THE GENERATION OF KNOWLEDGE · EPISTEMOLOGY

# The Limits of Knowledge and the Role of Uncertainty

Why all empirical knowledge is uncertain, the different kinds of uncertainty that researchers face, how people misjudge what they know, and how honest treatment of uncertainty makes knowledge more useful, not less.

ModQu research article · 2,158 words · 23 references (APA 7th edition)

At his trial in 399 BCE, Socrates told the jury that he was wiser than other people in only one respect: he did not think he knew what he did not know (Plato, 1997). Two and a half thousand years later, the point remains one of the most practical lessons in the theory of knowledge. Every piece of empirical knowledge has limits. Measurements have error, samples may not represent the whole, models simplify, and the future may not resemble the past. The question is not whether knowledge is uncertain, but whether its uncertainty is understood, measured where possible and honestly communicated.

This article examines the main sources of limits on knowledge, the different kinds of uncertainty that result, the evidence that people systematically misjudge what they know, and how research organisations can treat uncertainty as part of the knowledge they produce.

## Why empirical knowledge is never certain

### The problem of induction

The deepest limit was identified by David Hume. All our reasoning about matters of fact beyond present observation, he argued, depends on the assumption that the future will resemble the past. But that assumption cannot itself be proved, either by logic, since its denial involves no contradiction, or by experience, since any argument from past experience already assumes it (Hume, 1748/2007). However many white swans we have observed, the next one may be black.

Popper (1959) accepted Hume's argument and drew a radical conclusion: scientific theories can never be verified, only tested and potentially falsified. A theory that survives many serious attempts to refute it is **corroborated**, not proved. On this view, all scientific knowledge is provisional, held until a better theory replaces it. Most philosophers of science today would not go as far as Popper in rejecting inductive support, but almost all accept that empirical knowledge is **fallible**: it can be well justified and still turn out to be wrong.

### The limits of observation and measurement

A second limit concerns measurement. Every measurement involves error, and many important quantities cannot be observed directly at all. Attitudes, intentions, satisfaction and trust are measured through questions whose answers depend on wording, context and the respondent's willingness and ability to report accurately. Groves and Lyberg (2010)

described the **total survey error** framework, which identifies the many points at which errors can enter a survey estimate: in defining the population, building the sampling frame, drawing the sample, achieving responses, measuring the answers and processing the data. Only one of these, sampling error, is captured by the conventional margin of error.

The consequences are measurable. Shirani-Mehr et al. (2018) analysed thousands of pre-election polls in the United States and compared them with actual results. They found that the average total error was about twice as large as the error implied by the polls' reported margins, because of errors that the margins do not include. The polls were not wrong in a careless way; they simply knew less than their margins suggested.

### The limits of published research

A third limit concerns the body of research itself. Ioannidis (2005) argued that, given the small samples, flexible analyses and incentives for positive results common in many fields, a large share of published research findings are likely to be false. The Open Science Collaboration (2015) put this to the test by repeating 100 published psychology experiments. Although 97% of the original studies had reported statistically significant results, only 36% of the replications did, and the replication effects were on average about half the size of the originals. Knowledge that appears settled in the literature can be much less secure than it looks.

### The limits of prediction

A fourth limit concerns prediction. Even when the rules governing a system are known exactly, its future may be unknowable in practice. The meteorologist Edward Lorenz found that a simple model of atmospheric convection produced wildly different results when he restarted a run from numbers rounded to three decimal places instead of six (Lorenz, 1963). Tiny differences in starting conditions grew until the forecasts bore no relation to each other. This **sensitive dependence on initial conditions**, later popularised as the "butterfly effect", places a hard limit on how far ahead some systems can be forecast, however good the data and the model.

Social systems add a further complication: people respond to knowledge about themselves. Merton (1948) described the **self-fulfilling prophecy**, in which a false belief leads people to act in ways that make it come true, as when rumours of a bank's insolvency cause a run that makes the bank insolvent. The reverse also happens: a forecast of a shortage can prompt action that prevents it. Published research can change the behaviour it describes. An election poll that shows one candidate far ahead may affect turnout; a widely reported consumer trend may be amplified or reversed by the companies that read about it. Knowledge of social life is therefore always at risk of being overtaken by its own effects.

Taleb (2007) drew attention to a third problem: rare, high-impact events that lie outside normal expectations, which he called **black swans**. Because such events are rare, past data contain few examples of them, and models fitted to past data tend to underestimate their likelihood. After they happen, they are often explained as if they had been predictable, which hides how little was actually known beforehand. Taleb's argument is a modern form of Hume's: the past is an imperfect guide to the future, and least reliable precisely where the consequences are greatest.

## Kinds of uncertainty

Not all uncertainty is the same, and treating it as if it were leads to mistakes.

## Risk and uncertainty

The economist Frank Knight drew a famous distinction between **risk** and **uncertainty** (Knight, 1921). Risk applies when the possible outcomes and their probabilities are known, as with the roll of a fair die or the mortality rates used by insurers. Uncertainty, in Knight's sense, applies when the probabilities themselves cannot be known, because the situation is new or unique. Keynes (1921), writing in the same year, argued that many probabilities cannot be measured numerically at all.

Ellsberg (1961) showed that people treat these situations differently. In his experiments, most people preferred to bet on drawing a ball from an urn with a known mix of colours than from an urn with an unknown mix, even when the bets were logically equivalent. This **ambiguity aversion** shows that people distinguish between knowing the odds and not knowing them, even if formal decision theory often treats the two alike.

## Levels of uncertainty

Walker et al. (2003) proposed a scale of uncertainty for decision-makers who use models, running from complete certainty to total ignorance (Figure 1). **Statistical uncertainty** can be described with probabilities and confidence intervals. **Scenario uncertainty** applies when the range of possible outcomes can be described but not their probabilities. **Recognised ignorance** applies when we know that we do not understand the mechanisms involved. **Total ignorance** applies when we do not even know what we do not know. Smithson (1989) had earlier argued that ignorance itself has several forms, from simple error to topics that are deliberately treated as off-limits.

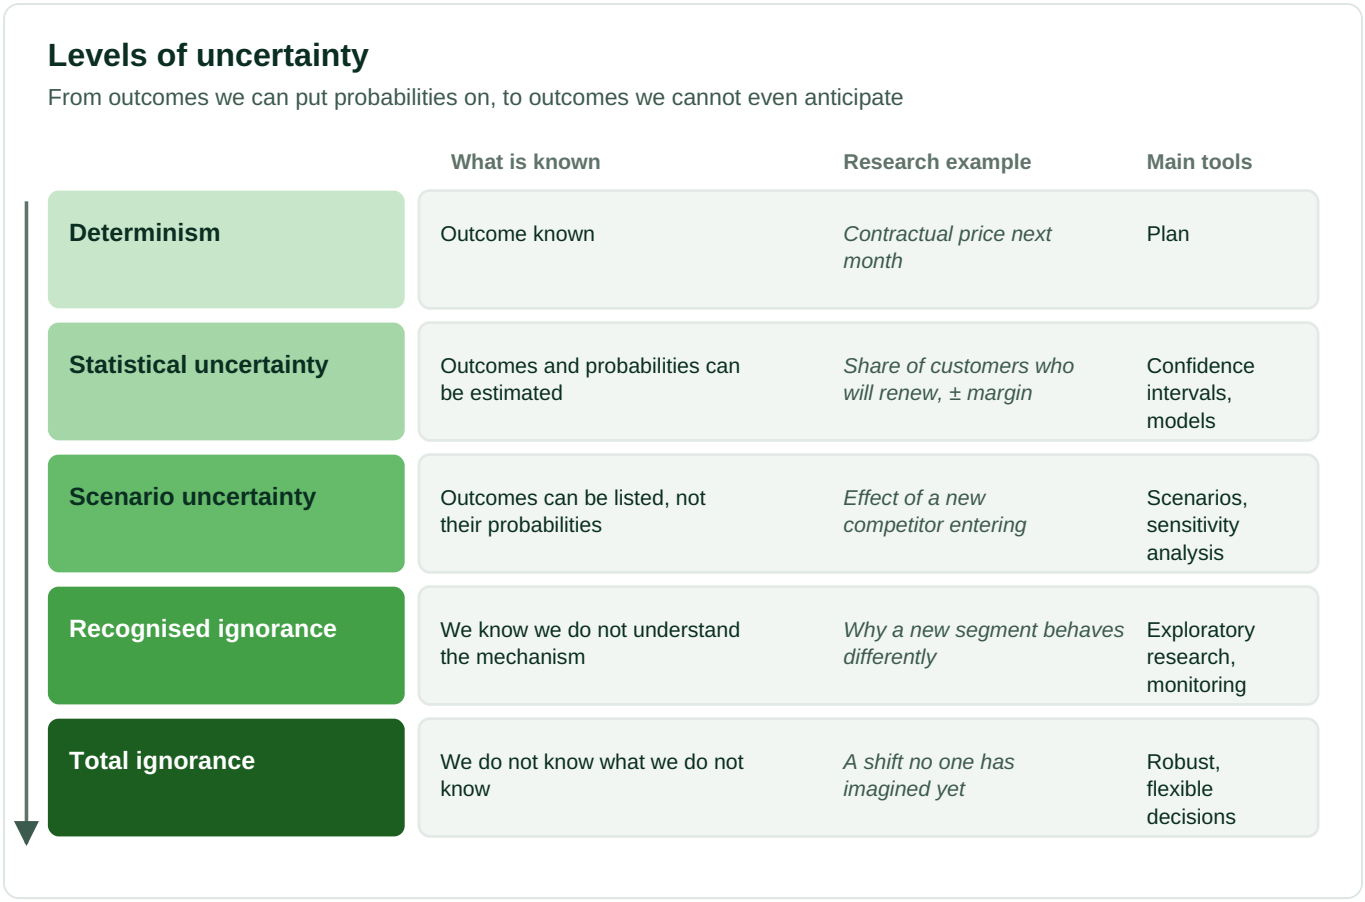



Figure 1. Levels of uncertainty, from outcomes that can be described with probabilities to those that cannot even be anticipated. Each level calls for different tools. Adapted from Walker et al. (2003); examples and diagram by ModQu.

Spiegelhalter and Riesch (2011) made a related distinction for risk analysis. Some uncertainty concerns outcomes within a model; some concerns the model's parameters; some concerns whether the model is adequate at all; and some concerns the possibility of events that no model has considered. They argued that analysts should acknowledge these deeper uncertainties explicitly, for example by describing the quality of the evidence behind a model, rather than hiding them behind precise-looking numbers.

## When the stakes are high and the facts uncertain

Funtowicz and Ravetz (1993) argued that some questions, such as environmental and health risks, combine high decision stakes with high uncertainty about the facts and disputed values. For these questions, which they called **post-normal science**, traditional expert methods are not enough. Quality assurance must include a wider range of perspectives, and the uncertainty itself must be openly discussed.

## How people misjudge what they know

Uncertainty would be easier to manage if people judged it accurately. Psychological research shows that they often do not.

Tversky and Kahneman (1974) showed that people judging uncertain events rely on mental shortcuts, such as how easily examples come to mind or how closely a case resembles a stereotype, and that these shortcuts produce predictable errors. Moore and Healy (2008) distinguished three forms of overconfidence: overestimating one's own performance, believing one is better than others, and being too sure that one's beliefs are accurate. The last, which they called **overprecision**, is the most persistent. When people are asked to give a range that they are 90% sure contains a true value, the true value often falls outside the range far more than 10% of the time.

Kruger and Dunning (1999) found that people with the least skill in a domain tended to overestimate their performance the most, partly because the skills needed to perform well are also needed to recognise good performance. Tetlock (2005) followed hundreds of experts making thousands of political and economic forecasts over many years and found that many performed little better than simple rules, and that experts who were most confident in a single big idea tended to be the least accurate.

These findings matter because the limits of knowledge are often invisible from the inside. People and organisations rarely feel uncertain about things they are wrong about. Recognising the limits of knowledge requires deliberate methods, not just good intentions.

## Communicating uncertainty

One common worry is that admitting uncertainty will undermine trust or make findings less useful. Research suggests the reverse is often true. Van der Bles et al. (2019) reviewed how uncertainty about facts and numbers can be communicated and proposed a framework covering what is uncertain, in what form the uncertainty is expressed, and how its source is described.They argued that uncertainty should be communicated as a matter of course, and that the form chosen, whether a range, a verbal qualifier or a description of the quality of evidence, should suit the audience and the decision.

Manski (2015) argued that official economic statistics, such as estimates of growth and employment, are routinely reported as if they were exact, even though they are revised substantially later and are subject to errors beyond sampling. He called for statistical agencies to report the uncertainty in their figures systematically, so that users would not mistake estimates for facts.

## Working with the limits of knowledge

For organisations that produce and use research, accepting the limits of knowledge leads to practical disciplines.

**Report all the error you can, not only sampling error.** A margin of error describes only one kind of uncertainty. A credible report also discusses coverage, nonresponse and measurement, and states plainly which of these could not be assessed.

**Match the method to the level of uncertainty.** Statistical tools work well for statistical uncertainty. For scenario uncertainty, scenario planning and sensitivity analysis are more honest than a single forecast. For recognised ignorance, the right response may be further research, monitoring or decisions that keep options open.

**Calibrate judgements.** Forecasters and analysts can be trained to give ranges that contain the truth as often as they claim. Keeping score of past judgements, and comparing stated confidence with actual accuracy, is the most direct way to find and correct overprecision.

**Prefer replicated and triangulated findings.** Given the evidence on replication, a single study should rarely be treated as settled knowledge. Findings that have been repeated, or that are supported by several different methods, deserve more confidence.

**Plan for surprises.** Total ignorance cannot be eliminated, but its effects can be reduced. Decisions that are robust to a range of possible futures, monitoring systems that detect unexpected changes early, and a willingness to revise conclusions quickly all help organisations cope with what they did not know they did not know.

**Communicate uncertainty clearly and without apology.** A finding presented with an honest range, a clear statement of its main limitations and a description of the evidence behind it is more useful to a decision-maker than a precise number that may mislead. Uncertainty is information. Leaving it out does not make knowledge more certain; it only makes it less reliable.

## Conclusion

Socrates's claim to wisdom was not that he knew a great deal, but that he knew the limits of what he knew. Modern research gives that insight practical form. Hume and Popper showed that empirical knowledge is fallible; survey methodologists and statisticians have shown how to measure some of its error; psychologists have shown how easily people overrate what they know; and decision scientists have mapped the different kinds of uncertainty that remain. Knowledge is not the absence of uncertainty. It is the best available understanding, together with an honest account of how far it can be trusted.

## References

Ellsberg, D. (1961). Risk, ambiguity, and the Savage axioms. *The Quarterly Journal of Economics*, *75*(4), 643–669. https://doi.org/10.2307/1884324

Funtowicz, S. O., & Ravetz, J. R. (1993). Science for the post-normal age. *Futures*, *25*(7), 739–755. https://doi.org/10.1016/0016-3287(93)90022-L

Groves, R. M., & Lyberg, L. (2010). Total survey error: Past, present, and future. *Public Opinion Quarterly*, *74*(5), 849–879. https://doi.org/10.1093/poq/nfq065

Hume, D. (2007). *An enquiry concerning human understanding* (P. Millican, Ed.). Oxford University Press. (Original work published 1748)

Ioannidis, J. P. A. (2005). Why most published research findings are false. *PLoS Medicine*, *2*(8), Article e124. https://doi.org/10.1371/journal.pmed.0020124

Keynes, J. M. (1921). *A treatise on probability*. Macmillan.

Knight, F. H. (1921). *Risk, uncertainty and profit*. Houghton Mifflin.

Kruger, J., & Dunning, D. (1999). Unskilled and unaware of it: How difficulties in recognizing one's own incompetence lead to inflated self-assessments. *Journal of Personality and Social Psychology, 77*(6), 1121–1134. https://doi.org/10.1037/0022-3514.77.6.1121

Lorenz, E. N. (1963). Deterministic nonperiodic flow. *Journal of the Atmospheric Sciences*, *20*(2), 130–141. https://doi.org/10.1175/1520-0469(1963)020<0130:DNF>2.0.CO;2

Manski, C. F. (2015). Communicating uncertainty in official economic statistics: An appraisal fifty years after Morgenstern. *Journal of Economic Literature, 53*(3), 631–653. https://doi.org/10.1257/jel.53.3.631

Merton, R. K. (1948). The self-fulfilling prophecy. *The Antioch Review, 8*(2), 193–210. https://doi.org/10.2307/4609267

Moore, D. A., & Healy, P. J. (2008). The trouble with overconfidence. *Psychological Review, 115*(2), 502–517. https://doi.org/10.1037/0033-295X.115.2.502

Open Science Collaboration. (2015). Estimating the reproducibility of psychological science. *Science*, *349*(6251), Article aac4716. https://doi.org/10.1126/science.aac4716

Plato. (1997). Apology (G. M. A. Grube, Trans.). In J. M. Cooper (Ed.), *Plato: Complete works*. Hackett.

Popper, K. R. (1959). *The logic of scientific discovery*. Hutchinson.

Shirani-Mehr, H., Rothschild, D., Goel, S., & Gelman, A. (2018). Disentangling bias and variance in election polls. *Journal of the American Statistical Association, 113*(522), 607–614. https://doi.org/10.1080/01621459.2018.1448823

Smithson, M. (1989). *Ignorance and uncertainty: Emerging paradigms*. Springer. https://doi.org/10.1007/978-1-4612-3628-3

Spiegelhalter, D. J., & Riesch, H. (2011). Don't know, can't know: Embracing deeper uncertainties when analysing risks. *Philosophical Transactions of the Royal Society A, 369*(1956), 4730–4750. https://doi.org/10.1098/rsta.2011.0163

Taleb, N. N. (2007). *The black swan: The impact of the highly improbable*. Random House.

Tetlock, P. E. (2005). *Expert political judgment: How good is it? How can we know?* Princeton University Press.

Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science*, *185*(4157), 1124–1131. https://doi.org/10.1126/science.185.4157.1124

van der Bles, A. M., van der Linden, S., Freeman, A. L. J., Mitchell, J., Galvao, A. B., Zaval, L., & Spiegelhalter, D. J. (2019). Communicating uncertainty about facts, numbers and science. *Royal Society Open Science, 6*(5), Article 181870. https://doi.org/10.1098/rsos.181870

Walker, W. E., Harremoës, P., Rotmans, J., van der Sluijs, J. P., van Asselt, M. B. A., Janssen, P., & Krayer von Krauss, M. P. (2003). Defining uncertainty: A conceptual basis for uncertainty management in model-based decision support. *Integrated Assessment, 4*(1), 5–17. https://doi.org/10.1076/iaij.4.1.5.16466