CRITICAL THINKING, SCIENCE, AND PSEUDOSCIENCE

WHY WE CAN'T TRUST OUR BRAINS

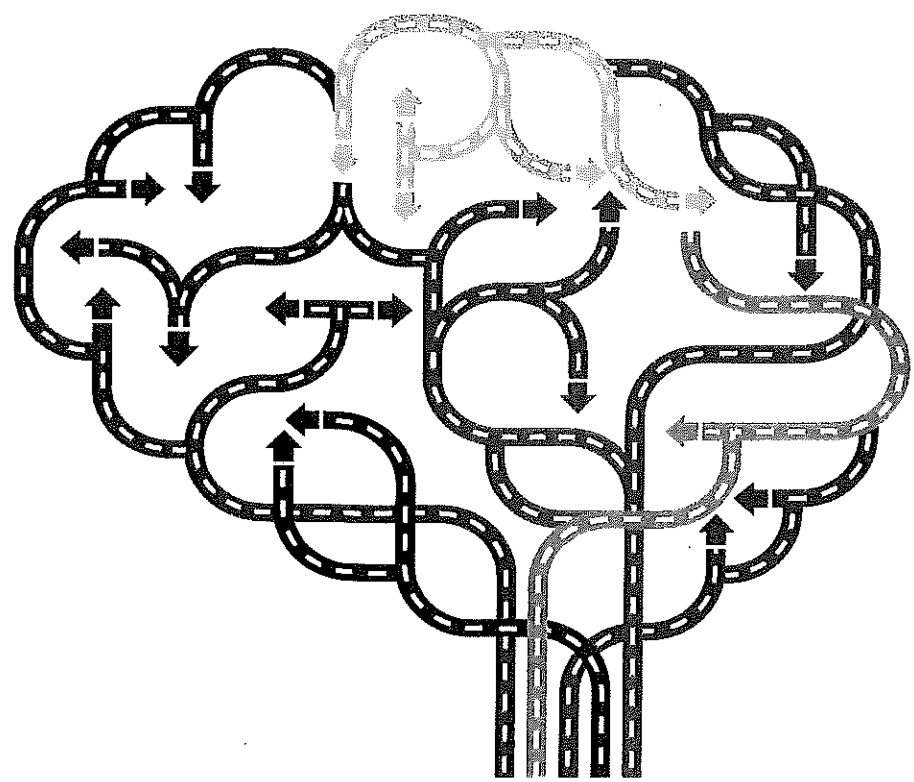

CALEB W. LACK & JACQUES ROUSSEAU

SPRINGER PUBLISHING COMPANY

CHAPTER 2

WHAT IS SCIENCE?

HYPE VERSUS HYPOTHESES

We're bombarded with claims made in the language of science every day of our lives. Newspapers tell us what's safe or unsafe to eat, popular magazines routinely feature colorful pictures of our brains, all in an effort to explain—or gesture at explanations of—claims regarding things like what makes us happy or what we're addicted to (or could become addicted to, if we're not careful). But a general problem with these media representations of scientific activity is that they often make things appear much simpler than they actually are. Science doesn't advance one press release at a time, despite appearances to the contrary (Goldacre, 2014)—you might never hear of significant breakthroughs, and what you hear about might in turn be misrepresented or deeply contested. Despite these complexities, there are some simple principles underlying scientific endeavors, as outlined in this chapter. However, the actual work of doing the science can be slow and uncertain, and the conclusions reached can be far more tentative than the headlines might lead you to believe.

Perhaps most important, the public typically won't get to hear of all the uncertainties and qualifications generated by scientific inquiry, and it is unlikely that the public will be aware of the fact that a significant proportion of published research findings—perhaps even the majority of them—are in fact false (Ioannidis, 2005). This is primarily because these complexities don't provide as compelling a narrative. Who wants to read about how "scientists are still uncertain about the strength of the causal connection between total cholesterol in the blood and heart disease," versus "eat as much fat as you like! New study shows that cholesterol is not a risk factor in heart disease!" The more hyperbolic and dogmatic a claim, the more it seems to dictate what appears on magazine covers and in dinner-table conversations. But this phenomenon, which we might describe as a "sexying up" (via simplification), in which titillation is what matters most, presents the danger of having us mistake the controversial or particularly exciting cases as being representative of how science *typically* works, and of the scientific method in general. The reality is that our knowledge of the world is always contingent

on what we know—and can know—at any given time. By contrast, popular media debates on scientific matters can make it appear that we're constantly replacing one dogmatic conclusion with another, in light of some new piece of data or study.

Instead, we need to remember that it's more accurate to regard new evidence as *tipping the scales* one way or the other, with a conclusion becoming more or less likely as new evidence emerges. It's usually a mistake to speak of being certain of some conclusion, both because certainty is very rarely (if ever) available to us, and second because it encourages an attitude of dogmatism, rather than the skeptical and curious attitude that is the hallmark of good scientific practice (Popper, 2002). Of course some conclusions are so well established that they might as well be regarded as certain, but the fact remains that contrary evidence will still tip those scales, even if to a trivial degree. Openness to correction is a key virtue that distinguishes science from pseudoscience (see Chapter 3 for more on this), so it's important that our language reflects that we are able to change our minds.

None of the preceding is an objection to making science compelling and interesting, even sexy and cool, because it is! Part of the purpose of this book is to inspire people, especially young people, to get involved in scientific work, as well as to help everyone understand scientific conclusions and live evidence-driven lives. But making science "cool" should not come at the cost of dumbing it down, or making it appear simpler than it is. We shouldn't be encouraged to think that getting clear positive results from a study is the norm, for example. Much scientific research will achieve few concrete results in terms of telling us something new or confirming a counterintuitive hypothesis, but could nevertheless help us to rule out various other hypotheses as *not* worth pursuing. This would still be a significant contribution to the body of knowledge, in that it would allow for future researchers to use their time more productively rather than chasing what have now been established as dead ends. But although the identification and elimination of these dead ends does constitute "new" knowledge, it is knowledge that will be of more use to future scientists than it would be for the lay public, and might not achieve the same prominence in the public media.

The general point is that much of science involves a slow, long slog and is conducted by people outside of the limelight, who often shun the limelight because they care more about the work than about becoming celebrities. The difficulties and challenges inherent in some fields of inquiry is exactly what motivates some scientists to work so hard to find the answers to questions posed in those fields. As examples like Carl Sagan or Neil deGrasse Tyson show, there's certainly no necessary tension between celebrity and science (Fahy, 2015), it's just that they tend to not be strongly associated with each other, thanks to the complexity of the work involved.

But even in cases in which the work is difficult—even perhaps completely obscure to nonscientists—there are nevertheless some general principles underlying scientific inquiry that we can all benefit from understanding. This chapter seeks to summarize those principles, with the aim of facilitating a broad understanding of why some apparently scientific claims, and people purporting to be scientists, are worthy of your attention and trust, and others less so.

SCIENCE: NOT ALWAYS "COMMON SENSE"

Our beliefs and attitudes are informed by the experiences we have, but it is easy to forget that those experiences might not be representative of what's typical for humans more generally. What this means is that a very small set of data (we are only one data point among billions) can be used to develop conclusions that perhaps don't accurately reflect the beliefs and attitudes of others. Now, of course it's true that our circumstances might be atypical, and therefore it's often true (and entirely reasonable) that our experiences and expectations *should* be different from the norm, or from those of other people we encounter.

But precisely because of how powerful our own experiences are to us—being the only ones we have direct access to—we can also be led to overstate the value of personal experience and forget that we *shouldn't* expect the patterns of our lives to be typical for others—or even reliable when it comes to understanding our own lives (Mlodinow, 2009). To take a simple example: If you have great success in losing weight while eating or not eating some sort of food (perhaps bread), you might think this would be true for everyone. This conclusion would be unduly hasty, for at least two clear reasons.

First, although human biology is of course similar, there is nevertheless enough variation among us that generalizing from your (single) case to a population of billions is exceedingly risky. Second, your anecdotal experience isn't usually as reliable as you might think it is. As mentioned in Chapter 1, our tendency is to overemphasize evidence that confirms something we already believe to be the case, while underemphasizing contradictory evidence: We'd typically be more likely to give credit to the diet while perhaps forgetting (or underplaying) the extent to which the dieter also simply ate less, or exercised more, or just contracted a tapeworm. What we perceive as the most *proximate* (closest in relationship to what we're experiencing; or immediately apparent as related) cause of something we experience—in this case weight loss—might be obscuring our recognition that something else better explains that experience.

The general point is that we tell ourselves stories to make sense of our lives and the world. And this general point reveals a key difference between how science operates (or should operate) and how common sense operates.

In science, the evidence is the evidence, and leads us to a particular conclusion (even if that conclusion is "we don't know"). Common sense, in contrast, tends to lead us away from uncertainty, and toward beliefs that fit in with other things we already believe to be true. We're programmed to fit information into structures and patterns (Shermer, 2008), and although doing so can often be useful (consider medical diagnoses as example, where understanding symptoms as representing some broader condition can help the physician to identify what ails you), this pattern-seeking behavior (or *patternicity*) can also lead us to errors such as the "gambler's fallacy," in which we might keep laying money onto the poker table because we think we see some fortuitous pattern in the cards that have been dealt in previous hands.

A scientific outlook is not personally invested in certain outcomes or subjective perceptions of data—it will only take your anecdotal experience for what it's worth, and what it's worth is not much, given that it's merely one data point among billions. It's also not a reliable data point, given that the observations themselves can't be externally validated for accuracy. In the language of science, our (perfectly natural) habit of deriving conclusions from personal experience amounts to an *uncontrolled* experiment, in which the variables involved are difficult to impossible to identify or quantify (Chabris, 2011).

Think further about some of the rules of thumb that inform our lives, many of which identify principles or strategies for decision making that could be described as "common sense" (Lilienfeld, Lynn, Namy, & Woolf, 2013). To express the idea of being prudent, we have "Look before you leap." But we also have "He who hesitates is lost." The contradiction between these is clear, and leaving aside the fact that some people are naturally more risk averse than others (and will therefore tend to look before they leap), they lead us to recognize that these general rules always need interpretation and contextually aware application—simply applying a rule without regard for context is unlikely to result in an optimal outcome. There are dozens of similar examples, like "You're never too old to learn" versus "You can't teach an old dog new tricks," or "Birds of a feather flock together" versus "Opposites attract." Part of what the scientific outlook gives us is a way to help us navigate confusions or contradictions like these in a principled and fairly reliable manner by giving us the tools to explain and predict the relationship between causes and effects.

We say "fairly reliable" because we typically cannot guarantee the truth of a conclusion, the best we can do is to reach conclusions as responsibly as possible, and accept that following a sound methodology vastly increases the likelihood of those conclusions being true. It's no discredit to science that it gets things wrong on occasion—the point is, there's no *better* way to reach conclusions, and the scientific method gets things right more often than it gets things wrong.

WHAT IS SCIENCE?

This simple question has a relatively complex answer. An entire subfield of philosophy, the philosophy of science, is in fact dedicated to trying to answer that question. This text aims at a general explanation of issues in critical reasoning and scientific inquiry, so in contrast to specialist texts in the philosophy of science, we will not be assessing the relative merits of realism, antirealism, empiricism, rationalism, and the like.[1] Suffice it to say that your authors adopt a viewpoint that is broadly *realist* or *naturalist*, in that we believe that scientific theories typically should—and in fact often do—represent what is actually real. Broadly, you can contrast this perspective with antirealism, which holds that scientific theories are not dependent on, or obliged to, represent metaphysical reality, and that they can work to describe observations regarding things like electrons or black holes without necessarily being accurate.

We also adopt an *empirical* stance, which holds that observations and the generalizations we make on the basis of those observations are central to scientific inquiry. Because scientific theories are empirical, this means that they are responsive to new observations, even when those observations conflict with existing hypotheses or theories (terms that we'll define in a moment). One cannot dogmatically hold on to theories, especially when they do not accommodate or explain new observations. As will be explained later, this means that theories are always open to being *falsified*.

Another way of putting this point is to highlight the importance of *fallibilism*, which is the knowledge, as a scientist, that new data might still come to light which puts all of your current convictions into doubt. We therefore embrace the possibility of error, and reject a dogmatic attitude toward what we know or think we know (which, as we saw in Chapter 1, may or may not be justified). In fact, the words "know" or "knowledge" are more accurately understood as referring to things that are true to the best of our (current) knowledge, while acknowledging that we might have to change our minds about what is true at some future point. The notion of fallibilism reminds scientists that certainty isn't available to us, and also it reminds us that it's *good to be wrong*, because "nothing obstructs access to the truth like a belief in absolute truthfulness" (Deutsch, 2013). Given those preliminaries, how might we define *science*? A definition that we find particularly useful is:

> A set of methods designed to describe and interpret observed or inferred phenomena, past or present, and aimed at building a testable body of knowledge open to rejection or confirmation (Shermer, 2002, p. 145).

---

[1] Those interested in an overview of these distinctions and debates could usefully consult a text such as Okasha's *Philosophy of Science: A Very Short Introduction* (2002).

Shermer's definition captures the essence of what science or the scientific method entails. Science offers us a toolkit, and it's then up to us to use those tools correctly and responsibly. Another feature of this definition worth highlighting is the fact that scientific hypotheses need to be *testable*. As we'll see later in this chapter, if there's no way to test a claim (even if the manner of testing is currently hypothetical), that claim currently falls outside the scope of science, and always might.

This principle of testability is vital to the health of science, in that it allows for science to be self-correcting. We can only discover our mistakes if there's a way of discovering them, and that means that claims must be testable. On a personal level, thinking scientifically is useful for a similar reason: We're very good at fooling ourselves, and one way to minimize fooling yourself is to test what you believe against the facts. And, if there's no way of testing one or more of your beliefs, then those beliefs might be false, and you would have no way of knowing.

A concern for being able to test hypotheses points to a related issue, namely, *bias in methodology*. Given that we know, in advance of even attempting to test a hypothesis, that we are prone to confirmation and other biases (as discussed extensively in Chapter 5), it is incumbent on us to design our experiments and other scientific activities in ways that attempt to ensure that we correct for these biases. The fact that we know how important it is that we allow ourselves to be wrong doesn't itself offer any guarantees that in the moment, we'll happily accept error. It's not comfortable to discover that we hold false beliefs, and this makes it vital that our reasoning, and our experiments, are not set up in ways that minimize the chances we'll be proved wrong.

Good science—and good scientists—tries to exemplify the virtue of dispassionate truth seeking. This does not mean the scientist does not care about his work, but something else entirely. Being a good scientist requires integrity, in the form of being honest and scrupulous with data—by not misrepresenting what it says, and by not ignoring data that is inconvenient to one's hypothesis (Goldacre, 2012). Good science requires a kind of disinterestedness, in that you would try your best to not be influenced by personal commitments or financial interests, such as pressures from those who are funding your research.

Most important, perhaps, is that good science requires that we're willing to be wrong and even that we seek out opportunities to discover where we are wrong. In doing so, beliefs that are tested and *survive* that testing can be regarded as more reliable than they were before surviving the test. Because we're so good at fooling ourselves, this also requires that we collaborate with each other as scientists—other people might spot issues that you are unable to see, perhaps because of bias, or perhaps because you've just spent too much time looking at things in one way, and a fresh pair of eyes can see

them in a different way. Science, in other words, is usually collaborative or communal (Adams, 2013). As a scientist, the point of publishing scientific findings in peer-reviewed journals is that the community can benefit from your findings, but also so that they can help you test them, and discover whether they can be replicated or not.

BUILDING BLOCKS OF THE SCIENTIFIC METHOD

Words like "hypothesis," "law," "theory," and "fact" are encountered fairly frequently in everyday conversation, even by nonscientists. But, these usages might sometimes be misleading or idiosyncratic when compared to their technical definitions. As such, let's start by briefly defining them as used when speaking about the scientific method.

A *hypothesis* is a testable statement that accounts for a set of observations. The task of testing it is made easier by stating the hypothesis in clear language. Also, the hypothesis should express something unambiguous enough that it is clear how you might go about testing it in order to confirm or disconfirm it. Consider this statement: "Childhood obesity is linked to the number of sugary drinks consumed daily." The observation being described here is, of course, childhood obesity, and we're given a suggestion as to one of its possible causes. Crucially, though, we're also given a statement that we can test in various ways. You might test it in an observational way, by comparing childhood obesity rates in populations that drink fewer sugary drinks to populations that drink more sugary drinks; or you might test it in an experimental way by designing a study in which you control as many other relevant variables as possible, to highlight the role that sugar consumption could play in childhood obesity.

By contrast, consider the central hypothesis of a book like Rhonda Byrne's *The Secret*, summarized as the "Law of Attraction." This "law" tells us (in life coach Bob Proctor's formulation): "Whatever is happening in your mind, you will attract it in your life" (De Fretes, n.d.). Is that hypothesis testable? No, because there is no way in which *it can fail to be true*. As hard as I (JR) might *think* I have massive wealth and fame "happening in my mind" right now, Byrne and her ilk can simply tell me that I'm not thinking hard enough, or not visualizing that wealth and fame in a sufficiently productive way. And then, of course, whenever good things might happen to me, they can still claim the credit because I *was* wishing it to be so. Books and movements like *The Secret* get to count their successes, and ignore their failures—the "law of attraction" is true in all instances because it isn't responsive to the *full* body of evidence (or any evidence at all) in the way that we would want a scientific hypothesis to be (Wheen, 2004).

Now, once a hypothesis is well enough established, we refer to it scientifically as a *law*. The "law of attraction" mentioned previously is an

example of someone spuriously using the term "law" to give the illusion of credence to mumbo-jumbo, because, as we have seen, the hypothesis in question is not only not well established, but never could be. By contrast, the hypothesis that "things dropped from your hand will fall to the floor" is sufficiently well established that we feel entitled to treat it as a law, and in fact do so in calling it the "law of gravity."

Next, a scientific *theory* is something that tends to be broader in scope than hypotheses and laws—it is a set of well-tested and well-supported hypotheses and laws, which in combination explain many more events, and make many more predictions, than the hypotheses or laws on their own. Here, for example, we would all be familiar with the "theory of evolution," which contains many hypotheses on one related theme, namely, the idea that natural selection guides the development of living organisms (Coyne, 2010). This is very dissimilar to how many people use "theory" in everyday terms, which generally just means "an idea that I have" (such as "Well, I have a theory as to why Nancy and John broke off their engagement").

Finally, and returning to some of the discussion around what "knowledge" means from Chapter 1, when we speak of scientific *facts* we're not referring to things that we *know* to be true with 100% certainty but instead referring to conclusions that are confirmed to such an extent that it would be reasonable to offer provisional agreement, and unreasonable to deny agreement. As suggested earlier in this chapter, we can—and should—change our minds when new evidence comes along that is stronger than competing evidence we've had to date. This doesn't mean that the facts themselves change—it instead means that we were simply wrong about what the facts were, and have had the opportunity to learn about that error.

SCIENTIFIC REASONING

We use scientific reasoning every day of our lives, whether or not we're doing so deliberately. When you hear the weather report calling for rain, but look out of the window to confirm whether or not it is raining, you're using scientific reasoning in the very basic sense that you're *testing a claim* against the evidence. Before we talk about the process by which we do so, let's look at a (slightly) more complicated example, which demonstrates certain key steps in the process of scientific reasoning. Those steps can be stipulated as:

1. Identifying a problem, or observation in need of explanation
2. Gathering information about the problem or observation
3. Formulating explanations (hypotheses) regarding the problem or observation
4. Conducting tests or experiments to see which, if any, of the hypotheses provide a resolution for the problem or explain the observation

5. Deriving a conclusion that accurately captures the resolution or observation (and, ideally, gives us guidance in terms of this or relevantly similar situations in the future)

An Example

Let's say that you're trying to charge your mobile phone, but the screen tells you that the phone is not charging, even though you confirm that you've plugged it in (and note that in confirming this, you've *already* started reasoning scientifically, by testing for the most obvious explanation). So here we're at Step 1: We have a problem or situation in need of explanation.

Step 2, information gathering, already started when you checked that you had plugged the phone in correctly. Step 2 might also involve elements like thinking back to when you last charged your phone. Everything seemed to work fine on the previous evening, so you know that something has changed, or gone wrong since then—now you just need to figure out what that is.

So (Step 3), what could explain the fact that it's not charging? Well, the power could be out in your neighborhood (that's one hypothesis), or just in your house (a competing hypothesis). Perhaps the problem is even more localized, and the power outlet you're using is faulty. The phone's charging unit or cable could also be a problem, or (and you don't really want to think about this one!), your $500 phone might be faulty.

You already know what to do in Step 4 because, as I say, we engage in this sort of reasoning process all the time, even if often unconsciously. You might be able to test the "power out in the neighborhood" and "power out in your house" hypothesis simultaneously, through realizing that you can hear the television playing in the next room. Then, you might rule out the power outlet being faulty, because your laptop is still receiving power despite being plugged into the same outlet. So now (in our stripped-down example), you're left with two hypotheses, and you'd probably next try testing a different cable—if the phone charges, you know the cable was faulty (which gives you Step 5).

If the phone still doesn't charge, you might think that this tells you the phone is faulty (also an answer to Step 5), but not quite yet—the second cable might also be faulty. So a possible next step would be to test both of these cables with a different phone to verify that they are able to charge a different device. It might just (phew!) be that all the cables you have available at the moment are faulty, and that there's nothing wrong with your phone at all.

This example uses the building blocks of hypotheses and evidence to reach conclusions, and is a completely familiar scenario to us because we do it or something close to it every day. But these building blocks can fit

together in various ways as we go about the business of scientific thinking. The two most common ways of generalizing the process of scientific reasoning and how it proceeds are the methods of *induction* and *deduction*, both of which merit further discussion.

INDUCTION AND DEDUCTION

Our past experiences often form the basis of our expectations in the future. If you've had a couple of good meals at a particular restaurant, you'd start recommending it to friends because you expect that both you, and they, are likely to have good meals there in future. In other words, specific observations in the past are used to derive a general principle, which involves events in the future. This, in short, is *induction.*

The restaurant example, drawn from everyday life, is generated by a process of induction and works in exactly the same way in scientific reasoning—it amounts to the derivation of a general principle, even a law, from specific observations (Holland, Holyoak, Nisbett, & Thagard, 1989). Notice that for these observations to generate reliable principles that can accurately predict events in the future, we need to make an important assumption, namely, that nature is (to a significant extent, at least) *uniform.* On that assumption, we are identifying regularities and using those as fairly stable or consistent variables for understanding the world around us and for predicting future events.

How are these general principles justified? First, the number of observation statements must be fairly large for us to be able to have confidence in them not being outliers or accidents. Second, the observations must be repeated under a wide variety of conditions. To go back to our everyday example, the more meals *you* eat at a particular restaurant, the more strongly justified *your* conclusion is that it's a good restaurant. And second, if a number of different people (with different tastes and standards than you) have eaten at that restaurant at various times (in season, out of season, for lunch and dinner, and so forth) and also enjoyed it, the strength of the conclusion is also positively affected.

But as many observation statements as we have—even when they have the sort of diversity described previously—the truth of an inductive conclusion can never be guaranteed. This is because we are predicting events in the future based on ones in the past, and all sorts of unknown or unpredictable variables could still intervene to make our conclusion false (Goodman, 1983). Inductive conclusions are therefore not *truth-preserving*, in that even though your observation statements might be entirely true, they could nevertheless generate an inductive conclusion that ends up making false predictions about the future.

By contrast, the method of *deduction* tends to be better at being truth-preserving, but this can come at the cost of being somewhat narrower in scope. This is because the conclusions of deductive processes are never more general than the observations or evidence used to generate those conclusions. Deduction works from the general to the specific (or top-down), whereas induction is reasoning that works from specific to the general (or bottom-up).

Induction is therefore often more suited to *developing* hypotheses, and deduction better for *testing* them. Take the example of visiting your doctor when suffering from a fever after having recently visited an area where malaria is common. By induction, your doctor could reason that it's more likely than it would normally be for you to have malaria, seeing as (a) you're exhibiting symptoms common to those who have malaria and (b) you've recently been exposed to risk of contracting malaria. But seeing as it's flu season also, and your doctor has seen a number of patients suffering from flu-related fevers during this week, she now has two competing hypotheses, both formed by induction. The way to tell them apart, and to find the real cause of your discomfort, is to use deduction—in this case, a blood test that will indicate the very *particular* detail of whether or not the relevant parasites have infected your system. From their presence or absence, we can *deduce* that you either have or don't have malaria, and then treat you accordingly.

VERIFICATION AND FALSIFICATION

The nature of induction means that any number of confirming instances can never actually *prove* a theory to be true (because you're generalizing about the future from a limited sample of past occurrences). Any one falsifying instance can *refute* a theory, in that we learn that the theory cannot accommodate all observations, and needs revision in order to do so (Popper, 2002). This can introduce a conflict between scientific thinking and our common understanding of what sorts of conclusions we should regard as reasonable in our day-to-day lives. When we reason about our own lives, we're inclined to make fairly confident predictions about the future based on what we've experienced in the past, and to do so by *verifying* our beliefs through confirmatory experiences—in fact, we might even be more inclined to restrict ourselves to testing the hypotheses we think will turn out to be true, rather than the ones we're not so sure of (Klayman & Ha, 1987). Furthermore, our personal anecdotes can be treated as reliable data, rather than subjective and possibly inaccurate information. And although it is true that the larger and more diverse the number of observational statements we have supporting a conclusion, the more likely it is to be true, a key element of science is the search for *disconfirming* instances rather than verification through more

confirming instances. This is because any number of theories could perhaps explain the confirming instances, and it is thus more useful to try to rule out the weaker theories.

This importance of falsification (rather than verification) is a key aspect that differentiates scientific thinking from our more natural way of thinking. Robust scientific theories need to be open to *falsification*, and scientific theories should ideally offer us—in the specification of the theory itself—what might count as falsifying conditions for the theory. It's important to make a distinction between two separate (though closely related) concepts here: falsifiability *in principle* versus a hypothesis actually having been *falsified* in practice.

From the anthropomorphized perspective of the hypothesis itself, it is of course not a good outcome for a hypothesis to be defeated by the evidence. But, as a scientist, you would hope to approach the evidence objectively, and believe what is most likely to be true, regardless of personal loyalties and preferences. So, from the imagined point of view of knowledge and scientific progress, it's actually a *good* outcome that a hypothesis is falsified in practice, in that we learn about another way in which we were previously wrong about something, and are (sometimes) given a clue as to where the correct answer might lie. To put it another way, for a hypothesis to be falsified rules out one potential source of error and unmuddies the waters of inquiry (at least to some extent).

For a hypothesis to be falsifiable *in principle* asks us to consider a different issue, namely, that we need to be able to imagine ways in which a hypothesis is *potentially* falsifiable. We need to have some mechanism (even if it is a hypothetical one) for testing it, to see whether it can survive those tests. This is a key distinction between science and pseudoscience (as will be further discussed in Chapter 3), in that claims made in fields like astrology tend to be so nonspecific that you cannot even imagine a method of putting them to a test. There are no possible falsifying instances of a general claim like "December will be a good month for you professionally" because even if you lost your job in December, the astrologer could say that this proves her point—you were, after all, stagnating in your current job, and are now free to find a different job, and one that can allow you to discover your *true* potential.

This tendency to look for *verification* of our beliefs, as the astrologer does in the preceding example, is intuitively sensible and it also has quite a heavyweight pedigree in the philosophy of science. In the 1920s and 1930s, the famous Vienna Circle developed the position known as "logical positivism," whereby a statement should be considered meaningful in only two sorts of situations: either it is a *formal* statement, such as you'd find in mathematics (2 + 2 = 4), or if it's capable of empirical verification (Sarkar, 1996).

However, as the astrology example shows, it's easy to find evidence that something is true. But finding confirming evidence doesn't tell us whether we were testing the correct hypothesis. The evidence could instead be confirming some other hypothesis that we haven't even considered yet (e.g., the hypothesis that we are superstitious and prone to believing all sorts of nonsense).

The Vienna Circle had the noble goal of trying to offer a clear principle to distinguish science from pseudoscience. They wanted to make it clear that metaphysical and theological statements were not cognitively or scientifically meaningful. But as Karl Popper pointed out in 1934 (Popper, 2002), the criterion of verifiability was too strong, and also too easy to satisfy. First, because some statements in science are useful and cannot (yet) be verified, like the ancient Greek notion of atoms could not be at the time; and second, because although it's relatively easy to verify a hypothesis, it's far more telling when a hypothesis is falsified, allowing us to rule it out as being true.

This is why the modern scientific method places a significant emphasis on attempts to falsify, rather than attempts to verify. To put it crudely, we take our various hypotheses, develop tests that would demonstrate those hypotheses to be false, and then see which hypotheses survive those tests. The ones that do are considered to be best justified and, in ordinary language, true (at least until replaced by a superior hypothesis, that is). We *triangulate on the truth* via eliminating falsehoods. It's also important to note another consequence of falsification: our conclusions are always provisional, in the sense that they haven't been falsified *yet*. New data could in future come to light that falsifies what we currently regard as true.

CONCLUSIONS

Figure 2.1 summarizes how, ideally, scientific thinking happens. You move from information gathering and hypothesis generation through inductive and deductive processes to help build ever-improving theories about both how the world is and why it is that way. A long and complex process, science continually puts its ideas to the test. Using this kind of method, we give up *certainty* in exchange for increased *confidence* that our beliefs are the best justified ones currently available to us, in full knowledge and humility that we might be getting aspects—or all of—the story wrong at the present moment. Pseudoscience, by contrast, starts with an inflated, and unwarranted, confidence in its conclusions being correct—and then proceeds in a manner that is completely immune to being corrected, because nothing can ever prove it wrong. This is the focus of our next chapter.

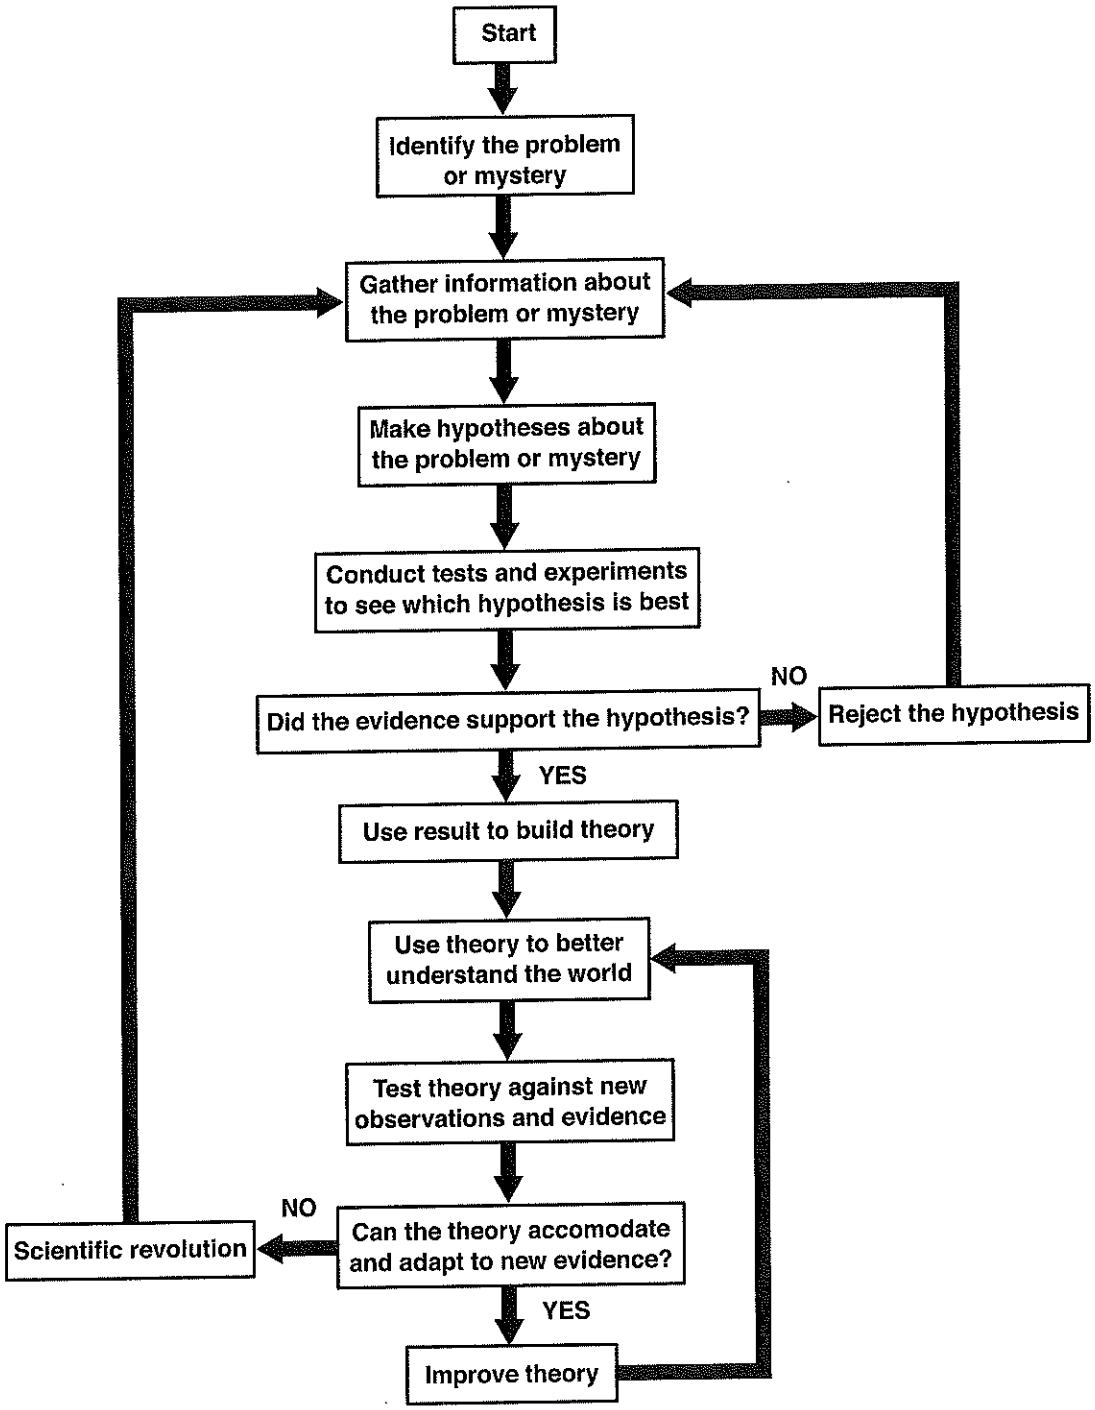

FIGURE 2.1 Science flowchart.

QUESTIONS FOR REFLECTION

1. *What do you regard as the most obvious difference between scientific thinking and our everyday ways of resolving questions? Should we always try to think scientifically about everyday problems? Why or why not?*
2. *Can you imagine ways in which the death of the traditional newsroom, with dedicated science editors and journalists, could positively impact the reporting of scientific developments?*

3. *If a belief cannot be investigated using the tools of science, can we be confident that it's a false belief? Why or why not? Can you think of some examples?*
4. *Do you ever consider whether your personal beliefs can be falsified? If not, do you now think that should you do so?*

REFERENCES

Adams, J. (2013). Collaborations: The fourth age of research. *Nature, 497*, 557–560.

Chabris, C. (2011). *The invisible gorilla: How our intuitions deceive us*. New York, NY: Harmony.

Coyne, J. (2010). *Why evolution is true*. London, UK: Penguin.

De Fretes, F. (n.d.). The secret *by Rhonda Byrne*. Retrieved from http://www.livingggood.com/book-reviews/bookreview-secret-rhonda-byrne

Deutsch, D. (2013). Why it's good to be wrong. *Nautilus, 2*. Retrieved from http://nautil.us/issue/2/uncertainty/why-its-good-to-be-wrong

Fahy, D. (2015). *The new celebrity scientists: Out of the lab and into the limelight*. Lanham, MD: Rowman & Littlefield.

Goldacre, B. (2012). *Bad pharma: How drug companies mislead doctors and harm patients*. London, UK: Fourth Estate.

Goldacre, B. (2014). *I think you'll find it's a bit more complicated than that*. London, UK: Fourth Estate.

Goodman, N. (1983). *Fact, fiction, and forecast*. Cambridge, MA: Harvard University Press.

Holland, J. H., Holyoak, K. J., Nisbett, R. E., & Thagard, P. R. (1989). *Induction: Processes of inference, learning, and discovery*. Cambridge, MA: MIT Press.

Ioannidis, J. (2005). Why most published research findings are false. *PLoS Medicine, 2*(8), e124. doi:10.1371/journal.pmed.0020124

Klayman, J., & Ha, Y. (1987). Confirmation, disconfirmation, and information in hypothesis testing. *Psychological Review, 94*(2), 211–228.

Lilienfeld, S. O., Lynn, S. J., Namy, L. L., & Woolf, N. J. (2013). *Psychology: From inquiry to understanding* (3rd ed.). New York, NY: Pearson.

Mlodinow, L. (2009). *The drunkard's walk: How randomness rules our lives*. London, UK: Vintage.

Okasha, S. (2002). *Philosophy of science: A very short introduction*. Oxford, UK: Oxford University Press.

Popper, K. (2002). *The logic of scientific discovery*. London, UK: Routledge.

Sarkar, S. (Ed.). (1996). *The legacy of the Vienna circle: Modern reappraisals*. New York, NY: Garland.

Shermer, M. (2002). *Why people believe weird things*. New York, NY: Holt.

Shermer, M. (2008). Patternicity: Finding meaningful patterns in meaningless noise. *Scientific American 299*(6), 48. Retrieved from http://www.scientificamerican.com/article/patternicity-finding-meaningful-patterns/?page=1

Wheen, F. (2004). *How mumbo-jumbo conquered the world: A short history of modern delusions*. London, UK: Harper.