

The Committee That Never Disagrees

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 — draft 4

Companion essay to “What a Memory Must Preserve: Machine Reasoning over Uncertain Observations” (September 2026) and to “Measure-Relational Algebra: The Core” (September 2026; doi:10.5281/zenodo.22999911). Canonical reference: realql.org.

The ox

In 1907, Francis Galton reported an experiment from a livestock fair. Nearly eight hundred visitors had paid sixpence each to guess the weight of an ox. Almost none of them was right. The *average* of their guesses missed the true weight by less than one percent — closer than nearly every individual guess.

The result is famous as the wisdom of crowds, and it is real. It is also conditional, and the condition is the entire story. The crowd is wise because its members are wrong *independently*. Each guess is truth plus a private error, and eight hundred private errors, pointing every which way, cancel. Double the crowd and the residual error shrinks again. The arithmetic is merciless and benign: average enough independent guesses and you can get as close to the truth as you like.

Now change one thing. Let the fairgoers, before guessing, all consult the same confident acquaintance — a man who has looked the ox over and formed an opinion. Each visitor still adds private judgment, still differs from the others, and the crowd still looks pleasingly diverse. But every guess now contains a common ingredient: the acquaintance’s error. Private errors still cancel. The shared error does not. It is the same error, in every guess, and averaging identical things reproduces them exactly. Grow the crowd to eight thousand, eight million: the residue never shrinks below the one mistake everybody inherited. The crowd converges — confidently, reproducibly, with beautiful statistics — on the acquaintance.

Every organization that has sat through a unanimous meeting knows this failure by name. Groupthink is not people being stupid; it is people being coupled. A committee whose members read the same briefing, defer to the same authority, or watch each other’s faces is a committee whose votes share an ingredient, and no enlargement of the committee dilutes it. The vote count grows; the information does not. Political scientists proved the formal version long ago: the celebrated theorem that says large juries converge on the truth requires independent jurors,

and can fail — quantifiably — when the jurors are correlated. A committee that never disagrees is not a committee that is certainly right. It is usually a committee with one member.

Hold that shape. We are about to build it into everything.

A machine that studies the minutes

Start with the honest, boring fact underneath all record-keeping: a measurement is a state of knowledge, not a number. The instrument reports a value *and* an error structure; the record, almost everywhere, keeps the value and discards the structure. The accompanying mathematical work proves exactly when that is free — when everything downstream is linear, sums and averages and scalings — and exactly when it is not: everywhere else, computing on the collapsed record is not an approximation of the correct computation but a different one, capable of manufacturing structure the world does not contain. The showcase is a 2026 *Nature* result: forty years of apparent “saturation” in Earth’s magnetic response to solar storms turned out to be reproducible by observational uncertainty alone — the data providing no statistical evidence that the ceiling is physical. Correct for the uncertainty the record never carried and the ceiling vanishes — and the estimated impact of an extreme storm roughly doubles.

Now put a learning machine in front of that record, because that is what we now do with records. Here is the uncomfortable theorem, in plain words. A machine trained to predict outcomes from recorded values does exactly what it is built to do: it converges on the best possible summary of the record. If the record’s numbers were written down as exact when they were not, that best possible summary *is the artifact*. In a stripped-down model of the geomagnetic case one can write it in closed form: the function the learner is aiming at is precisely a drooping, saturating curve — a ceiling that is not there. Not approximately, not as a small-sample accident. The artifact is the target.

This inverts the usual comfort. More data does not fix it; more data means the machine finds the artifact faster and holds it more firmly. Better architecture does not fix it; a better learner is one that gets to the target more efficiently, and the target is wrong. The machine is the model student of the committee minutes: diligent, consistent, increasingly precise — about what the minutes say. The error was not made by the machine. It was made by the record, at the moment of writing, and the machine’s virtue is to preserve it perfectly.

The classical defense against this was a person: the statistician who suspects the data, reaches outside the record for the error model, and applies a century-old correction. That defense was always artisanal. The training pipeline is its industrial replacement, and the pipeline’s defining property is that it does *not* suspect the record — the record is the interface; fitting it faithfully is the design goal. Where the vigilant analyst was an unreliable defense, the pipeline is a reliable non-defense.

One acquaintance, consulted by everyone

So far this is one machine and one record. The reason it belongs in the same essay as Galton’s ox is what happens at deployment scale — and here the reader should know the register changes: what follows is a pattern, laid out plainly because its shape is already visible, not a body of finished results.

A handful of very large models are becoming the shared consultant of the developed world. The same model reads scans in hospital chains, flags defects in chip foundries, screens transactions in banks, summarizes intelligence for governments, imputes the missing fields in a thousand corporate records. Each institution, individually, is behaving sensibly: the model is capable, the alternative is more expensive, the outputs look good. But every one of those outputs carries a common ingredient — the model’s own error, its particular blind spots, its characteristic way of being wrong — and that ingredient is now the confident acquaintance at the fair, consulted by everyone, cited by no one.

Two consequences follow from the shape, and both are invisible from inside any single institution.

First, **averaging stops buying truth**. Aggregate a thousand model-derived estimates — across branches, across hospitals, across an industry — and the private noise cancels while the shared component stands. The statistics look better and better; the floor under them never moves. A million records produced through one model are, in the part that matters, a million copies of the same ignorance.

Second, **institutions that never exchanged a byte can be correlated**. Hospital and reinsurer, foundry and fund: if both consume the same model’s outputs, their errors are coupled, and any calculation that pools across them assuming independence — risk aggregation, meta-analysis, redundancy engineering — is computing with a fiction. Engineering redundancy is the starkest case, because redundancy is precisely a *purchase of independence*: three sensors instead of one, two pilots, dissimilar backup systems, on the theory that they will not all fail the same way at the same moment. Route the three sensor feeds through one perception model and the purchase is silently refunded. The three opinions were one opinion wearing three uniforms. Finance has already run this experiment once at scale: in 2008, institutions that believed themselves diversified found that their risk assessments had leaned on shared upstream models, and the correlations nobody had recorded arrived all at once. That episode is offered here as a shape, not a proof — but it is the shape.

The starkest case has lately acquired a name. Frontier laboratories now dispatch thousands of agents — copies of one model — to work a problem together, and the people closest to the practice describe it as oversight lost in complexity. The complexity is real but it is not the problem. A thousand copies of one model are the committee in its purest form: a thousand voices, one acquaintance. They can vote, cross-check, and converge; what they converge on is the model. Ask the swarm whether its own work is safe and you have asked one judge a thousand times, and the record of its answers does not carry the number — how much of the error is shared — that would tell you so.

There is an exchange rate, and it is worth knowing. One judge whose errors are independent of the model’s — a person who has not read the swarm’s summary, a second model raised on different data, an instrument pointed at the world — is worth more than the whole swarm the moment that judge’s own error is no larger than the shared part of the model’s. Not more useful in some vague sense: more, in the arithmetic of variances, than an unbounded number of copies. That is what “a human in the loop” buys, and it states the condition on which the purchase depends. A human who reads the machine’s summary first has joined the committee.

And the loop closes. Model outputs are written back into records as plain numbers — a prediction stored as a fact, an imputed value indistinguishable from a measured one — and the next generation of models trains on records partly written by the last. The committee does

not merely never disagree; it takes its own minutes, and studies them. Where that leads is now worked out: if the minutes keep the original outcomes, every generation re-learns the first generation’s mistake exactly; if each generation’s minutes are the last generation’s conclusions, the committee forgets, one generation at a time, that the outcome ever depended on the evidence — and its own records report, truthfully, that it has never agreed with itself more.

What the record would have to carry

None of this is an argument against using powerful models, any more than Galton’s fair is an argument against crowds. The crowd was *wise*; the mathematics of when it is wise is exact. The argument is about what the record keeps.

Everything above has a single root: the coupling is real, is often knowable at the moment of writing — *these ten thousand cells passed through model M, version v* — and is stored nowhere. Data lineage will not find it: two records can be coupled through a shared model while sharing no data at all, the way two fairgoers share an acquaintance without sharing a guess. What would find it is a record that carries, alongside each value, the state of knowledge it came from and the *cause* it came through — the instrument, the calibration, the model version — so that the shared ingredient is a first-class, queryable fact rather than an ambient condition. Then the arithmetic that today silently assumes independence can instead *check*; the aggregation that today reports a vanishing error bar can instead report the floor; and when the model is updated — the acquaintance changes his mind — the records that inherited his old opinion are findable, because his signature is on them.

The mathematical framework for records of this kind exists — that is what the accompanying papers construct and pose: values stored as states of knowledge, dependence carried explicitly, its causes proposed as first-class entities, collapse to a single number an explicit operation rather than a silent default, and the exact theorem for when the shortcut is free. What does not yet exist is the engineering, and it is offered openly as work, under CC BY, with no patents and nothing to license.

The claim, plainly

A crowd is wise only insofar as its members are wrong independently, and we are quietly rebuilding civilization’s record-keeping around a small number of shared consultants whose errors everyone inherits and no one writes down. Each machine trained on such records converges, diligently and precisely, on whatever the records got wrong; each institution consuming the same model is coupled to every other through whatever errors the model shares among them, invisibly; and the outputs return to the record as plain numbers, ready to teach the next machine. Every piece of this is the oldest failure of collective judgment — the committee that never disagrees — executed at machine speed, with the one aggravation that a human committee can sometimes feel its own conformity, and a record system cannot feel anything it was never built to represent.

The repair is not vigilance, which does not scale, and not abstinence, which will not happen. It is structure: records that keep what we actually know — the uncertainty, and the shared causes running through it — so that independence becomes something a calculation can verify

instead of something it assumes. Crowds made us wise for a century where measurements were cheap and errors were private. Whether the next century's crowds of machines make us wise or merely unanimous depends on a decision that looks, from the outside, like a boring question of data formats. Most load-bearing decisions do.

*The papers: M. P. Laing, “Measure-Relational Algebra: The Core,” realql.org, September 2026 (doi:10.5281/zenodo.22999911), and “What a Memory Must Preserve: Machine Reasoning over Uncertain Observations,” realql.org, September 2026. The motivating case: N. Sivadas et al., “Regression to the mean can explain saturation of geomagnetic storms,” *Nature* **655**, 1143–1147 (2026). Galton’s account: “Vox Populi,” *Nature* **75**, 450–451 (1907).*