

The Cost of Being Uncertified

Why Palmquist Moved Heavy-Metal Certification From the Marketing Budget to the Balance Sheet

Karen Pendergrass

Standards Architect, the Paleo Foundation, Heavy Metal Tested & Certified (HMTc) Program

June 2026

This briefing is written for the two people who decide whether a heavy-metals certificate is worth its price: the Chief Financial Officer and the General Counsel. It makes one argument, in numbers you can re-run with your own inputs — that a Supreme Court decision handed down this winter quietly created a cost most brands in this category have not yet written on any spreadsheet, and that the cost lives in five places at once.

A new line in the ledger

Before February 24, 2026, the cost of being uncertified to a heavy-metals standard was a marketing question. After February 24, it is a balance-sheet question.

On that date a unanimous Supreme Court, in *Hain Celestial Group, Inc. v. Palmquist* [1], took a heavy-metal injury case that a manufacturer had carried all the way to trial and won outright, and erased the win. Not because the science failed. Not because a product was proved defective. Because the case had been litigated in the wrong courthouse. The federal court that tried it never had jurisdiction; the flaw, in the Court's words, "lingered through judgment"; and the verdict — a clean defense win, the kind a general counsel reports to the board — was vacated and sent back to Texas state court to be tried again, this time against the manufacturer and the retailer both. Years of defense were converted, in a single opinion, into a verdict worth nothing and a trial that must be held a second time.

Handed a certification proposal, most CFOs ask the same reasonable question: what does it cost? That is the wrong question, and asking the wrong question with great precision is how careful people walk into trouble. Charlie Munger built a career on the opposite instruction, borrowed from the mathematician Carl Jacobi: invert, always invert. So invert this one. Do not ask what certification costs. Ask what being the uncertified defendant costs — now, after *Palmquist*. That is the number that belongs on the risk register, and almost no brand has put it there, for a plain reason: it does not sit inside any one department's

spreadsheet. It is part litigation, part settlement, part evidence, part insurance, part procurement. A cost that lives in five departments tends to live in none. The purpose of this briefing is to add it up.

Start with the premise, stated plainly, because it is the thing everyone in this category already knows and would rather not say: the metals are already in the food. FDA's own Total Diet Study found lead in 15% of samples, arsenic in 43%, cadmium in 33%, and mercury in 10% [2]. The question was never whether your product carries a measurable metal — it almost certainly does, and so does your competitor's. The question is what you can show about it when someone with a subpoena asks. *Palmquist* changed who gets to ask, and where. Here is what that did to five separate numbers.

1. The litigation differential

The value of a lawsuit is not fixed. It is a function of the room it is tried in. Change the room and you change the price.

For years a manufacturer sued in state court held a reliable counter-move: argue that the in-state retailer who sold the product was a sham defendant, remove the case to federal court on the strength of that argument, then drop the retailer. Federal court is the better room for a defendant — more demanding expert-evidence rules, a wider jury pool, a slower and more deliberate docket. *Palmquist* disarmed the move. Name the in-state retailer, keep a colorable claim against it, and the case stays in state court; and the old remove-then-drop maneuver is now not merely useless but hazardous, because a judgment won after a retailer was wrongly dismissed can be vacated years later with the litigation clock reset to zero. The retailer is no longer a party you remove and forget. It is an anchor that holds the case in the forum you least wanted it in.

It is worth being clear-eyed here, in the way Munger insisted on being clear-eyed, because the honest version of the argument is also the stronger one. The plaintiffs' science is not robust. No court has found that the trace metals in mainstream baby food cause the neurodevelopmental injuries these suits allege, and the plaintiffs' own theory — that any detectable quantity can contribute to harm — is the theory you reach for precisely when you cannot prove a dose. But that is exactly why the forum matters so much. When the merits are uncertain, the case is not decided on the merits. It is decided by everything arranged around them: the jury, the judge, the speed of the docket, the cost of discovery, the pressure to settle. *Palmquist* moved every one of those against the uncertified defendant. A case you might have won in federal court still costs more to defend, and is likelier to be paid to go away, once it is anchored in a state courtroom.

Put a number on it the way you would on any contingent liability. The portable figure is the cost of defending a single major case: a widely cited federal-courts survey put mean outside legal fees at roughly \$2.0 million, with discovery alone frequently north of \$600,000^[3]. That is the cost of fighting, before a dollar of settlement or verdict. *Palmquist* raises the share of your cases that reach that posture and lengthens the ones that do. If it lifts your expected frequency of a fully defended major case from one-in-a-hundred years to one-in-fifty, you have just added about \$20,000 a year of pure defense cost to your run rate — and you have added it whether or not you ever lose, whether or not the science is ever good. That is the litigation differential. It is the smallest of the five numbers, and the one most brands have quietly set at zero.

2. The settlement-value shift

Show me the incentive, Munger liked to say, and I will show you the outcome. Palmquist changed the incentives on both sides of the settlement table at once.

A settlement is priced in the shadow of trial — what each side believes the case is worth if it goes the distance, discounted by the odds and the cost of getting there. *Palmquist* lengthened that shadow for the defendant in three ways. First, the forum is now more often a state court with a home-state jury, which raises the plaintiff's estimate of a verdict. Second, the retailer is locked in as a co-defendant, adding a party with its own insurance tower and its own reasons to settle — and, in a private-label program, concentrating retailer and manufacturer exposure in one company with no upstream party to point to. Third, and this is the genuinely new one, even winning is no longer final: a defense verdict obtained in a court that turns out to have lacked jurisdiction can be vacated, as Hain's was, with the case reset to zero years later. When a win can be unwound, the rational defendant pays more to make the matter disappear rather than risk paying for the same trial twice.

The direction is certain; the magnitude is yours to model. Take a case you would value today at \$500,000 to settle. Raise the plaintiff's probability-of-recovery estimate by even ten points, add the cost and risk of a possible second trial, and seat a co-defendant beside you, and the same case clears materially higher — call it \$750,000 to \$1,000,000 for illustration. Then remember that a Proposition 65 or consumer-class plaintiff does not bring one case; he slices your portfolio into SKUs, flavors, and lots and brings many. The Prop 65 economics already show the machine: one published settlement ran \$28,000, of which \$25,000 was the plaintiff attorney's fee and \$3,000 the actual penalty^[4] — an apparatus built to be run again and again. *Palmquist* raises the price of every turn of the crank.

3. The evidentiary asymmetry

A margin of safety is the evidence you built before you needed it. The flimsiest exhibit in any courtroom is the one assembled after the complaint arrived.

This is the component that sets the other four, and it is where the cost is largest. When a heavy-metal claim reaches a jury, the plaintiff does not have to prove that one extraordinary lot poisoned one child. He has to tell a story: the metals were present, the company knew or should have known, the company did not do enough. “They didn't test” is a sentence a jury understands without an expert's help — and it was available in the baby-food cases because the lead defendant had not tested its finished products for heavy metals until 2019^[5].

The defense against that story is evidence of a particular kind, and it has three properties. It must be *contemporaneous* — generated routinely, before any complaint, so it cannot be waved away as litigation-built. It must be *independent* — sampled and analyzed by parties who do not profit from a clean result, under ISO/IEC 17025 accreditation^[7] and a documented chain of custody. And it must be *continuous* — repeated on a schedule, across lots and across time, so that the lot the child actually consumed sits inside a demonstrated pattern of compliance instead of beside a single fortunate number. Most companies, and most ordinary internal quality testing, produce at most one of the three. A drawer full of internal certificates of analysis is none of them: it is not independent, and assembled after the fact it is not contemporaneous either.

Here is the asymmetry, and it is the cruelest fact on the list. The plaintiff builds his record on your clock, with your own documents, in discovery. You can only build yours in advance. Once the complaint lands, the single most valuable thing — a test record with a timestamp that predates the lawsuit — can no longer be created at any price. This is the one cost here that money cannot buy back after the fact.

The honest objection follows immediately: what if we test, and the record is not clean? But the option to have no record is already gone. California's AB 899 mandates the testing and makes the results public^[6]. The only remaining question is whether your record is the defensible kind. A staged path exists for exactly

this — enroll, go under intensified surveillance, and document an improving trajectory — and it converts the plaintiff's best sentence, "they knew and did nothing," into a far worse one for him: "they knew, enrolled in an independent program, and hold timestamped, third-party-verified evidence of continuous improvement that predates this complaint." That record is discoverable, and you should expect it to be discovered; the objective was never to avoid a record, only to make it the kind that helps. The enrollment date itself is an appreciating asset: a date stamped before any complaint is filed is worth a little more every month a complaint does not arrive.

Price it as what it is — an option you can buy only before the event, whose payoff is the swing it produces in the other four numbers: a narrowed claim, a denied class, a motion that lands, a settlement that clears lower, a verdict that holds. Against the mid-brand event range below of \$8 million to \$33 million, even a one-third improvement in posture on a single serious matter is worth several million dollars, set against an annual record-building cost measured in the low six figures. The record does not have to win the case to pay for itself. It only has to move the settlement.

4. Insurance repricing

An underwriter is a professional skeptic who prices only what he can observe. Give him something to observe and the price moves.

Product-liability carriers are repricing heavy-metal exposure in real time as *Palmquist* filters into their underwriting models, because the ruling raises both the frequency of cases that reach an unfavorable forum and the severity of each one. An underwriter cannot price your good intentions; he cannot see them. He can see an auditable, continuous, independent testing record, because that is exactly the observable that separates a managed risk from an unmanaged one. Several carriers are already evaluating reduced-premium endorsements for suppliers certified to recognized third-party heavy-metal standards. The direction is not in doubt: the certified supplier will pay less than the uncertified one for the same coverage, for the same reason a sprinklered building pays less for fire cover than the one next door without sprinklers.

Be precise about the mechanism, because it is easy to muddle. Certification is not insurance. It does not pay your loss; an underwriter does. What certification does is lower what the underwriter charges, because it lowers what he has to fear. The base is your own program, so the exact figure is yours to pull, but the structure is simple: if your product-liability and general-liability premiums run in the low-to-mid six figures, an underwriting credit of even 5% to 15% for documented certification is several thousand to a few tens of thousands of dollars a year — recurring, and tilting further your way at each renewal as the uncertified pool reprices upward around you. Ask your broker to model the number. This is the one component of the five that already appears as a line item your CFO reconciles every year.

5. Retailer-mandate risk

When the retailer's incentive stops matching yours, the contract is what moves.

Palmquist repositioned the retailer most of all. A retailer named as a co-defendant used to count on early dismissal, or on riding the manufacturer's removal into federal court. Both exits have narrowed. The retailer is now a permanent state-court defendant for products it did not formulate, manufacture, or test — and a private-label retailer carries the exposure twice, as the retailer for purposes of joinder and the manufacturer for purposes of liability. A party that can no longer get out of the courtroom does the next best thing: it lowers the odds and the size of the verdict by putting only demonstrably defensible product

on its shelves. In practice that means supplier-qualification language written against a published standard — and the brands that cannot meet it do not get the shelf.

This is the largest and most immediate number for a growing brand, because it is not a probability-weighted tail at all. It is a gate. The 2027 retailer buying cycles are being drafted in calendar 2026; the brands that can show certification by the third quarter of 2026 are the ones eligible for the first mandate-compliant supplier lists, and the retailers writing those lists this cycle will set the category's price of admission for the cycles that follow. Model the cost as the discounted revenue of the channel you cannot enter, or the channel you lose. In the event scenarios below, delisting and listing-suppression alone is modeled at roughly 15% of annual revenue — about \$0.38 million for a \$5 million brand, \$3.75 million for a \$50 million brand, \$37.5 million for a \$500 million brand — and that is merely the cost of losing one major channel for part of one year, not the cost of never qualifying for it at all. For a brand whose growth thesis rests on a national retailer's premium or private-label tier, the number is the entire value of that tier. Like the litigation differential, and unlike a lawsuit, this cost arrives whether or not anyone is ever sued.

Five models, one balance sheet

Notice what just happened. The cost of being uncertified did not resolve into a single figure, because it does not come from a single discipline. It came from five: civil procedure, settlement theory, the law of evidence, insurance underwriting, and retail procurement. This is Munger's latticework made concrete — no one mental model captures the exposure, and a CFO who runs only the litigation model, or only the channel model, will undercount it badly. Add the five as what they are: four probability-weighted tails and one near-certain gate, all drawing on the same balance sheet. Underneath them sits the figure that frames the whole question. A single public contamination event — recall, delisting, or serious lawsuit cluster — is modeled at a one-year cost of \$1.6 million to \$6.6 million for a small brand at \$5 million in revenue, \$8 million to \$33 million for a mid-sized brand at \$50 million, and \$57 million to \$187 million or more for a large brand at \$500 million^[8,9].

Now set that against the cost of the thing that addresses all five at once. HMTc license fees are revenue-tiered — a small brand pays an entry tier, a mid-brand a mid tier, a large brand more — and laboratory testing is a separate, un-marked-up pass-through to an accredited lab, disciplined by the lab market rather than by the program. The program deliberately publishes no single sticker price, because the fee scales with the brand. What it does state is the number it has to stay under to pay for itself. For a mid-sized brand, on downside-risk avoidance alone — before a dollar of channel value, before the insurance credit — that break-even is about \$112,500 a year, and the all-in annual cost runs comfortably under \$200,000^[8]. Buffett's formulation is the one to hold here: price is what you pay; value is what you get. The price of certification is a known, small, recurring number. The value is the five costs you would otherwise carry alone.

Then be exact about what the number is *not*, because a claim that overreaches is worth less than one that holds. Certification does not make a brand safe, and a serious program will refuse to let you say it does — the mark is a compliance signal under continuous surveillance, not a safety guarantee, and a credible program will not shield a licensee from a regulator. It does not eliminate the contamination that is endemic to the food supply, and it does not turn a bad lot into a good one. What it does is change the evidence you hold when the contamination that sits in everyone's supply chain turns up in yours. That is a narrower claim than "safe," it is the honest one, and it is the only one worth a CFO's signature.

One word in the title is doing quiet work: *credible*. A certificate frozen at the four metals of the 2021 congressional report^[5] — arsenic, lead, cadmium, mercury — is silent on the nickel, aluminum, tin, and

hexavalent chromium that regulators from Brussels to Beijing to California already limit or list. A four-metal mark is therefore an uncovered exposure in its own right: a blind spot the brand pays for without being told it is there. A credible program is one whose scope tracks where the regulators already stand, not where they stood five years ago. That is why the standard underlying this analysis certifies against eight metals, expressed as ten analytes once the species that matter are split out^[8]. Reading the leading indicator is the whole point; the lawsuit is the lagging one.

The budget line moved

Here is the briefing in one line.

Before February 24, 2026, certification competed for a share of your marketing budget, and it lost as often as it won — because a marketing dollar is supposed to bring back a customer, and this one only brought back a claim on a label. After February 24, it competes for a share of your risk reserve, and there it wins, because it is the only entry on that page whose cost you set yourself, in advance, and small.

Munger told a generation of executives to invert the question: do not ask how to succeed, ask how to fail, and then decline to go there. The uncertified path to failure in this category is now drawn on a map. A state courtroom. An anchor retailer beside you. A plaintiff holding your own missing test record as his best exhibit. A carrier repricing your premium upward. A buyer rewriting the shelf without you. You do not have to be brilliant to avoid that path; you have to be, in Munger's phrase, consistently not stupid — and that has a price, and the price is a rounding error against the cost of being the uncertified defendant when the tide goes out.

Certification is now a litigation-cost-reduction expense, a channel-access expense, and an insurance-premium expense. It is not a marketing expense. *Palmquist* raised the value of all three at once and retired the fourth. The budget line did not get bigger. It moved — from the page where you spend to be noticed to the page where you spend not to be ruined.

The dollar figures in this briefing are illustrative models built on published anchors, not predictions; every range is one a brand can re-run with its own inputs.

Karen Pendergrass is the Standards Architect of the Heavy Metal Tested & Certified program at the Paleo Foundation. She can be reached at karen@paleofoundation.com. This briefing is informational and is not legal advice; nothing in it should be read as a representation about the outcome of any specific matter.

References

[1] *Hain Celestial Group, Inc. v. Palmquist*, 607 U.S. ____ (2026) (No. 24-724), decided Feb. 24, 2026; below, *Palmquist v. Hain Celestial Group, Inc.*, No. 23-40197 (5th Cir.).

[2] U.S. Food and Drug Administration, Total Diet Study, FY2018–2020 (toxic-element detection frequencies).

[3] Federal Judicial Center / U.S. Courts, “Litigation Cost Survey of Major Companies,” 2008 (mean outside legal fees approximately \$2.0M per major case; discovery costs frequently exceeding \$600,000).

[4] Consent Judgment, *People v. Promix Nutrition*, Proposition 65 settlement, California Office of the Attorney General (\$28,000 total; \$25,000 attorneys' fees and related costs; \$3,000 civil penalty).

- [5] U.S. House of Representatives, Committee on Oversight and Reform, Subcommittee on Economic and Consumer Policy, “Baby Foods Are Tainted with Dangerous Levels of Arsenic, Lead, Cadmium, and Mercury,” Staff Report, Feb. 4, 2021.
- [6] California Assembly Bill 899 (2023), baby-food heavy-metal testing and public disclosure; California Attorney General enforcement advisory, Mar. 6, 2026.
- [7] International Organization for Standardization, *ISO/IEC 17025:2017, General Requirements for the Competence of Testing and Calibration Laboratories*.
- [8] K. Pendergrass, *HMTc Infant and Child Foods Program Manual*, 2026 Edition, the Paleo Foundation, 2026. doi: [10.5281/zenodo.20270512](https://doi.org/10.5281/zenodo.20270512).
- [9] K. Pendergrass, “The Cost of Operating Without Credible Third-Party Heavy-Metal Certification,” *Journal of Food Metallomics*, 2026. doi: [10.5281/zenodo.18903738](https://doi.org/10.5281/zenodo.18903738). Recall-economics and market-impact anchors (GMA/Covington/EY; Allianz; Thomsen & McKenzie; Pozo & Schroeder) are documented there; this briefing updates the argument for the post-*Palmquist* forum.