

Optimal climate policy with demographic transitions

Elisa Belfiori

Universidad Torcuato Di Tella

ebelfiori@utdt.edu

Manuel Macera

Universidad Torcuato Di Tella

mmacera@utdt.edu

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. 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 ageing 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.

We thank the editor and two anonymous referees for their insightful remarks and suggestions that improved the paper. We are grateful to Rick van der Ploeg, Andy Neumeyer, Marten Hillebrand, Hugo Hopenhayn, and Facundo Piguillem for valuable discussions. We also thank seminar participants at the 2024 SED Winter Meetings and the Macroeconomic Perspectives on Climate Change Workshop for helpful comments, and Yanel Llohis and Tomás Marco for excellent research assistance. This paper subsumes some results that were originally part of a working paper version of our work entitled “Demographic Transitions and Social Discounting.”

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 ageing. 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 paper 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 characterises 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 internalises 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 ageing 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, understood as population growth and population ageing, affects this cost through two channels. The first is a scale effect: larger populations produce more output and generate more emissions, raising climate damages; ageing leaves this channel unchanged, as agents are identical across generations. The second, novel channel operates through discounting: population growth raises future consumption needs and makes society more patient, whereas population ageing, combined with a higher welfare weight on younger cohorts, makes society more impatient and reduces the present value of future damages.

To formalise this mechanism, we derive a generalised 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 ageing 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 generalised 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](#) to parametrise 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 parametrise 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 β 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 ageing society, which lowers their present value and keeps the social cost of carbon close to the benchmark level.

To isolate the role of ageing, 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 ageing. 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 ageing: when future generations are weighted more heavily, the rise in the social discount rate induced by population ageing has a greater impact. In extreme cases, low-fertility scenarios can produce higher peak temperatures than high-fertility ones, as weaker mitigation incentives in ageing 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) 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 ?, 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 welfare analysis within an overlapping-generations framework builds on Calvo and Obstfeld (1988) and Arrow *et al.* (2004), as well as more recent contributions by Eden (2023), Schneider *et al.* (2012), and Belfiori and Macera (2025). As in Belfiori and Macera (2025), demographic transitions in this setting endogenously shape social discounting: as the population ages, the social discount factor declines, and society effectively behaves as if it were more impatient. Unlike in dynastic models, where the discount factor is fixed, this framework implies that the effect of ageing 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 formalise this mechanism through a generalised Ramsey rule that incorporates ageing, 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 generalised 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 modelling 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.* \(2024\)](#) 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.* \(2024\)](#), in turn, analyzes robust Pareto improvements in decentralised 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 organised 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, labour, and energy input according to a constant returns to scale technology $\tilde{F}(K_t, N_t^Y, E_t)$. The energy sector is labour-intensive, with a linear production function $E_t = \psi N_t^E$.

Energy production releases carbon into the atmosphere. The carbon stock S_t evolves according to

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

where $\bar{S}$ is the pre-industrial stock of carbon, $1 - d_s$ is the amount of carbon emissions left in the atmosphere s periods into the future, and η captures the carbon content of energy production. For expositional purposes, in this section we focus on the special case $d_s = 1 - (1 - \gamma)^s$, which allows us to write

$$S_{t+1} = \gamma \bar{S} + (1 - \gamma) S_t + \eta E_t. \quad (2)$$

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

$$Y_t = F(S_t, K_t, N_t^Y, E_t) = [1 - x(S_t)] \tilde{F}(K_t, N_t^Y, E_t), \quad (3)$$

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

Aggregate feasibility requires

$$N_t^E + N_t^Y = N_t, \quad (4)$$

where N_t represents aggregate labour, and

$$C_t + K_{t+1} = F(S_t, K_t, N_t^Y, E_t) + (1 - \delta_k)K_t \quad (5)$$

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

Demographics. In each period t , a new cohort of identical individuals enters the population. The survival probability to age h for an individual born at t is denoted $\pi_{h,t}$. We set $\pi_{0,t} = 1$ and $\pi_{h,t} = 0$ whenever $h > H$, so that no individual lives beyond $H + 1$ periods. As usual, a law of large numbers applies to this type of economy and, thus, the proportion of agents born at time t who survive to age h is given by $\pi_{h,t}$.

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

$$N_{h,t+h} = \pi_{h,t}N_{0,t} \quad (6)$$

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 the total population grows according to:

$$N_{t+1} = (1 + n_{t+1})N_t, \quad (7)$$

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

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

$$\omega_{h,t} \equiv \frac{N_{h,t}}{N_t} = \frac{\pi_{h,t-h} N_{0,t-h}}{N_t}. \quad (8)$$

Preferences. A consumption plan $c = \{c_{0,t}, c_{1,t}, \dots, c_{H,t}\}_{t=0}^{\infty}$ specifies consumption for each individual in each period. For a cohort born at τ , lifetime utility from plan $\mathbf{c}$ is

$$U_{\tau}(\mathbf{c}) = \sum_{h=0}^H \beta^h \pi_{h,\tau} u(c_{h,\tau+h}), \quad (9)$$

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

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

$$C_t = \sum_{h=0}^H N_{h,t} c_{h,t} \quad (10)$$

for every period t .

¹Assuming 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.

3 The Social Planning Problem

Following [Calvo and Obstfeld \(1988\)](#) and [Arrow *et al.* \(2004\)](#), we adopt a discounted-utilitarian social welfare criterion. The planner's objective is

$$W(\mathbf{c}) = \sum_{\tau=-H}^{\infty} \alpha_{\tau} U_{\tau}(\mathbf{c}), \quad (11)$$

where α_{τ} is the Pareto weight on generation τ . These weights are defined by

$$\alpha_{\tau} = \frac{\delta^{\tau} N_{0,\tau}}{\sum_{s=-H}^{\infty} \delta^s N_{0,s}}, \quad \tau \geq -H. \quad (12)$$

The planner uses the generational discount factor δ to discount the well-being of future generations, and we assume $\delta > \beta$.² When $\beta = 0$, the framework collapses to the standard dynastic model with no altruism. In that case, the discount factor δ governs how the planner values welfare over time and is naturally related to the planner's pure rate of time preference.

The next result provides an equivalent representation of (11), in which the planner assigns weights to cohorts within each period rather than across generations. A proof is given in the appendix.

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

$$W(\mathbf{c}) = \sum_{t=-\infty}^{\infty} \delta^t N_t \sum_{h=0}^H \left(\frac{\beta}{\delta}\right)^h \omega_{h,t} u(c_{h,t}). \quad (13)$$

²This 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 \(2007\)](#)). In Integrated Assessment Models à la [Nordhaus \(2007\)](#), δ corresponds to a pure rate of time preference of about 1.5% per year, chosen so that—together with long-run growth via the Ramsey rule—it aligns with observed real interest rates; [Stern \(2007\)](#) instead adopts a near-zero rate on ethical grounds. Typical annual private discount factors are $\beta \approx 0.96$: in RBC calibrations (e.g., [Kydland and Prescott, 1982](#)), β is set to match steady-state real interest rates, while in structural micro work (e.g., [Gourinchas and Parker, 2002](#)) it is estimated by method of simulated moments to fit life-cycle consumption profiles.

The *socially optimal allocation* consists of a consumption plan $\mathbf{c}$ and a path for aggregate consumption, capital, labour, energy, and carbon stock, $\{C_t, K_t, N_t^Y, N_t^E, E_t, S_t\}_{t=0}^\infty$, that maximises social welfare (13), subject to the carbon cycle (2), feasibility constraints (5) and (4), and aggregate consistency (10). To characterise 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, labour, 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

$$\max_{\{c_{h,t}\}} \sum_{h=0}^H \left(\frac{\beta}{\delta}\right)^h \omega_{h,t} u(c_{h,t}) \quad (14)$$

subject to (10). The first-order condition with respect to individual consumption is:

$$\left(\frac{\beta}{\delta}\right)^h \omega_{h,t} u'(c_{h,t}) - \Lambda_t N_{h,t} = 0 \quad (15)$$

for all h , where Λ_t is the Lagrange multiplier on the constraint (10). Using the functional form for the instantaneous utility function and equation (8), we can solve for individual consumption:

$$c_{h,t} = \left(\frac{\beta}{\delta}\right)^{h/\nu} (N_t \Lambda_t)^{-1/\nu} \quad (16)$$

Thus, aggregate consumption in any period t equals

$$C_t = (N_t \Lambda_t)^{-1/\nu} \sum_h N_{h,t} \left(\frac{\beta}{\delta} \right)^{h/\nu} = (N_t \Lambda_t)^{-1/\nu} N_t \theta_t \quad (17)$$

where θ_t is defined as:

$$\theta_t \equiv \sum_{h=0}^H \omega_{h,t} \left(\frac{\beta}{\delta} \right)^{h/\nu} \quad (18)$$

We can express individual consumption as follows:

$$c_{h,t} = \left(\frac{\beta}{\delta} \right)^{h/\nu} \frac{C_t}{\theta_t N_t} \quad (19)$$

Plugging this expression into (14) delivers the following result.

Lemma 2 *[Ageing as a Preference Shock] The indirect utility function of the intratemporal social planning problem is given by*

$$U(C_t) = \theta_t' u(C_t / N_t). \quad (20)$$

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 ageing – meaning a persistent increase in the share of older individuals – then θ_t declines over time when $\delta > \beta$. ageing 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, labour, and carbon stock of the socially optimal allocation solve the following problem*

$$\max_{C_t, N_t^Y, N_t^E, K_{t+1}, S_{t+1}} \sum_{t=0}^{\infty} \delta^t N_t^\gamma \theta_t^\gamma u(C_t) \quad (21)$$

subject to [\(2\)](#), [\(4\)](#) and [\(5\)](#), 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: $\delta \left(\frac{N_{t+1}}{N_t} \right)^\gamma \left(\frac{\theta_{t+1}}{\theta_t} \right)^\gamma$. 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 cost of carbon emissions under full commitment.³

³The 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 \(1977\)](#)).

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 characterise this cost, let $\delta^t \sigma_t$ and $\delta^t \lambda_t$ denote the Lagrange multipliers associated with the carbon cycle (2) and the feasibility constraint (5), respectively. The first-order condition with respect to the global carbon stock is:

$$\delta^t \sigma_t = (1 - \gamma) \delta^{t+1} \sigma_{t+1} + \delta^{t+1} \lambda_{t+1} \frac{\partial F(S_{t+1}, K_{t+1}, N_{t+1}^Y, E_{t+1})}{\partial S_{t+1}} \quad (22)$$

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*

$$\hat{\sigma}_t \equiv \frac{\sigma_t}{\lambda_t} = \sum_{j=1}^{\infty} (1 - \gamma)^{j-1} \delta^j \left(\frac{N_{t+j}}{N_t} \right)^\nu \left(\frac{\theta_{t+j}}{\theta_t} \right)^\nu \frac{u'(C_{t+j})}{u'(C_t)} \frac{\partial F(S_{t+j}, K_{t+j}, N_{t+j}^Y, E_{t+j})}{\partial S_{t+j}}. \quad (23)$$

Intuitively, the SCC is the present discounted value of future output losses caused by an extra tonne 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

⁴Under 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.

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 23](#), the rate at which society discounts future climate damages from period $t + j$ back to t is given by

$$\left(\frac{1}{1 + r_{t,t+j}^s} \right)^j \equiv \delta^j \left(\frac{N_{t+j}}{N_t} \right)^\nu \left(\frac{\theta_{t+j}}{\theta_t} \right)^\nu \frac{u'(C_{t+j})}{u'(C_t)} \quad (24)$$

where $r_{t,t+j}^s$ denotes the annualised social discount rate from t to $t + j$. Log-linearizing this expression yields the generalised Ramsey formula:

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

$$r_{t,t+j}^s = \rho + \nu(g_{t,t+j} - n_{t,t+j} - \mu_{t,t+j}) \quad (25)$$

where $\rho \equiv \delta^{-1} - 1$ is the generational discount rate, and $g_{t,t+j}$, $n_{t,t+j}$, and $\mu_{t,t+j}$ 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. Faster population growth reduces per-capita consumption for future generations, which raises the marginal utility of their consumption. This strengthens the incentive to save and lowers the social discount rate—society therefore behaves more patiently when population growth is high.

Population ageing introduces an additional channel. As shown in [Lemma 2](#), the age-structure term θ_t acts as a taste shifter, altering how the planner weights the marginal utility of future generations. When the generational discount factor δ exceeds the private one β , the planner gives relatively more weight to future cohorts. As θ_t declines in an ageing population, the weighted marginal utility of future generations falls because future cohorts contain a larger share of elderly individuals. This weakens the incentive to save

and mitigate emissions, and the social discount rate increases—society effectively acts as if it were more impatient.

As population growth slows, its downward effect on the social discount rate weakens, while the ageing 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 modelled 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

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

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

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, (22) now becomes:

$$\delta^t \sigma_t = (1 - \gamma) \delta^{t+1} \sigma_{t+1} + \delta^{t+1} \lambda_{t+1} \frac{\partial F(S_{t+1}, K_{t+1}, N_{t+1}^Y, E_{t+1})}{\partial S_{t+1}} + \frac{\partial W(\mathbf{C})}{\partial S_{t+1}} \quad (26)$$

This expression now incorporates the first derivative of the social welfare function with respect to the stock of carbon in the atmosphere, which can be expressed as follows:

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

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 generalised version of the SCC:

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

$$\sigma_t^\diamond = \hat{\sigma}_t + v \sum_{j=1}^{\infty} (1 - \gamma)^{j-1} VS LY_{t+j} \left\{ \frac{\partial (N_{t+j} \theta_{t+j})}{\partial S_{t+j}} \frac{1}{N_{t+j} \theta_{t+j}} \right\} \quad (27)$$

where $\hat{\sigma}$ is given by (23), and

$$VS LY_{t+j} \equiv \delta^j \left(\frac{N_{t+j}}{N_t} \right)^v \left(\frac{\theta_{t+j}}{\theta_t} \right)^v \frac{u(C_{t+j})}{u'(C_t)}. \quad (28)$$

The additional term in the expression for the social cost of carbon includes $VS LY_{t+j}$, 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

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:

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

where b_t denotes the birth rate in period t , and $p_{h,t} \equiv \pi_{h,t}/\pi_{h-1,t}$ is the probability of surviving to age h conditional on being alive at age $h-1$. Using this, the semi-elasticity can be written as:

$$\left\{ \frac{\partial (N_{t+j} \theta_{t+j})}{\partial S_{t+j}} \frac{1}{N_{t+j} \theta_{t+j}} \right\} = \hat{\omega}_{0,t+j} \frac{\partial b_{t+j}}{\partial S_{t+j}} \frac{1}{b_{t+j}} + \sum_{h=1}^H \hat{\omega}_{h,t+j} \frac{\partial p_{h,t+j-h}}{\partial S_{t+j}} \frac{1}{p_{h,t+j-h}} \quad (29)$$

where

$$\hat{\omega}_{h,t+j} = \frac{N_{h,t+j} \left(\frac{\beta}{\delta} \right)^{h/v}}{\sum_{a=0}^H N_{a,t+j} \left(\frac{\beta}{\delta} \right)^{a/v}}, \quad (30)$$

can be interpreted as a modified population share of h -years old individuals.⁸ Therefore,

⁷This 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.,

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

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

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

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.

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 summarise the climate and economic parameters.

5.1 Parameterising Demographic Transitions.

Our analysis uses global demographic data from the United Nations' [World Population Prospects](#), 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 parametrise 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 modelling, 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.

occurs if $\beta > \delta$.




Figure 1: Global Population Growth

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

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 mean n^* and initial condition n_0 . Specifically, the model takes the form

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




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.

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^* .

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 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 com-

pared 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 global population remains positive. 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.

⁹The dip is assumed to occur off-sample and reflects structural forces beyond the observed range, so its parameters are imposed rather than forecast from historical data. We model the dip as a multiplicative modifier:

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

where d is depth, μ timing, and σ width. We use $d = 150$, $\mu = 240$, and $\sigma = 45$ (with $t = 0$ in 1960). These values were selected to (i) match the WPP low-fertility trajectory through 2100, (ii) generate a smooth but sufficiently negative growth episode, and (iii) maintain the non-negative long-run growth implied by the logistic component. The large d reflects that the baseline growth near $t = \mu$ is already close to zero, so a sizeable depth is needed to push net growth below zero while keeping the path smooth.




Figure 3: Global Survival Probability Profiles

Notes: Survival probability profiles from 1960 to 2023 (dark coloured lines) and projections through 2100 (light coloured 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 (grey dots), and 2100 (light grey dots, projected).

Survival Probabilities Survival probability data is organised 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 summarised 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 formal demography due to its ability to capture the exponential increase in mortality risk

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

¹¹This 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.* \(2001\)](#).

with age (see, e.g., [Preston *et al.* \(2001\)](#)). This implies a survival function of the form

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

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 $\{a_t, b_t\}_{t=1}^T$ summarises 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. We use this procedure because it provides a parsimonious decomposition with a fixed age pattern and a single time factor, so age-specific rates move coherently over time, and the common trend can be extrapolated with a simple univariate time-series model. 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:

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

where $\pi_{h,t}$ 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 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** summarises 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




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-specific survival probabilities and restricts attention to the adult population, conditional on survival to age 25.

date, our calculations assume that optimal policy internalises 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); 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

¹²When 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.

Production and damage functions follow [Goloso *et al.* \(2014\)](#). We use a Cobb-Douglas specification for total output,

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

with $\alpha = 0.34$. We set $\chi = 0.06$, which implies an energy share of $(1 - \alpha)\chi = 0.04$.¹³ This choice implies a labour share of $1 - \alpha = 0.66$, standard in macroeconomic calibrations. Damages enter multiplicatively as an output-loss share

$$1 - x(S_t) = \exp(-\xi(S_t - \bar{S})),$$

with $\xi = 2.379 \times 10^{-5}$.¹⁴

Labour 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 labour is allocated to energy production in the initial period, labour productivity ψ is set to equate this observed output to total energy production per unit of energy labour input. The level of total factor productivity A_0 is then calibrated to match global GDP in 2010.

The quantitative model uses the general carbon cycle in (1), setting $1 - d_s = \phi_L + (1 - \phi_L)\phi_0\phi^s$. In this specification, a share ϕ_L of emissions is long-lived (permanent); a share $(1 - \phi_L)\phi_0$ enters a slow-decaying pool with geometric decay factor ϕ ; and the remaining share $(1 - \phi_L)(1 - \phi_0)$ is removed immediately. The parameters ϕ_L, ϕ_0, ϕ and the initial stock $\bar{S}$ are calibrated following [Goloso *et al.* \(2014\)](#), adjusted to an annual frequency.¹⁵

To calibrate the carbon content of energy, we use data on energy use in 2010 reported

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.

¹⁴See [Goloso *et al.* \(2014\)](#) for calibration details.

¹⁵All theoretical expressions in the previous sections remain valid after replacing $(1 - \gamma)^s$ with $1 - d_s$.

Table 1: Climate and Economic Parameters

Parameter	Symbol	Benchmark
<u>Preferences and Technology</u>		
CRRA parameter	ν	1.0
Private discount factor	β	0.960
Generational discount factor	δ	0.985
Total Factor Productivity	A_0	1133
Labor Productivity in Energy	ψ	205
<u>Carbon Cycle</u>		
Pre industrial Stock of Carbon	$\bar{S}$	581 GtC
Initial Stock of Carbon	S_0	806 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
<u>Production and Damage Function</u>		
Capital share	α	0.34
Energy share	χ	0.04
Capital depreciation rate	δ_k	0.10
Damage Function Sensitivity	ξ	2.4e-05

in Golosov *et al.* (2014), originally based on the International Energy Agency, 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 increase over time and converges to 1 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 2125.¹⁶

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 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 dis-

¹⁶This 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.

counting of future damages in the low-fertility scenario.

To isolate the effect of ageing, 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 are reported in [Table 2](#).¹⁷ In the table, fertility scenarios reflect differences in population growth, while comparing fixed versus time-varying age distributions isolates the impact of population ageing.

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

¹⁷Note 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 ageing when computing the social cost of carbon and temperature.

Table 2: Demographic transitions versus Stationary scenarios

Demographic scenario	Initial SCC		Max SCC		Max Temp. in C°
	USD/tC	% RA	USD/tC	% RA	
<i>Representative agent (No Population Growth)</i>	49.5	100	58.3	100	3.1
Fixed Age Distribution					
Low Fertility	54.1	109	64.5	111	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	48.5	98	57.9	99	3.3
Medium Fertility	59.4	120	79.6	136	3.6
High Fertility	77.5	157	120.6	207	4.0
<i>Stationary 1% Population Growth</i>	129.3	261	286.5	491	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 Representative Agent (RA) benchmark, the peak SCC, its percentage relative to the RA benchmark, and the peak temperature.

are included here for comparison. Holding the age distribution constant over time yields higher initial SCCs across all fertility cases than the corresponding time-varying versions. Consistent with [Proposition 3](#), population ageing 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 ageing 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 highlights how different assumptions about intergenerational welfare weights can affect the present value of climate damages as the population evolves.

The results are summarised 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 ageing. 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

Table 3: Initial SCC and Peak Temperature

Demographic scenario	Generational Discount Factor δ				
	0.965	0.975	0.985 [†]	0.995	0.999 [‡]
<i>Initial SCC</i>					
<i>Representative agent (No Population Growth)</i>	23.2	31.3	49.5	129.4	465.5
Fixed Age Distribution					
Low Fertility	26.7	35.8	54.1	118.8	351.2
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	26.3	33.9	48.5	97.6	267.2
Medium Fertility	27.9	37.8	59.4	150.9	514.4
High Fertility	30.7	44.1	77.5	254.6	1098.1
<i>Peak temperature</i>					
<i>Representative agent (No Population Growth)</i>	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	2.2

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).


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 ageing on the social discount rate – widens as δ increases. In other words, the impact of ageing 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 $\delta = 0.999$, 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 ageing.

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 labour 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 labour allocated to energy produc-

A. Share of labour allocated to Energy



B. Savings Rate




Figure 6: Optimal Choices across Fertility scenarios.

tion, 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-




Figure 7: Optimal Emissions across Fertility Scenarios.

mark), the population effect dominates: higher fertility means a larger population, which more than offsets the reduction in labour 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: they always decline as fertility increases, but the decline is steeper when δ is high. This difference leads to higher total emissions as population increases when δ is low. 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 ageing 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 modelling 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 decentralise the optimal allocation in a market economy could yield further insights for climate policy design.

References

Aguiar, M., Amador, M. and Arellano, C. (2024). ‘Micro risks and (robust) pareto-improving policies’, *American Economic Review*, vol. 114(11), p. 3669–3713, doi:10.1257/

aer.20230128.

Arrow, K., Cropper, M., Gollier, C., Groom, B., Heal, G., Newell, R., Nordhaus, W., Pindyck, R., Pizer, W., Portney, P., Sterner, T., Tol, R.S.J. and Weitzman, M. (2013). 'Determining benefits and costs for future generations', *Science*, vol. 341(6144), pp. 349–350, doi:10.1126/science.1235665.

Arrow, K., Dasgupta, P., Goulder, L., Daily, G., Ehrlich, P., Heal, G., Levin, S., Mäler, K.G., Schneider, S., Starrett, D. and Walker, B. (2004). 'Are we consuming too much?', *Journal of Economic Perspectives*, vol. 18(3), p. 147–172, doi:10.1257/0895330042162377.

Barcons, S., Dávila, E. and Schaab, A. (2025). 'Intergenerational welfare assessments', Working Paper.

Barrage, L. (2019). 'Optimal dynamic carbon taxes in a climate–economy model with distortionary fiscal policy', *The Review of Economic Studies*, vol. 87(1), pp. 1–39, ISSN 0034-6527, doi:10.1093/restud/rdz055.

Belfiori, E., Carroll, D.R. and Hur, S. (2024). 'Unequal Climate Policy in an Unequal World', Federal Reserve Bank of Dallas.

Belfiori, E. and Macera, M. (2025). 'Demographic changes and social discounting', *Economics Letters*, vol. 256, p. 112584, ISSN 0165-1765, doi:https://doi.org/10.1016/j.econlet.2025.112584.

Bohn, H. and Stuart, C. (2015). 'Calculation of a population externality', *American Economic Journal: Economic Policy*, vol. 7(2), p. 61–87, doi:10.1257/pol.20130012.

Bourany, T. (2024). 'Climate change, inequality and the optimal climate policy', Working Paper.

- Bressler, R.D. (2021). 'The mortality cost of carbon', *Nature Communications*, vol. 12, p. 4467, doi:10.1038/s41467-021-24487-w.
- Calvo, G.A. and Obstfeld, M. (1988). 'Optimal time-consistent fiscal policy with finite lifetimes', *Econometrica*, vol. 56(2), pp. 411–432, ISSN 00129682, 14680262.
- Carleton, T.A., Jina, A., Delgado, M., Greenstone, M., Houser, T., Hsiang, S.M., Hultgren, A., Kopp, R.E., McCusker, K.E., Nath, I., Rising, J., Rode, A., Seo, H.K., Viaene, A., Yuan, J. and Zhang, A.T. (2023). 'Valuing the global mortality consequences of climate change accounting for adaptation costs and benefits', *The Quarterly Journal of Economics*, vol. 137(4), pp. 2037–2105, doi:10.1093/qje/qjac020.
- Cruz, J.L. and Rossi-Hansberg, E. (2023). 'The Economic Geography of Global Warming', *The Review of Economic Studies*, vol. 91(2), pp. 899–939, ISSN 0034-6527, doi:10.1093/restud/rdad042.
- Douenne, T., Hummel, A.J. and Pedroni, M. (2023). 'Optimal fiscal policy in a climate-economy model with heterogeneous households', University of Amsterdam.
- Eden, M. (2023). 'The cross-sectional implications of the social discount rate', *Econometrica*, vol. 91(6), pp. 2065–2088.
- Fried, S., Novan, K. and Peterman, W.B. (2024). 'Understanding the inequality and welfare impacts of carbon tax policies', *Journal of the Association of Environmental and Resource Economists*, vol. 11(S1), pp. S231–S260, doi:10.1086/732842.
- Gerlagh, R., Lupi, V. and Galeotti, M. (2023). 'Fertility and climate change', *The Scandinavian Journal of Economics*, vol. 125(1), pp. 208–252, doi:10.1111/sjoe.12520.
- Golosov, M., Hassler, J., Krusell, P. and Tsyvinski, A. (2014). 'Optimal taxes on fossil fuel in general equilibrium', *Econometrica*, vol. 82(1), pp. 41–88, doi:10.3982/ECTA10217.

- Golosov, M., Jones, L.E. and Tertilt, M. (2007). 'Efficiency with endogenous population growth', *Econometrica*, vol. 75(4), pp. 1039–1071, doi:<https://doi.org/10.1111/j.1468-0262.2007.00781.x>.
- Gourinchas, P.O. and Parker, J.A. (2002). 'Consumption over the life cycle', *Econometrica*, vol. 70(1), pp. 47–89.
- Hillebrand, E. and Hillebrand, M. (2019). 'Optimal climate policies in a dynamic multi-country equilibrium model', *Journal of Economic Theory*, vol. 179, pp. 200–239, ISSN 0022-0531, doi:<https://doi.org/10.1016/j.jet.2018.11.001>.
- Krusell, P. and Smith, Anthony A, J. (2022). 'Climate change around the world', National Bureau of Economic Research, doi:10.3386/w30338.
- Kydland, F.E. and Prescott, E.C. (1977). 'Rules rather than discretion: The inconsistency of optimal plans', *Journal of Political Economy*, vol. 85(3), pp. 473–491, ISSN 00223808, 1537534X.
- Kydland, F.E. and Prescott, E.C. (1982). 'Time to build and aggregate fluctuations', *Econometrica*, vol. 50(6), pp. 1345–1370.
- Lee, R.D. and Carter, L.R. (1992). 'Modeling and forecasting u.s. mortality', *Journal of the American Statistical Association*, vol. 87(419), pp. 659–671.
- Nordhaus, W.D. (2007). 'A review of the stern review on the economics of climate change', *Journal of Economic Literature*, vol. 45(3), p. 686–702, doi:10.1257/jel.45.3.686.
- Nordhaus, W.D. and Boyer, J. (2003). *Warming the world: economic models of global warming*, MIT Press.
- Preston, S.H., Heuveline, P. and Guillot, M. (2001). *Demography: Measuring and Modeling Population Processes*, Malden, MA: Blackwell.

- Ramsey, F.P. (1928). 'A mathematical theory of saving', *The Economic Journal*, vol. 38(152), pp. 543–559, ISSN 00130133, 14680297.
- Schneider, M.T., Traeger, C.P. and Winkler, R. (2012). 'Trading off generations: Equity, discounting, and climate change', *European Economic Review*, vol. 56(8), pp. 1621–1644, ISSN 0014-2921, doi:<https://doi.org/10.1016/j.euroecorev.2012.08.006>.
- Stern, N. (2007). *The Economics of Climate Change: The Stern Review*, Cambridge University Press.
- Stern, N. (2008). 'The economics of climate change', *American Economic Review*, vol. 98(2), pp. 1–37, doi:10.1257/aer.98.2.1.
- Van der Ploeg, F. and Rezai, A. (2021). 'Optimal carbon pricing in general equilibrium: Temperature caps and stranded assets in an extended annual dsge model', *Journal of Environmental Economics and Management*, vol. 110, p. 102522, ISSN 0095-0696, doi:<https://doi.org/10.1016/j.jeem.2021.102522>.
- Weitzman, M.L. (2007). 'A review of the stern review on the economics of climate change', *Journal of Economic Literature*, vol. 45(3), p. 703–724, doi:10.1257/jel.45.3.703.
- World Population Prospects (2024). 'World population prospects 2024', Online Edition, accessed May 7th, 2025.

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_{t+1}^0 be the growth rate of newborns between period t and $t + 1$. This means that:

$$N_{0,t+1} = (1 + n_t^0)N_{0,t}. \quad (31)$$

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

$$N_{h,t+h} = \pi_{h,t}N_{0,t}. \quad (32)$$

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

$$N_t = \left((1 + \eta_{t-H}^t) + (1 + \eta_t^{t+H-1})\pi_{1,t-H} + \dots + \pi_{H,t-H} \right) N_{0,t-H} \quad (33)$$

for all t , where η_t^{t+s} denotes the *cumulative* population growth rate of newborns between period t and $t + s$. Therefore, (33) 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 (12) into the social welfare function (11) yields

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

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

$$\sum_{t=0}^{\infty} \sum_{h=0}^H \frac{\delta^{t-h} N_{0,t-h}}{\sum_{s=0}^{\infty} \delta^s N_{0,s}} \beta^h \pi_{h,t-h} u(c_{h,t}) \quad (35)$$

Using (6), 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 (8) 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 normalise 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 (13).

C Proof of Proposition 3

Using the CRRA assumption and taking logs in (24) delivers:

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

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.