

# The Half-Trace

*One term, many names, no home.*

Michael P. Laing

`mllaing@post.harvard.edu`

ORCID 0009-0004-5593-5759

Written with Anthropic Claude Fable 5.1

August 2026; revised September 28, 2026

*Companion essay to “Measure-Relational Algebra: The Core” (September 2026; doi:10.5281/zenodo.22999911). Canonical reference: realql.org.*

There is a term that applied mathematics cannot stop discovering.

Take any smooth calculation  $f$  and feed it an uncertain input — a measurement, an estimate, anything known only up to a distribution with covariance  $\Sigma$ . Then the answer you get by computing  $f$  at the best single value is not the best single value of the answer. The two differ, to second order, by

$$\frac{1}{2} \operatorname{tr}(\nabla^2 f \Sigma) \text{ — half the curvature of the calculation, weighted by the uncertainty of the input.}$$

That is the whole object: one half, times curvature, times spread. If the calculation is a straight line, the curvature is zero and the term vanishes — averaging and calculating commute, and no harm is done by pretending the input was exact. The moment the calculation bends, the term is nonzero, and it has three properties that make it dangerous. It has a *sign*, fixed by the direction of the bend. It *grows* with the uncertainty you discarded. And it is *immune to averaging* — more data does not shrink it, more data polishes it into something that looks exactly like an empirical law.

What follows is a short history of this term being found, named, domesticated, and forgotten — field by field, over nearly three centuries — and an argument about where it should finally live.

## A tour of the rediscoveries

**Astronomy and geodesy found it first**, because they were the first to compute seriously with imperfect measurements. The French astronomer Nicolas-Louis de La Caille was propagating observational error through spherical trigonometry in 1741; Lambert did it systematically in the 1760s; and Gauss, in his 1823 *Theoria Combinationis Observationum*, gave the general first-order rule for how error moves through a function — the founding document of what every physics student learns as *propagation of error*. The second-order term — our term, the bias — rides one Taylor order behind the variance and has been rediscovered alongside it ever since.

**Psychometrics found it in 1904,** one notch before Jensen. Charles Spearman noticed that measurement error in an input doesn't just blur a correlation, it *attenuates* it — pulls it toward zero, with a sign — and published the correction. Epidemiology reuses this discovery to this day under the name *regression dilution*: measure blood pressure noisily, and the estimated effect of blood pressure on mortality is systematically too small. The correction factor is taught, when it is taught, as a specialist's repair.

**Mathematics named the sign in 1906.** Johan Jensen proved that for convex  $f$ , the expectation of the function exceeds the function of the expectation — always. The difference  $\mathbb{E}[f(X)] - f(\mathbb{E}[X])$  is now called the *Jensen gap*, and a modern literature computes sharp bounds for it. Jensen's inequality is arguably the most load-bearing elementary inequality in analysis; it is also, read operationally, a statement about data processing: *plugging in the average is wrong in a known direction*.

**Statistics industrialized it and then lost the receipt.** The technique of expanding a statistic around its mean to extract variance and bias became the *delta method* — used by Kelley in 1928, formalized by Doob in 1935 — and so thoroughly absorbed into practice that its own paternity became a research question: the field required two forensic papers, “Who Invented the Delta Method?” (Ver Hoef, 2012) and “Who Invented the Delta Method, Really?” (Gorroochurn, 2020), to trace the lineage back through Gauss to La Caille. A tool can hardly be more domesticated than that: everyone uses it, and nobody remembers where it came from.

**Stochastic analysis found that the term never sleeps.** Itô's lemma (1944) — the fundamental theorem of stochastic calculus — is ordinary calculus plus exactly one new piece: a  $\frac{1}{2}\sigma^2 f''$  correction that appears because a diffusing quantity is never a point, even over an instant. The entire difference between the calculus of certainty and the calculus of uncertainty *is this term*, applied continuously in time.

**Finance found it at least three times, and pays it daily.** The mean of a lognormal exceeds the exponential of the mean by a factor  $e^{\sigma^2/2}$  — the  $\frac{1}{2}\sigma^2$  that sits inside the Black–Scholes drift. Portfolio managers meet it as *volatility drag*: the geometric return trails the arithmetic return by about half the variance, a gap that compounds into real money. And fixed-income desks quote a *convexity adjustment* — one half, times the second derivative of price in yield, times the squared move — as a matter of trading routine; it is priced, hedged, and printed in CFA study guides. Finance is the one field that not only found the term but put it on the invoice.

**Metrology wrote it into law.** The *Guide to the Expression of Uncertainty in Measurement* — the international standard governing calibration laboratories — mandates first-order propagation and prescribes the second-order terms when curvature is significant. The half-trace is, quite literally, codified best practice for anyone certifying an instrument.

**Machine learning runs on it.** The gap between the log-likelihood and the evidence lower bound (ELBO) — the objective actually optimized in variational inference and variational autoencoders — *is* a Jensen gap, and tightening it is an active research industry. Elsewhere, “second-order” extensions of belief propagation push delta-method means and covariances through inference networks. The field's daily bread is baked with the term, under yet another local name.

And the nearest thing to a general theory came from an unexpected direction: **Nicolas Bouleau’s error calculus**, built on Dirichlet forms, which propagates uncertainty through computations with two operators — one carrying variances, one carrying biases — where the bias recursion is precisely first-derivative-times-old-bias *plus half the second derivative times the variance*. Our term, elevated to an operator calculus, with Gauss correctly credited for the first-order part. Bouleau even framed it as a *language* for error. It is the closest ancestor of the position taken here, and it deserves to be far better known.

## The pattern

Eight fields, three centuries, one term. Now notice what every one of these discoveries has in common.

Each is *analyst-side*. Astronomer, psychometrician, statistician, trader, metrologist, machine-learning researcher: in every case, the correction is applied at the point of analysis, by a person who already suspects the problem, using uncertainty information they had to fetch from somewhere outside the data itself. Even Bouleau’s calculus — the most systematic treatment ever given — is mathematics offered *to* the analyst. The half-trace has been discovered by every field that computes with measurements, and domiciled by none of them in the place the measurements actually live.

Meanwhile the record systems underneath all of these fields — the databases, the file formats, the tables — store points. They have no slot for  $\Sigma$ . So the term’s input is discarded at ingestion, and each field’s hard-won correction survives only as professional folklore: applied where a vigilant expert happens to stand, silently omitted everywhere else.

What happens where no expert stands? Space physics just told us. For forty years, the apparent saturation of Earth’s geomagnetic response to the solar wind was accepted as physics and explained by at least nine theoretical mechanisms. In 2026, Sivadas and colleagues showed in *Nature* that the observations give no statistical evidence for the ceiling: the curve is reproduced, quantitatively, by the half-trace’s signature alone — uncertain upstream measurements, stored as exact values, pushed through a nonlinear analysis; regression dilution at planetary scale. Undo the collapse and the relationship is linear throughout the observed range, and the implied impact of an extreme solar storm roughly doubles. Nobody was careless. The uncertainty was documented. The correction was a century old. There was simply no structural connection between the two, and vigilance did not happen to stand at that spot for forty years.

## Where the term should live

The accompanying paper adds two things to this long history, and it is worth being precise about how little — and how much — they are.

The first is a boundary. The folklore says nonlinearity causes trouble; the theorem says exactly where trouble is impossible: collapsing a distribution to its mean commutes with a calculation, *for every possible dependence among the inputs*, if and only if the calculation is affine (a straight-line combination — sums, scalings, weighted averages) — or the value was exact to begin with. Not approximately safe, not usually safe: the precise and complete extent

of the free region. Everything outside it is a categorically different computation, and the half-trace is the leading term of the difference.

The second is a change of address. If the safe region is exactly the affine one, then a data system can *know*, from the structure of a query alone, where collapse is licensed. Means and covariances move through affine steps exactly — no approximation, closed form — and arrive at each curved step as precisely the  $\Sigma$  that the half-trace needs. The term stops being a correction a suspicious analyst applies afterward and becomes a **price list the system computes by default**: this query, on this data, at this collapse point, costs this much bias, in this direction. Spearman’s attenuation, the trader’s convexity adjustment, the metrologist’s second-order terms, and the forty-year geomagnetic artifact are then not seven cousins in seven dialects but one audited quantity, evaluated where the data lives.

Three centuries of rediscovery is not a story about a difficult idea. Half, times curvature, times spread — it fits on a napkin. It is a story about an idea with no home: every generation of every quantitative field walks into the same doorframe, writes a local memo about the doorframe, and moves on. The door does not need another memo. It needs to be moved — out of the analyst’s vigilance and into the structure of the record, where what we know about a number can finally travel with the number itself.

---

*The paper: M. P. Laing, “Measure-Relational Algebra: The Core,” [realql.org](https://realql.org), September 2026, <https://doi.org/10.5281/zenodo.22999911>. The motivating case: N. Sivadas et al., *Nature* **655**, 1143–1147 (2026). On the delta method’s lineage: J. M. Ver Hoef, *Am. Stat.* **66**, 124–127 (2012); P. Gorroochurn, *Math. Intelligencer* **42**, 46–49 (2020). On error calculus: N. Bouleau, *Error Calculus for Finance and Physics: The Language of Dirichlet Forms* (de Gruyter, 2003).*