Massachusetts Institute of Technology

# Two-Thirds of the World Doesn't Exist Yet:

## How Population Projections to 2150 Redraw the Global Tobacco Epidemic

Meyline Szczepanski

Advised by Professor Julian Jamison

Research conducted at Nuffield College, University of Oxford

Laidlaw Scholars Programme

August 2026

**Abstract**

Long-run health policy weighs costs paid now against benefits realised decades later, and discounts the difference at a fixed rate. Two quantities that method depends on are rarely measured: how many future people a policy reaches, and how many of them do not yet exist. The UN World Population Prospects, the reference dataset for such work, stops at 2100. This study extends the WPP 2024 medium variant to 2150 for ten world regions, blending a fertility-driven growth path continuously with the demographic momentum already present at 2100 while the age structure evolves under a cohort-component model. A cohort-accounting layer then splits the living population in each projected year into people already born in 2024 and people born later. Of the 1.25 trillion person-years lived worldwide between 2024 and 2150, 68.3% belong to people not yet born, ranging from 38.9% in China to 83.3% in Sub-Saharan Africa. Applied to a regional model of smoking-attributable mortality, 639 million tobacco deaths are projected for 2025 to 2100, peaking at 9.16 million a year in 2058 even as prevalence falls from 17.6% to 7.6%. A 3% discount rate cuts the future share of person-years to 38.0% and moves measured tobacco burden away from Sub-Saharan Africa and the Middle East, where the epidemic is still growing. Discounting is not only a temporal operation but a distributional one.

## I. Introduction

Tobacco control, pension design and climate mitigation all require comparing a cost paid today against a benefit received by someone who may not be born for another fifty years. Economics supplies a standard instrument for that comparison, the exponential discount factor, and a long argument about what rate to use; Ramsey held that discounting welfare purely because it is future is indefensible, and the Stern and Nordhaus dispute over climate policy turned largely on the same point. Population ethics asks a related but distinct question: whether people who do not yet exist have claims on us at all, and whether those claims are as strong as the claims of the living (Parfit, 1984; Broome, 2004; Greaves and MacAskill, 2021).

Both literatures rest on a quantity neither measures. To know what discounting does, one needs to know how many people sit on each side of the discount factor. That is unavailable past 2100, because the UN World Population Prospects terminates there, and analysts working beyond it either freeze the population or extrapolate one global growth rate, neither of which preserves the regional divergence that dominates the century. The cutoff is arbitrary from a policy standpoint rather than a demographic one. A child born today may still be alive in 2110, and a cohort that takes up smoking in 2090 will die of it well into the twenty-second century, so the questions that most need long-run demography are precisely the ones that run past the point where the data stops.

A second gap sits inside the UN's own horizon. Projections are reported as totals, and a total conceals the fact that the population of 2075 consists overwhelmingly of different individuals from the population of 2025. An analyst reading that world population is near 10 billion in both years may conclude that little has changed, when the entire moral constituency has been replaced. Whether that replacement matters is precisely what separates person-affecting views from totalist ones, and it cannot be assessed without the cohort composition underneath the total.

This study extends WPP 2024 to 2150 for ten world regions in a way that stays continuous with the UN model at the 2100 boundary; adds a cohort-accounting layer that splits each projected year into people already born at a reference year and people born after it, weighted separately from the discount factor; and applies the result to smoking-attributable mortality, asking what share of projected tobacco deaths falls on people who do not yet exist and how discounting redistributes that burden.

## II. Methodology

**Data.** Population figures come from UN World Population Prospects 2024, medium variant, by single year of age and sex, 2024 to 2100 (United Nations, 2024). Countries were aggregated into the ten Lancet Commission on Investing in Health regions, restricted to UN member states with a 2023 population above 300,000, with fifteen territories assigned to the region containing them. Twenty-one small member states and forty-three non-member territories fall outside the scheme, so world totals are the sum of the ten regions. Smoking prevalence comes from the Institute for Health Metrics and Evaluation's Global Burden of Disease smoking tobacco use estimates, reported by sex, age group and year for 204 countries and territories in fifteen five-year age brackets (IHME, 2020); the historical series used here runs from 1990 to 2022, with projection beginning in 2023.

**Extending to 2150.** Total population and age structure are extended separately so fertility can vary without the pyramid becoming implausible. For total population, the growth rate each region actually achieved over the final decade of the UN projection is taken from the data as $r_{2100} = (P_{2100} / P_{2090})^{(1/10)} - 1$, capturing the momentum, fertility and migration the UN had already built in. A target fertility rate F converts to a long-run growth rate as $r\_target = 0.005 \times (F - 2.1)$, where 2.1 is replacement level and the coefficient encodes the approximation that each unit of TFR above or below replacement corresponds to about half a percentage point of annual growth. Rather than switching at 2101, the rate moves linearly over a transition period T:

$$r(t) = r_{2100} + (r\_target - r_{2100}) \times \min(t/T,\ 1),$$

with t years since 2100 and T set to 50 by default, after which population compounds year on year. The linear blend matters more than it may appear. A step change at 2101 would produce a visible kink in every regional series and would attribute to fertility assumptions a discontinuity that is really an artefact of where one dataset ends and another begins; the fifty-year transition instead lets momentum decay at the speed real populations decay. Each region's default target TFR is the UN's own projected value at 2100, so the extension continues the UN model rather than departing from it, while F and T remain editable for scenario comparison. Age structure is handled by a full cohort-component model run from 2101 to 2150 at the UN's age-specific mortality rates at 2100, held constant; the resulting age shares are applied to the total, which responds live to the fertility inputs. Migration is not modelled after 2100, being embedded in the starting growth rate and negligible beside fertility and mortality here.

**Cohort accounting.** A person alive in year Y at age a was born in Y − a. For reference year X = 2024 this divides each projected year into C(Y), those already born, and D(Y), those born afterwards, so that weighted population is $W(Y) = C(Y) + \alpha \cdot D(Y)$. Here α is the moral weight on people who do not yet exist: 1 under a totalist view, 0 under a strict person-affecting view. It is held separate from the discount factor, so that weighting by existence and discounting by delay can be told apart rather than folded into one rate. Combining the two operations, the quantity a planner evaluating a policy over a horizon of T years maximises is

$$V = \Sigma_{t=0}^{T} \sigma^t \,[\, C(X+t) + \alpha \cdot D(X+t) \,]$$

so that σ governs how heavily delay is penalised and α governs how heavily existence status is weighted. The two enter multiplicatively and independently, which is the point of holding them apart: a given value of V can be reached either by discounting steeply and weighting the unborn fully, or by discounting lightly and weighting them near zero, and these are different ethical positions that a single rate would render indistinguishable. In the tobacco application the same form applies with deaths in place of population, C and D partitioning each year's deaths by whether the decedent was born before or after the reference year.

**Tobacco.** Historical prevalence is taken from the GBD file. From 2023 it is projected at each region's own implied trend, the compound annual change in that region's aggregate prevalence between 2012 and 2022, computed separately by sex; the observed decline already reflects real cessation behaviour, so no extra quit-rate assumption is layered on top. Smokers are prevalence multiplied by projected population, and deaths in year $t$ are smokers in year $t-1$ multiplied by an age- and sex-specific fatality rate. The base schedule rises from roughly 0.02% a year at ages 10 to 14 to roughly 6% at 80 and above, calibrated to the excess mortality among current smokers reported by Jha et al. (2013) and Pirie et al. (2013), who find hazard ratios near 2.5 to 3.0 in older ages. A maturity multiplier scales it to 1.0 for mature epidemics, 0.75 for intermediate and 0.5 for newer ones, reflecting later initiation and shorter duration, with a sex split applied by tier.

## III. Results

**Population stops growing but never stops changing.** World population peaks at 10.28 billion in 2084 and falls to 9.75 billion by 2150, only 19% above its 2024 level. That near-stability hides extreme divergence (Table 1, Figure 1). Sub-Saharan Africa triples, from 1.29 to 3.90 billion, while China falls to 27% of its current size and Central and Eastern Europe contracts by 42%. By 2150 Sub-Saharan Africa alone holds 40% of world population, against 16% today.

***Table 1.*** *Regional population under the extended projection, in millions. Target TFR is the UN's projected fertility rate at 2100, carried forward as the default long-run assumption.*

| **Region** | **Target TFR** | **2024** | **2050** | **2100** | **2150** |
|---|---|---|---|---|---|
| World | 1.83 | 8,158 | 9,660 | 10,176 | 9,746 |
| Sub-Saharan Africa | 2.03 | 1,293 | 2,178 | 3,488 | 3,895 |
| Central Asia | 1.95 | 390 | 578 | 807 | 848 |
| Middle East and North Africa | 1.78 | 604 | 790 | 914 | 887 |
| United States | 1.65 | 345 | 381 | 421 | 414 |
| India | 1.69 | 1,451 | 1,680 | 1,505 | 1,269 |
| North Atlantic | 1.57 | 472 | 470 | 431 | 391 |
| Western Pacific and SE Asia | 1.67 | 1,169 | 1,282 | 1,132 | 963 |
| Latin America and Caribbean | 1.64 | 662 | 729 | 612 | 500 |
| Central and Eastern Europe | 1.57 | 321 | 286 | 220 | 185 |
| China | 1.34 | 1,451 | 1,286 | 646 | 395 |

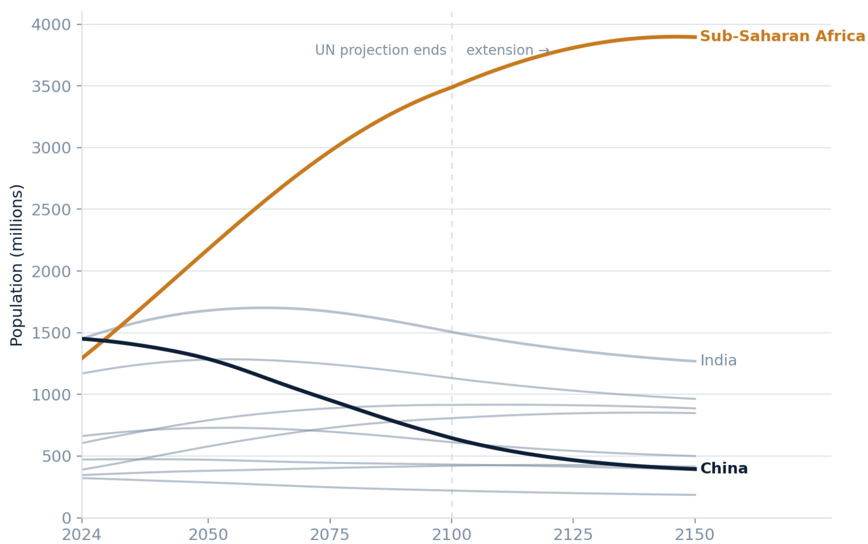


***Figure 1.*** *Regional population, 2024 to 2150. Everything right of the dashed line is produced by the extension.*

**Turnover outruns growth.** In 2030, 9.0% of living people were born after 2024; the majority crossover falls in 2064, and by 2100 the share is 87.3% (Figure 2). By 2125 nobody alive was born before 2024. The crossover date varies by nearly two decades across regions, from 2048 in Sub-Saharan Africa to beyond 2100 in China. That spread is itself a finding: a single global figure for how much of the future is already alive would misdescribe both ends of the distribution, understating turnover where fertility is high and overstating it where populations are ageing and contracting. Across the whole horizon, 1.25 trillion person-years are lived between 2024 and 2150, of which 852 billion, or 68.3%, belong to people not yet born, ranging from 38.9% in China to 83.3% in Sub-Saharan Africa.

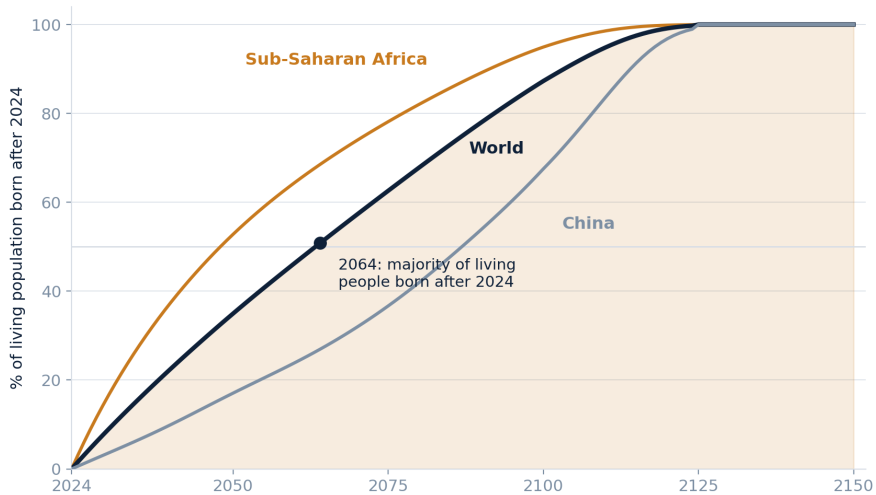


***Figure 2.*** *Share of the living population born after 2024. The shaded area is person-years belonging to people who do not yet exist at the reference year.*

**Discounting removes most of the future.** Applying an exponential factor to those person-years cuts the future share from 68.3% to 57.3% at 1%, 38.0% at 3% and 26.0% at 5% (Table 2). The reduction is uneven: Sub-Saharan Africa loses 40 percentage points of its future share between 0% and 5%, but because regions start from different bases, the ordering of regional weight is itself altered.

***Table 2.*** *Share of 2024 to 2150 person-years belonging to people born after 2024, by discount rate, in percent.*

| **Region** | **0%** | **1%** | **3%** | **5%** |
|---|---:|---:|---:|---:|
| World | 68.3 | 57.3 | 38.0 | 26.0 |
| Sub-Saharan Africa | 83.3 | 75.3 | 57.6 | 43.1 |
| Central Asia | 79.0 | 69.9 | 51.3 | 37.3 |
| Middle East and North Africa | 70.9 | 60.4 | 41.2 | 28.6 |
| United States | 65.5 | 53.8 | 33.5 | 21.5 |
| India | 62.5 | 51.5 | 33.8 | 23.4 |
| Western Pacific and SE Asia | 61.9 | 50.6 | 32.7 | 22.2 |

| Region | 0% | 1% | 3% | 5% |
|---|---|---|---|---|
| Latin America and Caribbean | 59.6 | 48.5 | 31.2 | 21.3 |
| North Atlantic | 59.1 | 47.0 | 28.0 | 17.6 |
| Central and Eastern Europe | 55.6 | 43.5 | 25.6 | 16.2 |
| China | 38.9 | 29.1 | 16.7 | 10.8 |

**Tobacco deaths peak long after prevalence falls.** Global smoking prevalence declines throughout, from 24.7% in 1990 to 17.6% in 2022 and 7.6% in 2100. Annual smoking-attributable deaths nonetheless keep rising for three more decades, peaking at 9.16 million in 2058 before falling to 6.99 million by 2100 (Figure 3), with 639.0 million deaths projected after 2024. The divergence is demographic: falling prevalence acts on a population that is growing and ageing, and deaths in a year are generated by smokers in the previous one who may have started forty years earlier. Regional peaks are staggered accordingly, from India in 2047 and China in 2057 to Sub-Saharan Africa and MENA, both still rising at 2100.

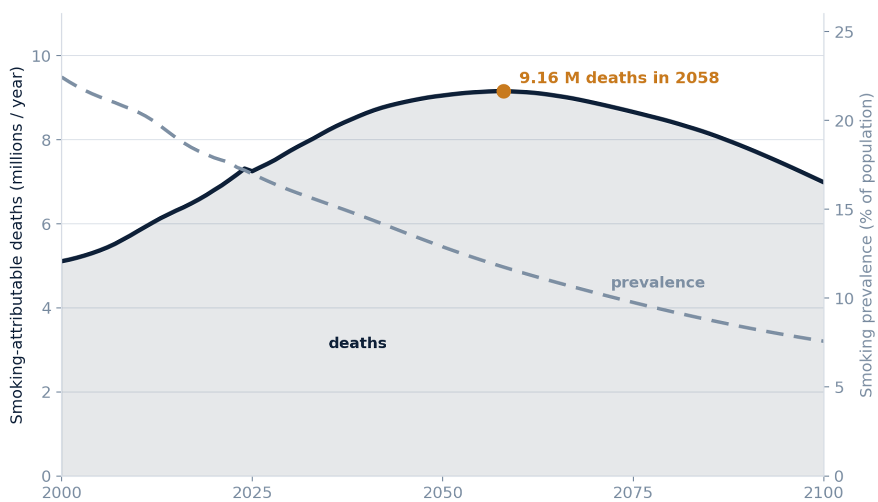


***Figure 3.*** *Global smoking-attributable deaths against prevalence, 2000 to 2100. Prevalence falls throughout; deaths rise for a further 34 years.*

**Discounting reallocates the burden.** Of the 639 million deaths projected for 2025 to 2100, 64.4 million, or 10.1%, fall on people not yet born in 2024, a share varying sixfold across regions: 31.2% in Sub-Saharan Africa and 21.1% in MENA against 4.9% in China. Discounting to 2024 reduces the total to 447 million at 1% and 250 million at 3%, and shifts its distribution (Figure 4). At 3%, Sub-Saharan Africa's share of global burden falls from 4.8% to 3.8% and MENA's from 6.9% to 5.6%, while the North Atlantic rises from 4.9% to 6.0%. Because epidemic timing correlates with region, a purely temporal operation produces a geographic redistribution.

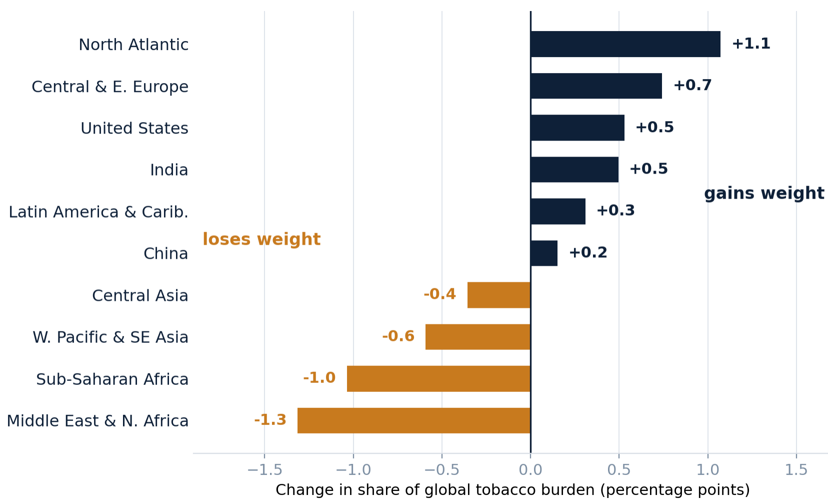


***Figure 4.*** *Change in each region's share of global smoking-attributable mortality when deaths projected for 2025 to 2100 are discounted to 2024 at 3%, in percentage points.*

## IV. Discussion

The identity of a population changes far faster than its size. World population in 2150 sits within 20% of its 2024 value, yet nobody alive in 2150 is alive today. Any evaluation that treats the world population as a continuous object across a century is describing a different set of people at each end without registering the substitution.

This matters because a single discount rate is doing two jobs at once, compounding a time preference, a probability that the benefit fails to arrive, and an implicit judgement about people who do not yet exist. Separating the existence weight from the discount factor shows the two are not substitutes: the weight scales the future cohort uniformly, while the factor penalises it as a function of delay, so a person born in 2030 and one born in 2120 are treated identically by the first and differ by four orders of magnitude at 3% under the second. Which is intended is rarely stated, and the choice is consequential. At 3% the future share of person-years falls from roughly two-thirds to roughly one-third, so a rate chosen on financial opportunity-cost grounds delivers, incidentally, something close to a 50% discount on unborn people as a class.

The tobacco case sharpens this twice over. The 34-year lag between falling prevalence and peaking mortality means cessation policy judged on near-term prevalence understates its own benefit, because the deaths avoided accrue to cohorts that have not started smoking. More troublingly, the regional reallocation runs in the direction of existing inequity: Sub-Saharan Africa and MENA are the two epidemics still growing, carry the largest share of deaths among the unborn, and lose the most weight under discounting. A funder allocating on discounted burden would move resources away from the epidemics yet to peak and toward those already in decline, a distributional consequence of a parameter defended on grounds of intertemporal consistency alone.

The two views this framework is built to distinguish give sharply different answers to the same policy question. Under a strict person-affecting view, a tobacco policy is evaluated almost entirely on its effect on people who exist when it is enacted, which for the 2025 to 2100 window means setting aside a tenth of the deaths outright and, in Sub-Saharan Africa, nearly a third. Under a totalist view the unborn count fully and the ranking of regional priorities inverts. Neither position is settled in the philosophical literature, but the size of the gap between them can now be measured rather than asserted, which is the contribution this study makes to a debate that has largely proceeded without quantities attached.

It is worth being explicit that the demographic and distributional results do not depend on α at all. The extension to 2150, the cohort decomposition, the person-year shares in Table 2 and the regional reallocation in Figure 4 are statements about who is alive in which year and how an exponential factor weights them; they are computed without any assumption about how much unborn people morally count. A reader who rejects the weighting framework entirely still inherits the finding that 68.3% of the person-years at stake belong to people not yet born, that conventional discount rates remove roughly half of them, and that doing so shifts measured tobacco burden away from the regions whose epidemics are still growing. α determines what one should conclude from those facts, not what the facts are, and the framework is usable as a descriptive tool before any position on population ethics is taken.

The framework leaves α free deliberately. Its value is an empirical question about what people actually believe, and a companion survey experiment now in pre-registration is designed to elicit it directly, presenting respondents with trade-offs between existing and potential lives and testing whether stated principles survive reflection. The demographic work here supplies the other half: a defensible count of how many potential people any given α is applied to.

## V. Limitations

Migration is not modelled after 2100, and the pre-computed age shares were generated at the default TFR, so scenarios far from it will misstate age structure, though mortality dominates pyramid shape over fifty years. The fertility-to-growth coefficient is a calibrated approximation rather than a fitted parameter. On the tobacco side the constraints are firmer: deaths for 2000 to 2022 are modelled estimates rather than observed data, pending a direct pull from the GBD Risk Factors tool, and while the age and maturity components of the fatality schedule follow the cited cohort studies, the sex-ratio multipliers are directional estimates rather than figures verified against a published table, and are the input most in need of validation. World totals exclude roughly 29 small states, and 2023 population is approximated by the 2024 figure. None of these affect the structural findings, which turn on cohort turnover rather than on the level of any particular rate.

A further limitation is that the projection holds the policy and product environment fixed. Prevalence is carried forward on each region's observed 2012 to 2022 trend, which embeds whatever tobacco control was in force during that decade and assumes it neither intensifies nor lapses. It therefore contains no representation of e-cigarettes, heated tobacco or nicotine pouches, whose long-run health consequences and substitution patterns are still uncertain, nor of tax rises, advertising restrictions, plain packaging or generational sales bans of the kind recently debated in the United Kingdom. Any of these could move prevalence substantially within a decade, and product substitution in particular could break the link between the current-smoking measure used here and the fatality schedule calibrated to combustible cigarettes. The projections should therefore be read as the trajectory implied by continuation of recent trends rather than as a forecast robust to policy change.

**VI. Conclusion**

Extending the UN projections to 2150 and decomposing them by birth cohort shows that 68.3% of the person-years lived between now and 2150 belong to people who do not yet exist, and that discounting at conventional rates removes roughly half of them from view. Applied to tobacco, it projects 639 million deaths after 2024, a tenth of them among the unborn, and shows discounting shifting burden away from the regions whose epidemics are still growing. Making the weight on future people explicit does not resolve the ethical question, but it reveals how much a modelling convention is deciding on its own.

**Acknowledgements**

I am grateful to Professor Julian Jamison for his supervision throughout this project, to Nuffield College and the University of Oxford for hosting the research, and to the Laidlaw Foundation and the Laidlaw Scholars Programme at MIT for making it possible.

**References**

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

Greaves, H., and MacAskill, W. (2021). The case for strong longtermism. *Global Priorities Institute Working Paper*, University of Oxford.

Institute for Health Metrics and Evaluation (2020). *Global Burden of Disease Study 2019 (GBD 2019) Smoking Tobacco Use Prevalence 1990–2019*. Seattle, WA: IHME. Retrieved from https://ghdx.healthdata.org/record/ihme-data/gbd-2019-smoking-tobacco-use-prevalence-1990-2019

Jha, P., Ramasundarahettige, C., Landsman, V., et al. (2013). 21st-century hazards of smoking and benefits of cessation in the United States. *New England Journal of Medicine*, 368(4), 341–350.

Nordhaus, W. D. (2007). A review of the *Stern Review on the Economics of Climate Change*. *Journal of Economic Literature*, 45(3), 686–702.

Parfit, D. (1984). *Reasons and Persons*. Oxford University Press.

Pirie, K., Peto, R., Reeves, G. K., Green, J., and Beral, V. (2013). The 21st century hazards of smoking and benefits of stopping: a prospective study of one million women in the UK. *The Lancet*, 381(9861), 133–141.

Ramsey, F. P. (1928). A mathematical theory of saving. *The Economic Journal*, 38(152), 543–559.

Reitsma, M. B., Kendrick, P. J., Ababneh, E., et al. (2021). Spatial, temporal, and demographic patterns in prevalence of smoking tobacco use and initiation among young people in 204 countries and territories, 1990–2019. *The Lancet*, 397(10292), 2337–2360.

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

United Nations, Department of Economic and Social Affairs, Population Division (2024). *World Population Prospects 2024: Summary of Results*. New York: United Nations. Retrieved from https://population.un.org/wpp/

Vollset, S. E., Goren, E., Yuan, C. W., et al. (2020). Fertility, mortality, migration, and population scenarios for 195 countries and territories from 2017 to 2100. *The Lancet*, 396(10258), 1285–1306.

Wittgenstein Centre for Demography and Global Human Capital. Human capital-based population projections.