7

KNOWLEDGE

Epistemology is the theory of knowledge. Knowledge, evidence, arguments, truth, and belief are the meat and potatoes of epistemologists. While these topics may initially sound abstract or remote, the issues of what you ought to believe and why affect every part of your life. The questions of how we can come to have knowledge, and how far our knowledge can extend, are so basic that epistemology is considered one of the most fundamental philosophical enterprises. Let's start off with a question that, like many so far discussed in this book, is deceptively simple on the face of it. 7.1

The Value of Truth

Do you have a right to your own opinion? It's a safe bet that you are ready to indignantly insist that of course you have a right to your opinion. Like most questions in philosophy, though, a little reflection shows that that it's more complicated than it first appears. One way to understand the question is this: do you have the right to express your own opinion? Put this way it looks like a mere matter of law. In the United States you mostly do have a right of free expression (thanks to the First Amendment to the Constitution). In other countries expression may be curtailed in various ways. In Germany it's illegal to declare your love for the National Socialist Party, in Saudi Arabia it is unlawful to insult the prophet Muhammad, in France it is illegal to boo the national anthem. The issue of expressing your opinions is largely a matter for governments and lawyers, not philosophers. There's 7.2

This Is Philosophy: An Introduction, First Edition. Steven D. Hales.

a way to understand the question of whether you have a right to your own opinion which is clearly philosophical, though:

- Is it OK to believe whatever you want?

Regardless of whether the government might crack down on you for expressing your beliefs, is it all right for you to believe whatever you choose to in the privacy of your own mind? Before you answer, consider the following question, which looks awfully similar:

- Is it OK to do whatever you want?

The answer to this question is obviously *no*. It's not OK to do whatever you want. The reason is obvious: some things are *wrong* to do, there are things you should not do. You could still do those wrong things—perhaps no one's stopping you—but you would be making a moral mistake. Chapters 1 and 2 on ethics in the present book deal with the ins and outs of what you shouldn't be doing. Analogously, maybe there are things that you should not believe. You could believe those wrong things anyway, but by doing so you would be making a mistake. Not a moral mistake, perhaps, but at least an intellectual one.

The rational principle

7.3 So what kinds of things would be the wrong thing to believe? How about this: *false* things. It's an intellectual mistake to believe false things; you shouldn't believe them. Instead (hold onto your hat!) you should believe true things. Perhaps that doesn't sound too radical to you. Let's formulate it as a principle.

The rational principle: You should gain truth and avoid error.

In other words, we should do whatever we can to have only true beliefs; we don't want any false beliefs. As rational thinkers we should prune our garden of beliefs, weeding out the false, the mistaken, the erroneous, the bogus and foolhardy. Instead we keep what is right, true, and real. Now, you may well ask how we can tell the difference between the true and the false, how we can tell the flowers from the weeds. That's an excellent question. But let's hold off on that a bit; we'll get to it shortly.

You might think that the rational principle isn't right, since sometimes 7.4
mistakes are useful and by messing up we can figure out what the right thing really is. In which case we shouldn't avoid error at all. Making errors is a useful step along the road to the truth. It's like if you're learning to play tennis—you hit a lot of shots out before you learn to hit them in consistently.

Well, sometimes things work that way, but often people believe ridicu- 7.5
lous things and never get one inch closer to giving them up and finding the truth. If we stop caring about avoiding errors, that's a recipe for giving up an active search for truth entirely and instead just passively hoping that we'll eventually see our mistakes in the fullness of time. Supposing that we will always see our mistakes for what they are, the principle still holds: we don't *want* to make errors, even if they are inevitable. The goal is to get rid of false beliefs and gain true beliefs. The tennis example is the right one. The first commandment of tennis is *you should hit your shots in and not hit them out*, a mighty fine principle to adopt, even if you're going to miss a lot of shots as you learn the game. A more serious challenge to the principle that you should gain truth and avoid error we'll call the hedonist's challenge.

The hedonist's challenge

The rational principle looks cold and puritanical, just the sort of boring 7.6
edict you'd expect from friendless eggheads. We're each going to kick around on the planet for 80 years or so—who cares whether what we believe is true or false? If you want to believe in space aliens, worship the **Flying Spaghetti Monster**,[1] think that there are energy chakras, or admire talk-show hosts for their insight and wisdom, knock yourself out. It just doesn't matter. If believing something makes you happy, if it gets you through the day, then go for it. If hunting down the truth floats your boat, then go for that instead. But if you prefer the tabloids to ***The New York Times***,[2] that's just as good.

The hedonist principle: You should believe whatever makes you happy.

Hedonists aren't opposed to the truth, they're merely indifferent to the truth. The key difference between the rational principle that you should gain truth and avoid error and the hedonist principle that you should believe whatever makes you happy has to do with what the

value of truth is supposed to be. A little philosophical joke will help illustrate this difference. (There aren't a lot of philosophical jokes; we need to enjoy all the ones we have).

> Socrates was widely lauded for his wisdom. One day the great philosopher came upon an acquaintance who ran up to him excitedly and said, "Socrates, do you know what I just heard about one of your students?"
>
> "Wait a moment," Socrates replied. "Before you tell me I'd like you to pass a little test. It's called the Test of Three."
>
> "Test of Three?"
>
> "That's right," Socrates continued. "Before you talk to me about my student let's take a moment to test what you're going to say. The first test is Certainty. Have you made absolutely sure that what you are about to tell me is true?"
>
> "No," the man said, "actually I just heard about it."
>
> "All right," said Socrates. "So you don't really know if it's true or not. Now let's try the second test, the test of Goodness. Is what you are about to tell me about my student something good?"
>
> "No, not really, Socrates."
>
> "So," Socrates continued, "you want to tell me something bad about him even though you're not certain it's true?"
>
> The man shrugged, a little embarrassed.
>
> Socrates continued. "You may be able to tell me though, because there is a third test, that of Usefulness. Is what you want to tell me about my student going to be useful to me?"
>
> "No, not really."
>
> "Well," concluded Socrates, "if what you want to tell me is neither Certain nor Good nor even Useful, why tell it to me at all?"
>
> The man was defeated and ashamed. This is the reason Socrates was a great philosopher and held in such high esteem. It also explains why he never discovered that his wife was stepping out with Plato.

What is the value of truth, or the value of attaining it? There are two possibilities.

1. The value of truth is *intrinsic*. Truth is valuable in itself, for its own sake, regardless of whether knowing it produces happiness or any other valuable thing.
2. The value of truth is *instrumental*. Truth is valuable insofar as knowing it allows us to survive, achieve our goals, and makes us happy. That is, the truth is no more than a useful tool to help us get what has

> intrinsic value. The Socrates joke treats truth as instrumentally valuable only (the test of Usefulness). Socrates was not harmed in any way by his false belief in his wife's fidelity. The truth doesn't really matter so long as there is no downside to false belief and believing makes you happy. There's no point in caring about truths that lead to unhappiness; in fact in those cases it's better to believe in a nice soothing lie.

There's no doubt that truth does have instrumental value. When considering medical treatment, it is best to know the truth about your condition, and know whether your surgeon trained at **Johns Hopkins**[3] or at Hollywood **Upstairs Medical College**.[4] If you're standing in the middle of a highway, you'd be better off possessing the truth that the truck seems to be getting larger because it is rapidly bearing down on your position, instead of believing the falsehood that sometimes trucks grow rapidly in size. When picking mushrooms for dinner, knowing how to spot the differences among chanterelles (delicious), amanita phalloides (poisonous), and amanita muscaria (hallucinogenic) is a valuable skill to have. The question is whether the *sole* value of truth is instrumental; if truth also has intrinsic value, then we should covet it for its own sake. We should want the truth, in the German philosopher **Fichte's** clarion cry, "even though the heavens fall."[5]

The hedonist principle takes the value of truth to be solely instrumental, 7.7
whereas the rational principle assumes that truth also has intrinsic value. "Gain truth and avoid error" advises pursuing the truth regardless of what it might do for you, or whether it would benefit you somehow. "Believe what makes you happy," on the other hand, recommends the truth only occasionally, just in those cases where it is a nice pleasant truth to have. Otherwise, forget it. How shall we decide which way to go?

Let's consider the following thought experiment and see whether you 7.8
think that truth has intrinsic value as well as instrumental.

Suppose that your boyfriend or girlfriend is cheating on you. Often. 7.9
Further imagine that there are two possible paths the future might take. Path 1: You find out about the cheating. The usual recriminations, crying, accusations, arguments, blowups, and breakups ensue. Path 2: You never find out about the cheating and nothing bad ever happens. Let's be as clear as possible—no one gets an STD, no one gets pregnant, there are never any rumors or suspicions, and from your perspective, everything is going just fine in your relationship.

7.10 Which path do you want to take? Do you want to find out about the cheating or not? If you do want to know, consequences be damned, then you believe that the value of truth is intrinsic. The case was specifically designed so that there was no downside to remaining cheerfully ignorant, and learning the truth only led to unhappiness. If you wanted the truth about your cheating partner anyway, it was not because it made you happy. It was because, to quote **Emil Faber**,[6] knowledge is good. If you preferred the second path, in which your partner keeps on cheating and your never find out about it, then you think that the value of truth is solely instrumental; it is only good to have the truth if it produces something useful or valuable for you have, like happiness. Like Socrates in the joke, if there is no good payout for learning the truth, then forget about it.

7.11 If you are like most people, then you would want to know if your partner is cheating on you. Which means that you think that possessing the truth has intrinsic value. Thus the rational principle is right about our intellectual duties: we should gain truth and avoid error. Even if the rational principle is correct, that doesn't mean that you are necessarily interested in every topic under the sun, or care about what's true in every area of human inquiry. Do you care who will win the next World Cup or American Idol? Do you care whether **Goldbach's conjecture** is true?[7] Does it matter how many angels can dance on the point of a **needle**?[8] Maybe not. Nevertheless, there may be value in things that hold no personal interest for you. As the fidelity example above illustrates, some of that value is intrinsic.

7.12 But how can we pursue this goal? How can we tell whether our beliefs are true or false, or when we should go ahead and believe some claim or proposal?

The Value of Evidence

7.13 In the 1980s **Michael Shermer** was a professional marathon cyclist.[9] In the **Tour de France**, cycling's most famous race, riders churn out up to 140 miles in a single day.[10] Marathon riders do more than twice that distance, for days on end. Shermer once completed the **Race Across America** (3100 miles) in 10 days, 8 hours.[11] Only Iron Man triathletes and ultramarathon runners approach this level of relentless, body-punishing competition. Shermer and his fellow marathon cyclists were ready to try anything to improve their performance, and keep them strong in the saddle. After all, reasoned Shermer, what did they have to lose? If someone had a theory, why not take it on faith and try it out?

Fraud and quackery

Shermer once fasted for a week on a diet of nothing but a potion made of water, cayenne pepper, garlic, and lemon. Halfway through a long ride he collapsed, violently ill. He went out to a health spa for a mud bath that was supposed to suck the toxins out of his body, and found his skin was dyed red for a week. Shermer set up a negative ion generator in his bedroom that would charge the air to give him more energy. It only turned the walls black with dust. An iridologist studied the irises in his eyes, and told him that the little green flecks meant there was something wrong with his kidneys. Shermer's never had a kidney problem before or since. He set up a pyramid in his apartment to focus energy, and only got strange looks from guests. Shermer then had a Rolfing massage, which is a really deep tissue massage, and it hurt so much that he never went back. During one race, he slept under an "Electro-Acuscope" which was promised to measure his brain waves and put him into an alpha state for better sleeping, rejuvenate his muscles, and heal his injuries. Instead it did nothing he could detect. Finally, a nonaccredited "nutritionist" advised taking handfuls of assorted vitamins and minerals every six hours during the Race Across America. They were so disgusting that Shermer could barely choke them down. But he did, and wound up with nothing but the most expensive and colorful urine in America. It was then that he decided maybe he should not believe every extravagant claim and snake oil salesman that came his way. He went back to college and ultimately earned a Ph.D. in the philosophy of science, starting a second career promoting critical thinking. 7.14

The entire history of medical fraud and quackery is based upon sick folks who don't look into the scientific rationale for the claims that people make and instead choose to believe in remedies because they sound good, or conform to their own prejudices, or are slickly marketed. Pleasure is no proof of truth, though. In the nineteenth and early twentieth centuries, trained medical professionals used the electric vibrator to cure "female problems" like hysteria, nervousness, and weakness, by causing hysterical paroxysm (i.e. orgasms). In its advertising copy, the manufacturers of the **White Cross Electric Vibrator**[12] claimed that: 7.15

> Vibration is life. It will chase away the years like magic. Every nerve, every fibre of your whole body will tingle with the force of your own awakened powers. All the keen relish, the pleasures of youth, will throb within you. Rich, red blood will be sent coursing through your veins and you will realize

> thoroughly the joy of living. Your self-respect, even, will be increased a hundredfold. (Maines, 1999, p. 108)

Even earlier, the elixirs of patent medicines usually contained at least one of the following ingredients: alcohol, opium, or a laxative. All three of these had immediate effects of one sort or another, thus assuring the user that the medicine was "working." For example, Lydia Pinkham's Vegetable Compound, chiefly marketed to women, contained some useless vegetable root extracts and was 19 percent (38 proof) alcohol. This of course, was not mentioned in their advertising, which claimed "for all weaknesses of the generative organs of either sex, it is second to no remedy that has ever been before the public, and for all diseases of the kidneys it is the greatest remedy in the world." Needless to say, if you spent the afternoon knocking back some Lydia Pinkham's and going a few rounds with the White Cross, you'd feel a lot better. But that's a far cry from actually curing a disease or improving your lasting health.

7.16 Many quack remedies were, and are, downright dangerous. Consider the Testone Radium Energizer, which was produced and sold around 1900. Radium had only been recently discovered, and the principles of radiation were badly understood, especially by nonphysicists. The Testone Radium Energizer was essentially a jockstrap laced with 20 micrograms of refined radium, 200 times the tolerance dose set for workers at the Manhattan Project. Yet according to the advertising copy, it

> is a scientific means of applying the ENERGIZING GAMMA RAYS to the male gonada, or testes—those fountain-head of Manly Courage and Vigor . . . The Radium Pad comes into direct contact with the testes and completely envelopes them. In this manner, these vital sex glands may be KEPT CONSTANTLY under the strengthening influence of the radium rays—truly a greatly desired benefit . . .

Sounds great, right? Radium has a lot of energy, energy is good for you; it will put some pep in your step. Besides which, the company guaranteed the product. What could go wrong?

7.17 In the list of the world's worst inventions, the Testone Radium Energizer has to rank right up there; sterility, radiation poisoning, cancer, and death are likely side effects. It may not be as bad as the hydrogen bomb, but it's a good thing the Energizer was outlawed before Michael Shermer could buy special Testone Radium cycling shorts. They would surely give new meaning to saddle sores.

Ways we can go wrong

Quack medicine and fraudulent solutions to problems never go away. 7.18
Loads of people have gotten rich through Internet sales of **penis-enlargement creams** and supplements, despite the fact that none of their products work in the slightest.[13] The same thing is true of fad diets and various lotions, minerals, and massagers that are claimed to eliminate **cellulite**.[14] Newspapers regularly recount the plight of victims of stock market and investment scams. There are scores of ways in which we can wind up believing the wrong thing. We can

1. be taken in by swindlers and con artists
2. just look at one-sided evidence
3. refuse to consider the evidence against what we already believe
4. believe in something because we really want it to be true
5. be prone to psychological biases.

There are all kinds of ways that we might mess up in forming our beliefs, and the field of critical thinking is essentially applied epistemology, in which one learns the scientific method, studies how to assess evidence, and examines how to be on guard against shoddy reasoning and our own psychological foibles. For our purposes here, we can at least note this: *evidence matters*. Our rational goal is to gain truth and avoid error, and the best strategy is to look at the evidence. To be sure, there will be plenty of times that our evidence is incomplete, or misleading, and we will wind up believing something false, just as a novice gardener may sometimes pull up flowers instead of the weeds. There are no guarantees. Nevertheless, evidence is the signpost on the road to truth.

How Much Evidence Do We Need?

Part 1: We need a lot

Just how much evidence for a claim do we need before it becomes rational 7.19
to believe that claim? In the late nineteenth century, the English mathematician and philosopher **W. K. Clifford** considered this question and proposed a striking, and firm, answer.[15] Clifford declared that:

> *It is wrong always, everywhere, and for anyone, to believe anything upon insufficient evidence.*

It's not enough to have *some* evidence for a claim. Not just any old reason will do to make a belief a reasonable one. Clifford thought that only *sufficient* evidence was enough to justify believing. Admittedly, just what sufficient evidence means isn't completely clear, but minimally he means that you need considerable evidence to believe something; the evidence needs to obviously point to one thing before believing that claim to be true is the right choice. When there's no credible evidence, or the data is incomplete, or the arguments seem to lead to all sorts of different conclusions, then the smart action is to suspend belief until the evidence is in and we can tell what really is the most reasonable thing to believe.

7.20 Clifford offers a thought experiment. Imagine a shipowner who is about to send out an emigrant ship, carrying several families over the ocean to a new home. The ship he is sending out is rather old, and has a lot of nautical miles on her. People had raised some doubts about her seaworthiness, and suggested that maybe the ship should be dry-docked, inspected, and overhauled before going out again. The shipowner considered these suggestions, but managed to convince himself that the ship was fine and ready to go. He put out of his mind any ungenerous suspicions about the integrity of shipwrights and contractors, and told himself that the ship has already made so many successful voyages that surely the Lord will see the emigrant families safely to their new home. With this sort of reasoning, the shipowner convinced himself that the ship was sound and seaworthy.

7.21 The shipowner was not deceitful or dishonest. He sincerely believed in the soundness of the ship. It's just that his belief was not grounded in actual evidence (like one might get from an impartial inspection); instead he was engaged in rationalization and self-deception. Clifford gives two different endings to this case: in the first the ship sinks at sea with a loss of all hands. In the second the ship makes it safely over the ocean to port. Clifford argued that, from the point of view of rational belief and intellectual integrity, it didn't make a bit of difference whether the ship sank or not. The shipowner did the wrong thing by forming his belief on the basis of self-deception and wishful thinking instead of on the basis of evidence. So whether the shipowner fulfilled his intellectual duties had nothing to do with the truth or falsity of his beliefs about the ship. Instead it had everything to with whether his beliefs were based upon sufficient evidence, good reasons, and cogent arguments. The shipowner needed a lot of evidence to believe that the ship was seaworthy, and he didn't have it.

7.22 You might object that it didn't matter if the shipowner thought that the ship was sound. His beliefs were irrelevant—all that mattered was whether

he inspected the ship before sending it out. He could have believed that the ship was fine and still gone through the process of inspecting it. The problem with this approach, as Clifford himself notes, is that if you have already decided what the truth is before looking at the evidence, you aren't going to be impartial. Prejudice hinders rational assessment, which is why defense attorneys question prospective jurors before a trial (during ***voir dire***[16]) and dismiss those who already have an opinion about the defendant's guilt or innocence. Clifford advises keeping an open mind until all the evidence is in.

7.23 Another objection you might raise is that beliefs are essentially a private matter. Duties and obligations are to other people, not to ourselves; the very idea of an *intellectual* duty is senseless. We might have moral duties to act in certain ways, but that just underscores the idea that it's actions that matter, not what we believe in the privacy of our own minds. The shipowner did something wrong when he sent out a dodgy ship, but what he believed or didn't believe is irrelevant. So he didn't need any evidence at all to form his beliefs; beliefs are just a personal choice.

7.24 The problem with this objection, according to Clifford, is twofold: first, our beliefs invariably influence our actions. Letting careless reasoning and ignorant beliefs off the hook when they lead to negligent actions like that of the shipowner is about like letting a gunman off the hook and putting all the blame on the bullet instead. Second, and more important, we have an obligation to posterity to pass down the best beliefs we can. What we teach our children is the most important legacy that we can leave them, and we should do everything we can to make sure that we pass along the best knowledge of our time. We can fill our children's minds full of prejudice and superstitions or we can educate them and provide what wisdom we can. To be sure, what counts as knowledge to one generation may be overthrown by the discoveries of the next. But we should make sure that our successors do not have to start with nothing.

7.25 Clifford makes no distinction between beliefs like the shipowner's, which mattered considerably to the welfare of others, and insignificant, trivial beliefs that affect no one. He thought that no beliefs are truly trivial. Having unjustified beliefs leads to the bad habit of not caring about evidence and reason, and this habit only paves the way for more unjustified beliefs. It's like a lawman who decides that he'll take only small bribes to look the other way on minor crimes. Once that line has been crossed, small bribes have a way of growing until complete corruption sets in. Crooked cops undermine a civil society based on the rule of law, and sloppy, dishonest thinking

leads to losing reverence for truth and honesty. Not caring about the evidence, not caring about how to reliably find the truth, only means that you don't really believe in the intrinsic value of truth after all.

7.26 Getting things right is so important that not only must we insist on evidence for what we believe, but we need a lot of evidence. To believe, to make a judgment, is to enter into a sacred intellectual trust—it is not something to be done lightly. You've got to make sure that you've done everything you can to get it right.

Part 2: Go on, take a chance

7.27 The philosopher and psychologist **William James** mounted the biggest challenge to Clifford's view that we should only believe if we have sufficient evidence to do so.[17] James and Clifford agree in some important ways; both confirm that:

- Our goal as rational thinkers should be to gain truth and avoid error.
- If nearly all the evidence supports some claim P and there is practically no decent evidence against P, then the rational thing to do is believe P. P is most probably true.

Where they differ is in cases where the evidence is less than sufficient, where the evidence is mixed or ambiguous. For example, suppose that you're wondering whether it will rain later today. You check one newspaper in the morning and it says that it will rain. You look at another newspaper and it says that the rain will pass to the south and miss you completely. So you turn on the TV and the Weather Channel says that in fact the rain is going to hit you after all. But the weatherman on the local news channel reports that the rain will be elsewhere and you'll stay dry. There is credible evidence that it will rain on you this afternoon, and equally credible evidence that it will not rain on you this afternoon. Clearly you want to believe that it will rain if and only if it is actually true. So what should you do?

7.28 In any case where there is mixed evidence for some claim P, there are three possible choices.

1. Withhold belief about P. Just refuse to either believe or deny P until more evidence comes in that clearly settles the matter about P and you can make a more informed decision.

2. Go ahead and believe P anyway.
3. Go ahead and deny P anyway.

Obviously the Clifford approach is to choose (1)—we should suspend judgment about P until we amass enough evidence that clearly points to the truth (or falsity) of P. Wait until we have sufficient evidence before believing anything at all about P. As James points out, Clifford sets very high standards for belief: don't believe anything until there is a ton of evidence for it. By setting such high standards, Clifford is excellent at avoiding errors. If you don't believe anything until there is a mountain of evidence for it, you won't be wrong very often. On the other hand, you won't be believing much at all, since there will be many things for which you just don't have enough evidence to lead you one way or another. Remember that our goal is to *gain truth and avoid error*. The Clifford approach of option (1)—withhold belief until the evidence shakes out—is terrific at half of this goal. Namely, it's great for avoiding errors. It's not so good at gaining truth, because if you don't believe very many things, it follows that you don't believe very many true things either.

In fact, suppose that you didn't care at all about gaining truth, and the 7.29
only thing you wanted was to avoid error. The best approach would be to believe nothing at all. You're never right, but you're also never wrong! Some of the **ancient philosophers**, like Sextus Empiricus, recommended this idea.[18] On the other hand, suppose that all you wanted was to gain as much truth as you possibly can, even if that means making a lot of mistakes. Then you should believe everything you hear. Believe contradictions if you can manage it. Sure, you'll be wrong a lot of the time, but by believing so much, you'll be sure to scoop up all the truths you can. Perhaps no philosopher has defended the extremely gullible view of believe everything you hear. The point is that the two parts of our rational goal pull in opposite directions. *Gain truth* advises believing everything, whereas *avoid error* advises believing nothing. How can we do both at once? Clifford and James both recommend looking to the evidence. Clifford sets a high evidential bar for believing, thus emphasizing the *avoid error* half of the goal. In other words, Clifford adopts the following risk averse principle:

> *Risk averse principle*: Better off to miss out on some truths rather than add more errors.

James's objection is this. OK, the risk averse principle is one way we might go, but why should we think it is any better than a more risk positive

principle, one that promotes *gain truth* instead of emphasizing error avoidance? Such a principle is, he thought, just as rational to adopt as Clifford's risk averse principle.

> *Risk positive principle*: Better off to add more errors rather than miss out on some truths.

In a case of mixed and inconclusive evidence for P, the risk positive principle counsels us to go ahead and take a chance and believe that P is true. Yes, we will be wrong more frequently than with the risk averse principle, but we'll be right more often as well.

7.30 At this point you might well be saying that all these principles about belief, truth, and evidence sound awfully abstract and obscure; how could any of this stuff be relevant to your ordinary life? The answer is swift and perhaps surprising: you make decisions every day using analogues of the risk averse and risk positive principles.

7.31 Here's one example. Just like *gain truth and avoid error* is a good goal, so is *gain pleasure and avoid pain.* The problem is that many pleasurable things are risky or dangerous. Drinking a lot of alcohol may be fun, but there's the downside of hangovers. Skiing is a good time, but you could break your leg. Motorcycles, casual sex, cocaine, falling in love, gambling—all pleasurable, all with risks. If you only cared about avoiding pain, then you would sit on the couch, fasten your seatbelt, and take no chances. If you only cared about gaining pleasure, then you would stand naked on your head, popping a wheelie on your motorcycle at 100 mph down the highway while shooting China white. You're not going to go for either of those extremes. The two middle-of-the-road positions are:

> *Risk averse principle (pleasure/pain version)*: Better off to miss out on some pleasures rather than add more pain.

> *Risk positive principle (pleasure/pain version)*: Better off to add more pain rather than miss out on some pleasures.

So which will you, in general, prefer? Do you tend to avoid risk, or are you more willing to take some chances for a good time? There's not necessarily a right or wrong answer here, so much as a measure of your own personal attitude to taking chances. But, just as the twin goals of gain truth and avoid error are inversely proportional (you maximize one goal at the expense of the other), so too with pain and pleasure.

Here's another analogue, one that, like the pursuit of truth, has broader social ramifications than just your own pleasure and pain. One of the most fundamental objectives of the legal system is to *convict the guilty and set free the innocent*. As in the previous cases, the two parts of this goal are in conflict with each other. Clearly if we only cared about convicting the guilty, we would convict everyone charged with a crime, a plan sure to convict all the guilty parties. Yet if all that mattered was making sure that no innocent people were unjustly convicted, we would set everyone free, which would guarantee that no innocent person would go to prison. The middle positions are these: 7.32

> *Risk averse principle (guilt/innocence version)*: Better off to let some guilty go free rather than convict more innocent people.
>
> *Risk positive principle (guilt/innocence version)*: Better off to convict more innocent people rather than let any guilty people go free.

To ensure that we don't convict any innocent people, or as few as possible, we should set very high standards of evidence—the state should meet a robust and substantial burden of proof. By setting such high standards we won't be convicting very many innocent people, since it is hard to get that much evidence against someone who is in fact innocent. On the other hand, we will be letting many guilty people go free, because the prosecution just couldn't come up with enough evidence to sufficiently demonstrate their guilt. That's the risk averse principle.

To pursue the alternative aim of convicting the guilty we ought to lower our standards and make it easier to punish the accused; with a lower bar of evidence we'll be convicting more people, and therefore be convicting more guilty people. Of course, more innocent people will wind up in jail on the basis of flimsy or circumstantial evidence. That's the risk positive principle. 7.33

Which principle do you vote for in the guilt and innocence case, and why? If you're afraid of criminals harming you or your property, you'll likely opt for the risk positive version, to make sure that as many crooks as possible are behind bars. If you're more afraid of an omnipotent government robbing you of your civil liberties, then you'll want the risk averse principle, which will make it harder for an innocent citizen to get railroaded. There is no perfect solution that will maximize both objectives. 7.34

Recall that when it comes to gaining truth and avoiding error that Clifford defends the risk averse principle. James's point is not that the risk 7.35

positive principle is superior and that's the one we should adopt, but rather that the two are *equally rational*. There's no reason, in James's view, that we should privilege not screwing up over getting it right. We're entitled to believe anything "live enough for us" that it is tempting to believe; it's rational to be a little risky. If there is no evidence whatsoever for a hypothesis, that's beyond the pale; it's still wrong to believe something when there is nothing at all to back it up. But if there's some plausible evidence for a claim, even if it is inconclusive, then go ahead if you want—take a chance and believe it.

7.36 James's view does lead to some surprising results, though. If there is some evidence for and against some hypothesis or proposal P, then there is some evidence for and against not-P. The evidence that counts against P will count in favor of not-P, and vice versa. So on the risk positive principle, it would be just as responsible to go ahead and believe not-P as it is to believe P. James uses the example of "God exists" as a hypothesis he thinks there is some evidence for and some against, but neither conclusive. While James argues that it is rational to be risk positive and believe that God exists, it follows that it is just as rational to go ahead and believe that God does *not* exist when faced with conflicting evidence. So in the very best case, James shows that belief in P based on conflicting evidence is intellectually responsible. But he does not show that we ought to believe P when the evidence conflicts. Maybe we can't be criticized for believing P on grounds of reason and evidence, but we can't be criticized for denying P either. At the end we're left with a curious situation in which belief and disbelief are equally rational. At this point you might miss the good old option of suspending judgment and waiting for sufficient evidence.

7.37 Here's another odd result for James. Consider a case where there is conflicting evidence about some hypothesis. In fact, like our example above about whether it will rain this afternoon, let's suppose that the amount of evidence for and against is exactly the same. There are five data points, five good arguments, five equally decent reasons, (however you want to put it) both for the hypothesis and against it (Figure 7.1a).

7.38 In Case A, James is committed to saying that under the risk positive principle it is perfectly rational to believe that the hypothesis is true (there's evidence in its favor), and it is also completely rational to believe that the hypothesis is false (since there's evidence it is false). So far, no problem. But what about this case (Figure 7.1b)?

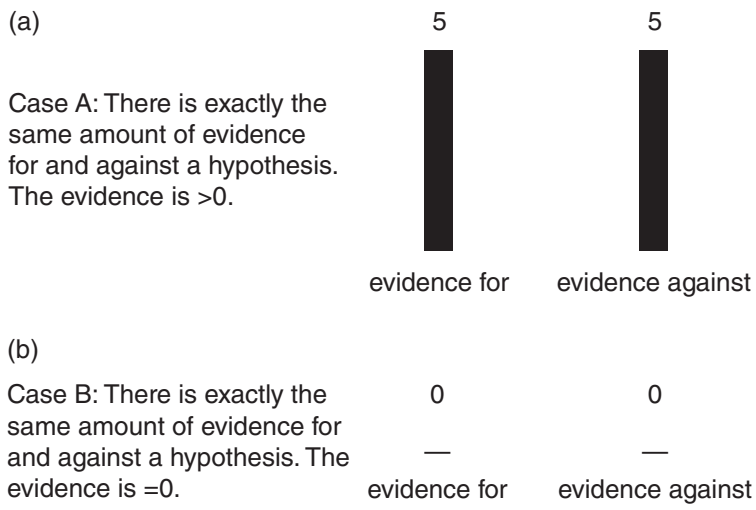

Figure 7.1 Balancing evidence

Suppose the hypothesis is that the next person who walks past you is 7.39
wearing white socks. You can't see the color of their socks, or even whether they're going sockless. So you have no reason to believe that they're wearing white socks, and also no reason to deny it. Is it reasonable to go ahead and believe that they are wearing white socks? Presumably not; believing a claim on the basis of literally nothing whatsoever is surely the very essence of irresponsible belief. James might be willing to take some chances and believe things when the evidence is inconclusive, but believing when there is no evidence at all is something else altogether. So in Case B, the right call seems to be to withhold belief and suspend judgment until some actual evidence shows up that justifies believing one thing or the other.

Here's the problem: why not think that Case B and Case A really come 7.40
to the same thing? That is, there's a sense in which the all the reasons to think it will rain this afternoon and the reasons to think it will stay dry *cancel each other out*. The scale of evidence is perfectly balanced in Case A, just as it is in Case B. So the mixed-evidence case is really the zero-evidence case after all. Which means that it's incoherent to think that belief is a rational choice in Case A, but irrational in Case B. Either you should think that you should withhold judgment in both cases, or it's fine to form a belief in both cases. This result is a predicament for James, who is fine with believing in the case of mixed evidence, but not in the case of zero evidence. So it could be that Clifford is right after all—high standards are the way to

go. Although then we're left with James's puzzle of why we should prefer to avoid error instead of gaining truth.

Sources of Evidence

Perception, testimony, memory, reason

7.41 Evidence for our beliefs can come from all sorts of places. You can gain evidence that James and Clifford once had a debate over the standards of beliefs by reading this chapter. Checking a site like www.weather.gov gives you evidence about **tomorrow's temperature**.[19] Your memory is evidence about what you need to do later; that is, you remember what your upcoming tasks are, and ground your beliefs in those memories. You might even have certain sorts of subconscious instincts or intuitions that are evidential. For example, you might instinctively know you can hit this fastball, or that you said something wrong to your girlfriend. And of course, your immediate perceptions can provide you with evidence—you know that the salsa picante is very spicy because you tasted it.

7.42 The traditional sources of knowledge include testimony (knowledge you get from good teachers and other reliable sources), memory (the things you directly recall to be true), and sensation. There is also the more abstract faculty of reason. For example, in the Sir Arthur Conan Doyle short story "A Scandal in Bohemia," Sherlock Holmes sees Dr Watson and tells him that he can see that Watson had gotten very wet recently, and has "a most clumsy and careless servant girl." Watson, constantly amazed at Holmes's powers of inference, asks how Holmes knew all this. Holmes replies,

> It is simplicity itself . . . My eyes tell me that on the inside of your left shoe, just where the firelight strikes it, the leather is scored by six almost parallel cuts. Obviously they have been caused by someone who has very carelessly scraped round the edges of the sole in order to remove crusted mud from it. Hence, you see, my double deduction that you had been out in vile weather, and that you had a particularly malignant boot-slitting specimen of the **London slavey**.[20]

Here Holmes uses reason to make an inference to the best explanation of what he observes about Dr Watson's shoe. This explanation, which is not the result of perception, memory, or testimony, takes the observation of

the cut and damaged shoe and reasons backwards to its most likely cause. Watson quickly admits that Holmes is right.

Any of these sources of knowledge might be mistaken, and are no guar- 7.43
antee of truth. However, if you do have some knowledge, it's likely that it came from one of these sources. In 1690, John Locke argued that all of these—perception, testimony, memory, and reason—are grounded in one fundamental method, a sort of grand unification theory of knowledge. According to Locke, all of our knowledge ultimately comes from just one source, namely experience, or our sense perceptions. This is his theory of empiricism: there is nothing in the intellect that is not previously in the senses. The bumper sticker version is *no conceiving without first perceiving*.

Empiricism

Locke thought that when we are born our minds are a blank slate, or a 7.44
tabula rasa. We have the capacity to learn, and the faculties for learning, but our minds are empty and waiting for nature to inscribe them. A more contemporary analogy than Locke's blank slate is that when you were born your mind was like a formatted, yet otherwise empty, computer hard drive. Then your experience of the world starts programming your head and filling it with information. When you remember something, you remember some earlier experience you had, some prior perception or sensation. When you rely on the testimony of others, you directly experience that testimony—for instance you read it in a newspaper, hear someone speaking, or look it up online. What's more, you assume that that the authority you've consulted has had direct experience of what they're talking about. If you believe that there was a car wreck on your commute because your friend told you about it, you assume that she saw the wreck herself, or, if she is just reporting what someone else said, that that person saw the wreck. It all goes back to someone's experience.

Even reason, like Holmes's inference to the best explanation above, is a 7.45
matter of the mind assembling and rearranging the basic ideas given to it by experience. Nature programs your mind, but you are able to make connections among your ideas, draw inferences, and reason out new conclusions. You had no innate ideas, though. You weren't born knowing anything; knowledge must come from experience of one form or another.

Empiricism has been a vital step to developing the scientific method. 7.46
Prior to the seventeenth century, knowledge was widely regarded as either

the product of pure reason alone or the result of religious revelation. In fact, a common medieval view is well expressed by Jorge of Burgos, a character in Umberto Eco's brilliant novel set in the fourteenth century, ***The Name of the Rose***, "There is no progress, no revolution of ages in the history of knowledge, but at most a continuous and sublime recapitulation."[21] After the fall from the golden age—whether Eden or the heights of classical Greece and Rome—there is no such thing as the advance of knowledge. Such a view is shocking only to us, heirs of empiricism and scientific inquiry, but not to Jorge's audience.

7.47 There have been many criticisms of Locke's empiricism. For one, it is hard to see how experience of the world leads to mathematical knowledge. Consider the Pythagorean Theorem: $A^2 + B^2 = C^2$. The length of the hypotenuse of a right triangle is the square root of the sum of the squares of the two sides. **Pythagoras** was reportedly so delighted when he discovered this theorem that he sacrificed several oxen to the gods.[22] But how could Pythagoras have discovered it empirically? No actual, physical, triangle has a perfect right angle or even perfectly straight sides, no matter how carefully you attempt to draw one. Therefore the Pythagorean Theorem can't be 100 percent accurate of any triangle in the world, even if it is 99.9 percent accurate. Yet all mathematicians regard it as an absolutely correct truth of Euclidean geometry. So if the theorem is true, it must be true of abstract, mathematically ideal triangles, not physical triangles that Pythagoras actually saw. Then how was his knowledge of the theorem grounded in empirical experience? It seems that it can't be.

7.48 More generally, there seems to be knowledge that, like knowledge of mathematics, is a priori. "A priori" is philosopher-speak for knowledge that is prior to experience, and independent from it. Consider ethical claims, like "you should keep your promises," or "capital punishment is immoral." Those might be true or they might be false, but either way it does not look like science, observation, and testing are going to sort it out. Philosophical reflection and argument, not empirical experience, is a more effective strategy.

7.49 Or how about these:

- All squares are rectangles.
- Everything that has a shape also has a size.
- Nothing completely red is completely blue.
- If anyone is a cyclist, then there are bicycles.
- All bachelors are unmarried.

You ask bachelor after bachelor for his marital status, and lo and behold, all tell you that they are wifeless. But you didn't have to do that—while experience can confirm that all bachelors are unmarried, you don't need to conduct such a poll to know that all bachelors are unmarried. Furthermore, no conceivable experience could undermine or refute the bulleted statements. They are true come what may; they must be true, they are *necessarily* true. Authentically empirical claims are not necessarily true. Consider these:

- The universe is 13.7 billion years old.
- Perpetual motion violates physical law.
- A meteor killed off the dinosaurs.
- My desk is 6 feet wide.
- Human beings share 98 percent of their DNA with chimpanzees.

Statements that are really empirical are contingent; they might be true or they might be false. We'll find out from experience, whether it is simple observation or a complicated scientific experiment, which it is. A priori statements are not contingent. If they are true, then they are necessarily true, and the most experience can do is to confirm the obvious, as in the bachelor example. If an a priori statement is false, then it is necessarily false. For example, "some circles have three corners" is necessarily false, and no matter how many circles you look at, you'll never find one with three corners. You don't even need to look at any to know that.

7.50 So one sort of criticism of Locke's empiricism is that there is knowledge that is apparently not grounded in experience. Another criticism is that a priori knowledge has a quality of necessity about it, which empirical knowledge does not. Again this suggests that at best some of our knowledge is ultimately grounded in experience, but not all. But what is knowledge exactly? We have so far examined the value of the truth, the value of evidence, the matter of how much evidence we need to justify belief, and looked at little at some sources of evidence. Let's now ask about knowledge itself.

The Nature of Knowledge

7.51 In the film *Men in Black*, Agent Kay (played by Tommy Lee Jones) revealed to Agent Jay (played by Will Smith) that there were extraterrestrial aliens

secretly living in New York City. Afterwards, as they sat on a public bench, Agent Kay said,

> 1500 years ago everybody 'knew' the Earth was the center of the universe. 500 years ago everybody 'knew' the Earth was flat. And 15 minutes ago you 'knew' that people were alone on this planet. Imagine what we'll '**know**' tomorrow.[23]

Kay's point is that we often believe that we know things that we do not in fact know. It might have been the case that 1500 years ago that everybody thought that the Earth was the center of the universe, and, had you asked them, would have claimed to know it. But they were wrong. They knew no such thing, for the simple reason that the Earth is *not* the center of the universe. Likewise, before Agent Kay showed him the truth, Agent Jay would have claimed to know that there were no extraterrestrials living in Manhattan. In other words, Jay was wrong; he didn't know the truth about the aliens. Kay advises humility at the end of the quotation above: tomorrow we may well find out that we don't really know the things we think we do.

7.52 How can you discover that you don't actually possess the knowledge you think you do? Like Agent Kay, it is by finding out that you were wrong and that what you believed was mistaken. Actual knowledge requires the truth; to realize that you were in error is to grasp that you didn't really know what you thought you did. It follows that part of what it is to legitimately and authentically know something is to be in possession of the truth. Knowledge requires that you believe something, and it is true. Here is a first proposal for the nature of knowledge.

Analysis of knowledge, first attempt

Knowledge = true belief

7.53 That can't be entirely correct, though, because you might have a true belief by luck alone, and in that case it doesn't sound right to say that you have knowledge. For example, suppose that you buy a lottery ticket for Powerball. You're just feeling lucky, and you manage to convince yourself that this time you are definitely going to win. Let's imagine that by a stroke of incredible fortune, you beat the one in 175 million odds and actually take home the prize. Did you know that you were going to win? Well, you

believed that you would win, and your belief was true. If knowledge is just true belief, then you did know. But that can't be right—no one knows that they will win Powerball. At best you could make a fortuitous guess, but then it is merely luck, not knowledge.

7.54 True belief is part of the story, but not all. What else do we need to convert true belief into knowledge? There needs to be something that *connects* your belief to the truth, some way in which belief is sensitive to the truth, and doesn't just stumble over it by accident. The traditional answer, going back to Plato, is that knowledge is true *justified* belief. Earlier we looked at the value of evidence in gaining the truth, and the idea here is that it is sufficiently strong evidence that knits belief and truth together into knowledge. How much evidence do we need before we can claim that a belief is justified? Well, this too was looked at earlier with the Clifford/James debate. Without returning to that thicket, let us just assume that there is some threshold of evidence that marks the boundary between beliefs that are justified and those that are not.

7.55 Note that a belief might be justified and still turn out false. Agent Jay's belief that there were no extraterrestrial aliens living in New York City was entirely justified on the evidence that he possessed. As Lisa put it to Homer in **The Simpsons** episode The Springfield Files, "It's just that the people who claim they've seen aliens are always pathetic low-lifes with boring jobs. Oh, and you, Dad."[24] But Agent Jay's belief that there were no aliens, while justified, was nevertheless false. Having good evidence does not guarantee that you'll hit the target of the truth. Thus a belief might be true but unjustified (the lucky guess), or justified but false.

7.56 Suppose you hit all three cherries on the slot machine: belief, justification, and truth. Do you win the payout of knowledge? Here's a revised conception of knowledge.

Analysis of knowledge, second attempt

Knowledge = justified true belief

7.57 The requirement of justification was added to escape the problem of luck, and for a very long time, philosophers thought it did the trick. It turns out, though, that epistemic luck is an insidious foe. Imagine a clock that is very precise and reliable—you depend on it frequently, and every time you've looked at it in the past, the clock has given you the right time. Unfortunately, while unknown to you, the clock has stopped working with its hands

left pointing to 3:00. You glance at the clock for the time, and as a result come to believe that it is 3:00. By sheer coincidence, you looked at the clock at exactly 3:00. Thus you believe that it is 3:00, it is true that it is 3:00, and you are justified in believing that it is 3:00 (since you came to have that belief on the basis of looking at a generally reliable clock). However, you don't *know* that it is 3:00—you can't know what time it is from looking at a stopped clock. You just got lucky that you looked at it at 3:00. It's the problem of epistemic luck all over again.

7.58 In the clock case you have a justified true belief, but you do not have knowledge. Therefore either knowledge is justified true belief plus some additional condition, or knowledge needs to be reconceived as something else altogether. There is no generally accepted view on the correct analysis of knowledge, although there are many creative and ingenious attempts. These are all beyond what can be addressed in the present volume. However, it is fair to say that justified true belief is generally regarded as being *close* to knowledge, even though not identical to it.

The Skeptic's Challenge

7.59 Whatever the standards are for evidence, no matter where we set the bar for justification, and regardless of the correct analysis of knowledge, there are those who think that it is *never* reasonable to believe anything, and that knowledge is perpetually elusive. Such people are known as skeptics. Before we get started on their arguments, it's important to note that "skepticism" can mean two different things.

Modest skepticism and radical skepticism

Modest skepticism. Modest skepticism is no more than critical thinking. It's the idea that you should demand evidence before you believe a claim, buy a product, join a religion, or vote for a candidate. And when you are offered reasons, you should scrutinize those reasons closely and consider opposing points of view. Make sure that the premises of the arguments you're considering really do support their conclusions, and that the premises themselves are acceptable ones. Be aware of the fact that smooth-talking charlatans will try to convince you of things that are dangerous, dumb, irrational, and all-around boneheaded.

Radical skepticism. We can neither avoid error nor gain truth. We're either incapable of eliminating error, invariably committed to circular reasoning, or should suspend judgment indefinitely. As far as the search for knowledge goes, we're basically screwed. If we have any true beliefs at all, we have them by accident. There's no trustworthy way to separate the true from the false, or, if there is, we can't figure out what it is.

Modest skepticism has been a theme so far in the present chapter: truth is 7.60
valuable, we should gain truth and avoid error, and getting some evidence is the best way to do so, even if there is debate about how much evidence justifies belief. The fun fact about radical skepticism is that just about no one agrees with the skeptical conclusion that knowledge is impossible. On the other hand, there's very little agreement about the best way to rebut the skeptic's ingenious arguments, or even if it is possible to do so at all.

So what are these arguments? There are many kinds of radical skepti- 7.61
cism, going back at least to Sextus Empiricus in the second century. The most famous skeptical argument attempts to show that there is an unbridgeable gulf between the truth about how the world really is, and any evidence we might marshal about it. The best known, and most discussed, of these skeptical arguments come from the French philosopher **René Descartes**.[25] In his little book *Meditations on First Philosophy*, Descartes offered some puzzling and disturbing thought experiments. Let's look at Cartesian-style skeptical reflections.

Dreamers, demons, and movies

Have you ever woken up in the morning, gotten out of bed, fixed yourself 7.62
some breakfast, brushed your teeth, and whatever else you do in the mornings, and then *really* woken up and realized that you had been just dreaming about having gotten up and starting your day? It's startling and disorienting, and for a while you're not sure what's real, whether you really have eaten a bowl of cereal or not. Now suppose that while you're trying to sort out your strange dream and get on with the morning, the exact thing happens—again you wake up and find yourself lying in bed. You had the same dream over again, just with the addition of a "waking-up" dream added to it. How many times would this have to happen to you before you wake up and say to yourself "All right, this is probably just another of those freaky waking-up dreams and in fact I'm still asleep in my bed." Of course, maybe you really did wake up for real this last time. How could you tell?

7.63 In fact, how can you ever tell if you're dreaming or awake? Dreams can be hyperrealistic, and, when you are in the middle of a dream, it feels as real as our waking lives. In our dreams we can be frightened, or amused, or sexually aroused. We have conversations with other people, and don't know what they are going to say next, just as we do awake. Sometimes you may be able to tell yourself "This is just a dream," and it really is. On the other hand, when some wonderful or terrible thing has happened to you in reality, you may tell yourself the very same thing ("I must be dreaming!"). Some people think that by pinching yourself you can determine if you are awake, a rather silly idea, since you could just as easily dream that you are pinching yourself.

7.64 In the movies, characters often dream and have no idea that they are dreaming. Dorothy's entire adventure in ***The Wizard of Oz*** is revealed at the end to have been nothing more than a dream that resulted from her being knocked unconscious by a tornado.[26] For Dorothy the dream of Oz was more vivid—and in brilliant Technicolor—than her drab real life in monochromatic Kansas. It never occurs to Dorothy, just as it rarely occurs to characters in other and-then-they-woke-up movie plots, that she has no reason to believe that Kansas is any more real than Oz. That is, the feeling of having woken up, and being surrounded by Auntie Em, Uncle Henry, and others could itself just be another dream. If the profoundly real-feeling Oz could have been just a dream, then the similarly real-feeling Kansas might be just a dream too. Even more, why wouldn't Dorothy think that vibrant, colorful Oz was reality, and thin, bloodless, sepia-toned Kansas was the shadow world of dreams?

7.65 In the fourth century BCE, the Chinese philosopher Chuang Chou (莊子, also known as Master Chuang, or Chuang Tzu; in pinyin, 莊子 is transliterated as **Zhuangzi**)[27] pondered the dream argument, and gave perhaps the first statement of it.

> Once Chuang Chou dreamt he was a butterfly, a butterfly flitting and fluttering around, happy with himself and doing as he pleased. He didn't know he was Chuang Chou. Suddenly he woke up and there he was, solid and unmistakable Chuang Chou. But he didn't know if he was Chuang Chou who had dreamt he was a butterfly, or a butterfly dreaming he was Chuang Chou. (Tzu, 1968, p. 49)

Master Chuang drew out the skeptical implications. If he couldn't tell whether he was a butterfly or a man, then he knew very little indeed.

> He who dreams of drinking wine may weep when morning comes; he who dreams of weeping may in the morning go off to hunt. While he is dreaming he does not know it is a dream, and in his dream he may even try to interpret a dream. Only after he wakes does he know it was a dream. And someday there will be a great awakening when we know that this is all a great dream. Yet the stupid believe that they are awake, busily and brightly assuming they understand things . . . (Tzu, 1968, p. 47)

Maybe you don't like the dream argument. Perhaps you think that there has got to be some sort of test to tell whether, at any given moment, you are dreaming. Pinching might not work, but surely there's something. Fine, says the skeptic. The dream argument is just one arrow in a vast quiver. There's no end to the skeptical scenarios that we can devise. Descartes himself offered another one, sometimes called the Demon argument. Here it is. Suppose there is an evil Marvel Comics supervillain called The Demon, who uses all his powers to delude you. The Demon is a trickster, an illusionist, a black magician who fools your senses with his powers of necromancy. What you see is unreal, what you hear is bogus, what you taste, touch, and smell is all a sham; everything your senses tell you is just an illusion conjured up by The Demon to keep you permanently deluded about how the world really is. The Demon never reveals himself, of course; he keeps his own existence as hidden as the real world.

7.66 Can you tell whether The Demon really exists or if he is just make-believe? Can you tell the difference between reality and the deceptive illusions of The Demon? Presumably the answer to both questions is *no*. If The Demon exists and you are his victim right now, nothing changes for you—the world appears just as it always has. The difference is that, in fact, those appearances have nothing to do with reality, and you are massively deceived about what's real. If The Demon doesn't exist, and your senses genuinely are putting you in reliable contact with reality, then again, as far as you can tell, nothing changes. It's just that you do have knowledge about the world instead of being stuffed full of false beliefs by The Demon.

7.67 The skeptic is not claiming that The Demon really deceives you, or that you really are dreaming at this moment. Rather, the claim is simply that you *might* be. You just can't tell. Since you might be massively fooled, and you have no way of telling, you can't actually claim to know anything about the world around you.

7.68 Once you see how to spin out these skeptical scenarios, it's easy to come with any number of them. The film *The Matrix* is just another version of skepticism. In the movie, the reality perceived by humans is generated by

computers who feed it directly into the brains of captive humans used as power sources for the machines. Instead of sense organs delivering electrical impulses, which the brain then interprets as sight, sound, smell, taste, and touch, the computers provide the electrical impulses instead. The resulting perceived world is known as "The Matrix." It feels and seems 100 percent genuine, but the neural simulacrum world of the Matrix in fact has almost no relation to the real world. As Morpheus, one of the main characters in the film, states, "The Matrix is everywhere. It is all around us. Even now, in this very room. You can see it when you look out your window, or when you turn on your television. You can feel it when you go to work, when you go to church, when you pay your taxes. It is the world that has been pulled over your eyes to blind you from the **truth**."[28]

7.69 How can you tell whether you are living in the **Matrix** right now?[29] The skeptic's challenge is that you can't. Anything you might offer as evidence that the world as you perceive it is real, and that your beliefs about that world are true ones, could all be the result of malevolent computer programming.

7.70 The skeptic draws a sharp break between the sensations and perceptions inside your mind and whatever might be the cause of those sensations. Normally, you assume that if you see a chair, an actual chair outside of your mind that your perception somehow resembles causes your perception of a chair. The skeptic's strategy is to argue that your perception of a chair could have been caused by any number of different things—dreams, demons, the Matrix, etcetera. A real chair could even have caused it, as you believe. Who knows?

The theater of the mind

7.71 One way to think about the skeptic's argument is to think of the mind as a kind of theater. On the screen are all the images and sounds of a world. You experience the noises and flickering lights, interpreting them as people, places, and things. But the cause of the images on the screen is completely unknown to you—all your experience is on the screen.

7.72 In Figure 7.2, you (whatever you are; see the chapter on personal identity) are depicted by the stick figure on the left. Your perceptions, not only sight, but sound, taste, touch, and smell, are represented by the movie screen. It is all enveloping, the complete sense-surround system of your life, all your experiences and sensations are there. The projector is whatever is causing the images on your mental screen. But what is it exactly? What's

Figure 7.2 The theater of the mind

the cause of all your sensations? That is the great mystery. Here are some hypotheses about what the projector could be. The projector is:

1. A world of objects external to the mind that the images on the screen resemble.
2. Your own dreaming mind.
3. The Demon.
4. The malevolent machines in the Matrix.
5. Any number of other skeptical scenarios that we might think up.

You believe that the correct explanation of the projector is option 1, a world of mind-external objects that cause representations of themselves to happen in your mind. The skeptic's point is that there is absolutely no rational evidence whatsoever to believe that. For all you can tell, one of the other options is causing those sensations. You can never get past the screen to figure out the nature of the projector, because any evidence you might give is *on* the screen, not behind it. All we have to go on to figure out the nature of reality is the sensations in our minds, but the projector—the true nature of the world—is forever hidden from view.

Let's put the skeptic's argument another way. Take any ordinary belief 7.73
you have about the world: you have a body, you're reading this chapter, you're wearing clothes, you live on planet Earth, whatever. If any of those

ordinary beliefs are true, then no skeptical scenario is correct. If your ordinary beliefs are right, then you're not just dreaming that you have a body, or deceived by The Demon into believing you're reading this chapter, or wearing clothes provided by the Matrix. In other words, the truth of your ordinary empirical beliefs is incompatible with any skeptical scenario being true. Let's call this:

> *The metaphysical principle*: If any ordinary claim about the world is true, then no skeptical possibility (dreaming, Demon, Matrix, etc.) is true.

The metaphysical principle is extremely plausible; it essentially states that if the projector in the theater of the mind is a world of objects, then the projector isn't anything else, like The Demon. Hard to see how we might fault that rather obvious point. In fact, let's go ahead and claim that we *know* the metaphysical principle to be true. Since this is a claim about knowledge, let's call it:

> *The epistemic principle*: We know that if any ordinary claim about the world is true, then no skeptical possibility is true.

The epistemic principle just says that we know the metaphysical principle. All fairly straightforward so far. Now, says the skeptic, I've got you! Let's just imagine for a second that you know that you are reading this chapter. That's an ordinary, routine claim about the world. By the epistemic principle, it follows right away that you know no skeptical possibility is true. But, says the skeptic, that exactly what you *don't* know (since, you might be dreaming, deceived, and so on). Thus you don't know you are reading this chapter. Here's the argument in outline:

1.	You know that you are reading this chapter.	premise
2.	You know that (if you are reading this chapter, then you are not merely dreaming that you are reading this chapter).	epistemic principle
3.	Therefore you know that you are not merely dreaming that you are reading this chapter.	From 1, 2
4.	But (3) is precisely what you don't know, according to the skeptic.	premise
5.	Since the assumption of (1) led to a contradiction, it must be false. You don't know that you are reading this chapter.	3 & 4 contradict

The first premise is assumed for the sake of the argument and could be any modest thing that you think you know about the world—you know you're reading this chapter, you know that the sky is blue, you know that Domino's delivers, whatever. According to the epistemic principle, which we defended earlier, if you know one of those simple facts about the world, then you know that you're not just dreaming the whole thing. You can't very well *know* that you are reading this chapter if you're unable to rule out the possibility that you might just be dreaming it, or living in the Matrix or something. It follows directly in (3) that you *can* rule out those skeptical scenarios. Great, right? Not so much, says the skeptic. Skeptical possibilities are the very things that you can't just write off. You *don't* know you're not dreaming, or deceived by The Demon. In other words, having ordinary knowledge of the world implies that you can dismiss the skeptic; but since you can't legitimately ignore the skeptic, you don't have ordinary knowledge of the world after all.

7.74 Descartes himself gave one sort of response to the skeptic. Even if the skeptic is right, and we don't know anything about the world outside of our minds, that does *not* mean that we don't know anything at all. There's still plenty that we do know. For example, you know that you exist. As Descartes famously wrote, "Cogito ergo sum." I think, therefore I am. In fact, try to imagine that you, right now, don't exist. Maybe you can imagine that there was a time before you existed, and a time after you cease to exist. But when you try to imagine that you don't exist right this minute, well, who's doing the imagining? You are! So you must exist. Another way to think about it is that you might be fooled or deceived about many things, but your existence isn't one of them. If The Demon tricks you, whom is he tricking? Again, it's got to be you. You have to exist in some manner to be the subject of deception. You can be certain, Descartes argued, of your own existence.

7.75 You know that you exist. But that's not all. Consider again the metaphor of the theater of the mind. The skeptic's challenge is that you don't know the nature of the projector since all you have access to is what on your mental screen. Yet you do have immediate apprehension of what's on the screen, that is, the contents of your own mind. One sort of thing in your mind is *phenomenal states*, that is, sensations and feelings. You can be sure that you are having a noisy, red tractorish sensation, even if you have no idea whether a red tractor outside of your mind is really causing it or you are just dreaming the whole thing. You can know that you are feeling happy, blue, melancholy, cheerful, angry, wistful, jealous, disappointed, or excited even if you do not know the causes of these sensations, or are seriously

mistaken about why you have those feelings. The Matrix may be able to fool you into believing you are eating a piece of steak, but cannot fool you about having a sensation of juicy deliciousness or being pleased about it. You are authentically having those sensations and feelings.

7.76 Besides phenomenal states, you also have *intentional states*. These are things like beliefs, hopes, desires, fears, wishes, loves, and hates. For example, you believe that you are wearing clothes right now. You might be mistaken about that (that's the skeptic's point), but one thing you are not mistaken about is that *you believe* that you are wearing clothes. You're not wrong about what it is that you believe. There is a subtle distinction to be made here—your beliefs can be wrong (you might be the victim of massive skeptical deception), but you aren't wrong about whether you have them. Suppose you believe that there is a Santa Claus. Your belief that there is a Santa is false (sorry), but your additional belief that *you believe that there is a Santa Claus,* that one's true. You might be afraid of spiders without there being any spiders, but you still know that you have that fear. You can know that you're in love with James Bond, despite the fact that he's fictional. These sorts of states are in your mind; they're on the mental screen, not behind it. The skeptic gets traction by proposing various hypotheses about the projector, but what's on the screen—that's something, Descartes thought, that you have direct, inerrant, access to.

7.77 Suppose that Descartes is right, and you do know the fact of your own existence, the nature of your phenomenal states, and your intentional states. That's not exactly an atomic pile driver move against the skeptic. All it really does is insist that despite skepticism we can still have knowledge of the contents of our minds, along with the fact that we have minds. The skeptic still holds the better hand, a straight flush against your pair of threes. You still know nothing about the projector behind the screen in the theater of the mind; you know nothing at all about the nature of extramental reality.

7.78 There have been many responses to Cartesian-style skepticism, as you might imagine. It is safe to report that none have been widely accepted as wholly convincing. A thorough survey of approaches to skepticism is well beyond the ambitions of the present book. Nevertheless, it's only fair to offer a little taste of how some have answered it. The English philosopher **G. E. Moore** gave one well-known response to the skeptic.[30] Moore's argument is very easy to state, and reactions to it tend to fall into one of two camps. People either think that the argument is a brilliant, common sense response to skepticism that settles the issue, or they think that Moore's

argument is childishly naive, and assumes the very thing it is supposed to prove. There isn't much of a middle ground. Here's the argument:

1. If skepticism is true, then we have no empirical knowledge.
2. But we do have empirical knowledge.
3. Therefore skepticism is false.

Simple, huh? What's more reasonable, Moore says, that you know that you have a hand, or that you know that the skeptic's scenarios really are possible or his reasoning is legitimate? We start with basic truisms like "I have a hand" or "I live on planet Earth." The rest of our knowledge rests on things like that. These truisms are what we know the best, says Moore, and we should be more confident in them than in anything the skeptic has to say. Yet the skeptic asks us to throw out our very starting point, the things that we know the best.

Essentially, the skeptic claims that the following are jointly inconsistent— 7.79
they can't all be true, and at least one has to be false.

a. *The epistemic principle*: We know that if any ordinary claim about the world is true, then no skeptical possibility is true.
b. We know ordinary claims about the world.
c. We do not know that no skeptical possibility is true.

The skeptic's answer to the inconsistency problem is to toss out (b). The epistemic principle is true, and it's true that we can't eliminate skeptical possibilities. The only choice is to admit that we don't know ordinary things about the world. Moore's response is to agree with the skeptic about the epistemic principle, but keep (b) and throw out (c). That is, Moore insists that we do know ordinary things about the world; he gives a long list of such things, such as "you know that you have a hand" and "you are, and have always been, very near planet Earth." It is far more reasonable, Moore thinks, to hold onto this commonsense knowledge than it is to concede that we can't summarily reject the skeptic's fanciful scenarios. If the choice is *deny that you know you have a hand* or *deny that you might be living in the Matrix*, Moore opts for the second. Therefore, we do know ordinary facts about an extra-mental world, and we know that skepticism is false.

You may well be persuaded by the theater of the mind, and think that 7.80
Moore is just missing the skeptic's point entirely. As noted, philosophers

are divided about his approach. But sometimes the best defense is a good offense.

The Counterfeit Detector

7.81 Skepticism of the sort just discussed leads us to the conclusion that we may be ignorant about the genuine nature of the world outside of our minds. There is a more comprehensive skeptical argument that aims to show that we know nothing whatsoever. It doesn't matter whether you are a fan of perception, introspection, or pure reason; no matter what method you use to gain your beliefs you will never be able to achieve knowledge. This conundrum is known as the problem of the criterion.

7.82 We're interested in getting some knowledge. Remember Faber's dictum from earlier: knowledge is good. But how can we tell whether we have a bit of knowledge, whether one of our beliefs achieves the exalted state of knowledge? In part the answer has something to do with whether we have the right evidence, and if we have enough of it to justify what we believe. No problem so far. But how can we tell whether we really do in fact have the right evidence and if we have enough of it to convert our true beliefs from a lucky accident into knowledge? That is, how do we know when to trust our evidence as authentically leading us to the truth? An example will help.

Genuine and counterfeit money

7.83 Imagine that you are in the **Secret Service**.[31] In the movies, Secret Service agents are always jumping in front of bullets to protect the president. To be sure, that's one of their jobs. But another one, and in fact the very reason that President Lincoln created the Secret Service, is to protect the nation's currency by tracking down phony bank notes and arresting counterfeiters. It's important to get funny money out of circulation because if there is enough of it floating around, it will lead to a devaluation of US currency and then to inflation. In a way, the presence of bad bills poisons the good ones.

7.84 Suppose someone hands you a $100 bill. How can you tell whether it is a fake? To tell whether a $100 bill is a counterfeit, you need some method of testing it. There are many ways we could try to tell the difference between legit $100 bills and phony ones. We could weigh them, for example, or we

could see how they respond to certain chemicals, or we could judge money based on our personal impressions of the character of the owner. There are good methods and bad methods. A good method will reliably pick out the true $100 bills and also reliably winnow out the counterfeits. As you can see, this is another analogue of *gain truth and avoid error*.

7.85 The Secret Service has a method: there are certain features that they look for to detect counterfeits. One of the things they examine is the paper used. **Crane and Co.** has made the paper for the Treasury since 1879,[32] and their proprietary paper contains "counterfeit deterrents, such as advanced security threads, watermarks, planchettes, security fibers, special additives, and fluorescent and phosphorescent elements." Other things to look for are the crispness and texture of the engraving, a watermark of Benjamin Franklin, and a color-shifting numeral 100. New notes also have a 3-D security ribbon that changes images as the bill is tilted and moved. In other words, the Secret Service has a whole system worked out to separate good money from bad.

Particularism and methodism

7.86 How can you be confident that the Secret Service method is a good one? The answer is that the government makes the money—they buy the fancy paper from Crane and Co. and print it themselves at the Bureau of Printing and Engraving. So they have perfect examples of genuine $100 bills fresh off the printing presses, and can be completely certain that those bills are real. All that needs to be done is to compare any $100 note in circulation with the new ones the Secret Service is certain are genuine, and see whether they pass the test. In other words, to be a counterfeit detector, you start with a particular example of a true $100 bill and by studying it come to derive a method for spotting fakes. Let's call this approach *particularism*.

> *Particularism*: Start with some examples of what you are positive is real and then from those figure out a reliable method to separate the true from the false.

Suppose that someone hands you a bunch of $100 notes. You're always broke and have never held a $100 bill in your hands before, much less worked for the Bureau of Printing and Engraving or the Secret Service. You're in no position to derive your own method of counterfeit detection. So how can tell whether the notes you're holding are real? You don't want

to be taken in by fakes. The solution is that you can just apply the methods discussed above—examine watermarks, the colors and engraving, the 3-D ribbon, etcetera. You're confident that the Secret Service knows what it is talking about and that those are good ways to check for fakes. The idea that to separate out the real from the phony you start with some method in which you're confident and just apply it is called *methodism*.

> *Methodism*: Start with some method you're sure is reliable and then use it to distinguish between the true and the false.

With particularism, you assume that you already have some examples of genuine $100 bills, and from there come up with a method for identifying fakes. Methodism works in the other direction. You start by assuming that you have a great method of identifying fakes and then use it to figure out which $100 bills are genuine. You may notice a bit of a problem here. Particularism assumes you can already tell the difference between real bills and phony ones. That's how you know you have some examples of genuine Benjamins. Which is to say, particularism works only by presupposing that you have some reliable method of picking out the true from the false. Methodism, on the other hand, presupposes that you already have some samples of real bills—that's how you're able to decide among different methods to detect counterfeits and figure out what system is the trustworthiest one. So particularism can work only if you already have a reliable method; in other words, methodism is logically prior to particularism. On the other hand, methodism will work only if you already have samples of bills you are positive are genuine; in other words, particularism is logically prior to methodism.

The wheel

7.87 The government has a way out of this circle (Figure 7.3): it prints the money. Therefore the $100 bills rolling off the assembly line are guaranteed to be genuine by federal decree. From there the method of detecting counterfeits can be developed. So far, so good. But suppose that instead of separating money into piles of real and fake, we have to separate our beliefs into two piles: true and false. As with currency, the bad beliefs tend to poison the good ones, so we need to identify, and throw away, the bad ones.

We don't make the truth like the Bureau of Printing and Engraving makes the money, so there is no independent guarantor of what's true. For

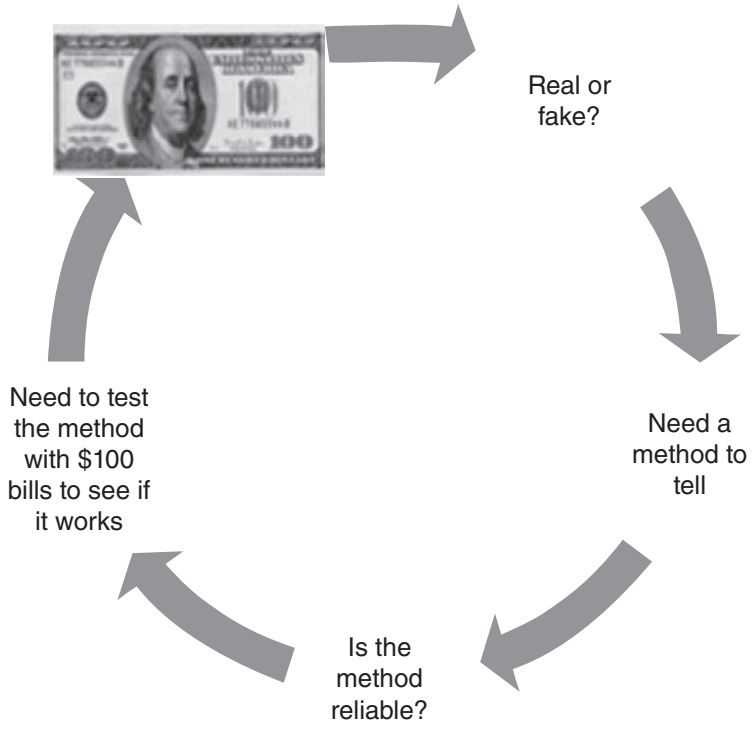

Figure 7.3 The wheel of money

any belief you care to offer, we can ask whether the belief is true or false. To figure it out we'll need a method of telling the true from the false. Of course, we want a trustworthy method that will reliably give us accurate results. Yet the only way we'll know whether our method is a reliable one is if it is correctly separating true beliefs from false ones, which requires that we *already* know which are which.

We are caught in the circle again (Figure 7.4). This time with no way 7.88
out. At least, there does not seem to be an escape that does not merely assume something for which we have no evidence. As we saw earlier, some philosophers, like Locke, hold that all of our knowledge ultimately comes from our senses and that the scientific method of experimentation and observation is the only way to get at the truth. Others maintain that it is through reflection on our own ideas and the analysis of concepts and their relations that we come to have knowledge. Both of these groups are methodists—they assume that they have the right method to separate truth out from error and then apply their belief-sorting strategy. Other philosophers claim that there are some facts that we know for certain (for example,

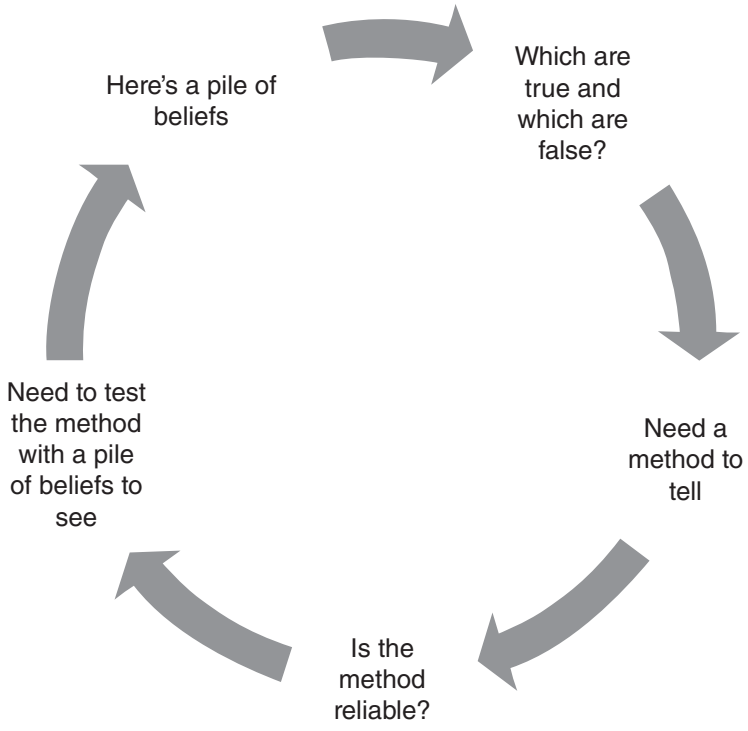

Figure 7.4 The wheel of belief

that you exist, that you have a body, that you are near the planet Earth, etc.) and, given that we know these things to be true, we can then figure out the best way to arrive at more truths. In other words, they are particularists. G. E. Moore was a particularist.

7.89 The problem is that both methodists and particularists simply assume that they have an answer to one of the stops on the wheel. That approach might have worked in the currency case, since the government creates the authentic bills, but, as we have already noted, it doesn't work here. Both methodism and particularism beg the question, that is, they assume the very thing that needs to be proved—a logical no-no. Yet it seems that there is no alternative except radical skepticism: since we cannot break out of the circle except by a sneaky, logically illegitimate move, we are never able to tell whether our beliefs are true or false. Genuine knowledge isn't possible.

7.90 Maybe radical skepticism isn't any better off than methodism or particularism, though. Here's why. The skeptic's position is:

> *The Wheel*: We can't know whether a belief is true unless we have some method to tell if it is true, but we can't know whether the method is a good

one unless we already know if it produces true beliefs. Each is logically prior to the other. Therefore we can't know either thing, and knowledge is impossible.

The skeptic is claiming that it is better to believe The Wheel than accept either methodism or particularism. If the skeptic claims to *know* that The Wheel is better, then the skeptic is offering The Wheel as an item of knowledge. If the skeptic is offering The Wheel as an item of knowledge, then the skeptic is a particularist. Therefore, the skeptic is a particularist. Or we can spin things around the other way. We can show that the skeptic is a methodist if she defends The Wheel by appeal to a general principle like this one: you can't know anything by **begging the question**.[33] This sets up nonquestion-beggingness as a requirement for knowledge, and is thereby a form of methodism. If the skeptic is *not* claiming to know that The Wheel is true, then there is no reason to fear it. Why should we worry about The Wheel if even its defenders don't claim to know it?

In the end, it might be that skepticism is not a real alternative to either methodism or particularism. Methodism vs. particularism may be the only real game in town. If that's right, then we are left with either assuming that we have some items of knowledge out of which we can build a method of inquiry, or we have to assume that we already have the correct procedure of gaining knowledge and then see what it gives us. 7.91

Now we have arrived at a deep mystery. Knowledge demands evidence. Yet we can't have knowledge unless we beg the question and accept either methodism or particularism without any evidence. We're therefore compelled to build our knowledge on something inherently and essentially impervious to evidence—surely a powerful motivation to accept extreme skepticism. Yet the claim of extreme skepticism, that The Wheel must be true, may not be an authentic alternative after all, as we've just seen in the last two paragraphs. Knowledge rests on paradoxical foundations indeed. 7.92

Annotated Bibliography

Chisholm, Roderick M. (1982) *The Foundations of Knowing* (Minneapolis: University of Minnesota Press). This collection contains Chisholm's explanation of the problem of the criterion, easily the best discussion of the problem in modern times.

Clifford, W. K. (1877) "The Ethics of Belief," *Contemporary Review* 29, full text available at www.infidels.org/library/historical/w_k_clifford/ethics_of_belief.html, accessed May 21, 2012. Clifford's classic defense of high evidential standards and why you should never form beliefs on insufficient evidence.

Descartes, Rene (1641) *Meditations on First Philosophy*, full text available at www.earlymoderntexts.com/de.html, accessed May 21, 2012. Descartes's most widely read work in epistemology, in which he worries about evil demons, dreaming, and presents the daunting problem of skepticism.

Gettier, Edmund (1963) "Is Justified True Belief Knowledge?" *Analysis* 23, 121–123, full text available at www.ditext.com/gettier/gettier.html, accessed May 21, 2012. A very short and lucid article with an incredible impact on subsequent epistemology. Gettier convinced everyone that the answer to his title question is "no."

James, William (1897) *The Will to Believe and Other Essays in Popular Philosophy* (New York: Longmans, Green). James's response to Clifford, and his defense of lowering one's standards for belief and taking some chances to gain more truths.

Littlejohn, Clayton (forthcoming) *This is Epistemology: An Introduction* (Oxford: Wiley-Blackwell). A more in-depth introduction to the theory of knowledge which examines perception, induction, testimony, justification, the analysis of knowledge, and how we might gain moral and religious knowledge.

Locke, John (1690) *An Essay Concerning Human Understanding*, full text available at www.earlymoderntexts.com/loess.html, accessed May 21, 2012. Locke's most famous work, in which he defends the view that all our knowledge comes from the evidence of the senses.

Maines, Rachel P. (1999) *The Technology of Orgasm* (Baltimore: Johns Hopkins University Press). The only published history of the vibrator, and a fascinating, scholarly look at one branch of medical quackery.

McCoy, Bob (2000) *Quack!: Tales of Medical Fraud from the Museum of Questionable Medical Devices* (Santa Monica, CA: Santa Monica Press). A constantly entertaining, yet very sobering, look at quack medical devices, products, and practices. Lots of excellent photos and old advertisements. There is a terrific chapter on radium "cures."

Moore, G. E. (1959) *Philosophical Papers* (London: George Allen and Unwin). This collection contains Moore's "A Defence of Common Sense," full text available at www.ditext.com/moore/common-sense.html, accessed May 21, 2012, and "Proof of an External World," his attempts to refute skepticism.

Pritchard, Duncan (2002) "Recent Work on Radical Skepticism," *American Philosophical Quarterly* 39, 215–257, full text available at www.philosophy.ed.ac.uk/people/full-academic/documents/RecentWorkonSkepticism.pdf, accessed May 21, 2012. Written for professionals, but a really fine overview of a variety of contemporary attempts to overcome radical skepticism.