Society
https://doi.org/10.1007/s12115-021-00572-0

BOOK REVIEW

John Kay and Mervyn King: *Radical Uncertainty: Decision-making for an Unknowable Future*

The Bridge Street Press, 2020. 528 pp., ISBN: 978-1408712603

Quassim Cassam [1]

In early 2011, President Obama faced one of the defining decisions of his presidency when the CIA shared with him a possible lead about the location of Osama bin Laden. US intelligence had identified a house in Abbottabad, Pakistan, that was occupied by the family of bin Laden's courier and two other families. The house was unusually secure, and its occupants went to extraordinary lengths to conceal their identities. When a surveillance drone spotted a tall man taking regular walks under the cover of trees in the garden, the CIA concluded that the Pacer, as they called him, was probably bin Laden. One analyst was 95% certain. Others were less confident and gave estimates ranging from 40 to 80%. Obama cut short a discussion of the merits of different estimates by observing "This is 50-50. Look guys, this is the flip of a coin" (p. 8). He could not base his decision on the notion that there was any greater certainty than that.

In their brilliant and wide-ranging account of the impact of radical uncertainty on decision-making, John Kay and Mervyn King are much impressed by the President's reaction. Indeed, the story of the discussions leading to the decision to raid the Abbottabad compound with US Navy SEALs is one they repeat many times in the course of 500 riveting pages. What impresses Kay and King is the President's impatience with bogus quantification. Either the Pacer was bin Laden, or he wasn't, and arguments about the *probability* that he was bin Laden were futile attempts to disguise the analysts' uncertainty. When Obama decided to send in the SEALs, "he did so not by probabilistic reasoning but by asking 'What is going on here?'" (p. 10).

A willingness to stand back and ask this banal-sounding question is, Kay and King argue, fundamental to our ability to cope with the fact that we have to make choices in "a radically uncertain world, in which probabilities cannot meaningfully be attached to alternative futures" (p. xvi). However, while the identity of the Pacer was uncertain prior to the raid, it was not *radically* uncertain. A radical uncertainty is one that cannot be resolved, but the identity of the Pacer was resolved when the SEALs located bin Laden in the Abbottabad compound. What is going on here? What does Obama's reasoning tell us about decision-making under *radical* uncertainty?

The notion of radical uncertainty is closely related to that of an "unknown unknown". Unknown unknowns, or what Nassim Nicholas Taleb calls true "black swans", are states of the world to which we cannot attach probabilities because we cannot conceive of these states. In these cases, Kay and King argue that it is not just that we do not know what will happen, we do not even know what kinds of things might happen. Yet, it was conceivable that the Pacer was Osama bin Laden and conceivable that he was not. Furthermore, the President had a fair idea of the kinds of things that might happen if the SEALs raided the compound. The identity of the Pacer was a known unknown rather than an unknown unknown – the President knew that he did not know whether that Pacer was bin Laden.

This shows that in practice, Kay and King have more than one conception of radical uncertainty. In one sense, a radical uncertainty is simply one that cannot be described in probabilistic terms applicable to games of chance, like the probability that a card drawn at random from a deck of cards will be a club. In this very broad sense, it *was* radically uncertain whether the Pacer was bin Laden, even though the uncertainty was ultimately resolved. In a stricter sense, a radical uncertainty is one that cannot be described in probabilistic terms because, unlike the presence or absence of bin Laden in Abbottabad, it is inconceivable. Before the wheel was invented, it would have made no sense to talk about the probability of the invention of the wheel since such a thing was at

✉ Quassim Cassam
q.cassam@warwick.ac.uk

[1] University of Warwick, Coventry, UK

Published online: 03 May 2021

that time inconceivable. Before the 9/11 attacks, it would have made no sense to talk about the probability of such an attack occurring.

Kay and King contrast radical with *resolvable* uncertainties, that is, ones that can be removed by looking something up or be represented by a known probability distribution. The President's uncertainty about the identity of the Pacer was not, in this sense, resolvable. Nor, for that matter, are many of the most pressing questions facing us today. The applicability of probabilistic reasoning to real-world problems is severely limited, as John Maynard Keynes and Frank Knight recognized many years ago. One of the missions of Kay and King's *tour de force* is to help their fellow economists, and the rest of us, to see clearly what Keynes and Knight saw.[1]

The lure of probabilities is compellingly described. A probability in the mathematical sense is a quantitative expression of the likelihood of one of a number of possible outcomes. So, for example, the probability of exactly 500 heads when a fair coin is tossed 1000 times is 2.523%. In the real world, matters are much less simple or quantifiable. The worlds of coin tosses and card games are "small" worlds in which the rules and possible outcomes are completely specified. We know what is going on here. In "large" and "non-stationary" worlds, where radical uncertainty is the norm, problems are rarely completely specified, the underlying processes are constantly changing, and our knowledge of these processes is imperfect. It is in these contexts that the application of the mathematics of probability is questionable and frequently spurious.

Of course, economists have their models, but "to make a statement about probability in a real world it is necessary to compound the probability derived from the model itself with the probability that the model is itself true" (p. 68). This simple but dazzling insight explains the failure of economic models to explain or predict events like the 2007–2008 financial crisis. On this point, Kay and King have some sympathy for Queen Elizabeth II. On a visit to the London School of Economics, she responded to a presentation on the financial crisis by asking "Why did no one see it coming?". Whatever their prestige, economists had "failed the test of providing useful insight" (p. 383). Economics, like dentistry, needs to be a practical subject, rather than a futile exercise in the construction of grand theories with little relevance to the real world.

Economists all too often make the mistake of "believing that you have more knowledge than you do about the real world from the application of conclusions from artificial models" (p. 109). The rot set in when Keynes and Knight lost the battle of ideas over the nature of uncertainty to Milton Friedman and others, and the concept of radical uncertainty was replaced in mainstream economics by the use of probabilities. Keynes and Knight distinguished between risks (events with known or knowable probabilities) and uncertainties (for which numerical probabilities cannot be specified). Writing in 1962, Friedman dismissed this distinction as invalid and urged economists to treat people as if they assigned numerical probabilities to every conceivable event.[2]

Readers who are not trained economists will not need much convincing by Kay and King that this is not a sensible assumption. Nor are they likely to be disconcerted by the observation that, contrary to what many economists would have us believe, real businesses and households do not optimize. They cope. They do not list possible courses of action and select the best available option. This is not a failure of rationality, just a sensible way of coping with an uncertain world. What mainstream economists understand by rationality is obedience to a set of a priori axioms, and *Radical Uncertainty* is a rich source of arguments against this view.

What is rational behaviour if not conformity with a set of axioms? Here, as elsewhere, Kay and King take ordinary usage as their guide. Rational judgements or actions are based on beliefs about the world that are reasonable, and they display internal logic and consistency. This type of practical rationality, which is similar to Aristotle's conception of deliberative excellence, calls into question the unflattering assessments of human behaviour offered by behavioural economists. Where behavioural economists see biases or other failures of rationality, Kay and King see real human beings doing their best to make sense of a complex, non-stationary world.

A good example of human sensemaking in practice is Daniel Kahneman's Linda problem, which is designed to bring out the prevalence of the conjunction fallacy, the fallacy of assuming that the probability of a conjunction can exceed the probability of its individual constituents. Suppose that all we know about Linda is that she is an outspoken philosophy major who is deeply concerned with discrimination and social justice. Is she more likely to be a bank teller or a bank teller who is active in the feminist movement? The latter, most people answer, but this is wrong in terms of probabilities. It cannot be more *probable* that Linda is a feminist bank teller than that she is a bank teller. To suppose that it can be more probable is to commit the conjunction fallacy. However, Kay and King note that subjects were not asked about *probabilities* but about *likelihood*, and these are not the same thing. It makes sense that Linda is a bank teller *and* a feminist because "people do not think of the Linda problem in terms of frequencies, or as an exercise in probabilistic reasoning" (p. 91).

Probability is one thing, likelihood is another. A third notion is confidence. The CIA analyst who told Obama that there was a 95% probability that the Pacer was bin Laden was

[1] Frank Knight, Risk, Uncertainty and Profit (New York: Houghton Mifflin, 1921); John Maynard Keynes The General Theory of Employment, Interest and Money (London: Macmillan and Co., 1936)

[2] Milton Friedman, Price Theory (New Brunswick, NJ: Transaction Publishers, 1962)

giving expression to his confidence that the Pacer was bin Laden. Such expressions of confidence are sometimes described as "subjective probabilities" that can supposedly be deduced by inviting people to bet on various outcomes. Kay and King have little time for the project of assigning probabilities to unique events or for the idea that observing how people gamble gives us insight into rational behaviour. It is also obscure how subjective probabilities relate to frequencies. The CIA analyst who thought that there was a 95% probability that the man in the Abbottabad house was bin Laden was presumably not saying that on 95% of similar occasions bin Laden would be found there.

Human reasoning tends to be narrative rather than probabilistic. Narratives or stories are the means by which we human beings "order our thoughts and make sense of the evidence given to us" (p. 211). Narratives are "an essential part of how we reason" (p. 230), and a "reference narrative" is "a story which is an expression of our realistic expectations" (p. 122). Narrative reasoning is "the most powerful mechanism for organizing our imperfect knowledge" (p. 410), and Obama's decision "was based not on a calculation of probabilities, but on weighing up the credibility and coherence of competing narratives" (p. 277). Narratives are our answer to the question "What is going on here?" When narratives are based on a priori reasoning rather than asking what is going on here, the results can be disastrous. The 2003 American invasion of Iraq might have been less disastrous, or avoided, if ideologues in the Bush administration had asked "What is going on here?" before taking the plunge.

The Iraq fiasco shows that there is no substitute for good judgement. What Kay and King have to say on this subject is reminiscent of Kant as well as Aristotle. There was no rule book for dealing with the situation in Iraq before the American invasion. Even if such a rule book existed, it would still have taken judgement to determine how the rules applied to the case at hand, and a regress threatens if we posit further rules for applying the rules in the rule book. As Kant noted, any rule demands guidance from judgement, and good judgement is a peculiar talent for which there is no substitute. Kay and King's way of making this point is to insist that good judgement "cannot be summarized in twelve rules for life, seven habits of effective people, or even twenty-one lessons for the twenty-first century" (p. 176). Nevertheless, questions remain about how a good judge should go about answering the question "what is going on here?"

There is a natural reading of the latter question on which it cannot be answered in some cases of radical uncertainty in the broad sense without engaging in probabilistic reasoning. Of special relevance is Bayes' theorem, which enables us to calculate *conditional* probabilities, that is to say, the probability that A will happen given that B has happened. There is the prior probability that the Pacer is bin Laden, and the posterior probability that the Pacer is bin Laden given that he is the same height as bin Laden, who was unusually tall. Imagine a Bayesian dial on the wall of the Situation Room. The dial moves with each new piece of information from Abbottabad. Surely we can talk about A being more probable given B than in the absence of B even if we do not regard conditional probabilities as precisely calculable. Kay and King insist that "there is no Bayesian dial" (p. 82), but the positing of such a dial is not so far-fetched when a person is presented with new information that changes the odds of a given proposition (e.g., "The Pacer is bin Laden") being true.[3]

The relevant form of reasoning in such cases is *abductive*. This type of reasoning, which seeks to provide the best explanation of an observed event, plays an important role in our attempts to understand the world. A willingness to engage in such reasoning is indispensable for anyone who wants to answer the question "what is going on here?", and there is every reason to suppose that Obama engaged in such reasoning in deciding what to do about the Pacer. In his recent memoir, Obama describes himself as "working the odds", and one wonders what he would make of the suggestion that his decision about the Abbottabad raid was not based on a calculation of probabilities.[4]

In a case like this, the challenge is to decide between competing narratives. The credibility of a narrative is its consistency with "real or imagined human experience" (p. 217). A narrative is coherent when its components are internally consistent. One narrative about the Pacer was that he was bin Laden. A different narrative was that he was a major drug dealer. Each narrative was internally consistent and consistent with human experience. Terrorists in hiding behave like the Pacer but so do drug barons in hiding. In what sense, then, was the bin Laden narrative more credible and coherent, or a better answer to the question "what is going on here"?

It is at this point that probabilistic reasoning comes into its own. The secretiveness of the Pacer did not favour the first narrative over the second. However, the Bayesian dial swung towards the first narrative given that the house had no landline or internet. Drug barons need to be able to communicate with the outside world. The relevant conditional probabilities might not be precisely calculable, but it still makes sense to describe the probability that the Pacer was bin Laden as higher than the probability that he was a drug dealer.

Sometimes King and Kay try to bring out the limitations of probabilistic reasoning by drawing attention to cases in which "the question is known, but the range of answers in unbounded" (p. 44). With a question like "What will happen in the Middle East in the next five years?" there is, as Keynes noted,

[3] For a more sympathetic account of the use of Bayesian reasoning in the world of intelligence, see David Omand, How Spies Think: Ten Lessons in Intelligence (London: Viking, 2020).

[4] Obama's own account of his thinking about the Abbottabad raid is worth reading in the light of Kay and King's discussion. See his A Promised Land (London: Viking, 2020), chapter 27.

no scientific basis on which to form any calculable probability whatever. It is the open-endedness of the question that leads Kay and King to conclude that in this case "there are no states to which we can sensibly attach probabilities" (p. 44).

Elsewhere, the focus is on the application of probabilistic reasoning to unique events, such as the result of the Kentucky Derby. If I think that the probability that Dobbin will win is 0.9, one interpretation of this is that if the race were to be run 100 times in identical circumstances, then Dobbin would win on 90 occasions. However, since the race will only be run once with the same runners and riders, the statement that "the probability that Dobbin will win is 0.9" is not a claim about frequency but about the speaker's subjective probability. Even if the relationship between frequentist and subjective probabilities is unclear, the latter can in this case be deduced by giving people the opportunity to bet on various outcomes.

A point that is obscured by Kay and King's discussion is that open-endedness and uniqueness are two quite different issues. The range of possible answers to the question "Is the Pacer bin Laden?" is not unbounded. This makes it a rather poor illustration of Keynes' point, which was specifically about open-ended questions. It is true, as Kay and King point out, that rational people will decline to participate in a wager when they do not have enough information to go on. However, sometimes they have no choice but to bet, and the challenge is to work out which way to go. This is ultimately a matter of judgement.

Kay and King make a number of suggestive observations about the nature of judgement, though one would be hard pushed to extract a theory of judgement from *Radical Uncertainty*. Perhaps this is too much to expect. Kay and King stress that "the exercise of judgement in the selection of narratives is eclectic and pragmatic" (p. 397). They are justifiably impressed by "fire chiefs whose judgements are venerated by their crews" (p. 151) and other professionals who have to make life and death decisions under severe time pressure. In these cases, good judgement does not depend on a person's ability to give an accurate account of the basis of their judgements. Rather, it might seem that good judgement is above all a matter of getting things right.

This is not how Kay and King see things. What they refer to as "resulting" mistakenly "judges the quality of a decision by its outcomes" (p. 265). Resulting is a mistake because good decisions can work out badly, while bad decisions can sometimes come out well. This shows that "to judge decision-making under uncertainty we need to review the process of decision-making itself" (p. 267). For example, President Obama is thought to have made a good decision when he sent the SEAL team to raid the bin Laden compound, whereas President Jimmy Carter is widely regarded as having made a bad decision when he authorized a failed attempt to rescue the Tehran hostages in 1979. What explains the difference?

According to Kay and King, Obama made a considered choice of action "by spending time listening to and communicating with experienced and knowledgeable advisers" (p. 277). However, there is no reason to think that Carter did anything different. The Tehran hostage rescue failed because three of the eight helicopters developed faults, and two American aircraft collided in the desert. Carter aborted the mission on the advice of his field commanders, and it is difficult to avoid the conclusion that he was simply less lucky than Obama. The fact that so many commentators continue to contrast Carter's supposedly bad decision with Obama's good one says much about the temptations of "resulting".

In reality, the quality of a decision is both a matter of outcome and the decision-making process. We talk about good judgement when the decision-making process displays a range of intellectual virtues *and* the judgement turns out to be correct. To the extent that it can be a matter of luck whether a judgement turns out to be correct, it can also be a matter of luck whether a person's judgement is good. Like football managers, Presidents and Prime Ministers are in the results business. They are judged by the results of their decisions rather than by how they came to them. If economic forecasters were in the results business, most of them would be out of a job.

Where does this leave Kay and King's account of the role of narratives in our thinking? They rightly cite Walter Fischer, the theorist of communication, whose ground-breaking work in the 1980s on what he called "narrative rationality" remains influential.[5] Fischer insisted that *Homo sapiens* is *Homo narrans*. Human beings are essentially storytellers for whom good reasons take the form of stories which vary in coherence and the extent to which they ring true. Telling stories is above all an exercise in sensemaking. It is a human response to the daunting task of making sense of reality and our place in it.

Kay and King adopt several aspects of Fischer's conception of narrativity, but questions remain about their conception of a reference narrative as a story which is an expression of our realistic expectations. The Bush administration had a reference narrative about Iraq, but its reference narrative did not consist of *realistic* expectations about the post-invasion scenario. Reference narratives do not have to be realistic. Reference narratives are not just, or even primarily, expressions of a person's expectations about the future. It makes more sense to think of them as stories that embody a set of fundamental principles or assumptions that shape one's thinking and planning. These assumptions embody one's conception of self, world, and one's place in the world.

The Bush reference narrative was that the world is divided into good and evil and that the USA is a force for good that is

[5] Walter R. Fischer, Human Communication as Narration: Towards a Philosophy of Reason, Value, and Action (Columbia: University of South Carolina Press, 1987)

entitled to use military force to improve the human condition. Rather than giving expression to the Bush administration's expectations, this narrative was a world view that *explained* its unrealistic expectations, including the unrealistic expectation that the USA would be seen by ordinary Iraqis as liberators and harbingers of liberal democracy in Iraq. Kay and King are critical of narratives based on a priori principles, but reference narratives are neither a priori nor empirical. They are closer to what Wittgenstein called "hinge propositions", that is, propositions that stand fast for us and shape the rest of our thinking and empirical enquiry.[6]

As well as characterizing reference narratives as expressions of our realistic expectations, Kay and King describe them as attempts to construct the best explanation of the complex world "from a myriad of little details and the knowledge of context derived from personal experience and the experience of others" (p. 410). However, experiences need to be interpreted, and every explanation is shaped by background assumptions. These background assumptions, rather than the explanations they shape, are the reference narrative. Without a reference narrative to control the interpretation of our experiences and inform our explanations, abductive reasoning would get us nowhere. The question "what is going on here?" can only be answered if the person who asks the question has a reference narrative that makes the question, and possible answers to it, meaningful to them.

As Kay and King note, reference narratives are not immune to challenge, though rejection of a reference narrative is a major step – something like a Kuhnian paradigm shift. A first-rate decision-maker confronted by radical uncertainty needs to "organise action around a reference narrative, while being open to the possibility that this narrative is false" (p. 285). Bush and his advisers failed to recognize the possibility that their reference narrative might be false. In addition, their narrative was neither robust nor resilient. Risk, as Kay and King understand it, is "failure of a projected narrative, derived from realistic expectations, to unfold as envisaged" (p. 123). The key to managing risk is to ensure that one's reference narrative has the properties of robustness and resilience.

How should these properties be understood? On one interpretation, robustness and resilience are about being adequately prepared for mishaps or things not working out as expected. When we plan a holiday, we expect a relaxing and enjoyable experience. If our narrative is robust and resilient, we think about the things which might go wrong and prepare for them; we leave ourselves extra time to reach the airport and take a supply of medication. On this account, a robust and resilient narrative is one that incorporates adequate contingency planning. However, such planning involves probabilistic reasoning. One does not take anti-malaria medication if one calculates that the probability of encountering mosquitos where one is going is low.

On a different view, robustness and resilience are epistemological notions. Epistemologically robust and resilient narratives are ones that, in virtue of their credibility, coherence, and other epistemological merits, are unlikely to be derailed. Here, as elsewhere, it is natural to think in terms of probabilities. The Bush administration's contingency planning for Iraq was so poor because it underestimated the probability of things going wrong. Perhaps it also failed to ask the question "what is going on here?". However, if it had asked this question and tried seriously to answer it, it would have found it difficult to avoid probabilistic reasoning.

Whether this is a problem for Kay and King depends on how radical their view is. Many of their criticisms of overreliance on flawed probabilistic reason are well justified, but there is the suspicion that in places their criticisms go too far. It remains an open question what role probabilistic reasoning can or should play in answering the question "what is going on here?". Whatever one makes of Kay and King's approach to this question, one thing is certain: this is a major contribution to our understanding of human understanding and decision-making. It should be read not just by economists and philosophers but by anyone who is interested in how to make good decisions in an uncertain world.

Quassim Cassam is Professor of Philosophy at the University of Warwick, UK, and author of several books, including most recently *Vices of the Mind* (2019, Oxford University Press) and *Conspiracy Theories* (2020, Polity).

[6] Ludwig Wittgenstein, On Certainty (Oxford: Basil Blackwell, 1969)