

How these scores are built, and what they cannot tell you.

Turning human suffering into a number is uncomfortable. We do it anyway because the alternative is worse: comparing crises by how much news coverage they happen to get, without ever saying so out loud. This page explains exactly how each score is built, so you can agree with it, disagree with it, or check our work.

The five factors

Every problem on this site is scored on five separate measures, each on a scale of 0 to 100.

These are not raw numbers. A war with ten million deaths and one with ten thousand deaths cannot both fit on the same simple scale and still leave room to compare anything in between. So each factor is a judgment call, scored against fixed reference points, not a direct translation of a raw count.

Loss of life (default weight: 30). Deaths directly caused by the problem, adjusted for whether the toll is annual or cumulative.

Lives impacted (default weight: 25). People displaced, injured, food insecure, or otherwise seriously harmed.

Financial cost (default weight: 20). Direct damage, lost economic output, and health system costs tied to the event.

Environmental impact (default weight: 15). Emissions, habitat loss, contamination, and how reversible the damage is.

Trajectory (default weight: 10). Whether the situation is getting worse, and how fast.

Where to find the actual numbers

The 0 to 100 scores are not tied to a specific year and are not the raw figures themselves. They are built so that the worst problems in the world land in the 90s, and small, contained problems land in the single digits.

If you want the real death toll, the real dollar figure, or the exact time period a number covers, go to the “Standardized Impact” section on that problem’s page. That is where the raw numbers, sources, and notes live.

A dash in that section means we do not have a reliable number, either because no credible source has published one or because that measure does not apply to that kind of problem. A

dash is not a zero, and it does not lower or raise the score.

The weighting

The main severity score is a weighted average of the five factors above, using our default weighting: 30 for loss of life, 25 for lives impacted, 20 for financial cost, 15 for environmental impact, and 10 for trajectory. Loss of life counts the most. Trajectory counts the least. That is an editorial choice, not a scientific one.

Because it is a choice, you can change it. Use the sliders on the index page to set your own weighting. Move loss of life to 50 and everything else to zero, and you get a pure death toll ranking. Weight environmental impact heavily, and slow, spreading damage will rise above sudden emergencies. Your weighting is saved in the page's link, so you can send anyone a version of the ranking built around what you think matters.

Scoring bands

0 to 34: Limited. 35 to 49: Serious. 50 to 64: Severe. 65 to 79: Critical. 80 and above: Catastrophic.

These bands use the same colors everywhere on the site.

Confidence levels

High. The numbers come from institutions with an established track record, such as the UN (data.un.org), WHO (who.int/data/gho), IPCC (ipcc.ch/reports), or a national statistics agency, and credible sources are close to agreement.

Medium. Credible sources exist, but they disagree with each other by a meaningful margin, or the count only covers part of the picture (for example, confirmed deaths are known but the true, higher total is not).

Low. The numbers are disputed, estimated indirectly, actively suppressed, or based on early reporting that is likely to change. Most news-derived entries start here.

Why don't these numbers match what I've seen elsewhere?

Because credible sources disagree with each other, often by a wide margin, and usually for a real reason rather than carelessness. Two published death tolls for the same war can differ several times over depending on whether they count only confirmed deaths or also estimate excess deaths, when the counting period starts, whether the number comes from verified reports or statistical modeling, and whether anyone tracking it actually has access to the affected area.

The same problem applies to "lives impacted," which depends entirely on where you draw the line for what counts as harm, and to financial cost, which depends on whether you

include indirect and future losses.

We do not average competing estimates into one number to make it look like sources agree when they do not. We pick one defensible published estimate, say where it came from, and label how confident we are in it. If a number here does not match a number you have seen elsewhere, the most likely reason is a different definition, not a mistake on our part.

If you want to report a problem with a specific entry, the useful things to flag are a broken source link, a figure that has since been replaced by a newer published number, or an entry filed under the wrong category or region. Reports are reviewed in batches, not answered one by one.

Deep-sourced versus news-derived entries

This site started with about a hundred entries covering the major ongoing global problems: wars, disease, famine, environmental collapse, and long-term economic harm. These were built in one careful research pass using named institutional reports from sources like the WHO (who.int/data/gho), the IPCC (ipcc.ch/reports), the World Bank (data.worldbank.org), and the UNDP (hdr.undp.org). We call these entries Deep-sourced.

After that, an automated process regularly scans world news, groups related headlines together, and either updates an existing entry or creates a new one. Anything created this way is labeled News-derived. This matters because a news-derived figure reflects what this week's reporting supports, not a fully settled count.

When a news sweep touches a Deep-sourced entry, the original researched number acts as a floor. A new figure from the news can only replace it if the new figure is higher. Reporting tends to revise death tolls upward over time far more often than downward, so we would rather understate a change than let one partial headline erase a carefully researched total.

Movement and the “what changed” digest

After every news sweep, we save each problem's rank under our default weighting. The up and down arrows on the index compare today's rank to the most recent saved snapshot that is at least seven days old. Movement reflects real changes in the underlying scores, or new entries appearing above or below a given problem. It is not a measure of how urgent something feels.

The “what changed” summary is built directly from that same snapshot comparison, plus a record of which figures were revised. It is generated from the data, not written or embellished by a person.

What our system is allowed to publish on its own

The automated news pipeline can publish routine changes by itself, but not the kind of claim that would seriously damage this site's credibility if it turned out to be wrong.

A brand new entry is held back from public view, invisible until a person reviews and approves it, if any of the following are true: it scores 45 or higher, it claims a thousand or more deaths, it claims a hundred thousand or more people displaced, or it falls under the conflict category, where who is responsible is often disputed.

Updates to existing figures follow a similar rule. A change only counts as confirmed, rather than noise, once it crosses both a percentage threshold and a fixed amount: deaths must rise by at least 25 percent and by at least 500 people; lives impacted by at least 25 percent and 100,000 people; displacement by at least 25 percent and 50,000 people; economic cost by at least 30 percent and one billion dollars. Below those thresholds, a figure updates automatically and shows up in the change digest. Above them, it waits for a person to approve it, and the published number does not move until that happens. The first time a previously empty field gets a value, it is always held for review, since that counts as a new claim rather than an update.

Entries that a person has personally reviewed carry a checked date. Every entry, reviewed or not, has a button to report a problem with it, and any report goes into the same review queue.

Provisional scores

Our search tool researches a problem that is not yet on the list and estimates where it would rank if it were added. These results are provisional. They come from a single research pass, they are not saved anywhere on the site, and they should be read as a rough sense of scale rather than an official ranking.

Known limits

Death counts are recorded inconsistently around the world, and confirmed counts can differ from broader excess mortality estimates by a factor of ten or more. Financial cost figures mix direct damage with lost economic output. Chronic, spread out problems tend to be far better documented in wealthy countries than poor ones, which biases any global comparison. Slow moving disasters are also undercounted compared to sudden ones, because news coverage drives how often a figure gets revised.

This is an index of reported impact. It is meant to help you argue about proportion and priority, not to serve as a final, authoritative record of human suffering.

Frequently asked questions

What does a dash mean? Is it the same as zero?

No. A dash means we do not have a reliable published figure for that measure, or the measure does not apply to that kind of problem. A zero would be a claim that the real value is zero. A dash makes no claim at all, and it neither raises nor lowers the severity score.

How do provisional scores work?

When you research a problem that is not yet on the list through Search a Threat, the result is a provisional score. It comes from a single research pass, it is not saved into the ranking, and it can shift once a person reviews it. Read it as a rough sense of scale, not an official position on the list.

Why is a score not the same as a death toll?

Each factor is a 0 to 100 relative severity index, scored against fixed reference points, so crises of wildly different magnitudes can share one scale. The raw death tolls, dollar figures, and time periods live in the Standardized Impact section on each problem's page.

Can I change the weighting?

Yes. The sliders on the index page let you set your own weighting, and it is saved in the page's link, so you can share your version of the ranking. The house weighting is 30 loss of life, 25 lives impacted, 20 financial, 15 environmental, 10 trajectory.

What is the difference between Deep-sourced and News-derived?

Deep-sourced entries were built in a careful research pass from named institutional reports. News-derived entries were created or updated by the automated news pipeline, so their figures reflect what current reporting supports rather than a settled count.

How often does the list update?

A news sweep runs daily. Routine changes publish automatically; larger claims — high severity, big death or displacement figures, or conflict entries — are held until a person reviews them.

I think a figure is wrong. What should I send?

Use the report button on the entry. The useful things to flag are a broken source link, a newer published figure, or an entry filed under the wrong category or region. Reports are reviewed in batches and are not answered individually.

thethreatlist.org/methodology