

“Negative Knowledge:” From Wicked Problems and Rude Surprises to Mathematics

Michael Smithson, The Australian National University

Posted 18/10/2010 on

http://www.bestthinking.com/thinkers/science/social_sciences/psychology/michael-smithson?tab=blog

It is one thing to know that we don't know, but what about knowing that we can *never* know something? Karin Knorr-Cetina (1999) first used the term *negative knowledge* to refer to knowledge about the limits of knowledge. This is a type of *meta-knowledge*, and is a special case of known unknowns. Philosophical interest in knowing what we don't know dates back at least to Socrates—certainly long before Donald Rumsfeld's [prize-winning remark](#) on the subject. Actually, Rumsfeld's “unknown unknowns” were articulated in print much earlier by philosopher Ann Kerwin, whose 1993 paper appeared along with mine and others in a [special issue](#) of the journal *Science Communication* as an outcome of our symposium on “Ignorance in Science” at the AAAS meeting in Boston earlier that year. My [1989](#) coinage, *meta-ignorance*, is synonymous with unknown unknowns.

There are plenty of things we know that we cannot know (e.g., I cannot know my precise weight and height at the moment I write this), but why should negative knowledge be important? There are at least three reasons. First, negative knowledge tells us to put a brake on what would otherwise be a futile wild goose-chase for certainty. Second, some things we cannot know we might consider important to know, and negative knowledge humbles us by highlighting our limitations. Third, negative knowledge about important matters may be contestable. We might disagree with others about it.

Let's begin with the notion that negative knowledge instructs us to cease inquiry. On the face of it, this would seem a good thing: Why waste effort and time on a question that you know cannot be answered? Peter Medawar (1967) famously coined the aphorism that science is the “art of the soluble.” A commonsensical inference follows that if a problem is not soluble then it isn't a scientific problem and so should be banished from scientific inquiry. Nevertheless, aside from logical flaw in this inference, over-subscribing to this kind of negative-knowledge characterization of science exacts a steep price.

First, there is what philosopher Jerome Ravetz (in the same journal and symposium as Ann Kerwin's paper) called *ignorance of ignorance*. By this phrase Ravetz meant something slightly different from meta-ignorance or unknown unknowns. He observed that conventional scientific training systematically shields students from problems outside the soluble. As a result, they remain unacquainted with those problems, i.e., ignorant about scientific ignorance itself. The same charge could be laid on many professions (e.g., engineering, law, medicine).

Second, by neglecting unsolvable problems scientists exclude themselves from any input into what people end up doing about those problems. Are there problem domains where negative knowledge defines the criteria for inclusion? Yes: *wicked problems* and *rude surprises*. The characteristics of wicked problems were identified in the classic [1973](#) paper by Rittel and Webber, and most of these referred to various kinds of negative knowledge. Thus, the very definition and scope of wicked problems are unresolvable; such problems have no definitive solutions; there are no ultimate tests of whether a solution works; every wicked problem is unique; and there are no opportunities to learn how to deal with them by trial-and-error. Claimants to the title of “wicked problem” include how to craft policy responses to climate change, how to combat terrorism, how to end poverty, and how to end war.

Rude surprises are not always wicked problems but nonetheless are, as Todd La Porte describes them in his [2005](#) paper, “unexpected, potentially overwhelming circumstances that

are likely to deliver punishing blows to human life, to political or economic viability, and/or to environmental integrity” (pg. 2). Financial advisors and traders around the world no doubt saw the most recent global financial crisis as a rude surprise.

As Matthias Gross (2010) points out at the beginning of his absorbing book, “ignorance and surprise belong together.” So it should not be, well, surprising that in an uncertain world we are in for plenty of surprises. But why are we so unprepared for surprises? Two important reasons are confirmation bias and the Catch-All Underestimation Bias (CAUB). Confirmation bias is the tendency to be more interested in and pay more attention to information that is likely to confirm what we already know or believe. As Robert Nickerson’s 1998 review sadly informs us, this tendency operates unconsciously even when we’re not trying to defend a position or bolster our self-esteem. The CAUB is a tendency to underestimate the likelihood that something we’ve never seen before will occur. The authors of the classic 1978 study first describing the CAUB pointed out that it’s an inescapable “out of sight, out of mind” phenomenon—After all, how can you have something in sight that never has occurred? And the final sting in the tail is that clever people and domain experts (e.g., scientists, professionals) suffer from both biases just as the rest of us do.

Now let’s move to the second major issue raised at the outset of this post: Not being able to know things we’d like to know. And let’s raise the stakes, from negative knowledge to negative meta-knowledge. Wouldn’t it be wonderful if we had a method of finding truths that was guaranteed not to steer us wrong? Possession of such a method would tame the wild seas of the unknown for us by providing the equivalent of an epistemic compass. Conversely, wouldn’t it be devastating if we found out that we never can obtain this method?

Early in the 20th century, mathematicians underwent the experience of expecting to find such a method and having their hopes dashed. They became among the first (and best) postmodernists. Their story has been told in numerous ways (even as a [graphic novel](#)), but for my money the best account is the late [Morris Kline](#)’s brilliant (1980) book, “Mathematics: The Loss of Certainty.” Here’s how Kline characterizes mathematicians’ views of their domain at the turn of the century:

“After many centuries of wandering through intellectual fog, by 1900 mathematicians had seemingly imparted to their subject the ideal structure... They had finally recognized the need for undefined terms; definitions were purged of vague or objectionable terms; the several branches were founded on rigorous axiomatic bases; and valid, rigorous, deductive proofs replaced intuitively or empirically based conclusions... mathematicians had cause to rejoice.” (pg. 197)

The tantalizing prospect before them was to establish the consistency and completeness of mathematical systems. Roughly speaking, *consistency* amounts to a guarantee of never running into paradoxes (well-formed mathematical propositions that nevertheless are provably both true and false) and *completeness* amounts to a guarantee of never running into undecidables (well-formed mathematical propositions whose truth or falsity cannot be proven). These guarantees would tame the unknown for mathematicians; a proper axiomatic foundation would ensure that any proposition derived from it would be provably true or false.

The famous 1931 paper by Kurt [Gödel](#) denied this paradise forever. He showed that if any mathematical theory adequate to deal with whole numbers is consistent, it will be incomplete. He also showed that consistency of such a theory could not be established by the logical principles in use by several foundational schools of mathematics. So, consistency would have to be determined by other methods and, if attained, its price would be incompleteness. But is there a way to ascertain which mathematical propositions are undecidable and which

provable? Alan [Turing](#)'s 1936 paper on "computable numbers" (in addition to inventing Turing machines!) showed that the answer to this question is "no." One of the consequences of these results is that instead of a foundational consensus there can be divergent schools of mathematics, each legitimate and selected as a matter of preference. Here we have definitively severe negative knowledge in an area that to most people even today epitomizes certitude.

"Loss of certainty" themes dominate high-level discourse in various intellectual and professional domains throughout the 20th century. Physics is perhaps the most well-known example, but one can find such themes in many other disciplines and fascinating debates around them. To give one example, historian Ann Curthoys' and John Docker's [2006](#) book "Is History Fiction?" begins by identifying three common responses to the book title's question: Relativists who answer in the affirmative, foundationalists who insist that history is well-grounded in evidence after all, and a third (they claim, largest) puzzled group who says "well, is it?" To give just one more, I'm a mathematical modeler in a discipline where various offspring of the "is psychology a science?" question are seriously debated. In particular, I and others (e.g., [here](#) and [here](#)) regard the jury as still out on whether there are (m)any quantifiable psychological attributes. Some such attributes can be rank-ordered, perhaps, but quantified? Good question.

Are there limits to negative knowledge—In other words, is there such a thing as negative negative-knowledge? It turns out that there is, mainly in the Gödelian realm of self-referential statements. For example, we cannot believe that we currently hold a false belief; otherwise we'd be compelled to disbelieve it. There are also limits to the extent to which we can self-attribute erroneous belief formation. Philosophers Andy Egan and Adam Elga laid these out in their delightfully titled [2005](#) paper, "I Can't Believe I'm Stupid." According to them, I can believe that in some domains my way of forming beliefs goes wrong all of the time (e.g., I have a sense of direction that is invariably wrong), but I can't believe that every time I form any belief it goes wrong without undermining that very meta-belief.

Dealing with wicked problems and rude surprises requires input from multiple stakeholders encompassing their perspectives, values, priorities, and (possibly non-scientific) ways of knowing. Likewise, there is no algorithm or sure-fire method to anticipate or forecast rude surprises or Nicolas Taleb's "[black swans](#)." These are exemplars of insoluble problems beyond the ken of science. But surely none of this implies that input from experts is useless or beside the point. So, are there ways of educating scientists, other experts, and professionals so that they will be less prone to Ravetz's ignorance of ignorance? And what about the rest of us—Are there ways we can combat confirmation bias and the CAUB? Are there effective methods for dealing with wicked problems or rude surprises? Ah, grounds for a future post!