THB Build your own climate scenario

Methodology: what the tool calculates, where every number comes from, and what it does not represent

6	**7**	**6**	**75**
assumptions	CMIP7 scenarios	Learn More pages	years, 2025 to 2100

Analysis by Roger Pielke Jr., The Honest Broker. 9 September 2026.

This is work in progress — caveat lector.

The tool this documents runs at scenarios.thehonestbroker.org. Every figure it reports is recomputed from the sources named in Part 3 rather than carried in prose, and every rate is rebuilt by a script that reads the primary series.

Contents

Part 1

The site

An overview of what Build your own climate scenario does, page by page, and how to read what it puts on the screen. The two parts that follow this one give the detail: **Methods** covers every calculation and every preset, **Data** covers where each number came from, when, and in what units.

What the tool does

The builder turns eight assumptions into one emissions path from 2025 to 2100, draws that path against the seven CMIP7 marker scenarios, and reports what the path adds up to.

Nothing about the exercise depends on agreeing with any of the eight numbers. The point is the arithmetic that connects them: a reader who thinks the world will decarbonise faster than the record can set that rate and watch what it does to 2100, on the same axis and against the same seven published scenarios as everybody else.

The eight assumptions

Four of them multiply together in the Kaya identity, and four sit outside it:

Slider	Sets	Range	Today
People in 2100	A level	6 to 14 billion	8.14 billion
Income growth per person	A rate that compounds	−0.5 to +3.5 %/yr	$21,393 in 2024
Energy needed per dollar	A rate that compounds	−4 to +0.5 %/yr	3.40 MJ per dollar
Carbon per unit of energy	A rate that compounds	−4 to +0.5 %/yr	65.18 $kgCO_2$ per GJ
Land use CO_2 in 2100	A level	−10 to +5 $GtCO_2$/yr	4.59 $GtCO_2$/yr
Methane in 2100	A level	80 to 600 Mt/yr	380 Mt/yr
Improvement delivered by 2062	A share of the century's technology improvement	5 to 95 %	50 % is a steady rate
Engineered CO_2 removal in 2100	A level	0 to 25 $GtCO_2$/yr	close to zero today

The first four decide fossil and industrial CO_2, cement included. Land use CO_2 adds on top as a separate term, running in a straight line from today to whatever the reader sets. Methane changes nothing about the CO_2 path and enters only through the warming figure.

The last two exist because the four Kaya factors fix where a path ends and say nothing about the route, and because multiplying positive factors never reaches a negative number. Timing redistributes the two technology rates across the century while holding their 2100 level exactly. Removal is the additive term that carries a path below zero, ramping as the square of elapsed time rather than as a line, which is both how the scenarios deploy it and what keeps it from duplicating the land use slider. Part 2 sets out both.

The base year is 2025, carrying the observed 2024 world, because 2025 has not closed in every source.

What the chart shows

One heavy line for the reader's scenario, seven lighter lines for the CMIP7 markers, and annual CO_2 including land use on the vertical axis.

That axis moves. Every render picks an axis that covers the reader's path and all seven markers together, on round steps, with zero always on it. A scenario far above the published range stays fully drawn instead of clipping, and the seven markers never leave the axis the reader's own line sits on.

Whatever the reader types into **Name your scenario** names the line on the chart, the column in the table, and every file the page hands back.

The four figures under the chart

- **Cumulative CO_2, 2025 to 2100.** Every year of the path, summed, in gigatonnes.
- **Warming in 2100 above 1850–1900.** A curve fitted to FaIR v2.2 runs of the seven markers, reading cumulative CO_2 and 2100 methane and nothing else. Indicative, not a model result, and never shown to more than two decimals.
- **Added warming from now.** The same figure measured against the 2015–2024 average of 1.24 °C rather than against 1850–1900.
- **Your 2100 world looks like.** The largest economy whose 2024 CO_2 per dollar of GDP sits within 5% of the reader's 2100 world, drawn from 66 economies large enough to measure well. Largest, rather than nearest, so the comparison lands on an economy the reader recognises instead of whichever small one happens to sit

on the number. Where nothing sits within 5% the tile says so instead of naming an economy, and names the one at the edge of the table instead: a world at 0.011 kg CO_2 per dollar resembles nothing on earth today. It compares that one ratio and says nothing else about the country.

Load one of the four CMIP7 presets and the first three figures report what that scenario published, rather than the reconstruction, until a slider moves.

The eight preset buttons

Four presets come from the record and from the two technology bounds; four load the Kaya factors a CMIP7 marker reports. Those four are the markers of the seven that publish a carbon-intensity rate.

Preset	Cumulative CO_2	Warming in 2100
Kaya at observed rates	4,207 Gt	3.22 °C
Trend continues	3,437 Gt	2.95 °C
Slowest technical progress	5,397 Gt	3.53 °C
Ausubel methane economy	1,614 Gt	2.23 °C

Every rate behind those four comes out of `scripts/build_data.py`, recomputed from the primary series rather than typed in, and the window that produced each one travels with it.

The four CMIP7 presets reproduce where a marker ends up far better than how it gets there, and three of them cannot reach the scenario they name at all. The interface reports the gap in place, computed on the spot, whenever one of them loads. **Methods** gives the three reasons — the shape of the path, the sign of the emissions, and the timing of the land-use sink — and the size of each miss. MEDIUM-to-LOW is the one that lands below its marker rather than above.

The Learn More pages

Eight pages, one per assumption. Each teaches the quantity, shows what the world has done, and hands back a value the reader built rather than guessed:

- **Population** adds up seven regional 2100 populations, each running between the UN's low and high variants.
- **Income per person** works in any of three directions: set a rate, set a 2100 level, or build the world average out of the three World Bank income groups.

- **Energy per dollar** takes a compound rate from any two years of the record, or a multiple of the observed rate, and says where the answer falls among the 35 twenty-five-year windows the world has actually run.
- **CO_2 per unit of energy** builds a 2100 fuel mix out of seven shares, burns it against the IPCC emission factors, adds the industrial CO_2 that no fuel switch touches, and takes the rate that lands there.
- **Land use** nets three flows — deforestation with other transitions and peat, existing regrowth, and newly restored area at a sequestration rate. Every land-based removal sits inside this net, which is why the page carries no engineered-removal control.
- **Methane** sets one control per source, scaled so today's five sources reproduce the 380 Mt the slider starts from.
- **When the improvement arrives** asks the record the question the slider asks, and the two technology factors answer it in opposite directions: energy per dollar banked 26% of its improvement by the midpoint, the fuel mix 79%.
- **Engineered CO_2 removal** sets what capture and storage takes back in 2100, and reports the sustained growth from today's 0.002 $GtCO_2$ a year that the level implies. Forests and soils belong to the land use page, so no tonne is set in both places.

A page shows the arithmetic at every step, not just the answer. **Use in my scenario** carries the value back to the builder, which names it, scrolls to that slider and lights it; **Back to my scenario** returns the eight numbers untouched.

What travels with a scenario

Eight numbers and a name, encoded in the address. Every link off every page carries them, so a reader who steps out to a Learn More page, the library or the bibliography comes back to the scenario they left. **Copy link to this scenario** hands over that same address.

The builder downloads a PNG and a PDF of the whole scenario sheet. Every figure on every page downloads two ways: as a PNG carrying the figure's title, the drawing and the numbers under it, and as a spreadsheet of those numbers. Where a figure rests on more rows than an image can hold, the PNG prints an evenly spaced sample and says so; the spreadsheet always carries every row.

What every figure carries

The Honest Broker mark, the source of the numbers, and the analysis credit, in that order, under every chart and every table on the site. The downloaded PNG draws the same three underneath the figure and the spreadsheet writes the last two under the table, because a figure that travels has to answer for itself.

The library and the bibliography

The **library** collects the posts at The Honest Broker that make the argument this tool grew out of, grouped by what each one argues. The **bibliography** lists every work the site draws on, each with a link and a plain-language note on what it supplied, and it generates the Learn More section of that list from the pages themselves, so a source cited on a page cannot go missing from the list.

What the tool does not do

Stated in full in **Methods**, and worth having up front:

- Nothing engineered removal does can come out of the four Kaya factors, which stay positive whenever there are people, income and energy. It is an additive term on a slider of its own, kept apart from the land use sink so that no tonne is subtracted twice, and it ramps on one fixed shape that no control bends.
- No nitrous oxide, no fluorinated gases, no aerosols as a term of their own.
- No carbon cycle feedbacks that depend on the path rather than the total.
- One warming number, not a range. The FaIR ensemble behind the fit spans well over a degree at any given cumulative total.
- Only 2100, not the path of warming to it or anything after it.
- No check that eight assumptions belong together. A handful of combinations no marker contains get named as an observation, not blocked.

How to check any of it

Every number the interface shows comes from a JSON file under `src/data/`, and nothing gets typed into the interface code. **Data** names the script that writes each file, the source it reads, and the vintage of that source. The scripts that can rebuild from a primary source say so; the few series carried over from the prototype say that instead, and name what would verify them.

Part 2

Methods

What the tool calculates, how the four presets were derived, and what it does not represent. Part 3 covers where the numbers come from.

The emissions calculation

Emissions follow the Kaya identity. Four factors multiply together, and land use CO_2 is added on top as a separate term:

```
CO₂ = population × (GDP ÷ population) × (energy ÷ GDP) × (CO₂ ÷ energy)
      + land use CO₂
```

The reader sets eight numbers: the 2100 population, three annual rates that compound from the base year, 2100 values for land use CO_2 and methane, when across the century the technology improvement arrives, and how much engineered CO_2 removal runs in 2100.

The last two joined the six because the first six fix where a trajectory ends and say nothing about the route there, and because a product of positive factors never reaches a negative number. Both gaps showed in the CMIP7 presets: MEDIUM landed within 1% of its marker's 2100 emissions and 12% above its century total, and the two deep-mitigation markers end net-negative where nothing built from the four Kaya terms can follow. See "Timing and removal" below.

Counting the stops each slider offers and multiplying them, the eight reach **2,412,174,456,707,806,791** settings. That is not the same as the number of scenarios. Energy per dollar and CO_2 per unit of energy enter the identity only through their product, so swapping one rate for the other leaves the path byte-identical, and about half of all settings repeat another. Methane changes no CO_2 point at all, but it does change the warming, so it counts as part of an outcome.

Counting the distinct products the two rates reach rather than the pairs they offer gives

```
population 81 x income 401 x technology rates 101,341 x land use 151
  x methane 521 x timing 91 x removal 51
```

which is **1,201,823,843,625,281,331** distinct scenarios, just over 1.2 quintillion. That is the figure the front page states, and the parenthetical beside it writes out the same seven factors.

One further collapse sits inside that and is not deducted. At a timing of exactly 50% income compounds over the same clock as the two technology rates, so all three collapse into a single product rather than two. Counting that one slice separately gives 1,200,628,952,606,294,829, which is 0.1% lower and rounds to the same figure. The page states the seven-factor product because a reader can multiply it back; `rateProducts.triples` in `src/data/config.json` carries the 36,961,039 three-rate products for anyone who wants the exact form.

`SCENARIO_COUNT` and `DISTINCT_SCENARIO_COUNT` in `src/model/config.ts` derive both from the slider definitions rather than carrying either as a literal, the front page prints what it derives, and `tests/config.test.ts` checks the figures in this paragraph still match. The distinct-product counts behind them come from `scripts/build_carried_data.py`, because the three-rate case is 81.5 million products and belongs to the build rather than to page load. Both count in BigInt: the settings product outgrew exact double arithmetic when timing and removal joined, coming out as ...806700 for a true ...806791, losing whole scenarios in the last digits without saying so. Both cover what the sliders themselves address: a hand-edited `?s=` link can carry a value between two stops, which the model clamps to range but does not snap, so links address a denser set again.

Far fewer answers than scenarios come out the other end. Across the whole space the century total spans -720 to 44,027 $GtCO_2$ and warming 1.37 to 6.11 °C, which at the precision the tiles print leaves at most 44,748 distinct totals and 475 distinct warming figures. Several factors multiplied together, and only the product shows.

The base year is 2025, holding the observed 2024 state, since 2025 is not complete in every source. The path runs one year at a time to 2100 and the cumulative total sums every year.

What the four factors cover

They cover **all fossil and industrial CO_2, cement included**. This matters, because the whole chart is the reader's line drawn against the seven CMIP7 marker lines, and those count all of it.

The prototype set base CO_2 per unit of energy from the Energy Institute's energy CO_2 plus flaring, at 60.5 kg per GJ. That leaves out cement and other industrial process emissions, roughly 1.9 $GtCO_2$ a year, and it started every path at 38.70 $GtCO_2$ in 2025 while the markers started between 42.4 and 43.8. On 2026-09-04 the base was recalibrated to 65.181 kg per GJ, taken from Global Carbon Budget fossil and industry CO_2 over Energy Institute total energy supply. Paths now start at 42.47, inside the marker range.

This changed results. Kaya at observed rates moved from 3,707 $GtCO_2$ to 4,084, which puts it above CMIP7 HIGH's 3,777 rather than below it. `presets.json` keeps the figures stated before the recalibration under `brief_stated`, and a test pins the size of the shift.

Land use followed on 2026-09-09. The base kept the brief's 3.83 $GtCO_2$ while the same block of `base.json` counted the Global Carbon Budget's 4.586 for the marker-comparable total, so the path started 0.76 $GtCO_2$ below the basis it was drawn against and the land use chart showed a step where the record met the reader's line. Land use now comes from the same source as the rest of the base year, and the 2025 start is 43.23 against a marker range of 42.4 to 43.8.

That improved three of the four CMIP7 presets against their published cumulative totals and made one worse: HIGH from −12 to +4 $GtCO_2$, MEDIUM from −3 to −2, MEDIUM-to-LOW from 0 to +2, and VERY LOW from +55 to +68. VERY LOW pays for it because its own land use is a sink of −4.71, so raising the base raises the whole early ramp against it. Every documented preset moved by about 29 $GtCO_2$ and by 0.01 °C or less. `presets.json` carries each earlier freeze under `superseded`, newest first, so a figure published against any of them stays traceable.

Population

Population interpolates between the real IIASA SSP1, SSP2 and SSP3 world trajectories on their 2100 endpoints, so the curve keeps a demographic shape: SSP1 peaks at 9.29 billion around 2060 and declines, SSP3 rises throughout. A straight line or a smoothstep between two endpoints could not do that.

Outside the SSP1 to SSP3 span, which is 8.09 to 12.98 billion in 2100, the nearest curve is scaled by the ratio of the target to its own endpoint. That keeps the shape and moves the level.

Scaling used to move the base year, and no longer does. Multiplying a curve multiplies its first point along with the rest, so at 6 billion the path began from 6.04 and at 14 billion from 8.79, against an observed 8.15. Both ends of the 6 to 14 billion slider reached it, and the reader saw the whole line lift or drop when the slider moved. 2025 is an observation and every other slider held it exactly, so this one now does too: the scaled curve carries an additive correction, at its full size in the base year and at nothing in 2100, which pins the start without moving the target. All three SSP curves share a base year of 8.15 billion, so inside the span there is nothing to correct and no path changes. `tests/kaya.test.ts` checks the base year against both ends and the default of all eight sliders, and `tests/population.test.ts` checks that the 2100 figure the slider asks for is still the figure it reaches.

Land use

Land use CO_2 runs in a straight line from 4.59 $GtCO_2$ today to whatever the reader sets for 2100. That figure is the Global Carbon Budget's land-use change CO_2 for the base year, the same source and vintage as the fossil and industrial CO_2 beside it.

A line is a placeholder. The marker paths carry land use folded into their total CO_2 and publish no separate land use series here, so there is no shape to borrow from them. If those series become available, the shape should follow them.

The two technology bounds

Both bound presets hold every technological trajectory at an extreme of the record. Two of the three rates come from window scans in `observed.json`, and one comes from a published paper.

Slowest technical progress takes the weakest 25-year window on energy per dollar, −1.36%/yr over 1990 to 2015, and the weakest 30-year window on CO_2 per unit of energy. On 2026-09-09 that second rate moved from −0.17%/yr over 1992 to 2022 to **−0.02%/yr over 1984 to 2014**. The first is the weakest window on the combustion series, CO_2 from energy over total energy supply; the second is the weakest window on the basis the slider itself moves, cement and industrial CO_2 included. The bound sets that slider, so it now reads off the quantity it sets. The preset's cumulative total went from 5,047 to 5,397 $GtCO_2$ and its warming from 3.47 to 3.53 °C.

Ausubel methane economy needs no equivalent correction. Its fuel-mix rate, −2.79%/yr, comes from the 1988 paper rather than from a window scan, and its energy-per-dollar rate has no CO_2 in it.

`observed.json` carries both scans, `carbon_per_energy_30y` and `carbon_per_energy_slider_30y`, so the difference stays visible.

Timing and removal

Two controls sit outside the Kaya identity and exist because the identity alone could not follow the scenarios.

Improvement delivered by 2062 redistributes the two technology rates across the century without moving where they end. The rate sliders fix the total change from the base year to 2100; this says how much of it lands by the midpoint. At 50% the annual rate is constant, which is what the model did before this control existed. Above 50% the improvement front-loads and slows later; below 50% it builds.

The weight on the annual rate decays geometrically rather than tilting linearly, so it never changes sign: a technology that improves cannot start worsening late in the century because the reader asked for an early push. Writing the effective years elapsed as

```
accumulated(t) = span * (1 - e^-Lt) / (1 - e^-L*span),  L = (2/span) * ln(s/(1-s))
```

for a midpoint share `s`, the endpoint is exact -- `accumulated(span) = span` for every `L` -- so the 2100 level the reader set on the rate slider stays put whatever the timing.

`accumulatedYears` in `src/model/rates.ts`.

Engineered CO_2 removal in 2100 is the additive term that carries a path below zero. The four Kaya factors multiply, so the fossil term approaches zero without ever crossing it, and no arrangement of the four reaches the net-negative emissions the deep-mitigation markers reach. Removal ramps as the square of elapsed time, close to nothing before the 2040s and accelerating after, which is how the scenarios deploy it.

That shape also earns the control its place. A straight ramp from zero to a 2100 level is arithmetically identical to moving the land use slider by the same amount -- checked, to zero difference across the path -- so a linear removal slider would have been a second way to say one thing and would have reached no scenario the six could not.

Both sit at no effect by default, so every result the model produced before they existed is unchanged.

Where removal is counted

Two of the eight controls can take carbon out of the air, and they must not take the same tonne out twice.

Land use is net CO_2 from land use, land-use change and forestry, on the Global Carbon Budget's definition. That number is already a net: 6.23 $GtCO_2$ a year of gross deforestation and 2.20 of other transitions and peat, less 4.40 of regrowth on land already recovering. The Global Carbon Budget's own net for the base year, 4.59, is what the model starts from; the decomposition above comes from the 2024 paper and nets to 4.03, and Part 3 covers the vintages. Every conventional removal — planting, restoration, soil carbon, durable wood products — lands inside it, which is what lets this one slider reach a sink of 10 $GtCO_2$ a year on its own.

Engineered removal is the removal that stores carbon outside that account: bioenergy with carbon capture, direct air capture, biochar, enhanced weathering. Bioenergy with capture competes for the same hectares, but the tonnes it stores go underground rather than into standing biomass, so the land-use flux does not measure them and adding them here counts them once.

The learn pages hold the same line. The land use builder nets three flows and carries no removal control; the removal builder has one control and carries no forest control; each page's prose points at the other.

One consequence for the CMIP7 ticks. `derive_timing` in `scripts/build_carried_data.py` fits each marker's removal **with that marker's own land use already in place**, so the fitted figure is what the published path needs beyond the land, not the marker's whole removal. MEDIUM-to-LOW assumes a land sink of 8.79 $GtCO_2$ a year and needs 5.5 on top of it; VERY LOW assumes 4.71 and needs 11.0. Reading either as that scenario's total removal would count its forests twice. HIGH's 1.5 is not removal at all: a constant rate overshoots HIGH's published path, ending 2100 at 56.0 against the marker's 55.0, and the fit takes the difference off on this slider for want of anywhere else.

Warming

Warming in 2100 above 1850 to 1900:

```
T = 1.7274 + 1.3100 × ln(1 + max(0, C) ÷ 1989.74) + 0.0012 × (CH4 − 380)
```

`C` is cumulative CO_2 from 2025 to 2100 in gigatonnes and `CH4` is methane in 2100 in megatonnes a year. Methane adds about 0.12 °C per 100 Mt a year.

This is a curve fitted to FaIR v2.2 runs of the seven CMIP7 markers, using the fair-calibrate v1.4.1 constrained ensemble. **It is indicative, not a model result**, and the interface says so. Warming is never shown to more than two decimal places.

How well it reproduces FaIR

Scenario	Cumulative CO_2 (Gt)	Methane (Mt/yr)	Emulator	FaIR	Error
VERY LOW	268	94.0	1.55	1.61	−0.06
LOW-to-NEGATIVE	373	216.0	1.76	1.68	+0.08
LOW	605	150.8	1.80	1.83	−0.03
MEDIUM-to-LOW	1,710	203.4	2.33	2.20	+0.13
HIGH-to-LOW	2,524	132.8	2.50	2.76	**−0.26**
MEDIUM	2,770	436.2	2.94	2.84	+0.10
HIGH	3,777	533.1	3.31	3.26	+0.05

Root mean square error 0.12 °C. Six of the seven fall within 0.15 °C. HIGH-to-LOW does not.

What the emulator cannot see

It reads cumulative CO_2 and methane, and nothing else. Two paths reaching the same total give the same answer however differently they got there. HIGH-to-LOW is where that bites: it climbs to 51.6 $GtCO_2$ a year in 2065, falls steeply after that, and carries the second-lowest methane of the seven. A fit with no term for the shape of the path has no way to tell it apart from a scenario that arrives at 2,524 Gt smoothly, and it lands 0.26 °C low.

Refitting for one structural outlier would move numbers already in print for the sake of a case the functional form cannot represent. The miss is documented and pinned by a test instead.

The four presets

Each loads all six inputs. The two bounds follow Ausubel's 1995 argument that technological trajectories move at rates steady enough to bound the future, and that a scenario halting them describes technical regression rather than business as usual. **Neither bound assumes poverty. Neither assumes a technology stops working.**

Preset	Population	Income	Energy/$	CO_2/energy	Land use	Methane	Cumulative	Warming
Kaya at observed rates	10.2	+1.91	−1.43	−0.15	2.0	380	4,207	3.22
Trend continues	10.2	+1.91	−1.62	−0.48	1.0	300	3,437	2.95
Slowest technical progress	11.4	+2.23	−1.36	−0.02	2.0	450	5,397	3.53
Ausubel methane economy	9.0	+1.91	−1.57	−2.79	−1.0	150	1,614	2.23

Kaya at observed rates. The 1990 to 2024 record projected forward unchanged. Income +1.91, energy per dollar −1.43 and CO_2 per unit of energy −0.15 all come straight out of the Energy Institute, World Bank and Global Carbon Budget series. The carbon-intensity figure was −0.21 until 2026-09-05, on a narrower basis; "The corrected carbon-intensity rate" below gives the change and what moved with it.

Trend continues. The 2015 to 2024 rates, faster than the long record on both technology terms: −1.62 and −0.48. Income stays at the long-record 1.91 rather than the recent decade's 2.02.

Slowest technical progress. Every technological trajectory held at the slowest sustained rate on record. Energy per dollar at −1.36 %/yr, the weakest 25-year window since 1990, which is 1990 to 2015. The fuel mix at −0.17 %/yr, the weakest 30-year window since 1965, which is 1992 to 2022. Income at +2.23, the fastest observed 25-year rate, 1994 to 2019, and population at 11.4 billion, the top of the UN 2024 95% range, so emissions run as high as slow technology permits.

Ausubel methane economy. Jesse Ausubel's 1988 published trajectory, which squeezes carbon out of primary energy to 0.06 tonnes of carbon per kilowatt-year by 2100, implying the fuel mix improving 2.79 %/yr. Energy per dollar at −1.57, its fastest observed 25-year rate, 1996 to 2021, and income still growing at the historical pace.

Every rate above is recomputed by `scripts/build_data.py` from the primary sources rather than typed in, and the window that produced each one is recorded alongside it.

The four CMIP7 presets

Four further buttons load the Kaya factors that a marker itself reports. These are the four markers of the seven that publish a carbon-intensity rate; the other three do not, because their 2100 fossil CO_2 runs negative. HIGH lands close to the scenario it names. The other three cannot:

Preset	Marker total	Preset reproduces	Gap
CMIP7 HIGH	3,777 Gt	3,838 Gt	+61
CMIP7 MEDIUM	2,770 Gt	3,094 Gt	+324
CMIP7 MEDIUM-to-LOW	1,710 Gt	1,230 Gt	−480
CMIP7 VERY LOW	268 Gt	1,298 Gt	+1,030

MEDIUM-to-LOW is the one that lands below the marker it names rather than above. The reason is in the next section. The interface says so in place when one of these is loaded rather than leaving the reader to notice.

The Learn More pages

Each of the eight assumptions gets a page that teaches the quantity, shows the record, and hands the reader a value built from assumptions they control. The pages share one scaffold (`src/ui/learn/page.ts`), one builder (`src/ui/learn/builder.ts`) and one plot component (`src/ui/plot.ts`), and each page module supplies only its own words, chart and arithmetic.

What every figure carries

The mark, the source of its numbers and the analysis credit, in that order, under every chart and every table on every page. The downloaded PNG draws the same three in a band beneath the figure, and the downloaded workbook writes the source and the credit under the table, because a figure that travels has to answer for itself.

The hand-off

A link on the top page carries the reader's scenario to a page as `/learn/<slug>/?s=<numbers>&n=<name>`, the same encoding share links use. The page holds that string untouched. Two exits lead back:

- **Back to my scenario** returns `/#s=<the identical string>`.

- **Use in my scenario** returns `/?applied=<field>#s=<state with one field replaced>`. The top page names the value, scrolls that slider into view, lights it for 1.5 seconds, and then rewrites the address bar to a clean `/#s=…` so a copied link carries no stale message.

A builder value passes through `roundToStep` and then `clampWithFlag` (`src/model/config.ts`), and the page prints the rounded value, and the clamp when one bites, before the reader commits. No page duplicates any of that arithmetic.

The population builder

Seven regional 2100 populations, added up:

```
world 2100 = Σ region 2100
```

Each control opens at the UN's medium variant and runs from the UN's low variant to its high variant. The UN builds those two by subtracting and adding half a child per woman at every date, and they add up across regions, so the sum reproduces the UN's own world figures exactly: 6.987, 10.180 and 14.395 billion. `scripts/build_wpp.py` asserts that at build time.

Three consequences worth stating, all of them visible on the page:

1. The builder cannot reach the slider's floor of 6 billion. Every region at the UN's low variant still gives 6.99 billion.
2. It can exceed the slider's ceiling of 14 billion, which reports 14.395 and clamps to 14.
3. The 95% prediction intervals do not add up this way, and the page says so rather than summing them: the regional lower bounds give 8.234 billion against the UN's world lower bound of 9.047.

The reader's curve on that page comes from `populationAt` in `src/model/population.js`, the same function the scenario itself uses, so the shape on the Learn More page and the shape behind the top page agree.

The energy per dollar builder

Two modes, both producing a rate in %/yr.

A window from the record. The reader picks two years and the builder takes the compound rate between the observed intensities of those years:

```
rate = ((intensity[end] / intensity[start]) ^ (1 / (end - start)) - 1) x 100
```

The second year is held at least ten years after the first, because a shorter window measures a business cycle rather than a trend.

A multiple of the observed rate. The rate observed from 1990 to 2024, −1.4277 %/yr, times a multiple the reader sets between 0 and 3.

Both modes report where the answer falls among the 35 windows of 25 years in the record, as a count and as a percentile, and what today's 3.40 MJ per dollar becomes by 2100 at that rate. The percentile counts windows that improved more slowly than the chosen rate, so a faster rate scores higher.

The series behind all of it splices two sources at 1990; Part 3 gives the method and the size of the difference between them.

The CO_2 per unit of energy builder

Seven fuel shares for 2100 plus one control for industrial CO_2, in four steps:

1. **Normalise.** The seven shares are scaled to sum to 100, so only their proportions matter.
2. **Burn.** Each share meets its emission factor, and the sum is scaled by 0.9201, the calibration that makes the 2024 mix reproduce the 2024 observed intensity. Part 3 explains where that factor comes from.
3. **Add what does not burn.** 5.237 kgCO_2 per GJ of cement, flaring and other industrial CO_2, times whatever the last control is set to.
4. **Take the rate.** From 65.181 kgCO_2 per GJ today to that number in 2100:

```
rate = ((intensity[2100] / 65.181) ^ (1 / 75) - 1) x 100
```

The page shows every intermediate number, because the interesting part of this builder is the arithmetic rather than the answer. A mix with no fossil fuel in it still leaves step 3, so the fastest rate reachable with industrial CO_2 unchanged is −3.31%/yr.

The income per person builder

Three modes.

Set a rate. The level in 2100 is `21,393.2 x (1 + rate/100) ^ 75`.

Set a 2100 level. The rate is the inverse: `((level / 21,393.2) ^ (1/75) - 1) x 100`.

By income group. Each group's average compounds at its own rate from its own 2024 level, and the world average is the population-weighted sum, scaled by 0.9967 so the three groups reproduce today's world figure. Population shares stay at today's values, which the page states, along with the direction of the bias: the UN projects the low-income share rising, so holding shares still overstates the world average a little.

The methane builder

One control per source, added up. Each source's 2024 level is scaled by 1.086, the ratio between the tool's 380 Mt base year and EDGAR's 350 Mt inventory total, so today's five sources reproduce the number the slider starts from. Part 3 explains why the two differ and why both sit inside the published range.

The page also reports what the answer is worth in the emulator, at 0.12 °C per 100 Mt against a reference of 380 Mt, and says plainly what that coefficient cannot represent: it comes from a straight-line fit to seven FaIR runs and ignores when the methane is emitted.

The land use builder

Three flows netted into one 2100 figure:

```
net = (deforestation + other transitions and peat) x scale
    - existing regrowth x scale
    - restored area x sequestration rate
```

with today's flows from the Global Carbon Budget: 6.23 $GtCO_2$ of gross deforestation, 2.20 of other transitions and peat, and 4.40 of regrowth. The restoration term is a straight area-times-rate product, in millions of hectares times tonnes of CO_2 a hectare a year, divided by a thousand to reach $GtCO_2$.

Engineered removal is not a fourth term here. It has its own slider, and the boundary between the two is set out under "Where removal is counted" above.

The timing builder

One control: what share of the century's technology improvement lands by 2062. The page measures the same quantity against the record, and the two factors disagree. Energy per dollar improved 0.54% a year over 1965 to 1994 and 1.53% over 1994 to 2024, banking 26% of its improvement by the midpoint: it accelerated. The fuel mix went the other way, improving 0.38% then 0.11% and banking 79% by the midpoint before stalling. `scripts/build_timing.py` derives both from series this site already carries, so the page adds no new download.

The removal builder

One control: what capture and storage takes back in 2100, on the same range as the slider it feeds. `scripts/build_removal.py` parses the State of Carbon Dioxide Removal executive summary rather than restating it, and fails loudly if the wording it matches on changes. Removal runs at 2.2 GtCO2 a year today, 5% of gross emissions, of which capture and storage is 0.002 GtCO2 a year, growing 40% a year against assessed pathways that scale it past 3.5 GtCO2 a year by 2050.

The builder reports the sustained growth from today's 0.002 GtCO2 a year that the level implies, because that is the claim a 2100 level makes. Reaching VERY LOW's

11.0 takes 12.2% a year for seventy-five years; the 3.5 the assessed pathways reach by 2050 takes 34.8% a year for twenty-five, close to the rate novel removal grows at now.

The builder had a second control, for what forests and soils take back, until 2026-09-09. It double counted: the land-use term nets regrowth and restoration against clearing already, so a reader could set the same carbon on both pages and subtract it twice. Forests stay on the land use page, machinery stays here.

What a CMIP7 preset does and does not reproduce

Loading a CMIP7 preset sets the eight sliders to that marker's own properties: the six Kaya factors it reports, plus the timing and removal values derived from its published CO_2 path in `scripts/build_carried_data.py`. The interface reports how closely the reconstruction follows, computed rather than asserted (`markerFidelity` in `src/model/flags.ts`).

Preset	2100 CO_2, this tool	Marker	Cumulative, this tool	Marker	2050, this tool	Marker
CMIP7 HIGH	54.5	55.0	3,781	3,777	48.3	47.1
CMIP7 MEDIUM	34.1	34.4	2,768	2,770	37.0	36.1
CMIP7 MEDIUM-to-LOW	−7.9	−9.2	1,712	1,710	35.0	35.1
CMIP7 VERY LOW	−11.1	−5.8	336	268	6.5	−1.2

Before timing and removal existed the same four presets missed their markers' century totals by −12, +325, −480 and +1,031 $GtCO_2$. Three causes drove that, and two of them are now controls rather than limitations:

Shape, solved. A constant rate spread one improvement evenly across 75 years while the markers bend. MEDIUM landed within 1% of its own 2100 emissions and still accumulated 12% more over the century, because MEDIUM cuts hardest in the 2030s and 2040s: its own implied rate runs −0.77%/yr to 2050 and −0.03%/yr after

1 Its timing value of 61% carries that.

Sign, solved. MEDIUM-to-LOW and VERY LOW remove more CO_2 than they emit before 2100, and four factors multiplied together stay positive. The removal term reaches those totals: 5.5 and 11.0 $GtCO_2$ a year by 2100.

Land use timing, still a line. Land use runs straight from today's source to the 2100 value, so a marker that builds its sink late is banked from 2025. With removal carrying the deep sinks this no longer dominates any preset, but the shape remains a placeholder for want of a published land use series.

What is left is not the reconstruction. VERY LOW keeps the widest absolute gap, 55 $GtCO_2$ on a total of 268, which the logarithmic emulator turns into three hundredths of a degree. The remaining distance between a preset's warming and its marker's published warming is the emulator's own residual, measured in "How well it reproduces FaIR" below: up to 0.26 °C when the emulator is fed the markers' own published totals, with no reconstruction involved at all. No arrangement of sliders reaches past that, which is why the tiles print the published figure beside the reconstruction rather than in place of it.

The chart draws the reconstruction in ink and picks the named marker out of the seven behind it whenever one of those presets is loaded, so the remaining divergence sits in front of the reader rather than in a footnote.

The corrected carbon-intensity rate

Until 2026-09-05 the mark under the CO_2-per-energy slider read −0.21 %/yr, which measures CO_2 from energy over total energy supply. The slider itself moves a wider quantity: the fossil and industrial CO_2 the CMIP7 markers count, cement included. On that basis the world improved **−0.15 %/yr** from 1990 to 2024, and the mark, the observed rate the notes compare against, and the "Kaya at observed rates" preset now all use it.

What moved with it:

	Before	After
Mark under the slider	−0.21 %/yr	−0.15 %/yr
"Kaya at observed rates", cumulative	4,083.9 $GtCO_2$	4,177.7 $GtCO_2$
"Kaya at observed rates", warming	3.19 °C	3.21 °C

`config.json` keeps the old figure under `supersededRates` and the preset keeps its old totals under `expected.superseded`, so anything published against the earlier numbers can be traced. The two technology bounds still use window rates measured on combustion alone, which `observed.json` derives from the Energy Institute series; moving those would change the bounds themselves, and that has not been done.

How far the marker factors can be checked

The seven markers' Kaya factors arrived with the prototype and have no primary source on this machine; verifying them needs the IIASA ScenarioMIP database, which requires an account. Two checks are possible without it, and both were run.

Populations reconcile exactly. HIGH takes SSP3's 12.98 billion, MEDIUM and three others take SSP2's 9.89, HIGH-to-LOW and VERY LOW take SSP1's 8.09 to 8.12. Those match `population.json` to the last digit.

The rates reconcile with each marker's own 2100 total, where a rate can exist. Compounding population, income and energy per dollar from the base year, then solving for the carbon intensity rate that lands on the marker's own 2100 fossil CO_2:

Marker	Published rate	Rate implied by its own 2100 total
HIGH	−0.08 %/yr	−0.10 %/yr
MEDIUM	−0.87 %/yr	−0.85 %/yr
HIGH-to-LOW	−2.78 %/yr	−3.50 %/yr
MEDIUM-to-LOW	−3.06 %/yr	no rate reaches it
LOW, LOW-to-NEGATIVE, VERY LOW	none published	no rate reaches it

HIGH and MEDIUM agree within 0.02 percentage points, which is as much corroboration as an internal check can give. HIGH-to-LOW is 0.72 points out.

The last four rows are the interesting result. Those four markers reach 2100 with fossil and industrial CO_2 **below zero** once their land-use term is separated: −0.4, −3.7, −23.2 and −1.1 $GtCO_2$ a year. A rate applied to a positive quantity cannot reach a negative one, which explains why three of them publish no carbon-intensity rate at all. It is the same structural limit the CMIP7 preset note reports on the front page.

What the tool does not represent

A removal path with a shape of its own. Engineered removal has a slider now, so a scenario that removes more carbon than it emits can be built here: five of the seven markers reach net negative CO_2, VERY LOW from 2050 and HIGH-to-LOW only in 2100, and the low ones rely on it. What the tool does not represent is the route. Removal ramps as the square of elapsed time for every scenario, which is a shape borrowed from how the markers deploy it rather than a shape any marker publishes, and no control bends it.

Gases other than CO_2 and methane. No nitrous oxide, no fluorinated gases.

Aerosols. Sulphate and other aerosols cool, and scenarios that burn less coal lose that cooling as well as the warming. The emulator absorbs whatever average relationship held across the seven markers and cannot separate it.

Carbon cycle feedbacks. Permafrost, weakening land and ocean sinks, and anything else that makes the airborne fraction depend on the path rather than the total. Whatever those contributed in the seven FaIR runs is baked into the fitted curve at their average.

Uncertainty. One number is shown, not a range. The FaIR ensemble the fit came from spans well over a degree at any given cumulative total. The emulator reports its central behaviour.

The timing of warming. Only 2100 is reported, not the path to it or anything after it.

Internal consistency. Nothing checks that the eight assumptions belong together. The coherence flags name a few combinations no marker contains, worded as an observation rather than a block, because an unexamined combination is not an impossible one. Nothing checks that the energy system implied by a given efficiency and fuel mix could be built, financed or fuelled.

The country comparison

The tile names the **largest** economy whose 2024 CO_2 per dollar of GDP sits within 5% of the reader's 2100 world, where largest means GDP, the denominator of the ratio being compared. It uses Energy Institute CO_2 over World Bank purchasing power GDP, restricted to 66 economies above 25 Mt CO_2 and $40 billion, so a small or poorly measured economy cannot become the answer. It compares one ratio and nothing else: it does not say the world would resemble that country in any other respect.

Taking the nearest economy of any size, which the tile did until 2026-09-07, kept returning whichever small economy happened to sit on the number: a 2100 world at 0.126 kg/$ came out as New Zealand, a $256 billion economy, while the Netherlands sat at 0.124 with five times the output. CMIP7 HIGH was one of the scenarios it hit. The 5% band is what keeps the likeness true; a wider one buys larger names at the cost of the claim, since at 10% a world at 0.16 kg/$ would "look like" Japan at 0.170.

Where no economy sits within 5%, the tile says so rather than naming one. It reads "no economy this clean", "no close match" or "no economy this carbon-intensive", and the note names the economy at the edge of the table so the reader keeps a bearing. Falling back to the nearest of any size, as it did until 2026-09-07, put Switzerland against a 2100 world at 0.011 kg/$ when Switzerland emits 0.043 — four times as much. Three of the eight presets land in that stretch, CMIP7 MEDIUM-to-LOW, CMIP7 VERY LOW and the Ausubel methane economy, so the tile was asserting a likeness that did not exist on every one of them.

Four stretches inside the tool's reachable range hold no economy within 5%: below 0.041 kg/$, and the gaps at 0.046–0.049, 0.071–0.076 and 0.381–0.404. The note prints the reader's own ratio in every case, so the gap stays visible. `tests/analogue.test.ts` walks the whole range and holds every rule.

The chart axis

The vertical axis is not fixed. It is chosen for each render to cover the reader's path and all seven markers together, on round steps, with zero always on it. Two things follow. A scenario that runs far above the published range stays fully drawn rather than clipping at the top of a fixed axis. And the seven markers stay on the same axis as the reader's line, so the comparison holds however far the sliders are pushed, at the cost of the markers compressing when the reader's path dwarfs them.

Rounding

Warming is shown to two decimal places, never more. Cumulative CO_2 is shown as a whole number. Rates carry their sign, so +1.91 and −1.43 read as a pair.

Part 3

Data

Every number the interface shows comes from `src/data/*.json`. Nothing is typed into the UI code. This file says where each series came from, when, and in what units.

Two scripts write those files, split by whether the numbers can be rebuilt from a primary source on this machine.

Script	Writes	Rebuildable
`scripts/build_data.py`	`observed.json`, `analogues.json`, `base.json`	Yes, from the sources below
`scripts/build_wpp.py`	`learn_population.json`	Yes, from the UN download
`scripts/build_energy_intensity.py`	`learn_energy_intensity.json`	Yes, from the sources below
`scripts/build_fuel_mix.py`	`learn_fuel_mix.json`	Yes, from the sources below
`scripts/build_income.py`	`learn_income.json`	Yes, from the sources below
`scripts/build_methane.py`	`learn_methane.json`	Yes, from the EDGAR download
`scripts/build_land_use.py`	`learn_land_use.json`	Yes, from the cached Global Carbon Budget
`scripts/extract_prototype.py` then `scripts/build_carried_data.py`	`config.json`, `emulator.json`, `markers.json`, `population.json`, `presets.json`, `notes.json`	No, carried from the prototype

Sources

Energy Institute Statistical Review of World Energy 2026

Downloaded by hand each June, held at `/home/rpielke/EI-Stats-Review-2026.xlsx` outside this repository because of its size. Sheets used:

- `CO₂ from Energy`, world row, 1965 to 2024, million tonnes CO_2. Covers CO_2 from energy only, with no cement or other industrial process emissions.
- `Total Energy Supply (TES) -EJ`, world row and country rows, 1965 to 2024, exajoules.
- `CO₂ from Flaring`, world row, 1975 to 2024, million tonnes CO_2. Used only in a base-year check.

The 2026 edition carries a 2025 estimate. It is excluded; the series stop at 2024, the last complete year.

One trap. The header row repeats the final year label three times: once for the value, once for the year-on-year percentage change, and once for the share of world total. Reading the header into a dictionary keyed by year silently returns the share column, which is 1.0 for the world row. `read_sheet` cuts the year block at the first non-monotonic repeat and asserts the world value is not 1.0.

World Bank

Fetched live from the public API, no key needed, cached in `scripts/_worldbank_cache.json`.

- `NY.GDP.MKTP.PP.KD`, GDP at purchasing power parity, constant 2021 international dollars. World aggregate 1990 to 2024, and one 2024 value per country for the analogue table. The series begins in 1990, which is why every GDP-derived rate starts there.
- `SP.POP.TOTL`, world population. This is the UN World Population Prospects 2024 revision as redistributed by the World Bank.

Country names in the Energy Institute workbook are joined to World Bank ISO3 codes through `scripts/country_iso3.py`, carried over from the decarbonization project where it was checked against the World Bank country list. It covers 78 of the 80 country rows. Taiwan and USSR have no entry, deliberately: the World Bank publishes no GDP series for either. Two names in that table have caused silent drops before and are handled explicitly, `Türkiye` (renamed from `Turkey` in the 2026 edition) and `Russian Federation`.

Venezuela drops out of the analogue table because the World Bank has no 2024 PPP GDP for it. The build prints every drop rather than swallowing it.

Global Carbon Budget

Fetched from the Our World in Data redistribution (`https://raw.githubusercontent.com/owid/co2-data/master/owid-co2-data.csv`), world rows only, cached in `scripts/_gcb_cache.json`.

Used for the base year, because the CMIP7 markers count all fossil and industrial CO_2 while the Energy Institute series covers energy and flaring alone. For 2024:

Component	$GtCO_2$
Fossil and industry, total	38.599
of which cement	1.473
of which flaring	0.416

Component	$GtCO_2$
of which other industry	0.424
Land use change	4.586
Total CO_2	**43.185**

That total is the basis the marker paths are on, which is what the marker values for 2025 confirm: they span 42.4 to 43.8 with a mean of 43.08.

ScenarioMIP CMIP7 marker runs

The seven marker emissions paths, their cumulative totals, their 2100 methane and land use values, and the Kaya factors that place each marker's tick under each slider. Five-yearly from 2025 to 2100, $GtCO_2$ including land use.

No primary source on this machine. These came into the prototype from ScenarioMIP and are lifted out of it by `scripts/extract_prototype.py`. Three markers (LOW, LOW-to-NEGATIVE, VERY LOW) publish no carbon-intensity rate, so those sliders carry four scenario ticks rather than seven.

UN World Population Prospects 2024

Two files, both fetched from the UN Population Division's own download area and both keyed to the 2024 revision.

- `WPP2024_TotalPopulationBySex.csv.gz`, dated 13 December 2024, about 17 MB compressed. Every location, every variant, 1950 to 2100, population in thousands. Used for the world series, the low, medium, high, momentum and constant-fertility variants, the 95% prediction interval, and the seven regional totals the population builder adds up.
- `WPP2024_Demographic_Indicators_Medium.csv.gz`. Total fertility for the world in 1950, 1990, 2024 and 2100, and for each of the seven regions in 2024.

Neither file is committed. `scripts/build_wpp.py` extracts the subset the site uses into `scripts/_wpp_cache.json`, which is committed, so a build works offline. Re-run with `--refresh` to fetch again.

One check the build enforces. The seven regions have to reproduce the UN's own world figures at the low, medium and high variants, to within 0.002 billion. They do: 6.987, 10.180 and 14.395 billion. The build exits rather than write a file whose regions sum to something the UN never published.

The 95% prediction intervals do not add up that way and the site says so on the page: the regional lower bounds sum to 8.234 billion against a world lower bound of 9.047, because the regions do not all land at the bottom of their own ranges at once.

Africa is split into sub-Saharan Africa (an SDG region in the UN's own hierarchy) and northern Africa (a subregion), which sum exactly to the UN's Africa figure at every variant. That keeps the region carrying most of the remaining growth on a control of its own.

IPCC 2006 Guidelines, default carbon content

`Volume 2 (Energy), Chapter 1, Table 1.3`, read from the published PDF. Values in kilograms of carbon per gigajoule, converted to CO_2 at 44/12:

Fuel	kgC/GJ	kgCO_2/GJ
Other bituminous coal	25.8	94.6
Crude oil	20.0	73.3
Natural gas	15.3	56.1

Nuclear, hydro, wind and solar carry none. Bioenergy carries none in energy-sector accounting, where biogenic CO_2 is reported under land use, which is also how the Energy Institute's CO_2 series treats it.

Applied raw, these overstate what the world emits. The 2024 mix gives 65.15 kgCO_2 per GJ against the 59.94 the Energy Institute reports from energy that year, mostly because part of the oil supply becomes plastics, lubricants and bitumen rather than exhaust. `build_fuel_mix.py` records the ratio, 0.9201, and the page applies it so today's mix reproduces today's intensity.

Maddison Project Database 2023

Downloaded from the Dataverse copy the Groningen Growth and Development Centre publishes, cached as `scripts/_mpd2023.xlsx` (5 MB, not committed) with the extracted world series in `scripts/_maddison_cache.json` (committed).

Used for one thing: world output before 1990, where the World Bank's purchasing-power series does not reach. The `Full data` sheet gives GDP per capita and population by country and year. The build sums GDP over the **151 countries with an unbroken 1965 to 2022 record**, 96.5% of world GDP in 1990, so each year compares the same economies as the last rather than a growing sample.

Only the growth rates are used. The World Bank level in 1990 is carried backwards on them, so every level from 1990 onward rests on the World Bank alone.

How much that choice decides. Over the 32 years the two sources share, the Maddison sample grows 3.46% a year and the World Bank series 3.19%. Carrying the level back on growth rates adjusted to close that gap gives a 1965 energy intensity of 5.86 MJ per dollar instead of 6.25, and a whole-record improvement of −0.92% a year instead of −1.03%. Both figures appear on the page.

IIASA SSP database v3.2 (June 2025 release)

World population trajectories for SSP1, SSP2 and SSP3, five-yearly 2025 to 2100, billions, aggregated from countries to world. Also carried from the prototype with no primary source here.

The UN 2024 95% prediction interval for 2100 population, 9.0 to 11.4 billion around a median of 10.2, is stated in the brief.

FaIR v2.2

The emulator coefficients, fitted to FaIR runs of the seven markers using the fair-calibrate v1.4.1 constrained ensemble. Carried from the prototype. See Part 2 for what the fit does and does not reproduce.

The files

`base.json`

The base-year state, all from 2024, the last complete year, applied at 2025.

Field	Value	Units	Source
`populationBn`	8.1409	billions	World Bank `SP.POP.TOTL`
`gdpPerPersonUsd`	21,393.2	constant 2021 international $	World Bank PPP GDP over population
`energyPerDollarMj`	3.4002	MJ per dollar	EI total energy supply over World Bank PPP GDP
`co2PerEnergyKgGj`	65.181	kg CO_2 per GJ	Global Carbon Budget fossil and industry CO_2 over EI total energy supply
`landUseGt`	4.586	$GtCO_2$/yr	Global Carbon Budget land-use change CO_2
`methaneMt`	380	Mt/yr	stated in the brief

`superseded` records what the prototype used and why it changed. The old value of 60.5 kg per GJ came from energy CO_2 plus flaring alone, which started every path about 4.4 $GtCO_2$ below all seven markers.

Land use moved on 2026-09-09 from the brief's 3.83 to the Global Carbon Budget's 4.586 for the same year, the source the rest of the base year already used. The `basis` block counted 4.586 for the marker-comparable total while the path started from 3.83, so the path ran 0.76 $GtCO_2$ below the basis it was drawn against and the land use chart stepped where the record met the reader's line. Part 2 records what the change moved.

`observed.json`

World history 1965 to 2024, the three Kaya factors derived from it, and the rate windows the notes quote.

Rate	Value	Period
Income per person	+1.90 %/yr	1990 to 2024
Energy per dollar	−1.43 %/yr	1990 to 2024
CO_2 per unit energy	−0.21 %/yr	1990 to 2024
Income per person	+2.02 %/yr	2015 to 2024
Energy per dollar	−1.62 %/yr	2015 to 2024
CO_2 per unit energy	−0.48 %/yr	2015 to 2024

Window extremes, each reported with the window that produced it:

Window	Highest	Lowest
Income, 25 years since 1990	+2.23 %/yr (1994–2019)	+1.86 %/yr (1990–2015)
Energy per dollar, 25 years since 1990	−1.36 %/yr (1990–2015)	−1.57 %/yr (1996–2021)
CO_2 per unit energy, 30 years since 1965	−0.17 %/yr (1992–2022)	−0.56 %/yr (1965–1995)

Reported as plain highest and lowest rather than slowest and fastest, because which end counts as slow progress flips with the sign of the series.

`learn_population.json`

Everything the population Learn More page draws and builds from, in the shape every Learn More data file uses: `series` for lines, `bands` for uncertainty, `parts` for builder controls, `constants` for the numbers the prose quotes.

Field	Value	Units
World, 2025	8.232	billions
World peak, medium	10.289 in 2084	billions
World 2100, medium	10.180	billions
World 2100, 95% interval	9.047 to 11.437	billions
World 2100, low and high variants	6.987 and 14.395	billions

Field	Value	Units
World 2100, momentum variant	9.596	billions
World 2100, constant fertility	18.193	billions
Total fertility, world	4.85 (1950), 3.31 (1990), 2.25 (2024), 1.84 (2100)	births per woman

The seven builder regions, 2100, billions:

Region	Low	Medium	High	2025
Sub-Saharan Africa	2.378	3.351	4.606	1.274
Northern Africa	0.319	0.462	0.648	0.276
Asia	3.074	4.613	6.677	4.835
Europe	0.413	0.592	0.824	0.744
Latin America and the Caribbean	0.400	0.613	0.905	0.668
Northern America	0.348	0.475	0.637	0.388
Oceania	0.053	0.073	0.098	0.047

Not yet here: IHME. The brief asks for the IHME projection alongside the UN and the SSPs on that chart. The IHME data portal was unavailable when the page was built, so the chart carries a line saying the projection joins it in a later revision. Nothing on the page reports an IHME number.

`learn_energy_intensity.json`

World energy intensity 1965 to 2024, and the rate of every 25-year window inside it.

Field	Value	Units
Intensity, 1965	6.2506	MJ per dollar
Intensity, 1990	5.5442	MJ per dollar
Intensity, 2024	3.4002	MJ per dollar
Rate, 1965 to 2024	−1.03	%/yr
Rate, 1990 to 2024	−1.43	%/yr
Rate, 2015 to 2024	−1.62	%/yr
Primary energy, 1965 and 2024	149.7 and 592.2	EJ
Energy growth, 1965 to 2024	+2.36	%/yr
Output growth, 1965 to 2024	+3.42	%/yr

The 2024 intensity of 3.4002 MJ per dollar reproduces `base.json` exactly, and the 1990 to 2024 rate of −1.4277 reproduces `observed.json`, because both rest on the same World Bank series over those years.

The 35 windows of 25 years run from **−1.5675 %/yr (1996 to 2021)** to **−0.4786 %/yr (1965 to 1990)**, with a median of −1.07. CMIP7 HIGH's −0.66 sits slower than 31 of the 35. The tool's own "Trend continues" rate of −1.62 sits faster than all 35.

`learn_fuel_mix.json`

Shares of world primary energy by fuel, 1965 to 2024, the emission factors above, and the two carbon-intensity series.

Share of primary energy	1965	2024
Coal	38.9%	27.9%
Oil	43.3%	33.6%
Gas	15.2%	25.1%
Nuclear	0.2%	5.2%
Hydro	2.2%	2.7%
Wind and solar	0.0%	2.9%
Bioenergy and other	0.2%	2.6%

Whatever the named fuels miss, biofuels most of it, joins bioenergy and other, so the shares add to the total energy supply exactly. The build exits if they do not sum to 100 in any year.

Two intensities, and they differ. On the basis the slider measures, which includes cement, flaring and other industrial CO_2 as the CMIP7 markers do, the world ran 75.56 kgCO_2 per GJ in 1965 and 65.18 in 2024. On combustion alone, 74.79 and 59.94. The 2024 slider-basis figure reproduces `base.json` exactly.

That gap changes the rate as well as the level: **−0.15%/yr from 1990 to 2024 on the slider's basis, −0.21%/yr on combustion alone**. The calibration mark under the front page's slider used −0.21 until 2026-09-05 and now uses −0.15, so the mark and the slider measure the same quantity; `config.json` keeps the old figure under `supersededRates`. Part 2 records what moved with it.

The non-combustion term, 5.237 kgCO_2 per GJ in 2024, breaks down as cement 2.487, other industry 0.717, flaring 0.702, and 1.331 for the difference between the Energy Institute and Global Carbon Budget inventories. Holding it where it stands puts a floor of 5.237 kgCO_2 per GJ under any fuel mix, which caps the improvement this page can reach at −3.31%/yr.

`learn_income.json`

World GDP per person 1965 to 2024, on the same spliced output series as the energy-intensity page, divided by World Bank world population, plus the three

World Bank income groups from 1990.

	Level, 2024	Growth since 1990	Past decade	Share of people	Share of output
World	$21,393	+1.90 %/yr	+2.02 %/yr		
High income	$58,645	+1.53 %/yr	+1.53 %/yr	17.4%	47.8%
Middle income	$15,025	+3.51 %/yr	+3.21 %/yr	73.4%	51.5%
Low income	$2,396	+0.75 %/yr	+0.22 %/yr	9.2%	1.0%

The 2024 world figure reproduces `base.json` exactly. The three groups cover 100.0% of world population and their population-weighted average comes to $21,464, 0.3% above the world figure; the builder scales by that ratio.

The current World Bank classification is applied across the whole record, so no country moves group mid-series.

`learn_methane.json`

World anthropogenic methane by source, 1990 to 2023, from EDGAR's 2024 release, summed from country and IPCC 2006 sector to five groups.

Source	1990	2023	Share	Growth
Fossil fuels	87.6 Mt	118.7 Mt	34.0%	+0.93 %/yr
Livestock	102.2 Mt	125.5 Mt	35.9%	+0.63 %/yr
Rice	39.1 Mt	36.8 Mt	10.5%	-0.19 %/yr
Waste	42.7 Mt	65.5 Mt	18.7%	+1.31 %/yr
Other anthropogenic	1.7 Mt	3.1 Mt	0.9%	+1.87 %/yr
Total	**273.2 Mt**	**349.6 Mt**		**+0.75 %/yr**

Rice is the one source that fell.

EDGAR's inventory and the tool's base year differ. EDGAR reaches 350 Mt for 2023; the tool uses 380 Mt. The Global Methane Budget 2000–2020 puts direct anthropogenic emissions at 369 Tg a year for 2010–2019, range 350 to 391, so both numbers sit inside the published range. The page scales each source by 1.086 so today's five reproduce the tool's base.

Natural wetlands, which the Global Methane Budget puts at 248 Tg a year together with inland fresh water, sit outside the slider and outside this file.

`learn_land_use.json`

The Global Carbon Budget land-use flux 1965 to 2024 with its uncertainty, its decomposition, and the forest growth rates the restoration control uses.

Field	Value
Flux, 1965	6.39 $GtCO_2$/yr
Flux, peak	8.02 $GtCO_2$/yr in 1997
Flux, 2024	4.59 $GtCO_2$/yr
Uncertainty	±2.57 $GtCO_2$/yr (±0.7 GtC, 1σ)
Gross deforestation, 2014–2023	6.23 $GtCO_2$/yr
Regrowth, 2014–2023	4.40 $GtCO_2$/yr
Other transitions and peat	2.20 $GtCO_2$/yr (the remainder)
Net, 2014–2023	4.03 $GtCO_2$/yr

Two figures for the same decade. The Our World in Data redistribution drawn on the chart averages 5.21 $GtCO_2$ over 2014–2023; the Global Carbon Budget 2024 paper reports 4.03 for that decade. The 1.18 $GtCO_2$ between them, most likely a vintage difference, sits well inside the ±2.57 uncertainty. The page says so rather than choosing one silently.

Forest growth rates come from IPCC 2006 Volume 4, Table 4.9, in tonnes of dry matter a hectare a year, converted at the 0.47 carbon fraction (Table 4.3) and 44/12:

Stand	t d.m./ha/yr	tCO_2/ha/yr
Tropical rain forest, South America, over 20 years	3.1	5.3
Young regrowth, Africa	10.0	17.2
Young regrowth, South America	11.0	19.0
Young regrowth, insular Asia	13.0	22.4

Above-ground biomass only; below-ground carbon and soil add more.

`analogues.json`

66 economies, their 2024 CO_2 and PPP GDP, and CO_2 per dollar in kilograms. Restricted to economies above 25 Mt CO_2 and $40 billion. The build asserts the count and prints anything dropped at the join. GDP carries the comparison as well as the ratio: the tile names the largest economy within 5% of the reader's 2100 world, not the nearest of any size, and says plainly that there is none when nothing sits that close. See Part 2.

`markers.json`, `population.json`, `emulator.json`, `config.json`, `presets.json`, `notes.json`

Carried over, as above. `config.json` also keeps `prototypeBase`, the base-year state the prototype used, so the two can be diffed.

Vintages

Source	Vintage	Checked
Energy Institute Statistical Review	2026 edition	2026-09-04
World Bank PPP GDP and population	live API	2026-09-04
Global Carbon Budget via Our World in Data	live file	2026-09-04
ScenarioMIP CMIP7 markers	as carried in the prototype	2026-09-04
IIASA SSP database	v3.2, June 2025 release	2026-09-04
UN World Population Prospects	2024 revision, file dated 2024-12-13	2026-09-05
Maddison Project Database	2023 release	2026-09-05
EDGAR greenhouse gases	2024 release, CH_4 to 2023	2026-09-05
Global Methane Budget	2000–2020, published 2025	2026-09-05
Global Carbon Budget paper	2024, published 2025	2026-09-05
IPCC 2006 Guidelines	Volumes 2 and 4	2026-09-05
FaIR calibration	v2.2, fair-calibrate v1.4.1	2026-09-04