the splurging motive dominates. In the decentralized market economy, present-biased households place excessive weight on immediate gratification, which leads to a severe under-accumulation of capital. Here, a commitment device acts as a disciplinary mechanism. By curbing short-sighted splurging motives, it sustains a higher equilibrium saving rate and promotes faster capital accumulation. As a result, the long-run welfare gains from a

[23] Because the household's value function J is homogeneous of degree $1-\gamma$ in wealth and $W = \mathbf{q}(\mathbf{s})\mathbf{K}$, this indifference condition simplifies directly to $\mathcal{V}(\mathbf{s}) = \frac{\hat{u}(\mathbf{s})\hat{\mathbf{q}}(\mathbf{s})}{u(\mathbf{s})\mathbf{q}(\mathbf{s})} - 1$, where the hat symbol ($\hat{}$) denotes the economy with an extended commitment horizon.

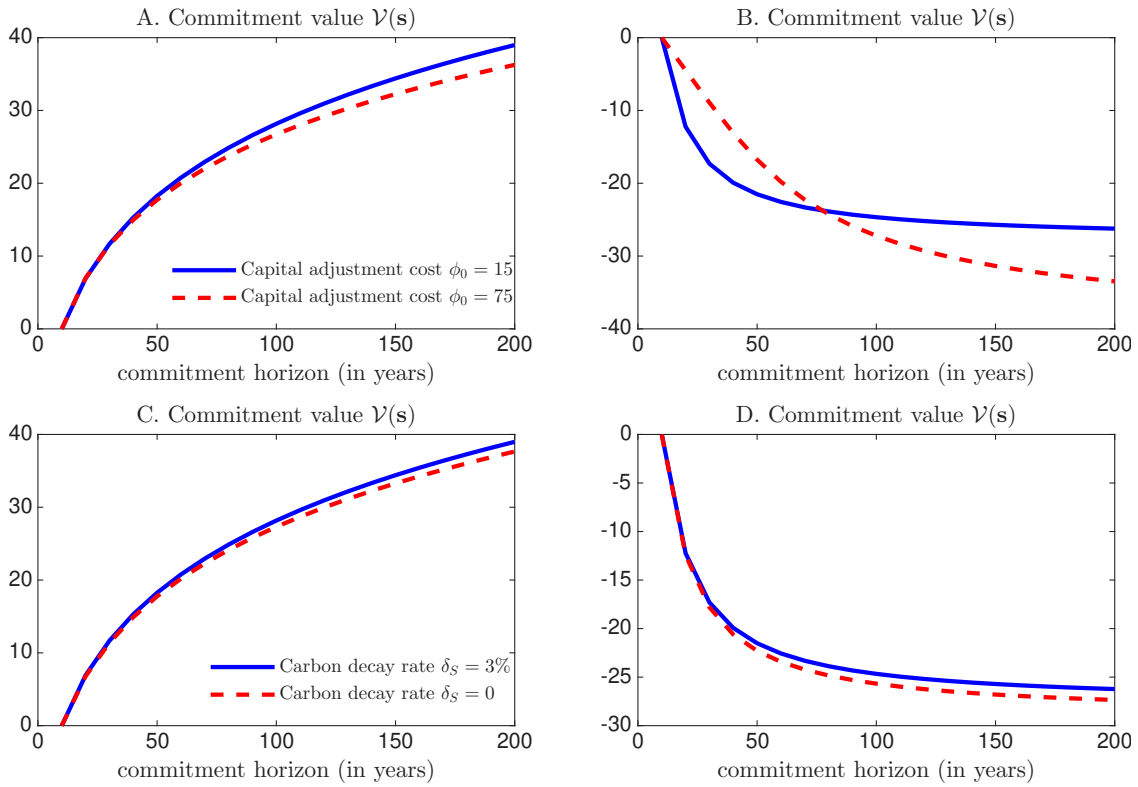


Figure 5: Value of commitment and asset illiquidity. Commitment value $\mathcal{V}(\mathbf{s})$ across commitment horizons under high ($\psi = 1.4$, left panels) and low ($\psi = 0.7$, right panels) EIS. Top and bottom panels illustrate the impacts of capital (ϕ_0) and environmental (δ_S) illiquidity, respectively. Solid and dashed lines denote baseline and heightened illiquidity.

larger capital buffer more than offset the short-run utility costs of reduced current consumption. Commitment thus improves household welfare, $\mathcal{V}(\mathbf{s}) > 0$.

Consider next the case in which the EIS is below one. In this regime, households are highly averse to consumption fluctuations, making the smoothing motive dominant. As a result, the flexibility to adjust saving and consumption over time plays a crucial role in determining household welfare. Extending the commitment horizon, however, imposes a rigid saving path, restricting the household's ability to respond to shocks and smooth consumption efficiently. Although commitment still promotes capital accumulation, the welfare cost of losing flexibility outweighs the long-run benefits of a larger capital stock. Consequently, commitment lowers welfare, $\mathcal{V}(\mathbf{s}) < 0$.

The role of illiquidity. Having established how the EIS governs the desirability of commitment, we now explore how asset illiquidity shapes these welfare effects. Within our framework, illiquidity arises along two dimensions: physical capital illiquidity, driven by capital adjustment costs (ϕ_0), and environmental illiquidity, driven by the natural decay rate of atmospheric carbon (δ_S).

Economically, a higher ϕ_0 makes capital accumulation more costly, whereas a lower δ_S slows the dissipation of atmospheric carbon. Both frictions increase the degree of illiquidity in the economy and thereby reshape the intertemporal trade-off faced by present-biased households.

We first analyze capital illiquidity (Panels A and B of Figure 5). Panel A shows that in a high-EIS regime—where the splurging motive dominates—the value of commitment strictly declines as the adjustment cost ϕ_0 increases. Intuitively, high adjustment costs restrict the household's ability to liquidate installed capital to finance immediate gratification. Consequently, physical capital illiquidity acts as an *endogenous substitute* for external commitment. Because this friction already curtails the incentive to splurge, the marginal welfare gain from extending the commitment horizon is substantially diminished. This mechanism echoes Acharya et al. (2026), who show that when borrowing constraints bind, asset illiquidity serves as a commitment technology that ties the hands of present-biased households and mitigates overconsumption.

Conversely, in a low-EIS regime, physical capital illiquidity exerts an asymmetric impact on the welfare costs of commitment. As depicted in Panel B, while higher adjustment costs (ϕ_0) mitigate the welfare losses associated with short-horizon commitments, they amplify the losses under extended horizons. To understand this, recall that an EIS below one implies a strong aversion to consumption volatility, making the smoothing motive dominant. Under a short commitment horizon, the commitment device forces present-biased households to rapidly increase their saving rates. Here, a higher ϕ_0 raises the marginal cost of installing capital, mechanically slowing the pace of accumulation. This friction prevents a sharp drop in short-term consumption, effectively smoothing the trajectory. Thus, capital illiquidity acts as a short-run buffer that attenuates the immediate welfare loss induced by forced savings. Under a long commitment horizon, however, the commitment device dictates a sustained, elevated level of capital. In this regime, a higher ϕ_0 forces the household to incur substantial adjustment costs to reach and maintain this capital stock, which severely crowds out consumption and exacerbates the welfare loss.

Next, we turn to the impact of environmental illiquidity (Panels C and D of Figure 5). In contrast to the dual role of capital illiquidity, a higher degree of environmental illiquidity—captured by a lower carbon decay rate (δ_S)—monotonically reduces the value of commitment across all EIS regimes. The economic intuition hinges on the interaction between carbon accumulation and asset returns. A lower δ_S slows the decay of atmospheric carbon, leading to a persistently higher aggregate carbon stock. This elevated carbon concentration exacerbates climate damages by accelerating capital depreciation, which reduces the marginal return on investment (Tobin's q). In

this environment, a commitment device that forces households to forgo current consumption in favor of future capital accumulation becomes economically inefficient. Consequently, environmental illiquidity reduces the welfare gains of commitment in the high-EIS regime (Panel C) and amplifies the welfare losses in the low-EIS regime (Panel D).

6 Discussion

6.1 Illiquidity as a Commitment Device Revisited

The interplay between asset illiquidity and the demand for commitment has been widely debated in the behavioral macro-finance literature (see, e.g., Laibson, 1997; Angeletos et al., 2001; Beshears et al., 2025). Within this literature, two recent studies provide useful benchmarks for our analysis. Maxted (2025) highlights a "present-bias dilemma": as long as households hold liquid wealth and borrowing constraints are slack, the fungibility of wealth renders illiquid assets ineffective as commitment devices. Building on this, Acharya et al. (2026) show that this ineffectiveness result breaks down only when borrowing constraints are strictly binding (e.g., for hand-to-mouth consumers). In such rigid environments, illiquidity ties households' hands, curbing overconsumption and restoring its commitment role.

Our general equilibrium framework bridges the gap between these two benchmark settings. Rather than relying on a fully unconstrained environment in Maxted (2025) or rigid, absorbing constraints emphasized by Acharya et al. (2026), our model features an endogenous soft constraint. Specifically, the aggregate resource constraint, together with convex capital adjustment costs (ϕ_0), makes capital liquidation costly without imposing strict absorbing boundaries, providing a more realistic characterization of macroeconomic capital dynamics. Crucially, we show that whether physical capital illiquidity acts as a commitment device depends entirely on the EIS. In a high-EIS regime, it serves as an endogenous commitment device that substitutes for external policy. In a low-EIS regime, however, where forced saving reduces welfare, capital illiquidity merely distorts consumption smoothing and lowers welfare.

More importantly, our model incorporates a novel dimension: environmental illiquidity. While heightened environmental illiquidity (captured by a lower carbon decay rate, δ_S) also reduces the value of commitment, it does not act as a commitment device. The economic distinction lies in the transmission mechanism. The distinction lies in the economic transmission mechanism. Physical capital illiquidity restricts households' ability to overconsume today, thereby forcing long-

run savings. In stark contrast, environmental illiquidity operates by lowering the return on capital. A lower carbon decay rate does not constrain the household's current consumption; instead, it leads to a higher aggregate carbon stock, which accelerates capital depreciation and reduces the marginal return on investment. In this situation, forcing present-biased households to save through a commitment device becomes inefficient. Thus, environmental illiquidity is not a commitment tool, but rather a fundamental friction that destroys the incentive to save for the future.

6.2 Temporal Resolution of Risk and Policy Distortion under Present Bias

Under recursive preferences, the temporal resolution of risk is a central issue for welfare analysis and policy design. Crucially, recursive utility disentangles relative risk aversion (γ) from the elasticity of intertemporal substitution (ψ), making the timing of uncertainty resolution relevant for decision-making. In this framework, γ governs the aversion to risk across states, while $1/\psi$ reflects the aversion to consumption fluctuations over time. A substantial literature has investigated how this property influences economic outcomes in the presence of long-run uncertainty. Specifically, Bansal and Yaron (2004) demonstrate that when $\gamma > 1/\psi$, agents exhibit a preference for the early resolution of uncertainty, making them highly sensitive to long-run growth risks. To quantify this effect, Epstein et al. (2014) define the "timing premium" as the fraction of consumption an individual is willing to give up to resolve long-run uncertainty early.[24]

The temporal resolution of risk is particularly critical for determining optimal climate policy. As Cai and Lontzek (2019) demonstrate, the separation of γ and $1/\psi$ significantly increases the SCC, as recursive preferences allow for a substantial climate risk premium that is absent under time-additive CRRA utility. Present bias, however, introduces a countervailing force. By prioritizing immediate gratification, a present-biased planner heavily discounts future climate damages. This short-term focus attenuates the precautionary saving motive driven by the preference for early resolution. In this section, we investigate how present bias reshapes the planner's attitude toward the early resolution of risk, and how this interaction depends on the magnitude of the EIS.

To quantify the impact of present bias on climate policy, we define the relative distortion in the SCC as:

$$\Delta_{SCC}(\mathbf{s}) = \frac{SCC^{TC}(\mathbf{s}) - SCC^{PB}(\mathbf{s})}{SCC^{PB}(\mathbf{s})}, \tag{34}$$

[24]The magnitude of the timing premium is closely linked to the persistence of the underlying risk process. Epstein et al. (2014) estimate a timing premium of 25–30% for the long-run risks model of Bansal and Yaron (2004). For rare disasters, Barro (2009) finds a premium of approximately 20% under i.i.d. shocks, whereas Wachter (2013) demonstrates that introducing time-varying disaster intensity significantly inflates the premium to roughly 40%.

where $SCC^{TC}(\mathbf{s})$ and $SCC^{PB}(\mathbf{s})$ represent the optimal carbon tax in the time-consistent and present-biased economies, respectively. This measure captures the extent to which present bias causes the planner to underestimate the SCC relative to the time-consistent benchmark.

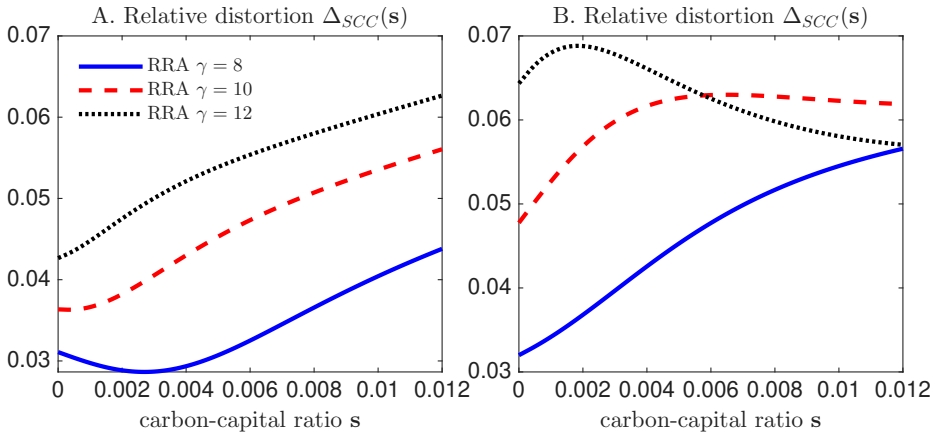


Figure 6: Relative distortion in the optimal carbon tax. Relative distortion for high EIS ($\psi = 1.4$, Panel A) and low EIS ($\psi = 0.7$, Panel B) across different levels of RRA γ. Solid, dashed, and dotted lines denote $\gamma = 8$, 10, and 12, respectively.

Figure 6 illustrates that the impact of risk aversion on carbon policy distortion hinges critically on the EIS. In a high-EIS regime (Panel A), the relative distortion increases monotonically with the RRA γ. This result confirms that present bias counteracts the preference for the early resolution of uncertainty. Intuitively, a high EIS implies a strong willingness to substitute consumption over time. A time-consistent planner then uses this flexibility to increase mitigation today in order to reduce long-run risks to economic growth. In contrast, a present-biased planner prioritizes immediate gratification and ignores these distant threats. As risk aversion rises, the time-consistent planner increases mitigation efforts significantly, while the present-biased planner remains short-term focused. This widening gap between their policies leads to a larger relative distortion.

In the low-EIS regime (Panel B), the relationship between risk aversion and policy distortion becomes non-monotonic and depends on the level of climate risk (s). When s is low, climate damages remain a distant threat. Here, the mechanism mirrors the high-EIS regime: higher risk aversion prompts the time-consistent planner to act earlier, while present bias causes the planner to ignore the threat, widening the policy gap. However, as s increases, climate risk becomes imminent. A low EIS implies a strong desire to smooth consumption over time. Because a climate disaster would trigger a sharp drop in future consumption, the fear of such fluctuations becomes the dominant force. This smoothing motive compels even a present-biased planner to increase mitigation to prevent a large decline in future utility. Consequently, the gap between the time-consistent and

present-biased policies narrows, and the relative distortion decreases.

In summary, while present bias generally suppresses the preference for early resolution, a strong consumption-smoothing motive can counteract this behavioral bias when climate risks become imminent.[25]

6.3 Robustness to Convex Damage Specifications

To assess the robustness of our baseline results, we consider an alternative specification where the climate damage function is convex. Specifically, we adopt the following functional form:

$$\mathcal{D}(\mathbf{s}) = \frac{\delta_1 \mathbf{s}^2}{1 + \delta_2 \mathbf{s}^2}, \tag{35}$$

where δ_1 governs the sensitivity of capital depreciation to carbon accumulation and δ_2 determines the curvature. This specification, motivated by the DICE models (Nordhaus, 2008, 2017), employs a quadratic formulation to capture accelerating economic losses.[26] Under the concave form (33), the initial unit of carbon emissions causes the largest marginal damage, and the incremental damage from further emissions gradually declines (Golosov et al., 2014). In stark contrast, the convex form (35) captures a tipping-point effect (Lemoine and Traeger, 2014).[27] Specifically, the climate system initially buffers carbon emissions with relatively low marginal damage. As carbon accumulates, however, marginal damages increase and become highly sensitive to additional emissions. This functional form therefore characterizes a system approaching critical carbon stock thresholds, beyond which further accumulation leads to sharp and potentially catastrophic capital destruction.

Replacing the concave damage function (33) with the convex specification (35) affects our results only quantitatively; the qualitative insights remain robust. We detail these quantitative differences across three dimensions, with corresponding figures provided in Appendix C.3. First, regarding optimal climate policies, both the optimal carbon tax (SCC) and the investment subsidy reach higher levels under the convex specification (see Figure 9). Under convex damages, the marginal damage to capital increases as the carbon stock grows. Consequently, the marginal benefit of

[25] In Appendix C.2, we extend our analysis to examine how the interplay between risk aversion and the EIS shapes the value of commitment. We specifically discuss how the tension between the splurging and smoothing motives determines the welfare implications of commitment devices across different preference regimes.

[26] Mathematically, for $\mathbf{s} < (3\delta_1)^{-1/2}$, we have $\mathcal{D}''(\mathbf{s}) > 0$, implying that the marginal climate damage is increasing in regions with a low-to-moderate carbon-capital ratio.

[27] Lemoine and Traeger (2014) incorporate stochastic tipping points into an integrated assessment model, demonstrating that the risk of crossing environmental thresholds significantly raises the optimal carbon tax even before the event occurs.

mitigation is higher because reducing emissions prevents the economy from entering a state of accelerating destruction, leading to a higher SCC. Similarly, since capital acts as a buffer against climate risk, its social value rises when climate risks accelerate. Thereby, the planner subsidizes investment more aggressively to build larger capital buffers.

Second, the welfare gains from implementing optimal climate policies are substantially larger under the convex specification (see Figure 10). In the decentralized market, firms fail to internalize the social benefits of emission mitigation and the social value of capital accumulation. This leads to a higher carbon-capital ratio compared to the planner's economy. When damages are convex, this excess carbon is much more harmful because capital depreciation accelerates rapidly as the climate state deteriorates. Therefore, policy intervention becomes more valuable, as it prevents the economy from falling into the high-damage regions of the convex curve.

Third, the convex specification increases the value of commitment across all preference regimes (see Figure 11). Because convex damages amplify future climate threats, the welfare losses caused by present-biased under-saving and under-mitigation become much more severe. A commitment device—which extends the household's planning horizon and forces higher investment and mitigation—delivers a larger benefit by curbing this accelerating capital depreciation. Consequently, the commitment value increases across different preference regimes: it becomes more positive in a high-EIS economy and less negative in a low-EIS economy.

7 Conclusion

We develop a continuous-time stochastic general equilibrium model to study climate risk mitigation in an economy with present-biased households. We model climate mitigation via an emission abatement technology that slows carbon accumulation, and we capture household preferences using quasi-hyperbolic recursive utility to disentangle present bias, risk aversion, and intertemporal substitution. The planner decentralizes the first-best allocation through a policy mix: a Pigouvian carbon tax anchored to the social cost of carbon corrects the emissions externality, while an investment subsidy realigns private capital accumulation. To balance the budget without distorting marginal decisions, this package is supplemented by a capital rebate and a lump-sum levy.

We calibrate household preference parameters to match the empirical term structure of discount rates, while parameters governing production and climate dynamics follow the macro-finance literature and integrated assessment models. Our quantitative analysis shows that optimal carbon

tax and investment subsidy rates are largely insensitive to the degree of present bias. This implies that policymakers can design robust climate policies using standard time-consistent frameworks without precisely calibrating behavioral biases. However, while optimal policy levels are robust, the welfare gains derived from these policies are highly sensitive to the household's commitment horizon. Decomposing these gains provides a guide for dynamic policy prioritization: the carbon tax contributes more when the carbon-capital ratio is low or when households are more long-term focused, whereas the relative importance of the investment subsidy increases as the carbon stock accumulates. Finally, we show that the value of commitment depends crucially on the elasticity of intertemporal substitution, which determines the trade-off between splurging and smoothing motives. In this environment, physical capital illiquidity (driven by adjustment costs) acts as an endogenous substitute for external commitment. In stark contrast, environmental illiquidity (driven by the persistence of atmospheric carbon) acts as a fundamental friction that depresses investment returns, thereby reducing the value of commitment across all preference regimes.

References

Acharya, Subas, David Jimenez-Gomez, Dmitrii Rachinskii, and Alejandro Rivera, "Present-Bias and the Value of Sophistication: Splurging vs. Smoothing," *Games and Economic Behavior*, 2026.

Angeletos, George-Marios, David Laibson, Andrea Repetto, Jeremy Tobacman, and Stephen Weinberg, "The hyperbolic consumption model: Calibration, simulation, and empirical evaluation," *Journal of Economic Perspectives*, 2001, *15* (3), 47–68.

Bansal, Ravi and Amir Yaron, "Risks for the long run: A potential resolution of asset pricing puzzles," *The Journal of Finance*, 2004, *59* (4), 1481–1509.

Barnett, Michael, William Brock, and Lars P. Hansen, "Pricing Uncertainty Induced by Climate Change," *Review of Financial Studies*, 2020, *33* (3), 1024–1066.

Barrage, Lint, "Optimal Dynamic Carbon Taxes in a Climate-Economy Model with Distortionary Fiscal Policy," *The Review of Economic Studies*, 2020, *87* (1), 1–39.

Barro, Robert J., "Rare Disasters and Asset Markets in the Twentieth Century," *The Quarterly Journal of Economics*, 2006, *121* (3), 823–866.

_ , "Rare Disasters, Asset Prices, and Welfare Costs," *American Economic Review*, 2009, *99* (1), 243–264.

Beshears, John, James J Choi, Christopher Clayton, Christopher Harris, David Laibson, and Brigitte C Madrian, "Optimal illiquidity," *Journal of Financial Economics*, 2025, *165*, 103996.

Cai, Yongyang and Thomas S. Lontzek, "The Social Cost of Carbon with Economic and Climate Risks," *Journal of Political Economy*, 2019, *127* (6), 2684–2734.

Dell, Melissa, Benjamin F. Jones, and Benjamin A. Olken, "Temperature and Income: Reconciling New Cross-Sectional and Panel Estimates," *American Economic Review*, 2009, *99* (2), 198–204.

_ , _ , and _ , "Temperature Shocks and Economic Growth: Evidence from the Last Half Century," *American Economic Journal: Macroeconomics*, 2012, *4* (3), 66–95.

den Bremer, Ton S Van and Frederick Van der Ploeg, "The risk-adjusted carbon price," *American Economic Review*, 2021, *111* (9), 2782–2810.

Duffie, Darrell, *Dynamic asset pricing theory*, Princeton University Press, 2010.

_ **and Larry G. Epstein**, "Stochastic Differential Utility," *Econometrica*, 1992, *60* (2), 353–394.

_ **and Yeneng Sun**, "The exact law of large numbers for independent random matching," *Journal of Economic Theory*, 2012, *147* (3), 1105–1139.

Epstein, Larry G. and Stanley E. Zin, "Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework," *Econometrica*, 1989, *57* (4), 937–969.

Epstein, Larry G, Emmanuel Farhi, and Tomasz Strzalecki, "How much would you pay to resolve long-run risk?," *American Economic Review*, 2014, *104* (9), 2680–2697.

Gerlagh, Reyer and Matti Liski, "Consistent Climate Policies," *Journal of European Economic Association*, 2017, *16* (1), 1–44.

Golosov, Mikhail, John Hassler, Per Krusell, and Aleh Tsyvinski, "Optimal Taxes on Fossil Fuel in General Equilibrium," *Econometrica*, 2014, *82* (1), 41–88.

Grenadier, Steven R. and Neng Wang, "Investment under Uncertainty and Time-Inconsistent Preferences," *Journal of Financial Economics*, 2007, *84* (1), 2–39.

Hambel, Christoph, Holger Kraft, and Eduardo Schwartz, "Optimal Carbon Abatement in a Stochastic Equilibrium Model," *European Economic Review*, 2021, *132*, 103642.

Harris, Christopher and David Laibson, "Instantaneous Gratification," *The Quarterly Journal of Economics*, 2013, *128* (1), 205–248.

Hassler, John, Per Krusell, and Conny Olovsson, "Presidential Address 2020: Suboptimal Climate Policy," *Journal of the European Economic Association*, 2021, *19* (6), 2895–2928.

Hayashi, Fumio, "Tobin's Marginal q and Average q: A Neoclassical Interpretation," *Econometrica*, 1982, *50* (1), 213–224.

Hong, Harrison, Neng Wang, and Jinqiang Yang, "Mitigating Disaster Risks in the Age of Climate Change," *Econometrica*, 2023, *91* (5), 1763–1802.

Iverson, Terrence and Larry Karp, "Carbon Taxes and Climate Commitment with Non-constant Time Preferences," *Review of Economic Studies*, 2021, *88* (2), 764–799.

Jermann, Urban J, "Asset pricing in production economies," *Journal of Monetary Economics*, 1998, *41* (2), 257–275.

Karp, Larry, "Global warming and hyperbolic discounting," *Journal of Public Economics*, 2005, *89* (2-3), 261–282.

_ and Yacov Tsur, "Time perspective and climate change policy," *Journal of Environmental Economics and Management*, 2011, *62* (1), 1–14.

Krusell, Per, Burhanettin Kuruşçu, and Anthony A. Smith, "Equilibrium Welfare and Government Policy with Quasi-geometric Discounting," *Journal of Economic Theory*, 2002, *105* (1), 42–72.

_ , _ , and _ , "Consumption-Savings Decisions with Quasi-Geometric Discounting," *Econometrica*, 2003, *71* (1), 365–375.

Laibson, David, "Golden Eggs and Hyperbolic Discounting," *The Quarterly Journal of Economics*, 1997, *112* (2), 443â478.

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

Maxted, Peter, "Present bias unconstrained: Consumption, welfare, and the present-bias dilemma," *The Quarterly Journal of Economics*, 2025, *140* (4), 2963–3013.

Nordhaus, William, *A question of balance: Weighing the options on global warming policies*, Yale University Press, 2008.

Nordhaus, William D., "Revisiting the Social Cost of Carbon," *Proceedings of the National Academy of Sciences*, 2017, *114* (7), 1518–1523.

O'donoghue, Ted and Matthew Rabin, "Doing it now or later," *American Economic Review*, 1999, *89* (1), 103–124.

_ and _ , "Choice and procrastination," *The Quarterly Journal of Economics*, 2001, *116* (1), 121–160.

Phelps, E. S. and R. A. Pollak, "On Second-Best National Saving and Game-Equilibrium Growth," *The Review of Economic Studies*, 1968, *35* (2), 185–199.

Pindyck, Robert S. and Neng Wang, "The Economic and Policy Consequences of Catastrophes," *American Economic Journal: Economic Policy*, 2013, *5* (4), 306–339.

Shigeta, Yuki, "Quasi-hyperbolic Discounting under Recursive Utility and Consumption-Investment Decisions," *Journal of Economic Theory*, 2022, *204*, 105518.

Stern, Nicholas, *The Economics of Climate Change: The Stern Review*, Cambridge, UK: Cambridge University Press, 2007.

Wachter, Jessica A., "Can Time-Varying Risk of Rare Disasters Explain Aggregate Stock Market Volatility?," *The Journal of Finance*, 2013, *68* (3), 987–1035.

Weitzman, Martin L., "Gamma Discounting," *American Economic Review*, 2001, *91* (1), 260–271.

Appendix

A Appendix for Section 3: The Market Economy

A.1 Household's Optimization Problem

Given the asset return process (10), the dynamics of the household's wealth are

$$\begin{aligned} dW_t = &\left(r_{t-}W_{t-} + (\mu_{\mathbf{Q},t-} - r_{t-})\Gamma_{t-} - C_{t-}\right) dt + \sigma^K_{\mathbf{Q},t-}\Gamma_{t-}d\mathcal{B}^K_t \\ &+ \sigma^S_{\mathbf{Q},t-}\Gamma_{t-}d\mathcal{B}^S_t + \left(\frac{\tilde{\mathbf{Q}}_t}{\mathbf{Q}_{t-}} - 1\right)\Gamma_{t-}d\mathcal{P}_t. \end{aligned} \tag{A.1}$$

The wealth dynamics (A.1) are driven by the risk-free return on saving, consumption, and the risky return on the aggregate stock portfolio. As implied by (10), the risky return contains two novel terms. First, stochastic fluctuations in the carbon stock drive changes in the aggregate equity value, and thus, exposes the risky portfolio to carbon risk, $\sigma^S_{\mathbf{Q},t-}\Gamma_{t-}d\mathcal{B}^S_t$. Second, the jump term, $\left(\frac{\tilde{\mathbf{Q}}_t}{\mathbf{Q}_{t-}} - 1\right)\Gamma_{t-}d\mathcal{P}_t$, captures the impact of present bias. Specifically, the arrival of a future self triggers a change in discount factor, causing an instantaneous revaluation of equity from $\mathbf{Q}_{t-}$ to $\tilde{\mathbf{Q}}_t$, thereby generating a discrete shock to household wealth.

By using the W_t process given in (A.1) and the $\mathbf{s}_t$ process given in (3), we obtain the HJB equation (11) for the current household's utility function. The FOCs for consumption (14) and the market portfolio allocation (15) can be obtained by maximizing the HJB equation (11) together with the conjecture (13).

To fully characterize the current household's optimization problem, we must also derive the household's continuation utility $\tilde{J}(\tilde{W}, \mathbf{s})$. Upon the arrival of a new self, the current household perceives the ex-dividend return on the aggregate stock market as

$$\frac{d\tilde{\mathbf{Q}}_t + \tilde{\mathbf{D}}_t dt}{\tilde{\mathbf{Q}}_t} = \tilde{\mu}_{\mathbf{Q},t}dt + \tilde{\sigma}^K_{\mathbf{Q},t}d\mathcal{B}^K_t + \tilde{\sigma}^S_{\mathbf{Q},t}d\mathcal{B}^S_t, \tag{A.2}$$

where the drift $\tilde{\mu}_{\mathbf{Q},t}$ and volatilities $\tilde{\sigma}^K_{\mathbf{Q},t}$ and $\tilde{\sigma}^S_{\mathbf{Q},t}$ are endogenous. Given this return process (A.2) and the future self's policies $(\tilde{C}_t, \tilde{\Gamma}_t)$, the current household's continuation wealth evolves as

$$d\tilde{W}_t = \left(\tilde{r}_t\tilde{W}_t + (\tilde{\mu}_{\mathbf{Q},t} - \tilde{r}_t)\tilde{\Gamma}_t - \tilde{C}_t\right) dt + \tilde{\sigma}^K_{\mathbf{Q},t}\tilde{\Gamma}_t d\mathcal{B}^K_t + \tilde{\sigma}^S_{\mathbf{Q},t}\tilde{\Gamma}_t d\mathcal{B}^S_t. \tag{A.3}$$

Using the dynamics of $\tilde{W}_t$ in (A.3), the continuation utility satisfies the equation (12).

A.2 Firm's Optimization Problem

By applying Ito's lemma, we obtain the following dynamics for the current firm's value

$$\begin{aligned} dQ_t =& \Big(\left(I_{t-} - \delta_K K_{t-} - \mathcal{D}(\mathbf{s}_{t-})K_{t-}\right) dt + \sigma_K K_{t-} d\mathcal{B}_t^K \Big) Q_K + \frac{\sigma_K^2 K_{t-}^2}{2} Q_{KK} dt \\ &+ \Big(\mu_{\mathbf{s}}(\mathbf{s}_{t-}; \mathbf{i}_{t-}, \mathbf{m}_{t-})\mathbf{s}_{t-} dt + (\sigma_S d\mathcal{B}_t^S - \sigma_K d\mathcal{B}_t^K)\mathbf{s}_{t-} \Big) Q_{\mathbf{s}} + \frac{\Sigma_{\mathbf{s}} \mathbf{s}_{t-}^2}{2} Q_{\mathbf{ss}} dt \\ &+ \left(\vartheta \sigma_K \sigma_S - \sigma_K^2\right) K_{t-} \mathbf{s}_{t-} Q_{K\mathbf{s}} dt + \left(\frac{\tilde{Q}_t}{Q_{t-}} - 1 \right) Q_{t-} d\mathcal{P}_t, \end{aligned} \tag{A.4}$$

where $\tilde{Q}_t = \tilde{Q}(K_t, \mathbf{s}_t)$ is the continuation value of the current firm. No arbitrage implies that the drift of the SDF-discounted process, $\mathbb{M}_t \left(AK_t - I_t - \Phi(I_t, K_t) - M_t\right) dt + d(\mathbb{M}_t Q_t)$, is zero under the physical measure:

$$0 = \max_{i_t, m_t} \frac{\left(A - i_t - \phi(i_t) - m_t\right) K_t}{Q(K_t, \mathbf{s}_t)} + \mathbb{E}\left[\frac{d\mathbb{M}_t}{\mathbb{M}_t}\right] \frac{1}{dt} + \mathbb{E}\left[\frac{dQ_t}{Q_t}\right] \frac{1}{dt} + \mathbb{E}\left[\frac{d\left\langle \mathbb{M}_t, Q_t \right\rangle}{\mathbb{M}_t Q_t}\right] \frac{1}{dt}. \tag{A.5}$$

Conjecture that $Q(K_t, \mathbf{s}_t) = q(\mathbf{s}_t) K_t$, then by (A.4), we obtain

$$\begin{aligned} \mathbb{E}\left[\frac{dQ_t}{Q_t}\right] \frac{1}{dt} =& i_t - \delta_K - \mathcal{D}(\mathbf{s}_t) + \left(\frac{\mathbf{e}(\mathbf{m}_t)}{\mathbf{s}_t} - \delta_S - \mathbf{i}_t + \delta_K + \mathcal{D}(\mathbf{s}_t) \right) \frac{q'(\mathbf{s}_t)\mathbf{s}_t}{q(\mathbf{s}_t)} \\ &+ \frac{\Sigma_{\mathbf{s}} \mathbf{s}_t^2}{2} \frac{q''(\mathbf{s}_t)}{q(\mathbf{s}_t)} + \xi \left(\frac{\tilde{q}(\mathbf{s}_t)}{q(\mathbf{s}_t)} - 1 \right). \end{aligned} \tag{A.6}$$

Moreover, the expected quadratic covariation between the SDF and the firm value is exactly the negative of the equilibrium risk premium defined in Equation (22):

$$\mathbb{E}\left[\frac{d\langle \mathbb{M}_t, Q_t \rangle}{\mathbb{M}_t Q_t}\right] \frac{1}{dt} = -rp(\mathbf{s}_t). \tag{A.7}$$

From the SDF dynamics (16), we have $\mathbb{E}\left[\frac{d\mathbb{M}_t}{\mathbb{M}_t}\right] \frac{1}{dt} = -r(\mathbf{s}_t)$. Substituting this, along with (A.6) and (A.7), into the no-arbitrage condition (A.5) yields the firm's HJB equation (17).

A.3 Recursive Competitive Equilibrium

We characterize the recursive competitive equilibrium in five steps.

1. Derive the aggregate stock return dynamics ($\mu_{\mathbf{Q}}$, $\sigma_{\mathbf{Q}}^K$, $\sigma_{\mathbf{Q}}^S$), which determine the endogenous wealth volatilities (Σ_W, $\Sigma_{W\mathbf{s}}$). These results identify the return dynamics ($\tilde{\mu}_{\mathbf{Q}}$, $\tilde{\sigma}_{\mathbf{Q}}^K$, $\tilde{\sigma}_{\mathbf{Q}}^S$) and wealth volatilities ($\Sigma_{\tilde{W}}$, $\Sigma_{\tilde{W}\mathbf{s}}$) in the continuation game.

2. Solve for the household's value functions, characterized by the current and continuation CEW, $u(\mathbf{s})$ and $\tilde{u}(\mathbf{s})$.

3. Determine the equilibrium SDF by deriving the market prices of risk ($\eta_{\mathbb{M}}^K$, $\eta_{\mathbb{M}}^S$) and the discount shock price (χ) from the household's optimal utility functions.

4. Characterize the aggregate firm value $\mathbf{q}(\mathbf{s})$, the risk-free rate $r(\mathbf{s})$, and the risk premium $rp(\mathbf{s})$, taking the equilibrium SDF as given.

5. Determine the continuation firm value $\tilde{\mathbf{q}}(\mathbf{s})$ using the continuation prices of risk $(\tilde{\eta}^K_{\mathbb{M}}, \tilde{\eta}^S_{\mathbb{M}})$ and the continuation risk-free rate $\tilde{r}(\mathbf{s})$ implied by Steps 3 and 4.

For ease of exposition, we restate the elasticity notations for a generic function $g(\mathbf{s}_t)$ used in the main text:

$$\epsilon_g(\mathbf{s}_t) \equiv \frac{g'(\mathbf{s}_t)\mathbf{s}_t}{g(\mathbf{s}_t)} \text{ and } \epsilon'_g(\mathbf{s}_t) \equiv \frac{g''(\mathbf{s}_t)\mathbf{s}_t^2}{g(\mathbf{s}_t)}.$$

Step 1. Equilibrium return on the aggregate stock market. In the competitive equilibrium, market clearing implies $\mathbf{K}_t = K_t$ and $\mathbf{Q}_t = Q_t$. Hence, the aggregate firm value is $\mathbf{Q}(\mathbf{K}_t, \mathbf{s}_t) = \mathbf{q}(\mathbf{s}_t)\mathbf{K}_t$, and its continuation value is $\tilde{\mathbf{Q}}(\mathbf{K}_t, \mathbf{s}_t) = \tilde{\mathbf{q}}(\mathbf{s}_t)\mathbf{K}_t$. Then, Equation (A.4) becomes

$$\begin{aligned}\frac{d\mathbf{Q}_t}{\mathbf{Q}_t} =& \Big(\mathbf{i}_t - \delta_K - \mathcal{D}(\mathbf{s}_t) + \big(\mu_{\mathbf{s}}(\mathbf{s}_t) - (\sigma_K^2 - \vartheta\sigma_K\sigma_S)\big)\epsilon_{\mathbf{q}}(\mathbf{s}_t)\Big)dt \\ &+ \frac{\Sigma_{\mathbf{s}}}{2}\epsilon'_{\mathbf{q}}(\mathbf{s}_t)dt + \sigma_K\big(1 - \epsilon_{\mathbf{q}}(\mathbf{s}_t)\big)d\mathcal{B}_t^K + \sigma_S\epsilon_{\mathbf{q}}(\mathbf{s}_t)d\mathcal{B}_t^S + \Big(\frac{\tilde{\mathbf{q}}(\mathbf{s}_t)}{\mathbf{q}(\mathbf{s}_t)} - 1\Big)d\mathcal{P}_t.\end{aligned} \tag{A.8}$$

Matching the drift and diffusion terms in (A.8) with the return process (10), we obtain the expected cum-dividend return

$$\mu_{\mathbf{Q}}(\mathbf{s}_t) = \mathbf{i}_t - \delta_K - \mathcal{D}(\mathbf{s}_t) + \Big(\mu_{\mathbf{s}}(\mathbf{s}_t) - (\sigma_K^2 - \vartheta\sigma_K\sigma_S)\Big)\epsilon_{\mathbf{q}}(\mathbf{s}_t) + \frac{\Sigma_{\mathbf{s}}}{2}\epsilon'_{\mathbf{q}}(\mathbf{s}_t) + \xi\Big(\frac{\tilde{\mathbf{q}}(\mathbf{s}_t)}{\mathbf{q}(\mathbf{s}_t)} - 1\Big) + \frac{\mathbf{y}_t}{\mathbf{q}(\mathbf{s}_t)}, \tag{A.9}$$

where $\mathbf{y}_t = A - \mathbf{i}_t - \phi(\mathbf{i}_t) - \mathbf{m}_t$ is the aggregate dividend per unit of capital. The aggregate stock return volatilities are:

$$\sigma^K_{\mathbf{Q}}(\mathbf{s}_t) = \sigma_K\big(1 - \epsilon_{\mathbf{q}}(\mathbf{s}_t)\big), \ \sigma^S_{\mathbf{Q}}(\mathbf{s}_t) = \sigma_S\epsilon_{\mathbf{q}}(\mathbf{s}_t). \tag{A.10}$$

Because investment and mitigation are functions of $\mathbf{s}_t$ in the recursive equilibrium, expressions (A.9) and (A.10) then relate Σ_W and $\Sigma_{W\mathbf{s}}$ to $\mathbf{s}_t$ through the equilibrium firm value and policies. Specifically, we have

$$\Sigma_W(\mathbf{s}_t) \equiv \frac{\text{var}(dW_t)}{dt} = \sigma_K^2 - 2(\sigma_K^2 - \vartheta\sigma_S\sigma_K)\epsilon_{\mathbf{q}}(\mathbf{s}_t) + \Sigma_{\mathbf{s}}\epsilon_{\mathbf{q}}(\mathbf{s}_t)^2, \tag{A.11}$$

$$\Sigma_{W\mathbf{s}}(\mathbf{s}_t) \equiv \frac{\text{cov}(dW_t \cdot d\mathbf{s}_t)}{dt} = -(\sigma_K^2 - \vartheta\sigma_S\sigma_K) + \Sigma_{\mathbf{s}}\epsilon_{\mathbf{q}}(\mathbf{s}_t). \tag{A.12}$$

The continuation return parameters $(\tilde{\mu}_{\mathbf{Q}}, \tilde{\sigma}^K_{\mathbf{Q}}, \tilde{\sigma}^S_{\mathbf{Q}})$ in (A.2) follow symmetrically by setting $\xi = 0$ and replacing current policies and valuations with their continuation counterparts $(\tilde{\mathbf{i}}, \tilde{\mathbf{m}}, \tilde{\mathbf{q}})$. Then, the squared volatility of the continuation wealth is $\Sigma_{\tilde{W}} = (\tilde{\sigma}^K_{\mathbf{Q}})^2 + 2\vartheta\tilde{\sigma}^K_{\mathbf{Q}}\tilde{\sigma}^S_{\mathbf{Q}} + (\tilde{\sigma}^S_{\mathbf{Q}})^2$, and the covariance between the wealth dynamics and the carbon-capital ratio is $\Sigma_{\tilde{W}\mathbf{s}} = -\sigma_K(\tilde{\sigma}^K_{\mathbf{Q}} + \vartheta\tilde{\sigma}^S_{\mathbf{Q}}) + \sigma_S(\tilde{\sigma}^S_{\mathbf{Q}} + \vartheta\tilde{\sigma}^K_{\mathbf{Q}})$. They are identical to (A.11) and (A.12) except with $\tilde{\mathbf{q}}(\mathbf{s}_t)$ replacing $\mathbf{q}(\mathbf{s}_t)$.

Step 2. Household's value functions. In equilibrium, goods and asset markets clear: $\mathbf{C}_t = C_t$ and $\Gamma_t = W_t = \mathbf{Q}_t = \mathbf{q}(\mathbf{s}_t)\mathbf{K}_t$. Symmetrically, for the continuation game, $\tilde{\Gamma}_t = \tilde{W}_t = \tilde{\mathbf{Q}}_t = \tilde{\mathbf{q}}(\mathbf{s}_t)\mathbf{K}_t$. Substituting these market clearing conditions and the value function conjecture (13) into the HJB

equation (11) yields:

$$\begin{aligned} 0 =& \frac{f(\mathbf{C}, J)}{(u(\mathbf{s})W)^{1-\gamma}} + \mu_{\mathbf{Q}}(\mathbf{s}) - \frac{\mathbf{c}(\mathbf{s})}{\mathbf{q}(\mathbf{s})} - \frac{\gamma \Sigma_W(\mathbf{s})}{2} + \Big(\mu_{\mathbf{s}}(\mathbf{s}; \mathbf{i}, \mathbf{m}) + (1-\gamma)\Sigma_{W\mathbf{s}}(\mathbf{s})\Big)\epsilon_u(\mathbf{s}) \\ &+ \frac{\Sigma_{\mathbf{s}}}{2}\left(\epsilon_u'(\mathbf{s}) - \gamma \epsilon_u(\mathbf{s})^2\right) + \frac{\xi}{1-\gamma}\left(\beta^\theta \left(\frac{\tilde{u}(\mathbf{s})}{u(\mathbf{s})}\frac{\tilde{\mathbf{q}}(\mathbf{s})}{\mathbf{q}(\mathbf{s})}\right)^{1-\gamma} - 1\right). \end{aligned} \tag{A.13}$$

The normalized aggregator in (A.13) can be expressed as:

$$\frac{f(\mathbf{C}, J)}{(u(\mathbf{s})W)^{1-\gamma}} = \frac{\rho}{1-\psi^{-1}}\left(\left(\frac{\mathbf{c}(\mathbf{s})}{\mathbf{q}(\mathbf{s})}\right)^{1-\psi^{-1}} \frac{1}{u(\mathbf{s})^{1-\psi^{-1}}} - 1\right). \tag{A.14}$$

Because the current household controls consumption C_t, the FOC (14) applies, which can be rewritten as:

$$\mathbf{c}(\mathbf{s}) = \rho^\psi u(\mathbf{s})^{1-\psi}\mathbf{q}(\mathbf{s}). \tag{A.15}$$

Combining (A.15) with (A.14), we obtain:

$$\begin{aligned} 0 =& \frac{1}{1-\psi^{-1}}\left(\frac{\mathbf{c}(\mathbf{s})}{\mathbf{q}(\mathbf{s})} - \rho\right) + \mu_{\mathbf{Q}}(\mathbf{s}) - \frac{\mathbf{c}(\mathbf{s})}{\mathbf{q}(\mathbf{s})} - \frac{\gamma \Sigma_W(\mathbf{s})}{2} + \Big(\mu_{\mathbf{s}}(\mathbf{s}; \mathbf{i}, \mathbf{m}) + (1-\gamma)\Sigma_{W\mathbf{s}}(\mathbf{s})\Big)\epsilon_u(\mathbf{s}) \\ &+ \frac{\Sigma_{\mathbf{s}}}{2}\left(\epsilon'_u(\mathbf{s}) - \gamma \epsilon_u(\mathbf{s})^2\right) + \frac{\xi}{1-\gamma}\left(\beta^\theta \left(\frac{\tilde{u}(\mathbf{s})}{u(\mathbf{s})}\frac{\tilde{\mathbf{q}}(\mathbf{s})}{\mathbf{q}(\mathbf{s})}\right)^{1-\gamma} - 1\right). \end{aligned} \tag{A.16}$$

Alternatively, substituting out $\mathbf{c}(\mathbf{s})/\mathbf{q}(\mathbf{s})$ in the first term, the above equation can be written as:

$$\begin{aligned} 0 =& \frac{\psi^{-1}\rho^\psi u(\mathbf{s})^{1-\psi} - \rho}{1-\psi^{-1}} + \mu_{\mathbf{Q}}(\mathbf{s}) - \frac{\gamma \Sigma_W(\mathbf{s})}{2} + \Big(\mu_{\mathbf{s}}(\mathbf{s}; \mathbf{i}, \mathbf{m}) + (1-\gamma)\Sigma_{W\mathbf{s}}(\mathbf{s})\Big)\epsilon_u(\mathbf{s}) \\ &+ \frac{\Sigma_{\mathbf{s}}}{2}\left(\epsilon'_u(\mathbf{s}) - \gamma \epsilon_u(\mathbf{s})^2\right) + \frac{\xi}{1-\gamma}\left(\beta^\theta \left(\frac{\tilde{u}(\mathbf{s})}{u(\mathbf{s})}\frac{\tilde{\mathbf{q}}(\mathbf{s})}{\mathbf{q}(\mathbf{s})}\right)^{1-\gamma} - 1\right). \end{aligned} \tag{A.17}$$

Following identical steps for the continuation utility (9), we have:

$$\begin{aligned} 0 =& \frac{f(\tilde{\mathbf{C}}, \tilde{J})}{(\tilde{u}(\mathbf{s})\tilde{W})^{1-\gamma}} + \tilde{\mu}_{\mathbf{Q}}(\mathbf{s}) - \frac{\tilde{\mathbf{c}}(\mathbf{s})}{\tilde{\mathbf{q}}(\mathbf{s})} - \frac{\gamma \Sigma_{\tilde{W}}(\mathbf{s})}{2} + \left(\mu_{\mathbf{s}}(\mathbf{s}; \tilde{\mathbf{i}}, \tilde{\mathbf{m}}) + (1-\gamma)\Sigma_{\tilde{W}\mathbf{s}}(\mathbf{s})\right)\epsilon_{\tilde{u}}(\mathbf{s}) \\ &+ \frac{\Sigma_{\mathbf{s}}}{2}\left(\epsilon'_{\tilde{u}}(\mathbf{s}) - \gamma \epsilon_{\tilde{u}}(\mathbf{s})^2\right). \end{aligned} \tag{A.18}$$

In a stationary MPE, $\tilde{\mathbf{c}}(\mathbf{s}) = \mathbf{c}(\mathbf{s})$. Applying the consumption FOC (A.15), the continuation aggregator becomes:

$$\begin{aligned} \frac{f(\tilde{\mathbf{C}}, \tilde{J})}{(\tilde{u}(\mathbf{s})\tilde{W})^{1-\gamma}} =& \frac{\rho}{1-\psi^{-1}}\left(\left(\frac{\mathbf{c}(\mathbf{s})}{\tilde{\mathbf{q}}(\mathbf{s})}\right)^{1-\psi^{-1}} \frac{1}{\tilde{u}(\mathbf{s})^{1-\psi^{-1}}} - 1\right) \\ =& \frac{1}{1-\psi^{-1}}\left(\rho^\psi u(\mathbf{s})^{1-\psi}\left(\frac{u(\mathbf{s})}{\tilde{u}(\mathbf{s})}\frac{\mathbf{q}(\mathbf{s})}{\tilde{\mathbf{q}}(\mathbf{s})}\right)^{1-\psi^{-1}} - \rho\right). \end{aligned} \tag{A.19}$$

Substituting (A.19) into (A.18) yields the ODE for the continuation utility:

$$\begin{aligned} 0 =& \frac{1}{1-\psi^{-1}}\left(\rho^{\psi}u(\mathbf{s})^{1-\psi}\left(\frac{u(\mathbf{s})}{\tilde{u}(\mathbf{s})}\frac{\mathbf{q}(\mathbf{s})}{\tilde{\mathbf{q}}(\mathbf{s})}\right)^{1-\psi^{-1}}-\rho\right)+\tilde{\mu}_{\mathbf{Q}}(\mathbf{s})-\frac{\mathbf{c}(\mathbf{s})}{\tilde{\mathbf{q}}(\mathbf{s})}-\frac{\gamma\Sigma_{\tilde{W}}(\mathbf{s})}{2} \\ &+\left(\mu_{\mathbf{s}}(\mathbf{s};\tilde{\mathbf{i}},\tilde{\mathbf{m}})+(1-\gamma)\Sigma_{\tilde{W}\mathbf{s}}(\mathbf{s})\right)\epsilon_{\tilde{u}}(\mathbf{s})+\frac{\Sigma_{\mathbf{s}}}{2}\left(\epsilon'_{\tilde{u}}(\mathbf{s})-\gamma\epsilon_{\tilde{u}}(\mathbf{s})^2\right). \end{aligned} \tag{A.20}$$

Step 3. Equilibrium SDF. Under quasi-hyperbolic recursive preferences, the household's SDF is:

$$\mathbb{M}_t=\begin{cases}\exp\left(\int_0^t f_J(C_u,J_u)du\right)f_C(C_t,J_t) & \text{if } t\in[0,\tau),\\ \beta^{\theta}\exp\left(\int_0^{\tau} f_J(C_u,J_u)du+\int_{\tau}^t f_J(\tilde{C}_u,\tilde{J}_u)du\right)f_C(\tilde{C}_t,\tilde{J}_t) & \text{if } t\in[\tau,\infty),\end{cases} \tag{A.21}$$

where τ denotes the arrival time of the future self. The first line in (A.21) shows that prior to the future self's arrival ($t<\tau$), the SDF evolves according to the standard formulation for stochastic differential utility, governed by the current self's value function J_t (Duffie and Epstein, 1992). The second line in (A.21) shows that upon the arrival of the future self, the SDF experiences a discrete downward jump by the present-bias factor β^{θ} and its evolution is determined by the future self's value function $\tilde{J}_t$, reflecting a shift in the valuation perspective.

The dynamics of the SDF in (A.21) are given by:

$$\frac{d\mathbb{M}_t}{\mathbb{M}_t}=f_J(C_t,J_t)dt+\frac{df_C(C_t,J_t)}{f_C(C_t,J_t)}+\left(\beta^{\theta}\frac{f_C(\tilde{C}_t,\tilde{J}_t)}{f_C(C_t,J_t)}-1\right)d\mathcal{P}_t. \tag{A.22}$$

In equilibrium, $C_t=\mathbf{C}_t$ and $\tilde{\mathbf{C}}_t=\mathbf{C}_t$. Then, the first term on the right-hand side of (A.22) becomes

$$f_J(\mathbf{C}_t,J_t)=\frac{\rho(1-\gamma)}{1-\psi^{-1}}\left(\frac{\psi^{-1}-\gamma}{1-\gamma}\left(\frac{\mathbf{c}(\mathbf{s})}{u(\mathbf{s})\mathbf{q}(\mathbf{s})}\right)^{1-\psi^{-1}}-1\right).$$

Turning to the second term in (A.22), the first-order condition for consumption (14) under the conjectured value function (13) implies:

$$f_C(\mathbf{C}_t,J_t)=u(\mathbf{s}_t)^{1-\gamma}\mathbf{q}(\mathbf{s}_t)^{-\gamma}\mathbf{K}_t^{-\gamma}.$$

Thus, $f_C(\mathbf{C}_t,J_t)$ is driven by the dynamics of the carbon-capital ratio $\mathbf{s}_t$ and the aggregate capital stock $\mathbf{K}_t$. By the stochastic product rule,

$$\frac{df_C(\mathbf{C}_t,J_t)}{f_C(\mathbf{C}_t,J_t)}=\frac{d\mathbf{K}_t^{-\gamma}}{\mathbf{K}_t^{-\gamma}}+\frac{du(\mathbf{s}_t)^{1-\gamma}\mathbf{q}(\mathbf{s}_t)^{-\gamma}}{u(\mathbf{s}_t)^{1-\gamma}\mathbf{q}(\mathbf{s}_t)^{-\gamma}}+\frac{d\left\langle u(\mathbf{s}_t)^{1-\gamma}\mathbf{q}(\mathbf{s}_t)^{-\gamma},\mathbf{K}_t^{-\gamma}\right\rangle}{u(\mathbf{s}_t)^{1-\gamma}\mathbf{q}(\mathbf{s}_t)^{-\gamma}\mathbf{K}_t^{-\gamma}}. \tag{A.23}$$

We compute the right-hand side of (A.23) term by term using Ito's Lemma. The first term is given by:

$$\frac{d\mathbf{K}_t^{-\gamma}}{\mathbf{K}_t^{-\gamma}}=-\left(\gamma(\mathbf{i}_t-\delta_K-\mathcal{D}(\mathbf{s}_t))-\frac{\gamma(\gamma+1)}{2}\sigma_K^2\right)dt-\gamma\sigma_K d\mathcal{B}_t^K. \tag{A.24}$$

For the second term, it is more convenient to define

$$h(\mathbf{s}_t) = \exp\Big((1-\gamma)\ln u(\mathbf{s}_t) - \gamma \ln \mathbf{q}(\mathbf{s}_t)\Big).$$

Applying Ito's Lemma, the dynamics of $h(\mathbf{s}_t)$ are:

$$\frac{dh(\mathbf{s}_t)}{h(\mathbf{s}_t)} = \epsilon_h(\mathbf{s}_t)\frac{d\mathbf{s}_t}{\mathbf{s}_t} + \frac{1}{2}\epsilon_h'(\mathbf{s}_t)\Sigma_{\mathbf{s}}dt,$$

where

$$\epsilon_h(\mathbf{s}_t) = (1-\gamma)\epsilon_u(\mathbf{s}_t) - \gamma\epsilon_{\mathbf{q}}(\mathbf{s}_t),$$

$$\epsilon_h'(\mathbf{s}_t) = \big((1-\gamma)\epsilon_u(\mathbf{s}_t) - \gamma\epsilon_{\mathbf{q}}(\mathbf{s}_t)\big)^2 + (1-\gamma)\big(\epsilon_u'(\mathbf{s}_t) - \epsilon_u(\mathbf{s}_t)^2\big) - \gamma\big(\epsilon_{\mathbf{q}}'(\mathbf{s}_t) - \epsilon_{\mathbf{q}}(\mathbf{s}_t)^2\big).$$

The third term, representing the quadratic covariation in (A.23), is:

$$\frac{d\big\langle h(\mathbf{s}_t), \mathbf{K}_t^{-\gamma}\big\rangle}{h(\mathbf{s}_t)\mathbf{K}_t^{-\gamma}} = \big((1-\gamma)\epsilon_u(\mathbf{s}_t) - \gamma\epsilon_{\mathbf{q}}(\mathbf{s}_t)\big)\gamma\big(\sigma_K^2 - \vartheta\sigma_K\sigma_S\big)dt.$$

Finally, differentiating the Epstein-Zin aggregator (8) yields $f_C(\mathbf{C}_t, J_t) = \rho\mathbf{C}_t^{-\psi^{-1}}(u(\mathbf{s}_t)\mathbf{q}(\mathbf{s}_t)\mathbf{K}_t)^{\psi^{-1}-\gamma}$ and $f_C(\mathbf{C}_t, \tilde{J}_t) = \rho\mathbf{C}_t^{-\psi^{-1}}(\tilde{u}(\mathbf{s}_t)\tilde{\mathbf{q}}(\mathbf{s}_t)\mathbf{K}_t)^{\psi^{-1}-\gamma}$. Thus, the jump coefficient in (A.22) is:

$$\beta^\theta \frac{f_C(\mathbf{C}_t, \tilde{J}_t)}{f_C(\mathbf{C}_t, J_t)} = \beta^\theta\left(\frac{\tilde{u}(\mathbf{s}_t)\tilde{\mathbf{q}}(\mathbf{s}_t)}{u(\mathbf{s}_t)\mathbf{q}(\mathbf{s}_t)}\right)^{\psi^{-1}-\gamma}. \tag{A.25}$$

Substituting (A.24) through (A.25) back into (A.22), we obtain:

$$\begin{aligned}\frac{d\mathbb{M}_t}{\mathbb{M}_t} =& \Big(f_J(\mathbf{C}_t, J_t) - \gamma(\mathbf{i}(\mathbf{s}_t) - \delta_K - \mathcal{D}(\mathbf{s}_t)) + \frac{\gamma(\gamma+1)}{2}\sigma_K^2\Big)dt \\ &+ \left(\Big(\mu_{\mathbf{s}}(\mathbf{s}_t) + \gamma\big(\sigma_K^2 - \vartheta\sigma_K\sigma_S\big)\Big)\epsilon_h(\mathbf{s}_t) + \frac{1}{2}\Sigma_{\mathbf{s}}\epsilon_h'(\mathbf{s}_t)\right)dt \\ &- \big(\gamma\sigma_K + \sigma_K\epsilon_h(\mathbf{s}_t)\big)d\mathcal{B}_t^K + \sigma_S\epsilon_h(\mathbf{s}_t)d\mathcal{B}_t^S + \left(\beta^\theta\left(\frac{\tilde{u}(\mathbf{s}_t)\tilde{\mathbf{q}}(\mathbf{s}_t)}{u(\mathbf{s}_t)\mathbf{q}(\mathbf{s}_t)}\right)^{\psi^{-1}-\gamma} - 1\right)d\mathcal{P}_t.\end{aligned}$$

Matching terms with the SDF in (16) yields:

$$\eta_{\mathbb{M}}^K = \gamma\sigma_K + \sigma_K\epsilon_h(\mathbf{s}_t), \quad \eta_{\mathbb{M}}^S = -\sigma_S\epsilon_h(\mathbf{s}_t), \quad \text{and} \quad \chi_t = \beta^\theta\left(\frac{\tilde{u}(\mathbf{s}_t)\tilde{\mathbf{q}}(\mathbf{s}_t)}{u(\mathbf{s}_t)\mathbf{q}(\mathbf{s}_t)}\right)^{\psi^{-1}-\gamma}. \tag{A.26}$$

Note that $\eta_{\mathbb{M}}^K$, $\eta_{\mathbb{M}}^S$, and χ_t are functions of the state variable $\mathbf{s}_t$.

Step 4. Firm value and equilibrium risk-free rate. As derived earlier, the aggregate firm value $\mathbf{q}(\mathbf{s})$ satisfies the HJB equation (17) with the SDF risk adjustments given in (A.26). Since the goods market clears in the competitive equilibrium, $\mathbf{c}(\mathbf{s}) = A - \mathbf{i}(\mathbf{s}) - \phi(\mathbf{i}) - \mathbf{m}(\mathbf{s})$, the HJB

equation (17) implies the equilibrium risk-free rate:

$$\begin{aligned} r(\mathbf{s}) =& \frac{\mathbf{c}(\mathbf{s})}{\mathbf{q}(\mathbf{s})} + \mathbf{i}(\mathbf{s}) - \delta_K - \mathcal{D}(\mathbf{s}) - \sigma_K(\eta_{\mathbb{M}}^K + \vartheta\eta_{\mathbb{M}}^S) \\ &+ \Big(\mu_{\mathbf{s}}(\mathbf{s}) - (\sigma_K^2 - \vartheta\sigma_S\sigma_K) + \sigma_K(\eta_{\mathbb{M}}^K + \vartheta\eta_{\mathbb{M}}^S) - \sigma_S(\eta_{\mathbb{M}}^S + \vartheta\eta_{\mathbb{M}}^K)\Big)\epsilon_{\mathbf{q}}(\mathbf{s}) \\ &+ \frac{\Sigma_{\mathbf{s}}}{2}\epsilon_{\mathbf{q}}'(\mathbf{s}) + \xi\chi(\mathbf{s})\Big(\frac{\tilde{\mathbf{q}}(\mathbf{s})}{\mathbf{q}(\mathbf{s})} - 1\Big). \end{aligned} \tag{A.27}$$

To determine the consumption-to-q ratio (the first term on the right-hand side of (A.27)), we use the household's HJB equation (A.16):

$$\begin{aligned} \frac{\mathbf{c}(\mathbf{s})}{\mathbf{q}(\mathbf{s})} =& \rho - (1-\psi^{-1})\left[\mu_{\mathbf{Q}}(\mathbf{s}) - \frac{\mathbf{c}(\mathbf{s})}{\mathbf{q}(\mathbf{s})} - \frac{\gamma\Sigma_W(\mathbf{s})}{2} + \Big(\mu_{\mathbf{s}}(\mathbf{s}) + (1-\gamma)\Sigma_{W\mathbf{s}}(\mathbf{s})\Big)\epsilon_u(\mathbf{s}) + \frac{\Sigma_{\mathbf{s}}}{2}\Big(\epsilon_u'(\mathbf{s}) - \gamma\epsilon_u(\mathbf{s})\Big)\right] \\ &- \frac{\xi(1-\psi^{-1})}{1-\gamma}\left(\beta^\theta\left(\frac{\tilde{u}(\mathbf{s})}{u(\mathbf{s})}\frac{\tilde{\mathbf{q}}(\mathbf{s})}{\mathbf{q}(\mathbf{s})}\right)^{1-\gamma} - 1\right). \end{aligned} \tag{A.28}$$

For later use, we define the term in the square brackets of (A.28) as:

$$\Delta(u,\mathbf{q}) \equiv \mu_{\mathbf{Q}}(\mathbf{s}) - \frac{\mathbf{c}(\mathbf{s})}{\mathbf{q}(\mathbf{s})} - \frac{\gamma\Sigma_W(\mathbf{s})}{2} + \Big(\mu_{\mathbf{s}}(\mathbf{s}) + (1-\gamma)\Sigma_{W\mathbf{s}}(\mathbf{s})\Big)\epsilon_u(\mathbf{s}) + \frac{\Sigma_{\mathbf{s}}}{2}\Big(\epsilon_u'(\mathbf{s}) - \gamma\epsilon_u(\mathbf{s})\Big). \tag{A.29}$$

Substituting (A.9), (A.11), and (A.12) into (A.29) and suppressing the argument $\mathbf{s}$ in the elasticities, we obtain:

$$\begin{aligned} \Delta(u,\mathbf{q}) =& \underbrace{\mathbf{i}(\mathbf{s}) - \delta_K - \mathcal{D}(\mathbf{s}) + \Big(\mu_{\mathbf{s}}(\mathbf{s}) - (\sigma_K^2 - \vartheta\sigma_K\sigma_S)\Big)\epsilon_{\mathbf{q}}(\mathbf{s}) + \frac{\Sigma_{\mathbf{s}}}{2}\epsilon_{\mathbf{q}}'(\mathbf{s})}_{=\mu_{\mathbf{Q}}(\mathbf{s}) - \mathbf{c}(\mathbf{s})/\mathbf{q}(\mathbf{s})} \\ &\underbrace{- \Big(\frac{\gamma}{2}\sigma_K^2 - \gamma(\sigma_K^2 - \vartheta\sigma_K\sigma_S)\epsilon_{\mathbf{q}}(\mathbf{s}) + \frac{\Sigma_{\mathbf{s}}}{2}\gamma\epsilon_{\mathbf{q}}(\mathbf{s})^2\Big)}_{=\gamma\Sigma_W(\mathbf{s})/2} \\ &+ \mu_{\mathbf{s}}(\mathbf{s})\epsilon_u(\mathbf{s}) \underbrace{- (1-\gamma)(\sigma_K^2 - \vartheta\sigma_K\sigma_S)\epsilon_u(\mathbf{s})\frac{\Sigma_{\mathbf{s}}}{2}(1-\gamma)2\epsilon_u(\mathbf{s})\epsilon_{\mathbf{q}}(\mathbf{s})}_{=(1-\gamma)\Sigma_{W\mathbf{s}}(\mathbf{s})\epsilon_u(\mathbf{s})} + \frac{\Sigma_{\mathbf{s}}}{2}\Big(\epsilon_u'(\mathbf{s}) - \gamma\epsilon_u(\mathbf{s})^2\Big) \\ =& \mathbf{i}(\mathbf{s}) - \delta_K - \mathcal{D}(\mathbf{s}) - \frac{\gamma}{2}\sigma_K^2 + \Big(\mu_{\mathbf{s}}(\mathbf{s}) - (1-\gamma)(\sigma_K^2 - \vartheta\sigma_K\sigma_S)\Big)(\epsilon_u(\mathbf{s}) + \epsilon_{\mathbf{q}}(\mathbf{s})) \\ &+ \frac{\Sigma_{\mathbf{s}}}{2}\Big((1-\gamma)2\epsilon_u(\mathbf{s})\epsilon_{\mathbf{q}}(\mathbf{s}) + \epsilon_{\mathbf{q}}'(\mathbf{s}) - \gamma\epsilon_{\mathbf{q}}(\mathbf{s})^2 + \epsilon_u'(\mathbf{s}) - \gamma\epsilon_u(\mathbf{s})^2\Big) \end{aligned} \tag{A.30}$$

Combining (A.28) and (A.30) with (A.27), we obtain the equilibrium risk-free rate:

$$\begin{aligned} r(\mathbf{s}) =& \rho + \psi^{-1}\Big(\mathbf{i}(\mathbf{s}) - \delta_K - \mathcal{D}(\mathbf{s})\Big) - \frac{\gamma(\psi^{-1}+1)}{2}\sigma_K^2 \\ & - \mu_{\mathbf{s}}(\mathbf{s})\big((1-\psi^{-1})(\epsilon_u + \epsilon_{\mathbf{q}}) - \epsilon_{\mathbf{q}}\big) - (\sigma_K^2 - \vartheta\sigma_S\sigma_K)\big(\psi^{-1}(1-\gamma)(\epsilon_u+\epsilon_{\mathbf{q}}) - \gamma\epsilon_{\mathbf{q}}\big) \\ & - \frac{\Sigma_{\mathbf{s}}}{2}\Big(\gamma(\psi^{-1}+1)\epsilon_{\mathbf{q}}^2 + (1-\psi^{-1})(\epsilon_u' - \gamma\epsilon_u^2) - \psi^{-1}(1-\gamma)2\epsilon_u\epsilon_{\mathbf{q}} - \psi^{-1}\epsilon_{\mathbf{q}}'\Big) \\ & + \xi\left(\chi\Big(\frac{\tilde{\mathbf{q}}(\mathbf{s})}{\mathbf{q}(\mathbf{s})} - 1\Big) - \frac{1-\psi^{-1}}{1-\gamma}\left(\beta^\theta\left(\frac{\tilde{u}(\mathbf{s})}{u(\mathbf{s})}\frac{\tilde{\mathbf{q}}(\mathbf{s})}{\mathbf{q}(\mathbf{s})}\right)^{1-\gamma} - 1\right)\right). \end{aligned} \tag{A.31}$$

Step 5. Continuation firm value. The continuation SDF of the current household is given by

$$\frac{d\tilde{\mathbb{M}}_t}{\tilde{\mathbb{M}}_t} = -\tilde{r}_t dt - \tilde{\eta}_{\mathbb{M}}^K d\mathcal{B}_t^K - \tilde{\eta}_{\mathbb{M}}^S d\mathcal{B}_t^S, \tag{A.32}$$

where $\tilde{\eta}_{\mathbb{M}}^K = \gamma\sigma_K + \sigma_K\epsilon_{\tilde{h}}$ and $\tilde{\eta}_{\mathbb{M}}^S = -\sigma_S\epsilon_{\tilde{h}}$ as in (A.26). After the arrival of a future household, the current household employs exponential discounting and the discount rate no longer changes. Hence, unlike (16), the current household's continuation SDF (A.32) contains no jump term.

In the competitive equilibrium, the current household uses (A.32) to price the representative firm in the continuation game. Given the stationary MPE in which $\tilde{\mathbf{i}} = \mathbf{i}$ and $\tilde{\mathbf{m}} = \mathbf{m}$, the continuation firm value $\tilde{\mathbf{q}}(\mathbf{s})$ satisfies

$$\begin{aligned} \tilde{r}\tilde{\mathbf{q}}(\mathbf{s}) =& A - \mathbf{i}(\mathbf{s}) - \phi(\mathbf{i}) - \mathbf{m} + (\mathbf{i} - \delta_K - \mathcal{D}(\mathbf{s}))\tilde{\mathbf{q}}(\mathbf{s}) - \sigma_K(\tilde{\eta}_{\mathbb{M}}^K + \vartheta\tilde{\eta}_{\mathbb{M}}^S)(1-\epsilon_{\tilde{\mathbf{q}}}(\mathbf{s}))\tilde{\mathbf{q}}(\mathbf{s}) \\ & + \Big(\mu_{\mathbf{s}}(\mathbf{s}) - \big(\sigma_K^2 - \vartheta\sigma_S\sigma_K\big) - \sigma_S(\tilde{\eta}_{\mathbb{M}}^S + \vartheta\tilde{\eta}_{\mathbb{M}}^K)\Big)\epsilon_{\tilde{\mathbf{q}}}(\mathbf{s})\tilde{\mathbf{q}}(\mathbf{s}) + \frac{\Sigma_{\mathbf{s}}}{2}\epsilon_{\tilde{\mathbf{q}}}'(\mathbf{s})\tilde{\mathbf{q}}(\mathbf{s}), \end{aligned} \tag{A.33}$$

where the continuation risk-free rate $\tilde{r}(\mathbf{s})$ is given by:

$$\begin{aligned} \tilde{r}(\mathbf{s}) =& \rho + \psi^{-1}\Big(\mathbf{i}(\mathbf{s}) - \delta_K - \mathcal{D}(\mathbf{s})\Big) - \frac{\gamma(\psi^{-1}+1)}{2}\sigma_K^2 \\ & - \mu_{\mathbf{s}}(\mathbf{s})\big((1-\psi^{-1})(\epsilon_{\tilde{u}} + \epsilon_{\tilde{\mathbf{q}}}) - \epsilon_{\tilde{\mathbf{q}}}\big) - (\sigma_K^2 - \vartheta\sigma_S\sigma_K)\big(\psi^{-1}(1-\gamma)(\epsilon_{\tilde{u}}+\epsilon_{\tilde{\mathbf{q}}}) - \gamma\epsilon_{\tilde{\mathbf{q}}}\big) \\ & - \frac{\Sigma_{\mathbf{s}}}{2}\Big(\gamma(\psi^{-1}+1)\epsilon_{\tilde{\mathbf{q}}}^2 + (1-\psi^{-1})(\epsilon_{\tilde{u}}' - \gamma\epsilon_{\tilde{u}}^2) - \psi^{-1}(1-\gamma)2\epsilon_{\tilde{u}}\epsilon_{\tilde{\mathbf{q}}} - \psi^{-1}\epsilon_{\tilde{\mathbf{q}}}'\Big). \end{aligned} \tag{A.34}$$

The continuation risk-free rate (A.34) shares the same functional form as the equilibrium risk-free rate (A.31), with two key distinctions. First, the continuation utility $\tilde{u}$ and firm value $\tilde{\mathbf{q}}$ replace their current counterparts u and $\mathbf{q}$. Second, (A.34) contains no jump term.

Summary. Equations (A.17), (A.20), (17), and (A.33) constitute a system of four differential equations that jointly determine the household's value, $u(\mathbf{s})$ and $\tilde{u}(\mathbf{s})$, and the aggregate equity value, $\mathbf{q}(\mathbf{s})$ and $\tilde{\mathbf{q}}(\mathbf{s})$. These valuation functions pin down the equilibrium investment $\mathbf{i}(\mathbf{s})$, consumption $\mathbf{c}(\mathbf{s})$, and portfolio choice Γ via the FOCs (18), (14), and (15), respectively.

B Appendix for Section 4: The Planner's Economy

B.1 First-Best Resource Allocation

Substituting the value function (27) into the FOC for the aggregate investment (25) and the FOC for the aggregate mitigation (26), we obtain

$$b(\mathbf{s}) = \mathbf{c}^*(\mathbf{s})^{\frac{1}{1-\psi}} \left(\frac{\rho\big(1+\phi(\mathbf{i}^*)\big)b(\mathbf{s})}{b(\mathbf{s}) - \mathbf{s}b'(\mathbf{s})} \right)^{-\frac{\psi}{1-\psi}}, \tag{B.1}$$

$$\rho \left(\frac{\mathbf{c}^*(\mathbf{s})}{b(\mathbf{s})} \right)^{-\psi^{-1}} = \lambda A \alpha_1^{\alpha_2} \alpha_2 \mathbf{m}^*(\mathbf{s})^{\alpha_2 - 1} b'(\mathbf{s}), \tag{B.2}$$

where $\mathbf{c}^*(\mathbf{s}) = A - \mathbf{i}^*(\mathbf{s}) - \phi(\mathbf{i}^*) - \mathbf{m}^*(\mathbf{s})$ implied by the aggregate resource constraint.

Next, by substituting the value function (27) into the HJB equation (23), we obtain

$$\begin{aligned} 0 =& \frac{f(\mathbf{C}^*, V)}{(b(\mathbf{s})\mathbf{K})^{1-\gamma}} + \mathbf{i}^*(\mathbf{s}) - \delta_K - \mathcal{D}(\mathbf{s}) - \frac{\gamma\sigma_K^2}{2} + \Big(\mu_{\mathbf{s}}(\mathbf{s}; \mathbf{i}^*, \mathbf{m}^*) - (1-\gamma)(\sigma_K^2 - \vartheta\sigma_S\sigma_K)\Big)\frac{b'(\mathbf{s})\mathbf{s}}{b(\mathbf{s})} \\ &+ \frac{\Sigma_{\mathbf{s}}\mathbf{s}^2}{2}\left(\frac{b''(\mathbf{s})}{b(\mathbf{s})} - \gamma\left(\frac{b'(\mathbf{s})}{b(\mathbf{s})} \right)^2 \right) + \frac{\xi}{1-\gamma}\left(\beta^\theta \left(\frac{\tilde{b}(\mathbf{s})}{b(\mathbf{s})} \right)^{1-\gamma} - 1 \right), \end{aligned} \tag{B.3}$$

and similarly for the continuation value equation (24),

$$\begin{aligned} 0 =& \frac{f(\tilde{\mathbf{C}}^*, \tilde{V})}{(\tilde{b}(\mathbf{s})\mathbf{K})^{1-\gamma}} + \tilde{\mathbf{i}}^*(\mathbf{s}) - \delta_K - \mathcal{D}(\mathbf{s}) - \frac{\gamma\sigma_K^2}{2} + \Big(\mu_{\mathbf{s}}(\mathbf{s}; \tilde{\mathbf{i}}^*, \tilde{\mathbf{m}}^*) - (1-\gamma)(\sigma_K^2 - \vartheta\sigma_S\sigma_K)\Big)\frac{\tilde{b}'(\mathbf{s})\mathbf{s}}{\tilde{b}(\mathbf{s})} \\ &+ \frac{\Sigma_{\mathbf{s}}\mathbf{s}^2}{2}\left(\frac{\tilde{b}''(\mathbf{s})}{\tilde{b}(\mathbf{s})} - \gamma\left(\frac{\tilde{b}'(\mathbf{s})}{\tilde{b}(\mathbf{s})} \right)^2 \right). \end{aligned}$$

Using the aggregator (8) and the value function (27),

$$\frac{f(\mathbf{C}^*, V)}{(b(\mathbf{s})\mathbf{K})^{1-\gamma}} = \frac{\rho}{1-\psi^{-1}}\left(\left(\frac{\mathbf{c}^*(\mathbf{s})}{b(\mathbf{s})} \right)^{1-\psi^{-1}} - 1 \right). \tag{B.4}$$

Together with the FOC (25), the HJB equation (B.3) becomes

$$\begin{aligned} 0 =& \frac{\rho}{1-\psi^{-1}}\left[\left(\frac{b(\mathbf{s}) - b'(\mathbf{s})\mathbf{s}}{\rho(1+\phi'(\mathbf{i}^*))} \right)^{1-\psi} - 1 \right] + \mathbf{i}^*(\mathbf{s}) - \delta_K - \mathcal{D}(\mathbf{s}) - \frac{\gamma\sigma_K^2}{2} \\ &+ \left[\mu_{\mathbf{s}}(\mathbf{s}; \mathbf{i}^*, \mathbf{m}^*) - (1-\gamma)(\sigma_K^2 - \vartheta\sigma_S\sigma_K) \right]\frac{b'(\mathbf{s})\mathbf{s}}{b(\mathbf{s})} \\ &+ \frac{\Sigma_{\mathbf{s}}\mathbf{s}^2}{2}\left[\frac{b''(\mathbf{s})}{b(\mathbf{s})} - \gamma\left(\frac{b'(\mathbf{s})}{b(\mathbf{s})} \right)^2 \right] + \frac{\xi}{1-\gamma}\left[\beta^\theta \left(\frac{\tilde{b}(\mathbf{s})}{b(\mathbf{s})} \right)^{1-\gamma} - 1 \right]. \end{aligned} \tag{B.5}$$

Because the current planner does not control $\tilde{\mathbf{i}}^*$ to maximize her continuation value, no first-order condition links $\tilde{\mathbf{C}}^*$ to $\tilde{b}(\mathbf{s})$. However, imposing the stationary MPE conditions ($\tilde{\mathbf{C}}^* = \mathbf{C}^*$, $\tilde{\mathbf{i}} = \mathbf{i}^*$,

and $\tilde{\mathbf{m}} = \mathbf{m}^*$) and substituting the optimality condition (B.1) into (B.4) yields:

$$\begin{aligned} 0 =& \frac{\rho}{1-\psi^{-1}}\left[\left(\frac{b(\mathbf{s})-b'(\mathbf{s})\mathbf{s}}{\rho(1+\phi'(\mathbf{i}^*))}\right)^{1-\psi}\left(\frac{b(\mathbf{s})}{\tilde{b}(\mathbf{s})}\right)^{1-\psi^{-1}}-1\right]+\mathbf{i}^*(\mathbf{s})-\delta_K-\mathcal{D}(\mathbf{s})-\frac{\gamma\sigma_K^2}{2} \\ &+\left(\mu_{\mathbf{s}}(\mathbf{s};\mathbf{i}^*,\mathbf{m}^*)-(1-\gamma)(\sigma_K^2-\vartheta\sigma_S\sigma_K)\right)\frac{\tilde{b}'(\mathbf{s})\mathbf{s}}{\tilde{b}(\mathbf{s})}+\frac{\Sigma_{\mathbf{s}}\mathbf{s}^2}{2}\left(\frac{\tilde{b}''(\mathbf{s})}{\tilde{b}(\mathbf{s})}-\gamma\left(\frac{\tilde{b}'(\mathbf{s})}{\tilde{b}(\mathbf{s})}\right)^2\right). \end{aligned} \tag{B.6}$$

Summary. Equations (B.1), (B.2), (B.5), and (B.6) form a system of four ordinary differential equations that jointly determine $b(\mathbf{s})$, $\tilde{b}(\mathbf{s})$, $\mathbf{i}^*(\mathbf{s})$, and $\mathbf{m}^*(\mathbf{s})$.

B.2 Asset Pricing Implications in the Planner's Economy

With present-biased preferences, the SDF $\{\mathbb{M}_t^* : t \geq 0\}$ implied by the planner's solution is:

$$\mathbb{M}_t^* = \begin{cases} \exp\left(\int_0^t f_V(\mathbf{C}_u^*, V_u)du\right) f_{\mathbf{C}}(\mathbf{C}_t^*, V_t) & \text{if } t \in [0,\tau), \\ \beta^\theta \exp\left(\int_0^\tau f_V(\mathbf{C}_u^*, V_u)du + \int_\tau^t f_V(\tilde{\mathbf{C}}_u^*, \tilde{V}_u)du\right) f_{\mathbf{C}}(\tilde{\mathbf{C}}_t^*, \tilde{V}_t) & \text{if } t \in [\tau,\infty), \end{cases} \tag{B.7}$$

where τ denotes the arrival time of the future self. The dynamics of the SDF in (B.7) are given by:

$$\frac{d\mathbb{M}_t^*}{\mathbb{M}_t^*} = f_V(\mathbf{C}_t^*, V_t)dt + \frac{df_{\mathbf{C}}(\mathbf{C}_t^*, V_t)}{f_{\mathbf{C}}(\mathbf{C}_t^*, V_t)} + \left(\beta^\theta \frac{f_{\mathbf{C}}(\tilde{\mathbf{C}}_t^*, \tilde{V}_t)}{f_{\mathbf{C}}(\mathbf{C}_t^*, V_t)} - 1\right) d\mathcal{P}_t. \tag{B.8}$$

We directly compute the terms on the right-hand side of (B.8). First, differentiating the Epstein-Zin aggregator (8) yields:

$$f_V(\mathbf{C}_t^*, V_t) = \frac{\rho(1-\gamma)}{1-\psi^{-1}}\left(\frac{\psi^{-1}-\gamma}{1-\gamma}\left(\frac{\mathbf{c}^*(\mathbf{s})}{b(\mathbf{s})}\right)^{1-\psi^{-1}} - 1\right). \tag{B.9}$$

Second, substituting the current value function (27) and the first-best Tobin's q (28) into the FOC for aggregate investment (25) yields:

$$f_{\mathbf{C}}(\mathbf{C}_t^*, V_t) = \frac{b(\mathbf{s}_t)^{1-\gamma}\mathbf{K}_t^{-\gamma}}{\mathbf{q}^*(\mathbf{s}_t)}.$$

Thus, $f_{\mathbf{C}}(\mathbf{C}_t^*, V_t)$ is driven by the dynamics of $\mathbf{s}_t$ and $\mathbf{K}_t$. Applying Ito's product rule gives:

$$\frac{df_{\mathbf{C}}(\mathbf{C}_t^*, V_t)}{f_{\mathbf{C}}(\mathbf{C}_t^*, V_t)} = \frac{d\mathbf{K}_t^{-\gamma}}{\mathbf{K}_t^{-\gamma}} + \frac{dh^*(\mathbf{s}_t)}{h^*(\mathbf{s}_t)} + \frac{d\langle h^*(\mathbf{s}_t), \mathbf{K}_t^{-\gamma}\rangle}{h^*(\mathbf{s}_t)\mathbf{K}_t^{-\gamma}}, \tag{B.10}$$

where $h^*(\mathbf{s}_t) \equiv \exp\Big((1-\gamma)\ln b(\mathbf{s}_t) - \ln \mathbf{q}^*(\mathbf{s}_t)\Big)$. Evaluating the right-hand side of (B.10) term by term using Ito's lemma, the first term is given by (A.24). The second term is:

$$\frac{dh^*(\mathbf{s}_t)}{h^*(\mathbf{s}_t)} = \epsilon_{h^*}(\mathbf{s}_t)\frac{d\mathbf{s}_t}{\mathbf{s}_t} + \frac{\Sigma_{\mathbf{s}}}{2}\epsilon'_{h^*}(\mathbf{s}_t)dt, \tag{B.11}$$

where $\epsilon_{h^*}(\mathbf{s}_t) = (1-\gamma)\epsilon_b(\mathbf{s}_t) - \epsilon_{\mathbf{q}^*}(\mathbf{s}_t)$, and

$$\epsilon'_{h^*}(\mathbf{s}_t) = \epsilon_{h^*}(\mathbf{s}_t)^2 + (1-\gamma)\big(\epsilon'_b(\mathbf{s}_t) - \epsilon_b(\mathbf{s}_t)^2\big) - \big(\epsilon'_{\mathbf{q}^*}(\mathbf{s}_t) - \epsilon_{\mathbf{q}^*}(\mathbf{s}_t)^2\big).$$

The third term, representing the quadratic covariation in (B.10), is:

$$\frac{d\big\langle h^*(\mathbf{s}_t), \mathbf{K}_t^{-\gamma}\big\rangle}{h^*(\mathbf{s}_t)\mathbf{K}_t^{-\gamma}} = \gamma\big(\sigma_K^2 - \vartheta\sigma_K\sigma_S\big)\epsilon_{h^*}(\mathbf{s}_t)dt.$$

Finally, imposing the stationary MPE condition $\mathbf{C}_t^* = \tilde{\mathbf{C}}_t^*$, the jump component in (B.8) becomes:

$$\beta^\theta \frac{f_{\mathbf{C}}(\tilde{\mathbf{C}}_t^*, \tilde{V}_t)}{f_{\mathbf{C}}(\mathbf{C}_t^*, V_t)} = \beta^\theta \left(\frac{\tilde{b}(\mathbf{s}_t)}{b(\mathbf{s}_t)}\right)^{\psi^{-1}-\gamma}. \tag{B.12}$$

We first characterize the first-best consumption-to-q ratio, which corresponds to the planner's marginal propensity to consume out of aggregate wealth. Following the exact same logic as in the decentralized market economy (specifically, the steps leading to Equation (A.28) in Appendix A.3), we can determine this ratio by substituting the planner's value function into the corresponding HJB equation. Since the algebraic steps are structurally identical to those in the market equilibrium, we omit the detailed derivation here. The first-best consumption-to-q ratio is given by:

$$\begin{aligned}\frac{\mathbf{c}^*(\mathbf{s})}{\mathbf{q}^*(\mathbf{s})} =& \rho + (\psi^{-1}-1)\Big(\mathbf{i}^*(\mathbf{s}) - \delta_K - \mathcal{D}(\mathbf{s}) - \frac{\gamma\sigma_K^2}{2}\Big) \\ &+ (\psi^{-1}-1)\left(\alpha_{\mathbf{s}}(\mathbf{s};\mathbf{i}^*,\mathbf{m}^*)\epsilon_b(\mathbf{s}) + \frac{\Sigma_{\mathbf{s}}}{2}\Big(\epsilon'_b(\mathbf{s}) - \gamma\epsilon_b(\mathbf{s})^2\Big)\right) \\ &- \frac{\xi(\psi^{-1}-1)}{1-\gamma}\left(1 - \beta^\theta\left(\frac{\tilde{b}(\mathbf{s})}{b(\mathbf{s})}\right)^{1-\gamma}\right).\end{aligned}$$

Similar to Equation (20) in the decentralized market, this ratio is determined by the planner's impatience, economic growth, precautionary savings driven by carbon dynamics, and the direct impact of the present bias.

Then, substituting (B.9) through (B.12) into (B.8) and using the equilibrium condition $\mathbb{E}\left[\frac{d\mathbb{M}_t^*}{\mathbb{M}_t^*}\right] = -r^*(\mathbf{s}_t)dt$, we obtain the first-best risk-free rate:

$$\begin{aligned}r^*(\mathbf{s}) =& \rho + \psi^{-1}\Big(\mathbf{i}^*(\mathbf{s}) - \delta_K - \mathcal{D}(\mathbf{s})\Big) - \frac{\gamma(\psi^{-1}+1)}{2}\sigma_K^2 \\ &- \mu_{\mathbf{s}}(\mathbf{s};\mathbf{i}^*,\mathbf{m}^*)\big((1-\psi^{-1})\epsilon_b(\mathbf{s}) - \epsilon_{\mathbf{q}^*}(\mathbf{s})\big) - (\sigma_K^2 - \vartheta\sigma_S\sigma_K)\big(\psi^{-1}(1-\gamma)\epsilon_b(\mathbf{s}) - \gamma\epsilon_{\mathbf{q}^*}(\mathbf{s})\big) \\ &- \frac{\Sigma_{\mathbf{s}}}{2}\Big((1-\psi^{-1})\big(\epsilon'_b(\mathbf{s}) - \gamma\epsilon_b(\mathbf{s})^2\big) - (1-\gamma)2\epsilon_b(\mathbf{s})\epsilon_{\mathbf{q}^*}(\mathbf{s}) - \epsilon'_{\mathbf{q}^*}(\mathbf{s}) + 2\epsilon_{\mathbf{q}^*}(\mathbf{s})^2\Big) \\ &+ \xi\left(1 - \beta^\theta\left(\frac{\tilde{b}(\mathbf{s})}{b(\mathbf{s})}\right)^{\psi^{-1}-\gamma} + \frac{\psi^{-1}-\gamma}{1-\gamma}\left(\beta^\theta\left(\frac{\tilde{b}(\mathbf{s})}{b(\mathbf{s})}\right)^{1-\gamma} - 1\right)\right).\end{aligned} \tag{B.13}$$

Finally, using the equilibrium condition $\mathbb{E}\left[\frac{d\langle\mathbb{M}_t^*, \mathbf{Q}_t^*\rangle}{\mathbb{M}_t^*\mathbf{Q}_t^*}\right]\frac{1}{dt} = -rp^*(\mathbf{s}_t)$, the first-best risk premium

$rp^*(\mathbf{s})$ is given by:

$$\begin{aligned} rp^*(\mathbf{s}) =& \Big(\sigma_K\big(1-\epsilon_{\mathbf{q}^*}(\mathbf{s})\big)+\vartheta\sigma_S\epsilon_{\mathbf{q}^*}(\mathbf{s})\Big)\eta^K_{\mathbb{M}^*}(\mathbf{s})+\Big(\sigma_S\epsilon_{\mathbf{q}^*}(\mathbf{s})+\vartheta\sigma_K\big(1-\epsilon_{\mathbf{q}^*}(\mathbf{s})\big)\Big)\eta^S_{\mathbb{M}^*}(\mathbf{s}) \\ &-\xi\left(\beta^\theta\left(\frac{\tilde{b}(\mathbf{s})}{b(\mathbf{s})}\right)^{\psi^{-1}-\gamma}-1\right)\left(\frac{\tilde{\mathbf{q}}^*(\mathbf{s})}{\mathbf{q}^*(\mathbf{s})}-1\right). \end{aligned} \tag{B.14}$$

The terms in the above expression are determined as follows. First, $\mathbf{Q}^*_t = \mathbf{q}^*(\mathbf{s}_t)\mathbf{K}_t$ denotes the shadow aggregate firm value. Because $\mathbf{Q}^*_t$ shares the exact same functional form as its market economy counterpart, its return dynamics follow (A.8), with $\mathbf{q}^*(\mathbf{s})$ replacing $\mathbf{q}(\mathbf{s})$. Here, the elasticity $\epsilon_{\mathbf{q}^*}(\mathbf{s})$ follows the generic definition provided in (19).

Second, the planner's SDF $\mathbb{M}^*_t$ implies the shadow market prices of risk. By matching the diffusion terms of the planner's SDF dynamics with the standard SDF representation, we obtain the first-best market prices of capital risk and carbon risk, respectively:

$$\eta^K_{\mathbb{M}^*}(\mathbf{s}) = \gamma\sigma_K + \sigma_K\epsilon_{h^*}(\mathbf{s}), \quad \text{and} \quad \eta^S_{\mathbb{M}^*}(\mathbf{s}) = -\sigma_S\epsilon_{h^*}(\mathbf{s}), \tag{B.15}$$

where $\epsilon_{h^*}(\mathbf{s})$ is defined in (B.11). Notice that these shadow risk prices are structurally isomorphic to those in the decentralized economy (A.26), evaluated at the first-best allocations.

Finally, the jump component in the risk premium depends on the continuation shadow firm value, $\tilde{\mathbf{q}}^*(\mathbf{s})$. Following the exact same logic as in Step 5 of Appendix A.3, $\tilde{\mathbf{q}}^*(\mathbf{s})$ is priced using the planner's continuation SDF $\tilde{\mathbb{M}}^*_t$. It satisfies a continuation valuation equation structurally identical to (A.33), evaluated at the first-best policies $(\mathbf{i}^*, \mathbf{m}^*)$ and the first-best continuation risk-free rate $\tilde{r}^*(\mathbf{s})$. Because the continuation game features no further preference shocks ($\xi = 0$), $\tilde{r}^*(\mathbf{s})$ shares the same functional form as the current first-best risk-free rate (B.13) but without the jump term, given by:

$$\begin{aligned} \tilde{r}^*(\mathbf{s}) =& \rho+\psi^{-1}\Big(\mathbf{i}^*(\mathbf{s})-\delta_K-\mathcal{D}(\mathbf{s})\Big)-\frac{\gamma(\psi^{-1}+1)}{2}\sigma_K^2 \\ &-\mu_{\mathbf{s}}(\mathbf{s};\mathbf{i}^*,\mathbf{m}^*)\big((1-\psi^{-1})\epsilon_{\tilde{b}}(\mathbf{s})-\epsilon_{\tilde{\mathbf{q}}^*}(\mathbf{s})\big)-(\sigma_K^2-\vartheta\sigma_S\sigma_K)\big(\psi^{-1}(1-\gamma)\epsilon_{\tilde{b}}(\mathbf{s})-\gamma\epsilon_{\tilde{\mathbf{q}}^*}(\mathbf{s})\big) \\ &-\frac{\Sigma_{\mathbf{s}}}{2}\Big((1-\psi^{-1})\big(\epsilon'_{\tilde{b}}(\mathbf{s})-\gamma\epsilon_{\tilde{b}}(\mathbf{s})^2\big)-(1-\gamma)2\epsilon_{\tilde{b}}(\mathbf{s})\epsilon_{\tilde{\mathbf{q}}^*}(\mathbf{s})-\epsilon'_{\tilde{\mathbf{q}}^*}(\mathbf{s})+2\epsilon_{\tilde{\mathbf{q}}^*}(\mathbf{s})^2\Big). \end{aligned}$$

B.3 Optimal Taxation in Market Economy Restores First-Best

In this Appendix, we demonstrate that the climate policy mix—comprising a carbon tax τ^c_t, an investment subsidy τ^i_t, and a capital rebate τ^{reb}_t—implements the first-best allocation in the decentralized market economy. The formal proof proceeds in four steps.

Step 1. Given $\mathbf{i}(\mathbf{s}) = \mathbf{i}^*(\mathbf{s})$ and $\mathbf{m}(\mathbf{s}) = \mathbf{m}^*(\mathbf{s})$, we show that $u(\mathbf{s})\mathbf{q}(\mathbf{s}) = b(\mathbf{s})$. The assumption that $\mathbf{i}(\mathbf{s}) = \mathbf{i}^*(\mathbf{s})$ and $\mathbf{m}(\mathbf{s}) = \mathbf{m}^*(\mathbf{s})$ holds true under the optimal policy and will be verified in the next step. We start with the representative household's FOC for consumption (A.15) and rewrite

it as

$$\begin{aligned} u(\mathbf{s})\mathbf{q}(\mathbf{s}) &= \mathbf{c}(\mathbf{s})^{\frac{1}{1-\psi}}\big(\rho\mathbf{q}(\mathbf{s})\big)^{-\frac{\psi}{1-\psi}} \\ &= \mathbf{c}(\mathbf{s})^{\frac{1}{1-\psi}}\Big(\rho\big(1+\phi'(\mathbf{i})\big)\big(1-\tau^i(\mathbf{s})\big)\Big)^{-\frac{\psi}{1-\psi}} \\ &= \mathbf{c}^*(\mathbf{s})^{\frac{1}{1-\psi}}\left(\rho\frac{1+\phi'(\mathbf{i}^*)}{1-\epsilon_b(\mathbf{s})}\right)^{-\frac{\psi}{1-\psi}} \\ &= b(\mathbf{s}), \end{aligned} \tag{B.16}$$

where the second equality substitutes the firm's investment FOC given the subsidy $\tau^i(\mathbf{s})$. The third equality uses the fact that $\mathbf{c}(\mathbf{s}) = A - \mathbf{i}^*(\mathbf{s}) - \phi(\mathbf{i}^*) - \mathbf{m}^*(\mathbf{s})$ under the optimal policy and the assumption that $\mathbf{i}(\mathbf{s}) = \mathbf{i}^*(\mathbf{s})$ and $\mathbf{m}(\mathbf{s}) = \mathbf{m}^*(\mathbf{s})$, then by the aggregate resource constraint, $\mathbf{c}(\mathbf{s}) = \mathbf{c}^*(\mathbf{s})$. Moreover, this equality also uses the formula for the optimal investment subsidy. The final equality follows from the planner's investment FOC (B.1). Consequently, $u(\mathbf{s})\mathbf{q}(\mathbf{s}) = b(\mathbf{s})$, meaning the household's certainty equivalent of capital coincides with that of the planner.

Step 2. We show that if $\tilde{u}(\mathbf{s})\tilde{\mathbf{q}}(\mathbf{s}) = \tilde{b}(\mathbf{s})$, then $\mathbf{i}(\mathbf{s}) = \mathbf{i}^*(\mathbf{s})$ and $\mathbf{m}(\mathbf{s}) = \mathbf{m}^*(\mathbf{s})$. Combining the current household's HJB equation (A.16) with (A.29) and (A.30), the HJB equation can be expressed as

$$\begin{aligned} 0 =& \frac{1}{1-\psi^{-1}}\left(\frac{\mathbf{c}(\mathbf{s})}{\mathbf{q}(\mathbf{s})} - \rho\right) + \mathbf{i}(\mathbf{s}) - \delta_K - \mathcal{D}(\mathbf{s}) - \frac{\gamma}{2}\sigma_K^2 + \Big(\mu_{\mathbf{s}}(\mathbf{s};\mathbf{i},\mathbf{m}) - (1-\gamma)(\sigma_K^2 - \vartheta\sigma_K\sigma_S)\Big)(\epsilon_u + \epsilon_{\mathbf{q}}) \\ &+ \frac{\Sigma_{\mathbf{s}}}{2}\Big((1-\gamma)2\epsilon_u\epsilon_{\mathbf{q}} + \epsilon'_{\mathbf{q}} - \gamma\epsilon_{\mathbf{q}}^2 + \epsilon'_u - \gamma\epsilon_u^2\Big) + \frac{\xi}{1-\gamma}\left(\beta^\theta\left(\frac{\tilde{u}(\mathbf{s})}{u(\mathbf{s})}\frac{\tilde{\mathbf{q}}(\mathbf{s})}{\mathbf{q}(\mathbf{s})}\right)^{1-\gamma} - 1\right). \end{aligned} \tag{B.17}$$

Given $u(\mathbf{s})\mathbf{q}(\mathbf{s}) = b(\mathbf{s})$, the elasticities satisfy (i) $\epsilon_b = \epsilon_u + \epsilon_{\mathbf{q}}$ and (ii) $\epsilon'_b = (\epsilon_u+\epsilon_{\mathbf{q}})^2 + \epsilon'_u - \epsilon_u^2 + \epsilon'_{\mathbf{q}} - \epsilon_{\mathbf{q}}^2$. Combining these relations gives:

$$\epsilon'_b - \gamma\epsilon_b^2 = (1-\gamma)2\epsilon_u\epsilon_{\mathbf{q}} + \epsilon'_u - \gamma\epsilon_u^2 + \epsilon'_{\mathbf{q}} - \gamma\epsilon_{\mathbf{q}}^2.$$

Therefore, the current household's HJB equation (B.17) is identical to the current planner's HJB equation (B.5). In addition, the carbon tax aligns the firm's FOC for mitigation with the planner's FOC for mitigation. Together, these conditions imply $\mathbf{i}(\mathbf{s}) = \mathbf{i}^*(\mathbf{s})$ and $\mathbf{m}(\mathbf{s}) = \mathbf{m}^*(\mathbf{s})$.

Step 3. Given $\mathbf{i}(\mathbf{s}) = \mathbf{i}^*(\mathbf{s})$ and $\mathbf{m}(\mathbf{s}) = \mathbf{m}^*(\mathbf{s})$, we show that $\tilde{u}(\mathbf{s})\tilde{\mathbf{q}}(\mathbf{s}) = \tilde{b}(\mathbf{s})$. To establish the result, we compare the current household's continuation utility (A.20) with the current planner's continuation value (B.6). Note that the expression (A.30) also applies to the continuation game in the market economy. That is,

$$\begin{aligned} \Delta(\tilde{u},\tilde{\mathbf{q}}) =& \mathbf{i}(\mathbf{s}) - \delta_K - \mathcal{D}(\mathbf{s}) - \frac{\gamma}{2}\sigma_K^2 + \Big(\mu_{\mathbf{s}}(\mathbf{s};\mathbf{i},\mathbf{m}) - (1-\gamma)(\sigma_K^2 - \vartheta\sigma_K\sigma_S)\Big)(\epsilon_{\tilde{u}} + \epsilon_{\tilde{\mathbf{q}}}) \\ &+ \frac{\Sigma_{\mathbf{s}}}{2}\Big((1-\gamma)2\epsilon_{\tilde{u}}\epsilon_{\tilde{\mathbf{q}}} + \epsilon'_{\tilde{\mathbf{q}}} - \gamma\epsilon_{\tilde{\mathbf{q}}}^2 + \epsilon'_{\tilde{u}} - \gamma\epsilon_{\tilde{u}}^2\Big). \end{aligned}$$

Combined with relation (B.16), equations (A.20) and (B.6) coincide. Hence, $\tilde{u}(\mathbf{s})\tilde{\mathbf{q}}(\mathbf{s}) = \tilde{b}(\mathbf{s})$.

Step 4. Show that $r(\mathbf{s}) = r^*(\mathbf{s})$. In this final step, we establish the equilibrium risk-free rate $r(\mathbf{s})$ is identical to the first-best risk-free rate $r^*(\mathbf{s})$. Comparing (A.31) with (B.13), the first lines in both

expressions (capturing impatience, the economic growth rate, and the precautionary saving motive) are identical. The second and third lines also match, given $u(\mathbf{s})\mathbf{q}(\mathbf{s}) = b(\mathbf{s})$ and $\mathbf{q}(\mathbf{s}) = \mathbf{q}^*(\mathbf{s})$. Thus, we focus on the fourth term, which captures present bias. Specifically, the present-bias term in (A.31) can be decomposed as:[28]

$$\xi\left(\text{terms identical to the first best} + \beta^\theta\left(\left(\frac{\tilde{u}(\mathbf{s})\tilde{\mathbf{q}}(\mathbf{s})}{u(\mathbf{s})\mathbf{q}(\mathbf{s})}\right)^{\psi^{-1}-\gamma}\frac{\tilde{\mathbf{q}}(\mathbf{s})}{\mathbf{q}(\mathbf{s})} - \left(\frac{\tilde{u}(\mathbf{s})\tilde{\mathbf{q}}(\mathbf{s})}{u(\mathbf{s})\mathbf{q}(\mathbf{s})}\right)^{1-\gamma}\right)\right),$$

and therefore, $r(\mathbf{s}) = r^*(\mathbf{s})$ holds if and only if:

$$\left(\frac{\tilde{u}(\mathbf{s})\tilde{\mathbf{q}}(\mathbf{s})}{u(\mathbf{s})\mathbf{q}(\mathbf{s})}\right)^{1-\psi^{-1}} = \frac{\tilde{\mathbf{q}}(\mathbf{s})}{\mathbf{q}(\mathbf{s})}. \tag{B.18}$$

We argue that this condition holds in the competitive equilibrium. From the current household's continuation utility (A.18) and (A.19), we have

$$(\tilde{\mathbf{q}}(\mathbf{s})\tilde{u}(\mathbf{s}))^{1-\psi^{-1}} = \frac{\rho\tilde{\mathbf{c}}(\mathbf{s})^{1-\psi^{-1}}}{\rho - (1-\psi^{-1})\Delta(\tilde{u},\tilde{\mathbf{q}})}. \tag{B.19}$$

Applying the household's consumption FOC (A.15) and imposing the stationary MPE condition ($\tilde{\mathbf{c}}(\mathbf{s}) = \mathbf{c}(\mathbf{s})$), Equation (B.19) simplifies to:

$$\left(\frac{\tilde{u}(\mathbf{s})\tilde{\mathbf{q}}(\mathbf{s})}{u(\mathbf{s})\mathbf{q}(\mathbf{s})}\right)^{1-\psi^{-1}} = \frac{1}{\mathbf{q}(\mathbf{s})}\frac{\mathbf{c}(\mathbf{s})}{\rho - (1-\psi^{-1})\Delta(\tilde{u},\tilde{\mathbf{q}})}.$$

Because the continuation firm value $\tilde{\mathbf{q}}(\mathbf{s})$ solves (A.33) with the discount rate $\tilde{r}(\mathbf{s})$, it satisfies $\tilde{\mathbf{q}}(\mathbf{s}) = \mathbf{c}(\mathbf{s})/\big(\rho - (1-\psi^{-1})\Delta(\tilde{u},\tilde{\mathbf{q}})\big)$.[29] Substituting this result into the right-hand side confirms that condition (B.18) holds. We thus conclude that $r(\mathbf{s}) = r^*(\mathbf{s})$.

C Additional Results

C.1 Insensitivity of Optimal Policies to the Present-Bias Factor

As a robustness check, we examine how the quasi-hyperbolic discount factor β influences the design of optimal climate policies. Consistent with the findings in Section 5.3 regarding the planning horizon $1/\xi$, Figure 7 demonstrates that the levels of the optimal carbon tax and investment subsidy are insensitive to variations in β.

[28] The terms identical to the first-best allocation are $1 - \beta^\theta\left(\frac{\tilde{u}(\mathbf{s})\tilde{\mathbf{q}}(\mathbf{s})}{u(\mathbf{s})\mathbf{q}(\mathbf{s})}\right)^{\psi^{-1}-\gamma} + \frac{\psi^{-1}-\gamma}{1-\gamma}\left(\beta^\theta\left(\frac{\tilde{u}(\mathbf{s})\tilde{\mathbf{q}}(\mathbf{s})}{u(\mathbf{s})\mathbf{q}(\mathbf{s})}\right)^{1-\gamma} - 1\right)$, given $u(\mathbf{s})\mathbf{q}(\mathbf{s}) = b(\mathbf{s})$ and $\tilde{u}(\mathbf{s})\tilde{\mathbf{q}}(\mathbf{s}) = \tilde{b}(\mathbf{s})$.

[29] The detailed derivation mirrors the steps from (A.27) to (A.31), which relate $\mathbf{q}$ and $\Delta(u,\mathbf{q})$ to the equilibrium risk-free rate $r(\mathbf{s})$, but with $\xi = 0$.

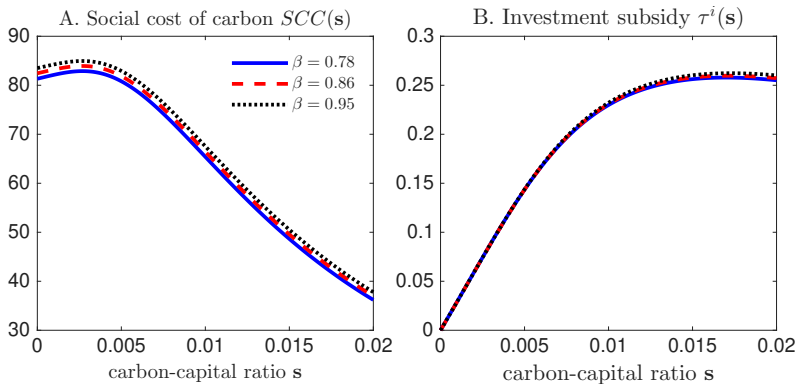


Figure 7: Robustness of optimal climate policies to the present-bias factor β. The optimal carbon tax (SCC, Panel A) and the optimal investment subsidy (Panel B) under $\beta = 0.78$ (solid lines), 0.86 (dashed lines), and 0.95 (dotted lines).

C.2 Temporal Resolution of Risk and Value of Commitment

This section examines the welfare implications of commitment through the lens of the temporal resolution of risk. As discussed in Section 6.2, recursive preferences disentangle the RRA (γ) from the EIS (ψ). When $\gamma > 1/\psi$, the planner exhibits a preference for the early resolution of risk, motivating her to tackle long-run climate threats by increasing current mitigation and saving. However, present bias weakens this precautionary motive by placing excessive weight on immediate consumption. The value of commitment, therefore, measures the welfare gain from correcting this present-biased distortion and realigning decisions with the preference for early resolution.

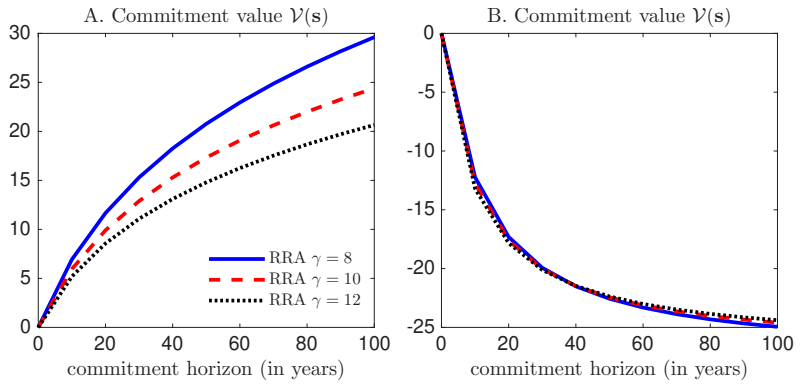


Figure 8: Effects of relative risk aversion on commitment value. Commitment value $\mathcal{V}(\mathbf{s})$ across commitment horizons under high EIS ($\psi = 1.4$, Panel A) and low EIS ($\psi = 0.7$, Panel B). Solid, dashed, and dotted lines correspond to the RRA $\gamma = 8$, 10, and 12.

Figure 8 depicts the commitment value across different levels of RRA γ. In the high-EIS regime (Panel A), the value of commitment increases with γ for any given horizon. Intuitively, a high EIS implies a strong willingness to substitute consumption over time, which lowers the welfare cost of shifting resources to the future. As γ increases, the precautionary motive associated with the preference for early resolution of risk strengthens. Consequently, a commitment device effectively forces present-biased households to curb current consumption, leading to higher equilibrium investment. The resulting larger capital buffer lowers the carbon-capital ratio and mitigates future climate risk.

Conversely, in the low-EIS regime (Panel B), increasing the RRA γ has a negligible impact on the commitment value. Intuitively, a low EIS implies a strong desire for consumption smoothing, making any deviation from a smooth consumption path highly costly in terms of welfare. Although higher risk aversion strengthens the preference for early resolution of risk, this effect is offset by the high cost of consumption fluctuations. Consequently, the dominant smoothing motive restricts equilibrium adjustments in consumption and investment, leaving the commitment value largely insensitive to variations in risk aversion.

C.3 Quantitative Results under Convex Damages

This section provides the quantitative results for the alternative convex damage specification discussed in Section 6.3. We set the sensitivity parameter $\delta_1 = 100$ and the curvature parameter $\delta_2 = 2000$. These parameters are calibrated to induce a sharp acceleration in climate damages at moderate carbon levels, while setting the maximum capital depreciation penalty at a plausible bound of 5% (δ_1/δ_2). Although the qualitative insights from our baseline model remain robust, the convex specification—which captures potential tipping points—significantly amplifies the magnitude of optimal policy interventions and their associated welfare gains.

First, Figure 9 illustrates the optimal carbon tax (SCC) and investment subsidy under convex damages. The SCC exhibits a hump-shaped pattern (Panel A): initially rising with the carbon-capital ratio $\mathbf{s}$, peaking, and eventually declining. Furthermore, the optimal investment subsidy is uniformly higher than in the concave baseline (Panel B) and increases monotonically with the carbon-capital ratio.

Second, Figure 10 depicts the welfare gains from implementing the optimal climate policies. Under the convex specification, the aggregate welfare gain is approximately 2.4%, compared to about 1.2% in the concave case (Panel A). A similar pattern holds when the carbon tax or the investment subsidy is considered in isolation (Panels C and D).

Finally, Figure 11 shows that introducing convex damages increases the value of commitment across different EIS regimes. Specifically, the commitment value becomes more positive in the high-EIS regime (Panel A) and less negative in the low-EIS regime (Panel B). Because convexity amplifies future climate risks, the welfare losses resulting from under-saving and under-mitigation become much larger. Consequently, the marginal benefit of a commitment device is significantly enhanced under the convex damage.

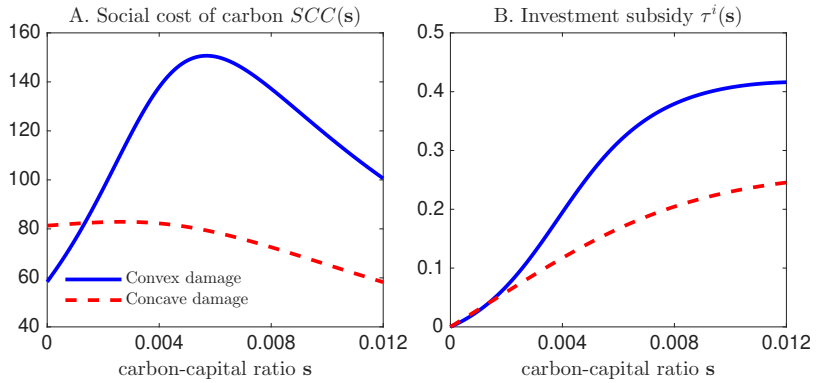


Figure 9: Optimal climate policies under convex damage. The optimal carbon tax (SCC, Panel A) and the optimal investment subsidy (Panel B) under convex (solid lines) and concave (dashed lines) damage specifications.

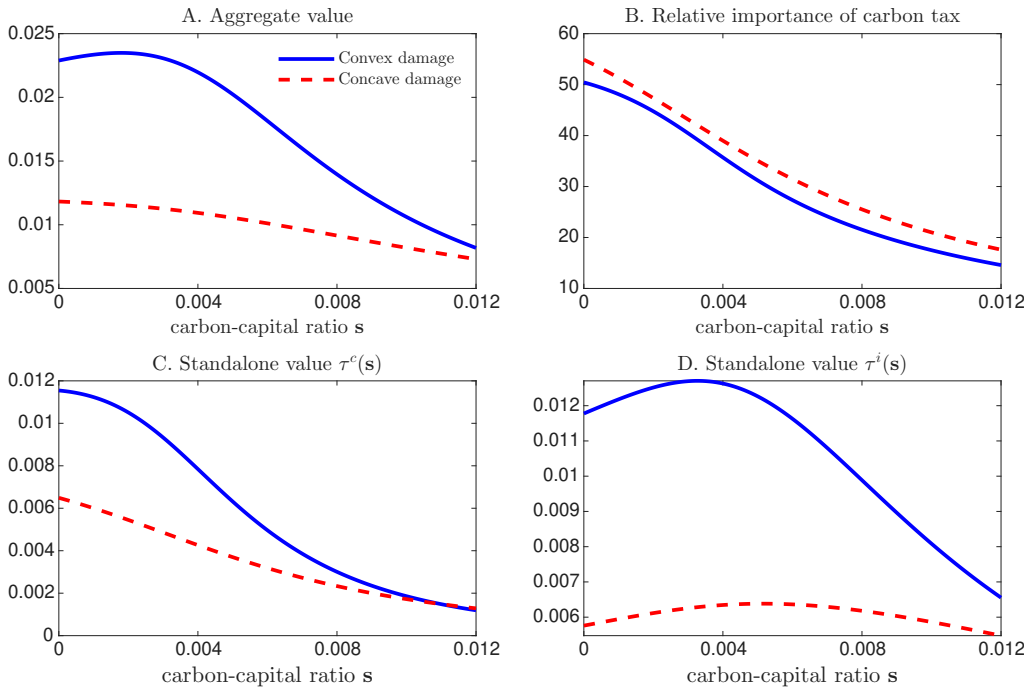


Figure 10: Welfare gains of climate policies under convex damage. Aggregate welfare gains (Panel A), relative importance of the carbon tax (Panel B), and standalone gains from the carbon tax (Panel C) and investment subsidy (Panel D) under convex (solid lines) and concave (dashed lines) damage specifications.

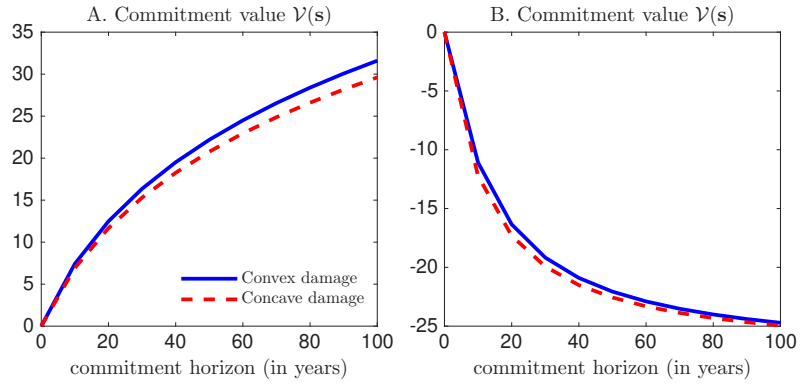


Figure 11: Value of commitment under convex damage. Commitment value $\mathcal{V}(\mathbf{s})$ across commitment horizons under high EIS ($\psi = 1.4$, Panel A) and low EIS ($\psi = 0.7$, Panel B) for convex (solid lines) and concave (dashed lines) damage specifications.

Online Appendix: Naive Beliefs in "Mitigating Climate Risk in a Present-Biased World"

Yuan Li Tak-Yuen Wong Siqi Zhao

The representative household in our baseline model is sophisticated in the sense of Krusell et al. (2002) and Gerlagh and Liski (2017): the current self correctly anticipates that future selves will remain present-biased. In this Online Appendix, we instead assume that the representative household and the social planner hold *naive* beliefs: the current self incorrectly anticipates that all future selves will act in a time-consistent manner. This change in belief formation is the sole departure from our baseline analysis. All other model ingredients—the technology and climate blocks, the continuous-time quasi-hyperbolic recursive utility framework of Shigeta (2022) (with the recursive utility curvature defined as $\theta \equiv (1-\gamma)/(1-\psi^{-1})$), and the baseline calibration—remain identical to those in the main text. Consequently, the economy is still characterized by the single aggregate state variable, the carbon-capital ratio $\mathbf{s} \equiv \mathbf{S}/\mathbf{K}$. Accordingly, the purpose of this appendix is not to re-derive the full model from scratch, but rather to isolate how naive beliefs alter the decentralized equilibrium, optimal policy design, and the welfare implications of commitment relative to our sophisticated benchmark.

OA.1 Belief Formation and Value Functions

Mathematically, this departure alters the interpretation of the expected value change induced by the arrival of a future self. In the sophisticated case, the current self correctly anticipates the Markov Perfect Equilibrium (MPE) strategies, making the continuation values the fixed points of an intrapersonal game. Under naivete, however, the current self evaluates continuation objects as if future choices were governed by a standard time-consistent economy without additional present-bias distortions. In the Hamilton-Jacobi-Bellman (HJB) equations for the decentralized household and the planner, the jump terms are respectively captured by:

$$\xi\left(\beta^{\theta}\tilde{J}(\tilde{W},\mathbf{s})-J(W,\mathbf{s})\right) \quad \text{and} \quad \xi\left(\beta^{\theta}\tilde{V}(\mathbf{K},\mathbf{S})-V(\mathbf{K},\mathbf{S})\right). \tag{OA.1}$$

While the functional forms of these jump terms remain identical to those in the main text, the continuation objects $\tilde{J}$ and $\tilde{V}$ are no longer the equilibrium outcomes generated by future present-biased selves, but rather the current self's naive beliefs.

Despite this shift in belief formation, the structural properties of the economy remain unchanged. Because the underlying technology and climate dynamics are identical, the value functions preserve their homogeneity of degree $1-\gamma$ in wealth and aggregate capital. The key objects

continue to take the form:

$$J(W,\mathbf{s}) = \frac{\left[u(\mathbf{s})W\right]^{1-\gamma}}{1-\gamma}, \qquad \tilde{J}(\tilde{W},\mathbf{s}) = \frac{\left[\tilde{u}(\mathbf{s})\tilde{W}\right]^{1-\gamma}}{1-\gamma}, \tag{OA.2}$$

and

$$V(\mathbf{K},\mathbf{S}) = \frac{\left(b(\mathbf{s})\mathbf{K}\right)^{1-\gamma}}{1-\gamma}, \qquad \tilde{V}(\mathbf{K},\mathbf{S}) = \frac{\left(\tilde{b}(\mathbf{s})\mathbf{K}\right)^{1-\gamma}}{1-\gamma}. \tag{OA.3}$$

In summary, the physical environment and the current decisions are identical to the sophisticated benchmark. The only difference is how the current self evaluates continuation wealth and continuation utility.

OA.2 Market Equilibrium under Naive Beliefs

The decentralized market equilibrium under naive beliefs preserves the exact same structure as the sophisticated benchmark. Households optimize consumption and portfolio allocations, while firms choose investment and mitigation. The economy remains fully characterized by the aggregate carbon-capital ratio $\mathbf{s}$. The household's optimal consumption rule is given by:

$$C(W,\mathbf{s}) = \rho^{\psi} u(\mathbf{s})^{1-\psi} W, \tag{OA.4}$$

and the first-order condition (FOC) for the firm's investment remains:

$$q(\mathbf{s}) = 1 + \phi'\left(i(\mathbf{s})\right). \tag{OA.5}$$

Crucially, while these FOCs retain their identical analytical forms, the equilibrium objects $u(\mathbf{s})$ (which governs the marginal propensity to consume) and $q(\mathbf{s})$ (Tobin's average q) are quantitatively distinct from the sophisticated case. This divergence arises because the jump terms triggered by the arrival of a future self are now evaluated under naive beliefs. Consequently, naivete changes how households value future capital and climate risks, altering current decisions without introducing new mechanical wedges.

Regarding climate risk mitigation, the public-good nature of emission abatement remains unchanged. Because individual firms are infinitesimal, they take aggregate mitigation as given and fail to internalize the social benefits of reduced climate damages. As a result, the classic free-rider problem persists, and the laissez-faire equilibrium continues to feature zero private mitigation:

$$m(\mathbf{s}) = 0. \tag{OA.6}$$

In sum, the decentralized economy under naive beliefs retains the baseline environmental externality. The behavioral friction of naivete operates exclusively through the misvaluation of continuation states rather than through direct alterations to the physical or market environment. Numerically, the solution method mirrors the main text, requiring solely the substitution of naive continuation values into the jump components of the valuation equations.

OA.3 The Naive Planner's Problem

OA.3.1 The First-Best Allocation

The naive social planner inherits the exact same recursive, present-biased preferences as the representative household. However, unlike the sophisticated planner in the baseline economy, the naive planner evaluates continuation utility under the false premise that all future planners will act in a time-consistent manner. Despite this shift in belief formation, the structure of the first-best problem remains intact. The planner continues to choose aggregate consumption, investment, and mitigation to maximize current welfare, fully internalizing both the public-good nature of emission abatement and the positive externality of aggregate capital accumulation on climate risk dilution.

Consequently, the planner's FOCs retain their analytical forms from the main text. The optimality condition for aggregate investment equates the marginal cost of capital accumulation to its marginal social benefit:

$$(1 + \Phi_I(\mathbf{I}, \mathbf{K}))\, f_{\mathbf{C}}(\mathbf{C}, V) = V_{\mathbf{K}}(\mathbf{K}, \mathbf{S}). \tag{OA.7}$$

Similarly, the optimality condition for aggregate mitigation equates the forgone marginal utility of current consumption to the marginal social benefit of reduced future climate damages:

$$f_{\mathbf{C}}(\mathbf{C}, V) = -\mathbf{E}_M(\mathbf{M}, \mathbf{K})\left(-V_{\mathbf{S}}(\mathbf{K}, \mathbf{S})\right). \tag{OA.8}$$

Exploiting the homogeneity property of the value function, the first-best Tobin's q under naive beliefs can again be expressed as:

$$\mathbf{q}^*(\mathbf{s}) = \frac{1 + \phi'(\mathbf{i}^*(\mathbf{s}))}{1 - \epsilon_b(\mathbf{s})}. \tag{OA.9}$$

As in the sophisticated benchmark, the elasticity term $\epsilon_b(\mathbf{s}) \equiv \mathbf{s}b'(\mathbf{s})/b(\mathbf{s})$ captures the social value of capital in buffering the economy against climate risk. This section, therefore, serves the identical benchmark role as Section 4 in the main text, with the critical distinction that the optimal allocation is now derived under naive rather than sophisticated beliefs.

Given the identical state space and homogeneity structure, our numerical procedure mirrors the sophisticated baseline. The central equilibrium objects are the naive planner's certainty-equivalent wealth function, $b(\mathbf{s})$, and the associated optimal policy rules, $\mathbf{i}^*(\mathbf{s})$ and $\mathbf{m}^*(\mathbf{s})$. These objects feed directly into the climate policy implementation and welfare comparisons presented in the subsequent sections.

OA.3.2 Policy Implementation

The logic for decentralizing the first-best allocation under naive beliefs mirrors the sophisticated benchmark, but with a critical distinction regarding the planner's expectations of future policies. In the sophisticated benchmark, the current planner correctly anticipates that future planners will share her present bias; thus, the policy rules chosen today coincide with those expected tomorrow (i.e., a MPE). Under naivete, however, the current planner incorrectly believes that future planners will act in a time-consistent manner. Consequently, while the current planner implements the current climate policies based on the valuation under present bias $b(\mathbf{s})$, she anticipates that future climate policies will be determined by the valuation under time consistency $\tilde{b}(\mathbf{s})$ and its elasticity $\epsilon_{\tilde{b}}(\mathbf{s})$.

Despite this intertemporal wedge in policy expectations, the analytical forms of the policy instruments deployed in the current period remain identical to the sophisticated benchmark. To

implement today's optimal allocation, the planner utilizes a policy mix comprising a Pigouvian carbon tax, a capital rebate, an investment subsidy, and a lump-sum levy.

To correct the emissions externality, the optimal carbon tax $\tau^c(\mathbf{s})$ is equated to the social cost of carbon (SCC), which retains its analytical definition:

$$SCC(\mathbf{s}) = -\frac{V_{\mathbf{S}}(\mathbf{K}, \mathbf{S})}{f_{\mathbf{C}}(\mathbf{C}^*, V)} = -\frac{b'(\mathbf{s})}{\rho}\left[\frac{\mathbf{c}^*(\mathbf{s})}{b(\mathbf{s})}\right]^{\psi^{-1}}. \tag{OA.10}$$

Because emission abatement exhibits diminishing marginal returns, the planner must rebate the excess tax revenue to prevent distorting private capital accumulation. This capital rebate $\tau^{reb}(\mathbf{s})$ is given by:

$$\tau^{reb}(\mathbf{s}) = \tau^c(\mathbf{s})\mathbf{e}(\mathbf{m}^*(\mathbf{s})) - \mathbf{m}^*(\mathbf{s}). \tag{OA.11}$$

To correct the capital accumulation externality, the planner provides an investment subsidy $\tau^i(\mathbf{s})$ that aligns private investment incentives with the social optimum:

$$\tau^i(\mathbf{s}) = \frac{\epsilon_b(\mathbf{s})}{\epsilon_b(\mathbf{s}) - 1}. \tag{OA.12}$$

To balance the budget for this subsidy without distorting marginal decisions, the planner simultaneously imposes a lump-sum levy $\mathbf{L}_t = \tau^i(\mathbf{s}_t)\left(\mathbf{I}_t^* + \mathbf{\Phi}_t^*\right)$, exactly as in the main text.

Finally, to evaluate the welfare implications of these policies, we replicate the policy decomposition from the main text using two partial-implementation counterfactuals: a carbon-tax-only economy and an investment-subsidy-only economy. These counterfactuals allow us to isolate each instrument's standalone welfare contribution under naive beliefs, enabling a direct comparison with our sophisticated benchmark.

OA.4 Quantitative Comparison: Naivete versus Sophistication

In this section, we quantitatively assess the implications of naivete by contrasting it against our sophisticated benchmark. We characterize this comparison around three primary dimensions: (i) the optimal levels of the carbon tax and investment subsidy, (ii) the magnitude and composition of the welfare gains from these policies, and (iii) the value of an external commitment device. Importantly, whether households are naive or sophisticated does not alter our qualitative results; the distinction between these belief structures yields only quantitative differences.

OA.4.1 Optimal Climate Policies

Figure OA.1 compares the optimal carbon tax and investment subsidy under naive and sophisticated beliefs. Strikingly, the optimal policies under the two beliefs are nearly identical. This demonstrates that the optimal levels of both instruments are remarkably robust to the household's belief formation.

To understand this near-irrelevance result, recall from Equation (OA.10) that the SCC is the ratio of the marginal disutility of carbon accumulation ($-V_{\mathbf{S}}$) to the marginal utility of current consumption ($f_{\mathbf{C}}$). A naive planner incorrectly believes that her future selves will be patient and accumulate more capital. This optimism generates two offsetting effects. First, it raises the perceived return on current savings, prompting the planner to consume less today and thereby increasing the marginal utility of current consumption. Second, because the current planner expects

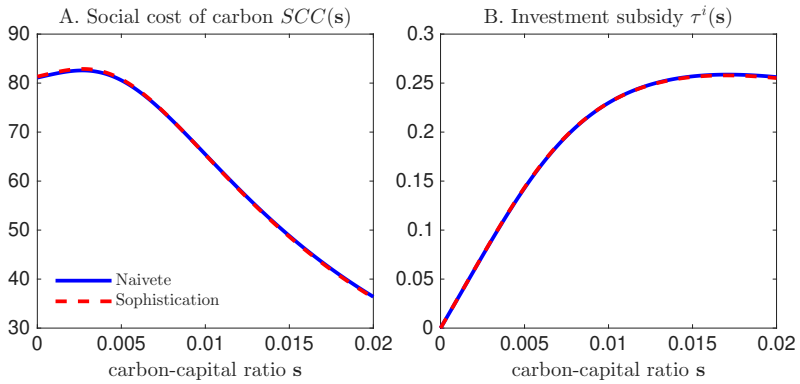


Figure OA.1: Optimal climate policies under naive and sophisticated beliefs. The optimal carbon tax $SCC(s)$ (Panel A) and the optimal investment subsidy $\tau^i(s)$ (Panel B). Solid and dashed lines correspond to naive and sophisticated beliefs, respectively.

a much larger future capital stock, an extra unit of carbon will destroy more future wealth through accelerated depreciation. This amplifies the marginal disutility of carbon accumulation. These two effects cancel each other out, leaving the optimal carbon tax largely unchanged. A similar mechanism leaves the optimal investment subsidy equally unaffected.

This finding significantly strengthens the policy implications derived in the main text. In Section 5.3, we demonstrate that optimal climate policies are remarkably robust to the *degree* of present bias. Here, we state that these policies are also robust to the *awareness* of this bias. This dual insensitivity implies that policymakers can rely on standard time-consistent frameworks when designing climate policies, without needing to precisely calibrate the magnitude of present bias or determine whether households hold naive or sophisticated beliefs.

OA.4.2 Welfare Gains from Climate Policies

The quantitative contrast becomes much sharper when we turn to welfare gains. Panel A of Figure OA.2 illustrates that while naivete does not alter the optimal policy levels, it significantly lowers the aggregate welfare gain from jointly implementing the carbon tax and investment subsidy.

To understand this decline, we decompose the total welfare gains into the standalone contributions of the carbon tax (Panel C of Figure OA.2) and the investment subsidy (Panel D of Figure OA.2). The decomposition reveals that while the standalone value of the carbon tax increases slightly under naivete, the welfare gain from the investment subsidy drops substantially. The intuition stems from the purpose of the subsidy, which is to correct the capital accumulation externality and encourages firms to build a physical buffer against climate risks. A sophisticated planner recognizes that future planners will under-save, making this subsidized capital buffer highly valuable today. In contrast, a naive planner incorrectly anticipates that future selves will be time-consistent and accumulate sufficient capital on their own. Consequently, she perceives current policy intervention as far less necessary, thus reducing the standalone value of the investment subsidy.

Crucially, this asymmetric impact fundamentally alters the relative importance of the two policy instruments. As shown in Panel B of Figure OA.2, the relative contribution of the carbon tax to total welfare gains increases under naive beliefs. This yields a crucial implication for dynamic policy prioritization: if policymakers face implementation constraints in an economy where households are predominantly naive, they should prioritize carbon pricing.

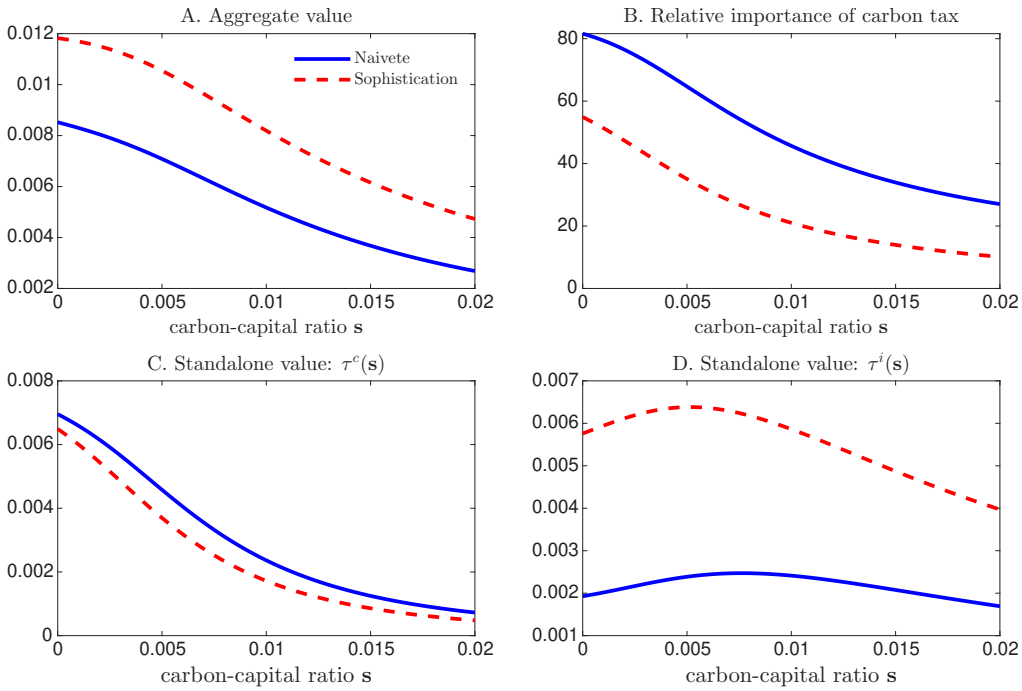


Figure OA.2: Welfare gains from climate policies under naive and sophisticated beliefs. Aggregate welfare gains from joint implementation (Panel A), relative importance of the carbon tax (Panel B), and standalone gains from the carbon tax (Panel C) and investment subsidy (Panel D). Solid and dashed lines denote naive and sophisticated beliefs, respectively.

OA.4.3 Commitment Value

The final quantitative comparison concerns the value of commitment. As in the main text, we define the commitment value $\mathcal{V}(\mathbf{s})$ as the equivalent variation in aggregate capital required to make households indifferent between a baseline economy with a 10-year horizon and a counterfactual economy with an extended horizon. Figure OA.3 compares this value under naive and sophisticated beliefs for both high-EIS and low-EIS scenarios. Qualitatively, the result remains unchanged, confirming that the tension between splurging and smoothing motives is robust to the belief specification. Quantitatively, however, naivete reduces the value of commitment. As Figure OA.3 demonstrates, naivete substantially lowers the welfare gains in the high-EIS economy (Panel A) and slightly amplifies the welfare losses in the low-EIS economy (Panel B).

The intuition behind this reduction stems from the planner's misperception of future behavior. Under a high EIS, the primary benefit of commitment arises from disciplining future selves against overconsumption. Because a naive planner incorrectly expects future selves to act patiently, she underestimates the need for such discipline, which reduces the welfare gains from commitment. Conversely, under a low EIS, households prioritize the flexibility to smooth consumption over time, and a naive planner optimistically believes that future selves will do so without intervention. In this environment, the rigid saving path imposed by commitment magnifies the perceived cost of lost flexibility, thereby deepening the overall welfare loss.

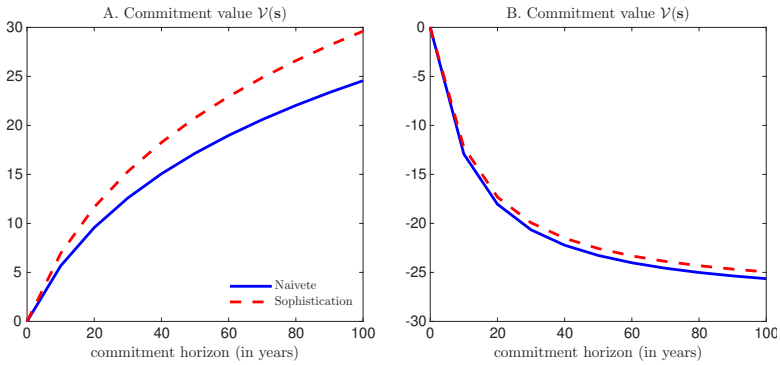


Figure OA.3: Commitment value under naive and sophisticated beliefs. Commitment value $\mathcal{V}(s)$ across commitment horizons under high EIS ($\psi = 1.4$, Panel A) and low EIS ($\psi = 0.7$, Panel B). Solid and dashed lines denote naive and sophisticated beliefs, respectively.

OA.4.4 Summary

Our quantitative analysis yields different implications for policy design and welfare evaluation. On the one hand, the optimal levels of the carbon tax and investment subsidy are robust to whether households are naive or sophisticated, confirming that policymakers can design climate interventions using standard time-consistent frameworks. On the other hand, welfare outcomes are highly sensitive to belief formation. Because a naive planner optimistically expects future selves to invest more and smooth consumption perfectly, she undervalues interventions targeting capital accumulation. Consequently, naivete reduces the welfare gains from the investment subsidy, elevates the dynamic priority of the carbon tax, and diminishes the value of commitment devices.

References

Gerlagh, Reyer and Matti Liski, "Consistent Climate Policies," *Journal of European Economic Association*, 2017, *16* (1), 1–44.

Krusell, Per, Burhanettin Kuruşçu, and Anthony A. Smith, "Equilibrium Welfare and Government Policy with Quasi-geometric Discounting," *Journal of Economic Theory*, 2002, *105* (1), 42–72.

Shigeta, Yuki, "Quasi-hyperbolic Discounting under Recursive Utility and Consumption-Investment Decisions," *Journal of Economic Theory*, 2022, *204*, 105518.