

Hi, I'm back again after a few weeks' travel (presenting papers at conferences). I've already [posted](#) material on this blog about the "ignorance explosion." Numerous writings have taken up the theme that there is far too much relevant information for any of us to learn and process and the problem is worsening, despite the benefits of the internet and effective search-engines. We all have had to become more hyper-specialized and fragmented in our knowledge-bases than our forebears, and many of us find it very difficult as a result to agree with one another about the "essential" knowledge that every child should receive in their education and that every citizen should possess.

Well, here is a modest proposal for one such essential: We should all become expert about experts and expertise. That is, we should develop *meta-expertise*.

We can't know everything, but knowing an expert when we see one, being able to tell the difference between an expert and an impostor, and knowing what it takes to become an expert can guide our search for assistance in all things about which we're ignorant. A meta-expert should:

1. Know the broad parameters of and requirements for attaining expertise;
2. Be able to distinguish a genuine expert from a pretender or a charlatan;
3. Know whether expertise is and when it is not attainable in a given domain;
4. Possess effective criteria for evaluating expertise, within reasonable limits; and
5. Be aware of the limitations of specialized expertise.

Let's start with that strongly democratic source of expertise: [Wikipedia's take on experts](#):

"In many domains there are objective measures of performance capable of distinguishing experts from novices: expert chess players will almost always win games against recreational chess players; expert medical specialists are more likely to diagnose a disease correctly; etc."

That said, the Wikipedia entry also raises a potentially vexing point, namely that "expertise" may come down to merely a matter of consensus, often dictated by the self-same "experts." Examples readily spring to mind in areas where objective measures are hard to come by, such as the arts. But consider also domains where objective measures may be obtainable but not assessable by laypeople. Higher mathematics is a good example. Only a tiny group of people on the planet were capable of assessing whether Andrew Wiles really had proven Fermat's Theorem. The rest of us have to take their word for it.

A crude but useful dichotomy splits views about expertise into two camps: Constructivist and performative. The constructivist view emphasizes the influence of communities of practice in determining what expertise is and who is deemed to have it. The performative view portrays expertise as a matter of learning through deliberative practice. Both views have their points, and many domains of expertise have elements of both. Even domains where objective indicators of expertise are available can have constructivist underpinnings. A proficient modern-day undergraduate physics student would fail late 19<sup>th</sup>-century undergraduate physics exams; and experienced medical practitioners emigrating from one country to another may find their qualifications and experience unrecognized by their adopted country.

What are the requirements for attaining deep expertise? Two popular criteria are talent and deliberative practice. Re deliberate practice, a much-discussed rule of thumb is the "10,000 hour rule." This rule was popularized in Malcolm Gladwell's book [Outliers](#) and some authors misattribute it to him. It actually dates back to studies of chess masters in the 1970's (see Ericsson, K. A., R. Th. Krampe, and C. Tesch-Römer, [1993](#)), and its generalizability to other domains still is debatable. Nevertheless, the 10K rule has some merit, and unfortunately it has

been routinely ignored in many psychological studies comparing “experts” with novices, where the “experts” often are undergraduates who have been given a few hours’ practice on a relatively trivial task.

The 10K rule can be a useful guide but there’s an important caveat. It may be a *necessary* but it is by no means a *sufficient* condition for guaranteeing deep expertise. At least three other conditions have to be met: *Deliberative* and *effective* practice in a domain *where deep expertise is attainable*. Despite this quite simple line of reasoning, plenty of published authors have committed the error of viewing the 10K rule as both necessary and sufficient. Gladwell didn’t make this mistake, but Jane [McGonigal’s recent book](#) on video and computer games devotes considerable space to the notion that because gamers are spending upwards of 10K hours playing games they must be attaining deep “expertise” of some kind. Perhaps some may be, provided they are playing games of sufficient depth. But many will not. (BTW, McGonigal’s book is worth a read despite her over-the-top optimism about how games can save the world—and take a look at her game-design collaborator [Bogost’s](#) somewhat dissenting review of her book).

Back to the caveats. First, no deliberation makes practice useless. Having spent approximately 8 hours every day sleeping for the past 61 years (178,120 hours) hasn’t made me an expert on sleep. Likewise, deliberative but ineffective practice methods deny us top-level expertise. Early studies of Morse Code experts demonstrated that mere deliberative practice did not guarantee best performance results; specific training regimes were required instead. Autodidacts with insight and aspirations to attain the highest performative levels in their domains eventually realise how important getting the “right” coaching or teaching is.

Finally, there is the problem of determining whether effective, deliberative practice yields deep expertise in any domain. The domain may simply not be “deep” enough. In games of strategy, tic-tac-toe is a clear example of insufficient depth, checkers is a less obvious but still clear example, whereas chess and go clearly have sufficient depth.

Tic-tac-toe aside, are there domains that possess depth where deep expertise nevertheless is unattainable? There are, at least, some domains that are deeply complex where “experts” perform no better than less trained individuals or simple algorithms. Psychotherapy is one such domain. There is a plethora of studies demonstrating that clinical psychologists’ predictions of patient outcomes are worse than simple linear regression models (cf. Dawes’ searing indictment in his [1994](#) book) and that sometimes experts’ decisions are no more accurate than beginners’ decisions and simple decision aids. Similar results have been reported for financial planners and political experts. In Philip Tetlock’s [2005](#) book on so-called “expert” predictions, he finds that many so-called experts perform no better than chance in predicting political events, financial trends, and so on.

What can explain the absence of deep expertise in these instances? Tetlock attributes experts’ poor performance to two factors, among others: Hyperspecialization and overconfidence. “We reach the point of diminishing marginal predictive returns for knowledge disconcertingly quickly,” he reports. “In this age of academic hyperspecialization, there is no reason for supposing that contributors to top journals—distinguished political scientists, area study specialists, economists, and so on—are any better than journalists or attentive readers of the *New York Times* in ‘reading’ emerging situations.” And the more famous the forecaster the more overblown the forecasts. “Experts in demand,” Tetlock says, “were more overconfident than their colleagues who eked out existences far from the limelight.” [Tetlock also claims](#) that cognitive style counts: “Foxes” tend to outperform “hedgehogs.” These terms are taken from [Isaiah Berlin’s popular essay](#): Foxes know a little about lots of things, whereas hedgehogs know one big thing.

Another contributing factor may be a lack of meta-cognitive insight on the part of the experts. A hallmark of expertise is ignoring (not ignorance). This proposition may sound less counter-intuitive if it's rephrased to say that experts know what to ignore. In an earlier post I mentioned Mary Omodei and her colleagues' [chapter](#) in a 2005 [book](#) on professionals' decision making in connection with this claim. Their chapter opens with the observation of a widespread assumption that domain experts also know how to optimally allocate their cognitive resources when making judgments or decisions in their domain. Their research with expert fire-fighting commanders cast doubt on this assumption.

The key manipulations in the Omodei simulated fire-fighting experiments determined the extent to which commanders had unrestricted access to "complete" information about the fires, weather conditions, and other environmental matters. They found that commanders performed more poorly when information access was unrestricted than when they had to request information from subordinates. They also found that commanders performed more poorly when they believed all available information was reliable than when they believed that some of it was unreliable. The disquieting implication of these findings is that domain expertise doesn't include meta-cognitive expertise

Cognitive biases and styles aside, another contributing set of factors may be the characteristics of the complex, deep domains themselves that render deep expertise very difficult to attain. Here is a list of tests you can apply to such domains by way of evaluating their potential for the development of genuine expertise:

1. Stationarity? Is the domain stable enough for generalizable methods to be derived? In chaotic systems long-range prediction is impossible because of initial-condition sensitivity. In human history, politics and culture, the underlying processes may not be stationary at all.
2. Rarity? When it comes to prediction, rare phenomena simply are difficult to predict (see my [post](#) on making the wrong decisions most of the time for the right reasons).
3. Observability? Can the outcomes of predictions or decisions be directly or immediately observed? For example in psychology, direct observation of mental states is nearly impossible, and in climatology the consequences of human interventions will take a very long time to unfold.
4. Objective or even impartial criteria? For instance, what is "good," "beautiful," or even "acceptable" in domains such as music, dance or the visual arts? Are such domains irreducibly subjective and culture-bound?
5. Testability? Are there clear criteria for when an expert has succeeded or failed? Or is there too much "wobble-room" to be able to tell?

Finally, here are a few tests that can be used to evaluate the "experts" in your life:

1. Credentials: Does the expert possess credentials that have involved testable criteria for demonstrating proficiency?
2. Walking the walk: Is the expert an active practitioner in their domain (versus being a critic or a commentator)?
3. Overconfidence: Ask your expert to make yes-no predictions in their domain of expertise, and before any of these predictions can be tested ask them to estimate the percentage of time they're going to be correct. Compare that estimate with the resulting percentage correct. If their estimate was too high then your expert may suffer from over-confidence.

4. Confirmation bias: We're all prone to this, but some more so than others. Is your expert reasonably open to evidence or viewpoints contrary to their own views?
5. Hedgehog-Fox test: Tetlock found that Foxes were better-calibrated and more able to entertain self-disconfirming counterfactuals than hedgehogs, but allowed that hedgehogs can occasionally be "stunningly right" in a way that foxes cannot. Is your expert a fox or a hedgehog?
6. Willingness to own up to error: Bad luck is a far more popular explanation for being wrong than good luck is for being right. Is your expert balanced, i.e., equally critical, when assessing their own successes and failures?