

ARTICLE

<https://doi.org/10.1038/s41467-019-09499-x>

OPEN

The impact of human health co-benefits on evaluations of global climate policy

Noah Scovronick^{1,2,17}, Mark Budolfson^{3,4,17} , Francis Dennig^{5,6,17}, Frank Errickson^{7,17}, Marc Fleurbaey^{2,8}, Wei Peng⁹, Robert H. Socolow¹⁰, Dean Spears^{11,12,13,14} & Fabian Wagner^{2,15,16,17}

The health co-benefits of CO₂ mitigation can provide a strong incentive for climate policy through reductions in air pollutant emissions that occur when targeting shared sources. However, reducing air pollutant emissions may also have an important co-harm, as the aerosols they form produce net cooling overall. Nevertheless, aerosol impacts have not been fully incorporated into cost-benefit modeling that estimates how much the world should optimally mitigate. Here we find that when both co-benefits and co-harms are taken fully into account, optimal climate policy results in immediate net benefits globally, overturning previous findings from cost-benefit models that omit these effects. The global health benefits from climate policy could reach trillions of dollars annually, but will importantly depend on the air quality policies that nations adopt independently of climate change. Depending on how society values better health, economically optimal levels of mitigation may be consistent with a target of 2 °C or lower.

¹Department of Environmental Health, Rollins School of Public Health, Emory University, 1518 Clifton Rd. NE, Atlanta, GA 30322, USA. ²Woodrow Wilson School of Public and International Affairs, Princeton University, Robertson Hall, Princeton, NJ 08544, USA. ³Gund Institute for Environment and Department of Philosophy, University of Vermont, Burlington, VT 05310, USA. ⁴Edmond J. Safra Center for Ethics, Harvard University, 124 Mount Auburn Street, Cambridge, MA 02138, USA. ⁵Social Sciences (Economics), Yale-NUS College, Singapore 138610, Singapore. ⁶NUS fellow at the Center for Advanced Study in the Behavioral Sciences, Stanford, CA 94305, USA. ⁷Energy and Resources Group, University of California Berkeley, 310 Barrows Hall, Berkeley, CA 94720, USA. ⁸University Center for Human Values, Princeton University, 304 Louis Marx Hall, Princeton, NJ 08544, USA. ⁹School of International Affairs and Department of Civil and Environmental Engineering, Pennsylvania State University, University Park, PA 16801, USA. ¹⁰Department of Mechanical and Aerospace Engineering, Princeton University, Olden Street, Princeton, NJ 08544, USA. ¹¹Department of Economics, University of Texas at Austin, 2225 Speedway, Austin, TX 78712, USA. ¹²Economics and Planning Unit, Indian Statistical Institute – Delhi Centre, 7, S.J.S Sansawal Marg, New Delhi 110016, India. ¹³IZA Institute of Labor Economics, Schaumburg-Lippe-Strasse 5-9, 53113 Bonn, Germany. ¹⁴Institute for Futures Studies, Hölländargatan 13, Stockholm, Sweden. ¹⁵International Institute for Applied Systems Analysis, Laxenburg 2361, Austria. ¹⁶Andlinger Center for Energy and the Environment, Princeton University, 86 Olden Street, Princeton, NJ 08544, USA. ¹⁷These authors contributed equally: Noah Scovronick, Mark Budolfson, Francis Dennig, Frank Errickson, Fabian Wagner. Correspondence and requests for materials should be addressed to N.S. (email: Scovronick@emory.edu)

Climate policies targeting CO₂ may also reduce air pollutant emissions—and the aerosols they produce—as the two share emission sources. Prior studies on the topic have quantified the associated health co-benefits of pre-defined greenhouse gas reduction scenarios^{1–6}, or estimated the economic impacts from reducing specific pollutants^{7–9}, but these types of impacts have not been fully incorporated into cost-benefit modeling that estimates how much the world should optimally mitigate. We move this literature forward by developing a comprehensive cost-benefit integrated assessment model based on William Nordhaus' Regionalized Integrated Climate Economy (RICE) model, where the new developments allow the model to weigh both the health co-benefits and the climate co-harms of aerosol co-reductions (co-harms exist because aerosols produce net cooling overall¹⁰); the latter in particular has been a largely neglected aspect of the co-benefits discussion. (We use the term co-harm to refer to the net climate harm of aerosol reductions, recognizing that reducing some species of emissions individually may produce different effects than the net of all species together; for example, although reducing black carbon produces a climate benefit, that effect may be outweighed by the climate harm of reductions in other species.)

These modeling developments, which account for the key air pollutant emissions and their individual properties, provide new capability to investigate fundamental policy questions that have not been answered by existing studies¹¹. This includes determining: (1) the optimal climate policy across time and how it is affected by independent air quality control, (2) whether climate policy produces immediate net benefits, or if there are inter-generational tradeoffs, and (3) if specified climate targets are justifiable on cost-benefit grounds.

To answer these questions, we modify the RICE optimization model¹² to include an empirically calibrated, regionally differentiated feedback mechanism whereby reducing CO₂ also reduces regional air pollutant emissions from co-emitting sources. We then quantify and monetize the impact on both health and radiative forcing throughout the world and compute the resulting optimal climate policy. We thus modify the standard tradeoff between CO₂ mitigation costs and climate damages with a more complete analysis that simultaneously weighs mitigation costs, climate damages from CO₂, and the health and climate consequences of changes in air pollutant co-emissions. The resulting model estimates optimal climate policy after jointly considering all these factors. As a robustness check, we also modify the widely used FUND (Climate Framework for Uncertainty, Negotiation and Distribution) model¹³ to include the same mechanisms.

We find that when both co-benefits and co-harms are taken fully into account, optimal climate policy results in immediate net benefits globally, which overturns previous findings from cost-benefit models that omit these effects. The global health benefits from climate policy could reach trillions of dollars annually, but their magnitude will importantly depend on the air quality policies that nations adopt independently of climate change. Depending on how society values better health, we show that economically optimal levels of mitigation may be consistent with a target of 2 °C or lower.

Results

Main results. We summarize the results of the new model (hereafter referred to as RICE + AIR, for RICE + Aerosol Impacts and Responses) in terms of the optimal fraction of business-as-usual CO₂ emissions that should be reduced over time. We refer to this quantitative reduction in CO₂ as the optimal decarbonization fraction and express it as a percentage. (What we call the decarbonization fraction is often referred to as the control rate.) The blue line in Fig. 1a shows the optimal decarbonization pathway if

the cost-benefit analysis only considers the climate impacts of CO₂ and associated aerosol co-reductions. Health impacts are not included. In this reference case, the optimal decarbonization fraction is 24% of business-as-usual emissions in 2030, rising to 35% in 2050 and ultimately reaching full decarbonization by 2130. It is similar to the optimal trajectory of the standard RICE model, which features exogenous aerosol forcing and excludes health co-benefits (Supplementary Fig. 1).

In contrast to the blue line, which considers only the climate impacts of CO₂ and aerosols, the red line also includes the health co-benefits of the aerosol reductions that result from decarbonization. This leads to increased optimal CO₂ mitigation. The difference between the red line and the blue line demonstrates the importance of the health benefits; roughly 45–60% more decarbonization is optimal over the next five decades (and 10–40% thereafter) compared to the reference case that only considers climate consequences. Full mitigation also occurs earlier in time. The additional emission reductions that result from the inclusion of the health co-benefits cumulatively amounts to ~270 GtC (Fig. 1a). Importantly, all of these results account for the damages from lost cooling attributable to the aerosol co-reductions.

The additional decarbonization justified by the health gains leads to a peak temperature 0.4 °C lower than the reference case (Fig. 1b). The carbon price pathways associated with Fig. 1 can be found in Supplementary Fig. 2.

Figure 2a shows that the optimal climate policy has immediate and continual monetized global net benefits when accounting for health co-benefits. This overturns the findings from standard cost-benefit optimization models, which ignore health co-benefits and thus imply that optimal climate policy has net costs for much of this century (Fig. 2b). The result is consistent with prior co-benefits studies that have analyzed specific emission reduction scenarios and reported high benefits relative to mitigation costs^{1,14}, but which generally do not also account for (monetized) climate-related impacts.

The distribution of health co-benefits by region for the full RICE + AIR optimum is displayed in Fig. 3a. In line with recent scenario-based studies^{1,2,15}, many of the co-benefits accrue in India and China in early periods, attributable to their large populations and high capacity for mitigation-related reductions in PM_{2.5}. China's benefits decline by mid-century due to relatively rapid economic development and a stabilizing population—which both act to constrain emissions—whereas those in India persist and are the major driver of increased decarbonization relative to the reference case (see below for a sensitivity test which corroborates India's importance). Towards the end of the century, sub-Saharan Africa replaces China as the second-largest beneficiary, as air pollution remains problematic due to lagging economic development accompanied by the world's largest population. Other regions also stand to benefit, including less populous regions that show important benefits per capita and/or per Gross Domestic Product (GDP) (Fig. 3b).

In the results presented above, the averted premature mortality from aerosol co-reductions produces annual monetized benefits in the hundreds of billions of dollars over the next few decades, rising to several trillion annually at the end of the century (Fig. 3c). To derive these numbers, we multiply, for each region-time pair, the total life-years gained by 2 years of per capita consumption. This approach produces different life-year monetizations for each region, which leads to the slight change in composition of monetized benefits (compare Fig. 3a, c). However, this does not imply that we assign life-years in poorer regions less value in the objective function, because the optimization accounts for the diminishing marginal utility of consumption through a concave relationship between wealth and well-being, as described below and in the Methods section (Eq. (1)). Below we also show the sensitivity of our results to alternative valuations of health benefits.




Fig. 1 Optimal decarbonization and temperature. **a** Decarbonization over time for the reference case optimal policy that excludes health co-benefits (blue line) and in the full RICE + AIR optimal policy that includes health co-benefits (red line). **b** Estimated global temperature rise above preindustrial levels that would occur given the decarbonization in **a**. Decarbonization is relative to a business-as-usual scenario without any climate action, and 100% decarbonization signifies zero net carbon emissions




Fig. 2 Costs and benefits of mitigation. Decomposition of the change in global consumption relative to the business-as-usual (BAU) scenario under **a** the full RICE + AIR optimal policy, and **b** the reference case optimal policy. Health co-benefits and benefits from avoided CO₂ damages are positive, while mitigation costs and aerosol co-harms (climate damage from the co-reduction of cooling aerosols) are negative. The black solid line displays the global net effect. **a** shows that the net effect on global consumption is immediately positive when health co-benefits are taken into account, in contrast to the reference case (**b**), which is representative of standard cost-benefit models that do not include health co-benefits and thus imply that optimal climate policy has net costs for much of this century. If health co-benefits are added to the reference policy in **b**—by adding the light-red bars displayed in **a**—the global net effect becomes immediately positive, and if health co-benefits were removed from **a**, the net effect would be negative for most of this century




Fig. 3 Health benefits of carbon mitigation. Life-years gained **a** overall and **b** per 100,000 population by region from the air quality improvements associated with the optimal decarbonization in RICE + AIR. **c, d** show the resulting monetized benefits in total, and as a percent of GDP, respectively. Note that if a region's $\text{PM}_{2.5}$ exposure (concentration) drops below $5.8 \mu\text{g}/\text{m}^3$, health benefits no longer accrue—this threshold assumption is common in other global air quality assessments¹⁶, and is tested in the sensitivity analyses

Independent air quality control. All results presented above assume the level of air quality control that occurs independently of climate policy will proceed approximately as projected in the coming decades, based on current and planned policies. We systematically test the importance of this assumption by implementing an environmental Kuznets-type approach where emission intensities decrease with increasing per capita income (described in detail in the Methods section, and in particular Eqs. (2)–(4) and associated discussion). Figure 4 reports results when we slow down and speed up this Kuznets-relevant income, allowing it to range between roughly 50% ($\chi = 0.5$) and 150% ($\chi = 1.5$) of the true (modeled) income; lower values imply less stringent air quality control and vice versa. In these model runs the true income is used in all other parts of the model.

The results demonstrate that although assumptions about air quality control do influence optimal mitigation levels, more decarbonization remains optimal compared to the reference case under all of these scenarios. However, the mechanisms driving the additional decarbonization vary. If the (pre-mitigation) air is dirtier

than projected ($\chi < 1$), extra mitigation occurs primarily to reap the health co-benefits. If stringent air quality control occurs in the future ($\chi > 1$), more decarbonization is also optimal compared to the reference case. In part this is to capture remaining health co-benefits, but also for two other reasons. First, the potential damages associated with the substantial loss of the aerosol cooling effect requires additional CO_2 reductions as a counterbalance. And second, when the air is cleaner, each marginal reduction in CO_2 has a greater net benefit because it is not coupled to as much cooling aerosol.

Another way of understanding the role of independent air quality measures on optimal climate policy is by comparing not only the different RICE + AIR cases, as in Fig. 4, but also the difference between the RICE + AIR cases and their corresponding reference cases, as reported in Table 1. The results confirm that as autonomous emission controls get stronger (represented by higher χ values), the less health co-benefits matter for climate policy.

One of the scenarios displayed in Fig. 4 is the fully co-optimal case that selects the ideal combination of both air quality and climate policies by introducing an additional policy lever that can




Fig. 4 Impact of independent air quality control. **a** Optimal decarbonization and **b** associated global temperature rise given different levels of autonomous air quality control. The reference and RICE + AIR cases from Fig. 1, where $\chi = 1$, are displayed as the blue and red solid lines, respectively

act directly on individual air pollutant emissions through end-of-pipe measures, rather than via CO₂ (described in Supplementary Note 1). When this lever is added, large reductions in air pollutant emissions occur when and where associated abatement costs are low, the reductions lead to relatively large decreases in exposure, and/or they produce few climate damages. This first-best policy where climate and air pollution policies are co-optimized leads to rapid air quality improvements but still increased decarbonization relative to the reference case, albeit by a lower margin than in the standard $\chi = 1$ case that underlies our main results. Note that we selected $\chi = 1$ as the basis for the main result in order to explore the effect health co-benefits have when regions act consistently with their current and planned air pollution policies (which are suboptimal). As this section demonstrates, the magnitude of the co-benefit effect is importantly influenced by the assumed level of independent air quality control.

Discounting. The RICE model has a discounted utilitarian objective, meaning that for optimal policy calculations, the objective of the model is to maximize the sum of discounted well-being (see Eq. (1) of the Methods and associated discussion). The discount rate for consumption is determined via the Ramsey rule, which adds the rate of pure time preference to the product of the elasticity of marginal utility and the economic growth rate. The rate of pure time preference is the rate at which the weight given to future well-being declines over time. The elasticity of marginal utility represents the lesser importance of each additional dollar to well-being as one gets richer. Unless otherwise noted, we assume an elasticity of marginal utility of 1.5, the default in RICE.

In all the results presented so far we have used a rate of pure time preference of 1.5% per year, which is the default value RICE, but close to the upper end of the range highlighted by the IPCC (Intergovernmental Panel on Climate Change)^{16,17}. This choice puts more weight on near-term impacts compared to those occurring further in the future. As a result, aerosol impacts, which occur more immediately than the climate impacts from CO₂, have a relatively outsized importance. Nevertheless, when we implement a

much lower rate that corresponds to a near-zero (0.1%) preference for the present over the future, optimal decarbonization remains substantially higher in RICE + AIR compared to the analogous reference case (Fig. 5). (Near-zero time preference is often used in the climate economics literature¹⁷, including in the Stern Review¹⁸).

Conversely, if we increase time preference to 3.5%, which implies a discount rate for consumption of 7%, the difference between RICE and RICE + AIR becomes even more stark. (To calculate the 7% discount rate via the Ramsey rule, we used the economic growth rate in the initial time period. We note that the exact discount rate may change over time due to the effect of climate damages on economic growth; the impact however, is negligible.) We chose the 7% value because it is the highest of the primary discount rates generally used in US federal cost-benefit analyses. (Although most experts consider lower rates to be more appropriate in an intergenerational context such as in the case of climate change-related analyses^{16,17}, the current US administration has signaled a preference for including the higher 7% rate¹⁹.)

Such high time preference places so much emphasis on the near-term health benefits associated with reducing CO₂ that optimal decarbonization is substantially higher than both the analogous reference case and the reference case with more moderate (1.5%) time preference. These results indicate that even after strongly discounting the future, a robust climate policy is still warranted. In Supplementary Fig. 3 we present further analyses showing results with a 3% discount rate, which is the lower value generally used for US federal cost-benefit analyses.

Monetizing and valuing health benefits. The results above assume that one life-year gained equals 2 years of per capita consumption in dollar terms. Two years of per capita consumption corresponds to the approach to monetizing health impacts used in key studies that form the basis of the RICE climate damage function^{12,20}, and is similar to survey-based life-year assessments²¹. However, this approach yields much lower monetary benefits than using the value of a statistical life (VSL), which is also widely adopted in the literature. If we assume that

Table 1 Impact of independent air quality control

	2030 Reference	2030 RICE + AIR	2050 Reference	2050 RICE + AIR	Diff. in peak temp (°C) ^a
$\chi = 0.5$	20%	39%	27%	57%	0.9
$\chi = 0.75$	22%	37%	33%	54%	0.5
$\chi = 1$ (main result)	24%	36%	35%	52%	0.4
$\chi = 1.25$	24%	35%	36%	51%	0.4
$\chi = 1.5$	25%	35%	37%	50%	0.3

Optimal decarbonization rates in RICE + AIR versus the corresponding reference case for different values of χ , which represent different levels of independent air quality control. $\chi = 1$ represents the assumption underlying the main results, whereas higher (lower) values indicate more (less) stringent air quality control in the future. The difference in peak temperature is the difference between the highest global average temperature in the RICE + AIR optimum versus the reference case given the assumed value of χ

^aPositive values indicate a lower peak temperature in RICE + AIR. Values are rounded




Fig. 5 The influence of pure time preference. Optimal decarbonization with low (0.1%), medium (1.5%), and high (3.5%) rates of pure time preference for the reference case (that excludes health co-benefits) and in the full RICE + AIR case (that includes health co-benefits). All runs have a consumption elasticity of marginal utility of 1.5

each adult death attributable to $PM_{2.5}$ results in approximately 10 years of life lost²², we would use roughly 8–16 years of per capita consumption per life-year gained, instead of two; thus, in what follows we use 8–16 years of per capita consumption as one possible approximation of a VSL, in addition to a more direct VSL-based approach described below.

The sensitivity of our results to alternative life-year monetizations is reported in Table 2. The monetized global health benefits in our main results discussed above would be roughly 4–11 times higher if we used VSL-like monetizations in the optimization, and the associated optimal level of decarbonization would likely be consistent with keeping the maximum global temperature rise to 2 °C. Two degrees is a target specified in the Paris Agreement and widely considered as necessary to avoid dangerous climate change, but one that has generally not been warranted according to previous cost-benefit assessments (using similar discounting parameters) that omit health co-benefits¹². Combining the low (0.1%) rate of time preference with a VSL-based approach may justify a target as low as 1.8 °C.

Table 2 reports the sensitivity of the results to differences in how life-years are monetized in dollar terms. In RICE's cost-benefit approach, an important second step occurs when the monetized health benefits are valued in well-being terms via the objective function (see Eq. (1) in the Methods section), which gives the model

the aim of maximizing the (discounted) sum of global well-being through time, as is standard in optimal policy modeling. A key feature of the objective function is diminishing marginal utility of consumption, which captures the core concept that an additional dollar generates more well-being when given to a poorer person than to a richer person. The result is that while life-years are assigned a lower dollar amount in poorer countries in absolute terms, they are actually assigned greater value in well-being/utility terms. We discuss this issue in more detail in Supplementary Note 2 and show in Supplementary Table 1 that the strong effect of adding health co-benefits persists whether life years are valued more highly in poorer regions, less highly, or exactly the same as in wealthier regions; however, the magnitude of the effect changes.

Additional sensitivities in RICE + AIR. Table 3 reports results for several other sensitivities, presenting the percent increase in optimal decarbonization rates in the RICE + AIR case compared to the corresponding reference case. The table is organized as follows, with variables in parentheses referring to the relevant parameter in the model equations reported in the Methods.

First (Test 1), we explore different air pollutant co-reduction levels (κ), with the range representing the high and low values after applying alternative estimates from the other four Shared Socio-economic Pathways (SSPs). Note that the SSPs are not ordered in terms of their co-reduction potentials, but instead reflect different possible futures across multiple socioeconomic dimensions. In the second test (Test 2), we substitute the TM5-FAst Scenario Screening Tool (TM5-FASST) source-receptor matrix (SRM) for the SRM based on simulations of the European Monitoring and Evaluation Program atmospheric chemistry and transport model. We confined this sensitivity to Asia because it was the only region analyzed as part of a recent project at the International Institute of Applied Systems Analysis. However, the four Asian regions account for ~85% of all life-years gained globally over the next century in the main analyses (Fig. 3), and therefore largely drive the findings. In Test 3 and Test 4 we assume that the relative risk for all-cause mortality (β) was the lower and upper bound, respectively, of the confidence interval in Forestiere et al.²⁴ rather than the central estimate. Test 5 lowers the $PM_{2.5}$ threshold below which there are no adverse health effects (τ) to $1 \mu g/m^3$ instead of $5.8 \mu g/m^3$. In Test 6 we assume there is no benefit from reducing air pollution at levels above $50 \mu g/m^3$. We include this counterfactual boundary case for two reasons. The first is to investigate the maximum possible concavity in risk functions at high levels of exposure; in reality however, recent empirical studies^{25,26} indicate that the marginal effect of air pollution remains positive at levels well above $50 \mu g/m^3$. The second reason is to provide additional emphasis on the importance of India and China, as this test effectively eliminates the impact of both of those nations (as well as the Middle East/North Africa region) over the next several

Table 2 Monetizing life-years/lives

	2030	2050	Peak temperature (°C) ^a
Reference (no co-benefits)	24%	35%	3.0
1 year of per capita income	30%	44%	2.9
2 years of per capita income (main result)	36%	52%	2.6
4 years of per capita income	46%	66%	2.3
8 years of per capita income	60%	85%	2.1
16 years of per capita income	78%	96%	1.9
VSL ^b	57%	85%	2.1
VSL with low discounting ^c	79%	100%	1.8

Optimal decarbonization rates in 2030 and 2050, as well as the maximum global temperature rise given different life-year monetizations for health co-benefits. (The last two rows use an alternative method that monetizes lives using the VSL.)

^aAbove preindustrial levels. Values are rounded

^bValue of a statistical life. Equation: $VSL = 9,000,000 \times (GDP_{pc}/54,000)^{1.4}$ as in Robinson et al.²³

^cIn this run, time preference = 0.1%, compared to 1.5% in all other runs

Table 3 Additional sensitivity analyses

	2030	2050	Diff. in peak temperature (°C) ^a
Main result	52%	49%	0.4
Test 1	7–49%	10–50%	0.1–0.4
Test 2	51%	48%	0.4
Test 3	33%	30%	0.2
Test 4	70%	65%	0.5
Test 5	52%	49%	0.4
Test 6	3%	13%	<0.1
Test 7	30%	29%	0.3
Test 8	82%	76%	0.7

Percent increase in optimal decarbonization rates in RICE + AIR versus the corresponding reference case given different modeling assumptions. The difference in peak temperature is the difference in the highest global average temperature

^aThe difference in peak temperature between the reference case and the RICE + AIR case within each test: positive values indicate a lower peak temperature in the RICE + AIR case. Values are rounded

decades and thus explores the model’s sensitivity to their exclusion. Test 7 assumes that climate damages are twice the standard values in RICE, while the final test (Test 8) uses the Finite Amplitude Impulse Response (FAIR) climate model (version 1.0)²⁷ as an alternative to RICE’s native climate model (Supplementary Note 3 describes how we integrate the FAIR model).

Exploring model uncertainty by linking AIR to the FUND model. Like RICE/DICE (the Dynamic Integrated Climate Economy model (DICE) is RICE’s global counterpart), the FUND model is another one of the three leading climate economy models used by the US Interagency Working Group to estimate the social cost of carbon²⁸. FUND has different world regions, a different economic framework, a different climate model, and a different specification of damages when compared to RICE, and thus provides an important opportunity to explore model uncertainty. Comparing Fig. 1 with Fig. 6 demonstrates that results for FUND + AIR and RICE + AIR are qualitatively similar, despite the well-known differences in the structure and policy recommendations of the two models^{28,29}; adding aerosol impacts leads to dramatically increased optimal levels of decarbonization. (Supplementary Note 4 describes how we link FUND to AIR.)

Discussion

We have developed a new modeling framework for analyzing the costs and benefits of the co-reductions in air pollutant emissions that result from CO₂ policy. We find that these impacts have a critical role in determining optimal decarbonization rates, as the

potential health co-benefits that result from improved air quality are large, occur quickly enough to be economically important, outweigh the near-term co-harms from lost cooling, and are concentrated in developing regions. This remains true even with high discount rates and relatively conservative valuations of improved health. However, as the size of the health co-benefits will be partially determined by future air quality policies, decision makers should jointly plan both types of interventions. Depending on how society values health, it may be economically optimal to limit temperature rise to 2 °C or lower, thus corroborating the climate targets from the Paris Agreement. Overall, optimal mitigation results in immediate net benefits globally.

Our findings should be interpreted in light of several factors. First, our sensitivity analyses identified key variables that meaningfully affect the optimal level of decarbonization in RICE + AIR. These include the shape and magnitude of the exposure-response functions that relate PM_{2.5} exposure to mortality, the assumed relationship between CO₂ emissions and air pollutant emissions, the discount rate, and the valuation of the health benefits. The first two factors are largely amenable to empirical inquiry, whereas the latter two depend partially on ethical judgments; all have a strong bearing on how much society should mitigate, and when.

Second, we follow standard convention in assuming a uniform global carbon price in our optimal policy calculations. Economists prefer this assumption because a uniform price minimizes the cost to the global economy of any particular level of emissions reductions, and is thus a necessary part of the first-best approach to climate policy. However, another necessary part of first-best policy is a non-climate equity policy involving redistribution to address the large economic inequalities that exist throughout the world; in the absence of such an equity-focused policy, a uniform global carbon price would ignore equity considerations, and arguably impose an unjustifiably heavy burden on developing countries. An important feature of our results relevant to this issue is that health co-benefits provide the most incentive for additional decarbonization in lower- and middle-income regions. An interesting extension of our research would explore how co-benefits affect mitigation policies in a range of different burden-sharing regimes and in second-best climate policy calculations^{30,31}.

Third, we follow common practice in the co-benefits literature in assuming that there is a background appetite for air pollution policy that is independent of climate policy, such that the same background levels of investment will be made in air quality regardless of CO₂ reductions (and so policies will ratchet up insofar as greater CO₂ reductions make the antecedent level of air quality mitigation less expensive); this is the basic assumption behind the co-reductions we estimate in our main results. As reported above in Fig. 4, we also co-optimize air quality and




Fig. 6 FUND + AIR results. Optimal decarbonization rates over time for the FUND reference case (that excludes health benefits) and the FUND + AIR case (that includes health benefits)

climate policies, which is another scenario (with different properties) in which there would be co-benefits from CO₂ reductions. A limitation is that other scenarios are also possible, such as CO₂ reductions in a context involving a fixed cap and trade regime for air pollutants where the cap remains fixed at the same level in conjunction with new CO₂ reductions; in this scenario, it is possible that there would not be many co-benefits from CO₂ reductions³².

Fourth, we were not able to fully assess uncertainties around key aerosol processes. For example, the aerosol indirect effect may be the single largest source of uncertainty in radiative forcing assessments^{10,33}. New evidence continues to accrue about key factors that contribute to the formation of new particles from precursor gases, and there are likely to be complex feedback mechanisms that occur in response to future temperature changes^{34,35}. Some of these processes are not yet adequately represented in even the most sophisticated Earth system models, let alone reduced-form versions. A related concern is that we have used separate models to estimate the health effects (TM5-FASST) and climate effects (Model for the Assessment of Greenhouse gas Induced Climate Change 6 (MAGICC6)) of air pollutant emissions, a decoupling that may also introduce uncertainty. Nevertheless, MAGICC and TM5 have distinct strengths that we harness accordingly; MAGICC takes account of the full aerosol load of the whole atmosphere, top to bottom, and at the level of hemispheres, while TM5 tells us about the concentration at ground level, where people live and breathe, in principle at much higher resolution. Both models work with the same global emission inventories.

Fifth, we have only explored the effect of co-reducing air pollutant emissions on PM_{2.5}-related mortality. Other co-benefits may occur, which would likely push further in the direction of increased decarbonization; these include morbidity impacts from PM_{2.5}—which are generally minor compared to those from mortality—as well as other more indirect impacts such as potential increases in crop yields, improvements in visibility, and health impacts from changing exposure to tropospheric ozone^{2,9,35,36}. In addition, cost-benefit models also miss some other impacts of reduced climate change, including the effects of methane on ozone, the effects of ocean acidification, and others³⁷.

Due to the uncertainties inherent in our modeling framework, we expect the accuracy of quantitative estimates of the co-harms and co-benefits of optimal climate policy to sharpen over time, particularly as our understanding of atmospheric science

progresses. Nevertheless, the novel modeling approach described here offers important new insights into how much we should mitigate and over what time period, and the sensitivity tests above indicate that we should not expect the qualitative story told by our results to change in light of improved empirical estimates. Our methods also enable investigation of other key questions beyond the scope of this study, including how the inclusion of health co-benefits influences optimal climate policy under different burden-sharing regimes and different worldviews about how much to prioritize the poor, future people, and citizens of other countries.

Methods

The RICE model. The RICE model was first developed in 1996 to analyze the tradeoffs between investing in climate mitigation, which incurs a cost relatively soon, and climate damages, which incur costs in the more distant future^{38,39}. RICE is the regionalized counterpart to the DICE model, which is one of three leading cost-benefit climate economy models used by researchers and governments for regulatory analysis, including to estimate the social cost of carbon²⁸. Here we describe the key aspects of the standard RICE2010 model; for a more extensive description of this open-access model, see ref. ¹². (Also see ref. ³⁹ for a more extensively documented, but earlier version of RICE.)

Briefly, RICE is a regionalized global optimization model that includes an economic component and a geophysical component that are linked. RICE divides the world into 12 regions, some of which are single countries, while others are groups of countries. Each region has a distinct endowment of economic inputs, including capital, labor, and technology, which together produce that region's gross output via a Cobb–Douglas production function. Pre-mitigation carbon emissions are a function of gross output and an exogenously determined, region-specific carbon intensity pathway. These carbon emissions can be reduced (mitigated) at a cost to gross output through control policies, set to equalize the marginal abatement cost in all regions. Local mitigation cost is borne by each region, and there are no inter-regional transfers. Any remaining (post-mitigation) carbon emissions are incorporated into the climate module where they influence global temperature and, ultimately, the future economy through climate-related damages. Future climate change affects regions differently, with poorer regions generally more vulnerable to climate damages. Damage estimates increase quadratically with a change in the global surface temperature and, like mitigation costs, are incurred directly as the loss of a proportion of gross output. Gross output minus the loss of mitigation costs and climate damages is what we refer to hereafter as GDP.

The model's optimization balances mitigation costs, which lower consumption at the time of mitigation, against climate damages which lower consumption in the future. (Regional consumption is the fraction of GDP that is not saved; mitigation cost and climate damage affect consumption only via their effect on GDP). The optimal tradeoff maximizes the sum of discounted well-being, W , which is a concave function of consumption as follows:

$$W(c_{it}) = \sum_{it} \frac{L_{it} c_{it}^{(1-\eta)}}{(1+\rho)^t 1-\eta}, \quad (1)$$

where L is population, c per capita consumption, ρ the rate of pure time preference, and η the consumption elasticity of marginal utility (inequality aversion). The subscripts i and t are the region and time indices, respectively. The model is solved by maximizing this global objective function. As a result, any factor affecting consumption, such as health impacts or climate damages, can be included in the model's optimization framework. Unless otherwise specified, we maintain RICE's default parameter values for time preference and inequality aversion of 1.5% and 1.5, respectively.

For this study, all simulations were run using the Mimi model development package in the Julia programming language. This Mimi/Julia version of RICE is fully faithful to the standard Excel version, but has more flexibility⁴⁰. We make four changes to this standard version of RICE, in addition to those directly related to the AIR module, which are described below. First, we update the population projections to those of the UN2017 medium variant⁴¹. This is a newer source of projections and also enables us to use internally consistent projections of deaths and life-years, as described further below. The impact of changing population has been comprehensively explored elsewhere^{42–44}. Second, we update the exogenous radiative forcing terms to the values used in RCP6.0⁴⁵, which is in line with the latest versions of DICE^{46,47}, which is the global (single-region) variant of RICE. These estimates are newer than those found in RICE2010 and are available in disaggregated form, thus allowing us to remove and endogenize the individual aerosol term for use in the AIR module, while maintaining the other non-CO₂ forcings (also described in more detail below). Third, we allow CO₂ concentrations and the global temperature to be endogenous in the second time-step (2015–2025)—it is fixed in standard RICE2010—as the now-endogenized aerosols will produce effects immediately after mitigation. And fourth, we use a modified objective function that avoids Negishi weights, which distort time preferences⁴⁸ and the inter-regional tradeoff in ways that are opaque and difficult to justify descriptively

and normatively⁴⁹. We have explained this latter change in more detail in a previous publication⁵⁰ and also in a sensitivity presented in Supplementary Table 2.

In the standard RICE2010 model, anthropogenic CO₂ is the only endogenous climate forcer. All other sources of radiative forcing, including land-use change, non-CO₂ gases, and aerosols are represented through a single exogenous forcing term that aggregates the individual trajectories of each source. This simplifying assumption is problematic when it comes to aerosols. Mitigation actions affecting CO₂ have the potential to strongly influence emissions of the air pollutants that produce aerosols, as the two types of emissions share many sources^{4,5}. Therefore, if carbon emissions are reduced, aerosols will tend to decrease simultaneously. A change in aerosols implies a change in radiative forcing as well as a change in ambient particulate air pollution⁵. Climate change and air pollution both affect well-being, and capturing the impact of these pathways was the motivation for developing the AIR module, which we now explain.

Overview of the AIR module. In this section, we provide a general overview of how we developed the AIR module, with a technical description—including all equations—in the sections that follow.

Broadly speaking, our approach consists of five steps. First, we estimate the baseline (before carbon mitigation) emissions of five air pollutant species (primary PM_{2.5}, oxides of nitrogen, sulfur dioxide, organic carbon, and black carbon). Emissions are estimated for each region-time pair with income-dependent emission intensity projections (emissions per unit GDP) based on the Greenhouse gas and Air pollution Interactions and Synergies (GAINS) model and specifically the ECLIPSE emission scenarios⁵¹. Our central case assumes air pollutant emissions in the coming decades follow the ECLIPSEV5a baseline scenario, which includes current and planned air quality legislation but no climate policy. In sensitivity analyses we alter this assumption, allowing for faster or slower independent air quality cleanup, including a case where we simultaneously co-optimize both air quality and climate policy. This co-optimization introduces policy levers for end-of-pipe technologies that act on individual air pollutants. These levers require associated cost curves that are also drawn from the GAINS model.

In the second step, we determine the change in air pollutant emissions that would result from a change in CO₂ emissions using information from the SSP project⁵². This provides an estimate of co-reductions based on empirically realistic projections about the regionally differentiated interaction between future climate and air pollution policies, and is consistent with theoretical results from economics that show that co-benefit effects could be different in scenarios with different properties³². Emission information in the SSPs is estimated from bottom-up integrated assessment modeling that includes regionally differentiated, spatially explicit representations of energy production and structure and, like in RICE, assumes that CO₂ policy occurs through a single global carbon price.

Third, we link changes in air pollutant emissions to changes in estimated average human exposure to PM_{2.5} by applying the source-receptor matrix (SRM) from the TM5-FASST air quality model⁵³. The TM5-FASST SRM was computed from simulations of the full TM5 chemical transport model for 56 source regions, which were aggregated to approximate the RICE regions⁵³. Once exposure is estimated, it is possible to calculate the number of life-years gained attributable to (reduced) air pollution by combining an exposure-response function with projections of future mortality and life expectancy⁵⁴, which we took from the UN World Population Prospects. We applied a (log)linear exposure-response function for mortality from all causes in adults based on a meta-analysis published in a recent World Health Organization report²⁴. We selected this approach for consistency with the UN projections—which only estimate mortality from all causes—and because recent epidemiological analyses indicate that the strong effects of air pollution occur at exposure (concentration) levels up to and including those most relevant to our study and that they likely affect a wide range of outcomes^{25,26}.

Fourth, we allow the change in air pollutant emissions to influence the global temperature using aerosol forcing coefficients derived from the MAGICC climate model⁵⁵. The coefficients incorporate both the direct and indirect effects represented in MAGICC, with the latter including those related to albedo and cloud responses. Aerosol forcing is added to the forcing from the other greenhouse gases in RICE's climate module to produce estimates of future climate change.

And fifth, for each region we monetize and then value the aerosol impacts. The average per capita health benefits are added to per capita consumption, whereas the climate effects—monetized using RICE's standard climate damage function—subtract from consumption¹². Consumption is transformed into well-being by a concave function in the optimization via RICE's objective/social welfare function (Eq. (1)). Once the impacts have been valued in this way, they then enter the optimization.

Figure 7 illustrates the different model components and their linkages. We now present a technical description of each of these steps.

Estimating baseline emissions. Air pollutant emissions from natural sources and open burning remain exogenous in our framework, following the trajectory based on RCP6.0⁴⁵. All other anthropogenic air pollutant emissions are endogenous, as follows.

The level of baseline (pre-mitigation) anthropogenic emissions of each aerosol precursor, E^0 , is a function of an emission intensity factor (emissions per unit

GDP), e , and the GDP, Y :

$$E_{ip}^0(Y_{it}, L_{it}, e_{ip}(Y_{it}/L_{it})) = e_{ip}(Y_{it}/L_{it}) \cdot Y_{it}. \quad (2)$$

The emission intensity is region (i), time (t), and pollutant (p) specific and characteristically depends on the per capita income level of the region, defined above as the GDP, Y , divided by the population, L . GDP is endogenously determined in RICE while population is exogenous, taking the medium variant estimate of the 2017 version of the UN Population Prospects through 2100⁴¹, and remaining constant thereafter.

We estimated emission intensities (and emissions) for five air pollutants: sulfur dioxide (SO₂), primary fine particulate matter (PM_{2.5}), oxides of nitrogen (NO_x), organic carbon (OC), and black carbon (BC). Emission intensities up to the year 2050 were derived from region-specific projections used in the GAINS model and specifically the ECLIPSEV5a scenario that reflects existing national and regional air pollution policies, but excludes decarbonization from climate mitigation⁵¹.

We fit the following functional form to the scenario data to extrapolate beyond 2050:

$$e_{ip}(Y_{it}/L_{it}, \chi) = \varphi_{1,ip} \cdot \left[\exp\left(-\Omega_{1,ip} \cdot I_{it}(Y_{it}/L_{it}, \chi)\right) + \varphi_{2,ip} \cdot \exp\left(-\Omega_{2,ip} \cdot I_{it}(Y_{it}/L_{it}, \chi)\right) \right], \quad (3)$$

where the Ω s and φ s are fitting parameters and χ is the Kuznets-relevant income, as described below. Resulting emission intensities are displayed graphically in Supplementary Fig. 4.

This functional form implies that emission intensities decrease with rising per capita income I , as observed in the projected emission scenario. This can be interpreted as a particular version of the environmental Kuznets curve, estimated with the ECLIPSE data. We write our Kuznets-relevant income $I(\chi)$ as:

$$I_{it}(Y_{it}/L_{it}, \chi) = \frac{Y_{i(t=2005)}}{L_{i(t=2005)}} + \chi \cdot \left(\frac{Y_{it}}{L_{it}} - \frac{Y_{i(t=2005)}}{L_{i(t=2005)}} \right). \quad (4)$$

For $\chi = 1$, the Kuznets-relevant income equals the true (modeled) per capita income in the region-time pair, producing our default best-guess emission intensity factors. Changing $I(\chi)$ allows us to explore assumptions of more or less stringent autonomous air quality policies: we can speed up ($\chi > 1$) or slow down ($\chi < 1$) the decrease in emission intensities over time accordingly.

Supplementary Figure 5 shows the baseline (pre-mitigation) level of emissions by region and time under the default of $\chi = 1$, where the Kuznets-relevant income equals the true (modeled) income.

Relationship of CO₂ mitigation to air pollutant emissions. The above method projects the level of air pollutant emissions assuming no CO₂ mitigation. Reducing emissions of CO₂ will typically also reduce the emissions of air pollutants, as they often stem from the same sources. In our framework, the percentage reduction in CO₂ relative to the business-as-usual (without mitigation) scenario is called the decarbonization fraction (control rate), μ_{it} , and it is associated through parameter κ_{ip} with a reduction in pollutant p relative to its baseline (pre-mitigation) level:

$$\frac{\Delta E_{ip}(\mu_{it})}{E_{ip}^0} = \kappa_{ip} \frac{\Delta E_{it,CO_2}(\mu_{it})}{E_{it,CO_2}^0} = \kappa_{ip} \cdot \mu_{it} \quad (5)$$

or,

$$\Delta E_{ip}(\mu_{it}, Y_{it}, L_{it}) = \kappa_{ip} \cdot \mu_{it} \cdot E_{ip}^0 = \kappa_{ip} \cdot \mu_{it} \cdot e_{ip}(Y_{it}/L_{it}) \cdot Y_{it}. \quad (6)$$

The parameter κ_{ip} describes the effectiveness of CO₂ mitigation in co-reducing emissions of pollutant p . The κ_{ip} parameter was estimated from the SSPs, which are a set of five storylines designed to analyze tradeoffs between climate change and socioeconomic factors⁵². Each of the five SSPs contains multiple sub-scenarios that differ only by their level of decarbonization; all socioeconomic factors remain constant. Therefore, each pairwise comparison of these sub-scenarios includes an implicit estimate of κ_{ip} at each time period for each region and pollutant. The emission information in the SSPs is estimated from bottom-up integrated assessment modeling that includes regionally differentiated, spatially explicit representations of energy production and structure, and like RICE, assumes that CO₂ policy occurs through a single global carbon price.

We fit a simple linear regression line through the multiple estimates of κ for each SSP, constrained to begin at the origin, and take the slope of that line as our estimate (Supplementary Fig. 6). We derive five estimates of κ (one for each SSP) for each region-pollutant pair, using the middle-of-the-road SSP2 as our standard case, with the alternative SSPs tested in sensitivity analyses. The total post-mitigation level of emissions is therefore:

$$E_{ip}(\mu_{it}) = E_{ip}^0 - \Delta E_{ip} = E_{ip}^0 \cdot (1 - \kappa_{ip} \cdot \mu_{it}). \quad (7)$$

We acknowledge that there are cases where the data derived from the SSP database appear to exhibit a non-linear rather than linear relationship between CO₂ reduction and air pollutant reduction (Supplementary Fig. 6). Our goal was to find a relationship that on average, and in particular at full mitigation, provides a reasonable approximation to the implied air pollution reduction. Having said this, we are aware that our assumption of linearity in the relation may affect our overall results: a higher co-reduction at low CO₂ mitigation offers incentives for further




Fig. 7 Diagram of the AIR module. Flow chart illustrating how the AIR module (rectangles) links with the RICE model (gray circle) to estimate emissions of air pollutants and their impacts. The model/method underlying each step in AIR is shown within parentheses, along with the relevant equations

reduction than in the linear case, and vice versa. Thus, the shape of the mitigation profile could be affected, though the effect is likely to be small if, as we assume, the carbon prices in the RICE regions are coupled to each other.

From air pollutant emissions to health impacts. The health co-benefits in RICE + AIR are calculated directly from the change in ambient population-weighted concentrations of PM_{2.5} attributable to CO₂ mitigation. In the next paragraph, we describe how we estimate this change in concentration. Further below we describe how we keep track of the absolute level of pre- and post-mitigation PM_{2.5} concentrations, which are used in the health impact calculations only to ensure that no health benefits accrue at exposures below given threshold values.

The change in ambient concentrations of PM_{2.5} attributable to CO₂ mitigation is a function of the change in aerosol precursor emissions—in this case primary PM_{2.5}, NO_x, and SO₂—as well as other factors such as meteorological conditions. To estimate this relationship, we extracted the SRM from the freely available TM5-FASST global atmospheric SRM⁵³. For each pollutant, the SRM provides an estimate of the change in population-weighted PM_{2.5} concentrations (hereafter referred to as exposure) given a unit change in emissions. The SRM from TM5-FASST was computed for 56 source regions from simulations with the full TM5 chemical transport model⁵³. Using an SRM is a practical alternative to running full atmospheric chemistry transport model simulations, which is infeasible in our optimization context. The TM5-FASST model is described in detail elsewhere, and has been used similarly in other projects^{3,53}.

The TM5-FASST interface allows the 56 source regions to be aggregated into larger regions that approximate the RICE regions. Due to the size of the RICE regions, we assume that the change in exposure to PM_{2.5}, ΔC, in region *i* depends only on the change in emissions within that same region (*i*) and that the change is estimated via the SRM (SR) that encodes atmospheric transport. We also assume that the SRM does not change over time and that changes in the three precursor pollutants are additive:

$$\begin{aligned} \Delta C_{it}(\mu_{it}, Y_{it}, L_{it}) &= \sum_{ip} SR_{iip} \cdot \Delta E_{iip}(\mu_{it}, Y_{it}, L_{it}) \\ &= \sum_p SR_{iip} \cdot \Delta E_{iip}(\mu_{it}, Y_{it}, L_{it}). \end{aligned} \tag{8}$$

As mentioned above, we keep track of the absolute level of exposure, a variable that is used only to ensure that no health benefits accrue at exposures below a given threshold level (described further below). The absolute exposure levels in the pre-mitigation case (where μ = 0) are calculated by translating the change in emissions relative to 2005 into a change in exposure using the SRM, and then subtracting it from the 2005 exposure

$$C_{it}^0(\mu_{it} = 0) = \max \left(C_{i,t=2005} - \sum_{ip} SR_{iip} \cdot (E_{i,t=2005,p} - E_{iip}^0), 0 \right). \tag{9}$$

Here the max function ensures PM_{2.5} concentrations do not drop below zero. Emissions and exposure in 2005 were taken from the EDGAR (Emission Database for Global Atmospheric Research) emission database⁵⁶ and Brauer et al.⁵⁷, respectively, and then aggregated into the RICE regions.

Mitigation of CO₂ (μ > 0) reduces exposure further:

$$C_{it}(\mu_{it}) = \max \left(C_{i,t=2005} - \sum_{ip} SR_{iip} \cdot (E_{i,t=2005,p} - (E_{iip}^0 - \Delta E_{iip})), 0 \right) \tag{10}$$

$$= \max \left(C_{i,t=2005} - \sum_{ip} SR_{iip} \cdot (E_{i,t=2005,p} - E_{iip}(\mu_{it})), 0 \right) \tag{11}$$

We define the health co-benefit as the avoided premature mortality resulting from reductions in PM_{2.5} exposure attributable to CO₂ mitigation (ΔC from Eq. (8)). This benefit can be quantified through the attributable fraction (AF)⁵⁴:

$$AF_{it} = \frac{RR_{it} - 1}{RR_{it}}, \tag{12}$$

where the relative risk, RR, for each region and each time period is a function of the change in exposure and a health impact function (β) that links a unit change in exposure to a change in the risk of adult (≥30) mortality from all causes:

$$RR_{it} = \exp(\beta \cdot \Delta C_{it}(\mu_{it}, Y_{it}, L_{it})). \tag{13}$$

We assume a log-linear relationship between PM_{2.5} exposure and all-cause mortality with a relative risk of 1.066 (95% confidence interval (CI) = 1.040, 1.093) for each 10 μg/m³ change in exposure, based on a meta-analysis published by the World Health Organization²⁴. However, we note that many recent assessments of ambient air pollution have used the cause-specific integrated exposure-response (IER) functions to estimate mortality impacts⁵⁸. Here we focus on all-cause mortality for two reasons. First, we have recently shown that population size/growth strongly affects estimates of optimal climate policy, including for reasons unrelated to human health^{42,43}. Therefore, we use the most recent long-term (to 2100) population projections provided by the UN Population Division, which does not publish corresponding estimates of cause-specific mortality. Second, important recent studies indicate that the IER functions may underestimate excess mortality^{25,26}, and suggest that mortality risks at the exposure levels seen in our regional analyses may fall within the 95% CI of the WHO estimate presented above and tested in sensitivity analyses²⁵.

In all analyses, we assume that components of PM_{2.5} are equally toxic²⁴, that health benefits accrue in the same 10-year time-steps as the improvement in air quality, and that population and mortality remains constant after 2100. We also confine the health impacts to mortality from PM_{2.5} exposure, though note that there is also concern about the health effects of exposure to smaller particles⁵⁹.

Multiplying the AF by the total number of life-years lost from all causes, Θ, yields the life years gained from the reduction in air pollution:

$$LY_{it}(\mu_{it}, Y_{it}, L_{it}) = AF_{it} \cdot \Theta_{it}, \tag{14}$$

$$= \frac{RR_{it} - 1}{RR_{it}} \cdot \Theta_{it}, \tag{15}$$

$$= (1 - \exp(-\beta \cdot \Delta C_{it}(\mu_{it}, Y_{it}, L_{it}))) \cdot \Theta_{it}. \tag{16}$$

Since $\beta \cdot \Delta C_{it}(\mu_{it}, Y_{it}, L_{it}) \ll 1$, we can write as an approximation:

$$LY_{it}(\mu_{it}, Y_{it}, L_{it}) = \beta \cdot \Delta C_{it}(\mu_{it}, Y_{it}, L_{it}) \cdot \Theta_{it}. \quad (17)$$

Theta (Θ) values are estimated from the UN data by multiplying the total deaths by the remaining life expectancy at the age of death. As UN life expectancy data is by exact age, reported at 5-year intervals, whereas mortality data is for 5-year age groups, remaining life expectancy for each 5-year age group was taken as the average of the group's bounding ages. For example, remaining life expectancy for all (averted) deaths in the 30–34 age group would be the average of the remaining life expectancy for a 30 year old and a 35 year old.

Health benefits in a given region can accrue until absolute exposure to $PM_{2.5}$ converges to some minimum level, representing either a point below which no additional health impacts occur (a threshold) or a theoretical minimum level where residual $PM_{2.5}$ consists only of natural sources. We followed recent global studies (e.g., ref. 60) and chose a lower threshold/theoretical minimum of $5.8 \mu\text{g}/\text{m}^3$. However, we acknowledge that there is a growing consensus that there may not be a safe level of $PM_{2.5}$ below which no adverse health effects occur^{59,61} and therefore ran sensitivities down to $1 \mu\text{g}/\text{m}^3$. If a reduction in a given time period brings a region's exposure below the threshold, the health co-benefit is calculated from the increment between the unmitigated level and the threshold:

$$LY_{it}^*(\mu_{it}, Y_{it}, L_{it}) = \beta \cdot \Theta_{it} \cdot \max[\min(\Delta C_{it}(\mu_{it}, Y_{it}, L_{it}), (C_{it}^0(\mu_{it}, Y_{it}, L_{it}) - \tau)), 0], \quad (18)$$

where τ is the value of the threshold level.

Radiative forcing and temperature effects from aerosols. Some air pollutants are climate forcers: BC is a warming agent while SO_2 , NO_x , and OC act to cool the atmosphere¹⁰. The net global forcing in each time period attributable to aerosols is taken as the sum of the individual contributions:

$$RF_t^{\text{aer}} = \sum_p \sum_i r_{ipt} \cdot E_{itp}(\mu_{it}, Y_{it}, L_{it}), \quad (19)$$

where r_{ipt} is a region-specific coefficient that relates the regional change in emissions of pollutant p to the change in average global forcing.

We used the MAGICC6 climate model⁵⁵ to derive the coefficients by determining the impact on forcing from a pulse change in emissions in each time period in each region:

$$r_{ipt} = \left[\frac{\partial RF_t^{\text{aer}}}{\partial E_{itp}} \right]_{\text{MAGICC6}}. \quad (20)$$

For this we ran MAGICC6 with a pre-defined representative concentration pathway (RCP) scenario and then again after having reduced the emissions of one pollutant in one region by a marginal amount. We repeated this procedure for each region, pollutant, time-step, and, finally, each RCP. We thus derived a reduced-form surface-response representation of aerosol forcing in MAGICC6. We separated the effects of the different pollutants to allow them to be controlled independently, as they are in reality (also see Supplementary Note 1). In this experiment, we observed a time dependence that accounts for changes in atmospheric dynamics as emissions accumulate:

$$r_{ipt} = u_{1,ip} \cdot t + u_{0,ip}. \quad (21)$$

The resulting coefficients incorporate both the direct and indirect forcing effects represented in MAGICC6, with the latter including those related to albedo and cloud responses. The time-dependence is independent of the initial conditions. We noted that r depends on the magnitude of the pulse change, but the effect is relatively small. Supplementary Table 3 reports values of r_{ipt} and Supplementary Fig. 7 shows the radiative forcing over time for three scenarios. Here the reader will note that we have used two separate models for estimating the health effects (TM5-FASST) and climate effects (MAGICC6) of air pollutant emissions, a decoupling that may introduce uncertainty. Nevertheless, MAGICC and TM5 have distinct strengths that we harness accordingly; MAGICC takes account of the full aerosol load of the whole atmosphere, top to bottom, and at the level of hemispheres, while TM5 tells us about the concentration at ground level, where people live and breathe, in principle at much higher resolution. Both models work with the same global emission inventories.

Aerosol forcing affects global mean atmospheric temperature just as CO_2 forcing does, so that the added atmospheric temperature flow is equal to:

$$\Delta T_t^{\text{atm}} = \xi (RF_t^{CO_2} + RF_t^{\text{aer}}), \quad (22)$$

where $RF_t^{CO_2}$ is CO_2 forcing and ξ is the decadal speed of adjustment for atmospheric temperature (equal to 0.208).

Aerosol feedbacks on the economy. As described, aerosol impacts occur from a change in radiative forcing and from a change in air quality. Changes in radiative forcing are transferred to RICE's climate module where they influence the global average surface temperature, which is the basis of the monetized climate damage estimates³⁹.

Changes in air quality are monetized as the health co-benefit, B , by multiplying the number of life years gained by the value of a life-year (VOLY). We follow the same approach taken in early versions of RICE's climate damage function where a VOLY is assumed to equal 2 years of regional per capita consumption, c^{39} .

$$B_{it}(\mu_{it}, c_{it}) = \text{VOLY}_{it} c_{it}^{\text{pre-health}} \cdot LY_{it}^*(\mu_{it}, Y_{it}, L_{it}). \quad (23)$$

A VOLY of 2 years of per capita consumption is generally the same order of magnitude as empirical estimates based on willingness-to-pay surveys²¹, but we also test several alternative values in sensitivity analyses. Final (post-health) per capita consumption, c_{it} is calculated as:

$$c_{it} = c_{it}^{\text{pre-health}} + B_{it}/L_{it}. \quad (24)$$

With monetized aerosol impacts now included in the economic framework, RICE can follow its normal optimization procedure to find the decarbonization pathway that maximizes the objective (Eq. (1)).

Additional information on select sensitivity analyses. The Supplementary Information contains additional information on the co-optimization of air quality and climate policies (Supplementary Note 1 and Supplementary Fig. 8), life-year monetization and valuation (Supplementary Note 2), integrating the FAIR climate module into RICE + AIR (Supplementary Note 3), and the development of FUND + AIR (Supplementary Note 4).

Data availability

The authors declare that all data supporting the findings of this study are available within the article and its Supplementary Information files.

Code availability

The AIR model is licensed under the open source MIT license. The code is freely available at <https://github.com/Environment-Research/AIR>.

Received: 24 July 2018 Accepted: 6 March 2019

Published online: 07 May 2019

References

1. Markandya, A. et al. Health co-benefits from air pollution and mitigation costs of the paris agreement: a modelling study. *Lancet Planet. Health* **2**, e126–e133 (2018).
2. Shindell, D., Faluvegi, G., Seltzer, K. & Shindell, C. Quantified, localized health benefits of accelerated carbon dioxide emissions reductions. *Nat. Clim. Change* **8**, 291 (2018).
3. Rao, S. et al. A multi-model assessment of the co-benefits of climate mitigation for global air quality. *Environ. Res. Lett.* **11**, 124013 (2016).
4. West, J. J. et al. Co-benefits of mitigating global greenhouse gas emissions for future air quality and human health. *Nat. Clim. Change* **3**, 885–889 (2013).
5. Shindell, D. et al. Simultaneously mitigating near-term climate change and improving human health and food security. *Science* **335**, 183–189 (2012).
6. Shindell, D. et al. Climate, health, agricultural and economic impacts of tighter vehicle-emission standards. *Nat. Clim. Change* **1**, 59 (2011).
7. Ikefuji, M., Magnus, J. R. & Sakamoto, H. The effect of health benefits on climate change mitigation policies. *Clim. Change* **126**, 229–243 (2014).
8. Bollen, J., van der Zwaan, B., Brink, C. & Eerens, H. Local air pollution and global climate change: a combined cost-benefit analysis. *Resour. Energy Econ.* **31**, 161–181 (2009).
9. Shindell, D. The social cost of atmospheric release. *Clim. Change* **130**, 313–326 (2015).
10. Myhre, G. et al. in *Climate change 2013: The Physical Science Basis* (eds Stocker, T. et al.) Ch. 8 (Cambridge University Press, Cambridge and New York, 2013).
11. Clarke, L. et al. in *Climate Change 2014. Mitigation of Climate Change. Contribution of Working Group III to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change* (eds Edenhofer, O. et al.) (Cambridge University Press, Cambridge, 2014).
12. Nordhaus, W. D. Economic aspects of global warming in a post-copenhagen environment. *Proc. Natl Acad. Sci. USA* **107**, 11721–11726 (2010).
13. Waldhoff, S. T., Anthoff, D., Rose, S. & Tol, R. S. The marginal damage costs of different greenhouse gases: an application of fund. *Economics* **8**, 1–33 (2014).
14. Saari, R. K., Selin, N. E., Rausch, S. & Thompson, T. M. A self-consistent method to assess air quality co-benefits from us climate policies. *J. Air Waste Manag. Assoc.* **65**, 74–89 (2015).
15. Li, M. et al. Air quality co-benefits of carbon pricing in china. *Nat. Clim. Change* **8**, 398–403 (2018).
16. Kolstad, C. et al. in *Climate Change 2014: Mitigation of Climate Change. Contribution of Working Group III to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change* (eds Edenhofer, O. et al.) (Cambridge University Press, Cambridge, 2014).

17. National Academies of Sciences, Engineering, and Medicine. *Valuing Climate Damages: Updating Estimation of the Social Cost of Carbon Dioxide* (National Academies Press, Washington, 2017).
18. Stern, N. H. *The Economics of Climate Change: The Stern Review* (Cambridge University Press, Cambridge, 2007).
19. Executive Office of the President. *Executive Order 13783: Promoting Energy Independence and Economic Growth* (2017).
20. Tol, R. S. The economic effects of climate change. *J. Econ. Perspect.* **23**, 29–51 (2009).
21. Desaiques, B. et al. Economic valuation of air pollution mortality: a 9-country contingent valuation survey of value of a life year (vol). *Ecol. Indic.* **11**, 902–910 (2011).
22. Committee on the Medical Effects of Air Pollutants. The mortality effects of long-term exposure to particulate air pollution in the United Kingdom. Health Protection Agency Report (Health Protection Agency, London, 2010).
23. Robinson, L. A., Hammitt, J. K. & Okeeffe, L. *Valuing Mortality Risk Reductions in Global Benefit-cost Analysis*. Working Paper No. 7: Guidelines for Benefit-Cost Analysis Project (2018).
24. Forestiere, F., Kan, H. & Cohen, A. *Background Paper 4: Updated Exposure-Response Functions Available for Estimating Mortality Impacts* (World Health Organization Europe, Copenhagen, 2014).
25. Yin, P. et al. Long-term fine particulate matter exposure and nonaccidental and cause-specific mortality in a large national cohort of chinese men. *Environ. Health Perspect.* **125**, 117002 (2017).
26. Burnett, R. et al. Global estimates of mortality associated with long-term exposure to outdoor fine particulate matter. *Proc. Natl Acad. Sci. USA* **115**, 9592–9597 (2018).
27. Millar, R. J., Nicholls, Z. R., Friedlingstein, P. & Allen, M. R. A modified impulse-response representation of the global near-surface air temperature and atmospheric concentration response to carbon dioxide emissions. *Atmos. Chem. Phys.* **17**, 7213–7228 (2017).
28. Interagency Working Group on Social Cost of Greenhouse Gases. *Technical Update of the Social Cost of Carbon for Regulatory Impact Analysis Under Executive Order 12866* (United States Government, 2016).
29. Adler, M. Priority for the worse-off and the social cost of carbon. *Nat. Clim. Change* **7**, 433–449 (2017).
30. Chichilnisky, G. & Heal, G. Who should abate carbon emissions?: An international viewpoint. *Econ. Lett.* **44**, 443–449 (1994).
31. Anthoff, D. *Working Paper: Optimal Global Dynamic Carbon Abatement* (University of California, Berkeley, 2011).
32. Fullerton, D. & Karney, D. H. Multiple pollutants, co-benefits, and suboptimal environmental policies. *J. Environ. Econ. Manag.* **87**, 52–71 (2018).
33. Boucher, O. et al. in *Climate Change 2013: The Physical Science Basis. Contribution of Working Group I to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change* (eds Edenhofer, O. et al.) (Cambridge University Press, Cambridge, 2013).
34. Fuzzi, S. et al. Particulate matter, air quality and climate: lessons learned and future needs. *Atmos. Chem. Phys.* **15**, 8217–8299 (2015).
35. Silva, R. A. et al. The effect of future ambient air pollution on human premature mortality to 2100 using output from the accimp model ensemble. *Atmos. Chem. Phys.* **16**, 9847–9862 (2016).
36. Chang, K. M. et al. Ancillary health effects of climate mitigation scenarios as drivers of policy uptake: a review of air quality, transportation and diet co-benefits modeling studies. *Environ. Res. Lett.* **12**, 113001 (2017).
37. Scovronick, N. et al. Human health and the social cost of carbon: a primer and call to action. *Epidemiology* (2019). In press.
38. Nordhaus, W. D. & Yang, Z. A regional dynamic general-equilibrium model of alternative climate-change strategies. *Am. Econ. Rev.* **86**, 741–765 (1996).
39. Nordhaus, W. D. & Boyer, J. *Warming The World: Economic Models of Global Warming* (MIT Press, Cambridge, 2000).
40. Anthoff, D., Errickson, F. & Rennels, L. Mimi-rice-2010.jl v1.1.0.
41. United Nations Population Division. *World Population Prospects: 2017 Revision* (United Nations Department of Economic and Social Affairs, 2017).
42. Scovronick, N. et al. Impact of population growth and population ethics on climate change mitigation policy. *Proc. Natl Acad. Sci. USA* **114**, 12338–12343 (2017).
43. Budolfson, M. et al. Optimal climate policy and the future of world economic development. *World Bank Econ. Rev.* **33**, 21–40 (2018).
44. Gillingham, K. et al. Modeling uncertainty in integrated assessment of climate change: a multi-model comparison. *J. Assoc. Environ. Resour. Econ.* **5**, 791–826 (2018).
45. Meinshausen, M. et al. The RCP greenhouse gas concentrations and their extensions from 1765 to 2300. *Clim. Change* **109**, 213–241 (2011).
46. Nordhaus, W. & Sztorc, P. *Dice 2013R: Introduction and User's Manual* 2nd edn (2013). http://www.econ.yale.edu/%7Enordhaus/homepage/documents/DICE_Manual_100413r1.pdf.
47. Nordhaus, W. D. Revisiting the social cost of carbon. *Proc. Natl Acad. Sci. USA* **114**, 1518–1523 (2017).
48. Dennig, F. & Emmerling, J. *A Note on Optima with Negishi Weights*. Working Paper (2017).
49. Stanton, E. A. Negishi welfare weights in integrated assessment models: the mathematics of global inequality. *Clim. Change* **107**, 417–432 (2011).
50. Dennig, F., Budolfson, M. B., Fleurbaey, M., Siebert, A. & Socolow, R. H. Inequality, climate impacts on the future poor, and carbon prices. *Proc. Natl Acad. Sci. USA* **112**, 15827–15832 (2015).
51. Stohl, A. et al. Evaluating the climate and air quality impacts of short-lived pollutants. *Atmos. Chem. Phys.* **15**, 10529–10566 (2015).
52. Rao, S. et al. Future air pollution in the shared socio-economic pathways. *Glob. Environ. Change* **42**, 346–358 (2017).
53. Van Dingenen, R. et al. Tm5-fasst: a global atmospheric source-receptor model for rapid impact analysis of emission changes on air quality and short-lived climate pollutants. *Atmos. Chem. Phys.* **18**, 16173–16211 (2018).
54. Pruss-Ustun, A., Mathers, C. C., Corvalan, C. & Woodward, A. *Assessing the Environmental Burden of Disease at National and Local Levels*. World Health Organization Report (World Health Organization, Geneva, 2003).
55. Meinshausen, M., Raper, S. & Wigley, T. Emulating coupled atmosphere-ocean and carbon cycle models with a simpler model, magicc6-part 1: model description and calibration. *Atmos. Chem. Phys.* **11**, 1417–1456 (2011).
56. European Commission Joint Research Centre. Emission Database for Global Atmospheric Research (EDGAR) (2016).
57. Brauer, M. et al. Ambient air pollution exposure estimation for the global burden of disease 2013. *Environ. Sci. Technol.* **50**, 79–88 (2016).
58. Cohen, A. J. et al. Estimates and 25-year trends of the global burden of disease attributable to ambient air pollution: an analysis of data from the global burden of diseases study 2015. *Lancet* **389**, 1907–1918 (2017).
59. World Health Organization Regional Office for Europe. Review of evidence on health aspects of air pollution: Revihaap project. Technical Report (2013).
60. Apte, J. S., Marshall, J. D., Cohen, A. J. & Brauer, M. Addressing global mortality from ambient pm2.5. *Environ. Sci. Technol.* **49**, 8057–8066 (2015).
61. Di, Q. et al. Air pollution and mortality in the medicare population. *N. Engl. J. Med.* **376**, 2513–2522 (2017).

Acknowledgements

We thank Ciara Burnham and the Climate Future's Initiative at Princeton University for support. M.B. was also supported by a Catalyst Award from the Gund Institute for Environment at the University of Vermont.

Author contributions

M.B., F.D., F.E., M.F., W.P., N.S., R.H.S., D.S., and F.W. helped design the research, interpret results, and edit the manuscript. M.B. and F.D. led the economic modeling. F. D., F.E., and F.W. led model development. N.S. led health impact modeling. F.W. led emission and climate modeling. N.S. and M.B. prepared the manuscript.

Additional information

Supplementary Information accompanies this paper at <https://doi.org/10.1038/s41467-019-09499-x>.

Competing interests: The authors declare no competing interests.

Reprints and permission information is available online at <http://npg.nature.com/reprintsandpermissions/>

Journal peer review information: *Nature Communications* thanks Marcus Sarofim, Anil Markandya and the other anonymous reviewer(s) for their contribution to the peer review of this work.

Publisher's note: Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2019