HealthRisk

Human health risk from contaminated groundwater – in-browser app • Hydrogeological Repository

What it does

HealthRisk turns a measured concentration in groundwater into what it means for the people who drink it. One exposure scenario – who is exposed, for how long, and by which routes – drives four assessments: **toxic**, **carcinogenic**, **microbiological** and **radiogenic**. Because all four read the same scenario, their answers can be set beside each other, which is the question that usually matters on a contaminated site: of everything in this water, what is the thing to deal with first.

The equations are those of the US EPA *Risk Assessment Guidance for Superfund*. Reference doses and cancer potency factors come from the IRIS database, the dose–response parameters from the published pathogen literature, and the risk coefficients from US EPA Federal Guidance.

Everything runs in your browser

No installation, no account, and no dataset to download. The whole parameter database is 10 kB compressed and is fetched once when the page opens; every dose, quotient and probability after that is computed on your own machine, so changing a concentration and re-reading the answer is immediate – and nothing about the site you are assessing is transmitted anywhere.

The four assessments

Assessment	What it computes	How to read it
Toxic	The average daily dose over the published reference dose for that contaminant – the hazard quotient – summed across contaminants and pathways into a hazard index.	Below 1 the dose sits inside the exposure estimated to be without appreciable risk over a lifetime. At or above 1 it has left it.
Carcinogenic	The lifetime average daily dose times the cancer potency factor, through the Poisson model Risk = $1 - e^{-LADD \times CPF}$.	A probability of an additional cancer case. 10^{-6} is one case per million people exposed at that concentration.
Microbiological	Probability of infection from a dose of organisms, by the beta-distributed model or – for Giardia and Cryptosporidium – the single-hit exponential model, compounded over a year of exposures.	A probability of infection, which is not the same as illness. It rises towards certainty as the dose rises.
Radiogenic	An activity intake times a US EPA risk coefficient, for ingestion, inhalation or submersion, reported as both mortality and morbidity.	Mortality is the chance of dying from a radiation-induced cancer; morbidity is the chance of developing one, so it is always the larger.

What you get

Output	Detail
Hazard quotients	Per contaminant and per pathway, against the oral reference dose or the inhalation reference concentration, with the hazard index totalled.
Dose by pathway	The same doses stacked by route rather than divided by a reference dose. This is the chart that says what to do about an answer that is too high.

Output	Detail
Carcinogenic risk	Per contaminant and combined, against the conventional 10^{-6} to 10^{-4} band, expressed both as a probability and as odds.
Dose–response curve	The infection model for one organism across eight decades of dose, with the current exposure marked on it, and the concentration at which the annual risk reaches 1 in 10 000.
Radiogenic risk	Mortality and morbidity side by side per nuclide, with the intake or the concentration–time product that produced them.
The database	All 78 contaminants, 16 micro-organisms and 321 nuclides, searchable, as published.
Exports	CSV for each assessment – carrying the whole scenario alongside the results – and for each database table, plus chart images.

The exposure scenario

Everything below step 1 reads the same scenario, and the app is explicit about where each figure in it came from. The source document states five: two litres of water a day, twenty cubic metres of air a day, a seventy-year lifetime, a 70 kg adult and a 10 kg child. Every other figure an exposure calculation needs is supplied by the app from the conventional defaults, and is marked as such on the page, in the tables and in the exported CSV.

Two of the three exposure routes rest on a factor the source document never gives. It names inhalation as “contaminant transferred from water to vapor in air” and dermal sorption as “sorption through skin in baths and showers”, and stops there. Turning either into a dose needs a volatilisation factor or a permeability coefficient, and both are supplied here rather than taken from the document. The defaults are conventional, but they are assumptions: the volatilisation default assumes complete transfer from water to air, which is right for a volatile solvent and badly wrong for a metal.

One consequence of the equations as the document writes them is worth knowing before it is discovered. Dose is C × IR × ED and the average daily dose divides it by BW × ED, so the exposure duration cancels: the whole toxic assessment is independent of how long the exposure lasts. Only the lifetime average daily dose keeps it, which is why the exposure duration moves the carcinogenic and radiogenic answers and leaves the hazard quotients exactly where they were.

The database

Table	Holds	Records
Contaminants	Oral cancer slope factors, oral reference doses and inhalation reference concentrations.	78 contaminants: 17 with a potency factor, 75 with a reference dose, 6 with an inhalation value.
Micro-organisms	The α and β of the beta-distributed model, or the r of the single-hit exponential model.	16 organisms – bacteria, viruses and protozoa.
Nuclides	Mortality and morbidity risk coefficients for ingestion, inhalation and submersion.	321 nuclides, from H-3 to Cf-252.

Values are carried through exactly as published, misspellings included, so the table in the app can be checked against the source document line by line. Searching for the correct spelling still finds a record whose printed name is wrong. Where the source is internally inconsistent – a nuclide with a coefficient in one table and not its pair, a reference dose

with the magnitude of a slope factor – the source is carried unchanged and the app says so beside the table.

Limitations worth knowing

- This is an exposure and risk calculation, not a site model. It does not transport a contaminant from a source to a receptor, does not decide what concentration a receptor will see, and does not age or degrade a plume. Where the concentration comes from, and how representative it is, is the larger part of any real assessment and none of it happens here.
- Every default is a population-level figure rather than a measurement of the people at your site, and reference doses and potency factors carry uncertainty factors of ten to a thousand by design. These numbers decide whether something warrants a proper assessment; they do not make a compliance determination.
- A cancer potency factor is an extrapolation from animals at high doses to people at low ones, spanning several orders of magnitude and assuming no threshold below which the risk is zero. That extrapolation is usually a larger uncertainty than the concentration that was measured.
- The appendix carries one oral slope factor per carcinogen and no inhalation or dermal one, so where those pathways are switched on the same factor is applied to their share of the dose. Real inhalation unit risks differ from oral slope factors, sometimes by an order of magnitude.
- The hazard index adds quotients for contaminants that may do entirely different things. That is defensible where two attack the same organ by the same mechanism and is arithmetic without meaning where they do not, so a hazard index above 1 is a signal to read the contributions individually.
- The microbiological models give the probability of infection, not of illness, and the annual figure assumes each exposure is independent – ignoring both acquired immunity and the tendency of contamination to arrive in a few bad days rather than evenly through the year.
- One scenario applies for its whole duration: one body weight, one set of intake rates, one concentration. A child who becomes an adult during a thirty-year exposure is two scenarios, and the way to handle that is to run both.

The benchmarks drawn on the charts are conventions, not the document's. The 10^{-6} to 10^{-4} band for an incremental lifetime cancer risk, a hazard quotient of 1, and the WHO tolerable annual probability of infection of 1 in 10 000 are all drawn as reference lines because a risk plotted without one is unreadable. The source document gives its interpretation in words and sets no thresholds.

Credit

The dose, average daily dose and lifetime average daily dose equations, and the carcinogenic and toxic risk formulations, are those of the **US EPA**, *Risk Assessment Guidance for Superfund*. Reference doses and cancer potency factors are from the EPA **IRIS** database. The radiogenic risk coefficients are the revised coefficients published by the **US EPA** in its Federal Guidance. The beta-distributed and single-hit exponential dose–response models are the standard formulations of quantitative microbial risk assessment. This app is an independent implementation of those published methods and claims no authorship of any of them.

Beta. HealthRisk is published for evaluation and feedback. Check any figure you rely on against an independent calculation before using it in an assessment, and never present a screening number from it as a finding about a real population.