Sep/Oct 2008

Enlarge cover

Back Issues

Search

Forum

Calendar

Links

Books

Interview

Nassim Nicholas Taleb

Problems of Knowledge

Mediocristan and Extremistan

Nassim Nicholas Taleb has had a run-away success with *The Black Swan*, a book about surprise run-away successes. **Constantine Sandis** talks with him about knowledge and scepticism.

CS: "The fox knows many things, but the hedgehog knows one big thing," the Greek poet Archilochus once wrote. Isaiah Berlin famously used this saying to introduce a distinction between two different kinds of thinkers: those who "pursue many ends, often unrelated and even contradictory," and those who relate everything to a "single central vision… a single, universal, organizing principle in terms of which alone all that they are and say has significance." Nassim, I think you would see yourself as an intellectual hedgehog. I have heard you describe your big idea, which you have had from childhood, as the view that the more improbable an outcome is, the higher its impact (and vice versa*). Many controversial corollaries follow from this one claim, and you have recorded some of the most important in numerous essays, as well as in your books* Fooled by Randomness *and* The Black Swan*. These thoughts range from remarks on the perils of socioeconomic forecasting, to views on epistemology, agnosticism, explanation and understanding, the dynamics of historical events, and perhaps most importantly, advice on where to live and which parties to attend.*

As I find myself in the awkward position of wishing to endorse almost all of the corollaries while fiercely disagreeing with the large idea from which they stem, I would like to ask you whether you really think that all *improbable events have a high impact, and that* all *high impact events are improbable, or whether this is simply a misleading way of describing your insight? Are you happy to allow for counterexamples?*

NNT: My core idea is about the effect of non-observables in real life. My focus is on the errors which result: how the way we should act is affected by things we can't observe, and how we can make decisions when we don't have all the relevant information.

My idea concerns decision-making under conditions of uncertainty, dealing with incomplete information, and living in a world that has a more complicated ecology than we tend to think. This question about justifiable responses to the unknown goes way beyond the conventionally-phrased 'problem of induction' and what people call 'scepticism'. These classifications are a little too tidy for real life. Alas, the texture of real life is more sophisticated and more demanding than analytical philosophy can apparently handle. The point is to avoid 'being the turkey'. [*There is a philosophical parable about a turkey who on the basis of daily observations concludes that he's always fed at 9am. On Christmas Eve he discovers this was an overhasty generalisation – Ed.*] To do so you have to stand some concepts on their head – like your concept of the use of beliefs in decision making.

Let me explain this. First, when you are making a decision, you don't act on your 'raw beliefs' about the situation in question, or on what probabilists call 'degrees of beliefs'. Instead you base your decision on something that involves possible *consequences* as you see them. For instance, you can *believe* with a high degree of certainty, say 99%, that a plant is not poisonous, yet refuse to eat it; or that the plane won 't crash – yet refuse to ride in it. Yet you may accept a bet with a mere 0.1% degree of belief that you will win. *Perception of impact*, that is, of consequences, is everything in decision-making – *not likelihood of truth*. I insist, likelihood of truth is too crude a notion, rarely useable. And potential impact is vastly easier to figure out than the probability of outcome, for mathematical reasons related to sample insufficiency. So I focus only on the high impact domains I call 'Extremistan', in which a single event can be massively consequential.

Beware beliefs. Although you may not believe with any conviction that you are in Extremistan, you might still act as if you were there, if the consequence of an event might be too large. I will repeat until I die: *beliefs* have little to do with *what we do in real life*.

The problem of induction also deals with non-observables. It asks, How can you generalize from particular observations of the world to theories that you can use to make predictions? But having spent my life taking decisions in a highly random environment – I was a Wall Street trader – I felt alienated from modern philosophical treatments of induction, which I dub 'Mickey Mouse' owing to their highly naïve attributes.

The typical formulation of the problem of induction is attributed to Hume, although incidentally, it is much, much older. Hume got it from Huet, Foucher, Bayle, and a collection of sophisticated and forgotten French thinkers, who got it wholesale from earlier sources, including Algazel. Hume said that induction presupposes that nature behaves in a uniform fashion, but that this belief has no defence in reason – it just reflects our mental habits resulting from our experiences so far. But having reached this sceptical position, Hume started thinking of the problem as ridiculous. He left it in the philosophical cabinet, as nothing to do with real life, making it something 'academic', in the bad sense of the word. Modern philosophers calls such ivory tower theorising 'the problem of insulation' – in *The Black Swan* I present this as the problem of 'domain dependence'. The problem is that what academics do in front of a blackboard has little bearing on what they do in real life. And, I insist, *it is real life that matters to me*: I'm interested in the ecology of uncertainty, not induction and deduction.

However, a certain class of sceptics took the problem of uncertainty into vastly more operational territory. They were the medical doctors of the *Empirical* sect to which Sextus Empiricus was supposed to belong. For them scepticism was a very practical problem, unrelated to sterile Humean scepticism, and also unrelated to the Pyrrhonians, who took their scepticism to absurd extremes.

Finally, let me answer your question. I have two points of disagreement with you. Firstly, I do not consider myself a hedgehog, but a fox: I warn against focusing ('anchoring') on a single possible rare event. Rather, be prepared for the fact that the next large surprise, technological or historical, will not resemble what you have in mind (big surprises are what some people call 'unknown unknowns'). In other words, learn to be abstract, and think in second order effects rather than being anecdotal – which I show to be against human nature. And crucially, rare events in Extremistan are more consequential by their very nature: the once-every-hundred-year flood is more damaging than the 10 year one, and less frequent. You have fewer billionaires than millionaires. You get far fewer wars that kill more than 20 million people than wars that kill a few thousand. There are far fewer bestselling authors than authors. So, empirically, the rate of occurrence of events tends to decline with their impact.

CS: Your distinction between Mediocristan and Extremistan is unconcerned with the fact that as far as logic is concerned, any domain can switch from Mediocristan to Extremistan (or vice versa*) in a split second. For example, there could be a complete lack of bestsellers in next year 's book market. Such switches would have a sufficiently large impact to count as Black Swans. Yet your sound advice about which sort of situations to seek out and which ones to avoid is unfazed by the logical possibility of a switch – in my opinion quite rightly.*

I agree that Hume seems to have a general insulation problem when he switches from scepticism to non-scepticism about the external world. What exactly is going on there is open to several interpretations, as is his reaction to the problem of induction. On one reading – supported by his sceptical attitude towards reports of miracles – Hume concludes that induction is justified because all we can mean *by causation is [experience of] past regularity creating an expectation of continuing regular activity – which is the assumption made for induction. As Wittgenstein would later put it, such regularity just* is *what we call a ground for behaving in certain ways. Justified belief does not require* logical *certainty.*

I think you believe something similar, Nassim: not being a turkey is a matter of learning from past irregularities. Yet to do this, one must treat empirical facts about the past as grounds for decision-making. Would you agree, then, that your kind of scepticism is not about induction, but, on the contrary, is grounded on inference based on past irregularities*? If so, is there some extent to which you are, after all, only concerned about impact in relation to probability? At any rate, that was the reasoning behind calling you a 'hedgehog'.*

NNT: Total scepticism, like that advocated by Pyrrho of Elis, is not workable. Take old Pyrrho himself. He was said to doubt the existence of obstacles, and so walked into them. He doubted (or claimed to doubt) everything. It doesn't take long for someone who doubts physical obstacles to exit the gene pool. Such indiscriminate scepticism is no different from total gullibility. However the *adoxastism* [scepticism] of Sextus Empiricus took into account the 'commemorative signs' – what a scar is to a wound or a wall to a passage. Sextus would have avoided stepping off a cliff. He was practical.

There is an art to being a sceptic – it requires both commitment and sophistication. Let me return to the problem of induction. There exists an asymmetry here: Extremistan can fool you into believing it is Mediocristan, but never the reverse. Things can look pretty average until the extreme instance comes along; but once we've had the extreme instance, we know we're not in Mediocristan any more. At least we know we're in Extremistan, and can avoid the traps there.

This is related to the theory of 'falsification' generally attributed to Karl Popper. This theory basically says that you should ask about any erstwhile scientific theory whether or not it could be disproved. You can discover more easily what's wrong than what's right. You can know that 'not all swans are white' more easily and confidently than 'all swans are white'. But falsification wasn't invented by Popper. The ancient empirical doctors knew it well, practiced it, and called it the 'negative instance'. The negative instance was considered vastly more informative than confirmatory observations.

Popper didn't discuss the ancient Empiricists. How did Popper rediscover falsification? He was trying to escape the false problem of induction'. But in doing so he reformulated a weak version of the anti-Aristotelian ideas exposed in Sextus Empiricus ' work and others.

CS: Hume felt similarly to you about the Pyrrhonian sceptics. He writes "the great subverter of Pyrrhonism *or the excessive principles of scepticism is action, and employment, and common life " which, in combination with human nature, causes the Pyrrhonian absolutely sceptical principles to vanish like smoke. In their place Hume offers "a more mitigated scepticism… which may be both durable and useful… of advantage to mankind… and which may be the natural result of Pyrrhonian doubts and scruples." He then proceeds to describe this sensible scepticism as "the limitation of our enquiries to such subjects as best adapted to the narrow capacity of human understanding."*

I can picture Hume asserting that your distinction between Mediocristan and Extremistan is justified by your knowledge of human nature rather than pure mathematics or abstract logic. But I'm glad you mentioned Popper. His 'falsifiability' test is meant to distinguish genuine science from pseudo-science. Do you think that the sorts of predictions made by social scientists and economists are falsifiable? One worry about falsification here might be that so-called experts could easily hide behind the Quine-Duhem thesis of 'confirmation holism': that is, every time they fail to predict something, they hide behind the view that no theory can be tested in isolation. Would such tactics count as an instance of the kind of intellectual cowardice and/or arrogance you attack in your books?

NNT: Taken seriously as a theory, Popperian falsification does not cut the mustard in an environment as complicated as ours. Popper was an improvement on Hume; but he was still writing to Mickey Mouse philosophers, and did not go all the way. I wrote that his really potent contribution was his presentation of the inherent limits of forecasting.

Quine's response about falsification, and its sequels, concerns the multiplication of degrees of freedom in what I call the inverse problem: given a set of observations, in a complex world an infinity of theories will be compatible with the dataset, many producing uncorrelated and vastly different forecasts. But there is an art to how to think conservatively and cautiously. The more your problem and your errors are in Extremistan, the less dependant you can be on the variations.

Let me mention another problem with current philosophical works on uncertainty. Almost all I've seen focus on things that are inconsequential to everyday life, but which analytical philosophers can discuss 'rigorously' – like quantum indeterminacy, which in everyday terms amounts to just plain Mediocristan randomness. Although individual quantum behaviour is random, my table does not jump three foot to the left because its quanta are *collectively predictable* – we can predict what philosophers call the 'uncertainty' to twenty-decimal precision. To me, *real* uncertainty is the poorly-predicted climatic problem, or what will happen with food prices, not the properties of some subatomic particle our eyes can't see.

To conclude, my aim with *The Black Swan* was to make philosophy more ecologically rigorous: less pseudo-rigorous, less naïve, less Mickey Mouse; more real, more relevant. This is what philosophy was originally meant to be.

• *Nassim Nicholas Taleb was formerly a Wall Street trader, He's the author of two bestsellers about knowledge and randomness:* Fooled by Randomness *and* The Black Swan.

• *Constantine Sandis is a Senior Lecturer in Philosophy at Oxford Brookes University and NYU in London.*

© Philosophy Now 2008. All rights reserved.