The Two Pathways

Karen Pendergrass · ORCID 0000-0002-2348-7259

Karen Pendergrass. "The Two Pathways". https://heavymetalcertified.com/program/two-pathways. Retrieved 2026-09-21.

This is Method v2. Why Method v1 was retired is on the method page. Occurrence characterises the market here. It does not invent a claim cell.

Most HMTc limits are not calibrated at all. The default standard for any subcategory and analyte is **the strictest maximum level set by a credible government regulator**, the EU, Codex, the FDA, WHO/JECFA, FSANZ, or Health Canada, converted to the product's native basis. That is a lookup against sovereign law, not a computation, and it is where most of the panel lands: nickel, tin, aluminium, and hexavalent chromium adopt the government limit directly, and where the law itself splits a limit by formulation, nickel and aluminium in infant formula, the standard adopts the law's own split and cites it.

The two pathways describe what the program does in the space the government limit leaves open: the four priority toxics, where a clean market can beat the law, and the cells that no government regulates at all. In those cells the program sets the number from the occurrence distribution, and the question every category asks differently returns: *how clean is achievable here, today, without certifying nothing or certifying everything?* The choice between the two pathways is the most consequential decision in setting one of those numbers, so it is worth being explicit about both the intent and the mechanism. It is how a number is set. It is not a second certification track: every claim analyte is equal at the gate.

Lead with the pass rate, not the percentile

Where the program sets an occurrence-based number, the pathways are easiest to understand as **pass-rate targets**:

- A **clean** subcategory is one where good product is the norm and contamination is the exception. The standard should pass roughly the cleanest **90%** of the market and exclude the dirtiest tail. The mechanism that delivers that is the **97th percentile (P97)** of the subcategory's occurrence distribution.
- A **dirty** subcategory is one where contamination is endemic, to the soil, the commodity, or the process, so that demanding near-zero on day one would certify almost nothing. The standard should pass roughly the cleaner **40%** of the market and put the rest on notice. The mechanism is the **45th percentile (P45)**.

The percentile is not the goal; the pass rate is. P97 and P45 are the levers calibrated to hit ~90% and ~40% against real occurrence distributions: targets confirmed at 94% and 42% against FDA per-sample lot data. For a priority toxic the occurrence value is floored against the government maximum, so the published limit is the *stricter* of the two: where the market cannot beat the law, the law governs; where it can, the tighter market value governs and the ratchet keeps room to work.

Why a dirty pathway is not a loophole

A P45 limit looks permissive next to P97, and out of context it could read as the certifier lowering its standards where the data is bad. It is the opposite. In an endemic category, a P97 limit would pass nearly everything: it would certify the status quo and change nothing. A P45 limit fails the majority of current product on purpose. It is the ratchet's first turn: it rewards the cleaner end of a dirty category, creates a commercial reason to source and process toward cleaner inputs, and tightens across editions as the category responds. The honest framing, *this category is dirty, here is how dirty, here is the limit and why it sits where it does*, is exactly the record that holds up when the standard is challenged.

Independent rows default to clean

Within the occurrence-set cells, a row with no within-pair partner is treated as **clean by default** and set at P97, with an override slot where the evidence justifies dirty treatment. Clean-versus-dirty is a category-level judgment resolved in governance review, not a per-row convenience: which is why a row whose pairing is unsettled is held at gate one rather than shipped on a guess.

The pool is the certification-relevant market

An occurrence percentile estimates a population, so it matters which population. The pool is restricted to the **certification-relevant market**: the jurisdictions that enforce binding maximum levels for the product class with comparable oversight. The admission test is **regulatory-equivalence, not geography**: a market with binding limits and equivalent surveillance is in, and no certified brand ships into a market the pool does not represent. Documented seasonal and geographic variance is priced into the distribution rather than granted as a per-brand waiver.

After publication: ALARA and the ratchet

Neither pathway is a resting place, and neither is the government-limit default. Once a limit publishes, an ALARA review ("as low as reasonably achievable") and an 80th-percentile ratchet trigger drive tightening in later editions as the category's distribution shifts cleaner. A standard that never moved would, over time, certify a market that had moved on without it.

For the breadth side of the same argument, why the panel grew from four metals to ten, and why a one-time test is weaker than a surveillance record, see Four Metals Is No Longer Enough and The Surveillance Protocol.