Impact of population growth and population ethics on climate change mitigation policy
SI Appendix

Although the description of our method allows for all the DICE results to be reproduced exactly – the model and population estimates are freely available – we also include as Supporting Information the full DICE modeling and results in editable Excel spreadsheets.

Table of Contents

Section 1. Formal implementation of average and total utilitarianism

In general, many social objectives, also called welfare functions (SWFs), are available in the population ethics literature, with varying approaches to social valuation of population size. We study two forms, designed to be representative:

- **Average utilitarianism (AU):** the type of SWF which does not regard extra people as a social benefit, all else equal, but instead averages well-being over the set of people which exist.
- **Total utilitarianism (TU):** the type of SWF which considers total well-being, and is therefore increasing in the size of the population and in the level of average well-being. TU would regard extra people as a social benefit, all else equal, such that a larger population represents a social improvement, provided that the extra people are living lives worth living (adding to the total stock of well-being) and unless it is sufficiently offset by a decline in average well-being.

As O'Neill and Wexler have discussed in a paper considering the climate externalities of having children, "in general, comparing welfare across different population sizes introduces profound theoretical issues" (10 p. 314). Both AU and TU have advantages and disadvantages, which are well-explored in the economics and philosophy literature on social choice and welfare.

Despite this plurality in the population ethics literature, cost-benefit climate-economy models essentially share the same TU equation for calculating present-discounted social welfare, W, at a given point in time t:

$$W^{TU} = \sum_t \left(\frac{L_t}{(1+\rho)^t} \times \frac{c_t^{1-\eta} - 1}{1-\eta} \right) \qquad \text{(Equation S1: TU)}$$

Here L_t is the population in period t, ρ the pure rate of time preference, and c_t the average consumption in time t, scaled to wellbeing by the inequality aversion parameter η. Note that in Equation S1, increasing the future population has an effect conceptually (and formally) similar to decreasing the rate of pure time preference, ρ, and thus gives more weight to the future in the calculation of the optimal trade-off between mitigation and damages. This is the equation that we use for TU.

In contrast with the straightforward equation for TU, in an intertemporal context, "average utilitarianism" is ambiguous (19, p. 213-214). This ambiguity is amplified when time discounting is introduced. Consequently, we use a form of average utilitarianism that involves applying time discounting to the average period wellbeing calculated by the model, with the objective of maximizing the sum of these discounted average period utilities – this is the SWF we describe as AU, which is related to the version used by Asheim (20):

$$W^{AU} = \sum_t \left(\frac{1}{(1+\rho)^t} \times \frac{c_t^{1-\eta}-1}{1-\eta}\right) \qquad \text{(Equation S2: AU)}$$

An alternative, more orthodox, implementation of average utilitarianism would simply divide the total sum of wellbeing by the total size of the present and future population; such an approach would ignore discounting, and show no pure time preference. Another alternative, called "dynamic average utilitarianism" by Dasgupta (19) and used by Arrow, et al. (21), would divide the discounted TU objective function by the total discounted population of all present and future generations; because such a dynamic AU SWF would deliver the same results as TU in our exogenous-population analyses, and the same result as our version of AU when population size is endogenous, we do not further consider this approach.

However, in a further sense, even these alternatives would not be the most theoretically pure version of AU, even ignoring the complicating role of time discounting. This is because AU in its most pure form would average the wellbeing of the entire human population across history in the future and past. On this view, whether or not creating future people is a social improvement depends, in part, on whether or not their lives are better than the lives of people who lived and died millennia ago.* TU – like the implementation of AU that we use in our paper – does not have this dependency on the distant past because its additive structure separates the evaluation of future wellbeing from the evaluation of past wellbeing.

Some generalized versions of TU involve a "critical level," which is a minimal level of wellbeing for adding an extra life to be regarded as a social improvement by the social welfare function: an extra life at a level of wellbeing below the critical level (even if it has positive utility) would lower the SWF. Our implementation of TU assumes a critical level equal to a GDP per capita of $3,775 per year (the default value in DICE). We note that any of our TU results for the SCC or optimal peak temperature would be numerically identical if a critical level were changed to any other (typically positive) number, because the additive critical level would not appear in the first order condition of the optimization. This is because DICE's SCC optimization procedure does not endogenize the size of the population. Critical levels are discussed further in Extension 4 below.

* "On this view, whether or not a person should have a child now depends in part on how well off – and how numerous – people were in the Stone Age, and it is difficult to see why this sort of consideration should be relevant." p. 115. McMahan, Jefferson. "Problems of population theory." *Ethics*. (1981): 96-127.

Section 2. Methods: further details

DICE and RICE modeling details

We use the Excel version of the DICE2013 climate-economy model as downloaded from William Nordhaus's website, which is freely available online (http://aida.wss.yale.edu/~nordhaus/homepage), and has been described in detail elsewhere (15, 24).

Briefly, DICE is a global (single-region) optimization model that includes an economic and a geophysical component that are linked. Economic activity produces emissions, which are a function of GDP (output) and a time-varying ratio of emissions-to-output, as well as emission control policies; carbon intensity is exogenous. Population influences emissions by influencing output via a Cobb-Douglas production function.

If unmitigated, emissions affect the future economy through climate-related damages, which increase with temperature and are incurred as the loss of a percentage of output. More specifically, DICE's climate damages are based on 13 published studies that have estimated the magnitude of monetized impacts that would occur given a specified, discrete increase in temperature (36). The continuous (with temperature) DICE damage function was parameterized by assuming that damages are a quadratic function of temperature and that there are no tipping points. The number and types of impacts assessed in the underlying studies varied, but normally included changes in agricultural productivity, human health (mortality and/or morbidity from climate-sensitive diseases), and damages from extreme weather events and sea-level rise, amongst other factors.

The key tradeoff in the model, therefore, is between mitigation, which incurs a cost relatively soon, and climate damages, which incur costs in the more distant future. The standard optimization maximizes the sum of discounted wellbeing across time (i.e., has a TU SWF as in Equation S1). When a temperature target is added, the optimization is constrained to first ensure that temperature rise in the future stays at or below that target, and then the objective is maximized subject to that constraint.

Like all leading climate economy models, the size of the population is an exogenous variable in DICE (meaning that it is predetermined and that the model therefore has perfect foresight in this regard).

Other than the changes to population that are essential to our experiments, the DICE2013 model is unchanged (including default parameter values, such as a climate sensitivity of 3.2) with one exception: we specify an exogenous savings rate of 25.8%, which can be interpreted as the optimal savings rate of private savers with a time-separable and discounted objective with a logarithmic utility function (33,39). To explain this approach, we first note that there are two alternative treatments of savings in the climate-economy modeling literature: one approach that assumes that economic agents endogenously look forward to climate damages and policies and optimally adjust their planned savings (a leading example is in the official versions of DICE/RICE), and another approach that assumes that savings do not so respond to climate policy optimization (leading examples are FUND and PAGE; in a DICE/RICE framework, see (33). We prefer the second approach, which leads to our departure here from the official DICE model, on the grounds that we find it more realistic to assume that society has a fixed appetite for savings that is essentially insensitive to climate change and climate policy decisions. In Figure S1, we demonstrate that our results are substantively identical and quantitatively very similar if we instead endogenize optimal savings as in the official DICE model.

In certain modeling runs, we also make additional changes to DICE in order to investigate our research questions. Most importantly, to generate results for the runs maximizing AU, we altered the SWF accordingly as described in Equation S2 of the previous section. To explore the relationship of population growth and the rate of pure time preference (Figure S6 below), we changed the latter to the value that would make the optimal mitigation trajectories of the UN-medium and UN-low population scenarios provide the closest fit to that of the UN-high scenario. Here, closest fit is understood as the assumption for time preference that minimizes the differences between carbon prices squared in the resulting optimum and the default UN-high optimum, summed over all time periods.

For all analyses, DICE2013 initializes in 2010 with identical carbon prices (= $1). Although the model has a 300 year time horizon, we do not present results over the whole period since full mitigation occurs much earlier in almost all analyses.

In a few analyses found only here in the Supporting Information (not in the main text), we present results derived from a regionalized version of DICE known as RICE. As opposed to DICE, which is a single-region (global) model, RICE is regionalized and therefore has region-specific parameters including, for example, total factor productivity and emission intensities as well as population growth, which we aggregated from UN country-level data. Climate damages are also region-specific and taken from reference (24). The carbon price however, is globally harmonized.[†] The standard implementation of RICE has been described in detail elsewhere (24, 15). We make three additional modifications to RICE2010 (the most recent version of RICE). The first is to eliminate the use of Negishi weights, which were introduced originally to restrict redistribution and ensure that it does not become a policy tool. In our version of RICE, we do not use Negishi weights, as described in a previous publication (33). Second, we fix the savings rate at 25.8%, with the same rationale as above for our DICE modeling. In light of this, our implementation might more aptly be referred to as "RICE with fixed savings and without Negishi weights". Third, we substitute a reduced form version of the RICE2010 sea level rise damage module in our implementation; see (33) for details. The countries comprising the twelve regions can be found in the Supporting Spreadsheet, as can all summary results from these analyses using our variant of RICE.

[†]Although a global harmonized carbon price is considered the most economically efficient carbon price solution, and is the default in RICE, studies have modified cost-benefit climate-economy models to allow for regionally determined carbon prices. For two examples, see Anthoff, D (2009) Optimal Global Dynamic Carbon Taxation, (ESRI), WP278, and Budolfson M, Dennig, F (in press) Optimal Global Climate Policy and Regional Carbon Prices. In Chichilnisky, G, Sheeran, K, & Rezai, A eds., Handbook on the Economics of Climate Change, (Edward Elgar).

Population projections

In this study, we compute optimal carbon taxes and thus mitigation paths by exogenously specifying a variety of population trajectories, based primarily on the 2015 revision of the United Nations' World Population Prospects, and also on the Shared Socioeconomic Pathway project (30, 37, 38). Both of these sources provide a range of population estimates through 2100. To project beyond 2100 in cases where we use a projection that stops at 2100, we assume that the population growth rate in the time-step ending in 2100 tapers linearly to zero between 2100 and 2195 and remains constant thereafter.

This assumption means that there are only relatively modest changes in population after 2100. Therefore, to explore the possibility of a more dramatic decline in population, we also include an additional "Ultra low" population scenario based on results in Basten et al. (31) that assumes a life expectancy converging on 100 and a fertility of 1.5, which is slightly lower than the European Union's fertility today (1.58). The other scenarios reported in the main text are based on the UN-medium scenario, which is the scenario considered most likely by the UN, and the UN-high and UN-low scenarios, where fertility is projected to be 0.5 above and below the UN-medium, respectively. All 16 of the scenarios analyzed are presented graphically in Fig. S2.

Computation of mitigation cost savings

To compute mitigation cost savings (as displayed in Figure 2 of the main text and Tables S2-S4), we program DICE with the relevant AU or TU objective and then maximize that objective subject to the constraint that temperature must never rise above the target.

To calculate the per capita mitigation cost savings between a higher and lower population scenario in a given time period, we first calculate the difference in cost as a percentage of per capita GDP and then multiply that difference by the per capita GDP in the lower scenario, as follows:

$$\left(\frac{Mpc_{higher}}{GDPpc_{higher}} - \frac{Mpc_{lower}}{GDPpc_{lower}}\right) * \ GDPpc_{lower} \qquad \text{(Equation S3)}$$

where Mpc refers to the mitigation cost per capita and $GDPpc$ refers to the GDP per capita. To calculate total mitigation cost savings, the per capita cost savings is then multiplied by the population in the lower scenario (see Table S2 for an example calculation). We believe a comparison normalized to the percentage of per capita GDP in this way is more appropriate than one based on absolute cost differences, because the two societies would have different sized economies (including in per capita terms, for example because of a different capital-labor ratio) that would in turn be shared by a different number of people. For example, without normalization, the difference in the absolute total cost of abatement in 2050 between UN-medium and UN-low with a 2°C constraint is approximately 370% higher than reported in the top section of Table S3, which is misleading because costs are shared by a larger population in UN-high than in UN-low. Similarly, without normalization ,the difference in per capita costs in 2050 is approximately 45% lower, which is misleading because costs are borne by a population that is poorer per capita in UN-medium than in UN-low. (These calculations with normalized cost savings are used to generate numbers displayed in Figure 2 of the main text and Tables S2-S4. A comparison of cost savings with and without normalization can be found in Table S3, which demonstrates that our preferred normalization is quantitatively very similar to a simple per capita normalization .)

The "consumption effect" elasticity: The effect of population on emissions pressure

The consumption effect of a larger population on emissions, through expanded economic activity, is a largely empirical question with an important existing literature. Our paper – with its focus on social valuation – builds upon this literature, which we review briefly here along with its implications for our modeling framework.

O'Neill et al. (7) summarize the results on the relationship between population and emissions that were obtained by various estimations of the STIRPAT (Stochastic Impacts by Regression on Population, Affluence and Technology) equation. The authors conclude from the existing evidence that CO_2 emissions from energy respond roughly proportionally to population, but argue in favor of improved regression methodologies going forward, and in particular for panel regressions taking account of time series effects. Adopting this methodology, Liddle (6) and Casey & Galor (5) find that population may have a much greater influence on emissions than income per capita.

One issue therefore, is whether the economic effect of population on emissions is conceived in a reduced form way that incorporates all mechanisms, or whether the effect is estimated through particular economic pathways. We use the reduced-form DICE model, in which the elasticity of emissions is equal to one with respect to income and to 0.7 with respect to population (this is the coefficient of labor in the production function). However, the Liddle (6) and Casey & Galor (5) results may be more comparable with ours if what drives their estimates is the fact that emission intensity is influenced by income growth (via learning and technology replacement effects) and by population trends (via education being associated both with lower population and lower emission intensity). These mechanisms are imperfectly captured in our model via a trend on intensity. Our results can therefore be taken as capturing the most basic effects of population and may underestimate the potential climate benefits of reducing population.

Section 3. Further discussion of optimal peak temperature results

Our finding that optimal peak temperatures are *lower* given *higher* population under the standard TU social objective is a striking result that contrasts with an important previous study, (1), that used a much older vintage of DICE, DICE1994. The explanation of the sign difference between our results and the previous study (1) is that due to a modeling error, DICE1994 erroneously assumed that the cost as a proportion of GDP of mitigating a given fraction of business-as-usual emissions was increasing over time, whereas the correct assumption as in DICE2013 is that it decreases. As acknowledged over a decade ago by the architect of DICE, William Nordhaus, the older mitigation cost assumption was a mistaken modeling assumption ("The basic functional form for the abatement-cost function follows the structure assumed in the earlier DICE models. However, the structure has been reformulated over time to correct for an earlier modeling mistake. …The prior version used a functional form that implicitly and mistakenly assumed that the cost of the backstop technology increased over time. … Although this new specification makes little difference in the short run (to the tactics of climate policy, so to speak), it turns out that it makes a major difference over the long run (to the strategy or vision)" (42, p. 52-3)).

This error in the earlier model explains the sign difference between our results and (1), as importing the older, mistaken backstop assumptions (and as a direct consequence, mistaken mitigation cost assumptions) into DICE2013 yields results analogous to (1). In our results with the corrected mitigation cost assumptions, increasing the future population increases the SCC under both TU and AU, which accelerates the date of full mitigation (100% decarbonization), causing peak temperature rise to be lower and occur sooner. This effect is stronger under TU, because of the population weighting effect. When the pre-1999 vintages of DICE are run, or if their mistaken mitigation cost assumptions are imported into DICE2013, full mitigation never occurs, which accounts for the sign difference. If one mistakenly built into the modelling assumptions that full mitigation would never occur, then higher population would indeed lead to higher peak temperature, as the higher social cost of carbon would then not have the correct result of entailing faster full mitigation.

Section 4. Evidence and debates on population policy

A large literature studies the effectiveness of population and development policies on population size and growth, in ways that are relevant to climate change. One way to think about the relevance of this literature is that population policy can only be applicable to climate policy if it indeed can influence population growth rates, at acceptable costs and without other undesirable consequences.

Our paper, which does not do original econometric analysis of empirical data, employs estimates (especially of cost savings) from this literature. Because space is limited in the paper, we elaborate here on the evidence and debate in some papers in this literature.

There is considerable debate about the causes of the fertility reduction of the demographic transition, and in particular about the extent to which fertility reduction has been caused or accelerated by policy efforts, rather than by endogenous changes in behavior, such as responses to mortality decline or the importance of human capital (4). Connelly (41), describing the history of population policy, describes how some coercive policies have caused harm, especially to populations in developing countries. We abstract from much of this debate: our results are not intended to quantify full the social welfare implications of such programs, inclusive of any such important social costs.

Rather, in the spirit of Abel, et al.'s (43) investigation of a subset of the Sustainable Development Goals, we refer the reader to evidence that development programs that promote human development, child health, social equality, and women's social status also can lead parents to freely choose lower fertility. They find that the achievement of the SDGs would lead to a world population in 2100 of around 8.2-8.7 billion. This is substantially below the UN-lower 95th projection, leading Abel, et al to suggest that the UN prediction range is therefore too narrow. Wheeler and Hammer (11) show that family planning and women's education is a very cost-effective mitigation policy, via reduced population growth, compared to other options such as solar or wind. O'Neill (9) and O'Neill and Wexler (10) make similar computations and arguments about fertility policy.

What matters for policy is not merely that reducing population has climate consequences, but also the quantitative comparison of costs. We cannot know these hypothetical future costs in detail, but build upon estimates in the literature to understand their general magnitude. The spending shortfall for providing all women in the developing world with access to modern contraception was recently estimated at about $5 billion per year (44) and providing all children in low and lower-middle income countries a quality education at $40 billion annually (45). Both of these policies would be non-coercive, leaving women no less free than before to choose their fertility, and perhaps freer.

Along those lines, a fundamental aspect that we have not discussed in detail is that the non-climate benefits of policies that lower fertility – including better maternal and child health, improved gender equality, and more human capital – could be large, and could exceed the benefits from any avoided mitigation costs. Our observation, joining this existing literature and quantifying it in the context of a leading CEM, is that any successful low-fertility program would have a reduction in climate mitigation costs and damages to count among its other benefits, many of which are highly desirable in their own right.

Section 5. Supplementary exhibits (Figures and Tables)

This section presents all the supplementary figures and tables referred to in the main article. Additional exhibits from several extensions are reported in the next section. For all exhibits, data for the full multi-century time horizon can be found in the Supporting Spreadsheet, which is editable and optimizable by the user.

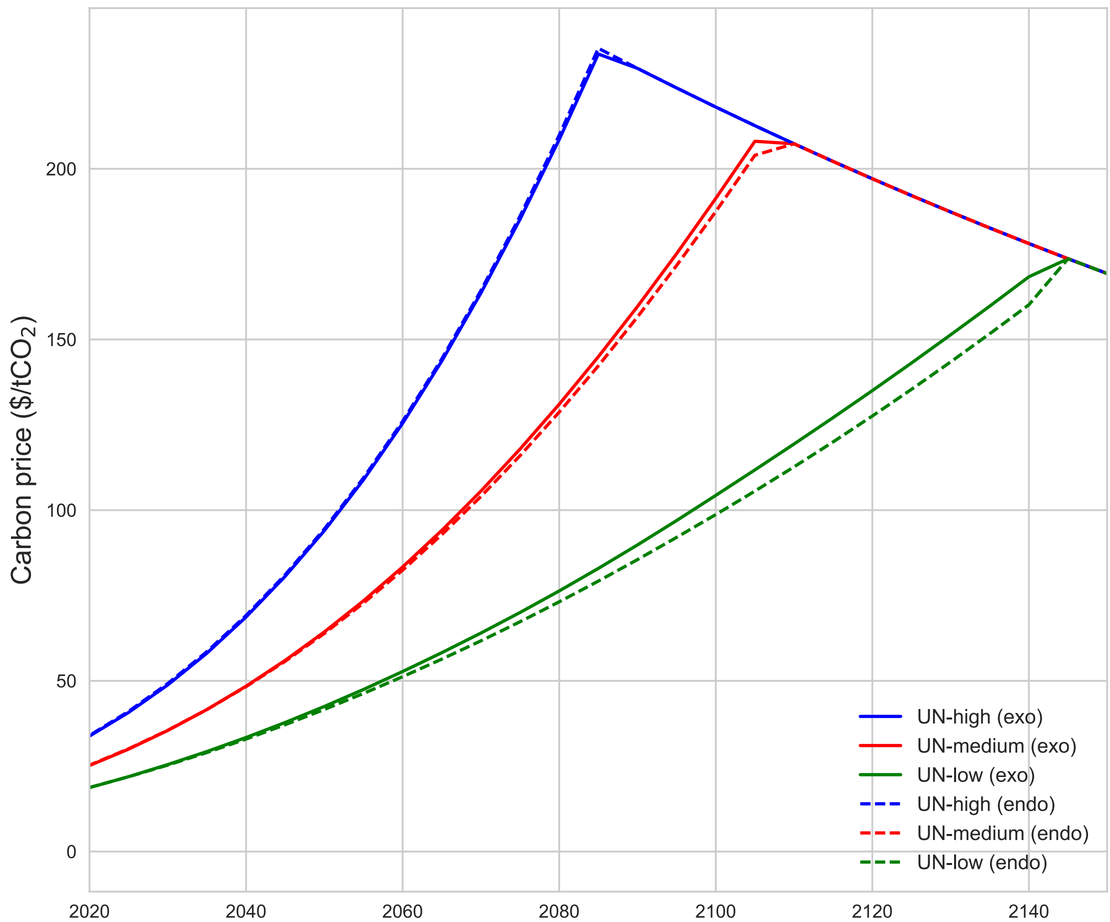

Fig. S1. Optimal carbon prices with an exogenous (exo) savings rate fixed at 25.8% and an endogenously (endo) determined savings rate under TU.

Notes on Fig. S1

In all analyses in the paper, we use a fixed savings rate, but standard DICE uses an endogenously determined savings rate (see Methods). This figure demonstrates the similarity of the two approaches, illustrated with a total utilitarian objective.

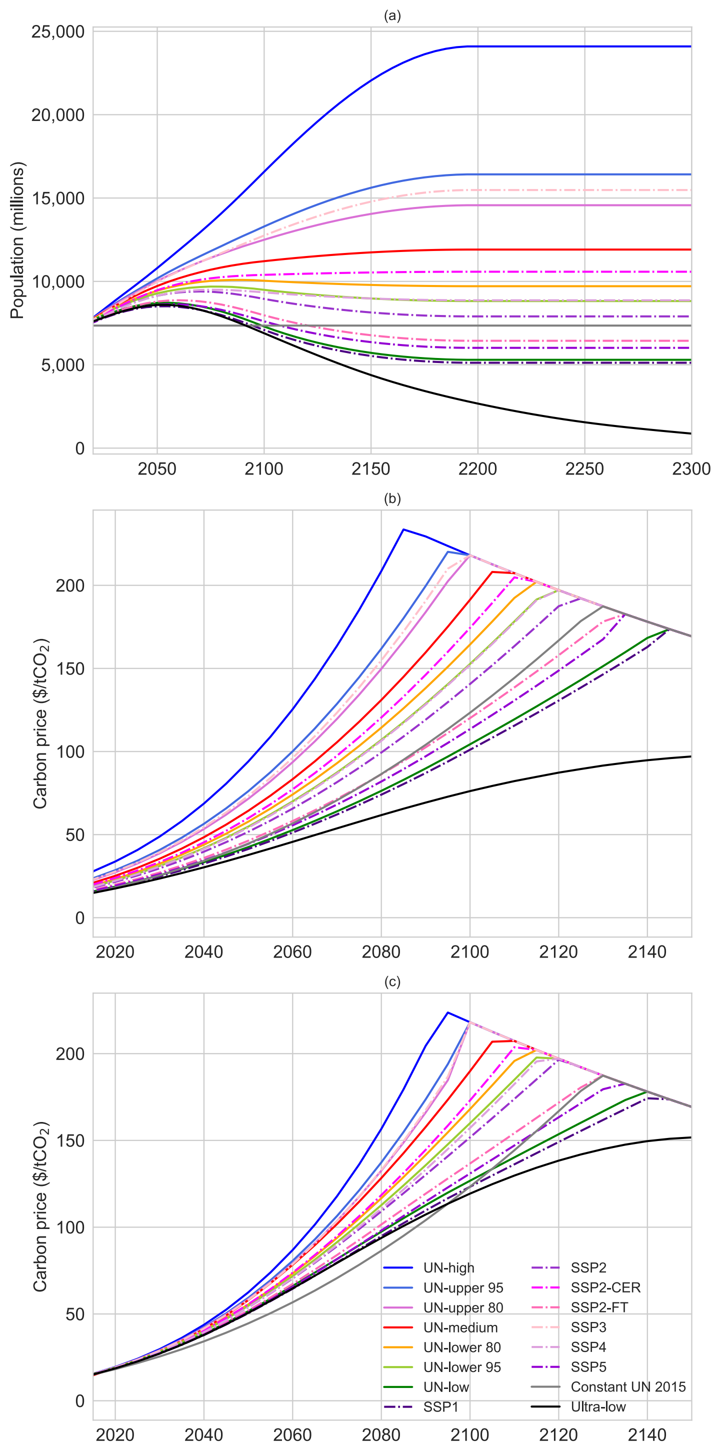

Fig. S2. (a) Population projections analyzed in DICE, based on data from the United Nations (UN) and the Shared Socioeconomic Pathway project (SSP) and an "Ultra-low" scenario, (b) optimal carbon prices with a TU social objective and (c) optimal carbon prices with an AU social objective.

Notes on Fig. S2: The UN and SSP data can be found in references (30,37) while the "Ultra-low" scenario is based on reference (31) and assumes that global fertility converges to 1.5 with a life expectancy of 100. Optimal carbon prices are reported with TU objective only for illustration.

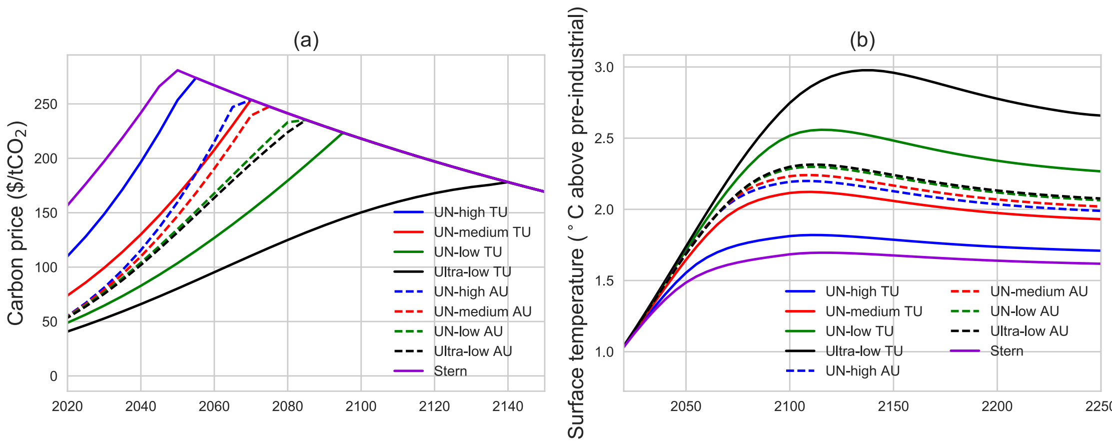

Fig. S3. (a) Optimal carbon taxes and (b) temperature rise assuming near-zero time preference (0.1% per annum).

<u>Notes on Fig. S3</u>

Here the "Stern" tax path is based on the UN-medium population with the parameter values assumed in the Stern Review (16), which includes near-zero time preference as well as low inequality aversion (1.01) and uses a TU social objective. These results show that even with UN-high population, mitigation effort with near-zero time preference never reaches Stern's rates in the near-term when inequality aversion is relatively high (1.45 vs. Stern's value of 1.01), but yet at the same time, the point of full mitigation is similar (2050 vs 2055 for Stern and UN-high, respectively).

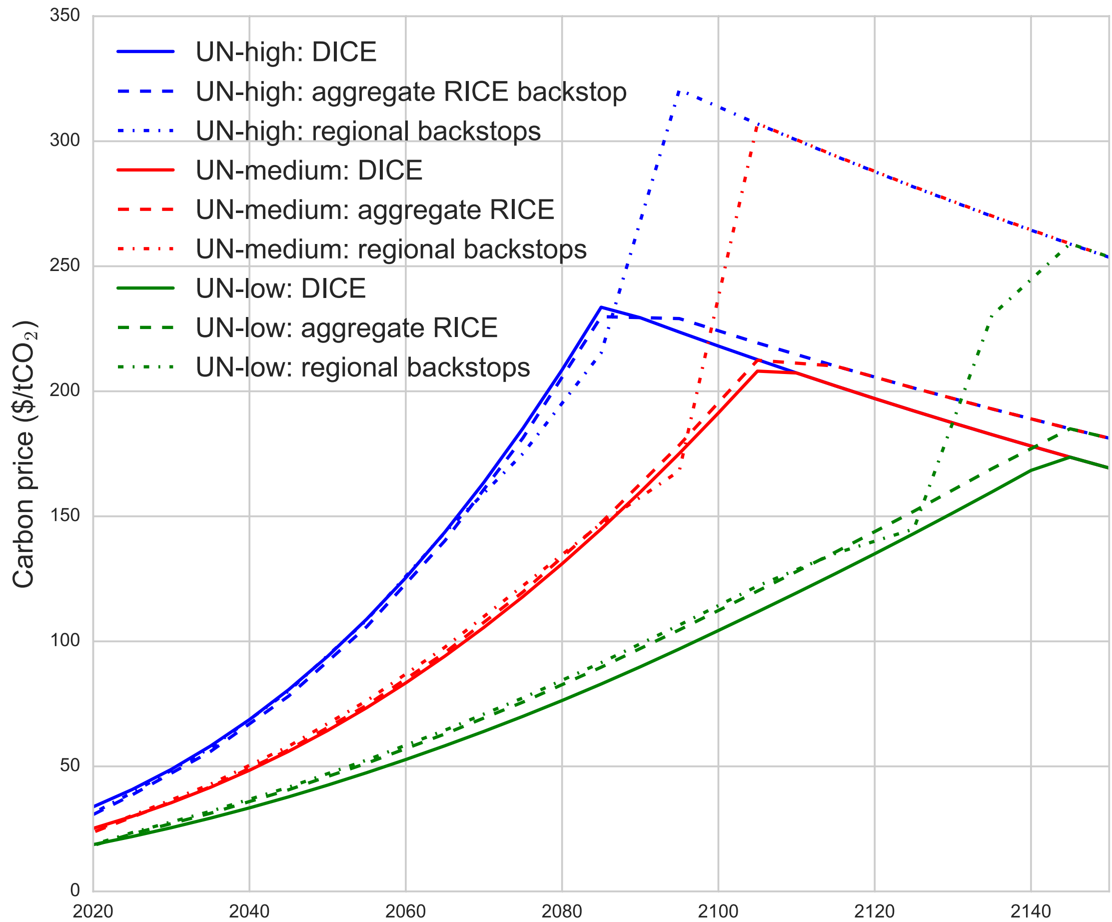

Fig. S4. Comparison of optimal carbon price pathways for the UN-high, -medium, and –low population scenarios in DICE vs. the regionalized version of DICE known as RICE. Here we compare the results in our implementation of DICE2013 (as presented in the main text) to results in a modified version of RICE2010 (modifications described above).

Notes on Fig. S4

Here we present two sets of results from our implementation of RICE2010. The first uses the standard RICE assumption that each region has its own backstop price (one of which is substantially higher than DICE's global backstop). The second uses a single globally aggregated backstop price from RICE and is presented for comparative purposes only, as it is similar to (but slightly higher than) the DICE2013 backstop. All of our regionalized analyses use the standard RICE assumption of different regional backstops. This is robustness illustration uses the standard total utilitarian objective.

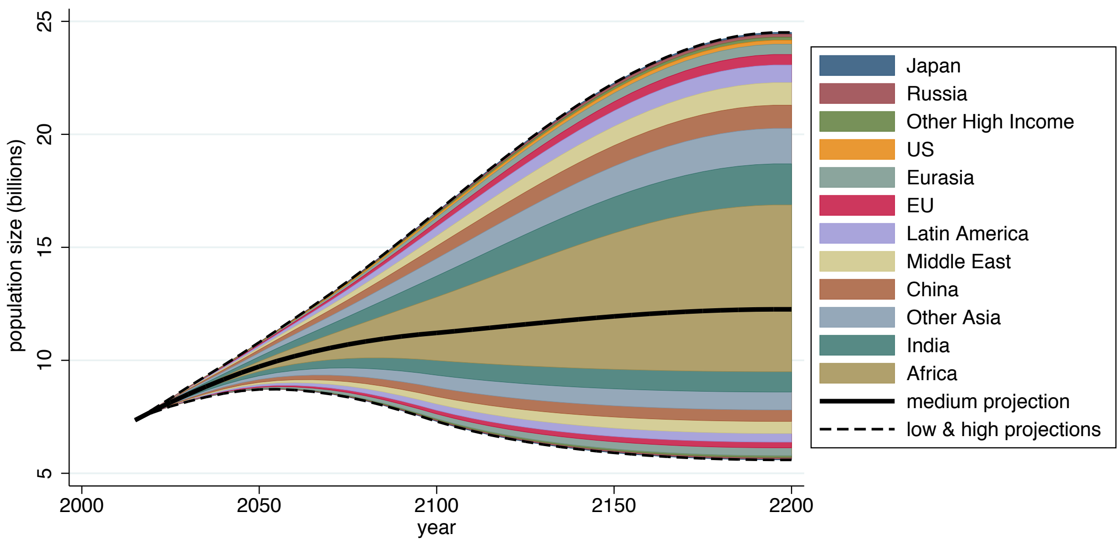

Fig. S5. Global population as projected in the UN-high, - medium and –low scenarios. Color bands show the contribution of each region represented in RICE to the difference between the three scenarios.

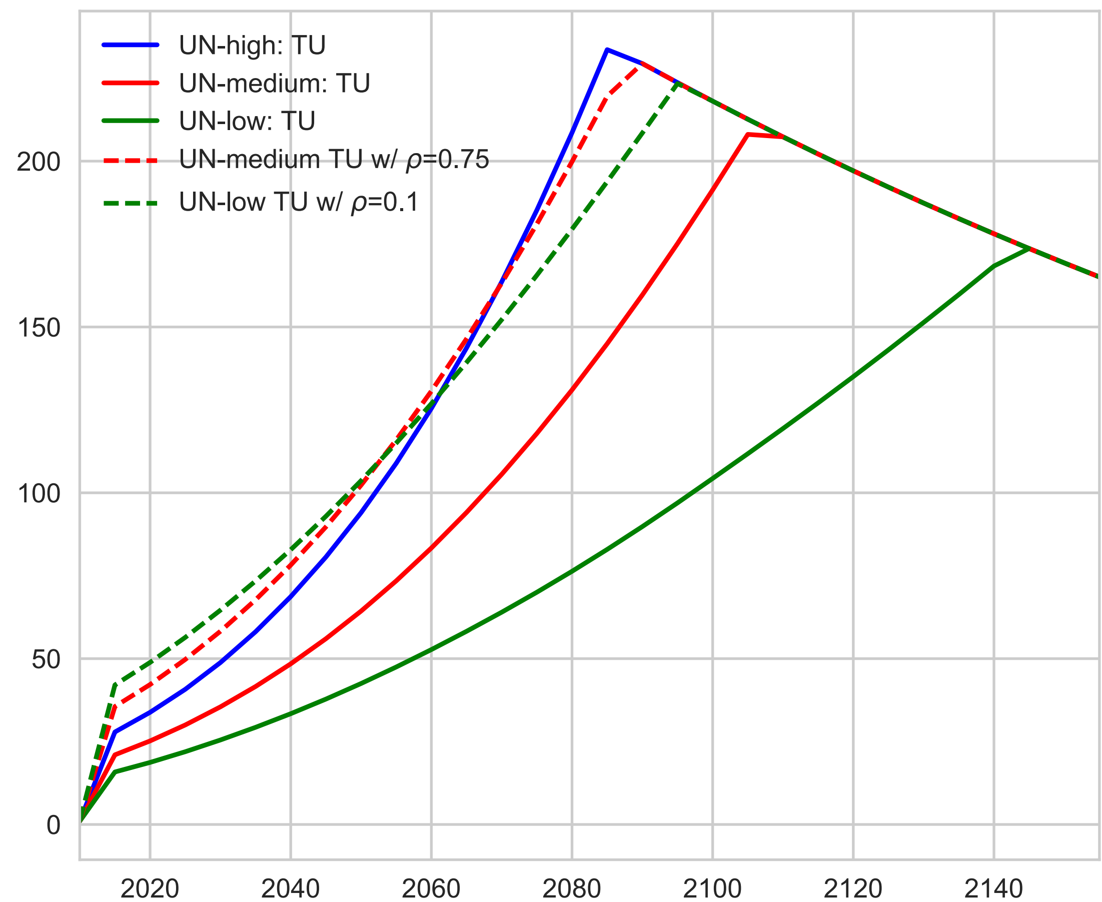

Fig. S6. Optimal carbon prices for three UN population scenarios estimated with the standard TU social objective and using the default rate of pure time preference (ρ) of 1.5% per annum, as well as two cases with different rates of pure time preference (0.75% and 0.1%) for comparison (while keeping inequality aversion = 1.45).

<u>Notes on Fig. S6</u>

To explore the relationship of population growth and the rate of pure time preference, we changed the latter to the value that would make the optimal mitigation trajectories of the UN-medium and UN-low population scenarios provide the closest fit to that of the UN-high scenario. Here, closest fit is understood as the assumption for time preference that minimizes the differences between carbon prices squared in the resulting optimum and the default UN-high optimum, summed over all time periods.

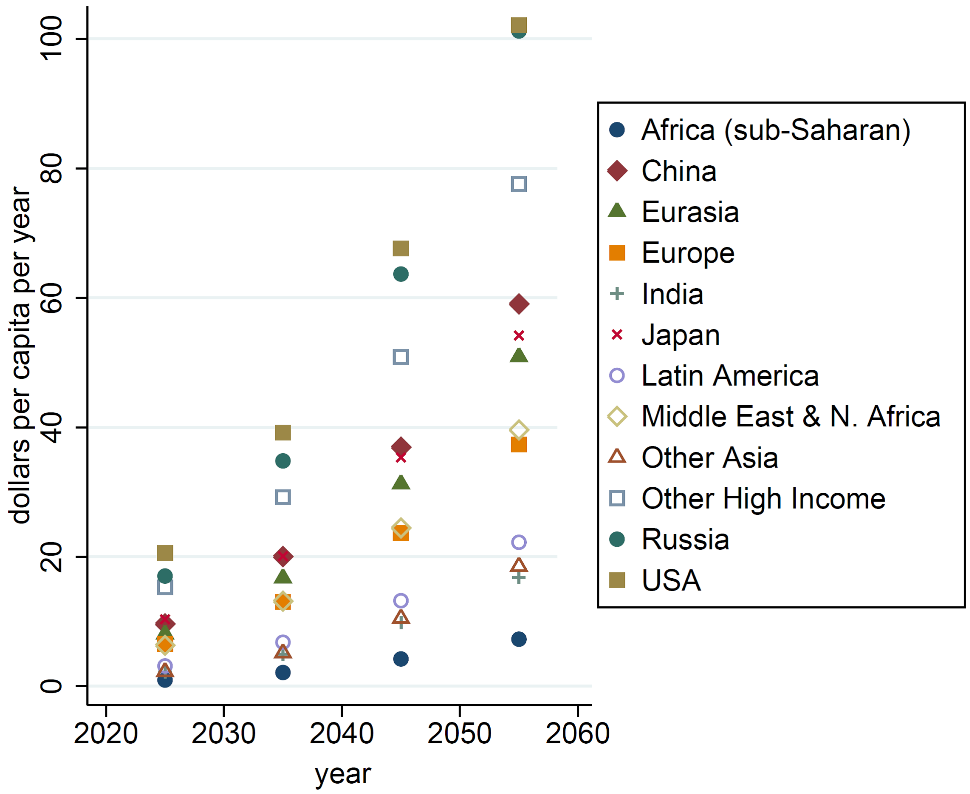

Fig. S7. Per capita abatement cost savings by region under the UN-low compared to the UN-medium population scenario given a 2 °C temperature target. (We show results only with a TU objective because there are negligible cost differences in the near term with an AU objective.)

Table S1. Optimal carbon prices (\$/ton CO_2) for four select population scenarios and the change in the UN-medium price when only specified regions follow the UN-low scenario. Results derived from a variant of the RICE model using the standard TU objective.

	2025	**2045**	**2065**
Price when all regions have UN-medium population	$30.37	$58.01	$97.63
Price when all regions have UN-low population	$23.53	$41.66	$64.45
Price when only developing regions have UN-low population*	$23.97	$43.23	$68.50
Price when all regions have SSP1 population**	$21.47	$37.96	$59.16
Percent change from UN-medium carbon price (above) if only indicated region has UN-low population while all others retain UN-medium population			
Africa (sub-Saharan)	-10.9	-10.9	-16.2
China	-2.5	-2.5	-4.1
Eurasia	-0.1	-0.1	-0.1
Europe	-0.5	-0.5	-1.2
India	-3.0	-3.0	-2.6
Japan	-0.1	-0.1	-0.3
Latin America	-0.6	-0.6	-1.2
Middle East and North Africa	-1.3	-1.3	-2.3
Other Asia	-2.1	-2.1	-1.9
Other High Income	-0.2	-0.2	-0.6
Russia	-0.1	-0.1	-0.4
USA	-0.4	-0.4	-1.9

* Here "developing" refers to Africa, China, India, Latin America, Middle East/North Africa and Other Asia. All other regions have UN-medium population.

** Applies to the global total population and the regional composition.

Table S2. The information in this table illustrates the calculation of the normalized avoided mitigation costs between scenarios (here given a constraint of a 2 °C increase in global temperature, and with a total utilitarian social objective). Population is in millions and GDP in 2005 US Dollars. Values rounded. (We illustrate with a total utilitarian objective because near-term cost differences are negligible with an average utilitarian objective, as shown in Figure 2 of the main text.)

		2020	2030	2040	2050
UN-medium	Population	7,758	8,501	9,157	9,725
	GDP per capita*	11,649	14,570	17,936	21,691
	Abatement cost per capita	42	99	224	481
UN-low	Population	7,689	8,180	8,532	8,710
	GDP per capita*	11,685	14,704	18,163	22,045
	Abatement cost per capita	38	93	214	471
UN-lower 80th	Population	7,732	8,411	8,984	9,440
	GDP per capita*	11,661	14,605	17,996	21,785
	Abatement cost per capita	41	98	221	478
UN-lower 95th	Population	7,718	8,360	8,889	9,284
	GDP per capita*	11,669	14,625	18,029	21,838
	Abatement cost per capita	41	97	220	477
SSP2	Population	7,612	8,256	8,772	9,140
	GDP per capita*	11,718	14,648	18,048	21,868
	Abatement cost per capita	40	95	216	471
SSP2FT	Population	7,567	8,131	8,546	8,792
	GDP per capita*	11,740	14,696	18,127	21,987
	Abatement cost per capita	39	93	213	465
SSP1	Population	7,535	8,024	8,357	8,504
	GDP per capita*	11,756	14,741	18,196	22,090
	Abatement cost per capita	38	91	209	460

Example calculation: To calculate the normalized per capita cost savings in the UN-medium compared to the UN-low scenario in 2050, as described in Section 2 above, the calculation is as follows:

$$\left(\frac{481}{21691}-\frac{471}{22045}\right)*\ 22045=\sim\$18$$

The total global savings is that number (~18) multiplied by the population in the UN-low scenario (~8710 million).

* All GDP estimates are from DICE, including those using SSP populations.

Table S3. Abatement cost savings of achieving the UN-low versus UN-medium population under the total utilitarian optimal pathway that meets different temperature targets for: the default parameterization with normalization (top section), without normalization (middle panel), and normalized with near-zero discounting (0.1% per annum) (bottom section). Values are rounded.

		2020	2025	2030	2035	2040	2045	2050
Default parameterization with normalization								
2°C	**Per capita ($)**	4	5	7	10	13	16	18
	Total (billion $)	30	43	61	84	111	137	157
3°C	**Per capita ($)**	3	4	6	8	12	16	21
	Total (billion $)	22	33	49	70	99	135	180
Default parameterization without normalization								
2°C	**Per capita ($)**	4	5	7	8	10	11	10
	Total (billion $)	32	52	85	140	227	363	577
3°C	**Per capita ($)**	3	4	6	8	11	15	19
	Total (billion $)	23	36	55	83	123	178	252
Near-zero time preference with normalization								
2°C	**Per capita ($)**	8	10	12	14	16	16	16
	Total (billion $)	61	78	96	116	132	141	137
3°C	**Per capita ($)**	29	40	55	75	101	134	176
	Total (billion $)	220	318	451	629	861	1159	1530

Notes on Table S3

The top section of the table reports the same numbers presented in the TU lines of Figure 2 of the main text. The cost savings can be viewed as a type of budget that is available to decision makers to implement policies capable of moving the world's population from a trajectory like the UN-medium variant to one like the UN-low; staying within budget would pay for itself entirely through avoided climate abatement costs, even without accounting for the other primary benefits generally associated with a faster peak in global population (34,46). With this budget in mind, it is worth comparing the avoided costs in the above table with spending shortfalls for two human development interventions associated with fertility decline: ensuring access to family planning (modern contraception) and quality education. Recent studies estimate the shortfalls for these programs, respectively, at about $5 billion and $40 billion per year (44,45). The results therefore suggest that the avoided mitigation costs could cover the whole funding shortfall for family planning by 2020, and both programs soon thereafter.

However, it is important here to acknowledge that the UN-low population path is defined by the mechanistic assumption that fertility in all regions would be 0.5 children lower, which is a crude assumption. It would be difficult to achieve this level of decline in many wealthy regions and almost certainly not through additional investments in education or family planning, as these are already supplied to a large extent.However, studies indicate that roughly this level of *global* fertility decline – which is what is important in DICE – may be achievable given large and sustained policy intervention (2,47,48), although there is debate about the full effects of such policies. Additionally, Table S1 shows that most of the change in the carbon prices from the UN-medium to UN-low result from changes in the developing regions.

Two additional notes to this analysis should be mentioned. First, a successful human development policy could also increase average per capita emissions pressure – even while reducing population size – by increasing the economic productivity of a more educated population with more female labor force participation. A second is that some historical population policies have harmed wellbeing, sometimes without actual substantially changing population growth (41), which is a note of caution.

An additional reason for presenting this table is to demonstrate that our normalization procedure (see Section 2 above) makes little difference overall, but is slightly more conservative than directly taking the difference in mitigation costs between the UN-medium and UN-low scenarios.

(We do not report equivalent results with an average utilitarian objective because near-term cost differences are negligible as shown in Figure 2 of the main text.)

Section 6. Extensions

In this section we extend the results analyses in the main text through a series of additional experiments described below.

Extension 1. Mixed-population paths: Isolating the role of population in consumption and in social welfare weighting of time periods

In this section we explain the mechanisms through which population growth affects mitigation policy by separating the two primary factors linking population growth to carbon prices. First, if there is a TU social objective and future population is greater, then greater weight is given to future damages, as there are more future people to suffer them. In this sense population growth acts like lowering the discount rate, increasing the *current* incentive to reduce emissions by increasing the importance of the future. Second, a larger future population puts upward pressure on future emissions and corresponding damages, making emissions reductions more costly to achieve. We refer to the first factor as the “Weighting effect” (i.e. population’s weighting of the future in the SWF), and the second factor as the “Consumption effect” (i.e. more emissions and climate damages in the absence of mitigation).

Understanding the full consequences of population for climate policy requires understanding the separate roles of these two separate mechanisms, which we decompose in Figure S8. In particular, we compute two mixed-assumption scenarios:

- **High population consumption, low population weighting in the SWF:** here economic activity (and thus unmitigated emissions) is computed as though the population were on the high trajectory but wellbeing is calculated and optimized as though population were on the low trajectory;
- **Low population consumption, high population weighting in the SWF:** here we do the reverse, using the low population trajectory for computing economic activity (and thus unmitigated emissions) and the high trajectory in the weighting and optimization of the SWF.

This experiment with artificial scenarios is meant to clarify the role of the two population effects on the recommended climate policy. It is not very different from investigating non-artificial scenarios: what would be the optimal policy if we believe that a UN-high quantity of future people will be exposed to climate damages, but if we also believe that those people will be more frugal and only cause the emissions profile entailed by a UN-low quantity of future people? Or conversely, what is the optimal policy if we believe that a UN-low quantity of future people will be exposed to climate damages, but they are profligate and cause the emission profile entailed by a UN-high quantity of people?

Perhaps more importantly for the main result of our paper, this mixed-scenario experiment clarifies the mechanisms by which the future population path has its effect under AU and TU. This is because, even while using a TU SWF, the "low population in SWF, high population consumption" path behaves similarly to AU using the high population path. Under AU, the impact of the high population path on the economy is considered, but it has no effect on social welfare weighting; in the mixed path, future population growth has *some* effect on future population weighting because population is not constant under the low path, but it is substantially reduced. Therefore, the eventually increasing mitigation costs due to eventually emissions pressure of greater population under this mixed path is conceptually similar to the initially low but eventually increasing mitigation costs under AU with a high future population.

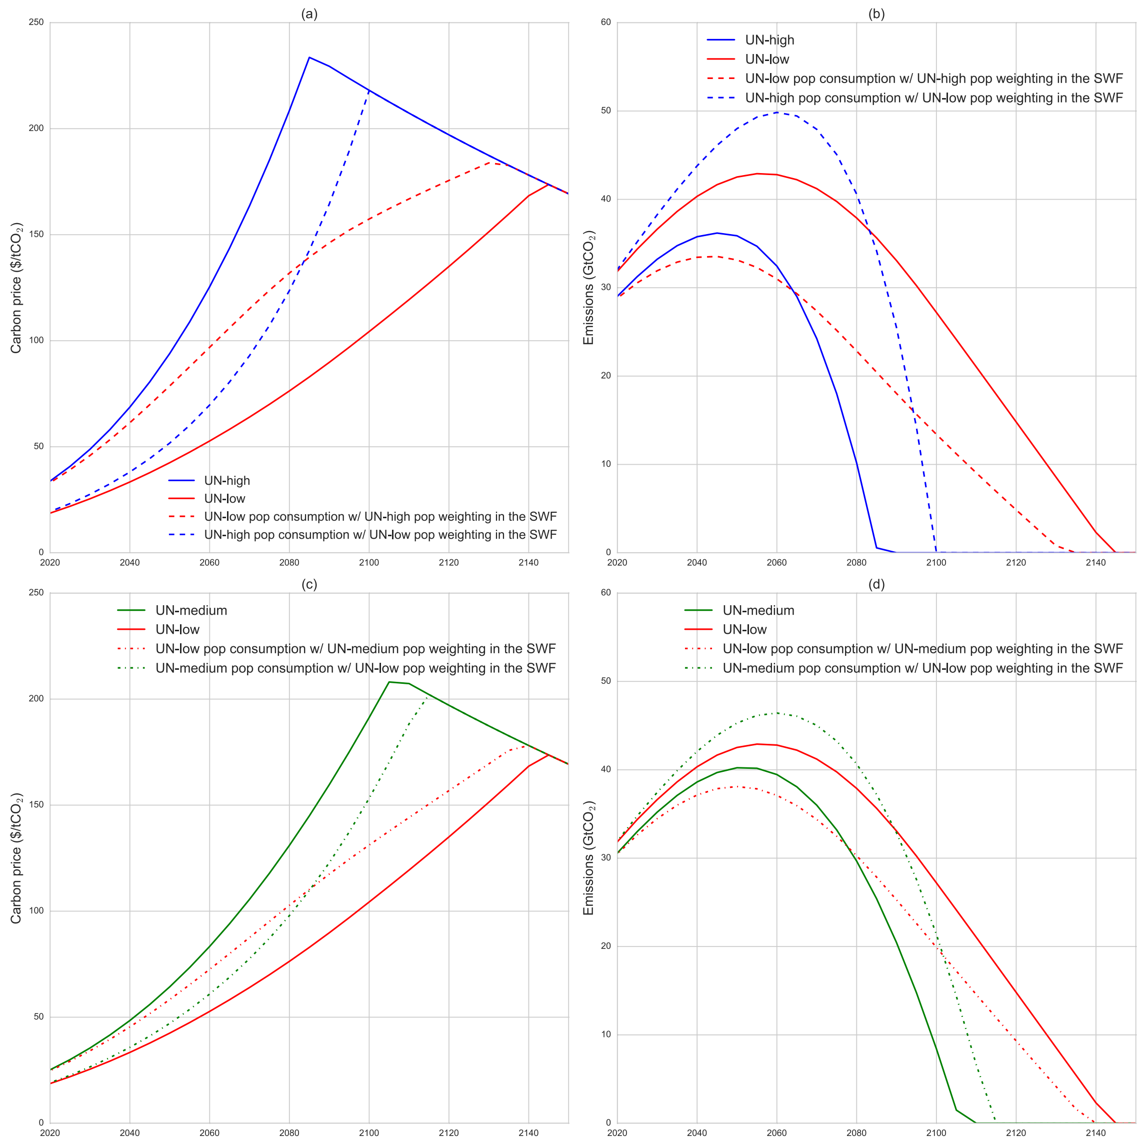

Fig. S8. Optimal carbon price and emissions pathways for standard population scenarios and specified mixed assumption scenarios. The mixed assumption experiment is designed separate the role of population growth in the weighting of social welfare from its role in contributing to unmitigated emissions and (potential) climate damages via consumption. This figure is only informative given a TU objective because with an AU objective there would be no social welfare weighting effect.

The upper left of panel (panel “a”) of Fig. S8 presents the results of this experiment. In the figure, four lines are plotted: the recommended tax paths of these two hypothetical, mixed-assumption scenarios are plotted as dashed lines alongside the original TU paths under the non-mixed UN-high and -low population scenarios, repeated from Fig. 1b of the main text as solid lines. The comparative importance of the two mechanisms changes over time. In the short run, the weighting effect dominates: both mixed-assumption optima closely match the optima of whichever population assumption is used in the SWF. For example, in early decades, the mixed scenario with high population in the SWF resembles the simple high-population scenario. This short-run result is because, during the nearest decades, the population implications with regard to emissions remain in the future, but there is still a mechanism for future population growth to influence near-term taxes through how the future is valued. Panel (b) of Figure S8 reports the emissions associated with the carbon price path.

In the long run of about a century, the economic consumption effect comes to dominate, a reversal visible in the crossing of the dashed lines. The long-run reversal occurs when faster population growth eventually translates into (increasingly) larger populations with (increasingly) greater total consumption and unmitigated emissions. It is at this point that the population paths used in the economy become more important. Therefore, the two mixed-assumption optimal tax paths cross in the late 21st century. For further analyses, compare the upper and lower left panels (panels “a” and “c”) of Fig. S8, which show how the result is qualitatively the same when comparing the UN-medium vs UN-low scenarios (or any higher vs lower scenario under this parameterization.

Extension 2. Exploring mitigation costs for temperature targets that are sub-optimally high (4 °C and 5 °C)

In Table S4 we report mitigation cost savings under temperature targets that entail less than optimal mitigation, namely 4 °C and 5°C. In these cases, we maximize the objective function subject to the relevant temperature constraint assuming constant growth of carbon prices from the starting values in the 3°C cases (by which we mean, starting in the UN-low variant from the initial carbon price in the UN-low 3°C case, and starting in the UN-medium variant from the initial carbon price in the UN-medium 3°C case). The cost savings are calculated as described in Section 1 above.

Table S4. Abatement cost savings under sub-optimally high temperature targets of achieving the UN-low versus UN-medium population under the total utilitarian optimal pathway that meets different temperature targets for: the default parameterization with normalization (top section), without normalization (middle panel), and normalized with near-zero discounting (bottom section). Values are rounded.

		2020	2025	2030	2035	2040	2045	2050
Default parameterization with normalization								
4°C	**Per capita ($)**	3	5	8	10	14	18	22
	Total (billion $)	26	42	62	87	118	154	195
5°C	**Per capita ($)**	3	5	7	9	11	14	17
	Total (billion $)	25	38	54	73	96	122	152
Default parameterization without normalization								
4°C	**Per capita ($)**	3	5	7	10	13	17	21
	Total (billion $)	27	44	68	98	137	186	245
5°C	**Per capita ($)**	3	5	6	8	11	14	17
	Total (billion $)	25	39	58	80	108	142	182
Near-zero time preference with normalization								
4°C	**Per capita ($)**	23	26	30	34	38	42	47
	Total (billion $)	179	210	245	282	323	365	410
5°C	**Per capita ($)**	23	26	29	32	35	39	43
	Total (billion $)	176	205	235	268	303	339	376

Extension 3. Assume equal mitigation effort, then compute the resulting difference in emissions and the savings from avoided climate damages

This subsection describes an alternative method for conceptualizing the climate benefits of a lower population in connection with suboptimal mitigation effort and a TU social objective.

The key assumption behind the mitigation cost savings analysis described in the earlier sections and displayed in Figure 2 of the main text and Tables S2-S4 is that the world will be willing to spend more on mitigation in the higher population (UN-medium) scenario because there is a temperature target to meet. An alternative assumption is that society may have a fixed appetite for mitigation, regardless of population. In Table S5, we show the implications of this alternative assumption for different levels of mitigation effort (different "appetites"), ranging from higher than optimal effort, to much less than optimal effort, to no effort at all; using DICE we quantify the resulting difference in emissions between the UN-medium and UN-low scenarios for each level of effort, and quantify the resulting difference in climate damages. In other words, in this section the benefit of a lower population appears (by the design of the experiment) in the form of avoided *climate damages* with a lower versus higher population, rather than in the form of avoided *mitigation cost*.

To model this alternative assumption that society might have a fixed appetite for mitigation, we define a level of mitigation effort *at a specified time* as the mitigation cost as a percentage of per capita GDP, and a level of mitigation effort *pathway* as a particular level of mitigation effort at each time point into the future. With these definitions in hand, we consider the four mitigation effort pathways that yield the four UN-medium temperature target scenarios reported in Tables S3 and S4 – i.e., we consider the mitigation effort pathway that yields peak warming of 2°C given UN-medium population, the mitigation effort pathway that yields peak warming of 3°C given UN-medium population, and so on. We also consider the mitigation effort pathway that yields the UN-medium Baseline scenario, where there is no climate mitigation.

Therefore, for each of the five UN-medium scenarios, we determine the control rate at each time period that *with the UN-low population* would yield the same mitigation cost as a percentage of per capita GDP at each time and then run the model accordingly.

Table S5 shows the difference in emissions between these five pairs of scenarios; again, each pair has the same mitigation effort pathway in the sense just described. Because there is by design almost no difference in *mitigation cost* as a percentage of per capita GDP, the savings implied by a lower population are here represented only by avoided *climate damages* in the further future. Such avoided climate damages are due to lower emissions, and the difference increases in magnitude as mitigation effort becomes less and less stringent.

Table S5. Avoided damages under UN-low rather than UN-medium population with a TU social objective given equal mitigation effort (for different mitigation effort pathways that yield different peak temperatures given UN-medium population).

		Total change in emissions	2020	2030	2040	2050	2150	2200
2°C	Per capita ($)	-3%	0	0	0	2	17	23
	Total ($billion)		0	0	4	14	98	119
	Change in emissions		0.2	0.7	1.2	0.9	-	-
3°C	Per capita ($)	-11%	0	0	1	3	216	274
	Total ($billion)		0	1	5	23	1,233	1,452
	Change in emissions		0.2	1.0	2.2	3.5	-	-
4°C	Per capita ($)	-22%	0	0	1	3	870	1,436
	Total ($billion)		0	1	6	24	4,973	7,608
	Change in emissions		0.2	1.0	2.3	3.9	5.6	0.0
5°C	Per capita ($)	-31%	0	0	1	3	1,314	3,342
	Total ($billion)		0	1	6	25	7,509	17,702
	Change in emissions		0.2	1.1	2.4	4.4	22.5	7.2
Base	Per capita ($)	-37%	0	0	1	4	2,275	6,945
	Total ($billion)		0	1	7	34	13,000	36,786
	Change in emissions		0.2	1.4	3.3	6.1	53.1	38.1

Extension 4. Critical levels in the social welfare function

Table 1 in the main text considers whether different population paths would represent overall improvements in wellbeing, as valued by the social welfare functions we use. As described above and in the main text, our implementations of AU and TU are two focal functional forms, but other functional forms are also consistent with the basic principles of valuing, not valuing, or conditionally valuing additional lives or a larger population. One alternative functional form within the general family of TU incorporates a "critical level": a level of personal utility (in our case consumption, c, is the only input into utility, $u(c)$) above which an additional life increases social welfare and below which it does not. This form of totalist utilitarianism replaces $[u(c) * n]$ with $[\big(u(c) - u(a)\big) * n]$ within each period, where a is the critical level and n is the population size (49).

All totalist SWFs implicitly assume a critical level – if nothing is subtracted, a critical level of zero is assumed. DICE's default functional form implicitly assumes a critical level equal to a GDP per capita of $3,775 per year. Therefore, if there were lives below this level of consumption, they would reduce social welfare; however, per capita GDP is greater than this level in every period in the model.

Tables S6 and S7 analyze alternative critical levels in percentage and absolute terms, respectively. As the critical level is reduced, the pattern is qualitatively the same as Table 1 of the main text: lowering population reduces total wellbeing. The effect is quantitatively dampened in percentage terms (Table S6) because the value of social welfare goes up when the critical level is reduced, but the absolute difference increases (Table S7).

As a verification of our mechanism, we also include a case where the critical level is very high ($10,000/year) – in our view, inappropriately high (because it implies that it makes the world worse for people to exist who have an income level comparable to the average citizen of Morocco or Ukraine, for example). Here, a larger population is still preferable to a smaller one; even though additional lives are regarded as a social worsening over the coming several decades when global per capita GDP has not yet reached $10,000 per year, the cumulative total is higher over the model's full time horizon. A critical level of $43,915 (roughly the current per capita GDP in the UK) is the point in DICE where the model has an equal amount of wellbeing whether assuming the UN-high or UN-low population.

As a final note, optimal carbon taxes, emissions, and temperature change are all entirely independent of the choice of critical level. This is a theoretical implication of the fact that that when population is exogenous (as in our model), the additive critical level term drops out of the first order condition for the optimum. In other words, assuming a particular future population path, the optimal emissions reduction pathway and associated outcome for world is the same regardless of what assumption is made about the critical level – the choice of the critical level only affects how society should value that future world.

Table S6: Social valuation of population paths under several different TU critical levels at optimal climate policies, as a percent of the UN-medium path.

Critical level ($/year)	UN-high	UN-medium	UN-low
$0.00	133.6%	100%	77.7%
$730.50	144.0%	100%	71.1%
$3,775.50 (DICE default value)	192.3%	100%	40.4%
$10,000.00*	--	--	--
* Total wellbeing is negative given a $10,000 critical level (Table S7), so we don't report percentages for this scenario.			

Table S7: Social valuation of population paths under several different TU critical levels at optimal climate policies, in units of wellbeing (utility).

Critical level ($/year)	UN-high	UN-medium	UN-low
$0.00	713,892	534,410	415,145
$730.50	25,763	17,894	12,719
$3,775.50 (DICE default value)	6,262	3,257	1,315
$10,000.00	-63	-1,491	-2,384

References

1. Gaffin SR, O'Neill BC (1997) Population and global warming with and without CO_2 targets. *Popul Environ* 18(4):389–413.

2. O'Neill BC, et al. (2010) Global demographic trends and future carbon emissions. *Proc. Natl. Acad. Sci. U.S.A.* 107(41):17521–17526.

3. O'Neill BC, MacKellar FL, Lutz W (2005) *Population and climate change*. (Cambridge).

4. Lam D (2011) How the world survived the population bomb: Lessons from 50 years of extraordinary demographic history. *Demography* 48(4):1231–1262.

5. Casey G, Galor O (2017) Is faster economic growth compatible with reductions in carbon emissions? the role of diminished population growth. *Environ Res Lett* 12(1):014003.

6. Liddle B (2011) Consumption-driven environmental impact and age structure change in OECD countries: A cointegration-STIRPAT analysis. *Demogr Res* 24:749–770.

7. O'Neill BC, et al. (2012) Demographic change and carbon dioxide emissions. *Lancet* 380(9837):157–164.

8. Spears D (2015) Smaller human population in 2100 could importantly reduce the risk of climate catastrophe. *Proc. Natl. Acad. Sci. U.S.A.* 112(18): E2270.

9. O'Neill BC (2000) Cairo and climate change: a win–win opportunity. *Glob Environ Change* 10(2):93–96.

10. O'Neill BC, Wexler L (2000) The greenhouse externality to childbearing: a sensitivity analysis. *Clim Change* 47(3):283–324.

11. Wheeler D, Hammer D (2010) The economics of population policy for carbon emissions reduction in developing countries, (Center for Global Development), Report.

12. Broome J (2012) *Climate Matters: Ethics in a Warming World*. (Norton).

13. Interagency Working Group on the Social Cost of Carbon (2013) *Technical Update of the Social Cost of Carbon for Regulatory Impact Analysis Under Executive Order 12866*. (United States Government).

14. Weyant J (2017) Some contributions of integrated assessment models of global climate change. *Review of Environmental Economics and Policy* 11(1), 115-137.

15. Nordhaus WD, Sztorc P (2013) *DICE 2013R: Introduction and user's manual*. (Yale).

16. Stern N (2006) *Stern Review: The economics of climate change*. (Her Majesty's Treasury, London) Vol. 30.

17. Hope C (2011) The social cost of CO_2 from the PAGE09 model, (Cambridge), Discussion Paper 2011-39.

18. Anthoff D (2009) *Equity and climate change: Applications of FUND.* (Max Planck Institute).

19. Dasgupta P (2001) *Human well-being and the natural environment*. (Oxford).

20. Asheim GB (2004) Green national accounting with a changing population. *Economic Theory*, 23(3): 601-619.

21. Arrow KJ, Dasgupta P, & Mäler KG (2003) The genuine savings criterion and the value of population. *Economic Theory* 21(2): 217-225.

22. Spears D (2017) Making people happy or making happy people? Questionnaire-experimental studies of population ethics and policy. *Social Choice and Welfare* 49(1): 145-169.

23. Arrhenius G (forthcoming) *Future generations: A challenge for moral theory*. (Oxford University Press).

24. Nordhaus WD, Boyer J (1998) What are the external costs of more rapid population growth?, Theoretical issues and empirical estimates *in 150th Anniversary Meeting of the American Association for the Advancement of Science, Philadelphia, February.* Vol. 15.

25. Gillingham K, et al. (2015) Modeling uncertainty in climate change: A multi-model comparison, (Cowles), Discussion Paper 2022.

26. Nordhaus WD (2007) A review of the 'Stern review on the economics of climate change'. *J Econ Lit* 45(3):686–702.

27. Dasgupta P (2007) The Stern Review's economics of climate change. *Natl Inst Econ Rev* 199:4–7.

28. Weitzman ML (2007) A review of the Stern Review on the economics of climate change. *J Econ Lit* 45(3):703–724.

29. Adler M, Anthoff D, Bosetti V, Garner G, Keller K, & Treich N (2017) Priority for the worse-off and the social cost of carbon. *Nature Climate Change* (7): 443–449.

30. Population Division (2015) *World Population Prospects: 2015 Revision*. (United Nations Department of Economic and Social Affairs).

31. Basten S, Lutz W, Scherbov S (2013) Very long range global population scenarios to 2300 and the implications of sustained low fertility. *Demogr Res* 28:1145–1166.

32. Nordhaus WD, Boyer J (2000) *Warming the world: economic models of global warming.* (Cambridge: MIT Press).

33. Dennig F, Budolfson MB, Fleurbaey M, Siebert A, Socolow RH (2015) Inequality, climate impacts on the future poor, and carbon prices. *Proc. Natl. Acad. Sci. U.S.A*. 112(52):15827–15832.

34. Lutz W, Muttarak R, Striessnig E (2014) Universal education is key to enhanced climate adaptation. *Science* 346(6213):1061–1062.

35. Lutz W, Striessnig E (2015) Demographic aspects of climate change mitigation and adaptation. *Popul Stud* 69(S1):S69–S76.

36. Tol RS (2009) The economic effects of climate change. *J Econ Perspectives* 23(2):29–51.

37. Lutz W, Butz WP, Samir K (2014) *World population and human capital in the twenty-first century*. (Oxford University Press, Oxford). http://www.wittgensteincentre.org/dataexplorer.

38. Gerland P, et al. (2014) World population stabilization unlikely this century. *Science* 346(6206):234–237.

39. Golosov M, Hassler J, Krusell P, Tsyvinski A (2014) Optimal taxes on fossil fuel in general equilibrium. *Econometrica* 82(1):41–88.

40. Arrow K, et al. (2004) Are we consuming too much? *J Econ Perspectives* 8(3): 147-172.

41. Connelly MJ (2008) *Fatal misconception*. (Harvard, Cambridge).

42. Nordhaus WD (2008) *A question of balance: Weighing the options on global warming policies*. (Yale University Press).

43. Abel GJ, Barakat B, Samir K, Lutz W (2016) Meeting the sustainable development goals leads to lower world population growth. *Proc. Natl. Acad. Sci. U.S.A*. 113(50):14294–14299.

44. Singh S, Darroch JE, Ashford LS (2014) Adding it up: the costs and benefits of investing in sexual and reproductive health, (Guttmacher Institute), Report.

45. Education for All (2015) *Global Monitoring Report: Pricing the right to education*. (UNESCO).

46. Cleland J, Conde-Agudelo A, Peterson H, Ross J, Tsui A (2012) Contraception and health. *Lancet* 380(9837):149–156.

47. Bradley S, Croft TN, Fishel JD, Westoff CF (2012) Revising unmet need for family planning, (USAID), Report.

48. Bongaarts, J, Cleland J, Townsend JW, Bertrand JT, Gupta MD (2012) Family planning programs for the 21st century: rationale and design, (Population Council), Report.

49. Broome, J (2004) *Weighing lives*. (Oxford University Press).