The Gambling Animal

The Gambling Animal

Humanity's Evolutionary Winning Streak – and How We Risk It All

Glenn Harrison
Don Ross

Profile Books

First published in Great Britain in 2025 by
Profile Books Ltd
29 Cloth Fair
London
EC1A 7JQ
www.profilebooks.com

1 3 5 7 9 10 8 6 4 2

Typeset in Swift by MacGuru Ltd
Diagrams by Jeff Edwards

Printed and bound in Great Britain by
CPI Group (UK) Ltd, Croydon CR0 4YY

The images on pp. 349 & 350 are reproduced with kind permission from Adventures with Elephants.
The image on p. 67 is reproduced courtesy of photosoup/iStock.

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

HB ISBN 978 1 78816 362 0
TPB ISBN 978 1 80522 584 3
eISBN 978 1 78283 618 6

MIX
Paper | Supporting responsible forestry
FSC® C171272
www.fsc.org

G. H.: To the wise matriarchs Lisa and Linnéa,
to Fabian, and of course to Logan

D. R.: In memory of Dan Dennett, who spotted
the goals and forged the paths to them

Contents

Introduction

The Winning Animal (For Now)

Human Pride and Human Shame

Humans, comparing themselves with other animals, have seldom failed to find sources of pride. Based on our achievements in engineering, science and art, we seem to be the smartest species. The size of our brain, relative to the size of our body, apparently backs up that abstract assessment on the basis of measurable physical hardware.

We are undeniably Earth's dominant ecological species. This doesn't mean we're the most numerous – there are far more houseflies and nematodes, for example – but we have transformed most of the niches on Earth to suit our needs for space, food and energy, and in doing so we have driven all our natural predators, except some parasites, to peripheral environments in which few people want to live. Recently, we've found ways to intervene against diseases so effectively that even in poor countries, majorities of people expect to live into old age, a privilege few other animals enjoy.

This comparison also makes many people feel ashamed of their species. Historically, it was often said that only humans engage in sadistic, gratuitous violence towards one another, though as we've come to understand the richness of other species' behaviour better over recent decades, we have learned that such sweeping generalisations are glib at best. What *is* true in our collective deficit column is that we've caused the extinctions of more species than any other (by far), and that we are the trigger for an epochal *mass* extinction now in progress.[1] This is largely because our industrial and agricultural activities have been generating changes in the basic conditions of the planetary environment that are too rapid for adaptive biological systems to keep up with. We seem, indeed, to be so violently disrupting the basic living conditions on Earth that we are putting our *own* future well-being at deep risk. Some even fear that one of the last species we'll push to extinction is *Homo sapiens*.

In the standard conception, the sources of pride and shame in humanity are two sides of one coin. We dominate the world due to our superior intelligence and the ingenuity it fosters; we threaten the stability of nature, and all the life that depends on it, because that ingenuity has allowed us to become too numerous and too powerful. Implicit in this straightforward view is a critique: surely our intelligence is less than it's cracked up to be if we're now proving to be too stupid to save our planet from our own actions. A distinction is often made between *practical* intelligence and *moral* intelligence: it is said that we don't have enough of the second to harness our mere cleverness about building things. A related idea is that we're deficient in self-control or willpower.

In this book we'll develop an alternative narrative about human success and the threats it creates that is less moralistic and more scientific. We'll explain both our ecological dominance and our struggle to contain its consequences in terms of a single concept: risk management.

There are many other books about risk management, and most of them are found in the business sections of stores and libraries. But risk management is a *natural* phenomenon, something all animals engage in, and which humans have been wrestling with since long before anyone explicitly named the concept. The book is a natural history of risk management, from the prehistoric origins of humans in Africa to the current crises of rapid climate change and biodiversity collapse. We think that this is a powerful way to understand the general career of *Homo sapiens* and our relationship to the planet and its other inhabitants, and we hope that after reading the book, you will think so too.

Who We Are

Any reader of a book about science naturally wants to know something about the perspective from which the authors write. In the case of this book, some readers might be surprised to find two economists writing about natural history. Aren't economists supposed to be preoccupied with markets and finance?

Many economists are indeed especially interested in these topics. We've often researched and written about them ourselves. But the deep subject matter of economics is more general: it's about how people and other agents (including non-human animals) allocate their always limited resources among possible actions they could take that might improve their level of flourishing. The word 'might' is important here. Interesting economic problems for agents are those that involve some degree of uncertainty or risk. This allows for choices that are sensible *in advance* of making them – *ex ante*, as we say – but could nevertheless work out badly *ex post*. Everyone is familiar with lots of everyday cases. For example, the weather forecast says there is a low probability of rain today, so you set off on a long walk in light clothing with no protection against the elements; but that low probability wasn't *zero*, and today you're reminded of that fact when, despite your reasonable planning, you end up soaked and cold.

Some economists mainly study risky choices like this by developing abstract theory. We do that too, and have researched and modelled a range of phenomena including addiction, public health, international trade, public infrastructure and environmental policy. But we primarily specialise in designing, administering and analysing *experiments* with volunteer subjects. We lead an international team of economists that has (so far) conducted experiments with about 40,000 participants from almost every continent and from all levels of wealth, income and education. What we know about risky human choice comes from analysing these experimental data (along with those of many colleagues) under the discipline of economic theory. To give the reader a general feel for the scientific expertise from which our reflections and arguments come, we'll state a few general principles we follow in all our experimental studies.

First, we don't measure risk responses only as *averages* in groups of our subjects. We begin by measuring them at the level of patterns in the choices of individual people, where the patterns in question are identified in terms of theory. Then, when we ascend to the level of talking about groups, the groups in question are picked out by variables, such as age, sex, income, wealth and education, which come from our data. We thereby avoid commitments to any special prior expectations about what kinds of differences these variables should make to people's risky choices; we let careful, controlled observation tell us these things.

Second, we always give the participants in our experiments real and meaningful incentives, usually cash payments, that are designed to give them strong motivation to express their true attitudes to risk, as opposed to attitudes they think that other people, including us, might expect them to have or think they *should* have.

Third, we don't assume that our experimental participants have indefinitely powerful and fast computers in their heads. We know, from evolutionary biology, psychology and neuroscience, that they do not. The theory we use is flexible enough to allow for their errors and, even more importantly, for the special, personal levels of anxiety or adventure they associate with risk. Sometimes our experiments are designed to require the participants to show us how they choose when they are not able to use information from other people. But equally often we give them controlled access to social information. This is because the 'computers' that real people outside our lab use to make risky decisions go beyond the ones inside their skulls.[2] They assess their options using information that's been gathered, often through trial-and-error experience, by other people across time and space, and recorded for general consumption in manuals, advice blogs, textbooks and rules of thumb. Some principles for risky choice aren't left up to individuals at all, but are factored into laws and regulations.

Fourth, we don't measure our subjects' risk attitudes as if these were revealed independently of their other thinking. People have special beliefs about how probabilities of risky events are distributed, and we don't just assume that we can know what these are in advance. Therefore the people who come into our lab do multiple experiments – always with real rewards at stake – and the results of these are combined, using the special body of applied mathematics that economists call *econometrics*, to compose a picture of the motivational structure of each subject as they handle risky decisions.

Experimental economics, like all experimental science, is a never-ending pursuit of more exact and nuanced truths. We don't pretend to know everything there is to be known about human risky choice, and we don't imagine that anyone ever will achieve such omniscience. One of the best things about science is that it always invites surprises and new discoveries. But, without being falsely modest, we do think that we now know *quite a lot* about how people handle risk. On that base of knowledge, we have expanded our focus beyond humans, and are currently running experiments in which the participants making risky decisions are African elephants.

We chose elephants as our first non-human subjects for two main reasons. First, they are intelligent enough, in the narrow sense of having big, powerful information processors between their ears, to cope with differences in basic statistical patterns. For all we can tell, they may be as smart as people in this narrow sense.* Second, they share information with one another about some decisions, because they face many important risky choices, mainly about where to look for food and water, as groups. Of course, our elephant subjects are not interested in cash incentives. But they readily apply their minds to trying to earn as many apples and oranges from us as they can.

Comparing human with non-human risky choices, under the discipline of experiments, is in our opinion the best way to try to understand the role played by social risk management in the ascent to ecological dominance of *Homo sapiens*. This understanding, as we will explain, is crucial to assessing the prospects for less disruptive exercise of our dominance. The human record of gambling with the natural environment currently faces potential gigantic losses, for ourselves and many other species, including elephants, caused by the collective bets we have made over the past few centuries.

---

*We refer here to people without the aid of artificial technology. Elephants obviously can't run large regressions, for example; but neither, without external equipment, can people.

What's Ahead

The story to come is built around three main scientific themes that we'll weave together. One of these themes is the *scientific modelling of risk and risk management*. The second is *convergent evolution*, the comparison of different strategies that natural selection finds for solving shared challenges encountered by different lines of animal (and plant) descent. The third is the force of *cultural* evolution, which can radically change behaviour at the scale of populations and create social pressures that change the risk environments to which individuals must adapt.

We'll start to weave these strands together by reflecting critically on the idea that humans are uniquely intelligent because we have unusually large brains. As we'll see, this is at best a very loose approximation of the truth. Some other animals also have large brains, and a few – elephants, for example – have larger brains than we do. In fact, the size of the human brain relative to the human body reflects the standard ratio for a primate.[3] (Other large primates, the surviving great apes such as gorillas, have brains that are *smaller* than would result from scaling up monkeys.) What *is* special about the human brain is that it has a very high total number of neurons in a particular part of the brain, frontal cortex. But is that the source of our ecological dominance? If so, why should that be? Elephants also have unique brains, with more neurons in a different part, the cerebellum, than other animals. Why do lots of frontal cortical neurons generate human-style intelligence, and lots of cerebellar neurons generate elephant-style intelligence? And what does it even mean to talk about two (or more) 'styles' of intelligence?

The answer, as we'll explore in detail, is that a large battery of frontal cortical neurons causes humans, to a much greater extent than other animals, to imagine versions of the world that differ from the actual one. This has a major downside: both individual humans, and groups of humans collectively, readily get caught in fantasy spaces that interfere with their grip on reality. But it also has a major upside: it allows them to collectively quantify, and then efficiently *distribute*, risks that other animals can manage only by deriving collective risk directly from individual risk. For them, individual conservatism therefore ensures collective conservatism. Risk always lies in a perceiver's future, so even managing risk by being very conservative requires that it be estimated on the basis of currently available evidence.

Elephants, as we'll see, rely on unusually stable and accurate memories. This allows them to be highly competent collective risk managers. But their style of risk management, unlike the human one, doesn't generate *new* risks that then need further management. Compared with humans, they are like investors who only buy very safe assets that already exist around them. Elephants do not generate 'derivative' assets, as our species does. Humans recurrently make big collective gambles, more daring than most individuals are willing to make for themselves.

This point reflects the most basic relationship in the science of risk. Larger risks imply larger potential returns *and* larger potential disasters. Early humans collectively made some very risky gambles, of a magnitude that elephants would not (and so did not). As a result, we might well have gone extinct long before we started taking over the planet. Indeed, all the many human and proto-human ('hominin') species that have existed except one – you know who! – *did* go extinct. Elephant species, by contrast, have only gone extinct when human pressure on them compounded baseline risk from climatic factors. But the direct ancestors of modern *Homo sapiens* got lucky, mainly because some large-scale climate-change events unfolded at just the right pace and in the right sequence to reward their evolutionary gambles. They consequently won the jackpot of ecological dominance. This book will take the reader through the remarkable story of this 'big night' in the natural casino.

One complication in thinking about people as natural risk managers is that *individual* humans are not systematically more risk-tolerant than other animals. We have learned this over the years in our experimental risk laboratory, where we measure the risk responses of people from multiple life stages, education levels, and parts of the world. The special human risk management style arises at the *population* level.

The obvious human disposition to become addicted to gambling, and to very risky drugs, might seem to be counter-evidence to our claim that individual people are not more risk-tolerant, in general, than other animals. The person who destroys her wealth playing slot machines typically does so alone. As we'll see in detail, however, all mammal brains share the same design features that create vulnerability to addiction. Non-human animals avoid addiction not through willpower in avoiding those risks, but because there is no one to build addictive environments to trap them. Individual people become addicted to risky activities because some *groups* of people – for example, casino and cigarette companies – collectively manage *their* risk by creating addiction traps to drain resources from other people.

Thus the main focus in this book is on risk management at the scale of populations. This is often where so-called *ecological risk* arises, where the word 'ecological' has a wider meaning than it does in everyday speech. Here, it means roughly 'at the scale of the general environment', where 'general environment' takes in *both* large-scale natural *and* cultural aspects.

There have been many other books written on factors that were crucial to the expansion of the human species, factors such as language, cooking, tool-making, singing and partnership with dogs. The story told in this book does not compete with these others. Instead, it aims to unify these partial stories by understanding all these factors, and more, as complementary pieces of a general risk-management package.

Human industrial civilisation, the way of life that is now deeply stressing the planet and our fellow creatures, evolved as a series of ratchets for managing ecological risk. Capitalism, by which we specifically mean pooling future-indexed financial resources while simultaneously decentralising decisions about how to use them, is only a very recent turn of this ratchet that began with small populations of hunter-gatherers in Eastern and Southern Africa.

Was the human path to ecological domination through industrial civilisation the only *possible* path? Nothing can be said about this with any scientific certainty. But much can be learned about what *actually* happened by seriously examining whether an alternative story could have played out. Elephants are an ideal comparative species to examine in detail, for two reasons. First, they co-evolved with humans in the same environment, facing the same general ecological challenges. Like humans, they managed these well enough to spread across multiple continents. Second, their brains evolved to give them the opposite strategy to humans in the trade-off between accurate modelling and wide-scope modelling. Our interest in elephants as a comparator species to humans extends to our experimental work. Currently, a group of six elephants in South Africa is working through the same risky choice experiments we've conducted with tens of thousands of people. The reader will be introduced to these elephants, and to our work with them.

Our history of human ecological risk management will conclude by looking towards the future. Can attention to our past track record tell us anything informative about how, at the scale of our world population, we might end up handling the huge new risks we've created? We will argue that it can – and that lesson has both some gloomy and some encouraging aspects. Our objective here is not to deliver a sermon about what we *wish* would be done to manage climate change and the threats to biodiversity. It is, rather, to bring to bear the science of risk on what we can realistically expect might happen. The goal is absolutely not a *deterministic* prediction. People build institutions, often deliberately. Institutions *choose* policies. These choices are produced by political processes, which reflect chosen actions by teams of people. But all the choices are risky gambles. That is the story of our species: we are nature's boldest gambling animal. We can learn a lot about how we'll be likely to place our bets in the immediate future by studying how we've placed them up to now.

And the fate of elephants, along with many others, depends not on how they manage the new risks, but on how we do.

1

Life Is Risky

Risk: A Fundamental Biological Concept

Four billion years ago, before life arose on Earth, there was no risk. Something can only face risk if there are some possible states of the world that are better for it than others. A rock faces no risk, because no rock ever tries to achieve anything that might be more likely under some conditions than others. But all living things have what economists call a 'value function'. Even the simplest organism, for example an amoeba, aims to resist dissolution by forces such as extreme heat or the digestive chemicals of a predator, and it aims to reproduce.* These aims, like all aims of every organism, require energy, so the amoeba must also find suitable food, in adequate quantities, and eat it. Thus temperature, the frequency of predators and the distribution of potential food are all sources of risk for an amoeba.

These conditions that create the existence of risk are, therefore, roughly the same as the conditions that distinguish living things from inanimate material. The distinctive processes of life are those that capture and convert energy so as to allow living forms to resist entropy, the natural tendency of physical structure to dissipate. Rocks don't do that: they just sit there and erode. We should not say that a rock faces risk of erosion, because the rock can do nothing to limit it or slow it down. The pace of erosion does involve varying probabilities: for example, a rock will erode faster in a stream where water rushes over it than buried in soil with low acidity. But the rock can

* Having no brain, an amoeba doesn't know that it has a value function.

marshal no energy to influence those probabilities. Only living things do that – that is, indeed, what it is to be alive.*

This tight association between life and risk reflects an insight that can be expressed as a nice slogan: *to live is to manage risk.*

If we *were* to think of the rock as facing erosion risk, then we'd have to say that it does nothing to manage that risk. But that is why it would be pointless to say that erosion is a risk for a rock in the first place: we would then be committed, for consistency, to saying that everything that happened to everything expressed risk, and the concept of risk would become meaningless. So this gets us to another nice slogan: *all risks are someone's management problem*. (For example, the amoeba manages its risk of starvation by following rising oxygen gradients that increase its probability of being in food-rich environments.)

This extremely close conceptual linkage among life, risk and risk management explains why, as scientists who specialise in risk management at both individual and population scales, we have a particular perspective on the history of the human species, and of other species with which ours can be compared. At the individual scale, an organism is most likely to flourish, relative to others of its kind, to the extent that its special patterns of risk management tend to be effective. At the population scale, a group will tend to spread its ecological footprint, relative to other groups with which it competes for scarce resources, to the extent that its members *coordinate* their risk management with one another in ways that tend to make it more efficient.

The notion of 'group' here should be understood broadly. At the large and long scales studied by theorists of biological evolution a relevant group could be a whole species. At the scales studied by

* We are here conceiving of life *functionally.* So if at some point in the future there are computers or robots that try to preserve themselves and pursue sub-goals they discover *in order to* preserve themselves, then they will be alive, even if their processes of energy capture and conversion are not biochemical. What is important to being alive is what something *does*. We don't know, at the present stage of technological development, whether it is possible to perform the processes of life without biochemistry; but we will likely know soon.

ecologists a group might be a local population of a species, or a troop, band, flock or school of specially related individuals. And in the case of humans, who culturally organise and sustain institutions, a relevant group could be a professional guild, a company, a labour union, a church, a university, a hospital, a political movement or party, a legally grounded jurisdiction such as a country, a community of collaborating scientists or many more.

Risk and Probability

To tell the story of how humans evolved, both biologically and culturally, through managing sequences of risks, we first need to explain what risk means scientifically. In common parlance, a risk refers to a prospect for trouble: we talk of the risk of cancer, or of an electrical fire, or of nuclear war. Such habits of reference frame risks as *threats*. However, in economics the concept has a broader interpretation, referring to the entire range of uncertain possible outcomes associated with an action or circumstance.

Some such ranges, what we might call *pure* threats, indeed involve only downside risk: there is no upside, to any person, from choking on food. And some risk scenarios feature only upsides. In finance, some assets represent absolutely safe investments, in the sense that the investor is guaranteed at least to get their principal back. Here the uncertain prospect that they might earn (a bit) *more* than the principal makes the technical concept of risk applicable. But most risks – and all real investments that carry significant probability of high positive returns – might yield either an upside or a downside.

A couple of simple examples illustrate how widely this point applies. The bubonic plague that swept fourteenth-century England had enormous downside risk for most people: it killed between 30 and 40 per cent of the population. But for survivors it had significant upsides in the form of higher wages and lower taxes afterwards. A dangerous repair task that might land you in hospital can have an upside if you are hoping to meet and marry a doctor or a nurse.

Many characterisations of the behaviour of animals frame risk management solely in terms of avoiding bad outcomes – the risk of becoming someone's supper, or the risk of not getting enough food and water to get through the night.* But ecologists integrate evaluation of those downside risks with the upside risks of various good outcomes occurring. When an animal creeps out of the cover of forest at night to get some water, predatory risk is there: but it is balanced, sometimes well and sometimes not so well, against the upside risks of getting a little water or, better, a lot of water. This book is an extended story of these kinds of balancing acts, across species and over time.

It is generally, and correctly, understood that risk involves competing probabilities. However, probability itself is a complex idea. As we'll see later, even scientists and mathematicians, at least in Europe, didn't understand it until the scientific revolution had been under way for a century and a half. Many philosophers think that we *still* don't understand probability, because there is more than one concept of it and they argue with one another over which of these concepts is the correct one.[1] Other philosophers take the view, which we share, that the word 'probability' is used to refer to two different concepts, which are related to one another in numerous important ways but, in the end, are fundamentally different things.†

The word 'probability' is often used to refer to *frequencies.* The classic example is the humble coin toss. If a coin is 'fair', meaning that it is not significantly heavier or more air-resistant on one side than on the other, any sequence of flips longer than about twenty is highly likely to show a mean ratio of heads to tails of around 1:1, and as the sequence gets longer the ratio will converge ever more exactly closer to that mean. In the limit – a physically imaginary point at which the coins have been flipped an infinite number of times – the ratio will be effectively indistinguishable from 1:1.

---

* For small animals with high metabolic rates, such as songbirds, this is the largest source of everyday risk.

† One of these relationships is that both of the things called 'probability' respect the same principles of numerical calculation. This is what fuels the arguments of those philosophers who think that the two main concepts of probability are *rivals*. In our view, that the axioms of calculation are shared by the two probability concepts says more about the nature of mathematics, in general, than about the nature of either concept.

Frequencies are very often important to estimate, in science, business, sports and everyday life.* But they do not furnish a general model of probabilities, of the kind that we need in this book. We will often be interested in the probabilities of events that occur only occasionally, or indeed only once – for example, the extinction of a species.

It is not *impossible*, using ingenious logic, to conceive of such probabilities as frequencies. Suppose, for example, you are considering the probability, before the event, that the dodo bird would go extinct. This event happened once and could not, as a matter of biological principle, happen more than once. You could, in theory, run many simulations of the history of life on Earth, and observe the frequency with which dodos go extinct in your population of simulated histories.

But this is really just a point about abstract conceptual logic, of little practical relevance. Its ineffectiveness would be revealed if someone actually tried to carry out the procedure described above on dodo extinction. Unless they put extreme, and accurate, constraints on (at least) hundreds of millions of parameters in the model that generated the simulations, in vanishingly few of them would dodos evolve in the first place: all species, like all lottery winners, are highly improbable. If the simulator had such vast and detailed knowledge of the genetics of birds, and correspondingly prodigious knowledge of all the past ecological circumstances of birds, that they could accurately set all those parameters and get dodos coming and going across many simulations, we would not believe their reported probability of dodo extinction based on the frequency they observed; we would instead believe it because they knew the probability to begin with by knowing all those parameters, knew it because they were evidently the greatest expert on bird evolution in the history of ornithology. Thus the imagined simulation, if it were really possible, would be pointless.

This example illustrates the concept of probability that we will use throughout this book. By 'probability' we will refer to the kind of estimate our imaginary super-ornithologist reliably furnished: a *belief*

* For example, batting averages in cricket and baseball are frequencies.

based on the subjective knowledge of birds and the mass of evidence about evolutionary conditions they had accumulated. Such beliefs are not about frequencies.

Again: frequencies, of events that really do happen often and with low sensitivity to variations in their causes, *are* very often useful to estimate, and when we want to refer to them we'll say 'frequency'. We'll reserve 'probability' to refer to someone's subjective belief about how relatively likely some outcomes are – which, as in the case of beliefs about probabilities in coin tossing, or solidly established scientific facts, might be shared by everyone.

Probabilities of this kind go by various synonyms in technical discussions. Philosophers call them 'credences'. The scholar who did most to introduce their importance to modern economists, Leonard Savage, called them 'personal probabilities'.[2] In current general usage they are most often called 'subjective probabilities' – a name that, for reasons we'll explain later, has some seriously misleading connotations.

Whatever label one uses for these probabilities – which, to repeat, we will just call plain 'probabilities' – there is a special body of mathematics for inferring them from evidence. These mathematics are based on Bayes' Rule,* a formula that anyone with a standard statistics package (such as R) can launch with a simple command on their tablet or computer. Bayes' Rule is the abstract form of any sound algorithm that tells someone how to rigorously quantitatively measure the effect that incoming data have on their initial expectations about outcomes of some process.

For example, if you were estimating an amoeba's probability of conserving enough energy to reproduce by division tomorrow, and then observed its environment becoming more acidic, Bayes' Rule would tell you, based on specific parameters about amoeba biochemistry, how to adjust the probability down. Processes that generate outcomes also generate, along the way, information about the evolving probability of the outcome in question. In the context of Bayesian inference we therefore refer to them as data-generating processes.

---

*An appendix explains what Bayes' Rule is more formally, and another appendix discusses the historical discovery of Bayes' Rule (see pp. 353 and 356).

The model for a Bayesian reasoner is, explicitly, a gambler. The Bayesian imagines that she is choosing among possible bets she could make on the outcome of interest, where different bets are offered by different bookies at different prices. We expect the Bayesian's reasoning to reflect the *consequences*, with respect to potential wins and losses, that would result from making decisions on the basis of her beliefs. This is why Savage referred to Bayesian probabilities as 'personal': they reflect the reasoner's 'skin in the game'. This is in contrast to frequencies, which are 'impersonal' in that they are objective facts that are independent of anyone's special perspective – or attitudes towards risk. The coin-toss frequency is merely an abstract fact until some people decide to condition outcomes that matter to them on some particular instance or sequence of instances; then their knowledge of the frequency becomes evidence they can use in estimating probabilities.

The Bayesian modeller starts from her beliefs about outcome risks, reflected in the bets she would be willing to make, before she sees any new data emerging from the process. These are called her *prior* beliefs. They might be based on similar processes she has observed in the past or read about, or on some theory she thinks might be relevant, or on a hunch whose origins she can't identify. Then she observes some new data that bear on the outcome of her betting. Bayes' Rule tells her how to *modify* her prior beliefs to reflect these data, again to be thought of as changes in the bookies with whom she chooses to wager. The modified expectation after the new data are taken into account is called the *posterior* belief. In a sequence of coin tosses, your prior might be based on the known frequency with a fair coin, but if you observe a failure of convergence to 50 per cent heads after many tosses, you should adopt a posterior belief in which the probability that the coin is fair is adjusted downwards.

Bayes' Rule states the mathematically correct way to update prior beliefs in light of evidence and derive posterior beliefs.[3] Different people can all follow the Rule, concerning the same events and observing the same data-generating process, and yet arrive at different posterior beliefs. These differences can have two sources:

different priors, or different weightings assigned to the impact of evidence on the priors. Bayes' Rule structures inferences – again, best understood as changes in preferred bets – as *conditional* on the reasoner's special priors and weightings, which she does not derive *from* the Rule, but plugs into it.

Bayes' Rule thus allows the reasoner wide latitude for personal discretion – again we see the basis for Savage's semantics. In any given case, a Bayesian modeller could have no prior beliefs at all, and simply let the new data they see completely determine their posterior beliefs. Or the person might have firm prior beliefs, and need a lot of new data before posterior beliefs change much from prior beliefs. Another person, or the same person in a different circumstance, might have dogmatic beliefs about something being almost sure to happen, and attach zero weight to any observation that violates this expectation. Such dogmatism is unwise more often than not, and obviously guarantees that information will be wasted, but sometimes it is warranted. For example, we would not change our bets on an outcome if we encountered new data bearing on it that violated established laws of physics, such as a report that some signal had reached its destination faster than light; we would instead conjecture that there was something wrong with the data.

There are of course some important relationships between frequencies and probabilities. The frequency of heads and tails in coin tosses can be so reliably estimated that it constitutes the best possible evidence on which anyone's prior expectation should be based. As long as observed coins are fair, experience will not lead to revised posterior beliefs. However, the expectation based on the frequency should not be adhered to dogmatically in a world where some people have incentives to manufacture biased coins.

Strong evidence has accumulated that biological brains implement a close approximation to Bayesian learning.[4] This is not an 'add-on' feature of brains, but basic to the way they work in general. Brains, unlike digital computers made by human engineers, do not have any facility for hard storage of information. Their memories are *dynamic dispositions* to reproduce previous response patterns, not static records. These dispositions are implicit in patterns of neural connectivity and the varying weights* on the likelihood of one neuron's response triggering a response across synapses to other neurons. Neuroscientists therefore model the brain's information-processing functions using statistical models, and it has recently become evident that the best such models are roughly Bayesian in form.

Brains scan the organism's external environment – and, in the case of animals with complex brains, such as mammals, the internal environment of other brain areas – based on patterns they have encountered in the past. These are their prior expectations.† What they search for are not confirmations of these priors, but events that depart from them, in both minor and major aspects. That is, they seek surprises. These trigger attention, which might or might not lead to action by the organism, and correction of the priors – in Bayesian terms, the formation of new posterior expectations. As with any Bayesian system, updates don't involve *erasing* priors. After all, they encoded valuable information. The learning process *adjusts* priors over time, partly in light of the statistical probability of the surprising input given *other* related priors. That is Bayesian learning.‡

Thus people – and all animals with brains – update basic expectations in the Bayesian style even though they're not aware of doing so.

---

* A weight (or weighting) is an assignment of frequencies to each element in a distribution of alternatives, with all weights summing to 1. For a simple example (with just two weights), consider the classic investment advice to people building retirement savings that they weight their portfolio as 0.6 equities and 0.4 bonds by value.

† In this context of neural information processing, we avoid saying 'prior *beliefs*' because we don't want to encourage the widespread error of thinking that people have beliefs 'in their brains' – beliefs are relationships between whole people, those with whom they compare perspectives and their environments. This has been explained over the past several decades by a leading school of thought in the philosophy of cognitive science, whose most important founding figure is Daniel Dennett.

‡ Though brains are thus truly Bayesian in their *manner* of learning, they don't rigidly follow Bayes' literal algorithmic rule, in the way that human-designed Bayesian inference software does. Though natural brains and engineered computers often process similar functions, brains don't store and execute hard rules like traditional computers do. In this respect, the 'deep learning' AI systems currently generating much commercial and social excitement are more like natural brains than traditional computers – as often in the past, engineers have learned some tricks from nature.

But here again we must stress that their beliefs – especially about culturally evolved objects such as football matches, murder suspects or financial assets – don't carve up information-processing spaces into categories in the way that the brain does. (The units of information in the brain can't be described except in the technical terminology of neuroscience.)[5] Beliefs about these kinds of things are anchored in *social* information processing, the kind of which people *are* aware because they talk about it among themselves and compare notes.

We will say more about such *social cognition* – the kind on which people can consciously reflect – later. For now, we'll just make the point that people unknowingly engage in Bayesian reasoning about the human-constructed world because cues and guidelines for doing so have culturally evolved along with the objects of reasoning themselves. When you face a new source of risk that you fear you don't understand well, and consult an expert or search the internet, you are behaving as a Bayesian, trying to sensibly update your (diffuse) prior belief.

Thus people's Bayesian brains give them a platform for general learning capacity that builds on experience, and in the complex environments built by human societies, social and cultural learning encoded in institutional practices and public messaging gets them the rest of the way.

These mechanisms are by no means always reliable. Learning in brains can be hijacked by environments for which natural selection didn't prepare them – later we'll examine a very common form of such hijacking, which causes addiction. On the social scale, people very often rely on innocently deceptive or deliberately malignant, manipulative and misleading sources. As we'll see in detail throughout the book, human brains feature imaginative capacities not shared by other brains, and these can lead them, both individually and collectively, to beliefs that are fantastically disconnected from reality.

Another way in which Bayesian learning can go wrong is by misapplication to what Savage called 'large worlds'.[6] This is commonly, but mistakenly, thought to refer to any environment about which we don't currently know enough to be able to attach any precise quantities to our expectations. But that standard interpretation doesn't get

Savage's meaning quite right.[7] If we can only apply very imprecise estimates around our expectations, that is simply expressed in the prior as involving wide distributions around estimated values. Such a 'diffuse' prior isn't very informative, but it's not nothing, and it can still function just fine as a basis for subsequent updating. We then remain in the opposite of a large world – called, naturally, a 'small world' – and Bayesian inference works soundly. What makes a world count as 'large' is if our best (prior) model of it identifies a data-generating process that only applies on a particular timescale, but then we consider a situation in which the same variables and modelled effects occur on a *different* timescale. That implies a different data-generating process. In that case, our prior isn't merely diffuse; it's likely to be *misleading*, and we shouldn't anticipate its revision to be just a matter of increasing precision. Rather, a whole new model will be required.

Now let's examine details that unite the concepts of probability and risk. Following the modelling of probabilities by reference to real or hypothetical gambles, economists refer to any choice over future outcomes as a *lottery*. To a typical non-economist, this word might conjure thoughts of the familiar government-run lotteries where people can win very large amounts of money with very small probabilities, or private lotteries such as the ubiquitous casino slot machines, which we'll discuss in detail later. But in economists' way of speaking, every instance of risk is a lottery. So when you commit to a new personal relationship, that is a lottery. When you drive over the speed limit, that is also a lottery, where possible outcomes – *prospects*, in the language of economists – are both monetary (a speeding fine) and non-monetary (reduced travel time, the fun of going fast, injury or death). Economists also refer to a special case, where only one outcome has a probability of 1, as a lottery, often calling it a 'completely safe lottery'; this clearly marks the economist's lottery concept as a special technical one, since no one (not even an economist) would talk that way in everyday life.*

* Scientists do not depart in this way from normal use of language just to establish their solidarity as nerds. They treat choices over sure things as special cases of lotteries in order to be able to frame *all* choices within a single mathematical modelling framework.

The *expected value* of a lottery is what a gambler would receive if the lottery were played out many times, and she was paid the average outcome. Suppose a bookie offers a gambler a lottery on one toss of a fair coin that pays the gambler $100 on heads and $25 on tails. (It is of course unusual for a bookie to offer a lottery that gives the gambler money no matter what happens – but recall that some risks can involve only upsides, as is the case for the gambler here. If it helps, you can imagine that the bookie sells tickets to this nice lottery, perhaps to launder ill-gotten money.) If the coin were flipped twenty times and the result was nine heads and eleven tails, the gambler would receive $1,175 which is ($100 × 9) + ($25 × 11). So the average outcome is $58.75, which is $1,175 ÷ 20. We can let a computer toss the virtual coin thousands of times, and the average will settle down at $62.50, which is ($100 × ½) + ($25 × ½). Notice that expected value as just defined is based on frequencies, not probabilities. Because everyone knows the frequency associated with fair coin tosses, they can use this frequency to inform their expectation about a *single* bet on a coin toss. *If* the gambler weights the value *to her* of every dollar, whether she loses or gains it, exactly the same, then she should assess the expected value of the lottery as $62.50. (The actual outcome from the one coin toss will of course be either $100 or $25.)

Imagine that you have bought this lottery – it doesn't matter for present purposes how much you paid for your ticket – and now someone offers you $70 for it. If you accept their offer you walk away with $70 for certain, with no risk at all. You give up the chance of winning $100, but you also avoid the much less attractive chance of winning only $25. Here is where the expected value is important, since it is a measure of what you give up, in expectation, if you agree to take the $70. You replace $62.50 in expectation with $70 for sure. Should you sell your ticket?

You might think this decision is a no-brainer. But what about an offer to give up the lottery for $40 instead? That is lower than the expected value, but higher than the $25 you would get if the toss comes up tails. By asking a number of these simple questions, varying the certain amount of money being offered in exchange for the lottery, we learn something fundamental about the *risk preferences* of the person answering the questions. It is at this point that probabilities enter the picture – risk always involves probabilities and probabilities always imply risk.

People, and other animals, vary in their risk preferences – or, as economists say, vary in the *utility* they associate with possible outcomes as a result of the risks they must accept to try to get them. They typically factor this personal 'utility function' into their decisions. Economists model this by estimating the quantitative weights people or animals apply to objective values, such as dollar amounts, or expected calories in a food option. People and other animals choose as if these weighted values, their *expected utilities*, are their probabilities.

For all organisms, the most pressing risks that must be managed as preconditions for coping with all others are the risks of not maintaining steady supplies of water and food. When economists model the behaviour and ecology of an animal, they typically frame it in terms of an energy budget that must be allocated to harvest a scarce food supply efficiently enough to support the continuance of life, for at least long enough to produce and sustain reproduction and offspring.

For example, the foraging behaviour of bumblebees has been carefully studied from this point of view,[8] as they are unusually convenient research subjects that can be released to harvest predictable resources, tracked in daylight conditions (often in research greenhouses) and always come back to the home from which they started, carrying most of what they found. The bumblebee must balance the energy needed to fly to food sources, the energy used in hovering at flowers, the returns on time and energy investments of varying quantities and qualities of nectar, and the risk of predation (at least outside researchers' greenhouses). This behavioural information can be readily translated into the mathematics of risk expectations and risk preference.

Where people are concerned, identifying risk preferences is harder in some ways and easier in others. It is harder because people have constructed vastly more complicated ecologies of risk sources than other animals must cope with, so a single human action might involve trading off dozens of risks. It is easier because people long ago culturally evolved an instrument, money, that they use as a uniform index of most of their risk valuations; thus we can get rich information on their preferred trade-offs by comparing the amounts of money they are willing to pay for different lotteries.

The *ultimate* prizes in important human lotteries should not be thought of as monetary amounts, since the value of money lies in the goods and services it is used to procure; furthermore, adults in most cultures are skilled at using money prices as proxies for their most basic scarce resource, their time. That said, in experimental settings where we investigate the *general* characteristics of human risk preferences, it is convenient to set up situations in which all the subjects are just choosing among outcomes with varying expected values in money.

In our previous example of the all-upside coin toss lottery, if somebody is willing to take any offer less than the expected value, and will give up the risky lottery in favour of the completely safe lottery being offered in exchange, we say they are *risk averse*. When someone takes the completely safe lottery in this case, they effectively leave money, in expectation, on the table. Say they agree to walk away with $55 for certain, for example, when it is offered. The difference between the expected value of the prospect and the money taken for certain, in this case $7.50, is money left behind in expectation, otherwise known as the *risk premium* for this person.

If one person has a risk premium of $7.50 for this simple lottery, and someone else has a risk premium of $10 for the same lottery, we can say that both people are risk averse *and* that the second person is *more* risk averse than the first. When people can trade risk with each other, this difference is central. Similarly, if a third person has a risk premium of just $1, then that person is still risk averse, just less risk averse than the first two people, and so on. We will see this simple idea played out in many social settings throughout the book, and all it relies on is people having different levels of risk aversion.

What if someone would have to be paid extra money for certain to give up the lottery? This is a person that would only accept a completely safe lottery greater than $62.50, let's say $67. Then this person has a negative risk premium of –\$4.50, and is called a *risk lover* or risk seeker – they pay to gamble.

Finally, what if someone just wonders why anyone would leave money on the table in expectation? We call such individuals *risk neutral*, since they behave as if they do not care about whether they get the \$62.50 for certain or with some risks of getting more or less than \$62.50, as long as the expected value of the risky lottery remains \$62.50. We might say that the only information that motivates their choices is relative expected values. Recall that in the case of a coin toss, where we know the relevant frequencies, we use this knowledge to compute the expected value. Thus, for risk-neutral choosers, the probabilities reflected in their choices will tend to mirror known frequencies.

This is one of the reasons that probabilities and frequencies often get assimilated in people's thinking. In the many cases where people make professional decisions that mainly affect other people's welfare – a financial adviser making investments on behalf of clients, or a judge deciding on a sentence, or an executive making decisions on behalf of her company – we socially insist that they behave as if they were risk neutral.* In the coin-toss example it is easy to identify risk neutrality because it will align probability exactly with frequency (so expected utility will map exactly onto expected value).

However, as we pointed out earlier, in many risky choices frequencies that apply to future cases are not known and *cannot* be known, because we have no run of observed data-generating processes that we are confident will remain stable across new circumstances. If you are offered a new job, tempting you away from a current one that you like, you cannot sensibly ask yourself, 'What is the frequency with which new jobs offered to me are better than exactly the job I have now?' The question barely makes sense, let alone having a useful answer. And there is also no value in trying to decide as if you were risk neutral, if in fact you are not, because then you'll simply

*If the people whose welfare is at stake are not risk neutral, then we should want the professional to take this into account. But we don't want the professional's *own* risk aversion or love of risk to play a role.

be ignoring something that actually matters to you, for no evident reason.

In general, economic theory offers no advice to people about risk preferences. It is in no sense 'irrational' to be risk averse or risk loving. As a matter of fact, experimentally derived evidence shows that the large majority of people are moderately risk averse.[9] There is no reason why this should be regarded as a problem, for them or for society.

The variability of risk preferences plays a central role in the story we'll be telling throughout this book. Our discussion to this point has focused on the risk preferences of individuals. But of key importance to the coming story are differences between individual risk preferences and risk preferences of groups to which these individuals belong. A major theme of the book is that our species *collectively* engages in bold gambles *despite* the risk aversion of most individual people.

Risk Management by Groups and Populations

The idea that a group of people often – indeed, *typically* – has risk preferences that differ from the individuals that make up that very group might seem surprising. Indeed, it might seem strange to think of groups as having risk preferences at all. After all, Bayesian probabilities are often referred to as *subjective* probabilities, and in everyday psychology subjectivity is usually associated with subjective *feelings;* but groups don't literally have feelings.

When theorists refer to 'subjective' probabilities, what they mean to emphasise is that these probabilities aren't objective frequencies – they depend on information that varies from agent to agent, and on people's risk preferences.* Sometimes such variation reflects psycho-

*This statement skates over some technical points that matter a lot to economists. If two people's patterns of choice behaviour are best described by the mathematical model known as Expected Utility Theory (EUT), then all differences in their choices (in identical lotteries) result only from differences in their levels of risk aversion or differences in the information they have. But majorities of people *weight* their probability estimates depending on logical differences, but sometimes it rests only on varying social or cultural contexts, or ways of processing information. Thus the kind of subjectivity involved in risk preferences needn't involve emotional states of the kind that only individual organisms have.

We will refer to a range of tools used to manage risks by individual people (and elephants), as well as households, companies, government agencies and other collective agents, each of them expending resources to change the risks they face, as modifying the agents' *risk-environment niches*. In much of the book we will apply this perspective on risk management to the level of whole species. Indeed, one way to understand what we will do is to see us as sketching the history of humans – and also elephants – in terms of the evolution of socially constructed risk-environment niches.

In much of the book ahead we'll therefore be focused on risks encountered at *population levels* – for example, by all the humans living in Southern and Eastern Africa in the late Pleistocene, or all the mammoths in North America during the last ice age, or even the entire species *Homo sapiens* or the entire family *Elephantidae*. This is apt to seem a bit puzzling when conjoined with the idea of 'management'. What sense can it make to say that a population, or a whole species, perceives a risk and manages it? They cannot all hold conferences to make joint plans. Most of the individuals in a population never encounter one another, or even know who's in the population.

Collective risk management isn't mysterious in cases where everyone is organised in a hierarchy, with a boss or small committee at the top who makes decisions and passes them down to people with well-defined institutional roles. But this form of collective risk response can work only for problems that are explicitly conceived in advance. Most of the major risks encountered by humanity as a whole,

whether the outcomes they're considering are unusually good or unusually bad, according to their own utilities. Such people's choices are modelled by the mathematical model known as Rank-Dependent Utility (RDU). Technically, an EUT pattern can be represented as a special case of an RDU pattern in which weights are one, so the economist can capture everyone's behaviour in a general RDU model that has flexible parameters for the weights (allowing them to be estimated as one for EUT people). This is part of the reason why we say that probabilities – unlike frequencies – are sensitive to risk preferences.

or by whole human cultural communities, were not explicitly considered as collective challenges by anyone before the fact. And strict hierarchies can't generally solve modern, big-scale collective risk problems, because hierarchical structures become hopelessly inefficient when they involve very large numbers of people. Bottom-up, 'self-organising' *cultural evolution* of problems and solutions is much more general and important.

When thinking about evolution, in the case of socially intelligent animals such as humans and elephants, we must always be clear about differences between, along with relations between, genetic evolution and cultural evolution. Many basic risks encountered by people and other animals have been solved by the development of our genes. The risk of getting small irritants in your eyes is (imperfectly) solved by eyelashes, which were provided for you by natural – genetic – selection. Later we'll devote attention to the evolution of brain structures, in both people and other animals, especially elephants. Of course, the basic anatomy of brains, like eyelashes, derives from information processed by DNA and RNA. However, even here, as we'll see in detail, natural selection didn't operate separately from, or in historical advance of, problems and solutions that emerged from human choices. Therefore cultural evolution played a role in the origins of what is distinctive about the human brain.

At the most general possible level, cultural evolution is driven by the competition among *ideas* to capture the attention of human minds. Ideas that become influential through cultural evolution do so under the influence of reasons, but the macro-scale reasons at the population level generally aren't the same as the micro-scale reasons of which individuals are aware.[10] For example, when the British 'decided' to drive on the left, and the French 'decided' to drive on the right, there was no collective deliberation about these decisions. Nor did each individual formulate their own preferred policy and then try to persuade everyone else to favour it by appealing to reasons. On the other hand, that both countries 'decided' one way or the other, and thereby managed risk of accidents once they had carriages that were hard to manoeuvre quickly, *is* explained by reference to a shared reason: once a majority had drifted into favouring one side of the road, each individual could quickly spot the value of following the most widespread habit. Thus each culture as a whole decided *without deliberation* on a risk-management solution that was then available for individuals to take up and be aware of.

The fact that human attention, at *both* individual and collective levels, is a scarce resource creates competitive pressure among the possible objects of that attention. Celebrities compete in this way: not everyone can be famous, and the fame of some crowds out the fame of others.* What counts as an 'idea' here is as varied as the things people can remember and bring to the attention of others. There are numerous Christmas carols that most people have never heard, and a few that almost everyone can sing by heart from childhood. These aren't necessarily the *best* songs from an individual's point of view. Many people would prefer *not* to have 'Jingle Bells' play in their heads and would be pleased if they could forget it. If this includes you, we apologise for what we have just inflicted on you, and hope that we can provide relief by now mentioning 'Silent Night', a more palatable earworm (to most).

Competition between Christmas songs is not very important to human welfare, but competition between technological ideas and between political ideologies certainly is. For people, ideas are major sources of risk. Just as with 'Jingle Bells', ideas can be successful in cultural evolution while harming many or even all people that attend to them. For example, racism, the idea that some human gene lines are morally, emotionally and cognitively superior to others, is a cultural evolutionary 'success' in its own terms, in that it spreads tenaciously and suppresses competing ideas; but for the people whose minds it colonises it carries huge downside risk and little upside risk. Like a resilient garden weed, there is no human culture in which racism won't tend to spread unless significant resources are devoted, continuously, to suppressing it, yet racism isn't good for the welfare of anyone, including, in the long run, racists. To cite a less egregious example, the word processing package Microsoft Word drove the once-popular WordPerfect to near extinction towards the end of the

---

* Sometimes these relationships are complementary instead of competitive; the members of Monty Python all helped to make one another famous.