

Uncertainty: Naming the Sources Before Adding Them

IDE-017 · Fundamental Ideas

Concept essays for *The Universe and Its Observers*

Author: Juan Segurolo

Version: v0.5 EN · 19 September 2026

Status: Publication version

Article form: Concept essay

Scale / context: From instrument budgets to policy-relevant projection

Boundary: Uncertainty as a set of sources and how they are separated, propagated and communicated. Probability as a calculus belongs to IDE-008; the metrological uncertainty budget of a measurement result belongs to IDE-014

Related works: IDE-018, IDE-008, IDE-014

Epistemic profile: Established knowledge with interpretive boundaries

Licence at publication: All Rights Reserved

Article Abstract

Uncertainty is not one thing with one number attached. A quantity can be uncertain because the instrument scatters, because a parameter was estimated from limited data, because the model omits a mechanism, or because the future course of events is not determined. These sources behave differently: some shrink with more data and some do not, and combining them into a single interval hides which is which. This article separates the sources, argues that the familiar division into aleatory and epistemic depends on where the modeller drew the boundary, shows that parameter conclusions are conditional on the adopted model and its assumptions, and takes seriously cases where quantification is unavailable.[3,4]

In 60 Seconds

Start by asking what would reduce the uncertainty. If more measurements of the same kind would narrow it, it behaves like an estimation problem; if they would not, something structural is responsible. That test does most of the work the aleatory and epistemic labels are supposed to do, and it avoids their main trap: whether variability counts as inherent randomness or as missing knowledge depends on the model chosen, not on the world, so the same quantity can be classified either way by two competent analysts.[2] A second warning follows. Sensitivity and parameter conclusions are conditional on the adopted model and its assumptions.[3] Finally, not everything is quantifiable. Knight distinguished measurable risk from uncertainty that cannot be reduced to a distribution, and pretending otherwise produces numbers with no warrant.[4,5]

Why This Question Matters

Uncertainty is where a scientific result meets a decision, and the way it is reported determines what the decision can be. A single interval invites the reader to treat all the doubt as one kind, which hides the

two questions that matter for action: would more of the same evidence help, and what would have to be wrong for the answer to change. It also makes a common failure invisible, in which a precisely computed interval around a parameter is offered as though it bounded the error in the conclusion, when the dominant uncertainty lies in the structure that produced it.

1. Sources Behave Differently

It is worth listing the sources separately because they respond to different remedies. Measurement uncertainty comes from calibration, resolution, environmental effects and the measurement model, and is the one with the most developed formal treatment.[1] Parameter uncertainty comes from estimating values from finite data. Structural uncertainty comes from the choice of model form and from mechanisms left out. Scenario uncertainty comes from conditions that are not determined at all, such as future choices.

The useful diagnostic is what would reduce each. More observations of the same kind reduce parameter uncertainty and, up to the definitional floor, measurement uncertainty. They do not reduce structural uncertainty, which requires a different model, and they do nothing at all about scenario uncertainty, which is not a state of ignorance but a genuine openness.

Reporting them together as one interval is therefore not conservative; it is uninformative. A result given as a value plus or minus a number invites the reader to assume that the number bounds the doubt, when it may bound only the most tractable component of it.

Where That Leaves Us

Measurement, parameter, structural and scenario uncertainty respond to different remedies, and the diagnostic question is what would reduce each. Collapsing them into one interval hides which component dominates. The most common attempt to organize them is a two-way split, and that split turns out to be less absolute than it first appears.

2. Aleatory and Epistemic, and Why the Line Moves

The standard distinction separates aleatory uncertainty, treated as inherent variability, from epistemic uncertainty, treated as incomplete knowledge that could in principle be reduced. It is a useful working split, and engineering practice depends on it because the two propagate differently through a reliability analysis.

Der Kiureghian and Ditlevsen argue that the distinction is a feature of the model rather than of nature. A variable treated as randomly varying in one model can be represented explicitly, and hence become epistemic, in a more detailed one; the classification therefore depends on the boundary the analyst drew.[2] What matters for a decision, they argue, is not the label but whether the uncertainty can be reduced within the time and resources available.

That reframing is more useful than the labels. Asking whether a component is reducible in practice, and at what cost, produces an answer that guides action, whereas asking whether it is aleatory in principle produces a dispute about modelling conventions.

Where That Leaves Us

The aleatory and epistemic split is a property of the model rather than of the world, since a variable treated as random in one description can be represented explicitly in a finer one. The question that survives is whether a component is reducible in practice and at what cost. A related and more dangerous confusion concerns which component dominates in the first place.

3. Parameter Precision Is Conditional on the Model

Sensitivity and parameter conclusions are conditional on the adopted model and its assumptions.
[3]

Structural uncertainty is harder to quantify and often larger. Inference: under misspecification, a fitted parameter can compensate for an omitted mechanism, so a narrow interval need not imply an adequate model.

Global sensitivity analysis is the standard tool for finding out what the answer depends on. Rather than varying one input at a time around a nominal point, it explores the input space jointly and apportions the variance of the output among the inputs and their interactions, which is what reveals that a result is driven by an assumption nobody was watching.[3]

Where That Leaves Us

Global sensitivity analysis is what identifies which assumptions the answer actually rests on. Beyond all of these lies the case where no distribution is available at all.

4. When Quantification Is Not Available

Knight distinguished measurable risk, where the distribution of outcomes is known or can be estimated, from uncertainty proper, where it cannot.[4] The distinction was made about economic decisions but the structure is general: some situations offer a well-defined reference class and others do not, and a probability computed without one is an assertion rather than a measurement.

This is not an argument against quantifying. It is an argument for saying which case you are in. A reported probability that rests on a defensible frequency or on an explicitly stated model earns its status; one produced because a number was required has the same appearance and none of the warrant. The Bayesian framework makes the assumptions explicit rather than removing them, which is a real gain, since the prior and the model are then visible and can be argued with.[6]

Where quantification is genuinely unavailable, the honest options are to describe the range of possibilities, to identify what would distinguish them, and to look for decisions that perform acceptably across the range rather than optimally under one assumed distribution.

Where That Leaves Us

Knight's distinction between measurable risk and unquantifiable uncertainty is about whether a reference class exists, and a probability produced without one is an assertion rather than a measurement. Making assumptions explicit is a gain even when it does not remove them. What remains is how to say all of this to someone who has to act on it.

5. Saying It Honestly

Funtowicz and Ravetz proposed that a quantitative result carried into policy should be accompanied by more than an error bar: by an account of the quality of the underlying information, including how the numbers were produced and how much of the result rests on judgement.[5] Uncertainty assessment includes the quality and pedigree of the underlying information, not only the width of a numerical interval. Inference: intervals of equal width can therefore rest on evidence of different quality.

Whatever notation is used, three things have to survive into the summary. Which sources of uncertainty were included and which were not; whether the dominant component is reducible by more of the same work; and what would have to be true for the conclusion to change. A result reported without these is not more cautious for having an interval attached.

The temptation runs the other way. Decision makers ask for a number, and an interval that omits structural uncertainty is narrower and looks more useful than one that includes it. Resisting that is the whole of the discipline, and it is easier when the sources were separated at the start rather than aggregated and then defended.

Where That Leaves Us

Uncertainty assessment includes the quality and pedigree of the underlying information, not only the width of a numerical interval. Narrower intervals are more attractive and often less complete. IDE-018 takes up a structural property that determines how uncertainty spreads through a system: how its parts are connected.

What We Know, and How

- Measurement, parameter, structural and scenario uncertainty respond to different remedies, and only some are reduced by more observations of the same kind. That four-way scheme is this article's own organising distinction rather than a result taken from a source.
- The aleatory and epistemic classification depends on the model boundary the analyst drew rather than on nature, so the operationally useful question is whether a component is reducible in practice.[2]
- Sensitivity and parameter conclusions are conditional on the adopted model and its assumptions.[3]
- Global sensitivity analysis explores the input space jointly and apportions output variance among inputs and their interactions, which is how the dominant assumptions are identified.[3]

What We Still Don't Know

- How to characterize and propagate structural uncertainty in a principled way when the set of plausible model structures cannot be enumerated.
- Whether and how deep uncertainty, in Knight's sense, can be represented formally without importing the unwarranted precision that the distinction was drawn to avoid.[4]
- Which uncertainty communication formats actually change decisions for the better, which is an empirical question about users rather than a question about the mathematics.

What Could Change Our Minds

- A demonstration that structural uncertainty in a class of models can be bounded rigorously from data alone, without enumerating alternative structures, would remove the main reason parameter intervals are considered incomplete.[3]
- Evidence that analysts independently classifying the same uncertainties as aleatory or epistemic agree at high rates regardless of their modelling choices would contradict the claim that the distinction is model-relative.[2]
- A worked case in which a probability assigned without any defensible reference class outperformed, over many repetitions, an explicit statement of the range of possibilities, would weaken the argument against quantifying deep uncertainty.[4,6]

If You Remember Only One Thing

Uncertainty has sources that behave differently, and an interval that does not say which were included conceals more than it reports.

Related Articles

- IDE-018 — Connection
- IDE-008 — Probability
- IDE-014 — Measurement

References

- [1] JCGM 100:2008. Evaluation of Measurement Data: Guide to the Expression of Uncertainty in Measurement (GUM). BIPM (2008).
- [2] Armen Der Kiureghian; Ove Ditlevsen. Aleatory or epistemic? Does it matter? Structural Safety 31(2), 105-112 (2009). DOI 10.1016/j.strusafe.2008.06.020.
- [3] Andrea Saltelli; Marco Ratto; Terry Andres; Francesca Campolongo; Jessica Cariboni; Debora Gatelli; Michaela Saisana; Stefano Tarantola. Global Sensitivity Analysis: The Primer. Wiley (2008).
- [4] Frank H. Knight. Risk, Uncertainty and Profit. Houghton Mifflin (1921).
- [5] Silvio O. Funtowicz; Jerome R. Ravetz. Uncertainty and Quality in Science for Policy. Theory and Decision Library A, volume 15. Kluwer Academic Publishers (1990). ISBN 978-0-7923-0799-0.
- [6] Edwin T. Jaynes; edited by G. Larry Bretthorst. Probability Theory: The Logic of Science. Cambridge University Press (2003). ISBN 978-0-521-59271-0.