

Data Viewer (user guide)

for Donor Compass and Cross-Cause Fund

Introduction to Data Viewer


We built a [Data Viewer](#) to enable users to inspect the cost-effectiveness data that informs Donor Compass recommendations and Cross-Cause Fund allocations. It is not the Donor Compass itself; it is a way to look under the hood at the numbers behind the tool.

It is intended for Rethink Priorities staff and reviewers who want to sense-check inputs, compare model versions, or understand why particular funds score the way they do under different assumptions. Fair warning: It's a lot of data.



Getting started

Select the **most recent dataset** in the drop-down menu on the left.



A note on the data

This data can look overwhelming at first, and that is normal. There are very large differences in scale between funds and between time periods. For example, some extinction-risk interventions produce values in the trillions or higher for distant time periods. In contrast, near-term global health interventions produce values in the thousands. These are not errors; they reflect genuine differences in scale under different moral and empirical assumptions.

As a result, charts may not look intuitive on first glance. A fund like GiveWell that performs strongly in the near term (T0–T2) may appear as a flat line compared to an extinction-risk fund whose T4–T5 values dominate the scale. The per-year toggle (which divides values by period length) and the fund and time period toggles can help make near-term effects more visible.

Sidebar controls

The left sidebar contains controls that apply across all tabs:


Funds	Toggle which funds appear in the charts. All funds are shown by default.
Unit mode	Total shows the raw value for the full time period. Per-year figures are divided by the period length, making short and long periods more comparable. Note: T5 (500+ years) uses 2,000 years as the divisor, which is a simplification since the period is actually unbounded.

What each tab shows

Compare Funds	One chart per effect type (e.g., Human Life-Years, Chickens DALYs). Each chart shows all selected funds as grouped bars across the selected time periods, for a single risk profile chosen at the top. Good for seeing which funds produce the most impact under a given worldview, and for spotting which effect types each fund affects. Impact is expressed per \$1M.
Compare Risk Profiles	Shows all 8 risk profiles as separate bar groups for each fund and effect type. All selected funds are shown, stacked below each other. Useful for understanding how sensitive a fund's score is to assumptions about tail risks, downside scenarios, and future discounting. Click profile names in the legend to hide/show individual profiles. [Note: only a selection of these risk profiles is used in our current advice]
Diminishing Returns	Shows how cost-effectiveness changes as cumulative donations to a fund increase. The y-axis is a multiplier: 1.0 means full effectiveness; 0.5 means each additional dollar buys half as much as the first dollar. The top chart overlays all funds for comparison; the per-fund charts below show more detail. The x-axis is cumulative spend in increments of the model's increment size.
A vs B	Only for reviewing changes made to the data. Compares two datasets side by side. Requires Dataset B to be selected. Shows raw numbers for A, raw numbers for B, and the difference (B minus A) as a heatmap where green indicates B is larger and red indicates B is smaller. Covers both diminishing returns and all effects across all risk profiles and time periods. Scroll horizontally within each table to see all columns.

Interacting with charts


- Click a fund or profile name in the legend above the charts to hide or show that series across all charts on the page. A strikethrough label indicates it is currently hidden.



- Hover over any bar to see the exact value in a tooltip. For capped bars, the tooltip shows the real uncapped value.



- Use the time period pills (on the Fund Charts and Risk Profiles tabs) to hide specific time periods.



- The A vs B tables scroll horizontally. Use the arrow buttons and drag bar below each table to navigate across columns.

	310M	320M	330M
	0.931	0.929	0.928
	0.931	0.929	0.928
	0	0	0
			→

Quick reference

The 6 time periods used in the model:

Label	Years	Per-year divisor
T0	0–5	5
T1	5–10	5
T2	10–20	10
T3	20–100	80
T4	100–500	400
T5	500+	2,000 (simplification — unbounded)

The risk profiles:

Neutral: Take the estimate at face value. If something has a 10% chance of saving 1,000 lives, count it as 100 expected lives saved. No adjustment for uncertainty or risk.

Standard expected value. Outcomes are weighted by their probability; no further adjustment. So if an outcome has value v and occurs with probability p , its contribution is $p \times v$

Upside Skeptical (Continuous): Don't trust the most optimistic scenarios. If a model says there's a small chance of enormous impact, we heavily discount that claim. This protects against inflated estimates driven by speculative best-case outcomes.

Outcomes up to the 97.5th percentile receive full weight; above that, the weight decays exponentially, falling to 1% at the 99th percentile and reaching zero by the 99.9th percentile. In practice, this nearly eliminates the contribution of extremely high-end outcomes while leaving the bulk of the distribution untouched. The formula for the weight between the 97.5th and 99.9th percentile is $\text{weight} = e^{\frac{-\ln(100)}{1.5} \times (p-97.5)}$. The average weight of all values standardizes this weight.

Downside Critical: Treat potential harms more seriously than potential benefits. A 10% chance of causing harm matters more to us than a 10% chance of doing equivalent good. Use this if you care more about avoiding harm than maximizing upside.

Gives amplified weight to harmful tail outcomes. Outcomes below zero are multiplied by 5, reflecting asymmetric aversion to causing harm.

Weighted Linear Utility (WLU): The bigger the stakes, the more cautious we get. We're less impressed by very large claimed benefits (they might not materialize) and more

concerned about very large potential harms (we really want to avoid those). The adjustment strength can be tuned, gentle or aggressive.

A formal risk adjustment method from [Bottomley & Williamson \(2023\)](#). WLU applies a stakes-sensitive penalty to worse outcomes via a coefficient c .

Formally, if an outcome i has value v_i and occurs with probability p , the weight assigned to it is:

$$w(i) = \frac{1}{1+|v_i|^c} \text{ if } v_i \geq 0$$

$$w(i) = 2 - \frac{1}{1+|v_i|^c} \text{ if } v_i < 0$$

The contribution C the outcome makes to the risk-weighted value of the intervention is as follows:

$$C(i) = \frac{\frac{w(i)}{n}}{\sum_{i=0}^n \frac{w(i)}{n}} \times p \times v_i$$

The weighted linear utility of the intervention I is thus the sum of all contributions from each value:

$$WLU(I) = \sum_{i=0}^n \left(\frac{\frac{w(i)}{n}}{\sum_{i=0}^n \frac{w(i)}{n}} \times p \times v_i \right)$$

At $c = 0$, WLU reduces to the expected value. Higher c means stronger down-weighting of bad outcomes.

Positive outcomes are weighted less as their magnitude grows (risk aversion). Negative outcomes get the mirror weight above 1 (loss aversion). We include different versions of this with different values of c .

Combined: Both skeptical of extreme upside and extra-concerned about downside. The most cautious option.

Applies the continuous upside skepticism decay described above (97.5th–99.9th percentile exponential decay) alongside Downside Critical loss aversion. Both adjustments apply simultaneously.

You can find more details in [Laura Duffy's 2023 RP report on risk aversion and cause prioritization](#).