CHAPTER

5 Judgment under Risk and Uncertainty

5.1 Introduction

The previous chapter showed how a powerful theory of probabilistic judgment can be built on the foundation of a limited number of relatively uncontroversial axioms. Though the axioms might seem weak, the resulting theory is anything but. In fact, the theory is open to criticism, especially on descriptive grounds. In this section, we consider whether the theory can serve as a descriptively adequate theory, that is, whether it captures how people actually make probabilistic judgments, and we explore a series of phenomena that suggest that it does not. The discrepancy suggests that a descriptively adequate theory of judgment must differ from the theory that we just learned. We will also continue our study of the building blocks of behavioral theory. In particular, we will continue the discussion of the heuristics-and-biases program in Chapter 3, by reviewing more heuristics and discussing the biases that these heuristics can lead to. Thus, this chapter gives a better idea of how behavioral economists respond to discrepancies between observed behavior and standard theory.

5.2 The gambler's fallacy

The notion of independence that we encountered in Section 4.3 is absolutely critical in economics and finance. If you are managing investments, for example, you are supposed to diversify. It would be unwise to put all of your eggs in one basket, investing all your money in a single asset such as Google stock. But in order to diversify, it is not enough to invest in two or more assets: if you invest in stocks that will rise and fall together, you do not actually have a diversified portfolio. What you should do is to invest your money in assets that are sufficiently independent. In real-world investment management, a great deal of effort goes into exploring whether assets can be assumed to be independent or not.

The notion of independence is also very important in fields like engineering. If you are in the process of designing a new nuclear power plant, you should include multiple safety systems that can prevent nuclear meltdown. But having five safety systems instead of one gives you additional safety only when a breakdown in one system is sufficiently independent from a breakdown in the other. If, for example, all safety systems are held together with one bolt, or plugged into the same outlet, a breakdown in the one system is not independent from a breakdown in the other, and your plant will not be as safe as it could be. In nuclear power plant design, and elsewhere, a great deal of effort goes into making sure that different systems (if they are all critical to the operation of the machine) are sufficiently independent from each other.

Exercise 5.1 The eggs and the basket Use the concept of independence to explain why you should not put all of your eggs in one basket.

Exercise 5.2 The alarms Knowing that alarm clocks sometimes fail to go off, some people set multiple alarms to make sure they will wake up in the morning.

(**a**) In order to work as intended, should the different alarms be dependent or independent?

(**b**) If you set multiple alarms on your smartphone, are they dependent or independent?

In principle, there are two ways in which you might make a mistake about independence: you may think that two outcomes are independent when in fact they are not, or you may think that two outcomes are not independent when in fact they are. People make both kinds of mistake. We will consider them in order.

Thinking that two outcomes are independent when in fact they are not happens, for instance, when people invest in stocks and bonds on the assumption that they are completely independent. In reality they are not. One of the important take-home lessons of the recent economic crisis is that a vast range of assets – US stocks, Norwegian real estate, etc. – are probabilistically dependent because of the highly international nature of modern finance and the complicated ways in which mortgages and the like are packaged and sold. Thinking that two outcomes are independent when in fact they are not also happens when people build nuclear power plant safety systems that have parts in common or that depend on the sobriety of one manager or the reliability of one source of electric power.

Thinking that two outcomes are dependent when in fact they are not occurs, for example, when people think that they can predict the outcome of a roulette wheel based on its previous outcomes. People cannot: these things are set up in such a way as to make the outcomes completely independent. And they are set up that way because it is a good way for the casino to make sure customers are unable to predict the outcomes. Nevertheless, many people believe that they can predict the outcome of a roulette game. For evidence, the internet offers a great deal of advice about how to beat the casino when playing roulette: "monitor the roulette table," "develop a system," "try the system on a free table before operating it for financial gain," and so on. Do an internet search for "roulette tips," and you will find a long list of webpages encouraging you to think of various outcomes as probabilistically dependent, when they are not.

Exercise 5.3 Winners Purveyors of lottery tickets are fond of posting signs such as that in Figure 5.1. Of the two mistakes we have identified in this section, which one are they hoping you will make today?

Figure 5.1 Winning ticket sold here

Exercise 5.4 Threes "Bad things always happen in threes," people sometimes say, pointing to evidence such as the fact that Janis Joplin, Jimi Hendrix and Jim Morrison all died within a few months of each other in late 1970 and early 1971. What sort of mistake are these people making?

One specific case of thinking that two outcomes are dependent when in fact they are not is the **gambler's fallacy**: thinking that a departure from the average behavior of some system will be corrected in the short term. People who think they are "due for" a hurricane or a car accident or the like because they have not experienced one for a few years are committing the gambler's fallacy. Here, I am assuming that hurricanes and car accidents are uncorrelated from year to year. It is possible that thinking you are due for a car accident makes you more likely to have one; if so, a number of accident-free years might in fact make it more likely for you to have an accident.

The following exercises illustrate how easy it is to go wrong.

Exercise 5.5 Gambler's fallacy Carefully note the difference between the following two questions:

(a) You intend to flip a fair coin eight times. What is the probability that you end up with eight heads?

(b) You have just flipped a fair coin seven times and ended up with seven heads. What is the probability that when you flip the coin one last time you will get another heads, meaning that you would have flipped eight heads in a row?

The gambler's fallacy is sometimes explained in terms of **representativeness**. We came across heuristics in Section 3.6 on anchoring and adjustment. According to the heuristics-and-biases program, people form judgments by following heuristics, or rules of thumb, which by and large are functional but which sometimes lead us astray. The **representativeness heuristic** is such a heuristic. When you employ the representativeness heuristic, you estimate the probability that some outcome was the result of a given process by reference to the degree to which the outcome is representative of that process. If the outcome is highly representative of the process, the probability that the former was a result of the latter is estimated to be high; if the outcome is highly unrepresentative of the process, the probability is estimated to be low.

The representativeness heuristic can explain the gambler's fallacy if we assume that a sequence like HHHHHHHH seems less representative of the process of flipping a fair coin eight times than a sequence like HHHHHHHT, which seems less representative than a sequence like HTTTHHTH. If you use the representativeness heuristic, you will conclude that the first sequence is less likely than the second, and that the second is less likely than the third. In reality, of course, the three are equally likely *ex ante*.

Exercise 5.6 Representativeness Which of the following two outcomes will strike people who use the representativeness heuristic as more likely: getting 4-3-6-2-1 or 6-6-6-6-6 when rolling five dice?

The representativeness heuristic might be perfectly functional in a wide variety of contexts. If it is used, for example, to infer that kind acts are more likely to be performed by kind people, and that mean acts more likely to be performed by mean people, the representativeness heuristic can protect us from adverse events. But because it can generate predictable and systematic patterns of mistakes, it can lead to bias, just as anchoring and adjustment can. For another example, consider the following case:

Exercise 5.7 Let us assume that whenever one gets pregnant, there is a 1/100 chance of having twins, and that being pregnant with twins once will not affect the probability of being pregnant with twins later.

(a) You are not yet pregnant, but intend to get pregnant twice. What is the probability that you will be pregnant with twins twice?

(b) You have just had a set of twins, and intend to get pregnant one more time. What is the probability that you will end up pregnant with twins again, that is, that you will have been pregnant with twins twice?

Again, having two sets of twins might strike a person as extraordinarily unrepresentative of the process that generates children. Thus, people relying on the representativeness heuristic will think of the probability of having a second set of twins conditional on having one set already as considerably smaller than the probability of having a set of twins the first time around. But by assumption, of course, these probabilities are equal.

Exercise 5.8 Mr Langford, cont. Suppose that Langford from Exercise 4.49 on pages 94–95 has just won two jackpots in a row and is about to play a third time. What is the probability that he will win a third time, so as to make it three jackpots in a row?

One way to explain the ubiquity of the gambler's fallacy is to say that people believe in the **law of small numbers**. That is, people exaggerate the degree to which small samples resemble the population from which they are drawn. In the case of the coins, the "population" consists of half heads and half tails. A believer in the law of small numbers would exaggerate the degree to which a small sample (such as a sequence of eight coin flips) will resemble the population and consist of half heads and half tails.

It is important to note, however, that there are games of chance in which outcomes are correlated. In Blackjack, for example, cards are drawn from a deck without being replaced, which means that the probability of drawing a given card will vary from draw to draw. In principle, then, you can beat the house by counting cards, which is why casinos reserve the right to throw you out if you do.

5.3 Conjunction and disjunction fallacies

We have already (in Section 4.3) come across the conjunction fallacy. "*A* AND *B*" is a conjunction; you commit the conjunction fallacy when you overestimate the probability of a conjunction. Consider the probability of winning the Lotto 6/49 (see Exercise 4.28 on page 85). When people learn the answer for the first time they are often shocked at how low it is. They are, in effect, overestimating the probability that the first number is right AND the second number is right AND ... and so on. Because they are overestimating the probability of a conjunction, they are committing the conjunction fallacy.

Example 5.9 Boeing aircraft A Boeing 747–400 has around 6 million parts. Suppose that each part is very reliable and only fails with probability 0.000,001. Assuming that failures are independent events, what is the probability that all parts work?

The probability that any one part works is 0.999,999, so the probability that all parts work is $(0.999{,}999)^{6{,}000{,}000} \approx 0.0025 = 0.25$ percent.

Given these numbers, the probability that all parts in a 747 work is only about a quarter of a percent! If this figure was lower than you expected, you may have committed the conjunction fallacy. Still, airplane crashes remain rare because planes are built with a great deal of redundancy, so that any one failure does not necessarily lead to a crash. That said, not all machines can be built in this way: some helicopters famously depend on a single rotor-retaining nut in such a way that if the nut fails, the whole machine will come crashing down. The term "Jesus nut" is sometimes used to denote a part whose failure would lead to a breakdown of the whole system. Presumably, the name is due to the only thing that can save you if the nut fails, though this assumes that a Jesus intervention is sufficiently independent of a nut failure.

The conjunction fallacy is particularly important in the context of **planning**. Complex projects are puzzles with many pieces, and typically each piece needs to be in place for the project to be successful. Even if the probability that any one piece will fall into place is high, the probability that all pieces of the puzzle will fall into place may be low. Planners who commit the conjunction fallacy will overestimate the probability of the conjunction – the proposition that the first piece is in place AND the second piece is in place AND the third piece is in place, and so on – meaning that they will overestimate the probability that the project will succeed.

The **planning fallacy** is the mistake of making plans based on predictions that are unreasonably similar to best-case scenarios. Many projects – senior theses, doctoral dissertations, dams, bridges, tunnels, railroads, highways, and wars – frequently take longer, and cost more, than planned. (Embarrassingly, that includes the book you are reading.) Here are two famous examples:

Example 5.10 The Sydney Opera House Many people consider the Sydney Opera House to be the champion of all planning disasters. According to original estimates in 1957, the opera house would be completed early in 1963 for $7 million. A scaled-down version of the opera house finally opened in 1973 at a cost of $102 million.

Example 5.11 Rail projects In more than 90 percent of rail projects undertaken worldwide between 1969 and 1998, planners overestimated the number of passengers who would use the systems by 106 percent. Cost overruns averaged 45 percent. Even more interestingly, estimates over the course of this period did not improve. That is, knowing about others' failed forecasts did not prompt planners to make theirs more realistic.

There is a related fallacy called the **disjunction fallacy**. "*A* OR *B*" is a disjunction; you commit the disjunction fallacy when you underestimate the probability of a disjunction. To illustrate, let us build upon Exercise 4.26(d) on page 82, in which you computed the probability of rolling at least one six when rolling two dice.

Example 5.12 Compute the probability of getting at least one six when rolling (**a**) one die, (**b**) two dice, (**c**) three dice, (**d**) ten dice.

(**a**) The probability of rolling at least one six when rolling one die equals one minus the probability of rolling a non-six, which equals $1 - 5/6 \approx 16.6$ percent.

(**b**) The probability of rolling at least one six when rolling two dice equals one minus the probability of rolling no sixes in two trials, which equals $1 - (5/6)^2 \approx 30.6$ percent.

(**c**) The probability of rolling at least one six when rolling three dice equals one minus the probability of rolling no sixes in three trials, which equals $1 - (5/6)^3 \approx 42.1$ percent.

(**d**) Finally, the probability of rolling at least one six when rolling ten dice equals one minus the probability of rolling no sixes in ten trials, which equals $1 - (5/6)^{10} \approx 83.8$ percent.

Notice how quickly the probability of rolling at least one six rises as the number of trials increases. The probability will never reach 100 percent, but it will approach it asymptotically, so that it will get closer and closer as the number of trials increases. If the resulting numbers here are greater than you expected, you may have committed the disjunction fallacy. For the probability of rolling at least one six in multiple trials equals the probability of rolling a six in the first trial, OR rolling a six in the second trial, OR ... And given the definition above, if you underestimate the probability of the disjunction, you are committing the disjunction fallacy.

Exercise 5.13 Hiking You plan to go on a hike in spite of the fact that a tornado watch is in effect. The national weather service tells you that for every hour in your area, there is a 30 percent chance that a tornado will strike. That is, there is a 30 percent chance that a tornado will strike your area between

10 am and 11 am, a 30 percent chance that a tornado will strike your area between 11 am and noon, and so on.

(a) What is the probability of a tornado striking your area at least once during a two-hour hike?
(b) What is the probability of a tornado striking your area at least once during a three-hour hike?
(c) What is the probability of a tornado striking your area at least once during a ten-hour hike?

Exercise 5.14 Flooding Imagine that you live in an area where floods occur on average every ten years. The probability of a flood in your area is constant from year to year. You are considering whether to live in your house for a few more years and save up some money, or whether to move before you lose everything you own in the next flood.

(a) What is the probability that there will no floods in your area over the course of the next two years?
(b) What is the probability that there will be exactly one flood in your area over the course of the next two years?
(c) What is the probability that there will be at least one flood over the course of the next two years?
(d) What is the probability that there will be at least one flood over the course of the next ten years?

Exercise 5.15 Terrorism Compute the probability that at least one major terrorist attack occurs over the course of the next ten years, given that there are 365.25 days in an average year, if the probability of an attack on any given day is 0.0001.

That last exercise illustrates an infamous statement by the Irish Republican Army (IRA), which for decades fought a guerilla war for the independence of Northern Ireland from the United Kingdom and a united Ireland. In the aftermath of an unsuccessful attempt to kill British Prime Minister Margaret Thatcher in 1984 by planting a bomb in her hotel, the IRA released a statement that ended with the words: "You have to be lucky all the time. We only have to be lucky once."

The disjunction fallacy is particularly important in the context of **risk assessment**. When assessing the risk that some complex system will fail, it is often the case that the system as a whole – whether a car or an organism – critically depends on multiple elements in such a way that the failure of any one of these elements would lead to a breakdown of the system. Even if the probability that any one element will fail is low, the probability that at least one element will fail may be high. Assessors who commit the disjunction fallacy will underestimate the probability of the disjunction – the proposition that the first element fails or the second element fails or the third element fails, and so on – meaning that they will underestimate the probability of a system breakdown.

There is an obvious symmetry between the two fallacies discussed in this section. According to de Morgan's law, $A\&B$ is logically equivalent to $\neg[\neg A \vee \neg B]$ (see Section 2.4). So if you overestimate the probability $\Pr(A\&B)$, this is the same as saying that you overestimate $\Pr(\neg[\neg A \vee \neg B])$. But by the NOT rule, that is the same as saying that you overestimate $1 - \Pr(\neg A \vee \neg B)$, which is to say that you underestimate $\Pr(\neg A \vee \neg B)$. In the context of the Linda example, overestimating the probability that she is a feminist bank teller is (according to de Morgan's law) the same as underestimating the probability that she is a non-feminist or a non-bank teller. In sum, if you adhere to de Morgan's law, then you commit the conjunction fallacy if and only if you commit the disjunction fallacy.

Both the conjunction and disjunction fallacies can be explained in terms of anchoring and adjustment (see Section 3.6). People overestimate the probability of conjunctions – and therefore commit the conjunction fallacy – if they use the probability of any one conjunct as an anchor and adjust downwards insufficiently. They underestimate the probability of disjunctions – and therefore commit the disjunction fallacy – if they use the probability of any one disjunct as an anchor and adjust upwards insufficiently. Here are more exercises:

Exercise 5.16 What is the probability of drawing at least one ace when drawing cards from an ordinary deck, with replacement, when you draw: **(a)** 1 card, **(b)** 2 cards, **(c)** 10 cards, and **(d)** 52 cards?

Exercise 5.17 The birthday problem Suppose that there are 30 students in your behavioral economics class. What is the probability that no two students have the same birthday? To make things easier, assume that every student was born the same non-leap year and that births are randomly distributed over the year.

Exercise 5.18 The preface paradox In the preface to your new book, you write that you are convinced that every sentence in your book is true. Yet you recognize that for each sentence there is a 1 percent chance that the sentence is false. **(a)** If your book has 100 sentences, what is the probability that at least one sentence is false? **(b)** What if your book has 1000 sentences?

Finally, an exercise about aviation safety.

Exercise 5.19 Private jet shopping Suppose you are fortunate (or delusional) enough to be shopping for a private jet. You have to decide whether to get a jet with one or two engines. Use p to denote the probability that an engine fails during any one flight. A "catastrophic engine failure" is an engine failure that makes the plane unable to fly.

(a) One of the jets you are looking to buy has only one engine. What is the probability of a catastrophic engine failure during any one flight with this plane?
(b) Another jet you are looking to buy has two engines, but is unable to fly with only one functioning engine. Assume that engine failures are independent events. What is the probability of a catastrophic engine failure during any one flight with this plane?
(c) Which jet strikes you as safer?
(d) What if the twin-engine jet can fly with only one functioning engine?

The answer to Exercise 5.19(b) is far from obvious. To help you out, consider constructing a table as in Figure 4.1 on page 83.

5.4 Base-rate neglect

One source of imperfect reasoning about probabilities results from the confusion between conditional probabilities $\Pr(A|B)$ and $\Pr(B|A)$. It might seem obvious that these two are distinct. As we know from Section 4.4, the probability that a randomly selected human being is a smoker is obviously different from the probability that a randomly selected smoker is a human being. However, there are contexts in which it is easy to mix these two up. In this section, we will consider some of these contexts.

Example 5.20 Mammograms Doctors often encourage women over a certain age to participate in routine mammogram screening for breast cancer. Suppose that from past statistics about some population, the following is known. At any one time, 1 percent of women have breast cancer. The test administered is correct in 90 percent of the cases. That is, if the woman does have cancer, there is a 90 percent probability that the test will be positive and a 10 percent probability that it will be negative. If the woman does not have cancer, there is a 10 percent probability that the test will be positive and a 90 percent probability that it will be negative. Suppose a woman has a positive test during a routine mammogram screening. Without knowing any other symptoms, what is the probability that she has breast cancer?

When confronted with this question, most people will answer close to 90 percent. After all, that is the accuracy of the test. Luckily, we do not need to rely on vague intuitions; we can compute the exact probability. In order to see how, consider Figure 5.2, in which C denotes the patient having cancer and P denotes the patient testing positive. Plugging the numbers into Bayes's rule (Proposition 4.42), we get:

$$\Pr(C|P) = \frac{\Pr(P|C) * \Pr(C)}{\Pr(P|C) * \Pr(C) + \Pr(P|\neg C) * \Pr(\neg C)}$$

$$= \frac{0.9 * 0.01}{0.9 * 0.01 + 0.1 * 0.99} \approx 0.08.$$

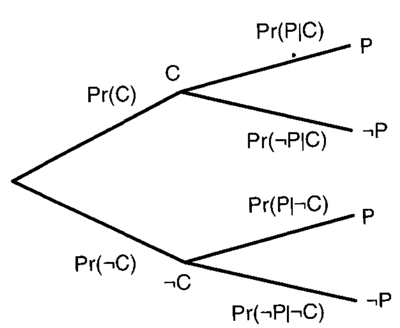

Figure 5.2 Breast cancer screening

Notice that the probability that somebody who has been identified as having cancer in fact has cancer is not equal to – in fact, not even remotely similar to – the accuracy of the test (which in this case is 90 percent). The probability that the woman has cancer is only about 8 percent – much lower than people think. This is paradoxical because we know that the test is reasonably good: it gives the correct outcome in 90 percent of cases. What we are forgetting is that relatively few people actually have cancer. Out of 1000 people, only about ten can be expected to have cancer. Of those, nine will test positive. Of the 990 women who do not have cancer, only 10 percent will test positive, but that is still 99 people. So only nine of the 108 people who test positive actually have cancer, and that is about 8 percent. Notice that this case is similar to the frisbee case: although the new machine has a lower failure rate than the old one, the average frisbee produced in the factory is more likely than not to come from the new machine, simply because it produces so many more frisbees.

The fraction of all the individuals in the population who have cancer (or some other characteristic of interest) is called the **base rate**. In the cancer case, the base rate is only one percent. One way to diagnose the mistake that people make is to say that they fail to take the base rate properly into account. Thus, the mistake is sometimes referred to as **base-rate neglect** or the **base-rate fallacy**. The judgment that we make in these situations should reflect three different factors: first, the base rate; second, the evidence; third, the conditional probabilities that we would see the evidence when the hypothesis is true and when it is false. We commit the base-rate fallacy when we fail to take the first of these three factors properly into account.

Incidentally, this example makes it clear why younger women are not routinely tested for breast cancer. In younger women, the base rate would be even lower; so the ratio of true positives to all positives would be even lower. Notice that in the previous example, when a woman from the relevant population gets a positive result the probability that she has cancer only increases from one percent to about eight percent, which is not a very large increase. If the base rate were even lower, the increase would be even smaller, and the conditional probability $\Pr(C \mid P)$ would not be very different from the base rate $\Pr(C)$. When this is the case, the test does not give the doctor any additional information that is relevant when producing a diagnosis, and so the test is said to be **non-diagnostic**. If the base rate were very high, the test would still not be diagnostic. In order for a test to be diagnostic, it helps if the base rate is somewhere in the middle.

Exercise 5.21 Mammograms, cont. Men can get breast cancer too, although this is very unusual. Using the language of "base rates" and "diagnosticity," explain why men are not routinely tested for breast cancer.

Testimony can be non-diagnostic, as the following classic example illustrates.

Exercise 5.22 Testimony A cab company was involved in a hit-and-run accident at night. Two cab companies, the Green and the Blue, operate in the city. You are given the following data: 85 percent of the cabs in the city are Green, 15 percent are Blue. A witness identified the cab involved in the accident as

Blue. The court tested the reliability of the witness under the same circumstances that existed on the night of the accident and concluded that the witness correctly identified each one of the two colors 80 percent of the time and failed 20 percent of the time. What is the probability that the cab involved in the accident was Blue rather than Green?

Exercise 5.23 Iron Bowl At an Auburn–Alabama game, 80 percent of attendees wore Alabama gear and 20 percent wore Auburn gear. During the game, one of the attendees apparently robbed a beer stand outside the stadium. A witness (who was neither an Alabama nor an Auburn fan) later told police that the robber wore Auburn gear. The witness, however, was the beer stand's best customer, and it was estimated that he would only be able to identify the correct gear about 75 percent of the time. What is the probability that the robber wore Auburn gear, given that the witness said that he did?

Here is a slightly different kind of problem.

Exercise 5.24 Down syndrome The probability of having a baby with Down syndrome increases with the age of the mother. Suppose that the following is true. For women 34 and younger, about one baby in 1000 is affected by Down syndrome. For women 35 and older, about one baby in 100 is affected. Women 34 years and younger have about 90 percent of all babies. What is the probability that a baby with Down syndrome has a mother who is 34 years or younger?

Base-rate neglect helps explain the planning fallacy: the fact that plans and predictions are often unreasonably similar to best-case scenarios (see Section 5.3). Notice that people who fall prey to the planning fallacy are convinced that their own project will be finished on time, even when they know that the vast majority of similar projects have run late.

From the point of view of the theory we explored in the previous chapter, the planning fallacy is surprising. If people updated their beliefs in Bayesian fashion, they would take previous overruns into account and gradually come up with a better estimate of future projects. We can, however, understand the optimistic estimates as a result of base-rate neglect. It is possible to think of the fraction of past projects that were associated with overruns as the base rate, and assume that people tend to ignore the base rate in their assessments.

The last problems in this section all relate to the war on terror.

Exercise 5.25 Jean Charles de Menezes In the aftermath of the July 21, 2005, terrorist attacks in London, British police received the authority to shoot terrorism suspects on sight. On July 22, plainclothes police officers shot and killed a terrorism suspect in the London underground. Use Bayes's rule to compute the probability that a randomly selected Londoner, identified by the police as a terrorist, in fact is a terrorist. Assume that London is a city of 10 million people, and that ten of them (at any given time) are terrorists. Assume also that police officers are extraordinarily competent, so that their assessments about whether a given person is a terrorist or not are correct 99.9 percent of the time.

The suspect, Jean Charles de Menezes, 27, was shot seven times in the head and once in the shoulder. He was later determined to be innocent. Notice, again, that the probability that somebody who has been identified as a terrorist is in fact a terrorist is not equal to – in fact, not even remotely similar to – the accuracy of the police officers' assessments. Notice, also, the time line: in this case, it is as though the police went out of their way to prove as soon as possible that they cannot be entrusted with the authority to execute people on sight.

Exercise 5.26 Behavior detection The following passage is from *USA Today:*

> Doug Kinsey stands near the security line at Dulles International Airport, watching the passing crowd in silence. Suddenly, his eyes lock on a passenger in jeans and a baseball cap.
>
> The man in his 20s looks around the terminal as though he's searching for something. He chews his fingernails and holds his boarding pass against his mouth, seemingly worried.
>
> Kinsey, a Transportation Security Administration [TSA] screener, huddles with his supervisor, Waverly Cousins, and the two agree: The man could be a problem. Kinsey moves in to talk to him.
>
> The episode this month is one of dozens of encounters airline passengers are having each day – often unwittingly – with a fast-growing but controversial security technique called behavior detection. The practice, pioneered by Israeli airport security, involves picking apparently suspicious people out of crowds and asking them questions about travel plans or work. All the while, their faces, body language and speech are being studied.
>
> The TSA has trained nearly 2,000 employees to use the tactic, which is raising alarms among civil libertarians and minorities who fear illegal arrests and ethnic profiling. It's also worrying researchers, including some in the Homeland Security Department, who say it's unproven and potentially ineffectual.

The government is unlikely to reveal data on the efficacy of this program, but we can make some reasonable assumptions. Every month, roughly 60 million people fly on US carriers. Let us imagine that 6 of them are terrorists. Let us also imagine that the TSA personnel are highly competent and will correctly identify a person as a terrorist or non-terrorist in 98/100 of cases. Questions:

(a) What is the probability that a passenger selected at random is a terrorist and is correctly identified as such by TSA personnel?
(b) What is the probability that a passenger selected at random is *not* a terrorist but is nevertheless (incorrectly) identified as a terrorist by TSA personnel?
(c) What is the probability that a passenger in fact is a terrorist conditional on having been identified as such by TSA personnel?
(d) Is this test diagnostic?

Notice that in the story above, the man was apparently a false positive, meaning that the story inadvertently ended up illustrating the lack of diagnosticity of the test.

Exercise 5.27 Diagnosticity Let us take it for granted that the behavior-detection test (from Exercise 5.26) is not diagnostic. The test may still be diagnostic in another setting, say, at a checkpoint at the US embassy in Kabul, Afghanistan. Explain how this is possible.

Recent evidence suggests that behavior-detection agents are not in fact very good at reading body language, in which case the program – which critics call "security theater" – would not work in Kabul either. On a related note: since 2004, the US Department of Homeland Security's US-VISIT Program, now called the Office of Biometric Identity Management, collects digital fingerprints and photographs from international travelers at US visa-issuing posts and ports of entry. Before long, the database might contain hundreds of millions of fingerprints. If a terrorist's fingerprint – recovered from a crime scene, perhaps – is found to match one in the database, what do you think the probability is that the match is the actual terrorist? If you find yourself caught up in this kind of dragnet, the only thing standing between you and the electric chair might be a jury's understanding of Bayes's rule. Good luck explaining it to them.

5.5 Confirmation bias

One striking feature of Bayesian updating in Section 4.6 is that John and Wes come to agree on the nature of the coin so quickly. As you will recall, after only about 15 flips of the coin, both assigned a probability of almost 100 percent to the possibility that the coin had two heads. People sometimes refer to this phenomenon as **washing out of the priors**. That is, after so many flips, John and Wes will assign roughly the same probability to the hypothesis, independently of what their priors used to be. This represents a hopeful picture of human nature: when rational people are exposed to the same evidence, over time they come to agree regardless of their starting point. (As is often the case, things get tricky when probabilities are zero; I continue to ignore such complications.)

In real life, unfortunately, people do not in general come to agree over time. Sometimes that is because they are exposed to very different evidence: conservatives tend to read conservative newspapers and blogs that present information selected because it supports conservative viewpoints; liberals tend to read liberal newspapers and blogs that present information selected because it supports liberal viewpoints. Yet sometimes people have access to the very same evidence presented in the very same way (as Wes and John do) but nevertheless fail to agree over time. Why is this?

Part of the story is a phenomenon psychologists call **confirmation bias**: a tendency to interpret evidence as supporting prior beliefs to a greater extent than warranted. In one classic study, participants who favored or opposed the death penalty read an article containing ambiguous information about the advantages and disadvantages of the death penalty. Rather than coming to agree as a result of being exposed to the same information, both groups of people interpreted the information as supporting their beliefs. That is, after reading the article, those who were previously opposed to the death penalty

US embassy in Kabul,

nts are not in fact very
rogram – which critics
ther. On a related note:
ty's US-VISIT Program,
nent, collects digital fin-
elers at US visa-issuing
might contain hundreds
– recovered from a crime
base, what do you think
rist? If you find yourself
anding between you and
of Bayes's rule. Good luck

on 4.6 is that John and Wes
ly. As you will recall, after
obability of almost 100 per-
s. People sometimes refer to
That is, after so many flips,
bility to the hypothesis, inde-
represents a hopeful picture
posed to the same evidence,
starting point. (As is often the
ero; I continue to ignore such

n general come to agree over
sed to very different evidence:
papers and blogs that present
servative viewpoints; liberals
at present information selected
metimes people have access to
same way (as Wes and John do)
is this?
logists call **confirmation bias**: a
prior beliefs to a greater extent
ipants who favored or opposed
mbiguous information about the
penalty. Rather than coming to
ame information, both groups of
orting their beliefs. That is, after
sly opposed to the death penalty

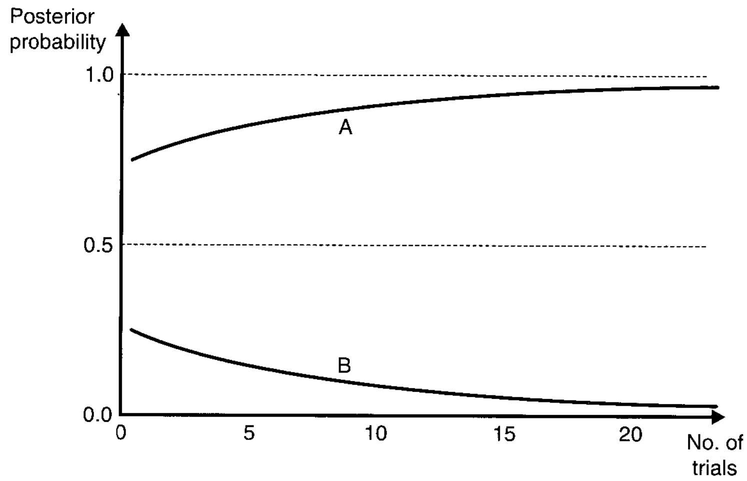

Figure 5.3 Confirmation bias

were even more strongly opposed and those who favored it were even more in favor. In the presence of confirmation bias, then, the picture of how people's beliefs change as they are exposed to the evidence may look less like Figure 4.5 on page 93 and more like Figure 5.3.

Exercise 5.28 Confirmation bias Imagine that John is suffering from confirmation bias. Which of the curves labeled A, B, and C in Figure 5.4 best represents the manner in which his probabilities change over time as the evidence comes in?

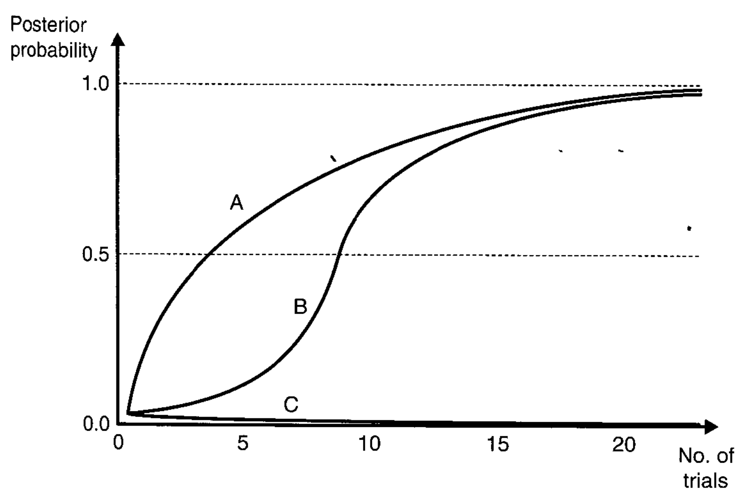

Figure 5.4 John's confirmation bias

Example 5.29 The jealous lover From literature or life, you may be familiar with the character of the jealous lover, who refuses to accept any evidence that his or her affections are reciprocated and who everywhere finds evidence fueling suspicions. As Marcel Proust, author of *In Search of Lost Time*, wrote: "It is astonishing how jealousy, which spends its time inventing so many petty but false suppositions, lacks imagination when it comes to discovering the truth." In more prosaic terms, the jealous lover exhibits confirmation bias.

Confirmation bias can explain a whole range of phenomena. It can explain why racist and sexist stereotypes persist over time. A sexist may dismiss or downplay evidence suggesting that girls are good at math and men are able to care for children but be very quick to pick up on any evidence that they are not. A racist may not notice all the people of other races who work hard, feed their families, pay their taxes, and do good deeds, but pay a lot of attention to those who do not. Confirmation bias can also explain why people gamble. Many gamblers believe that they can predict the outcome of the next game, in spite of overwhelming evidence that they cannot (they may, for example, have lost plenty of money in the past by mispredicting the outcomes). This could happen if the gambler notices all the cases when he did predict the outcome (and if the outcome is truly random, there will be such cases by chance alone) and fails to notice all the cases when he did not. The same line of thinking can explain why so many people think that they can beat the stock market, in spite of evidence that (in the absence of inside information) you might as well pick stocks randomly. Finally, confirmation bias can explain how certain conspiracy theories survive in spite of overwhelming contradictory evidence. The conspiracy theorist puts a lot of weight on morsels of evidence supporting the theory, and dismisses all the evidence undermining it.

Exercise 5.30 Reputation The fact that people exhibit confirmation bias makes it very important to manage your reputation – whether you are a student, professor, doctor, lawyer, or brand. Why?

Scientists, by the way, are not immune from confirmation bias. Philosopher of science Karl Popper noted how some scientists find data supporting their theories everywhere. He describes his encounter with Alfred Adler, the pioneering psychoanalyst. In Popper's words:

> Once, in 1919, I reported to him a case which to me did not seem particularly Adlerian, but which he found no difficulty in analysing in terms of his theory of inferiority feelings, although he had not even seen the child. Slightly shocked, I asked him how he could be so sure. "Because of my thousandfold experience," he replied; whereupon I could not help saying: "And with this new case, I suppose, your experience has become thousand-and-one-fold."

Popper's description makes it sound as though Adler is suffering confirmation bias in a big way. Sadly, it is easy to find similar examples in economics. Because of how easy it is to "confirm" just about any theory, Popper ended up arguing that the hallmark of a scientific theory was not the fact that it could be confirmed, but rather that it could at least in principle be **falsified** – shown to be false by empirical observation. A good question to ask yourself and others is: "What sort of evidence would make you change your mind?" If you cannot think of anything short of divine intervention, you are almost surely suffering confirmation bias.

Problem 5.31 Confirmation bias among economists

(a) Name two famous economists who in your view are suffering confirmation bias.
(b) Reflect upon the people you just named: if both are economists with whom you disagree, your response to (a) might itself be an expression of confirmation bias.

Psychological research suggests that confirmation bias is due to a number of different factors. First, people sometimes fail to notice evidence that goes against their beliefs, whereas they quickly pick up on evidence that supports them. Second, when the evidence is vague or ambiguous, and therefore admits of multiple interpretations, people tend to interpret it in such a way that it supports their beliefs. Third, people tend to apply a much higher standard of proof to evidence contradicting their beliefs than to evidence supporting them.

Exercise 5.32 Destroying America Explain how book titles such as *Demonic: How the Liberal Mob Is Endangering America* and *American Fascists: The Christian Right and the War on America* contribute to political polarization.

Problem 5.33 Preventing confirmation bias *In matters of politics, philosophy, religion, and so on, do you expose yourself to the ideas of people "on the other side" as you do to the ideas of people "on your side"? Are you paying as much attention to what they say? Are you applying the same standards of evidence? If you can honestly answer yes to all these questions, you belong to a small but admirable fraction of the population. If not, you might give it a try; it is an interesting exercise.*

5.6 Availability

When physicians examine children with headaches, one of the things they look for is brain tumors. The base rate is very low. Children with brain tumors are virtually certain to have headaches, but there are of course many other causes of headaches. As it happens, a simple examination successfully identifies the children with tumors in very many of the cases. That is, of all the children who have been properly examined and judged not to have a tumor, very few actually do.

Exercise 5.34 CT scans In some populations, brain tumors in children are rare: the base rate is only about 1/10,000. A child with a tumor is very likely to have occasional headaches: 99 out of 100 do. But there are many other reasons a child can have a headache: of those who do not have a tumor 1 in 10 have occasional headaches.

(a) Given that a child has occasional headaches (H), what is the probability that he or she has a brain tumor (T)?
(b) Let us say that among children with headaches, 999/1000 will ultimately be fine (F). Suppose that a physician using a simple test can correctly determine whether the child will be fine or not in 95/100 of children with headaches. Given that the doctor after performing the test gives the patient a green light (G), what is the probability that the child really will be fine?

As these exercises indicate, it is in fact very unlikely that the patient has a brain tumor provided that he or she has been properly examined. CT scans can determine almost conclusively whether the patient has a tumor, but they are prohibitively expensive. Knowing that patients who have been examined by a physician are unlikely to have a tumor, it is widely agreed that CT scans under these conditions are unjustified. However, once a physician happens to have a patient who turns out to have a tumor in spite of the fact that the examination failed to find one, the physician's behavior often changes dramatically. From then on, she wants to order far more CT scans than she previously did. Let us assume that the physician's experience with the last child did not change her values. Assuming that a drastic change in behavior must be due to a change in values or a change in beliefs, it follows that her beliefs must have changed. On the basis of what we know, did the physician update her beliefs rationally?

The story is, of course, far more complicated than it appears here. It is worth noticing, though, that the actual figures are widely known among medical doctors. This knowledge reflects the accumulated experience of far more cases than any one physician will see during her career. As a result, it seems unlikely that rational updating on the basis of one single case should have such a radical impact on a physician's behavior. So what is going on? Behavioral economists explain this kind of behavior in terms of **availability**: the ease with which information can be brought to mind when making a judgment. When the physician faces her next patient, though she at some level of consciousness still knows the figures that suggest a CT scan is uncalled for, chances are that the last case (in which she failed to find the tumor) will come to mind. It is particularly salient, in part because it happened recently, but also because it is highly emotionally loaded.

The **availability heuristic** is another prominent heuristic from the heuristics-and-biases program. When we rely on this heuristic, we assess the probability of some event occurring by the ease with which the event comes to mind. That is, the availability heuristic says that we can treat X as more likely than Y if X comes to mind more easily than Y. As pointed out in Section 3.6, heuristics are often perfectly functional. Suppose, for instance, that you happen to come across an alligator while walking to work. The chances are that images of alligators attacking other animals (including humans) will come to mind more easily than images of alligators acting cute and cuddly. If so, the availability heuristic tells you to assume the alligator is likely to be dangerous, which is obviously a helpful assumption under the circumstances. However, the availability heuristic (being a simple rule of thumb) can sometimes lead you astray, as in the case of the children with headaches. Thus, like anchoring and adjustment, availability can lead to bias.

Exercise 5.35 Contacts Your optometrist tells you that your new contacts are so good that you can wear them night and day for up to 30 days straight. She warns you that there is some chance that you will develop serious problems, but says that she has seen many clients and that the probability is low. After a week, you develop problems and see an ophthalmologist. The ophthalmologist tells you that he is the doctor who removes people's eyes when they have been injured as a result of improper contact use. He tells you that the probability of developing serious problems is high.

Use the concept of availability bias to explain how the optometrist and the ophthalmologist can report such different views about the likelihood of developing serious problems as a result of wearing contacts.

The availability heuristic can explain a variety of phenomena, including why people think violent crime is more common than other kinds of crime. Violent crime comes to mind so easily in part because images of violence can be particularly vivid, but also because it is covered so extensively in the press: "Area Man Not Shot Today" would make for a terrible headline. Availability can also explain why fears of airplane crashes, nuclear meltdowns, and terrorist attacks tend to increase dramatically shortly after such events occur, for the obvious reason that they come to mind particularly easily then. These considerations can help us see why anti-vaccination sentiments are so strong in spite of overwhelming evidence that vaccines are safe and effective. If you have a child, or even if you just hear of a child, who developed symptoms of autism shortly after being vaccinated, the dramatic series of events is likely to be highly salient and therefore strike you as more likely than it really is. Availability can also explain a variety of marketing practices. Advertising campaigns for grooming products, cigarettes, alcohol, and all sorts of other products depict users of the product as being particularly attractive and popular. If that is what comes to mind when potential buyers reflect on the consequences of using the products, availability might make attractiveness and popularity seem particularly likely outcomes. Availability can also explain why people repeat dangerous behaviors. If they once, for whatever reason, do something dangerous or reckless – drive drunk, do hard drugs, have unprotected sex, or whatever – and nothing bad happens, the salience of that event means that they might come to think the dangers have been exaggerated, which will make them more likely to engage in it again. Self-reinforcing cycles of this kind, especially when they involve multiple individuals whose beliefs and behaviors reinforce each other's, are called **availability cascades**.

The availability heuristic sheds light on the power of storytelling. As every writer knows, stories are often far more compelling than scientific data. If you doubt that, just ask a wolf. Wolves pose a trivial danger to humans: the number of verifiable, fatal attacks by wolves on humans is exceedingly low. And yet, fear of wolves runs deep. Part of the explanation is certainly that there are so many stories about big, bad wolves eating, e.g., little girls' grandmothers. As a result of all these stories, the idea of wolves attacking humans is highly salient, which means that people treat it as likely – even though the data establish it is not. Far more dangerous organisms, such as the *Salmonella* bacterium that kills some 400 people per year in the US alone, do not figure in the public imagination in the same way and consequently are not as feared as they probably should be. The power of storytelling can be harnessed to communicate risk information very effectively, but it can also do immense harm. A single story about an illegal immigrant committing a heinous crime can generate strong anti-immigration sentiments in spite of evidence of the beneficial welfare effects of migration.

Problem 5.36 Causes of death *According to the World Health Organization, the leading causes of death in the world are ischaemic heart disease, stroke, lower respiratory infections, and chronic obstructive lung disease. This makes the leading causes of death quite different from the leading sources of fear. The effect is that people spend relatively much time thinking about threats they can do little about (terrorist groups and epidemic diseases across the world, for example) and relatively little time thinking about things they can (such as cardiovascular health, which can be improved by exercise and diet). Availability can explain why people overestimate the danger posed to them by the former and underestimate that by the latter. But a no less interesting question is this: Can the power of availability and other heuristics be harnessed to encourage people to think more about things such as cardiovascular health?*

Availability bias also helps explain the base-rate fallacy. Consider the cancer case. Even if you know that there are false positives, so that a positive test does not necessarily mean that you have the disease, the chances are cases of true positives (people who were correctly diagnosed with cancer) are more likely to come to mind than cases of false positives (people who were incorrectly diagnosed with cancer). Insofar as you follow the availability heuristic, you will think of true positives as more likely than false positives. Because the actual probability of having the disease given a positive test is the ratio of true positives to all positives, an overestimation of the probability of a true positive will lead to an inflated probability of having the disease.

5.7 Overconfidence

Bayesian updating, as you know, can tell you precisely what probability you should assign to an event given the available evidence. Such probabilities are obviously relevant to statements of **confidence**: statements concerning how certain you are that various things will happen. There are many ways to express your confidence in a belief. After making an assertion, you can add: "… and I am 90 percent certain of that." Or, you can say: "It's 50–50," meaning that you are no more confident that the one thing will happen than you are that the other thing will happen. If you are in the business of providing professional forecasts, you may be used to offering 95 percent confidence intervals, which are ranges within which you expect the outcome to fall with 95 percent certainty. Thus, financial analysts might predict that a certain stock will reach $150 next year, and add that they are 95 percent certain that the value will be between $125 and $175; labor economists might predict that the unemployment rate will hit 6 percent and add that they are 95 percent sure that the actual unemployment range will fall in the 4–8 range. The more confident the analysts are in their predictive abilities, the narrower the confidence intervals will be.

When assessing statements of confidence, behavioral economists talk about **calibration**. Formally speaking, you are perfectly calibrated if, over the long run, for all propositions assigned the same probability, the proportion of the propositions that are true equals the probability you assigned. If you are calibrated and judge that something is 90 percent certain, nine times out of ten you will be right about that thing. Notice that you can be calibrated even if most of your predictions are wrong: if you are calibrated and 33 percent confident in your predictions, you will still be wrong two times out of three. Notice also that it is possible to be calibrated even when outcomes cannot be precisely predicted, e.g., because they are random. You do not know if an unbiased coin will come up heads or tails when you flip it; but if you predict with 50 percent confidence that it will come up heads, you will be perfectly calibrated.

All things equal, calibration seems to be a good thing. We certainly do expect competent people to be calibrated: if a structural engineer tells you he or she is 100 percent certain that a certain kind of house is safe to live in, you will be disappointed in him or her when half the houses of that kind collapse.

Yet, one of the most persistent findings in the literature is that people tend to exhibit **overconfidence**. That is, like the structural engineer in the previous paragraph, the certainty that people (including experts) express in their judgments tends to exceed the frequency with which those judgments are correct. In early studies, researchers asked undergraduates questions of the type "Absinthe is (a) a liqueur or (b) a precious stone?" and invited them to judge how confident they were that their answer was right. Participants who indicated that they were 100 percent certain that their answers were right were on the average correct 70–80 percent of the time. To test whether increased motivation would decrease the degree of overconfidence, the researchers asked participants to express their confidence in terms of odds (see text box on page 83) and offered participants to play a gamble based on those odds. When people said the odds that they were right were 100:1, in order to be well calibrated they should have said 4:1; when they said the odds were 100,000:1, they should have said 9:1. Though overconfidence was first studied in the lab, over time manifestations of systematic overconfidence have been found also among physicists, doctors, psychologists, CIA analysts, and others making expert judgments. Thus, overconfidence appears outside the laboratory, when knowledgeable judges make assertions within their field of specialization, and when they are motivated to provide accurate assessments.

Exercise 5.37 Meteorology Evidence suggests meteorologists are well calibrated and therefore an exception to the rule. This will strike many people as literally unbelievable. What heuristic might cause them to underestimate meteorologists' ability to offer calibrated predictions?

Studies indicate that overconfidence increases with confidence, and therefore is most extreme when confidence is high. Overconfidence is usually eliminated when confidence ratings are low, and when very low, people may even be underconfident. Overconfidence also increases with the difficulty of the judgment task. The more challenging the task is, the more likely a judge is to be overconfident; when it comes to very easy judgments, people may even be underconfident. Interestingly, overconfidence does not in general seem to decrease when people become more knowledgeable. In one famous study, the researcher asked participants in his study questions about the behaviors, attitudes, and interests of a real patient. As participants received more and more information about the patient's life, they assigned more and more confidence to their answers. Yet, their accuracy barely increased at all. Incidentally, the clinical psychologists who participated in the study – a majority of whom had PhDs – were no more accurate and no less confident than psychology graduate students and advanced undergraduates. The more educated people were actually *more* overconfident than the less educated people. It may well be true, as Alexander Pope said, that "[a] little learning is a dangerous thing" – and a lot of learning too.

Example 5.38 Apollo 11 On the 35th anniversary of the moon landing, CNN asked the crew of Apollo 11 what their biggest concern was at the time. Astronaut Neil Armstrong answered: "I think we tried very hard not to be overconfident, because when you get overconfident, that's when something snaps up and bites you. We were ever alert for little difficulties that might crop up and be able to handle those."

Is there anything you can do to be less overconfident and more calibrated? Research suggests that informing people about the prevalence of overconfidence makes little difference to their calibration. However, what does seem to help is to make highly repetitive judgments and to receive regular, prompt and unambiguous feedback. Moreover, overconfidence can be reduced by considering reasons that you might be wrong.

How is it possible for overconfidence to be so prevalent? For starters, many of our judgments are not repetitive, and we do not receive feedback that is regular, prompt, and unambiguous. Furthermore, even in the presence of outcome feedback, learning from experience is more difficult than one might think. Confirmation bias (Section 5.5) makes us overweight evidence that confirms our predictions and underweight evidence that disconfirms them; in this way, confirmation bias makes us blind to our failures. Availability bias (Section 5.6) does not help either. If the image of a situation where you were right when others were spectacularly wrong comes to mind easily and often, you might end up overestimating the probability that that sort of thing will happen again. And a phenomenon referred to as **hindsight bias** – that is, the tendency to exaggerate the probability that an event would occur by people who know that it in fact did occur – plays other tricks with our minds. Victims of the hindsight bias may never learn that past predictions were no good, because they misremember what they in fact predicted and see no need to be less confident in the future. Finally, people are very good at explaining away false predictions, e.g., by arguing that they were *almost* right, or that any failures were due to inherently unpredictable factors.

The overconfidence phenomenon receives indirect support from research on **competence**. For example, many studies suggest that people overestimate their competence in various practical tasks. The vast majority of drivers – in some studies, more than 90 percent – say that they are better than the median

The heuristics-and-biases program

According to the heuristics-and-biases program, we form judgments by using **heuristics** – functional but imperfect rules of thumb or mental shortcuts that help us form opinions and make decisions quickly. Here are four prominent heuristics:

- The **anchoring-and-adjustment** heuristic instructs you to pick an initial estimate (anchor) and adjust the initial estimate up or down (as you see fit) in order to come up with a final answer. Therefore, estimates will (imperfectly) track the anchors – even when irrelevant or uninformative.
- The **representativeness** heuristic tells you to estimate the probability that some outcome was the result of a given process by reference to the degree to which the outcome is representative of that process: the more representative the outcome, the higher the probability.
- The **availability heuristic** makes you assess the probability that some event will occur based on the ease with which the event comes to mind: the easier it comes to mind, the higher the probability.
- The **affect heuristic** gets you to assign probabilities to consequences based on how you feel about the thing they would be consequences of: the better you feel about it, the higher the probability of good consequences and the lower the probability of bad.

The heuristics-and-biases program does not say that people are dumb: to the contrary, it says that following heuristics is a largely functional way to make speedy decisions when it counts. Because the heuristics are imperfect, however, they can lead to **bias**, that is, systematic and predictable error.

Kahneman has proposed that the operation of heuristics can be understood in terms of **substitution**. When faced with a question we are unable to address directly, we sometimes replace it by an easier question and answer that instead. Rather than addressing the question "How likely is this airplane to crash?" we substitute the question "How easily can I imagine this airplane crashing?" Substitution allows us to come up with quick answers. But because the question we are actually answering is different from the original one, the answer too might be different: while airplane crashes are unlikely, mental representations of airplane crashes are easy to conjure. Thus, we come to exhibit availability bias and exaggerate the probability of a crash. (Similar stories can be told about the other heuristics.) The tricky thing is that we are unaware of the substitution and therefore of the difference between the answer we sought and the one we produced.

driver, which is a statistical impossibility. In one fascinating study, undergraduates whose test scores in grammar and logic put them in the bottom 25 percent of a group of peers, on the mean estimated that they were well above average. Even more surprising, perhaps, when participants received more information about their relative performance in the tests (by being asked to grade those of other participants), the strongest students became more calibrated, while the weakest students, if anything, became *less* calibrated. The results suggest the least competent are at a double disadvantage, in that their incompetence causes both poor performance and an inability to recognize their performance as poor: the *cognitive* skill required to perform a difficult mental task may well be tightly tied up with the *metacognitive* skill to assess the quality of our performance. This **Dunning–Kruger effect** may or may have been what comedian Ricky Gervais was getting at when he said: "When you are dead, you do not know you are dead. It is only painful for others. The same applies when you are stupid."

Exercise 5.39 Inevitability People think many things are inevitable. If you search for the expression "it was inevitable that" on Google News, you may get more than ten thousands hits. Which bias is reflected in the use of that expression?

Exercise 5.40 Adam Smith, once more What sort of phenomenon might Adam Smith have had in mind when he talked about the "over-weening conceit which the greater part of men have of their own abilities"?

5.8 Discussion

This chapter has explored phenomena that appear inconsistent with the theory of probabilistic judgment that we learned about in Chapter 4. While the theory of probability was never designed to capture the precise cognitive processes people use when forming judgments, there appears to be a wide range of circumstances under which people's intuitive probability judgments differ substantially, systematically, and predictably from the demands of the theory. As the examples have shown, differences can be costly. The phenomena are typically construed as undercutting the adequacy of probability theory as a descriptive theory. Yet some of these phenomena can also be invoked when challenging the correctness of probability theory as a normative standard. The fact that living up to the theory is so demanding – surely, part of the explanation for why people fail to live up to it in practice – is sometimes thought to undercut its normative correctness.

We have also discussed some of the theoretical tools used by behavioral economists to capture the manner in which people actually make judgments. Thus, we explored further aspects of the heuristics-and-biases program which is one prominent effort to develop a descriptively adequate theory of probabilistic judgment. Because heuristics are largely functional, it would be a mistake to try to eliminate reliance on them altogether. But heuristics can lead

cinating study, under-
ut them in the bottom
ed that they were well
n participants received
ne tests (by being asked
lents became more cali-
me *less* calibrated. The
advantage, in that their
n inability to recognize
d to perform a difficult
acognitive skill to assess
ger **effect** may or may
at when he said: "When
y painful for others. The

gs are inevitable. If you
Google News, you may
lected in the use of that

t of phenomenon might
t the "over-weening con-
abilities"?

inconsistent with the the-
ut in Chapter 4. While the
e the precise cognitive pro-
ppears to be a wide range
robability judgments differ
the demands of the theory.
costly. The phenomena are
of probability theory as a
can also be invoked when
s a normative standard. The
surely, part of the explana-
e – is sometimes thought to

al tools used by behavioral
le actually make judgments.
ristics-and-biases program,
riptively adequate theory of
gely functional, it would be a
ether. But heuristics can lead

us astray, and an awareness of the conditions under which this may happen might reduce the likelihood that they do. As the marketing applications make particularly clear, knowledge of behavioral economics in general, and heuristics and biases in particular, permits us to influence other people's judgment, for good or for evil. But knowledge of behavioral economics also permits us to anticipate and to resist other people's efforts to influence our judgment.

Exercise 5.41 Misguided criticism Some critics of the heuristics-and-biases program attack it for saying that human beings are irredeemably stupid. Thus, "the heuristics-and-biases view of human irrationality would lead us to believe that humans are hopelessly lost in the face of real-world complexity, given their supposed inability to reason according to the canon of classical rationality, even in simple laboratory experiments." Explain where these critics go wrong.

Needless to say, the chapter does not offer a complete list of phenomena that are inconsistent with standard theory or of the theoretical constructs designed to explain them. One important heuristic that has not come up yet is the **affect heuristic**. This heuristic tells you to assign probabilities to consequences based on how you feel about the thing they would be consequences of: the better you feel about it, the higher the probability you assign to good consequences and the lower the probability you assign to bad. As a result, those who love guns will think of scenarios in which guns are beneficial as relatively more likely, and scenarios in which guns are harmful as relatively less likely, than those who hate guns. Like other heuristics, the affect heuristic is often functional: when you see something you really do not like the look of in the back of your fridge, you will think the probability that you will get sick from eating it as relatively high – which is likely the proper response. But the affect heuristic can also lead you astray, by making beliefs about benefits and risks reflect feelings rather than the other way around.

The next part of the book will explore how the theory of probability can be incorporated into a theory of rational decision, and explores its strengths and weaknesses.

ADDITIONAL EXERCISES

Exercise 5.42 Probability matching Imagine that your friend Anne has a coin that has a 2/3 probability of coming up heads (H) and a 1/3 probability of coming up tails (T). She intends to flip it three times and give you a dollar for every time you correctly predicted how the coin would come up. Would you be more likely to win if you predicted that the coin will come up HHH or HHT?

If your prediction was HHT (or HTH or THH) you might have engaged in **probability matching**: choosing frequencies so as to match the

ADDITIONAL EXERCISES cont.

probabilities of the relevant events. Probability matching might result from the use of the representativeness heuristic, since an outcome like HHT (or HTH or THH) might seem so much more representative for the random process than HHH or TTT. As the exercise shows, probability matching is a suboptimal strategy leading to bias.

Exercise 5.43 Gender discrimination, cont. In Exercise 4.51 on page 95, we computed the probability that an editorial board of 20 members is all male by chance alone. If the answer strikes a person as low, what fallacy may he or she have committed?

Exercise 5.44 IVF In vitro fertilization (IVF) is a procedure by which egg cells are fertilized by sperm outside the womb. Let us assume that any time the procedure is performed the probability of success (meaning a live birth) is approximately 20 percent. Let us also assume, though this is unlikely to be true, that the probabilities of success at separate trials are independent. Imagine, first, that a woman has the procedure done twice.
(a) What is the probability that she will have *exactly* two live births?
(b) What is the probability that she will have *no* live births?
(c) What is the probability that she will have *at least* one live birth?
Imagine, next, that another woman has the procedure done five times.
(d) What is the probability that she will have at least one live birth?

Exercise 5.45 Mandatory drug testing In July 2011, the state of Florida started testing all welfare recipients for the use of illegal drugs. Statistics suggest that some 8 percent of adult Floridians use illegal drugs; let us assume that this is true for welfare recipients as well. Imagine that the drug test is 90 percent accurate, meaning that it gives the correct response in nine cases out of ten.
(a) What is the probability that a randomly selected Floridian welfare recipient uses illegal drugs and has a positive test?
(b) What is the probability that a randomly selected Floridian welfare recipient does not use illegal drugs but nevertheless has a positive test?
(c) What is the probability that a randomly selected Floridian welfare recipient has a positive test?
(d) Given that a randomly selected Floridian welfare recipient has a positive test, what is the probability the he or she uses illegal drugs?
(e) If a Florida voter favors the law because he thinks the answer to (d) is in the neighborhood of 90 percent, what fallacy might he be committing?

Exercise 5.46 CIA Intelligence services are deeply interested in how people think, both when they think correctly and when they think incorrectly. The following exercise is borrowed from the book *Psychology of Intelligence Analysis*, published by the US Central Intelligence Agency (CIA).

atching might result
ince an outcome like
representative for the
se shows, probability

ercise 4.51 on page 95,
d of 20 members is all
n as low, what fallacy

a procedure by which
b. Let us assume that
ty of success (meaning
so assume, though this
ss at separate trials are
procedure done twice.
tly two live births?
ive births?
east one live birth?
dure done five times.
east one live birth?

011, the state of Florida
f illegal drugs. Statistics
s use illegal drugs; let
s as well. Imagine that
hat it gives the correct

lected Floridian welfare
e test?
lected Floridian welfare
ertheless has a positive

lected Floridian welfare

welfare recipient has a
r she uses illegal drugs?
he thinks the answer to
what fallacy might he be

ply interested in how peo-
hen they think incorrectly.
k *Psychology of Intelligence*
ce Agency (CIA).

ADDITIONAL EXERCISES cont.

During the Vietnam War, a fighter plane made a non-fatal strafing attack on a US aerial reconnaissance mission at twilight. Both Cambodian and Vietnamese jets operate in the area. You know the following facts: (a) Specific case information: The US pilot identified the fighter as Cambodian. The pilot's aircraft recognition capabilities were tested under appropriate visibility and flight conditions. When presented with a sample of fighters (half with Vietnamese markings and half with Cambodian) the pilot made correct identifications 80 percent of the time and erred 20 percent of the time. (b) Base rate data: 85 percent of the jet fighters in that area are Vietnamese; 15 percent are Cambodian.

Question: What is the probability that the fighter was Cambodian rather than Vietnamese?

Exercise 5.47 Juan Williams In October 2010, National Public Radio (NPR) fired commentator Juan Williams after he made the following remark on Fox News: "When I get on a plane … if I see people who are in Muslim garb and I think, you know, they're identifying themselves first and foremost as Muslims, I get worried. I get nervous." Here, I will not comment on the wisdom of firing somebody for expressing such a sentiment or of expressing it in the first place. But we can discuss the *rationality* of the sentiment.

(a) In the United States, there are roughly 300 million people, of whom about 2 million are Muslims. Let us assume that at any one time there are ten terrorists able and willing to strike an airliner, and that as many as nine out of ten terrorists are Muslims. Under these assumptions, what is the probability that a randomly selected Muslim is a terrorist able and willing to strike an airliner?

(b) Use the notion of availability bias to explain how Juan Williams might overestimate the probability that a randomly selected Muslim is a terrorist able and willing to strike an airliner.

Exercise 5.48 Theories, theories Complete this sentence: "If all your observations support your scientific theories or political views, you are (probably) suffering…"

Exercise 5.49 Matthew Which heuristic is embodied in this line from Matthew 7:17–18: "So every good tree bears good fruit, but the bad tree bears bad fruit. A good tree cannot produce bad fruit, nor can a bad tree produce good fruit."

Exercise 5.50 Genetically modified organisms (GMOs) A person opposed to GMOs reads a compelling text about the benefits of such organisms and comes to quite like the thought of them. When asked about the risks, he says he has changed his mind and decided that not only are the benefits great, but the risks are trivial too – although he has acquired no new information about the risks. What heuristic might have been in play here?

ADDITIONAL EXERCISES cont.

Exercise 5.51 Schumpeter The Austrian economist Joseph Schumpeter claimed that he had set himself three goals in life: to be the greatest economist in the world, the best horseman in all of Austria, and the greatest lover in all of Vienna. He confessed that he had only reached two of the three goals. Suppose that he was wrong and had, in fact, reached none of them. Use each of the following ideas to explain how he might be so wrong: (**a**) confirmation bias, (**b**) availability bias, (**c**) overconfidence, (**d**) conjunction fallacy.

Exercise 5.52 Match each of the vignettes below with one of the following phenomena: *availability bias, base-rate neglect, confirmation bias, conjunction fallacy, disjunction fallacy, hindsight bias*, and *overconfidence*. If in doubt, pick the best fit.

(**a**) Al has always been convinced that people of Roma descent are prone to thievery. In fact, several of his co-workers have a Roma background. But he knows his co-workers are not thieves, and he does not think twice about it. One day during happy hour, however, an old acquaintance shares a story about two people "who looked like gypsies" stealing goods from a grocery store. "I knew it!" Al says to himself.

(**b**) Beth's car is falling apart. Her friends, who know these things, tell her that the car has a 10 percent probability of breaking down every mile. Beth really wants to go see a friend who lives about ten miles away, though. She ponders the significance of driving a car that has a 10 percent probability of breaking down each mile, but figures that the probability of the car breaking down during the trip cannot be much higher than, say, 15 percent. She is shocked when her car breaks down half way there.

(**c**) Cecile is so terrified of violent crime that she rarely leaves her house, even though she lives in a safe neighborhood. She is out of shape, suffers from hypertension, and would be much happier if she went for a walk every so often. However, as soon as she considers going for a walk, images of what might happen to innocent people quietly strolling down the sidewalk come to mind, and she is sure something horrible is going to happen to her too. As a result, she goes back to watching reruns of *Law and Order*.

(**d**) David, who has never left the country, somehow manages to get tested for malaria. The test comes back positive. David has never been so depressed. Convinced that he is mortally sick, he starts to draft his will.

(**e**) Because she has trouble getting up in the morning, Elizabeth often drives too fast on her way to school. After getting a speeding ticket on Monday last week, she religiously followed the law all week. Only this week is she starting to drive faster again.

ADDITIONAL EXERCISES cont.

(f) Fizzy does not think the US will want to start another front in the war on terror, so she believes that it is quite unlikely that the US will bomb Iranian nuclear facilities. She does, however, think that it is quite likely that the US will withdraw all troops from Afghanistan. When asked what she thinks about the possibility that the US will bomb Iranian nuclear facilities and withdraw all troops from Afghanistan, she thinks that it is more likely than the possibility that the US will bomb Iranian nuclear facilities.

(g) Georgina has trouble imagining an existence without Apple computers and iPads. Thus, she thinks it was inevitable that a man like Steve Jobs would appear and design such things.

(h) Harry lost all his luggage the last time he checked it. He is never going to check his luggage again, even if it means having unpleasant arguments with flight attendants.

Problem 5.53 *Drawing on your own experience, make up stories like those in Exercise 5.52 to illustrate the various ideas that you have read about in this chapter.*

FURTHER READING

A comprehensive introduction to heuristics and biases in judgment is Hastie and Dawes (2010). The gambler's fallacy and related mistakes are discussed in Tversky and Kahneman (1971). The conjunction and disjunction fallacies are explored in Tversky and Kahneman (1983) and in Tversky and Shafir (1992). The planning fallacy and the Sydney Opera House are discussed in Buehler et al. (1994, p. 366); Kahneman (2011) reports the data about rail projects. Base-rate neglect is examined in Bar-Hillel (1980). The *USA Today* story is Frank (2007). An extensive review of confirmation bias is Nickerson (1998); the quotation from *In Search of Lost Time* is Proust (2002 [1925], p. 402) and the study of confirmation bias in the context of the death penalty is Lord et al. (1979). Popper (2002 [1963], p. 46) describes meeting Adler. Availability is discussed alongside anchoring and adjustment and representativeness in Tversky and Kahneman (1974). The section on overconfidence draws on Angner (2006), which argues that economists acting as experts in matters of public policy exhibit significant overconfidence even within their domain of expertise. Neil Armstrong is cited in O'Brien (2004). The study by Fischhoff and colleagues is Fischhoff et al. (1977), the study involving a real patient is Oskamp (1982), and the study on competence is Kruger and Dunning (1999). Kahneman (2011, Chapter 9) discusses substitution and Smith (1976 [1776], p. 120) our "over-weening conceit." The critics of the heuristics-and-biases program are Gigerenzer and Goldstein (1996); the affect heuristic appears in Finucane et al. (2000). The CIA's intelligence analysis example comes from Heuer (1999, pp. 157–8) and the Juan Williams affair is described in Farhi (2010).

PART 3

Choice under Risk and Uncertainty

6 Rational Choice under Risk and Uncertainty

6.1 Introduction
6.2 Uncertainty
6.3 Expected value
6.4 Expected utility
6.5 Attitudes toward risk
6.6 Discussion

7 Decision-Making under Risk and Uncertainty

7.1 Introduction
7.2 Framing effects in decision-making under risk
7.3 Bundling and mental accounting
7.4 The Allais problem and the sure-thing principle
7.5 The Ellsberg problem and ambiguity aversion
7.6 Probability weighting
7.7 Discussion