

*Also by John Kay and Mervyn King*

**The British Tax System**

THE BRIDGE STREET PRESS

First published in Great Britain in 2020 by The Bridge Street Press

Copyright © John Kay and Mervyn King 2020

The moral right of the authors has been asserted.

All rights reserved.

No part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form or by any means, without the prior permission in writing of the publisher, nor be otherwise circulated in any form of binding or cover other than that in which it is published and without a similar condition including this condition being imposed on the subsequent purchaser.

Illustration credits: p. 104 Bank of England; p. 142 Shutterstock;  
p. 143 WikiArt and Hulton Archive/Getty Images.

A CIP catalogue record for this book  
is available from the British Library.

ISBN 978-1-4087-1258-0

The Bridge Street Press  
An imprint of  
Little, Brown Book Group  
Carmelite House  
50 Victoria Embankment  
London EC4Y 0DZ

An Hachette UK Company  
[www.hachette.co.uk](http://www.hachette.co.uk)

[www.littlebrown.co.uk](http://www.littlebrown.co.uk)

# CONTENTS

Acknowledgements

Preface

## **Part I**

### **Introduction: The Nature of Uncertainty**

1 The Unknowable Future

2 Puzzles and Mysteries

3 Radical Uncertainty is Everywhere

## **Part II**

### **The Lure of Probabilities**

4 Thinking with Probabilities

5 A Forgotten Dispute

6 Ambiguity and Vagueness

7 Probability and Optimisation

## **Part III**

### **Making Sense of Uncertainty**

8 Rationality in a Large World

9 Evolution and Decision-making

10 The Narrative Paradigm

11 Uncertainty, Probability and the Law

12 Good and Bad Narratives

13 Telling Stories Through Numbers

[14 Telling Stories Through Models](#)

15 Rationality and Communication

[16 Challenging Narratives](#)

## **Part IV**

### **Economics and Uncertainty**

17 The World of Finance

[18 Radical Uncertainty, Insurance and Investment](#)

[19 \(Mis\)Understanding Macroeconomics](#)

[20 The Use and Misuse of Models](#)

## **Part V**

### **[Living with Uncertainty](#)**

[21 Practical Knowledge](#)

[22 Adapting to Radical Uncertainty](#)

23 Embracing Uncertainty

Appendix: Axioms of Choice under Uncertainty

[Notes](#)

[Bibliography](#)

Further Reading

[Index](#)

*I prefer true but imperfect knowledge, even if it leaves much indetermined and unpredictable, to a pretence of exact knowledge that is likely to be false.*

—FRIEDRICH VON HAYEK,  
Nobel Prize lecture, 1974

## ACKNOWLEDGEMENTS

As the final words of this book were written in the summer of 2019 its two authors celebrated, if that is the word, fifty years of life as professional economists. Much has happened to economies and the study of economics over that period. The first Nobel Prize in Economic Sciences was awarded in 1969, and the announcement of the latest winner has become a regular part of the economist's calendar. In public, private and academic sectors, the number of economists has expanded greatly. We have had the good fortune to work with colleagues in all three domains. As we describe in this book, progress is made collectively and through communication with others. Our biggest debt of gratitude must, therefore, be to our professional colleagues, too numerous to mention, from around the world with whom we have interacted for half a century. Their influence on us, both conscious and unconscious, is immense.

Our immediate thanks are to those with whom we have discussed the ideas in this book during its writing: Rachel Barkow, Tim Besley, Amar Bhidé, David Bodanis, Alan Budd, Paul Collier, Samuel Issacharoff, Peter Kellner, Richard Pildes, Stuart Proffitt, Adam Ridley, Paul Seabright, Robert Skidelsky, Ed Smith, David Tuckett, and colleagues at NYU Stern School of Business and NYU Law School, as well as at the London School of Economics and the Bank of England. Some of the ideas were presented over several years to seminar participants at a number of places including All Souls College, Oxford, the Norwegian School of Economics in Bergen, NYU Law School and Yale Law School.

Others have made this venture feasible. Our personal assistants, Rachel Lawrence in England and Gail Thomas in New York, organised our lives to make it possible to work on the book. Andrew Wylie proved once again what an energetic and supportive literary agent he is. And our publishers – Tim Whiting and Zoe Gullen at The Bridge Street Press, and Drake McFeely at W.W. Norton – provided encouragement and advice at every stage.

Doris Nikolic and Matthew Ford not only provided excellent research assistance but persisted in asking difficult questions to which we felt impelled to try to provide answers. Their insights have greatly enriched the book.

Finally, our thanks must go to our long-suffering spouses, Mika Oldham and Barbara Melander-King, who not only gave us excellent comments and editorial suggestions but patiently allowed us the time to complete the manuscript.

## PREFACE

Forty years ago, we wrote a well-received book, *The British Tax System*, describing the failures – intellectual and practical – of the tax system. We were neither academic scribblers inventing a tax system from scratch nor tax accountants engrossed in excruciating detail. Instead, as young academics we set out to look carefully at how the tax system actually worked in practice and then to design improvements based on a small number of carefully thought through principles. Forty years later, we discovered that we had independently come to the view that economics as a whole faced a similar challenge and was in need of a fresh look. So we decided to collaborate again. This book is the result.

*The British Tax System* sold well and went into several editions. But our careers then went in different directions. John became the Director of the Institute for Fiscal Studies, started a successful consulting company focusing on business economics, and was the first director of the Saïd Business School at Oxford University and for twenty years a columnist for the *Financial Times*. Mervyn became an academic in various universities in the UK and US before joining the Bank of England as Chief Economist and later Governor from 2003 to 2013.

During those forty years we saw at first hand both the power of economics as a way of approaching practical problems, and also its limitations. As students and academics we pursued the traditional approach of trying to understand economic behaviour through the assumption that households, businesses, and indeed governments take actions in order to optimise outcomes. We learnt to approach



economic problems by asking what rational individuals were maximising. Businesses were maximising shareholder value, policy-makers were trying to maximise social welfare, and households were maximising their happiness or 'utility'. And if businesses were not maximising shareholder value, we inferred that they must be maximising something else – their growth, or the remuneration of their senior executives.

The limits on their ability to optimise were represented by constraints: the relationship between inputs and outputs in the case of businesses, the feasibility of different policies in the case of governments, and budget constraints in the case of households. This 'optimising' description of behaviour was well suited to the growing use of mathematical techniques in the social sciences. If the problems facing businesses, governments and families could be expressed in terms of well-defined models, then behaviour could be predicted by evaluating the 'optimal' solution to those problems.

Although much can be learnt by thinking in this way, our own practical experience was that none of these economic actors were trying to maximise anything at all. This was not because they were stupid, although sometimes they were, nor because they were irrational, although sometimes they were. It was because an injunction to maximise shareholder value, or social welfare, or household utility, is not a coherent guide to action. Business people, policy-makers and families could not even imagine having the information needed to determine the actions that would maximise shareholder value, social welfare or household utility. Or to know whether they had succeeded in doing so after the event. Honest and capable executives and politicians, of which there are many, try instead to make incremental decisions which they think will improve their business, or make the world a better place. And happy households are places where family members work together to ensure that tomorrow is at least as good as today.

Most economists would readily acknowledge that no one actually engages in the kinds of calculation which are described in economic models. But since the work of Paul Samuelson, economists have relied on the claim that if people observed certain

axioms which constituted 'rationality' they would – unconsciously – be optimising, rather as Molière's M. Jourdain had been talking prose for forty years without knowing it. And when this axiomatic approach is applied to consumer behaviour, as it was by Samuelson, the method is more fruitful than the sceptical observer might expect.

But we show in this book that the axiomatic approach to the definition of rationality comprehensively fails when applied to decisions made by businesses, governments or households about an uncertain future. And this failure is not because these economic actors are irrational, but because they *are* rational, and – mostly – do not pretend to knowledge they do not and could not have. Frequently they do not know what is going to happen and cannot successfully describe the range of things that *might* happen, far less know the relative likelihood of a variety of different possible events.

The financial crisis of 2007–08 brought home the intellectual failures of optimising models to capture the disruptive behaviour that results from confronting an unknowable future. But this is not another book about that financial crisis, or even another book about economics, although we believe that the implications for the study of economics are considerable. It is a book about how real people make choices in a radically uncertain world, in which probabilities cannot meaningfully be attached to alternative futures.

As we wrote this book, and discussed our ideas with friends and colleagues, we encountered very different reactions from general readers, on the one hand, and specialists, on the other. Most people find the concept of radical uncertainty natural and indeed obvious. For them, the challenge is not to accept the existence of radical uncertainty but to find ways of coping with it. We hope they will find answers to that challenge in the chapters that follow. Many people who have been trained in economics, statistics or decision theory, however, find it difficult to accept the centrality of radical uncertainty. And to these we need to add some who work in computer science and artificial intelligence – or who have simply read enough about these things to be caught up in the wave

of popular enthusiasm for the style of reasoning at which computers excel.

In trying to persuade those two different audiences of the importance of radical uncertainty, the risk is that one thinks we are flogging a dead horse; the other that we are flogging the winner of the Kentucky Derby by decrying a set of techniques which has transformed our thinking in economics, statistics, decision-making and artificial intelligence. We hope that general readers will nevertheless enjoy the spectacle of the flogging and that specialists will feel at least some of the sting of the lash.

# Part I

## Introduction: The Nature of Uncertainty

# 1

---

## THE UNKNOWABLE FUTURE

All we can know is that we know nothing. And that is the sum total of human wisdom.

—LEO TOLSTOY, *War and Peace*, 1867<sup>1</sup>

### **Unknownable unknowns**

In September 1812, Napoleon's troops defeated the Russian forces at Borodino, and cleared the way for the capture of Moscow. The French entered the now-deserted capital and razed its wooden structures to the ground. But the fruitless occupation marked the end of the Emperor's campaign of conquests. The freezing, starving, disease-ridden Grande Armée began its long retreat to Paris. Most of its soldiers never made it home.<sup>2</sup> Napoleon abdicated in April 1814.

Napoleon was the greatest political and military leader of the age, at the height of his power and achievement, and commander of the largest army the world had seen.<sup>3</sup> But he had little understanding of what was going on at Borodino. Nor, even today, do we truly understand why he was there. And the complex, many-faceted relationship between Russia and Western Europe remains unresolved two centuries later.

On 3 August 1492, Christopher Columbus sailed from Spain with the hope of finding a new route to the Indies. Most experienced sailors at the time believed that a westerly route to Asia was impracticable, given the distance involved and the problem of carrying sufficient food and water. And they were right. Against all reasonable advice, the Spanish Crown agreed to sponsor his expedition. Columbus did not know what his journey would entail, or how long it would take him, and after re-stocking in the Canary Islands, he landed in the Bahamas. He did not know that the New World, as it came to be called, even existed, and he did not know what he had found even after he had found it. He maintained that he had indeed landed in Asia, which is why America is named after his contemporary, Amerigo Vespucci, who had a better idea of where his own explorations had led. Whatever counted as cost-benefit analysis in the Spanish court took no account of the possibility of a New World; nor could it.

In February 1972, Richard Nixon met the Chinese leader Mao Tse-tung in Beijing. The meeting had long been planned, secretly, by the US President and his national security adviser, Henry Kissinger. Its purpose was by no means clear, although Nixon was struggling to extricate his country from Vietnam and both American and Chinese leaders wished to put distance between the two leaders of global communism, China and the Soviet Union. The much-feted meeting ended with a bland communiqué characteristic of global summits. Later that year, five men were arrested during a break-in at the Watergate complex in Washington and the subsequent cover-up led to Nixon's resignation in 1974. Two years after that, the ailing Mao died.

No one knew what the consequences of the Nixon/Mao meeting would be – it was not even certain when Nixon's plane arrived that Mao himself would meet the President. And, almost half a century later, no one really knows what the consequences were. Was its significance merely symbolic, a media event involving two leaders in failing political or physical health? Or was it a key milestone in the integration of China into the global economy, probably the most important economic development of the subsequent half century?

Emperors, explorers and presidents made decisions without fully understanding either the situation they faced or the effects of their actions. And so must we.

## **The global financial crisis**

On 9 August 2007, the French bank BNP Paribas announced that it was suspending transactions in three of its funds. The holdings of investors, which had been placed in securities linked to the US housing market, were effectively frozen. The failure of a small group of hedge funds was in itself a minor event. But within days the global financial crisis of 2007–08 was underway. That crisis peaked in September 2008 with the failure of Lehman Brothers, and only extraordinary efforts by central banks around the world prevented the collapse of the Western financial system. Nevertheless, the failures in financial markets which followed led to the most severe recession in the industrialised world since the Great Depression of the 1930s.

As the days passed and the position of the banks deteriorated, we were seeing events that had not been observed for several generations, if at all. Why was the banking system experiencing a crisis of confidence of the kind that we thought had been relegated to history? As the crisis passed and there was time for reflection, both authors sought to explain the background to the events that took us all by surprise. Mervyn King wrote *The End of Alchemy* and John Kay *Other People's Money*. From differing perspectives, the authors had arrived at a common view. The narrative of the pre-crisis period was falling apart. Markets in new, complex financial instruments were supposed to ensure that risks were placed with people who were best able to bear them. Or so the story had gone. As the crisis evolved, a new narrative was required; one which recognised that humans do not necessarily comprehend the consequences of their innovations. The risks had not been placed with those who had most capacity to handle them, but with those who did not understand them. And the institutions that sold those instruments understood little more and lacked the financial capacity to hold them when the market dried up.





or in charge of public bodies, such as governments and armies?

## **The war on terror**

Osama bin Laden masterminded the attack on Manhattan's World Trade Center on 11 September 2001, and US military preparations for a possible invasion of Iraq began almost immediately thereafter. In February of the following year, Donald Rumsfeld, the United States Secretary of Defense, gave a press briefing. Rumsfeld was asked to comment on reports that no evidence linked Baghdad to terrorist activity. His famous response was widely derided: 'There are known knowns; there are things we know we know. We also know there are known unknowns; that is to say we know there are some things we do not know. But there are also unknown unknowns – the ones we don't know we don't know.'<sup>6</sup>

Yet Rumsfeld was saying something important.<sup>7</sup> The follow-up question to Rumsfeld's musings is less well remembered than the observation that provoked it. The Defense Secretary was asked in which category – known knowns, known unknowns or unknown unknowns – did intelligence about terrorism and weapons of mass destruction fall? Rumsfeld's response was 'I am not going to say which it is'.<sup>8</sup> But no link between Iraq and the 9/11 attack has been established, and no weapons of mass destruction were found. The Iraq invasion was a military success: US forces rapidly reached Baghdad and toppled Saddam Hussein. But with hindsight the operation was a comprehensive failure of intelligence, judgement and contingency planning; a failure as politically damaging as the similar failures of intelligence, judgement and contingency planning in the financial sector were economically damaging. In the light of these failures, US government agencies were required to implement a more structured process for providing advice to the President.<sup>9</sup> Analysts were expected to quantify their confidence levels and express them as probabilities.

And so almost a decade later, in spring 2011, President Barack Obama met with his senior security advisers in the White House situation room to consider what he knew would be one of the defining decisions of his presidency. Should he approve the

proposed raid by US Navy SEALs on the Abbottabad compound in Pakistan where Osama bin Laden was believed to be hiding? Obama was well aware that a similarly daring plan in 1979 to rescue hostages from the Tehran embassy had ended in fiasco and may have cost Jimmy Carter a second term as President. 'John', the CIA team leader, was 95% certain that bin Laden was in the compound. But others were less sure. Most placed their probability estimate at about 80%. Some were as low as 40% or even 30%.

The President summed up the discussion. 'This is 50-50. Look guys, this is a flip of the coin. I can't base this decision on the notion that we have any greater certainty than that.'<sup>10</sup> Obama did not mean that the probability that the man in the compound was bin Laden was 0.5; still less that he planned to decide by tossing a coin.<sup>11</sup> His summary recognised that he had to make his decision without knowing whether the terrorist leader was in the compound or not. Obama would reflect on that discussion in a subsequent interview: 'In this situation, what you started getting was probabilities that disguised uncertainty as opposed to actually providing you with more useful information.'<sup>12</sup>

Either bin Laden was there, or he wasn't - though the plan involved many other risks and uncertainties, military, technological and political. Obama might have said to his colleagues something like 'if you guys can tell me there is a greater than 60% chance that the man in the compound is bin Laden, I will give the order to go ahead'. But he did not; to have done so would have been to pass responsibility for the decision from the Oval Office, where it belonged, to the intelligence agencies. But such passing on of responsibility did occur in the financial sector, where senior bank executives, such as Mr Viniar, had effectively delegated the management of uncertainty to the risk professionals and their models. Obama understood this issue; Viniar and his colleagues did not.

In the run-up to the global financial crisis, as in the preparations for the Abbottabad raid, policy advisers not only chose to express uncertainties as probabilities but were required to do so. Regulators prescribed the risk models used by financial institutions; Congress insisted on the quantification of judgements

based on intelligence reports. In both finance and politics, this expression of uncertainty was at best unhelpful and at worst actively misleading. The precision of the numbers presented to Viniar and Obama was spurious. Obama understood that he had to arrive at a decision on the basis of limited information, and made what, with hindsight, proved a good call. He did so not by probabilistic reasoning but by asking ‘What is going on here?’

In writing this book, we found inspiration in an anecdote from Richard Rumelt’s *Good Strategy/Bad Strategy*, by some distance the best book on business strategy written in the last decade. Rumelt describes a conversation with a colleague at UCLA who had observed some of his case-based MBA classes:

We were chatting about pedagogy . . . John gave me a sidelong look and said ‘it looks to me as if there is really only one question you’re asking in each case’. That question is ‘what’s going on here?’ John’s comment was something I’ve never heard said explicitly but it was instantly and obviously correct. A great deal of strategy work is trying to figure out what is going on. Not just deciding what to do, but the more fundamental problem of comprehending the situation.<sup>13</sup>

The question ‘What is going on here?’ sounds banal, but it is not. In our careers we have seen repeatedly how people immersed in technicalities, engaged in day-to-day preoccupations, have failed to stand back and ask, ‘What is going on here?’ We have often made that mistake ourselves.

## **Retirement planning**

Most readers will never be called upon to make decisions as momentous as those which Obama and Viniar faced – although they may be asked to judge the performance of those who make these critical decisions. But all of us make choices which require thinking about the future. House purchase and preparing for retirement are the most important financial decisions for most households. But few plan in any organised way. They decide, often

rather quickly, which house to buy based on a reaction rather than a checklist. Given choices – to opt in or out of pension schemes, to select one investment fund over another – they often prefer to avoid choices. The default option – which requires no action – is usually most popular.<sup>14</sup>

A few households approach the issue of retirement more systematically. There are software programs to help; some asset managers<sup>15</sup> offer free assistance and some economists have written commercial programs.<sup>16</sup> These calculate how much an individual or household should save and whether they are on track for a comfortable retirement. To start using these programs you must have a lot of information at your side: current facts, such as age and marital status, and information about your future self, such as the age at which you plan to retire, and how much you will then need to spend each year. You will need to disclose not just your current salary but what you expect it to be many years from now. You may even be asked to predict how long you will live. The principal uncertainty President Obama faced was binary – either bin Laden was there in Abbottabad, or he was not. Judging how long you will live is more difficult. You know that you are most unlikely to reach 125 years old and you can consult the carefully constructed life tables used by the actuaries of insurance companies. But the program may also ask you to express views on the economy and financial markets in the future. A historic series on inflation and investment returns will give you some sense of possible answers. You might sketch a range of possibilities even if you would struggle – and be unwise – to come up with an answer.

One answer to such questions the programs generally do not allow is ‘I don’t know’. But that is the answer which the authors, professional economists all our lives, would give to many of these questions. And frankly we think that is the answer you should give as well. There are many questions – such as ‘Who is the man in the compound in Pakistan?’ – to which the only sensible answer is ‘I don’t know’; and if this may be true of the present, it is even more true of the future. We all have to make decisions about retirement plans even though we do not know what our income or the inflation rate will be twenty years from now. If we knew the

answers to all the questions the programs pose, we could determine exactly how much we need to save to enjoy our retirement. But we do not make good decisions by professing knowledge we do not and cannot have. After we have consulted the programs, we at least know what we don't know, and we may be able to do a little to reduce our ignorance.

## **An intellectual failure**

The crisis of 2007–08 represented – obviously – a failure of economic analysis and economic policy. But while recognising the seriousness and cost of the financial crisis, economists have generally been reluctant to accept that their intellectual framework is in need of revision. Economists (used to) distinguish risk, by which they meant unknowns which could be described with probabilities, from uncertainty, which could not. They had already adopted mathematical techniques which gave the term ‘risk’ a different meaning from that of everyday usage. In this book we will describe the considerable confusion and economic damage which has arisen as a result of the failure to recognise that the terms ‘risk’, ‘uncertainty’ and ‘rationality’ have acquired technical meanings in economics which do not correspond to the everyday use of these words. And over the last century economists have attempted to elide that historic distinction between risk and uncertainty, and to apply probabilities to every instance of our imperfect knowledge of the future.

The difference between risk and uncertainty was the subject of lively debate in the inter-war period. Two great economists – Frank Knight in Chicago and John Maynard Keynes in Cambridge, England – argued forcefully for the continued importance of the distinction. Knight observed that ‘a measurable uncertainty, or “risk” proper, as we shall use the term, is so far different from an unmeasurable one that it is not in effect an uncertainty at all’.<sup>17</sup>

Keynes made a similar distinction. In an article summarising his magnum opus, *The General Theory of Employment, Interest and Money*, he wrote:

By 'uncertain' knowledge, let me explain, I do not mean merely to distinguish what is known for certain from what is only probable. The game of roulette is not subject, in this sense, to uncertainty; nor is the prospect of a Victory bond being drawn. Or, again, the expectation of life is only slightly uncertain. Even the weather is only moderately uncertain. The sense in which I am using the term is that in which the prospect of a European war is uncertain, or the price of copper and the rate of interest twenty years hence, or the obsolescence of a new invention, or the position of private wealth-owners in the social system in 1970. About these matters there is no scientific basis on which to form any calculable probability whatever. We simply do not know.<sup>18</sup>

The title of this book, and its central concept, is *radical uncertainty*. Uncertainty is the result of our incomplete knowledge of the world, or about the connection between our present actions and their future outcomes. Depending on the nature of the uncertainty, such incomplete knowledge may be distressing or pleasurable. I am fearful of the sentence the judge will impose, but look forward to new experiences on my forthcoming holiday. We might sometimes wish we had perfect foresight, so that nothing the future might hold could surprise us, but a little reflection will tell us that such a world would be a dull place.

We have chosen to replace the distinction between risk and uncertainty deployed by Knight and Keynes with a distinction between *resolvable* and *radical* uncertainty. Resolvable uncertainty is uncertainty which can be removed by looking something up (I am uncertain which city is the capital of Pennsylvania) or which can be represented by a known probability distribution of outcomes (the spin of a roulette wheel). With radical uncertainty, however, there is no similar means of resolving the uncertainty – we simply do not know. Radical uncertainty has many dimensions: obscurity; ignorance; vagueness; ambiguity; ill-defined problems; and a lack of information that in some cases but not all we might hope to rectify at a future date. These aspects of uncertainty are the stuff of everyday experience.

Radical uncertainty cannot be described in the probabilistic terms applicable to a game of chance. It is not just that we do not

know what will happen. We often do not even know the kinds of things that might happen. When we describe radical uncertainty we are not talking about ‘long tails’ – imaginable and well-defined events whose low probability can be estimated, such as a long losing streak at roulette. And we are not only talking about the ‘black swans’ identified by Nassim Nicholas Taleb – surprising events which no one could have anticipated until they happen, although these ‘black swans’ are examples of radical uncertainty.<sup>19</sup> We are emphasising the vast range of possibilities that lie in between the world of unlikely events which can nevertheless be described with the aid of probability distributions, and the world of the unimaginable. This is a world of uncertain futures and unpredictable consequences, about which there is necessary speculation and inevitable disagreement – disagreement which often will never be resolved. And it is that world which we mostly encounter. So the ramifications of radical uncertainty go well beyond financial markets; they extend to individual and collective decisions, as well as economic and political ones; and from decisions of global significance taken by statesmen to everyday decisions taken by the readers of this book.

For both Knight and Keynes, recognition of the pervasive nature of radical uncertainty was essential to an understanding of how a capitalist economy worked. Knight believed that it was radical uncertainty that created profit opportunities for entrepreneurs and that it was their skill and luck in navigating radical uncertainty which drove technical and economic progress. Fifteen years before the *General Theory*, Keynes had published *A Treatise on Probability*, and an appreciation of his evolving views on risk and uncertainty is necessary in interpreting his later work. But in the *General Theory* he re-expressed Knight’s thinking with characteristic literary flourish: ‘If the animal spirits are dimmed and the spontaneous optimism falters, leaving us to depend on nothing but the mathematical expectation, enterprise will fade and die.’<sup>20</sup> Keynes’ concern was less with the micro-economic drivers of innovation than with the macroeconomic factors behind the Great Depression. In his view it was (non-mathematical) expectations – ‘the state of confidence’ – which made it difficult to

achieve or restore the equilibrium which the classical economists had described.

But Keynes and Knight lost the battle to put radical uncertainty at the heart of economic analysis (in chapter 5 we explain why). Most economists today pay – at best – lip service to the difference between risk and uncertainty. The problem of radical uncertainty has supposedly been tamed by probabilistic reasoning. That belief has infected other areas of social science, including statistics, sociology and psychology, and even the law.

And so instead of recognising radical uncertainty, and adopting policies and strategies that will be robust to many alternative futures, banks and businesses are run with reliance on models which claim knowledge of the future that we do not have and never could have. Those models attempt to manage uncertainty by assuming that the analysis of commercial and financial risk is analogous to the analysis of roulette. We do not know how a particular spin will fall, but we do know the possible outcomes and the frequency of each of those outcomes if we played the game over and over. But uncertainty takes many forms, few of which can be represented in this way.

Three main propositions run through this book. First, the world of economics, business and finance is ‘non-stationary’ – it is not governed by unchanging scientific laws. Most important challenges in these worlds are unique events, so intelligent responses are inevitably judgements which reflect an interpretation of a particular situation. Different individuals and groups will make different assessments and arrive at different decisions, and often there will be no objectively right answer, either before or after the event. And because what we observe is not the outcome of a stationary process, conventional statistical inference rarely applies and forecasts are often based on shifting sands.

Second, individuals cannot and do not optimise; nor are they irrational, victims of ‘biases’ which describe the ways they deviate from ‘rational’ behaviour. The meaning of rational behaviour depends critically on the context of the situation and there are generally many different ways of being rational. We distinguish



axiomatic rationality, as used by economists, from evolutionary rationality, as practised by people. Many so-called ‘biases’ are responses to the complex world of radical uncertainty. Evolution in this uncertain world has led characteristics which are primarily adaptive to become embodied in human reasoning. Humans are successful at adapting to the environment in which they find themselves, and have not evolved to perform rapid calculations of well-defined problems at which computers excel. This is because the problems which humans face, whether sparkling at dinner party conversations or conducting international trade negotiations, are not well-defined problems amenable to rapid calculation.

Third, humans are social animals and communication plays an important role in decision-making. We frame our thinking in terms of narratives. And able leaders – whether in business, in politics, or in everyday life – make decisions, both personal and collective, by talking with others and being open to challenge from them. Humans, uniquely, produce artefacts of extraordinary complexity and are able to do so only by the successful development of networks of trust, cooperation and coordination. Market economies function only by virtue of being embedded in a social context.

Sensible – adaptive – public policy and business strategy cannot be determined by quantitative assessments of policies and projects, made by an industry of professional modellers using probabilistic reasoning. In this book we explain how it is that so many clever people came to believe otherwise – and why they are wrong. We reassert the distinction between risk and uncertainty, and suggest that if we control risk we can not only manage but positively enjoy uncertainty. If that seems paradoxical, read on.

---

## PUZZLES AND MYSTERIES

The worst historian has a clearer view of the period he studies than the best of us can hope to form of that in which we live. The obscurest epoch is to-day.

—ROBERT LOUIS STEVENSON<sup>1</sup>

In August 2004, NASA launched the probe MESSENGER from Cape Canaveral. Although Mercury is on average ‘only’ 60 million miles from Earth,<sup>2</sup> the rocket travelled 4.9 billion miles at a speed of 84,500 mph before it finally entered its investigative orbit of the planet, according to plan, in March 2011.<sup>3</sup>

This remarkable feat of computation was possible because:

- the equations of planetary motion have been comprehensively understood since the seventeenth century, thanks to Johannes Kepler and his successors;
- the equations of planetary motion are stationary, in the sense that those equations governed their motion for millions of years before Kepler’s discoveries and have continued to govern them since (‘stationary’ is a technical term in mathematics and statistics and does not relate to the movements of the planets themselves but to the underlying determinants of planetary motion, which do not change over time: we will make frequent

use of the word 'stationary' in this sense);<sup>4</sup>

- the motion of the planets is not significantly affected by human actions<sup>5</sup> or at all by human beliefs about their motion.

Accurate calculations, such as those undertaken by NASA scientists, can be made when the underlying process is more or less completely understood, when that process remains constant over time, and when the process is independent of our actions and beliefs. And extremely detailed forward planning – in this case, mapping the trajectory of a probe moving at rocket speed for years ahead – is then possible. MESSENGER entered its orbit of Mercury at the exact location NASA had anticipated six and a half years earlier.

## **Puzzles and mysteries**

NASA is a product of its time. We live in an age of enlightenment in which scientific reasoning has displaced argument from authority, whether religious or secular. Scientific evidence is the new authority. But is it possible to extend to other disciplines the methods of analysis which have led to so much progress in natural science? Can there be laws of human behaviour analogous to laws of physics?

Is anything, in either the physical or the political world, really random? 'God does not play dice,' said Einstein, expressing a belief that the world is, fundamentally, deterministic.<sup>6</sup> And, at some deep and unimaginable level of understanding, this may be true. But whatever the playwright of the universe intended, we actors are faced with uncertainty, either because of our ignorance or because of the changing nature of the underlying processes.

The risk managers at Goldman Sachs and the policy advisers at intelligence agencies whom we met in chapter 1 expressed their assessments in probabilistic language. In both cases these assessments were unhelpful, though for different reasons (at Goldman because the probabilistic estimates were taken seriously; in the White House because they were not). In neither case did the probabilities which were expressed provide the information the

decision-makers required. Where there is no adequate basis for formulating probabilities – and there was none at 200 West Street<sup>7</sup> or 1600 Pennsylvania Avenue – we face radical uncertainty.

The probabilists of financial institutions and intelligence agencies believed such conditions of radical uncertainty were rare and that they could estimate probabilities for most of the relevant contingencies. Since the seventeenth century it has been increasingly common to express uncertainty in probabilistic terms. The ‘probabilistic turn’ gathered pace in the twentieth century, and in the last two decades probabilistic reasoning has almost completely dominated the description and analysis of decision-making under uncertainty.

Other writers have made similar distinctions to the ‘known’ and ‘unknown’ unknowns described by Donald Rumsfeld. Greg Treverton, chairman of President Obama’s National Intelligence Council and for many years a senior figure in the US intelligence community, stressed the difference between ‘puzzles and mysteries’.<sup>8</sup> A puzzle has well-defined rules and a single solution, and we know when we have reached that solution. Puzzles deliver the satisfaction of a clear-cut task and a correct answer. Even when you can’t find the right answer, you know it exists. Puzzles can be solved; they have answers. But the solutions may be difficult to find. Economists have thrived on the difficulty of solving complex models of the economy precisely because they have been trained to tackle well-defined problems which have an answer. And (Nobel) prizes are awarded to those who solve the most difficult puzzles.

Mysteries offer no such clarity of definition, and no objectively correct solution:<sup>9</sup> they are imbued with vagueness and indeterminacy. We approach mysteries by asking ‘What is going on here?’, and recognise that even afterwards our understanding is likely to be only partial. They provide none of the comfort and pleasure of reaching the ‘right’ answer. Columbus thought he had landed in Asia. And even today, ‘What was going on here?’ in the global financial crisis, or during bin Laden’s sojourn in Pakistan, is hotly contested. What will be the future of the Middle East? Or the development of mobile computing, or the automobile industry?

Will banks as we know them survive? What is the future of capitalism, or democracy? A mystery cannot be solved as a crossword puzzle can; it can only be framed, by identifying the critical factors and applying some sense of how these factors have interacted in the past and might interact in the present or future. Puzzles may be more fun, but in our real lives the world increasingly offers us mysteries – either because the outcome is unknowable or because the issue itself is ill defined.

The political scientist Philip Tetlock has for three decades studied the performance of ‘expert’ forecasters, with mostly dispiriting results.<sup>10</sup> To find objective measures of the quality of expert judgement – and identify the determinants of good and bad judgement – Tetlock needs to specify problems with verifiable outcomes. In 2010 and 2011, the questions he set were of the kind, ‘Will Serbia be officially granted European Union candidacy by 31 December 2011?’ and ‘Will Italy restructure or default on its debts by 31 December 2011?’<sup>11</sup>

But these well-defined short-term questions are not really the questions to which policy-makers seek answers. More important is to know whether the United States and China will find a peaceful solution to the growing trade and military tensions between them. Or whether the European Union will continue its expansion, and what will be the shape of monetary union in five years’ time. Replacing complex mysteries with puzzles that have unambiguously right and wrong answers limits the interest and relevance of both problems and answers. While there are some problems for which the quantification of probabilities is an indispensable guide to solutions, most decisions in business, finance, politics and personal development, and their outcomes, are too complex and imprecisely defined to be approached in this way. They are subject to radical uncertainty.

## **Radical uncertainty and practical knowledge**

Treverton’s striking distinction between puzzles and mysteries is reproduced wherever practical decision-making is required. The urban planners Horst Rittel and Melvin Webber observed in 1973

that although the well-defined needs of communities for roads, sanitation, etc. had been met, their clients remained dissatisfied. More was needed; but the planners did not know what it was and the population was not able to articulate clearly its needs. Rittel and Webber distinguished the ‘tame’ problems which had been solved from the ‘wicked’ problems which might never be solved, and these terms are now often used in social policy and medicine.<sup>12</sup> A broken leg is a tame problem; but many patients present with symptoms whose cause is difficult to diagnose and which require treatment whose outcomes are uncertain. Their doctors must tackle wicked problems. Engineers also distinguish between puzzles and mysteries, and give them technical names – ‘aleatory’ and ‘epistemic’ uncertainty, respectively. Meteorological records will describe the regular tides and winds to which a bridge is likely to be exposed (aleatory uncertainty), but since every bridge and every bridge location is different the effect of these conditions on the structure is never completely known (epistemic uncertainty). The tides and winds are the subject of known frequency distributions (tables showing how frequent are particular values of tide and wind speed); uncertainty remains because every complex structure is necessarily idiosyncratic. This distinction between the uncertainty which can be described probabilistically and the uncertainty which surrounds every unique project or event is important in all applications of practical knowledge and central to the argument of this book.

Donald Rumsfeld was not the first person to describe ‘unknown unknowns’. British scientists invented the jet engine before the Second World War and by 1944 both Britain and Germany were able to produce jet fighters; the first commercial jets were Comets manufactured just outside London by the de Havilland company. The greater speed of jet propulsion promised – and in due course delivered – a transformation in international passenger travel. *American Aviation* magazine commented: ‘Whether we like it or not, the British are giving the US a drubbing in jet transport.’<sup>13</sup> But a BOAC (now British Airways) Comet disintegrated in mid-air in 1954 soon after taking off from Rome airport. The planes were grounded for almost three months during which modifications

rise to power he understood ‘what is going on here’, and because, when the war which he saw as inevitable came, he not only provided inspirational leadership but displayed a sure grasp of the central strategic issues.

Churchill recognised that British survival depended on securing American engagement in the war against Hitler, but he also appreciated the difficulty of bringing about that engagement. Like Roosevelt, he did not anticipate a Japanese attack on the United States, nor Hitler’s response to it. Though the most impatient of men, he waited on events, and was vindicated by history. Like President Obama, though in a much more complex situation, Churchill did not think that the uncertainties involved in big decisions could be represented by probabilities and (unlike Obama) no one suggested to him that they could. The environment was radically uncertain; the problems he faced were ‘wicked’, not ‘tame’.

Military campaigns are complex and evolve in unpredictable ways. The contrast between the five-and-a-half-year history of the Second World War, with innumerable unexpected twists and turns, and the six-and-a-half-year progress of MESSENGER to Mercury along a complex but completely anticipated trajectory, could hardly be more stark. Still, even war operations are responses towards defined objectives and generally proceed to a resolution, bloody though that may be. The Second World War was prompted by Nazi aggression, the issue was how that aggression was to be contained, and in 1945 there was a definitive outcome. Many issues in business, finance and politics are ill defined and never concluded. If anything, they are even more ‘wicked’ than those of military strategy.

## **What Steve Jobs knew**

Thomas Watson Jr followed his father, Tom Sr, as CEO of IBM and over half a century they built a company which came to dominate the global computer market from its beginnings in the 1950s until the personal computer revolution of the 1980s. The observation – widely attributed to Watson Jr – that there would be a world

market for only five computers is apocryphal. But these IBM machines were enormous. In the 1970s, a university or a large corporation would have a single computer. The authors remember carrying boxes of punched cards to the Oxford University computer, housed in a vast air-conditioned basement in Banbury Road, or to the Cambridge University computer at the appropriately named New Museums Site.

Ken Olsen was then chief executive of America's second largest computer company, Digital Equipment Corporation (DEC). In 1977, Olsen declared, 'There is no reason anyone would want a computer in their home.'<sup>17</sup> Olsen anticipated extensive use of computers. But, like many others, he assumed that millions of people would obtain computing capacity by plugging into a few very large central facilities, rather as they draw power from an electricity grid linking large power stations to their home appliances.

A different vision of the digital future imagined a range of smaller purpose-specific machines. By the early 1980s, most professional offices used word processors. Such machines made it easy to correct documents and cut and paste material, and could be linked to a high-quality printer. They transformed the job of the copy typist and almost eliminated the electric typewriter. Wang Laboratories was market leader in word processors. Small programmable calculators displaced the engineer's slide rule. Specialist machines from Hewlett-Packard and Casio which could calculate redemption yields or option values were substitutes for the instinct, or seat of the pants, of market traders.

These problem-specific machines were, however, superseded when the industry developed in a different way. In 1971, Intel developed a general-purpose chip, or microprocessor. As a result, a single small device could serve many functions. Intel's innovation paved the way for the construction of a mini-computer. At Xerox Parc in 1972, Butler Lampson built the Alto, a machine which differs little in appearance from a modern desktop computer. Lampson's team added many of the features we take for granted today. But it was years before the Xerox Corporation attempted to market a commercial version, and the company never succeeded in establishing a foothold in the computer business.



While Xerox was perfecting the Alto, personal computers were developed by hobbyists. The Altair desktop, a self-assembly kit with a price of \$400, was first advertised in *Popular Electronics* magazine in December 1974. Two young high-school friends in Seattle, Paul Allen and Bill Gates, adapted a simple programming language, BASIC, for the Altair. Some large companies outside the conventional computer industry recognised the potential of small computers. Home computers used tape cassettes for storage and television sets as monitors. AT&T and Sony sold desktop machines. All these initiatives failed.

Then in 1981, IBM launched a ‘personal computer’, immediately abbreviated to PC. IBM’s reputation and market presence were such that whatever the company supported commanded wide acceptance. It didn’t matter that many users thought the performance of the PC was inferior to that of machines already on the market. When software developers decided which format to use, IBM’s system was the obvious choice. Within months, ‘PC’ had become the generic term for a small computer.

In order to avoid its own slow-moving decision processes, and bypass managers who – rightly – feared that the innovation threatened their position, IBM had outsourced much of the development of the PC. For the operating system, the corporation turned to a small company, Microsoft, run by Gates and Allen. The pair in turn developed an off-the-shelf system they had bought for \$50,000. The computer giant had no real sense of the revolution it had launched and the rights to MS-DOS remained with Microsoft. When IBM attempted to regain control with a new and more sophisticated operating system, OS/2, it was too late. MS-DOS (powering Windows 3.1) was everywhere.

Meanwhile, Steve Jobs and Steve Wozniak began assembling Apple machines in 1976 in Jobs’s garage, now designated a historic site.<sup>18</sup> Although Gates and Microsoft had understood that ease of use was as important as technical sophistication to commercial success, Jobs extended this vision further and conceived a computer that you could use without understanding anything about computers. To achieve this goal, Jobs drew on another

invention from Xerox Parc – the graphical user interface. Apple machines had screens with icons which created the appearance of a desktop, and friendly aids such as a mouse and trash bin – innovations that seemed like gimmicks to the nerds who then predominated among computer users, but which opened computing to a much wider audience. Apple machines were more fun.

But you could access these capabilities only by buying Apple’s integrated software and hardware. Apple’s determination to maintain its proprietary system failed in the face of widespread adoption of the more open standard of the IBM PC: Windows, a combination of Apple’s graphical user interface with Microsoft’s ubiquitous MS-DOS, swept the world, and almost swept Apple from that world. By the mid-1990s, Apple was on the edge of bankruptcy, its market share falling, its innovations failing.

But 1997 was the year of the Second Coming of Steve Jobs (he had been forced out of the company a decade earlier by the board). The return of Jobs to the company he had founded twenty years earlier enthused the dwindling band of Apple devotees, but few in the business world had high expectations. In 1998, Dick Rumelt, the UCLA strategy professor we met in chapter 1, interviewed Jobs about his plans. The Apple CEO responded, ‘I’m going to wait for the next big thing.’<sup>19</sup> The ‘next big thing’ proved to be music. Music publishers resisted digital downloads, proclaiming them piracy. They sought to protect their established business of selling compact discs through record stores. Napster and other illegal file-sharing services flourished. Apple secured the rights to sell millions of downloadable tracks at 99 cents each through the iTunes Store, and launched the iPod in 2002. ‘One thousand songs in your pocket’, Jobs proclaimed.

The iPod prepared the ground for something much bigger still – the handheld computer. High-end portable devices for business people had been available since the turn of the century; the Palm Pilot was succeeded by the BlackBerry. But Apple aimed its products at consumers and then opened its systems to enable developers to provide ‘apps’. Combine the music player with the increasingly ubiquitous mobile phone, add a screen, and you could

devise almost unlimited applications for a gadget that would fit in your pocket. Steve Ballmer, CEO of Microsoft, laughed derisively when the iPhone appeared – who, he asked rhetorically, would pay \$500 for a phone?<sup>20</sup>

A lot of people would; ten years on, more than 1.5 billion smartphones had been sold. The smartphone changed the nature not just of entertainment, but of business communication. And by the time Jobs died in 2011, Apple had outstripped Microsoft to become the most valuable company in the world. Gates and his successor Ballmer were wrongfooted by the popularity of Apple's mobile devices, as was Nokia, the Finnish company which had become the world's largest supplier of mobile phones. In 2014, these two companies huddled together against the storm, as Microsoft acquired the remains of Nokia's handset division. Palm became part of Hewlett-Packard; BlackBerry is today a shadow of its former self.

But perhaps Olsen, who doubted that there would be a need for a home computer, was right after all. We no longer need or want computers in our home because we carry a computer wherever we go and plug into the limitless memory and processing power of the great servers in the cloud. But Olsen's company will not benefit from his belated vindication. The struggling Digital Equipment Corporation was absorbed into Compaq, which in turn was absorbed into Hewlett-Packard, which in turn was split in two in 2015 and is now known primarily as a manufacturer of printers. Wang, which had popularised the word processor, went bankrupt in 1992.

The history of personal computing combines extraordinary success in meeting the evolving needs of consumers with comprehensive failure by the corporations involved to anticipate how the market would develop. DEC failed to benefit from the company's leading position in a market for small computers which was about to experience exponential growth. Wang, Casio, Palm, BlackBerry and Nokia soared towards the sun and fell as quickly to earth. IBM pioneered a development which destroyed its established business. Apple's insistence on proprietary systems failed in the 1980s, but proved a success twenty years later.

board. But the lifeboat and the ship's papers were missing. As were the captain, his wife and daughter, and the crew of seven. Nothing was seen or heard of them again.

No sea mystery has ever received more attention. Seemingly plausible hypotheses, such as piracy or mutiny, appear inconsistent with the evidence, prompting speculation about outlandish ones, such as sea monsters. The incident is famous in part because Arthur Conan Doyle, the creator of Sherlock Holmes, wrote a fictional (and wholly implausible) account of what took place.<sup>24</sup> But his theory was to be the first of many.<sup>25</sup> What happened to the *Mary Celeste* will almost certainly remain a mystery, despite the efforts of successive generations of crime novelists to treat the question as a puzzle. Even though there is an answer, we will never know what it is. The claim of the modern science of decision theory is that most mysteries can be reduced to puzzles by the application of probabilistic reasoning. Such reasoning can provide solutions to puzzles, but not to mysteries. How to think about and cope with mysteries is the essence of managing life in the real world and is what this book is all about.

---

## RADICAL UNCERTAINTY IS EVERYWHERE

Time and chance happeneth to them all.

—Ecclesiastes 9:11

**F**orecasting is difficult. But as we have seen, the physicists and engineers of NASA were able, with uncanny accuracy, to predict the position of MESSENGER. NASA was dealing with a problem that was completely specified, comprehensively understood, and stationary. And that system was unchanged by human interaction with it. Its behaviour was not affected by what people understood of it or by what people did. If economic problems were like those faced by NASA, economists could have the same predictive capability as NASA.

But economic relationships change over time – the property of non-stationarity. And movements in the economy reflect our expectations. The sociologist Robert K. Merton<sup>1</sup> identified *reflexivity* as a distinctive property of social systems – the system itself is influenced by our beliefs about it. The idea of reflexivity was developed by the Austrian émigré philosopher Karl Popper and became central to the thinking of Popper’s student, the highly successful hedge fund manager George Soros.<sup>2</sup> And it would form part of the approach to macroeconomics of the Chicago economist

Robert Lucas and his followers, which we describe in chapter 19, although their perspective on the problem and its solution would be very different.

Reflexivity undermines stationarity. This was the essence of ‘Goodhart’s Law’ – any business or government policy which assumed stationarity of social and economic relationships was likely to fail because its implementation would alter the behaviour of those affected and therefore destroy that stationarity.<sup>3</sup> In an early illustration of reflexivity, Jonah prophesied the destruction of Nineveh, having received inside information concerning God’s plans to punish the city (his journey to Nineveh was interrupted by a bizarre encounter with a whale). But after his arrival the citizens repented on hearing his forewarning and the city was spared. This outcome ‘displeased Jonah exceedingly, and he was very angry’, feeling (unlike many modern forecasters) despondent at the very public refutation of his prediction. But God persuaded Jonah that the happy outcome was more important than the failure of his forecast.

The King of Nineveh wore sackcloth and sat in ashes; the titans of Wall Street had no similar opportunity or inclination. The collapse of Lehman Brothers on 15 September 2008 could not have been widely predicted because if it had been it would not have happened on that date. Either the bank would have collapsed earlier, or the regulators or Lehman itself would have taken steps to avoid, or at least minimise, the event. And because beliefs influence behaviour the economic system is forever changing.

## **The scope of probabilistic reasoning**

The Oxford Dictionary defines uncertainty as ‘the state of being uncertain’, and the meaning of uncertain as ‘not able to be relied on, not known or definite’.<sup>4</sup> Such uncertainty is a product of our incomplete knowledge of the state of the world – past, present or future. Or our incomplete knowledge of the connection between actions and outcomes. We talk about uncertainty only if incomplete knowledge leads to a state of doubt – we are only too familiar (some politicians come to mind) with people who are

*image*

*not*

*available*

intelligence and robots which run out of control. These are threats we can at least perceive. But the observation of a black swan was not a low-probability event; it was an unimaginable event, given European knowledge of swans. As the convict colonists boarded the First Fleet, no one would plausibly have offered, or accepted, a wager of the kind 'I bet you one thousand to one all the swans in Australia are white'. Natural phenomena are more likely than social ones to be the result of stationary processes – the structure of the physical world changes less than do global business, finance and politics. But the impact of a pandemic is determined as much or more by the state of medical knowledge as by the pathogens of disease. The Black Death will not recur – plague is easily cured by antibiotics (although the effectiveness of antibiotics is under threat) – and a significant outbreak of cholera in a developed country is highly unlikely. But we must expect to be hit by an epidemic of an infectious disease resulting from a virus which does not yet exist. To describe catastrophic pandemics, or environmental disasters, or nuclear annihilation, or our subjection to robots, in terms of probabilities is to mislead ourselves and others. We can talk only in terms of stories. And when our world ends, it will likely be the result not of some 'long tail' event arising from a low-probability outcome from a known frequency distribution, nor even of one of the contingencies hypothesised by Martin Rees and colleagues, but as a result of some contingency we have failed even to imagine.

In 1896, Lord Kelvin, one of the greatest physicists of his age, wrote that 'I have not the smallest molecule of faith in aerial navigation other than ballooning or of expectation of good results from any of the trials we hear of. So you will understand that I would not care to be a member of the aëronautical Society.'<sup>8</sup> His observation was followed in an embarrassingly short time by the first controlled flight, covering 300 yards and just under a minute long. And today, two centuries after the First Fleet reached Botany Bay with no expectation of observing black swans, an Airbus A380 weighing 360 tonnes can carry 550 passengers over 9000 miles from England to Australia. Something that would have been incomprehensible even a hundred years ago. The next hundred



Why has this seemingly obvious critique been so widely ignored? The hegemony of optimisation as the goal of decision-making is made possible by ignoring radical uncertainty. Building on the success of probabilistic reasoning in illuminating games of chance, the approach of decision theory bifurcates uncertainty into the unknown and unknowable, and the unknown but capable of being characterised by a known probability distribution. The practitioners of this approach wash their hands of the former, describing the unknown and unknowable as ‘shifts’ and ‘shocks’, as unpredictable and inexplicable as the Yucatán asteroid. Other uncertainties are treated as resolvable. There is no room for radical uncertainty.

But people routinely need to make decisions with imperfect information. Most of real life lies in between the opposing poles of randomness and black swans; we know something, but not enough, and the knowledge which is held collectively is widely and unequally distributed. Regulators and counterparties (firms which traded with the now defunct bank) might have known Lehman was poorly managed and inadequately capitalised, although even the latter was the subject of contention. They might have known that the bank was likely to fail, but not how or when. We see, but through a glass, darkly.

It is easy to understand why economists and statisticians, in search of clear and comprehensive solutions, have sought wide extension of the scope of probabilistic reasoning. The underlying mathematics has a certain simplicity and beauty, and in practice can be applied by those who have acquired the requisite modest technical skill. Arguably the two most brilliant economists of the post-war period, Paul Samuelson and Robert Solow, occupied adjoining offices at MIT for over half a century. As Samuelson relates, ‘When young he [Solow] would say, if you don’t regard probability theory as the most interesting subject in the world, then I feel sorry for you. I always agreed with that.’<sup>9</sup>

The appeal of probability theory is understandable. But we suspect the reason that such mathematics was, as we shall see, not developed until the seventeenth century is that few real-world problems can properly be represented in this way. The most

Sometimes the relevant state of the world, although a present fact, is not known to the decision-makers even after their best endeavours – is the man in the compound bin Laden? And sometimes the actual state of the world is a present or past fact, but one which is not known to anyone – what happened to the *Mary Celeste*? Or, in an issue of rather greater moment to the modern world, how many votes did Bush and Gore actually receive in Florida in the 2000 US presidential election?

Where the question is known, but the range of answers is unbounded, application of the mathematics of probability is questionable and the results ambiguous. When the question is known – ‘What will happen in the Middle East in the next five years?’ or ‘What will be the position of private property twenty years from now?’ – but the nature of the question means that the answers are ill defined, then there are no states to which we can sensibly attach probabilities.

## **Explaining uncertainty**

The results of most medical treatments are uncertain. Doctors learn frequency distributions from their own experience and that of the medical profession as a whole. But even if extensive data are available the circumstances of every patient are unique. The modern requirement for informed consent obliges doctors to communicate this uncertainty to their patients. But their patients typically crave certainties. They place great trust in the judgement of doctors, and want to believe that doctors understand more than they do. One London doctor obtained feedback from a patient group on the wording of an information leaflet for patients in intensive care who might wish to participate in a trial involving different antibiotic therapies. The lay patient group did not like the wording ‘doctors wish to carry out this trial as they do not know which antibiotic therapy is best’ and preferred the wording ‘doctors wish to carry out this trial to help decide which antibiotic is best’. One patient’s comment was that ‘everything is uncertain in intensive care and the last thing I want is the doctors not to know what to do’. The doctor’s reaction was that ‘expressing

response that people have ‘had enough of experts’.

## **Practical decision-making**

The different dimensions of uncertainty mean that the strategies we adopt to cope with risk and uncertainty will depend upon the particular problem with which we are faced. Most hypothetical questions – what is the capital of Pennsylvania? – are of no concern to most of the population of the world, and mostly the ignorance that results does not matter. The authors of this book have not the slightest interest in the question ‘Which horse will win the 2020 Kentucky Derby?’ They do not know who the runners and riders are, far less their form, or when the race will take place. Nor do they intend to bet on the outcome, or to find out the name of the winner when the finishing list is posted.

When decisions do matter, rational people delegate them to those who have, or are willing to invest in acquiring, relevant information and the capacity to interpret that information. For all that has recently been said about ‘the wisdom of crowds’, the authors prefer to fly with airlines which rely on the services of skilled and experienced pilots, rather than those who entrust the controls to the average opinion of the passengers.<sup>13</sup>

There is no general theory of how best to make decisions. Much of the academic literature on decision-making under uncertainty tries to frame the challenge as a puzzle. All decisions, it is assumed, can be expressed as mathematical problems. And potentially capable of being solved by computers. Your smartphone will tell you what restaurants are nearby, how to get there, and perhaps what you ate last night; but not where and what you want to eat now. The probably apocryphal but nevertheless illuminating story about a decision theorist contemplating whether or not to accept a job offer from a rival university illustrates this well: upon being urged by his colleague to apply tenets of rational decision-making under uncertainty and maximise his expected utility, as his academic papers suggested, he responded with exasperation, ‘Come on, this is serious.’<sup>14</sup>

Humans have evolved to cope with problems which are not

amenable to probabilistic reasoning – an issue to which we return in chapter 9. Our brains are not built like computers but as adaptive mechanisms for making connections and recognising patterns. Good decisions often result from leaps of the imagination. Creativity was the quality exhibited by that unknown Sumerian who invented the wheel, by Einstein, and by Steve Jobs. And, as Knight and Keynes emphasised, creativity is inseparable from uncertainty. By its nature, creativity cannot be formalised, only described after the event, with or without the help of equations.

## **The founding of California**

Johann Suter, a wannabe businessman, abandoned his creditors and his family in Baden (now part of Germany) in 1834. After much travel he reappeared in 1839 on the west coast of America as John Sutter. Ambitious to establish an agricultural empire, he settled in what we call today the Bay Area. At the time, San Francisco was a modest trading post with a harbour, and home to around a thousand people.

In 1848, the Treaty of Guadalupe Hidalgo ended the Mexican War and led to the annexation of California by the United States. And in the same year, one of Sutter's employees discovered gold on the property. Sutter tried to conceal the find, partly so that he could enjoy the full benefit himself, and partly because he sensed the negative implications for his farming interests. But concealment proved impossible. The *San Francisco Examiner* published a rumour of gold and in 1849 as many as a hundred thousand people are thought to have arrived in California. Some struck it rich, most did not. Others realised that a different, if less adventurous, route to success was to provide services to the '49ers. One of them was Leland Stanford. Stanford built a successful trading business, and served a single two-year term as Governor of the fledgling state. But he was most famous for his contemporaneous role in creating the Central Pacific Railroad in 1861. The Central Pacific then built a stretch of line from Sacramento to Promontory Point in Utah, where in 1869 it joined

sharp distinction between truth, which was axiomatic, and probability, which was merely the opinion of man. In premodern thought there was no such thing as randomness, since the course of events reflected the will of the gods, which was determinate if not fully known. The means of resolving uncertainty was not to be found in mathematics, but in a better appreciation of the will of the gods. Hence actions which seem to us absurd, such as inspecting the entrails of sacrificial animals or consulting an oracle, were employed for millennia. And traces of this approach linger on today, among those who follow astrology, read tea leaves, or attach weight to the predictions of gurus who are believed to have privileged access to knowledge of future events.

So it is not just the mathematical expression of the concept of probability that is recent, but the *concept* of probability itself as the quantitative expression of the likelihood of one of a number of possible outcomes. Even in the eighteenth century, Edward Gibbon could write of Hannibal's crossing of the Alps, 'although Livy's narrative has more of probability, yet that of Polybius has more of truth' and (in relation to a claim that the army of defeated Emperor Jovian was supplied with provisions by the victorious Persians), 'such a fact is probable but undoubtedly false'.<sup>3</sup>

What did Gibbon mean? The words 'prove', 'probable' and 'approve' have a common root. That relatedness is not apparent from the way we use these words today. But the relatedness was apparent to medieval writers – and to Gibbon – for whom 'probable' meant something like 'approved by most right-thinking people'. In an era in which truth was established by religious or secular authority, such right-thinking people might properly decline to look through Galileo's telescope on the grounds that the Church had decreed that what he claimed to see could not be there.<sup>4</sup>

When the Royal Society, Britain's premier scientific body, was founded in 1660, it took as its motto 'nullius in verba' – today translated informally as 'take nobody's word for it' – a forceful assertion of the primacy of experiment and discovery over argument from authority. The modern idea of probability was very much part of the development of scientific reasoning in the

seventeenth century; such reasoning was a prerequisite for the Industrial Revolution and the unprecedented economic growth it generated. The advance of probability theory would contribute to that economic development through the creation of markets in risk.

The first venues for managing risk were the coffee shops of London. Coffee had recently been imported from Arabia to Europe, and gentlemen met to consume the newly fashionable drink, converse, and do business. In the historic City of London, the insurance market began in Tom's coffee house, while securities were traded in Jonathan's – today regarded as the origin of the London Stock Exchange. The restoration of the monarchy after the dull morality of the Puritans led to an upsurge in gambling.

Mrs White's Chocolate House, in St James's near the royal palaces, evolved into the first of London's gentlemen's clubs, and served principally as a gaming venue. (And it may not be coincidental that fashionable St James's is today the centre of London's hedge funds, while the more business-oriented financial activities are found in the City.) The most famous coffee shop of all was that of Edward Lloyd, where patrons speculated on the weather, the tides and the fate of ships at sea, and merchants could lay off some of the risks associated with foreign trade. Established in 1688, Lloyd's of London is still globally pre-eminent in marine insurance. History casts a long shadow.

## **Mortality tables and life insurance**

As Pascal and Fermat were exchanging learned correspondence, an English cloth merchant, John Graunt, was trawling the records of the cemeteries of London. Graunt noted the reported causes of death and his data were used to observe, if not to prevent, the spread of plague. He compiled records of the incidence of death in different age groups, and his analysis is the precursor of the tables actuaries use today to compute appropriate prices for annuities and life insurance. Graunt worked with the assistance of his patron and friend Sir William Petty, whose *Statistical Account of England* foreshadowed the national accounts compiled today by

stationary process. The achievements of the great classical statisticians of that era had provided tools useful in many areas of both social and natural sciences. Statisticians had a secure place within the scientific community. The probabilistic turn led modern economists and other social scientists to proceed firmly down the probabilistic road.

### **The problem of points**

The question the Chevalier de Méré had put to Pascal, which led to the modern theory of probability, was ‘the problem of points’. Suppose a game of chance in the Chevalier’s salon is interrupted. What is a fair division of the stakes between the players, given the outcomes of the incomplete game? For example, two players contribute to a pot of 100 Louis d’or and agree that the winner of the most of seven games will scoop the pool. The Duke of A has won three games and the Marquis of B one. The Duke is summoned to see the King and the evening’s entertainment is abruptly terminated.

Before Pascal, the widely accepted solution to this problem gave three quarters of the pot to the Duke of A, recognising that he had won three out of four of the games actually played. This solution had been elaborated in the late fifteenth century by the Italian mathematician Luca Pacioli, widely regarded as one of the inventors of accounting, and at first sight seems plausible and fair.<sup>12</sup> But the Chevalier was not convinced that Pacioli had reached the right answer, and the two great mathematicians confirmed his doubt. If play had continued, the Marquis would have needed to win all three remaining games to succeed. If the chance of either player winning each game is one half, then the probability that the Marquis will succeed in winning all three is only one eighth. It follows that the probability that the Duke would have scooped the pool was seven eighths. So, the argument went, the pot should be shared in these proportions.

The Fermat–Pascal solution introduces three notions which are fundamental to all subsequent work. There is the mathematical concept of *probability* itself – the chance of winning any particular

game. There is the method of calculation of *compound probability* – the probability of winning three successive games is obtained from the probability of a single win, one half, raised to the power of three. And the solution introduces the idea of *expected value* – the amounts each player could have expected to win if the evening’s events had been repeated many times. And today we could programme a computer to simulate that scenario of multiple repetition and verify that the calculated expected value does indeed describe what would have happened if the evening’s events had occurred again and again. (And, in the Chevalier’s salon, they probably did.)

The solution of the problem of points was an early indication of the power of probabilistic reasoning. Pascal’s counter-intuitive answer is persuasive once you have understood his thinking. What matters is anticipating the future rather than analysing the past. If the Duke and Marquis had planned to play one hundred games, then the Duke’s three to one advantage at an early stage of the evening would have counted for little. But if there were to be only five games, the Marquis would certainly have lost – the result of the fifth game would be irrelevant, and it might not even have been played.

### **It pays to go Bayes**

The final step in the development of the new theory of probability was the achievement of an unlikely hero – an obscure eighteenth-century country Presbyterian clergyman in England. The Reverend Thomas Bayes is by chance buried in what is now the centre of London’s financial district. Among his papers he left a theorem that is one of the most widely taught ideas in statistics today.<sup>13</sup> Unknown in life Bayes may have been, but his name is known today throughout the world with branches of statistics and economics named after him. The term ‘Bayesian’, which describes not just a statistical technique but a school of thought, is the intellectual legacy of one man working in the Kent countryside.

Bayes’ theorem enables us to calculate *conditional probabilities*: what is the probability that A will happen, given that B has