RIDGE

RESEARCH INSTITUTE FOR DEVELOPMENT, GROWTH AND ECONOMICS

SINCE 2014

THE IMPACT OF NATURAL DISASTERS ON WEALTH IN THE UNITED STATES

James B. Davies, Rodrigo Lluberas

RIDGE Working Paper 029

This paper was presented at the
2025 RIDGE December Forum

Optimal Climate Policy with Demographic Transitions

Elisa Belfiori
Universidad Torcuato Di Tella
ebelfiori@utdt.edu

Manuel Macera
Universidad Torcuato Di Tella
mmacera@utdt.edu

July 12, 2025

Abstract

This paper develops a framework to integrate demographic change into climate policy analysis. We build an overlapping-generations climate-economy model with a rich demographic structure, allowing population growth and survival probabilities to vary over time. These features endogenously determine the age composition of society and shape aggregate preferences over time. A central result of our framework is an aggregation result: despite the heterogeneity across cohorts, the planner's problem can be recast as a representative-agent model with a time-varying social discount factor that reflects demographic dynamics. Notably, as the population ages, the discount factor declines, causing society to behave as if it were more impatient. We use this framework to study how current global demographic trends affect the social cost of carbon—the model-based measure of the economic cost of climate change. We find that this cost is significantly higher under medium- and high-fertility scenarios, driven by greater emissions from larger populations. These effects are partially offset by higher discounting in aging societies. Relative to a benchmark model with no demographic change, the initial social cost of carbon rises by 20% under medium fertility and by 57% under high fertility. Global temperatures also increase more sharply in these scenarios, suggesting that demographic trends will exacerbate the climate problem.

JEL-Codes: J11, H23, Q56, Q54.

Keywords: Climate change, Demographic transition, Social Discounting, Overlapping generations, Age-based Inequality.

1 Introduction

Climate change presents one of the defining policy challenges of our time, requiring societies to weigh present-day costs against long-term benefits. At the heart of this challenge lies a fundamental intergenerational trade-off: how much should the current generation invest in curbing emissions to protect the welfare of future generations? As climate change unfolds, the world is also undergoing significant demographic shifts. The global population has grown from 3.02 billion in 1960 to 8.09 billion in 2023, with projections reaching 10.18 billion by 2100. Yet growth is slowing, and societies are aging. Annual population growth rates have fallen from 1.6% to 0.9%, and in some scenarios are projected to turn negative by the end of the century, while the share of people aged 65 and over has doubled since 1960 and is projected to reach nearly one-quarter of the global population by 2100. These changes imply that future generations will differ markedly from today's — not only in size but also in age composition. This papers studies how climate policy should adjust in light of these demographic changes.

We develop an overlapping-generations climate-economy model with population growth and age-specific survival rates that both vary over time. These features generate a time-varying age distribution that characterizes real-world demographic transitions. In the model, production requires energy, which generates carbon emissions and causes climate damages — a standard climate externality.

We study the problem from the perspective of a benevolent social planner who internalizes climate damages and takes demographic trends as given. The planner evaluates welfare using a discounted utilitarian objective, placing direct weight on each generation. We first show that the demographic structure aggregates into social preferences in a way that resembles a negative preference shock. Because a planner who values future generations implicitly places more weight on the young, population aging reduces aggregate

welfare by increasing the relative share of older individuals.

We then provide an aggregation result that plays a central role in our analysis. We show that the planner's problem in our overlapping-generations model can be written as a representative-agent problem with a time-varying discount factor. In other words, the demographic transition maps into social discounting, making society more impatient as the population ages. Because this preference structure introduces time inconsistency, we study the planner's problem under full commitment.

The aggregation result allows us to derive tractable expressions for the social cost of carbon. Demographic change affects this cost through two channels. The first is a scale effect: larger populations produce more output and generate more emissions, raising climate damages. The second, novel channel operates through discounting: population aging raises the social discount rate, reflecting increased impatience, which lowers the present value of future damages.

To formalize this mechanism, we derive a generalized Ramsey formula that shows how demographic change affects the social discount rate along demographic transitions. Population growth lowers the social discount rate by raising future consumption needs, making society more patient. However, when population growth slows over time, this effect weakens. Population aging adds upward pressure to the discount rate, further reducing the present value of future damages. With the scale effect working in the opposite direction, determining the sign and magnitude of the net impact becomes an inherently quantitative question.

We also show that the model can be extended to account for climate-demographic feedbacks, where the stock of carbon in the atmosphere affects population dynamics. In this extension, we obtain a generalized expression for the social cost of carbon that includes an additional term capturing how changes in the carbon stock influence birth and survival rates. This provides a microfoundation for utility-based damage specifications anchored

in demographic patterns.

In the quantitative analysis, we use global data from the United Nations' World Population Prospects (2024) to parametrize demographic transitions for both population growth and survival probabilities. Given the long-term nature of climate change, we require demographic forecasts that extend beyond the UN's projection horizon of 2100. To this end, we rely on observed data to parametrize the transition and use the estimated parameters to generate long-run projections. To ensure robustness, we consider three population growth scenarios that track the UN's low-, medium-, and high-fertility variants, and retain these labels throughout. For survival probabilities, we employ a parametric specification and adapt the Lee and Carter (1992) method to construct long-run forecasts that match projected life expectancy through 2100. Overall, our calibration captures both the timing and the pace of the key features of the ongoing global demographic transition. Taken together, our calibration captures the timing and pace of the most salient features of the ongoing global demographic transition.

We compare climate outcomes under each fertility scenario to a representative-agent benchmark with no demographic change. The social cost of carbon is consistently higher under medium- and high-fertility paths, with the initial cost rising by up to 57% relative to the benchmark. Temperature outcomes also worsen with higher fertility, driven by greater emissions. Even in the low-fertility scenario, aggregate emissions exceed those in the benchmark, yet the social cost of carbon remains similar. The reason lies in social discounting: although future damages are larger under low fertility, they are discounted more heavily in an aging society, which lowers their present value and keeps the social cost of carbon close to the benchmark level.

To isolate the role of aging, we consider a counterfactual in which the planner acts as if the age distribution were constant over time. Across all fertility paths, this assumption raises the social cost of carbon by removing the discounting effect of aging. A simple

back-of-the-envelope comparison illustrates the magnitude of this bias: imposing a stationary 1% population growth rate with a fixed age structure yields a peak social cost of carbon more than twice as high as that in the high-fertility scenario—while resulting in virtually identical warming. This highlights how ignoring demographic transitions can significantly overestimate the social cost of carbon, distorting the design of optimal climate policy.

Finally, we examine how the results vary with the generational discount factor—the weight placed on future generations. We consider values spanning the Nordhaus and Stern benchmarks. As expected, a lower generational discount rate raises the social cost of carbon, confirming a standard result. But it also amplifies the effects of aging: when future generations are weighted more heavily, the rise in the social discount rate induced by population aging has a greater impact. In extreme cases, low-fertility scenarios can produce higher peak temperatures than high-fertility ones, as weaker mitigation incentives in aging societies outweigh the emissions advantage of slower population growth.

Related Literature

The paper contributes to a growing literature that incorporates different layers of heterogeneity into the representative-agent integrated assessment neoclassical growth model pioneered by Nordhaus and Boyer (2003) and further developed by Golosov et al. (2014), Barrage (2019), Van der Ploeg and Rezai (2021), and others. Recent contributions include Hillebrand and Hillebrand (2019), Krusell and Smith (2022), Belfiori and Macera (2024), and Bourany (2024), who examine inequality across countries; Fried et al. (2024), Belfiori et al. (2024), and Douenne et al. (2023), who focus on income inequality across households; and Cruz and Rossi-Hansberg (2023) and Conte et al. (2022), who analyze spatial inequality. While these papers enrich the integrated assessment framework along various dimensions, many existing models either abstract entirely from demographic dynamics,

as in Golosov et al. (2014), or incorporate only stationary population growth with a fixed age structure, as in Van der Ploeg and Rezai (2021). This paper extends the framework by introducing demographic heterogeneity driven by non-stationary transitions in both population size and age composition. We show that accounting for these dynamics alters how future climate damages are valued and has meaningful implications for climate policy.

Our approach to welfare analysis within an overlapping generations framework follows Calvo and Obstfeld (1988) and Arrow et al. (2004), as well as more recent work by Eden (2023) and Schneider et al. (2012). In this setting, we show that demographic transitions endogenously shape social discounting: as the population ages, the social discount factor declines, making society behave as if it were more impatient. Unlike in dynastic models, where the discount factor is fixed, this framework implies that the effect of aging on the valuation of future damages is amplified when the planner places greater weight on the welfare of future generations. This perspective complements and extends the influential work of Stern (2008), Nordhaus (2007), and Weitzman (2007) by linking social discounting more explicitly to demographic dynamics.

The idea that demographic change amplifies emissions connects to the population externality discussed in Bohn and Stuart (2015) and appears as well in Gerlagh et al. (2023) and Van der Ploeg and Rezai (2021). By contrast, the result that demographic change alters the valuation of future climate damages through its effect on the social discount rate is novel to this paper. We formalize this mechanism through a generalized Ramsey rule that incorporates aging, linking our analysis to the foundational work of Ramsey (1928) and the comprehensive review in Arrow et al. (2013).

This paper also contributes to the emerging literature linking climate change to demographic outcomes. Carleton et al. (2023) and references therein document how climate variability and extreme events affect fertility and mortality. We show how to accommodate such feedbacks within our framework by allowing demographic variables to depend

on the atmospheric carbon stock. In this extension, we derive a generalized expression for the social cost of carbon that includes an additional term capturing how changes in the carbon stock influence birth and survival rates. A related contribution is Bressler (2021), who develops a similar expression in a simpler setting without an explicit life-cycle structure and without modeling endogenous demographic dynamics, so our formulation nests this idea in a more general framework.

While our framework adopts a discounted utilitarian approach for analytical tractability, it shares with Barcons et al. (2025) and Aguiar et al. (2023) a concern for intergenerational welfare. Barcons et al. (2025) explores normative criteria beyond utilitarianism, offering a complementary perspective on how to weigh the welfare of different cohorts. Aguiar et al. (2023), in turn, analyzes robust Pareto improvements in decentralized environments, a direction we identify as promising for future work. These contributions, while methodologically distinct, align with the broader motivation of our paper to clarify how demographic structure shapes long-term policy evaluation.

The paper is organized as follows. Section 2 presents the model and describes the demographic dynamics. Section 3 sets up the social planning problem and derives the aggregation result. Section 4 analyzes the key demographic forces that shape the social cost of carbon. Section 5 outlines the calibration, and Section 6 presents the quantitative results. Section 7 concludes.

2 Model

Consider an overlapping generations, infinite horizon, climate-economy model. There is a single consumption good that is produced using capital, labor, and energy input according to a constant returns to scale technology $\tilde{F}(K_t, N_t^Y, E_t)$. Energy production releases carbon

into the atmosphere. The carbon stock S_t evolves according to

$$S_{t+1} = (1-\gamma)S_t + \eta E_t \tag{1}$$

where $\gamma \in [0,1)$ is the natural rate of carbon reabsorption and η captures the carbon content of energy production. We assume that the energy sector is labor intensive, with a linear production function, $E_t = \psi N_t^E$. The initial carbon stock, S_0, is given.

Carbon in the atmosphere generates a negative externality, which is modeled as an output loss. In particular, aggregate output in the economy is

$$Y_t = F(S_t, K_t, N_t^Y, E_t) = [1 - x(S_t)]\tilde{F}(K_t, N_t^Y, E_t), \tag{2}$$

where x represents the climate damage function. We assume that this function is increasing, convex and twice differentiable with $\lim_{S \to S_0} x(S) = 0$.

Aggregate feasibility requires

$$N_t^E + N_t^Y = N_t, \tag{3}$$

where N_t represents aggregate labor, and

$$C_t + K_{t+1} = F(S_t, K_t, N_t^Y, E_t) + (1-\delta_k)K_t \tag{4}$$

where δ_k denotes the depreciation rate of capital and C_t is aggregate consumption.

Demographics. Each period, a new generation is born, composed of a continuum of identical individuals. The probability that an agent born in period t survives to age h is $\pi_{h,t}$, with the convention that $\pi_{0,t} = 1$ for any t. In addition, we assume that $\pi_{h,t} = 0$ for any $h > H$, which implies that agents live for a maximum of $H+1$ periods. As usual, a law

of large numbers applies to this type of economy and, thus, the fraction of agents born in period t still alive at age h is also $\pi_{h,t}$.

Let $N_{h,t}$ denote the mass of population h-years old in period t. We then have:

$$N_{h,t+h} = \pi_{h,t} N_{0,t} \tag{5}$$

for every h, in every period t. Thus, total population is equal to

$$N_t = \sum_{h=0}^{H} \pi_{h,t-h} N_{0,t-h}$$

in each period t. We further assume that total population grows according to:

$$N_{t+1} = (1 + n_{t+1}) N_t, \tag{6}$$

where n_{t+1} is the exogenous population growth rate between period t and $t+1$.[1]

We use $\omega_t \equiv (\omega_{1,t}, \omega_{1,t}, ..., \omega_{H,t})$ to denote the age distribution in period t. The share of population aged h in period t, $\omega_{h,t}$, is given by

$$\omega_{h,t} \equiv \frac{N_{h,t}}{N_t} = \frac{\pi_{h,t-h} N_{0,t-h}}{N_t}. \tag{7}$$

Preferences. Let $\mathbf{c} = \{c_{0,t}, c_{1,t}, ..., c_{H,t}\}_{t=0}^{\infty}$ denote a consumption plan, which specifies an amount of consumption, $c_{h,t}$, for each member of the age group h in period t. The expected lifetime utility that an agent born in period τ derives from the consumption plan $\mathbf{c}$ is

$$U_\tau(\mathbf{c}) = \sum_{h=0}^{H} \beta^h \pi_{h,\tau} u(c_{h,\tau+h}) \tag{8}$$

[1] Assuming an exogenous growth rate for the entire population is without loss of generality. In Appendix A we show that, given a survival probability profile, there is a one to one mapping between choosing a growth rate for the newborns and a growth rate for the entire population.

where $\beta \in (0,1)$ denotes the private discount factor, and the instantaneous utility function is of the CRRA type, $u(c) = c^{1-\nu}/(1-\nu)$. This formulation rules out altruism across generations. All agents are endowed with one unit of labor at each age.

Aggregate consistency requires that aggregate consumption in period t is obtained from individual consumption as follows:

$$C_t = \sum_{h=0}^{H} N_{h,t} c_{h,t} \tag{9}$$

for every period t.

3 The Social Planning Problem

We adopt a discounted utilitarian social welfare function, similar to Calvo and Obstfeld (1988) and Arrow et al. (2004). The planner sets cohort-specific Pareto weights as follows:

$$W(\mathbf{c}) = \sum_{\tau=-H}^{\infty} \alpha_\tau U_\tau(\mathbf{c}), \tag{10}$$

where τ indexes a generation, and the Pareto weights are given by

$$\alpha_\tau = \frac{\delta^\tau N_{0,\tau}}{\sum_{s=-H}^{\infty} \delta^s N_{0,s}}, \tag{11}$$

for $\tau \geq -H$. The planner uses the factor δ to discount the well-being of future generations. This generational discount factor, δ, need not be equal to the private one, β, and in what follows we assume that $\delta > \beta$.[2] Note that when $\beta = 0$, the framework collapses to the standard dynastic model with no altruism. In this case, the discount factor δ governs how

[2]This assumption is consistent with standard calibrations in the literature (Schneider et al. (2012), Eden (2023)). Benchmark values for the social discount factor include $\delta \approx 0.985$ (Nordhaus (2007)) and $\delta \approx 0.999$ (Stern (2007)), whereas typical annual private discount factors imply $\beta \approx 0.96$ (e.g., Kydland and Prescott (1982); Laibson (1997); Gourinchas and Parker (2002)

the planner values welfare over time and is naturally related to the planner's pure rate of time preference.

The next result shows that it is possible to rewrite (10) in a more intuitive form, where one can think of the planner placing weights directly to the members of different cohorts at each period. The proof is in the appendix.

Lemma 1 *The Social Welfare Function can be written as*

$$W(\mathbf{c}) = \sum_{t=-\infty}^{\infty} \delta^t N_t \sum_{h=0}^{H} \left(\frac{\beta}{\delta}\right)^h \omega_{h,t} u(c_{h,t}). \tag{12}$$

It follows that the discounted utilitarian social welfare function (10), in which the planner sets weights to agents according to their date of birth, is equivalent to a social welfare function that is discounted utilitarian in the cross section and in the time series and treats agents at different points of their lifetimes as different selves.

The *socially optimal allocation* consists of a consumption plan $\mathbf{c}$ and a path for aggregate consumption, capital, labor, energy, and carbon stock, $\{C_t, K_t, N_t^Y, N_t^E, E_t, S_t\}_{t=0}^{\infty}$, that maximizes social welfare (10), subject to the carbon cycle (1), feasibility constraints (4) and (3), and aggregate consistency (9). To characterize this allocation, we divide the problem into two. We solve first the intra-temporal problem of allocating aggregate consumption in the cross-section, and then the intertemporal problem that determines the paths for aggregate consumption, capital, labor, energy and the carbon stock.

3.1 Intratemporal problem

Consider the problem of allocating resources in the cross section taking the path of aggregate consumption, $\{C_t\}_{t=0}^{\infty}$, as given. Each period, the planner must choose consumption

for each agent alive in order to solve

$$\max_{\{c_{h,t}\}} \sum_{h=0}^{H} \left(\frac{\beta}{\delta}\right)^{h} \omega_{h,t} u(c_{h,t}) \tag{13}$$

subject to (9). The first-order condition with respect to individual consumption is:

$$\left(\frac{\beta}{\delta}\right)^{h} \omega_{h,t} u'(c_{h,t}) - \Lambda_t N_{h,t} = 0 \tag{14}$$

for all h, where Λ_t is the Lagrange multiplier on the constraint (9). Using the functional form for the instantaneous utility function and equation (7), we can solve for individual consumption:

$$c_{h,t} = \left(\frac{\beta}{\delta}\right)^{h/\nu} (N_t \Lambda_t)^{-1/\nu} \tag{15}$$

Thus, aggregate consumption in any period t equals

$$C_t = (N_t \Lambda_t)^{-1/\nu} \sum_{h} N_{h,t} \left(\frac{\beta}{\delta}\right)^{h/\nu} = (N_t \Lambda_t)^{-1/\nu} N_t \theta_t \tag{16}$$

where θ_t is defined as:

$$\theta_t \equiv \sum_{h=0}^{H} \omega_{h,t} \left(\frac{\beta}{\delta}\right)^{h/\nu} \tag{17}$$

We can express individual consumption as follows:

$$c_{h,t} = \left(\frac{\beta}{\delta}\right)^{h/\nu} \frac{C_t}{\theta_t N_t} \tag{18}$$

Plugging this expression into (13) delivers the following result.

Lemma 2 ***[Aging as a Preference Shock]*** *The indirect utility function of the intratemporal social*

planning problem is given by

$$U(C_t) = \theta_t^{\nu} u(C_t/N_t). \tag{19}$$

This result highlights a key insight: age heterogeneity in the population aggregates into social welfare in a way that resembles a preference shock. The term θ_t captures how the planner accounts for shifts in the population's age composition.

Two-cohort example. To build intuition on how the age structure affects social preferences, consider the simplest overlapping generations model with only two cohorts: the young and the old. In this case,

$$\theta_t = 1 - \omega_{1,t}\left(1 - \left(\frac{\beta}{\delta}\right)^{1/\nu}\right)$$

If the demographic transition involves aging – meaning a persistent increase in the share of older individuals – then θ_t declines over time when $\delta > \beta$. Aging therefore acts as a negative preference shock from a societal perspective. A higher δ implies that the planner places greater weight on the young, so the welfare obtained from the distribution of aggregate consumption in the cross-section decreases as the share of older people rises. This logic extends to the general case with $H > 1$.

3.2 Intertemporal Problem

Using Lemma 2, it is possible to reduce the social planning problem to choosing only aggregate variables. We state this formally in the following aggregation result:

Proposition 1 (Aggregation Result) *The aggregate consumption, capital, labor, and carbon stock*

of the socially optimal allocation solve the following problem

$$\max_{C_t, N_t^Y, N_t^E, K_{t+1}, S_{t+1}} \sum_{t=0}^{\infty} \delta^t N_t^\nu \theta_t^\nu u(C_t) \tag{20}$$

subject to (1), (3) *and* (4), *with* K_0 *and* S_0 *given.*

Demographic dynamics are embedded in social preferences. While individuals discount future consumption at rate β, society as a whole applies a different social discount function: $\delta\left(\frac{N_{t+1}}{N_t}\right)^\nu\left(\frac{\theta_{t+1}}{\theta_t}\right)^\nu$. During demographic transitions, changes in population growth and age composition make these social preferences evolve over time, which can render the social planning problem dynamically inconsistent. Although this introduces interesting questions about commitment and time consistency, we do not pursue them here and instead focus on how demographic change affects the social discount rate and the social cost of carbon emissions under full commitment.[3]

4 The Social Cost of Carbon

Understanding the role of demographics in shaping social preferences is important because they directly affect the social cost of carbon (SCC)– the model-based measure of the economic cost of climate change.

To characterize this cost, let $\delta^t\sigma_t$ and $\delta^t\lambda_t$ denote the Lagrange multipliers associated with the carbon cycle (1) and the feasibility constraint (4), respectively. The first order condition with respect to the global carbon stock is:

$$\delta^t\sigma_t = (1-\gamma)\delta^{t+1}\sigma_{t+1} + \delta^{t+1}\lambda_{t+1}\frac{\partial F(S_{t+1}, K_{t+1}, N_{t+1}^Y, E_{t+1})}{\partial S_{t+1}} \tag{21}$$

[3]The solution without commitment would require solving a dynamic game across generations, where the resulting Markov equilibrium is not optimal but constrained-optimal: the best outcome achievable given the inability to commit to future policies (Kydland and Prescott (1977)).

By iterating forward on this condition, we obtain the following expression for the SCC:

Proposition 2 (The Social Cost of Carbon) *The social cost of carbon is equal to*

$$\hat{\sigma_t} \equiv \frac{\sigma_t}{\lambda_t} = \sum_{j=1}^{\infty} (1-\gamma)^{j-1} \delta^j \left(\frac{N_{t+j}}{N_t}\right)^{\nu} \left(\frac{\theta_{t+j}}{\theta_t}\right)^{\nu} \frac{u'(C_{t+j})}{u'(C_t)} \frac{\partial F(S_{t+j}, K_{t+j}, N^Y_{t+j}, E_{t+j})}{\partial S_{t+j}}. \tag{22}$$

Intuitively, the SCC is the present discounted value of future output losses caused by an extra ton of carbon emissions, where the expression accounts for the fact that carbon depreciates at a rate γ.[4] In our framework, demographic dynamics affect this cost through two main channels.

The first is a scale effect. A larger population implies greater total output, so the marginal damage of additional emissions operates over a larger base. As a result, the social cost of carbon rises with the size of the population. This effect aligns with the idea of the population externality discussed in Bohn and Stuart (2015): when economic activities generate externalities, population growth amplifies their consequences.

The second channel operates through social discounting. In Equation 22, the rate at which society discounts future climate damages from period $t+j$ back to t is given by

$$\left(\frac{1}{1+r^s_{t,t+j}}\right)^j \equiv \delta^j \left(\frac{N_{t+j}}{N_t}\right)^{\nu} \left(\frac{\theta_{t+j}}{\theta_t}\right)^{\nu} \frac{u'(C_{t+j})}{u'(C_t)} \tag{23}$$

where $r^s_{t,t+j}$ denotes the annualized social discount rate from t to $t+j$. Log-linearizing this expression yields the generalized Ramsey formula:

[4] Under standard functional form assumptions, the SCC can be expressed as a fraction of current output, as in Golosov et al. (2014). In our framework, this fraction is still determined by discounting, carbon depreciation, and climate damages. However, the discount factor now also reflects demographic dynamics. As a result, changes in population growth and age composition shape the time path of the SCC during demographic transitions, breaking its tight link to output.

Proposition 3 (Ramsey Formula) *The social discount rate between periods t and $t+j$ is:*

$$r^s_{t,t+j} = \rho + \nu(g_{t,t+j} - n_{t,t+j} - \mu_{t,t+j}) \tag{24}$$

where $\rho \equiv \delta^{-1} - 1$ is the generational discount rate, and g_t, n_t, and μ_t denote the average growth rates of consumption, population, and the age-structure term θ_t, respectively, from t to $t+j$.

The proof is in the appendix. Population growth affects social discounting through a standard channel: higher growth reduces future consumption per capita, which raises the marginal utility of future consumption and strengthens the incentive to save. This lowers the social discount rate, making society effectively more patient when population growth is high.

Demographic transitions introduce a novel channel through population aging. As shown in Lemma 2, the age-structure term θ_t effectively acts as a taste shifter that modifies how the planner weights the marginal utility of future generations. When the generational discount factor δ exceeds the private one β, the planner places greater relative weight on future cohorts. As the population ages and θ_t declines, the weighted marginal utility of future generations decreases because future cohorts contain a larger share of older individuals, which weakens the incentive to save. This raises the social discount rate, making society behave as if it is more impatient about future climate damages as aging proceeds.

As the transition unfolds, population growth slows and the downward effect of growth on the social discount rate weakens, while the aging channel pushes it upward. Taken together, these forces imply that the social discount rate rises as the demographic transition progresses, increasing the weight placed on near-term climate damages relative to long-term ones.

Overall, while population growth raises emissions through the scale effect, the rising social discount rate associated with demographic change reduces the present value of

future climate damages.[5] The net impact of demographic change on emissions and the SCC depends on the relative strength of these opposing forces—a quantitative question we turn to in the next section.

4.1 From Climate to Demographics.

Although our primary focus is on the effects of demographic change on climate outcomes, it is important to acknowledge the potential influence of climate on demographic dynamics. A growing body of research documents how climate variability and extreme events can affect fertility and mortality, thereby shaping the overall pace of population growth and age structure.[6] Our theoretical framework can readily accommodate this channel. In an extended version of the model, the demographic variables are no longer treated as exogenous but are instead modeled as functions of the atmospheric carbon stock, S_t. This specification allows demographic dynamics to respond endogenously to changes in climate, capturing potential feedback effects between the climate system and demographic processes.

The SCC must now account for the utility cost of emissions through their effect on the demographic variables that appear in the social discount factor. Consequently, an additional term arises in the first-order condition with respect to S_t to capture this channel. More precisely, (21) now becomes:

$$\delta^t \sigma_t = (1-\gamma)\delta^{t+1}\sigma_{+1} + \delta^{t+1}\lambda_{t+1}\frac{\partial F(S_{t+1}, K_{t+1}, N^Y_{t+1}, E_{t+1})}{\partial S_{t+1}} + \frac{\partial W(\mathbf{C})}{\partial S_{t+1}} \tag{25}$$

This expression now incorporates the first derivative of the social welfare function with

[5] A higher social discount rate also reduces aggregate savings, which further lowers capital accumulation and emissions, indirectly affecting the social cost of carbon.

[6] See Carleton et al. (2023) and references therein.

respect to the stock of carbon in the atmosphere, which can be expressed as follows:

$$\frac{\partial W(\mathbf{C})}{\partial S_{t+1}} = \nu \delta^{t+1} N_{t+1}^{\nu} \theta_{t+1}^{\nu} u(C_{t+1}) \left\{ \frac{\partial \left(N_{t+1} \theta_{t+1} \right)}{\partial S_{t+1}} \frac{1}{N_{t+1} \theta_{t+1}} \right\}$$

The term between curly brackets is the semi-elasticity of the demographic factors with respect to the amount of carbon in the atmosphere. Plugging this in the first-order condition and iterating forward delivers a generalized version of the SCC:

Proposition 4 (Generalized Social Cost of Carbon) *The social cost of carbon, accounting for feedback from climate to demographic variables, is given by*

$$\sigma_t^{\blacklozenge} = \hat{\sigma}_t + \nu \sum_{j=1}^{\infty} (1-\gamma)^{j-1} VSLY_{t+j} \left\{ \frac{\partial \left(N_{t+j} \theta_{t+j} \right)}{\partial S_{t+j}} \frac{1}{N_{t+j} \theta_{t+j}} \right\} \tag{26}$$

where $\hat{\sigma}$ is given by (22), and

$$VSLY_{t+j} \equiv \delta^j \left(\frac{N_{t+j}}{N_t} \right)^{\nu} \left(\frac{\theta_{t+j}}{\theta_t} \right)^{\nu} \frac{u(C_{t+j})}{u'(C_t)}. \tag{27}$$

The additional term in the expression for the social cost of carbon includes $VSLY_{t+j}$, naturally interpreted as the Value of a Statistical Life Year in period $t+j$, which measures the marginal rate of substitution between an increase in utility in year $t+j$ and consumption in period t. [7]

[7] This interpretation is consistent with the standard ex ante concept of the value of a statistical life-year (VSLY) in dynamic models. In such models, the VSLY is typically derived as the marginal rate of substitution between a small increase in the probability of surviving to period $t+j$ and a reduction in wealth (or consumption) in the current period. Under the expected utility framework, this is equivalent to the marginal value of additional utility in $t+j$ relative to the marginal utility of consumption at time t, i.e.,

$$\text{VSLY}_{t+j} = \frac{\partial V_t / \partial s_{t+j}}{\partial V_t / \partial C_t} = \frac{\delta_{t+j} u(C_{t+j})}{u'(C_t)},$$

where the discount factor δ_{t+j} implicitly includes the survival probability to period $t+j$. This interpretation preserves the logic of valuing small extensions to life within a utility-maximizing framework, while expressing the tradeoff in units of contemporaneous consumption. Bressler (2021) develops a related expression in a simpler setting without an explicit life-cycle structure and without modeling endogenous demographic

More broadly, this new term can be understood as a utility cost arising from the sensitivity of demographic fundamentals to climate. To relate this expression to empirical estimates of climate change effects on birth and mortality rates, note that:

$$N_t\theta_t = \sum_{h=0}^{H} N_{h,t}\left(\frac{\beta}{\delta}\right)^{h/v} = N_{0,t} + \sum_{h=1}^{H} N_{h,t}\left(\frac{\beta}{\delta}\right)^{h/v} = b_t N_{t-1} + \sum_{h=1}^{H} \pi_{h,t} N_{h-1,t-1}\left(\frac{\beta}{\delta}\right)^{h/v},$$

where b_t denotes the birth rate in period t. Using this, the semi-elasticity can be written as:

$$\left\{\frac{\partial\left(N_{t+j}\theta_{t+j}\right)}{\partial S_{t+j}}\frac{1}{N_{t+j}\theta_{t+j}}\right\} = \hat{\omega}_{0,t+j}\frac{\partial b_{t+j}}{\partial S_{t+j}}\frac{1}{b_{t+j}} + \sum_{h=1}^{H}\hat{\omega}_{h,t+j}\frac{\partial \pi_{h,t+j}}{\partial S_{t+j}}\frac{1}{\pi_{h,t+j}} \tag{28}$$

where

$$\hat{\omega}_{h,t+j} = \frac{N_{h,t+j}\left(\frac{\beta}{\delta}\right)^{h/v}}{\sum_{a=0}^{H} N_{a,t+j}\left(\frac{\beta}{\delta}\right)^{a/v}}, \tag{29}$$

can be interpreted as a modified population share of h-years old individuals.[8] Therefore, the semi-elasticity is a weighted average of the semi-elasticities of birth and survival rates with respect to the atmospheric carbon stock. The weight on each component depends on the relationship between β and δ. For instance, when $\beta < \delta$, relatively more weight is placed on the birth rate and survival rates at younger ages. This pattern is intuitive, as a higher δ implies greater concern for younger cohorts, which in turn places more emphasis on births and early-life survival in the social cost of carbon.

dynamics, so our formulation nests this idea in a more general framework.

[8] When $\beta < \delta$, the share of young people are up-scaled and that of old people down-scaled. The opposite occurs if $\beta > \delta$.

5 Calibration

We use our framework to quantify the impact of global demographic trends on the social cost of carbon. We first focus on the calibration of the demographic transition, which is less standard, and then summarize the climate and economic parameters.

5.1 Parameterizing Demographic Transitions.

Our analysis uses global demographic data from the United Nations' World Population Prospects (2024), hereafter WPP. This dataset provides observed demographic data up to 2023 and official projections through 2100. However, the long-term nature of climate change and its economic implications call for demographic forecasts that extend well beyond this horizon. To this end, we parametrize the transition dynamics of these variables, fit the models to observed data, and use the estimated parameters to produce long-run forecasts. The parametrization strategy is described below. Prior to modeling, all demographic data are smoothed using 10-year moving averages to reduce short-term fluctuations. We further assume that the initial data point represents a steady state, with subsequent observations capturing the unfolding demographic transition.

Population Growth Figure 1 displays historical and projected population growth rates under three fertility scenarios—high, medium, and low—as provided by the WPP. To capture the dynamics of population growth, we fit two parsimonious time series models to the data: an autoregressive model of order one (AR(1)) and a logistic specification. These are selected for their simplicity, providing robust yet tractable representations of demographic trends.

We begin with the AR(1) specification. This model assumes that the logarithm of the population growth rate, denoted n_t, follows a stationary AR(1) process with a long-run

Figure 1: Global Population Growth

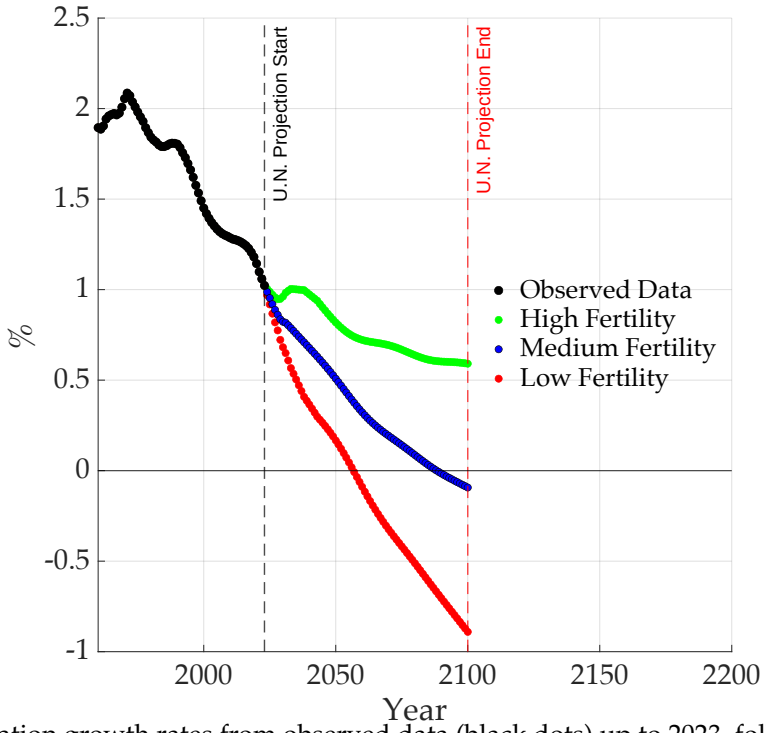

Notes: Global population growth rates from observed data (black dots) up to 2023, followed by projections under three fertility scenarios: high (green dots), medium (blue dots), and low (red dots). Vertical dashed lines mark the start and end of the United Nations' official projection period.

mean $n^\star$ and initial condition n_0. Specifically, the model takes the form

$$\log n_t = (1-\rho)\log n^\star + \rho \log n_{t-1},$$

where $\rho \in (0,1)$ governs the persistence of the process. We estimate three parameters: the persistence ρ, the initial level n_0, and the long-run mean $n^\star$.

As an alternative specification, we consider a logistic functional form to capture the transition in population growth rates. In this case, the model is given by

$$n_t = \frac{b_1}{1+\exp(-b_2(t-b_3))},$$

where b_1 represents the upper asymptote of the growth rate, b_2 controls the speed of the

Figure 2: Projected Global Population Growth

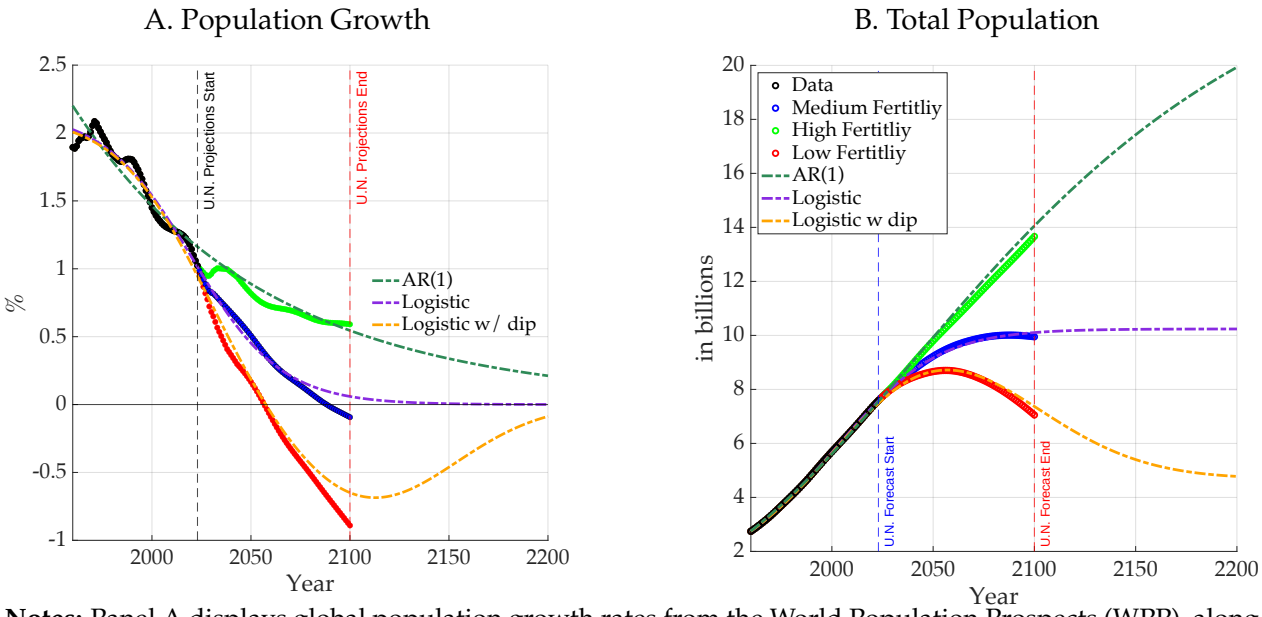

Notes: Panel A displays global population growth rates from the World Population Prospects (WPP), along with projection paths under alternative fertility scenarios. Dash-dotted lines depict the trajectories implied by three estimated processes: an AR(1) model, a logistic growth specification, and a logistic model augmented with a Gaussian dip. All models are estimated using observed data only, without incorporating projection information. Panel B shows the corresponding population levels under each specification. The Gaussian dip component is manually calibrated to approximate the dynamics of the low-fertility scenario in the WPP projections. All empirical data and projection scenarios are drawn from the WPP.

transition, and b_3 determines the inflection point—the time at which the growth rate is halfway between its initial level and the asymptote.

Both models are estimated via nonlinear least squares by minimizing the sum of squared deviations between the model-implied series and the observed data, and their fit is compared with WPP projections in Figure 2. The first panel displays population growth rates over time, showing that the AR(1) model closely tracks the high-fertility scenario, while the logistic specification aligns well with the medium-fertility projection. This is a notable outcome given that the models were estimated solely on historical data and did not incorporate WPP projections in their calibration.

To approximate the low-fertility scenario—which features persistent negative population growth—we modify the logistic specification by introducing a Gaussian-shaped dip.

The logistic parameters are re-estimated conditional on the presence of the dip, but the dip's characteristics are selected manually.[9] As shown in the second panel, which plots population levels over time, the augmented model tracks the WPP's low-fertility scenario with reasonable accuracy, offering a flexible extension to the baseline logistic form.

Our analysis assumes that population growth eventually stabilizes at a non-negative rate. This ensures that the demographic transition converges to a steady state in which a population remains. In this context, the logistic-with-dip specification offers a useful way to capture temporary episodes of population decline—such as those implied by the low-fertility scenario—while preserving long-run stability.

Throughout the remainder of the analysis, we refer to the AR(1) model as the high-fertility case, the logistic model as the medium-fertility case, and the logistic-with-dip model as the low-fertility case.

Survival Probabilities Survival probability data is organized as a matrix, where each row corresponds to a year t and each column to an age h.[10] The original age grid spans 0 to 100, but we restrict attention to the adult population and condition on survival to age $h_0 = 25$.[11] The data is summarized in Figure 3.

We forecast survival probabilities using only observed data and proceed in three steps. First, we assume that age at death follows a Gompertz distribution, a standard model in

[9]The dip is assumed to occur off-sample and reflects structural forces beyond the observed range. Its parameters are therefore not forecastable from historical data and must be imposed. The dip is modeled as a multiplicative modifier of the logistic function:

$$n_t = \left(\frac{b_1}{1 + \exp(-b_2(t - b_3))} \right) \cdot \left[1 - d \cdot \exp\left(-\frac{1}{2} \left(\frac{t - \mu}{\sigma} \right)^2 \right) \right]$$

where d is the dip depth, μ its timing, and σ its width.

[10]The WPP provides only a single projection scenario for survival probabilities, corresponding to its medium variant. No high or low mortality scenarios are reported.

[11]This lower truncation is standard in demographic studies using parametric models. Mortality in early life—particularly infant and child mortality—follows distinct patterns that are not well captured by models designed for adult mortality; see Preston et al. (2001).

Figure 3: Global Survival Probability Profiles

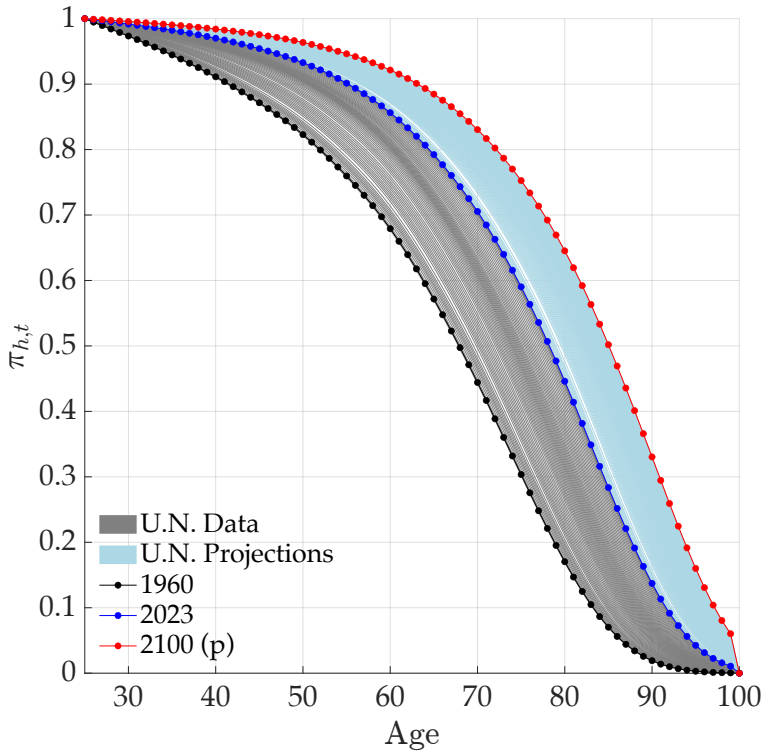

Notes: Survival probability profiles from 1960 to 2023 (gray lines) and projections through 2100 (light blue lines), based on data from the United Nations World Population Prospects (WPP). Each line represents the conditional probability of surviving to older ages, given survival to age 25. Three specific years are highlighted: 1960 (black dots), 2023 (blue dots), and 2100 (blue dots, projected).

formal demography due to its ability to capture the exponential increase in mortality risk with age (see, e.g., Preston et al. (2001)). This implies a survival function of the form

$$\pi_{h,t} = \exp\left(-\frac{b_t}{a_t}\left(e^{a_t h} - 1\right)\right), \quad h \geq h_0,$$

where (a_t, b_t) are period-specific parameters. These are estimated from the data for each year t by fitting the model to the empirical survival probabilities. This parametric approximation smooths over noisy empirical data, enables extrapolation, and yields a low-dimensional representation of adult survival. The resulting time series of estimated parameters $\{a_t, b_t\}_{t=1}^{T}$ summarizes the evolution of mortality and serves as the input for the next step.

In the second step, we take the time series of estimated Gompertz parameters and fit a separate AR(1) model to each. We focus on AR(1) specifications to ensure long-run stationarity, which rules out explosive or permanently shifting mortality dynamics. Using the fitted models, we generate predicted values for the Gompertz parameters over the sample period and reconstruct the corresponding survival curves. This step filters out transitory variation and smooths the year-to-year evolution of survival probabilities. The resulting fitted profiles serve as the input for the forecasting procedure described in Step 3.

Finally, we apply the Lee–Carter procedure (Lee and Carter, 1992) to the mortality profiles implied by the smoothed survival curves obtained in the previous step. Unlike the conventional approach, which applies Lee–Carter directly to raw mortality or survival data, we use as input the parametric survival profiles fitted in the previous step. We first convert the fitted survival probabilities into one-year mortality rates, defined as:

$$q_t(h) = 1 - \frac{\pi_t(h+1)}{\pi_t(h)},$$

where $\pi_t(h)$ is the probability of surviving to age h conditional on survival to age 25. We then model the log mortality rate at each age as:

$$\log q_t(h) = \alpha_h + \beta_h \kappa_t + \varepsilon_{t,h},$$

where α_h captures the average age profile, β_h reflects the sensitivity of age-specific mortality to changes in overall mortality conditions, and κ_t is a time-varying mortality index. We estimate the model using singular value decomposition and project κ_t forward using an AR(1) process. These projections, combined with the estimated age profiles, yield forecasts of age-specific mortality, which we then convert back into survival probabilities.

Figure 4 illustrates how the fitted survival profiles replicate key features of observed and projected mortality patterns across time (Panel A), while the implied life expectancy

Figure 4: Projected Survival Probabilities and Life Expectancy

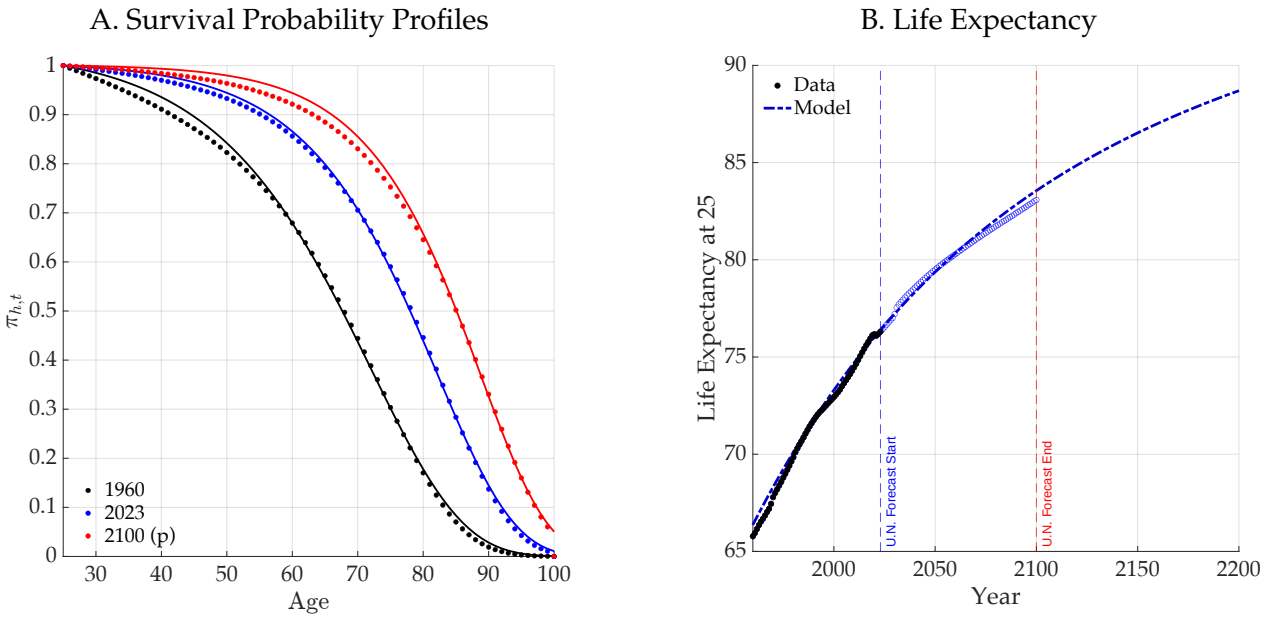

Notes: Panel A displays survival probability profiles for the years 1960, 2023, and 2100. Dots correspond to data from the World Population Prospects (WPP), with 2100 reflecting the medium-fertility projection. Solid lines represent fitted survival profiles for the same years, obtained using the estimation procedure detailed in the main text. Panel B shows life expectancy at age 25 in both the WPP data (observed and projected) and as implied by the fitted survival profiles. The estimation uses age-conditional survival probabilities and restricts attention to the adult population, conditional on survival to age 25.

series (Panel B) aligns closely with the data, capturing the timing and pace of the demographic transition embedded in the WPP projections.

5.2 Climate and Economic Parameters.

The calibration of economic and climate parameters follows closely the approaches of Golosov et al. (2014) and Krusell and Smith (2022), with adjustments to ensure consistency with the annual frequency of the demographic data. Table 1 summarizes the benchmark calibration. The remainder of this section provides details on the selection of each parameter.

We set the initial year to 2010. Since the demographic transition unfolds prior to this date, our calculations assume that optimal policy internalizes the ongoing trajectory of

these demographic trends from that point onward.

Preferences over consumption are logarithmic. The generational discount factor is set at $\delta = 0.985$, following Nordhaus (2007), which implies a social rate of time preference of about 1.5% per year. The private discount factor is calibrated at $\beta = 0.96$, consistent with standard values used in the real business cycle literature (e.g., Kydland and Prescott (1982); Laibson (1997); Gourinchas and Parker (2002)). This ensures that the long-run social discount factor exceeds the private rate, in line with conventional assumptions in the literature (e.g., Schneider et al. (2012); Eden (2023)). As a result, the benchmark calibration satisfies $\beta < \delta$.[12]

Production is modeled using a Cobb-Douglas specification,

$$\tilde{F}(K, N^Y, E) = A_0 K^{\alpha} (N^Y)^{(1-\alpha)(1-\chi)} E^{(1-\alpha)\chi},$$

with $\alpha = 0.34$. We set $\chi = 0.06$, which implies an energy share of $(1-\alpha)\chi = 0.04$, consistent with Golosov et al. (2014).[13] This choice implies a labor share of $1-\alpha = 0.66$, standard in macroeconomic calibration.

Labor productivity in the energy sector is pinned down by energy use data, reported in Krusell and Smith (2022). Summing across coal, oil, gas, and clean energy sources yields a total energy output of 12.27 Gtoe. Assuming that a fraction χ of labor is allocated to energy production in the initial period, labor productivity ψ is set to equate this observed output to total energy production per unit of energy labor input. The level of total factor productivity A_0 is then calibrated to match global GDP in 2010.

The quantitative model allows for both permanent and transitory components in the

[12] When choosing the generational discount factor, the long-run population growth rate must remain below the discount rate. This condition is satisfied in all the scenarios that we model.

[13] Krusell and Smith (2022) consider a higher energy share of 0.06, which would correspond to $\chi \approx 0.09$ in our framework. We find that our results are qualitatively robust to this specification.

Table 1: **Climate and Economic Parameters**

Parameter	Symbol	Benchmark
Preferences and Technology		
CRRA parameter	ν	1.0
Private discount factor	β	0.960
Generational discount factor	δ	0.985
Total Factor Productivity	A_0	1100
Labor Productivity in Energy	ψ	136
Carbon Cycle		
Pre industrial Stock of Carbon	$\overline{S}$	581 GtC
Initial Stock of Carbon	S_0	805 GtC
Share of permanent emissions	ϕ_L	0.25
Share of ephemeral emissions	ϕ_0	0.36
Decay rate of emissions	ϕ	2.3e-03
Carbon content of energy	η	0.69
Production and Damage Function		
Capital share	α	0.34
Energy share	χ	0.06
Capital depreciation rate	δ_k	0.10
Damage Function Elasticity	μ	2.4e-05

carbon cycle. Atmospheric carbon evolves according to

$$S_t = \bar{S} + \sum_{s=0}^{\infty}(1-d_s)E_{t-s},$$

where $1-d_s = \phi_L + (1-\phi_L)\phi_0\phi^s$ denotes the fraction of emissions that remain in the atmosphere s years after emission. The parameters ϕ_L, ϕ_0, ϕ, and the initial stock $\bar{S}$ are calibrated following Golosov et al. (2014), adjusted to an annual frequency.

To calibrate the carbon content of energy, we use data on energy use in 2010 reported in Golosov et al. (2014), originally based on IEA sources, and match it to global emissions in that year, measured in gigatonnes of carbon (GtC). The carbon content parameter η is

then given by the ratio of emissions to energy use.

Finally, we follow the approach in Krusell and Smith (2022) and assume the existence of a backstop technology—a form of exogenous carbon abatement—that leads to full adoption of clean energy at a fixed point in the future. Specifically, a time-varying parameter $\zeta_t \in [0, 1]$ captures the fraction of emissions abated through carbon capture and storage. This parameter follows a logistic decline over time and converges to zero in the long run, so that energy is effectively carbon-free beyond a given horizon. The logistic path is calibrated to match that used by Krusell and Smith: abatement reaches 1% by year 2020 and 50% by 2125.[14]

6 Implications of Global Demographic Trends

We now examine how demographic trends shape the social cost of carbon and global temperature. Figure 5 compares these outcomes under the three fertility scenarios labeled high, medium, and low. The labels reflect the processes introduced in the calibration section: the high-fertility case follows the AR(1) specification, the medium-fertility case follows the logistic specification, and the low-fertility case follows the logistic-with-dip specification. It also includes the corresponding paths that would prevail if demographic change were ignored, as in a representative-agent setup like Golosov et al. (2014).

The figure shows that the dynamics of these variables depend on the nature of the demographic transition. In panel A, the social cost of carbon is consistently higher under the medium- and high-fertility scenarios, with higher fertility translating into a higher SCC. A higher SCC reflects a greater motive to tax carbon, driven by the larger present value

[14] This functional form is also used in Golosov et al. (2014). An alternative approach, followed for example in Barrage (2019), assumes that the adoption of clean energy is endogenous but depends on an exogenous path for the cost of backstop technology. In practice, this approach shifts the calibration problem to specifying the price trajectory of clean energy. The effective role of the backstop remains exogenous in both cases.

Figure 5: Implications of Demographic Transitions

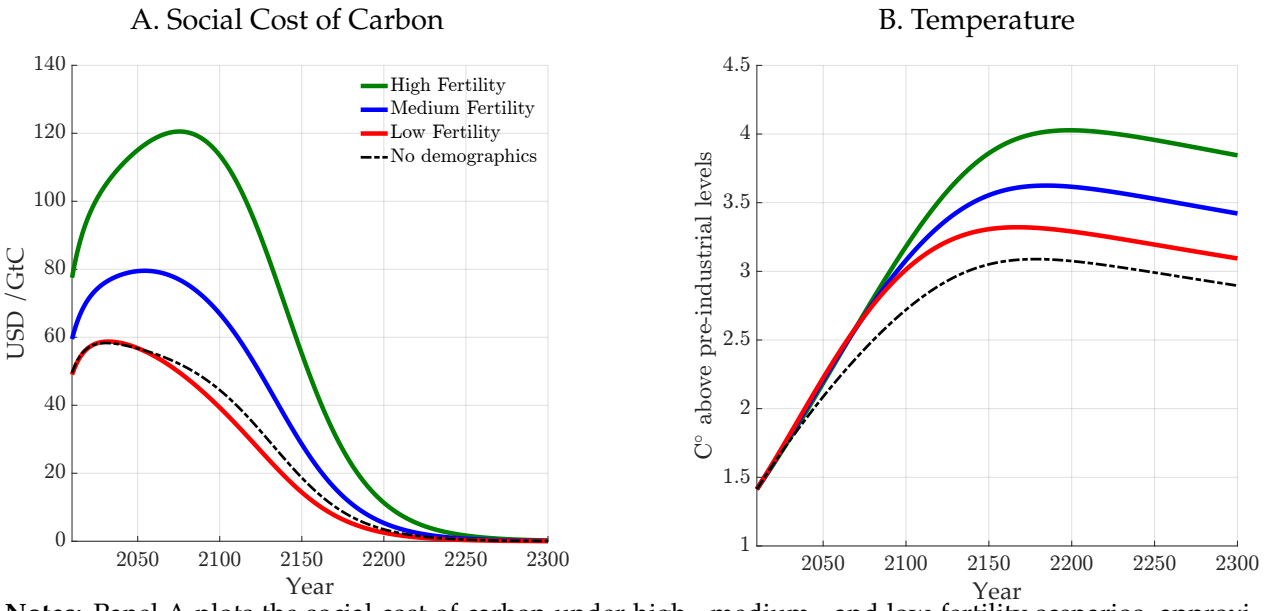

Notes: Panel A plots the social cost of carbon under high-, medium-, and low-fertility scenarios, approximated by AR(1), logistic, and logistic-with-dip trajectories, respectively. Panel B shows the corresponding temperature paths relative to preindustrial levels. Solid lines represent the demographic scenarios, while dash-dotted lines provide reference paths computed without demographic variation. In all cases, the transition to clean energy driven by backstop technology is assumed to be approximately completed by 2300.

of future damages. Relative to the representative-agent model with no demographics, the initial SCC is 20% higher in the medium-fertility case and 57% higher in the high-fertility case. By contrast, the low-fertility path lies very close to the representative-agent outcome.

Panel B shows that the temperature peak is higher under all three fertility scenarios than under the representative-agent benchmark, with the peak increasing as fertility rises. Even in the low-fertility case, the temperature path lies above the no-demographics outcome despite a similar SCC, which implies that aggregate emissions must be higher along the optimal path. Since emissions are higher, a similar SCC is explained by higher discounting of future damages in the low-fertility scenario.

To isolate the effect of aging, we solve the planner's problem under the restriction that the age distribution remains constant over time. This means holding θ_t fixed when solving the planner's problem in Proposition 1, and the results from this counterfactual

Table 2: **Demographic transitions versus Stationary scenarios**

Demographic scenario	**Intial SCC**		**Max SCC**		**Max Temp.**
	USD/tC	% RA	USD/tC	% RA	**in C∘**
Representative agent (No Population Growth)	*49.5*	*100*	*58.3*	*100*	*3.1*
Fixed Age Distribution					
Low Fertility	54.6	110	65.3	112	3.3
Medium Fertility	65.6	133	85.0	146	3.6
High Fertility	84.5	171	127.8	219	4.0
Time-varying Age Distribution					
Low Fertility	49.0	99	58.7	101	3.3
Medium Fertility	59.5	120	79.6	136	3.6
High Fertility	77.5	157	120.5	207	4.0
Stationary 1% Population Growth	*129.5*	*262*	*287.2*	*493*	*3.9*

Notes: Social cost of carbon and peak temperature under alternative demographic assumptions. The top row reports the representative-agent benchmark with no demographic structure. The next three rows hold the age distribution fixed over time (θ_t constant), under low, medium, and high fertility. The following three rows allow the age structure to evolve endogenously under the same fertility scenarios; these baseline cases correspond to those shown in Figure 5. The last row reports the stationary-population case with 1% constant growth for comparison. Columns give the initial SCC, its percentage relative to the RA benchmark, the peak SCC, its percentage relative to the RA benchmark, and the peak temperature.

are reported in Table 2.[15] In the table, fertility scenarios reflect differences in population growth, while comparing fixed versus time-varying age distributions isolates the impact of population aging.

The first column reports the initial SCC, including those shown in Figure 5 for the representative-agent benchmark (top row) and the time-varying age distribution scenarios. The counterfactual cases with a fixed age structure do not appear in the figure but are included here for comparison. Holding the age distribution constant over time yields

[15]Note that fixing θ_t is not the same as assuming fixed survival curves, since the age structure can still change with population growth. Treating θ_t as fixed corresponds to using a misspecified model that ignores the role of aging when computing the social cost of carbon and temperature.

higher initial SCCs across all fertility cases than the corresponding time-varying versions. Consistent with Proposition 3, population aging raises the social discount rate and lowers the SCC, whereas fixing the age structure removes this channel, making the planner more patient and increasing the present value of future climate damages. The third column shows that the difference between the representative agent case and the medium and high fertility case gets larger with time, reflecting the larger scale of future damages as population size grows. The last column shows that omitting the effect that aging has on discounting has a negligible impact on peak temperature.

The SCCs remain well below those under the stationary 1% population growth scenario, which holds the age structure fixed by construction (bottom row). This case yields a peak SCC more than twice that of the high-fertility scenario while leaving peak temperature nearly unchanged (3.9°C vs. 4.0°C). This contrast highlights why properly accounting for the demographic transition is essential: simply extrapolating recent average population growth would overestimate the SCC.

These results show that the impact of demographic change on the SCC and temperature paths depends on how the population age structure and size shape social preferences. Since this link hinges on the relationship between the generational discount factor, δ, and the private one, β, it is useful to examine how varying δ affects the results. The next section presents this sensitivity analysis and discusses its implications for the social cost of carbon, emissions, and temperature.

6.1 Robustness to Alternative Values of δ.

The generational discount factor, δ, determines how the planner values future generations relative to current ones. In this section, we maintain the assumption that $\delta > \beta$ and vary δ across a range that spans two prominent normative benchmarks: Nordhaus's calibration (see Nordhaus (2007)) and Stern's prescription (see Stern (2008)). Varying δ in this way

Table 3: **Initial SCC and Peak Temperature**

Demographic scenario	**Generational Discount Factor δ**				
	0.965	0.975	0.985†	0.995	0.999‡
			Initial SCC		
Representative agent (No Population Growth)	*23.2*	*31.3*	*49.5*	*129.4*	*465.6*
Fixed Age Distribution					
Low Fertility	26.8	36.0	54.6	121.0	360.9
Medium Fertility	28.3	39.6	65.6	182.8	679.2
High Fertility	31.0	45.8	84.5	308.2	1459.4
Time-varying Age Distribution					
Low Fertility	26.4	34.1	49.0	99.4	273.6
Medium Fertility	27.9	37.8	59.5	150.8	512.3
High Fertility	30.7	44.1	77.5	254.2	1092.8
			Peak temperature		
Representative agent (No Population Growth)	*3.2*	*3.2*	*3.1*	*2.8*	*2.2*
Fixed Age Distribution					
Low Fertility	3.4	3.4	3.3	3.1	2.5
Medium Fertility	3.8	3.7	3.6	3.2	2.3
High Fertility	4.3	4.2	4.0	3.3	2.1
Time-varying Age Distribution					
Low Fertility	3.4	3.4	3.3	3.1	2.5
Medium Fertility	3.8	3.7	3.6	3.3	2.4
High Fertility	4.3	4.2	4.0	3.4	2.3

Notes: Columns report outcomes for the specified values of the generational discount factor δ. Rows show demographic scenarios, including the representative-agent benchmark with no population growth and transitions under low, medium, and high fertility, assuming either a fixed or a time-varying age distribution. The top panel contains the initial social cost of carbon (SCC), and the bottom panel reports the corresponding peak temperature. The value $\delta = 0.985$ (†) corresponds to the discount factor used in Nordhaus (2007), and $\delta = 0.999$ (‡) corresponds to the discount factor used in Stern (2008).

highlights how different assumptions about intergenerational welfare weights can affect the present value of climate damages as the population evolves.

The results are summarized in Table 3, which displays initial SCCs and peak temperatures for different discount factor scenarios. As in Table 2, the fertility scenarios capture differences in population growth, while the distinction between fixed and time-varying age distributions isolates the role of population aging. Across columns, different values for the generational discount factor show how alternative assumptions about intergenerational welfare weights affect the results. For reference, the third column reproduces the results for $\delta = 0.985$, which serves as the baseline throughout the paper and aligns with Nordhaus's benchmark.

As expected, initial SCCs increase across all scenarios as more weight is placed on future generations. An interesting pattern emerges when comparing the fixed and time-varying age distribution cases: the gap between them – driven entirely by the effect of population aging on the social discount rate – widens as δ increases. In other words, the impact of aging on the discount rate, and therefore on the SCC, becomes stronger as the planner places more weight on the welfare of future generations. This pattern holds broadly across all fertility scenarios. Under the highest values of δ (as in Stern's calibration), the effect is so strong that in the low-fertility case the SCC falls to just over half the value implied by the representative-agent benchmark, whereas under lower values of δ (as in Nordhaus's calibration) the low-fertility SCC remains much closer to it.

The bottom panel reports peak temperatures for each scenario and discount factor. Differences between fixed and time-varying age structures remain small, except at very high values of δ. Although population size is held constant across columns, the planner's optimal choices regarding savings and mitigation vary, affecting emissions and, in turn, peak temperature. As δ increases, placing greater weight on future generations leads to lower peak temperatures by encouraging stronger mitigation. However, it is no longer the case that peak temperature necessarily increases with fertility. While it remains higher under all three fertility scenarios relative to the representative-agent benchmark, when

Figure 6: Optimal Choices across Fertility scenarios.

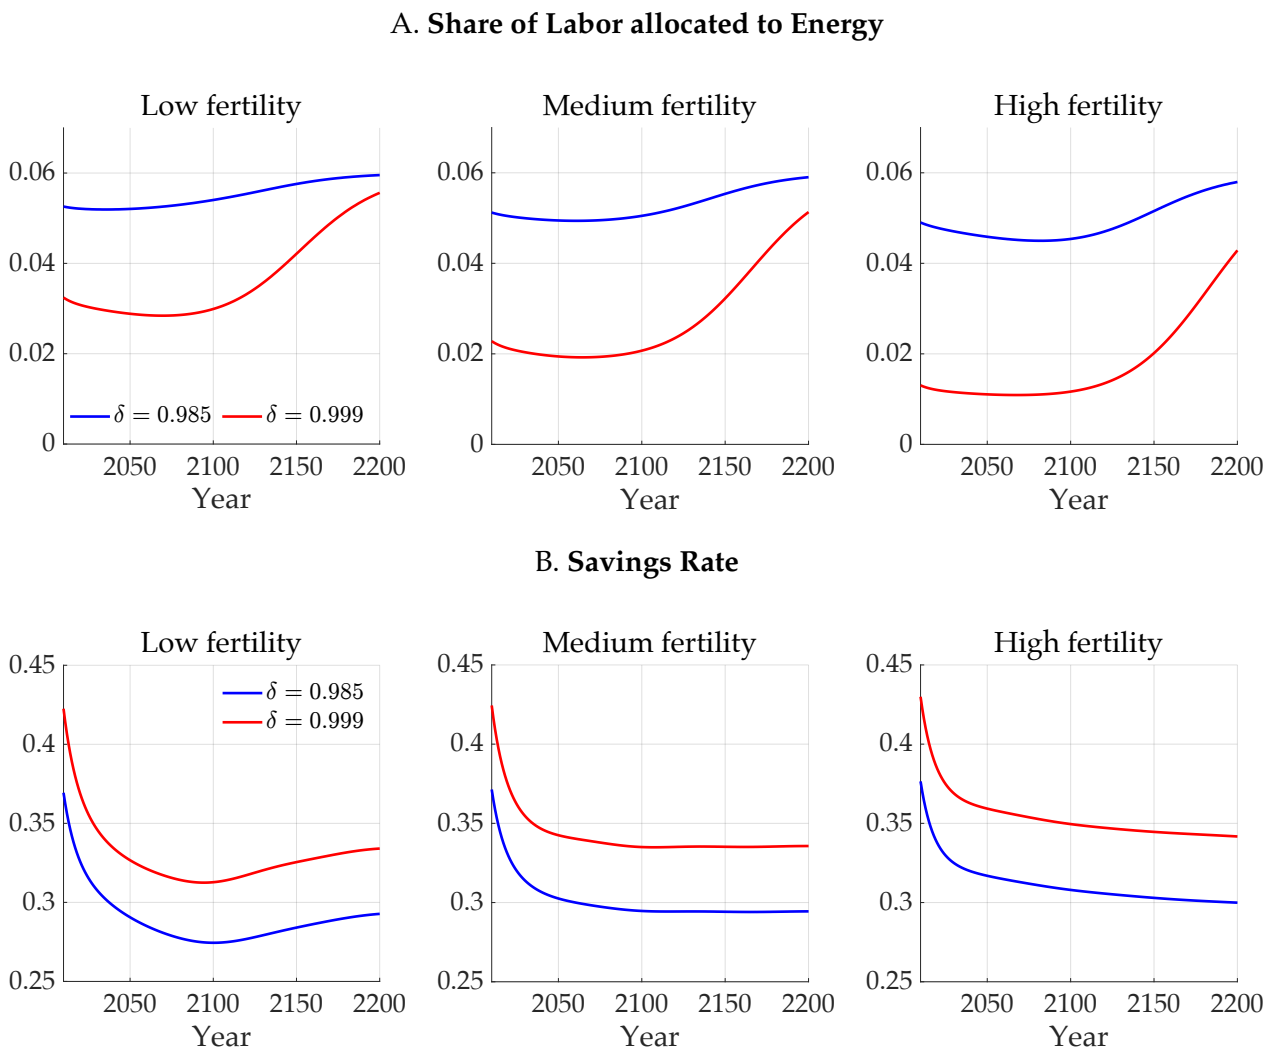

$\delta = 0.999$, the highest temperature occurs under the low-fertility scenario. Importantly, this reversal appears under both fixed and time-varying age structures, suggesting that the pattern is driven primarily by differences in population growth rather than by aging.

To understand this pattern, we examine how the planner's savings and mitigation choices respond to different fertility paths. Figure 6 shows the optimal savings rate and the share of labor allocated to the energy sector for each fertility scenario, focusing on the Nordhaus and Stern benchmarks for δ. The top panel shows that higher fertility induces the planner to mitigate more by reducing the share of labor allocated to energy production,

Figure 7: Optimal Emissions across Fertility Scenarios.

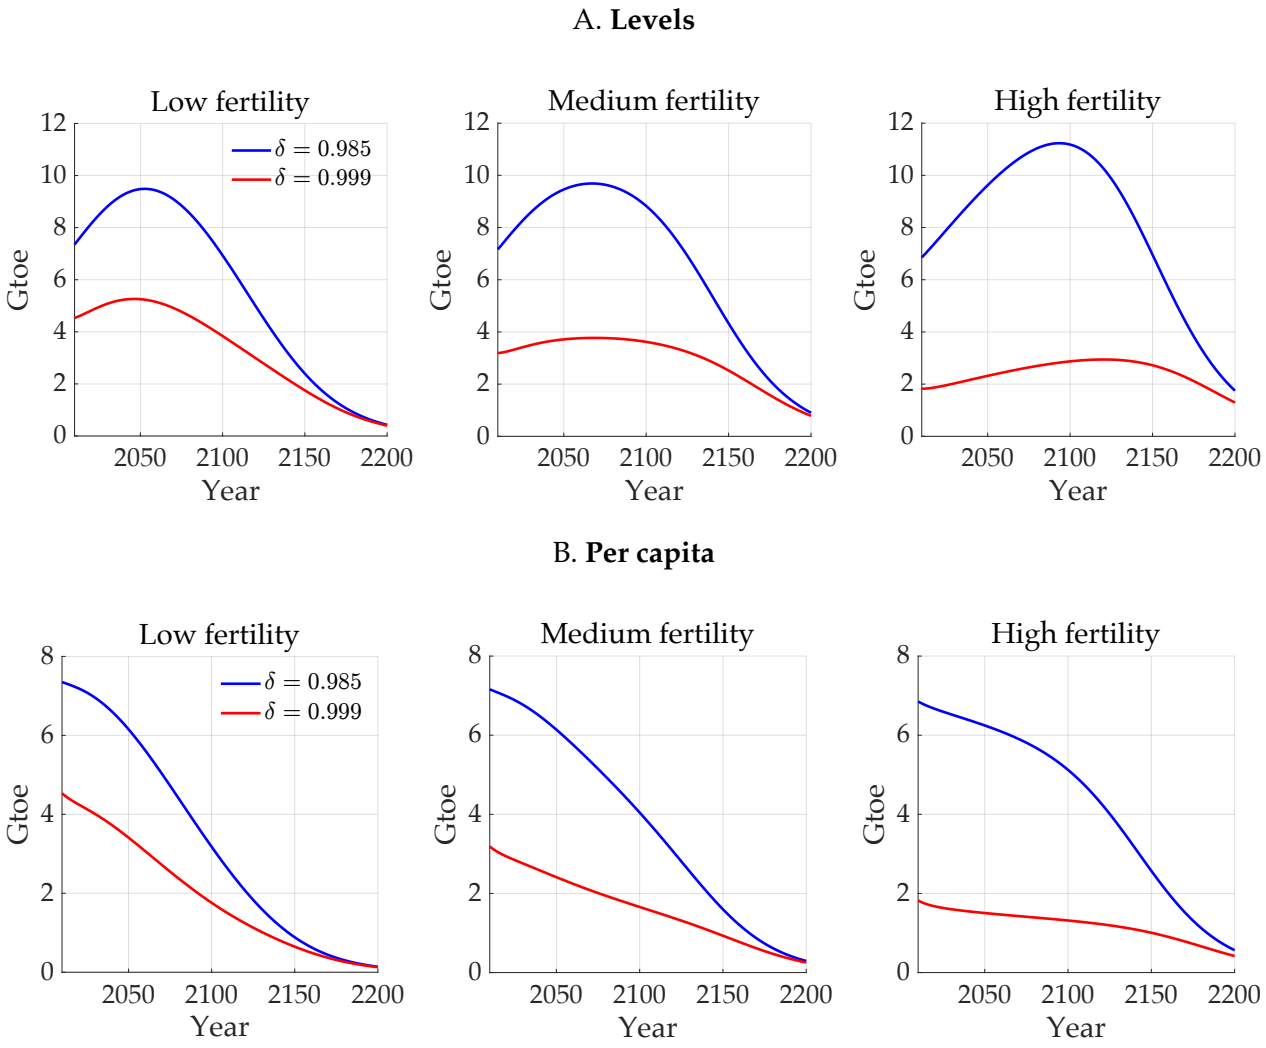

which lowers per-capita emissions. The bottom panel shows that the savings rate rises with fertility to smooth per-capita consumption over time. Both choices are consistent with Proposition 3: higher fertility implies higher population growth, which lowers the social discount rate and makes the planner more patient. A larger population requires more resources to maintain per-capita consumption, reinforcing the planner's incentive to sustain consumption levels over time, consistent with the Ramsey condition.

However, aggregate emissions depend not only on the planner's mitigation effort but also on the size of the population. When δ is relatively low (as in the Nordhaus bench-

mark), the population effect dominates: higher fertility means a larger population, which more than offsets the reduction in labor allocated to energy. As a result, total emissions and peak temperatures rise with fertility. When δ is high (as in the Stern benchmark), the planner places much greater weight on intergenerational equity, which amplifies the incentive to shift resources away from current energy use. Higher fertility still expands the population, but the stronger mitigation effort can offset this effect. By contrast, when fertility is low under a high δ, future cohorts are smaller, weakening the planner's motive to save and mitigate aggressively. This explains why, at high values of δ, the low-fertility scenario can yield higher emissions and temperatures than the high-fertility case.

This mechanism is illustrated in Figure 7, which shows how the planner's optimal choices translate into emissions outcomes. The top panel shows total emissions: when the generational discount factor δ is relatively low ($\delta = 0.985$), total emissions increase with fertility because the effect of a larger population outweighs the planner's mitigation efforts. By contrast, when δ is high ($\delta = 0.999$), stronger intergenerational equity concerns lead the planner to mitigate more aggressively, so total emissions actually decline with higher fertility. The bottom panel shows per-capita emissions. When δ is low, higher fertility lowers the initial level but flattens the decline, so per-capita emissions stay above the low-fertility path for some time. This slower decline, combined with a larger population, keeps total emissions higher overall. Together, these patterns confirm that the interaction between the discount factor, fertility, and population growth critically shapes the climate implications of optimal mitigation policy.

7 Conclusions

We develop a framework to study climate policy in the presence of demographic change, focusing on how shifts in population growth and age structure affect the social cost of

carbon emissions.

The key theoretical contribution is a formula for the social cost of carbon that embeds demographic variables into the social discount factor. This result builds on an aggregation result that allows us to express the planner's problem with overlapping generations in terms of a representative agent. Quantitatively, we find that accounting for demographic dynamics matters. In a model calibrated to match the ongoing global demographic transition, the social cost of carbon is substantially higher under medium- and high-fertility scenarios, compounding the effect of greater emissions from larger populations partially offset by higher discounting in aging societies. Global temperatures also increase more sharply in these scenarios, suggesting that demographic trends will exacerbate the climate problem.

We show that properly capturing the transitional path is essential as relying on long-run average population growth in a stationary framework can substantially overestimate the social cost of carbon. These findings highlight the need to integrate demographic dynamics into climate models. Policy frameworks that neglect these dynamics risk misestimating the costs of climate change and the appropriate level of mitigation efforts. Effective climate policy must reflect the evolving composition of society to ensure that mitigation efforts are aligned with the intergenerational trade-offs at the heart of the climate problem.

Our work opens several promising directions for future research. First, an important direction for future research is understanding how regional demographic divergence affects national assessments of climate damages. This could exacerbate the climate problem by introducing challenges in global climate policy coordination and resulting in unequal distributional effects across regions. Second, while we take initial steps toward modeling climate-induced demographic change, a full quantitative analysis of this feedback loop would be valuable. Third, we show that demographic dynamics endogenously generate a time-varying social discount rate, introducing time inconsistency in climate policy.

While this paper focuses on the solution under commitment, exploring outcomes without commitment would be a useful extension. Finally, since demographic dynamics create a wedge between social and private discount rates, studying how to decentralize the optimal allocation in a market economy could yield further insights for climate policy design.

References

AGUIAR, M. A., M. AMADOR, AND C. ARELLANO (2023): "Pareto improving fiscal and monetary policies: Samuelson in the new keynesian model," Tech. rep., National Bureau of Economic Research.

ARROW, K., M. CROPPER, C. GOLLIER, B. GROOM, G. HEAL, R. NEWELL, W. NORDHAUS, R. PINDYCK, W. PIZER, P. PORTNEY, T. STERNER, R. S. J. TOL, AND M. WEITZMAN (2013): "Determining Benefits and Costs for Future Generations," *Science*, 341, 349–350.

ARROW, K., P. DASGUPTA, L. GOULDER, G. DAILY, P. EHRLICH, G. HEAL, S. LEVIN, K.-G. MÄLER, S. SCHNEIDER, D. STARRETT, AND B. WALKER (2004): "Are We Consuming Too Much?" *Journal of Economic Perspectives*, 18, 147–172.

BARCONS, S., E. DÁVILA, AND A. SCHAAB (2025): "Intergenerational Welfare Assessments," .

BARRAGE, L. (2019): "Optimal Dynamic Carbon Taxes in a Climate–Economy Model with Distortionary Fiscal Policy," *The Review of Economic Studies*, 87, 1–39.

BELFIORI, E., D. R. CARROLL, AND S. HUR (2024): "Unequal Climate Policy in an Unequal World," Globalization Institute Working Papers 427, Federal Reserve Bank of Dallas.

BELFIORI, E. AND M. MACERA (2024): "Climate Inequality: Carbon Capture for Redistribution," *CESifo Working Paper No. 11239*.

BOHN, H. AND C. STUART (2015): "Calculation of a Population Externality," *American Economic Journal: Economic Policy*, 7, 61–87.

BOURANY, T. (2024): "Climate Change, Inequality and the Optimal Climate Policy," Tech. rep.

BRESSLER, R. D. (2021): "The mortality cost of carbon," *Nature Communications*, 12, 4467.

CALVO, G. A. AND M. OBSTFELD (1988): "Optimal Time-Consistent Fiscal Policy with Finite Lifetimes," *Econometrica*, 56, 411–432.

CARLETON, T. A., A. JINA, M. DELGADO, M. GREENSTONE, T. HOUSER, S. M. HSIANG, A. HULTGREN, R. E. KOPP, K. E. MCCUSKER, I. NATH, J. RISING, A. RODE, H. K. SEO, A. VIAENE, J. YUAN, AND A. T. ZHANG (2023): "Valuing the Global Mortality Consequences of Climate Change Accounting for Adaptation Costs and Benefits," *The Quarterly Journal of Economics*, 137, 2037–2105.

CONTE, B., K. DESMET, AND E. ROSSI-HANSBERG (2022): "On the Geographic Implications of Carbon Taxes," Working Paper 30678, National Bureau of Economic Research.

CRUZ, J.-L. AND E. ROSSI-HANSBERG (2023): "The Economic Geography of Global Warming," *The Review of Economic Studies*, 91, 899–939.

DOUENNE, T., A. J. HUMMEL, AND M. PEDRONI (2023): "Optimal Fiscal Policy in a Climate-Economy Model with Heterogeneous Households," Tech. rep., University of Amsterdam.

EDEN, M. (2023): "The Cross-Sectional Implications of the Social Discount Rate," *Econometrica*, 91, 2065–2088.

FRIED, S., K. NOVAN, AND W. B. PETERMAN (2024): "Understanding the Inequality and Welfare Impacts of Carbon Tax Policies," *Journal of the Association of Environmental and Resource Economists*, 11, S231–S260.

GERLAGH, R., V. LUPI, AND M. GALEOTTI (2023): "Fertility and Climate Change," *The Scandinavian Journal of Economics*, 125, 208–252.

GOLOSOV, M., J. HASSLER, P. KRUSELL, AND A. TSYVINSKI (2014): "Optimal Taxes on Fossil Fuel in General Equilibrium," *Econometrica*, 82, 41–88.

GOLOSOV, M., L. E. JONES, AND M. TERTILT (2007): "Efficiency with Endogenous Population Growth," *Econometrica*, 75, 1039–1071.

GOURINCHAS, P.-O. AND J. A. PARKER (2002): "Consumption over the Life Cycle," *Econometrica*, 70, 47–89.

HILLEBRAND, E. AND M. HILLEBRAND (2019): "Optimal climate policies in a dynamic multi-country equilibrium model," *Journal of Economic Theory*, 179, 200–239.

KRUSELL, P. AND J. SMITH, ANTHONY A (2022): "Climate Change Around the World," Working Paper 30338, National Bureau of Economic Research.

KYDLAND, F. E. AND E. C. PRESCOTT (1977): "Rules Rather than Discretion: The Inconsistency of Optimal Plans," *Journal of Political Economy*, 85, 473–491.

——— (1982): "Time to Build and Aggregate Fluctuations," *Econometrica*, 50, 1345–1370.

LAIBSON, D. (1997): "Golden Eggs and Hyperbolic Discounting," *Quarterly Journal of Economics*, 112, 443–478.

LEE, R. D. AND L. R. CARTER (1992): "Modeling and forecasting U.S. mortality," *Journal of the American Statistical Association*, 87, 659–671.

NORDHAUS, W. D. (2007): "A Review of the Stern Review on the Economics of Climate Change," *Journal of Economic Literature*, 45, 686–702.

NORDHAUS, W. D. AND J. BOYER (2003): *Warming the world: economic models of global warming*, MIT Press.

PRESTON, S. H., P. HEUVELINE, AND M. GUILLOT (2001): *Demography: Measuring and Modeling Population Processes*, Malden, MA: Blackwell.

RAMSEY, F. P. (1928): "A Mathematical Theory of Saving," *The Economic Journal*, 38, 543–559.

SCHNEIDER, M. T., C. P. TRAEGER, AND R. WINKLER (2012): "Trading off generations: Equity, discounting, and climate change," *European Economic Review*, 56, 1621–1644.

STERN, N. (2007): *The Economics of Climate Change: The Stern Review*, Cambridge University Press.

——— (2008): "The Economics of Climate Change," *American Economic Review*, 98, 1–37.

VAN DER PLOEG, F. AND A. REZAI (2021): "Optimal carbon pricing in general equilibrium: Temperature caps and stranded assets in an extended annual DSGE model," *Journal of Environmental Economics and Management*, 110, 102522.

WEITZMAN, M. L. (2007): "A Review of the Stern Review on the Economics of Climate Change," *Journal of Economic Literature*, 45, 703–724.

World Population Prospects (2024): "World Population Prospects 2024," Online Edition, accessed [insert date of retrieval].

Appendix

A From newborn to total population growth rates

We show here that, given a survival probability profile, there is a one to one map between choosing a growth rate for the newborns and a growth rate for the entire population. Let n^0_{t+1} be the growth rate of newborns between period t and $t+1$. This means that:

$$N_{0,t+1} = (1 + n^0_t) N_{0,t}. \tag{30}$$

for all t. Also, as pointed out in the main text we have that:

$$N_{h,t+h} = \pi_{h,t} N_{0,t}. \tag{31}$$

for all t and h. Combining these This implies that total population can be written as

$$N_t = \left((1 + \eta^t_{t-H}) + (1 + \eta^{t+h-1}_t)\pi_{1,t-H} + ... + \pi_{H,t-H} \right) N_{0,t-H} \tag{32}$$

for all t, where η^{t+s}_t denotes the *cumulative* population growth rate of newborns between period t and $t+s$. Therefore, (32) shows that given a profile of survival probabilities over time, there is a one-to-one map between the rate of total population growth defined as $n_{t+1} \equiv (N_{t+1} - N_t)/N_t$ and the rate of population growth of newborns.

B Proof of Lemma 1.

Replacing the expression for the Pareto weights (11) into the social welfare function (10) yields

$$\sum_{t=0}^{\infty} \frac{\delta^t N_{0,t}}{\sum_{s=0}^{\infty} \delta^s N_{0,s}} \sum_{h=0}^{H} \beta^h \pi_{h,t} u(c_{h,t+h}) \tag{33}$$

We can bring the all terms within the second summation to get:

$$\sum_{t=0}^{\infty} \sum_{h=0}^{H} \frac{\delta^{t-h} N_{0,t-h}}{\sum_{s=0}^{\infty} \delta^s N_{0,s}} \beta^h \pi_{h,t-h} u(c_{h,t}) \tag{34}$$

Using (5), we get:

$$\sum_{t=0}^{\infty} \sum_{h=0}^{H} \frac{\delta^t N_{h,t}}{\sum_{s=0}^{\infty} \delta^s N_{0,s}} \left(\frac{\beta}{\delta}\right)^h u(c_{h,t}).$$

Finally, use (7) to obtain:

$$\sum_{t=0}^{\infty} \sum_{h=0}^{H} \frac{\delta^t}{\sum_{s=0}^{\infty} \delta^s N_{0,s}} \omega_{h,t} N_t \left(\frac{\beta}{\delta}\right)^h u(c_{h,t}).$$

We can normalize total welfare to add up to $\sum_{s=0}^{\infty} \delta^s N_{0,s}$ and take this summation out of the expression so that:

$$\sum_{t=0}^{\infty} \delta^t N_t \sum_{h=0}^{H} \omega_{h,t} \left(\frac{\beta}{\delta}\right)^h u(c_{h,t}).$$

Normalizing $N_0 = 1$ and where N_t is the cumulative population growth rate. This corresponds to (12).

C Proof of Proposition 3

Using the CRRA assumption and taking logs in (23) delivers:

$$j \log(1 + r^s_{t,t+j}) = j \log \frac{1}{\delta} + \nu \left(\log \frac{C_{t+j}}{C_t} - \log \frac{N_{t+j}}{N_t} - \log \frac{\theta_{t+j}}{\theta_t} \right)$$

Using $\delta = \frac{1}{1+\rho}$ and that for any x_t, $\frac{1}{j} \log(x_{t+j}/x_t)$ is the average annual growth rate from t to $t + j$ delivers the expression in the main text.