http://doi.org/10.47967/TOR2022TRANS.VOL7.01

LEARNING TO LIVE IN A RADICALLY UNCERTAIN WORLD

BENJI WOOLF

The MRC Integrative Epidemiology Unit, University of Bristol
Department of Psychological Science, University of Bristol

ABSTRACT

The past few years have been fraught with uncertainty. In this essay, I review and reflect on Mervyn King and John Kay's topical book *Radical Uncertainty* (2020). In the book, King and Kay defend the usefulness of viewing some types of uncertainty as unquantifiable. In what follows, I will argue that this is a plausible characterisation of uncertainty, and I reflect on the extent to which it has helped me frame both my own personal experience of uncertainty and the uncertainty we have all faced due to the COVID-19 pandemic.

The 6th of August 2020 was an unremarkable day for me. At some point that day, I finished reading the penultimate chapter of Mervyn King and John Kay's book, which is the focus of this essay. As I went to bed at home, any predictions I could have made about the 7th of August; when I would finish the book; and the following months, were wrong. Within 24 hours I was admitted to intensive care. Within a week, I had gone into multiple organ failure and was in a chemically induced coma while on complete life-support. I remained in a coma until the 6th of September and did not come off life-support for a further month. During this period I was diagnosed as having one of the worst recorded cases of endocarditis (an infection of the heart, which statistically should have been fatal).

King, the former governor of the Bank of England, and Kay, a professor of economics at Oxford University, set out in their book, *Radical Uncertainty*, to understand the catastrophic failure of the economics profession to predict and respond to the 2008 financial crash. In this essay, I will first summarise the book and argue that they are essentially correct in their characterisation of uncertainty, and that the argument of the book is applicable to much more than just the economic and political decisions that they focus on. If my own - rather extreme - experience is anything to go by, their approach can be important for understanding life-changing events. Having argued for the soundness of their characterisation of uncertainty, I will then argue that the book diagnoses many of the flaws of the UK's early COVID-19 strategy, and the naive application of statistics more generally, before concluding with some more vernal reflections.

Radical Uncertainty is structured into six parts. The first describes the nature of uncertainty. King and Kay explain the history of the treatment of uncertainty in the economics literature. Broadly, there are two camps: one, which is now dominant in both psychology and economics, treats uncertainty as statistically quantifiable. An example of this from psychology is Signal Detection Theory, in which uncertainty is represented by the variability in the distribution of observations. The other school, which was defined by figures like Frank Knight (one of the founders of the Chicago School of economics) and John-Maynard Keynes, argues that uncertainty is not quantifiable. Radical uncertainty is, in essence, a defence of the importance of the latter understanding of uncertainty. However, unlike Knight, they do not define uncertainty in general as unquantifiable but instead note that some types of uncertainty are unquantifiable (what they dub uncertainty radically). This seems sensible. Critics of Knight have argued that there are tangible cases of uncertainty where it can be quantified: No-one can be certain about what will happen when we flip a coin, but we do know that whatever does happen has a 50% chance of occurring. However, these critics often then ignore that some instances of uncertainty are not obviously quantifiable. As King and Kay argue, President Obama had no way of quantifying what the outcome of the storming of bin Laden's compound would be, and nor would it make sense to assign a numerical probability to it.

The middle sections of the book explain King and Kay's opinion for why quantitative approaches to uncertainty have become dominant and why they think these approaches are incorrect. To do this, King and Kay cover the history of and philosophies behind modern economics, probability theory, and parts of cognitive and evolutionary psychology. This is all fascinating, especially to someone with interest in decision making, or the application of statistics to social sciences. Even without any formal training in economics, I found that the main concepts were clearly communicated and convincingly argued.

I found three arguments they make particularly compelling. Firstly, quantification of uncertainty requires correctly describing the data's probability distribution. If this is not the case, then both estimates of effect and uncertainty may be biased. Expected value theory, for example, says that decisions should be based on the mean of all possible outcomes weighted by their probabilities. This makes sense if the data is normally distributed, and the mean is a good description of the most plausible outcome. For example, imagine a lottery with a 50% chance of winning £2. The expected payment from this lottery is £1, so it would make sense to play the lottery if you can buy a ticket for no more than £1. On the other hand, imagine a lottery with a very small but finite chance, say one in a trillion, of winning an infinite amount of money. Because there is a potentially infinite pay-out, someone trying to maximise their 'expected' outcomes should be willing to pay as much as they could to play in this paradox, even though this would almost certainly lead to destitution. Keynes took this (St Petersburg) paradox to imply that 'expected' values are generally not of much use in decision making. Likewise, if a distribution is not stationary, i.e. its properties can change over time, then even if we can correctly describe the distribution at the current time, we do not necessarily know that this will generalise across time or localities, e.g. due to out of sample events. For example, when given a choice while hiking between a path on which most people got a minor injury due to poor maintenance, or one where there are no recent reports of injuries, most people would choose the second path. However, the first path would be more sensible if you knew that you were unaware of periodic but fatal flash floods on the second path.

The second argument is a related philosophical challenge to the widespread application of normal distributions, developed by N. N. Taleb in his book *The Black Swan* (2007). Taleb's book is a modernisation of Hume's 'problem of induction' for applied statistics. His central tenet is that the most important events in the real world are high impact and rare, such as flash floods. Black swans are defined as the subset of these events which we do not expect. Because important outliers are rare, a (random) sample of data is unlikely to include them and will therefore under-estimate how extreme the distribution is. To make matters worse, the rarer, and therefore more extreme, these outliers are, the less likely they are to be included in the sample. Without unrealistically large samples, it is therefore empirically impossible to guarantee that we have correctly described the true distribution of a set of observations using standard statistical methods when it is conceivable that there may be unobserved outliers more extreme than what we are currently expecting.

The final argument I will highlight is centred on the inability to have what are formally called 'well defined' outcomes. The classical, frequentist definition of the probability for some event is the proportion of times this event will occur over a suitably large number of observations. For example, to say that a die has a one in six chance of rolling a five is to say that if we were to roll it close to an infinite number of times, approximately one in every six rolls would be a five. This makes sense in the case of dice rolling. However, Kay and King argue that this logic is not applicable to many important real-world events. For example, Obama's advisors told him there was a 50% chance of success when raiding bin Laden's compound. For this to formally make sense, we would have to be able to estimate how many possible counterfactuals could occur and then estimate the number of these which had a chance of being successful. This is difficult because there is an infinite number of minor changes to any situation, which in a complex and chaotic system may have a meaningful impact on the outcome. We could try to approximate the probability by looking at the success rate of similar operations, but this only makes sense to the extent that we are comparing 'apples with apples.' We can obviously judge how similar two situations are, but it is unclear how we could create sufficiently accurate quantitative decision rules for this without inducing a sorites paradox (I would recommend https://plato.stanford.edu/entries/sorites-paradox/ if you are unsure about what a sorites paradox is). A similar line of reasoning leads King and Kay to make a distinction between 'big world' (of actual non-well-defined) problems and 'little world' (hypothetical well-defined) problems. Little world problems can be useful for understanding mechanisms that operate in the world but lack any radical uncertainty, and therefore are of limited use in making predictions and require qualitative interpretation when being applied. For example, quantitative evolutionary models are powerful tools for understanding the interplay of forces that influence evolutionary processes, but would not be of much use in predicting human evolution due to the large number of unpredictable high impact events which shaped our evolution (I'd recommend Richard Dawkins's books for an easy to read the example of the prior, and *A series of unfortunate events* by Sean Carroll to anyone interested in the role of chance in human evolution).

In the final part of the book, King and Kay present some practical applications of how to live with uncertainty. Much of this follows from previous arguments. If we accept the possibility of black swans, they argue that instead of trying to forecast the future, we should try to understand how systems work to build fail-safes to mitigate risk. This does not mean that quantitative analysis is useless. Understanding the mechanisms behind 'what is going on here' often requires using quantitative tools. However, we need to supplement them with qualitative knowledge to understand the limitations of generalising from the small world problems they solve to the big world problems we need to solve. (Interestingly, Judea Pearl reaches the same conclusion, for different reasons, about the requirement for qualitative knowledge when building quantitative models of causal mechanisms in his *Book of Why* (2018)).

As an epidemiologist, I found the COVID-19 pandemic an especially interesting test for these ideas. The early approach taken to the pandemic by the government and SAGE (Scientific Advisory Group for Emergencies) was to employ highly sophisticated (but opaque and sometimes unrealistic) models to predict the spread of the pandemic. The early March decision to let the virus 'rip' through the population, for example, was based on advice that assumed that COVID could be treated as an influenza pandemic. (Even the decision to lockdown was based on a forecast.) However, these early models ignored three important unknowns about COVID: long-term side effects, long-term immunity after infection, and the ability to truly isolate the most vulnerable from the virus. As we now know, COVID differs from this model in all the above aspects: immunity is not permanent, and there are long term side effects. In addition, almost every country which premised its response on protecting the vulnerable failed. Given these, letting the pandemic infect most of the population is not sensible. Had the government attempted to answer the question of 'what is going on here?' instead of 'what do the small world models say will happen?', this decision may have been aborted. The existing literature on coronaviruses, for example, already cast doubt on the potential for long term immunity, and the experience of every other country showed that stricter measures need to be imposed. (One rather depressing aside is that King's philosophy to modelling was influenced by his friend Robert May, who pioneered many of the mathematical methods used to model pandemics.)

I did not finish reading *Radial Uncertainty* until some point in October, during the prolonged period between leaving ICU and being discharged from the hospital. King and Kay conclude by suggesting that instead of forecasting the future, we should try to build resilient systems. While coming to terms with trauma and continued long-term medical uncertainties, I found this conclusion strangely therapeutic. However, this conclusion is, in my belief, also the weakest part of the book. Taleb argues in his sequel to *The Black Swan* that if we believe that the future is uncertain, we should not merely try to cope with uncertainty (i.e. be resilient). Instead, he argues, we should try to gain from uncertainty (i.e. to be 'Anti-Fragile'). Although I think anti-fragility may often be an unachievable ideal, I hope that my first-hand experience of radical uncertainty will improve my response to less existential future uncertainties.

REFERENCES

Kay, J.A. & M.A. King., 2020 *Radical uncertainty. Decision-making beyond the numbers.* Bridge Street Press.

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

Keynes, J.M., 1937 "The general theory of employment." *The quarterly journal of economics. 51*(2): 209-223.

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

Taleb, N.N., 2012 *Antifragile: how to live in a world we don't understand*. London: Allen Lane.

Dawkins, R. & N. Davis, 2017 *The selfish gene*. Macat Library.

Carroll, S.B., 2020 *A Series of Fortunate Events*. Princeton University Press.

Pearl, J. & D. Mackenzie, 2018 *The book of why: the new science of cause and effect*. Basic books.

BIOGRAPHY

Benjamin Woolf is a second year (part-time) PhD student in psychiatric and genetic epidemiology at the University of Bristol. His research mostly focuses on developing and applying quasi-experimental methods for casual inference, with an emphasis on Mendelian Randomisation.