

Unit 8

Probabilistic Reasoning

Introduction

Just as a course on critical thinking is not the ideal venue for an exploration of values, so it is also not suited to a full discussion of probability theory and statistics. But once again, a few things need to be said—the basic foundations needed for everyday reasoning should be laid out, and once that has been done, we will consider a few strategies for thinking about probability and statistics in everyday contexts.

Learning objectives

Upon completion of this unit you should be able to:

- describe the basics of probability theory;
- describe the basics of decision theory;
- estimate probabilities by reasonable methods;
- recognize some common errors in probabilistic reasoning; and
- recognize in reasoning about probability and statistics that basic arguments involving them can be better evaluated.

Assigned reading

Robert Baron and Donn Byrne. Excerpt from *Social Psychology: Understanding human interaction* (6th ed.).

How to proceed

1. Read the study notes for unit 8
2. Do the assigned reading
3. Answer the following questions.
 - Look at the graph labelled “figure 6.16” (on p. 253 in the original pagination). It is rather odd. The graph apparently shows that malls containing 86% positively rated stores are more attractive to people than malls containing 100% positively rated stores.
 - What could explain this finding?
 - Come up with several possible explanations, and decide which is most likely.
4. In the section entitled “Focus on Research: The Cutting Edge,” Baron and Byrne discuss research that seems to show that similar personalities contribute to attraction, just as similar attitudes do. But there is something deeply wrong with this research, at least, as described. What?

5. Once you have thought about these questions, submit your answers by clicking on **course discussion** on the course web site and compose your message. Fellow students will give feedback, and perhaps ask questions. In addition, your instructor will respond to comments from students, and may ask for clarification on student views stated in the discussion.

Study notes

Absolutely basic probability theory and Decision theory

This course manual is no place to deal with a proper introduction to probability theory or decision theory. However, you should know a few basic facts, with beginning with probability theory, and its application to critical thinking.

Probability theory

There are two ways in which probabilities can enter into a discussion of rational belief change. The first way is fairly straightforward: I might believe that there is a certain chance of something happening, and so be committed to accepting any claim entailed by that belief, such as the claim that the probability of the one thing happening is more or less than the probability of another thing happening. For instance, suppose I believe that the chance of rolling a 7 on two (six-sided) dice is $1/6$, and suppose I also believe that the chance of rolling a 6 is $5/36$. Then I ought to believe (or at least be ready to believe, if someone brings it to my attention) that it is easier to roll a total of 7 on two dice than to roll a total of 6 on two dice.

The second way in which probabilities may enter discussions of rational belief is more indirect. In addition to having beliefs about probabilities, one may also have beliefs held with less than perfect certainty but greater than perfect uncertainty. Probabilities may be helpful in modelling such beliefs. For instance, you might have a hunch that Sandra is now at the corner laundromat, and this might be modelled by attaching a certainty of $1/6$ to your belief in the model. Notice that, in this case, the probability isn't one you believe. Instead, the probability is a measure of your confidence in the non-probabilistic thing you do believe.

Epistemic rationality, of course, requires that your probabilistic beliefs not contradict or be inconsistent with each other, and that your senses of certainty cohere with each other. This is just an application of what was said earlier about rationality in general to the subject of probability in particular. But there are a few extra, helpful things that can be said about reasoning with probabilities, and they might as well be said here.

As everyone knows (or ought to know), if you think that the chance of one event is A/B , and the chance of another event is C/D , and the two probabilities are independent of one another, then you rationally ought to think the chance that both events will happen is AC/BD . For example, if you know the chance of rolling a six on one die is $1/6$, then the chance of rolling two sixes on two dice is $1/6 \times 1/6$, or $1/36$. If the probabilities are not independent, of course, then nothing so simple applies. For instance, suppose I think there is a 50% chance that my boss will be annoyed with me today, as a general rule of thumb, and that there is a 75% chance that I will get assigned pointless work, as a general rule of thumb. Still, if I also think

that, the more annoyed my boss is, the more likely I am to be assigned pointless work, then I cannot simply multiply the two figures to determine the chance that the boss will be annoyed *and* I will be assigned pointless work. In this case, simply multiplying the chances gives too low a probability, since boss annoyance increases the likelihood of pointless work assignment.

In addition to understanding how to calculate the probability that events A and B will *both* happen, you should also understand how to think about the probability that *either* A or B will happen. Suppose you are worrying about the safety of a nuclear reactor. You have reports that tell you that the chance of an earthquake breaching the reactor is 1/100,000 per year, the chance of a failure of all six heat-release mechanisms (which would cause an explosion) is 1/250,000 per year, and the chance of a terrorist attack of a sort causing a leak of radiation is 1/150,000 for the next year. Furthermore, these are the only dangers for the next year with a likelihood worth considering. Given this information, what is the chance of some event causing a leak of radiation within the year? You would not multiply these probabilities—that would give you the probability of all three disasters occurring at the same time! (Assuming that the probabilities are independent, which seems reasonable.)

To calculate the answer, you need to take advantage of a simple fact about logic and probability: the chance of A or B or C happening, if all three are independent, is the same as 1 minus the chance that not-A and not-B and not-C will all happen. In our example, the chance of one of the three disasters happening is equal to 1 minus the value you get by calculating the odds that there will not be an earthquake *and* there will not be a complete failure of the cooling system *and* there will not be a terrorist attack. This is $1 - \frac{99,999}{100,000} \times \frac{249,999}{250,000} \times \frac{149,999}{150,000}$, or roughly $1 - \frac{2}{100,000}$. In general, the chance of A or B or C or... happening, if A, B, C, etc. are independent, is best approached by thinking about the chance of avoiding A and avoiding B and avoiding C and..., which is equivalent to its *not* being the case that A or B or C or... And once you know the chance that *none* of A or B or C or... will happen, you just subtract that chance from 1 to find the chance that at least *one* of them *will* happen.

In thinking about the probability that a nuclear reactor accident will happen—or about other probabilities that involve the same abstract structure—people are often tempted to simplify things. If the odds of three bad things are 1 in 100,000, 1 in 250,000, and 1 in 150,000, and I want to know the odds of at least one happening, can't I just add up the three chances? Definitely not! To see this, just imagine the same sort of reasoning applied to flips of a coin. What is the chance of getting heads on either the first flip or the second? $\frac{1}{2}$ plus $\frac{1}{2}$ equal 1, but we all know perfectly well that the odds are less than perfect of getting a head on at least one of two flips. Of course, the right thing to do is the hard work: what are the odds of avoiding a head? $\frac{1}{2}$. To avoid heads twice, you have to get lucky twice, so there is a $\frac{1}{4}$ chance of avoiding heads twice. $1 - \frac{1}{4}$ is $\frac{3}{4}$, and so there is a $\frac{3}{4}$ chance of not avoiding heads twice—i.e., of getting heads at least once.

Decision theory

Decision theory is the other topic to discuss in this section. So far we have only considered probabilities as they figure in epistemic rationality, but probabilities are

also quite significant in practical rationality. Suppose you can take one of two courses of action: you can do A, which might get you one thing you want, or you can do B, which might get you something else you want. Suppose A has an 80% chance of working, and B has only a 10% chance of working. Should you do A or B? This question cannot be answered until more information is provided. Specifically, we need to know how badly you want each of the two possible outcomes. If action A is aimed at getting a drink of water, and action B is aimed at saving your life, and you have to choose between them, it hardly matters how low your chance of success is—the rational thing to do will be action B! No matter how low the odds, it is worth trying to save your life (generally) rather than get a drink of water, no matter how good the odds are of your getting the drink.

One standard way of weighing the rationality of various actions is to assign numerical values to various outcomes, representing their relative desirability, and then to assign probabilities to various actions, representing the likelihood that the action will lead to the outcome. Having done this, one can calculate the *expected utility* of any given action—the value of the outcome aimed at, multiplied by the probability of actually attaining it. The rational action to perform, then, is the one with the highest expected utility. In our example above, it was imagined that an extremely high value was attached to survival, and a much lower value was attached to getting a drink of water just then.

- | | |
|----------|---|
| Action A | 80% chance of getting, say, 10 units of value = 8 units of expected utility. |
| Action B | 10% chance of getting, say, 1,000,000 units of value = 100,000 units of expected utility. |

Our example had another simplifying assumption: that action A had a 0% chance of saving your life, and action B had a 0% chance of getting you a drink of water just then (a pretty reasonable pair of assumptions). But this need not always be the case. Consider a different scenario:

A politician wants both to decrease spending and to get re-elected. She evaluates two possible actions: a small cut in spending and a large cut in spending as follows:

Value of a small cut in spending: 20; of a large cut: 40; of getting re-elected: 50.

The chance of getting re-elected after a small cut: 70%; of getting re-elected after a large cut: 25%.

Given these valuations and estimates of probability, our politician can now calculate which action it is practically rational to take:

- | | |
|-----------|--|
| Small cut | 100% chance of getting 20 units of value + 70% chance of 50 units of value = 55 units of expected utility. |
| Large cut | 100% chance of getting 40 units of value + 25% chance of 50 units of value = 52.5 units of expected utility. |

Given these calculations, and the assumption that these are the only relevant factors to weigh, it would seem that our politician would be rational to choose the

smaller cut in spending: this maximizes the expected satisfaction of her desires for both fiscal restraint and getting re-elected.

Something the above example demonstrates is just how difficult it is to perform decision-theoretic calculations under real-life conditions. If in our example our politician had attached a 30% probability to getting re-elected after a large spending cut, rather than a 25% probability, then the two utility calculations would have come out exactly equal—and how could our politician be sure that her chance of re-election after a big cut would be 25% rather than 30%? Likewise, how could our politician be sure that getting re-elected means so much more to her than being the person who really cuts the budget back—maybe she would rather be a one-term Margaret Thatcher than a three-term Brian Mulroney. Attaching precise numerical values to one's desires and probabilistic expectations is usually very difficult, and thus decision theory is not a universally useful tool for decision-making.

Nonetheless, when one *can* put precise measures on certain outcomes (as when comparing financial outcomes), and precise probabilities, it can be very helpful.

Probabilistic guessing

It is not always possible to put a precise number on one's estimate of how likely something is. Precise numbers usually come from precise information, and in general we have precise information only about a few things: statistics we read in the newspapers, odds of throwing two sixes on two dice, medical statistics we get from doctors, and so on. For everything else, we are left guessing in very imprecise terms. And yet, somehow we manage to do just fine most of the time. The explanation is that, for many purposes, we do not need precise probabilistic or statistical information in order to form reasonable beliefs and make successful decisions; we can get by with a lot of uncertainty.

Consider, for instance, the decision millions of people make every day to drive a car somewhere. Very few of these people know what the odds are of being involved in a serious accident as a result of driving a car. Yet most of these people make reasonable decisions to drive even without this information. That is because, although they lack the precise information on how common car accidents might be, they know from experience and observation that car accidents are not all that common. They know that car accidents resulting in lasting injury are so uncommon that most people who drive cars get through an entire life without having such an accident. As a result of having this very imprecise, unquantified information, these people can make a rational choice to drive. All they need is the imprecise information they already have.

Compare this case to the case of someone considering a career flying a crop duster (an airplane that sprays chemicals on crops). The person considering work crop dusting might be uncertain about the precise risks of a career in crop dusting, but that does not mean that the potential crop duster pilot would be just as rational in choosing to fly her airplane as a potential car driver would be in choosing to drive his car. Not all uncertainty is the same! In the case of the car driver, he knows his experience is a good guide to the safety of driving a car. Imprecise, yes, but very informative anyway. In the case of the crop duster pilot, she is more likely to be ignorant of the risks of crop dusting. The risks of crop dusting, and the consequences of accidents, are not as widely publicized as the risks and

consequences of driving a car, and most people know far more car drivers than crop dusters. As a result, the person considering work as a crop duster pilot should probably check out the actual statistics before making a career choice, even though the car driver does not need to do the same.

Examples like these show that (a) people can sometimes make rational probability judgments without precise probabilistic information; but (b) people cannot always make rational probability judgments without probability information. Making decisions under uncertainty, like so much else in life, is simply a judgment call.

Tips for estimating probabilities

Like every other type of judgment call described in this course manual, however, imprecise judgments of probability are something everyone can do better or worse, and is something one can get better at by practice. Here are some tips for making the best probability estimates that you can.

1. Try to compare probabilities. Although we do not know the statistics for catching colds or breaking bones, we do know that it is much more common to catch a cold than it is to break a bone. If I need to do some reasoning that involves the probability that someone on my bowling team will break an arm or a leg, I can begin by seeing that the chance of that happening is definitely less than the chance of a person getting a cold.
2. Take a guess at a probability and consider what the consequences would be. Suppose I guess that a nuclear reactor has a 1% chance of causing a major radiation leak every year. Then it follows that, if the world contains one hundred nuclear reactors, the chance of a major radiation leak each year is close to 50% (remember: the relevant calculation is $0.99 \times 0.99 \times \dots$ fifty times). But if the odds were that high, then surely I would be hearing about a major nuclear accident every few years; yet I haven't heard about one in ages. And I know the world has at least one hundred nuclear reactors so my estimate of the dangers must be too high.
3. Use related facts you know to compare your guess to known probabilities. Imagine that I am wondering how common murders are in Los Angeles, and that I guess one person in a thousand gets murdered every year. Well, I know that Winnipeg has about 650,000 people. Every year in Winnipeg, around 20-30 people get murdered. So that means the murder rate in Winnipeg is about one in twenty thousand or one in thirty thousand. This suggests that perhaps I have over-estimated the frequency of murders in Los Angeles.

Practice exercise 8

For each of the following statistical or probabilistic claims, use your own knowledge of the world to determine whether it is true or false. Don't consult an expert reference source; just use your head. If you aren't sure, try using the above tips.

- a. Over half of Canada's population is under the age of thirty.
- b. If you spend your whole life in Manitoba, there is less than a 50% lifetime chance that you will face an armed robber who demands your money.

- c. The number of natural blondes in the world is less than half the number of natural non-blondes.
- d. Most parking lots reserve less than 10% of their spaces for handicapped parking.
- e. One out of every five or six NHL hockey games results in a fairly serious permanent injury to a player.
- f. Smoking exactly one cigarette is more dangerous to your health than being a member of the U.S. armed forces for one year.
- g. At least one out of every hundred scientists makes a major discovery of the sort that makes the world news or gets a Nobel Prize, each year.
- h. In major cities, people paid to work in restaurants must outnumber people paid to work in libraries by at least seven to one.
- i. The odds of having a car stolen in Winnipeg are at least ten times better than the odds of winning a million-dollar (or more) lottery.
- j. Canadians and Americans are both far more likely to die of cancer than of terrorism.

Common faults in probabilistic and statistical arguments

In the previous unit, we looked at some common traditional fallacies. In this unit, we will look at a few more common mistakes made in argumentation. Unlike the traditional fallacies, most of these common mistakes have not been described by philosophers, and most lack Latin names. Probably this is because the theory of probability was not developed until the 1800s, and so the ability to even describe errors in probabilistic reasoning has only recently come into existence. Yet these errors are, if anything, even more worthy of discussion, because they are even more common than more standard fallacies, and just as tricky to detect.

Availability biases

As the previous section suggested, it is often possible to come up with a rough estimate of how likely or unlikely something is based on one's general knowledge of the world. However, estimating likelihood is tricky, and it is easy to make mistakes. One of the most common is to estimate the likelihood of something based on how available examples of it are off the top of your head. For example, imagine that Edna is trying to guess how common it is to get breast cancer. She thinks for a moment, and cannot recall anyone she knows having breast cancer. So she concludes that it is not a common disease. That is probably a failure of critical thinking on her part. It is widely known that cancer in general is a common disease of older age, and that breast cancer is a common form of cancer among women. If Edna knows, somewhere in the back of her mind, that these things are true, then she ought to estimate that breast cancer is not so uncommon after all, though of course much less common than the flu, say. The fact that Edna could not think of anyone with breast cancer tricked her into thinking the disease itself was uncommon. In other words, the fact that she did not have any examples available to her off the top of her head convinced her that the objective occurrence of breast cancer was rare, when it should not have.

Relying on how available information is to you in order to judge likelihood is not always foolish, of course, any more than it is always irrational to make a quick generalization or argue by pointing out a slippery slope. Sometimes, a person is so knowledgeable about a topic that availability correlates directly to likelihood. For instance, the leading skin disease expert in Canada can probably make a reasonable comparison of the frequency of leprosy in Canada to psoriasis just by thinking how commonly the expert has heard of each. But for most of us, most of the time, estimating likelihood by how easy it is to think of examples is likely to be a bad idea. Instead, it is better to rely in the estimation techniques of the previous section.

Misunderstanding randomness

Suppose a friend has flipped a coin five times, and comes up with a sequence of heads (H) and tails (T). The friend has also just made up an imaginary sequence of heads and tails. Your job is to guess which sequence is real, and which is imaginary.

First sequence: HTHHT

Second sequence: TTTTH

Which sequence would you guess is the real one, and which the made-up one? Many people would say that the first sequence looks more random, and so more real. But this is a mistake. Both sequences are equally likely (or unlikely) to occur in a sequence of five coin tosses. After all, to get the random-looking HTHHT sequence you first need to throw heads, then tails—not a second head—then switch back to heads for exactly two throws—not more or fewer—and then end with heads. That is just as unlikely a sequence as throwing four tails in a row, followed by a head. In fact, any *exact* sequence is equally unlikely.

For the same reason, it is just as unlikely to be dealt a set of five cards that includes the five of clubs, the nine of hearts, the jack of hearts, the queen of hearts, and the two of spades as it is to be dealt all four aces and the jack of spades. It is a lot better for a gambler to be dealt the second hand, of course, and so luckier. But both hands are equally unlikely. (If you are not convinced, just imagine sitting down with a shuffled deck of cards and dealing out hands until you got the *exact* cards I described for the first hand. It would take a long, long time...)

People sometimes forget that looking uninteresting is not the same as being likely, or that being interesting is not the same as being especially unlikely. The exact number of blades of grass growing in a yard is *amazingly* distinctive, I suspect. If you counted the number of blades in every yard in a city, you would not find two yards with the exact same number, in all likelihood. Yet the normal response to this statistical distinctness is nothing. If a yard had 1,819,223 blades of grass in it, you would probably find that fact completely dull. But think: what are the odds that that particular yard would have exactly that many blades of grass? It is absolutely amazingly unlikely. But it does not seem strange. If, on the other hand, you found a yard that had exactly 2,000,000 blades of grass, you would think it was amazing. But it is in itself no more amazing than finding a yard with 1,819,223 blades of grass in it. One number “looks random” and the other one does not, so we naturally think that one number is pretty likely, while the other number is quite unlikely. But in fact, each number is equally unlikely.

Mistaking boringness for randomness is a common error when thinking about what is random and what is not. And it does not only occur when thinking about individual probabilities. It also occurs when thinking about overall patterns generated by supposedly random processes. Suppose again that a friend has thrown a coin five times and made up an imaginary sequence as well. But this time suppose that the friend does not tell you the exact sequences. Instead, suppose that you just get information about the overall distribution of heads and tails:

First set: 3 heads, 2 tails

Second set: 5 heads, 0 tails

Now, ask yourself again: which is more likely to be the real result of throwing a coin five times, and which is more likely to be made up? (Assume your friend is not trying to play mind-games with you, to simplify things!) Many people would say that the first set is more likely to be the result of a random process of flipping a coin than the second set. This time, most people would be right. When flipping an unbiased coin five times, it is rather more likely to throw a total of three heads and two tails than to throw a total of five heads. (There are ten ways to throw the first set, but there is only one way to throw the second set, out of the thirty-two possible ways of throwing the five tosses. So obviously the first set is more likely.) Our instincts about randomness are quite right in this case. Yet our instincts can easily be too strong. People have a tendency to think that a random set of coin tosses should have an overall distribution that is very close to fair, and that in general in any random process the overall distribution should be quite close to fair. However, this is quite an exaggeration. For instance, compare the probabilities of the two following events:

First event: flipped a coin six times, came up with 3 heads and 3 tails

Second event: flipped a coin six times, came up with something else

Which is the more likely event? As it turns out, one is much more likely to come up with an uneven distribution of heads and tails than an even one, even though an exactly even number of heads and tails is what seems most random.

The real nature of randomness is that it is not consistent, not even in its appearance of randomness. Arguments that forget this are prone to being bad arguments.

The gambler's fallacy

An especially well-known confusion over the nature of randomness is found in the gambler's fallacy. In the gambler's fallacy, one infers from the fact that a random process has been recently skewed in one direction to the unjustified conclusion that the random process will try to "correct" itself by tending, for a time, in the opposite direction. For instance, suppose you are betting on whether a coin will come up heads or tails. In the last three throws, the coin has come up heads all three times. If you would decide that now is an especially good time to bet on tails, since "the balance of probability must be restored" or something of the sort, then you would be making a bad inference, and committing the gambler's fallacy. After all, coin flips are completely random, and so have an equal chance of coming up heads or tails on every throw, no matter what a recent series has been.

Like most fallacies, there are closely related arguments that are good arguments. If I have reason to doubt that the process I am dealing with is completely random, then I can use information about what has happened before to predict a partially random process in the future. For another example, if I am playing several games of blackjack with a single deck of cards (and without shuffling between deals), and I notice that although high-scoring cards make up 16 of 52 cards (aces not included), none have them have been played so far and already half the deck is used (26 of 52), then I ought to be sure that a high percentage of the coming cards (16 of 26) will be high-scoring cards.

In general, we can call it a gambler's bet that, after a run of unlikely results, some other event will occur which ends the unlikely events or in some other way restores the appearance of statistical normality. A gambler's bet is reasonable when the person making the inference knows that the process producing the results is not purely random, and knows that an unlikely run of events has consequences that make it likely or inevitable that statistical normality will be restored in the end. It is unreasonable when the person making the bet knows that the process producing the events is purely random, and that each new event is uninfluenced by the past. It is only in this second case that a gambler's bet becomes the gambler's fallacy.

Base rate errors

Suppose the reliability of a certain medical test is fairly high, and it gives you bad news. Does that mean that you are probably sick? Not necessarily. Think about the following example, adapted from the work of psychologists Daniel Kahneman and Amos Tversky.

In a certain city there are two taxi cab companies, Pro Cabs and Duffer Taxis. Suppose that 85% of the cabs on the streets at any given time belong to Pro, on average, and 15% to Duffer (Pro is the larger company). Suppose further that a witness to a crime saw the criminal exit from a taxicab late at night, and swears that it was one belonging to the Duffer company. Being careful, the police test the witness using videotape of cabs seen at night, and discover that the witness is 80% reliable—right 4 times out of 5.

Does it follow from the high reliability of the witness that there is an 80% chance that the criminal actually got out of one of Duffer's taxis? Far from it. In fact, the odds are greater that the criminal got out of a Pro cab that the witness mistook for a Duffer. This can be seen once you realize that you have to think of all the possible scenarios in order to determine which is most likely. The criminal could have taken either taxi, and the witness could have made a correct or an incorrect identification. This gives us four possibilities.

1. Criminal takes Pro Cab (85%) and witness makes correct identification (80%)
2. Criminal takes Pro Cab (85%) and witness makes incorrect identification (20%)
3. Criminal takes Duffer Taxi (15%) and witness makes correct identification (80%)

4. Criminal takes Duffer Taxi (15%) and witness makes incorrect identification (20%)

Now, since the witness claims to have seen a Duffer Taxi, possible scenarios (1) and (4) are eliminated—these are the scenarios in which the witness would have claimed to see a Pro Cab. But scenarios (2) and (3) are both possible. The chance of (2) would seem to be 17%, and the chance of (3) would seem to be 12%. In other words, it is more likely that the witness made a mistake about a Pro Cab than it is that the witness made a correct judgment about a Duffer Taxi, even though the witness is 80% reliable! Now, these probabilities are not correct. Since the witness *does* claim to have seen a Duffer Taxi, they should add to 100%. The absolute actual chance of (2) is equal to 17 divided by $17 + 12$ —i.e., by the number of (2) scenarios divided by the number of possible scenarios, and similarly for the chance of (3). Thus the chance of scenario (2) is actually 59%, while the chance of scenario (3) is 41%. But we do not even need to make this final calculation to see that scenario (2) is more likely than (3). The moral of the story is, when one is given evidence for a claim with a certain reliability, one cannot know how likely the claim is to be true without also knowing the *base rate* for events of the type under consideration—the probability in general of the event happening.

The significance of base rates for medical tests should now be obvious. Even in the event that a very reliable test indicates a grave danger, it cannot be inferred that one's life is in grave danger—that inference must take into account the base rate of the danger in the population. Following the above example, we can see that an 80% reliable test for an illness found in less than 15% of the population does not give us stronger reason to think that the patient has the illness than that the patient does not have the illness. Of course, there is still reason to take the diagnosis seriously, but it would not yet be time to sell your worldly goods and take a final trip to your favourite beach.

Neglecting base rates can cause errors in probabilistic thinking in all sorts of contexts, not just medical and legal ones. What all these contexts have in common is that there are two important pieces of probabilistic or statistical information available to answer some question: one piece that gives background probabilities, and one piece that adds more specific information. To consider only the more specific information, and ignore the background probabilities, is to make a mistake based on neglecting base rates.

Let's try one more example. Suppose that you have baked a dozen chocolate chip cookies, and that you notice two are missing. The obvious suspects are your 6-year-old daughter and your 66-year-old father. Both deny taking the cookies, but your father has some cookie crumbs on his fingers, while your daughter does not. Should you suspect your father of being the cookie thief? Not necessarily. If your father, like many older people, finds cookies too sweet, while your daughter, like many younger people, loves cookies, then the base probability of your father taking the cookies is much lower than the base likelihood of your daughter doing so. That might justify overriding the physical evidence. After all, your father may have assisted your daughter in cookie theft, then helped her to eliminate the evidence from her face and hands while leaving traces on his own. In this case, you do not know any of the probabilities involved precisely enough to do a calculation to

determine which hypothesis is more likely, but you probably have a good enough estimate of the probabilities to see that your daughter should not yet be cleared of suspicion.

There is one further point to make about base rates. Use of base rate information is clearly rational in all the cases above. However, using base rate information is always a matter of judging particular events, not solely on the details you know about that particular event, but also on what you know about events of that kind. Using base rate information thus involves a kind of pre-judgment. A parent who suspects the crumbless daughter just as much as the cookie-crumbed father as a result of known base rates is clearly pre-judging the daughter, in part. If the parent is rational, then the pre-judgment will only be partial: it will only be probabilistic. But there will be a pre-judgment before the physical evidence is ever taken into account, and that pre-judgement will weigh against the physical evidence. This might sound like the beginnings of prejudice. And surely prejudice is not rational!

To see the point even more clearly, consider the position of border guards. Base rate information suggests that young men with long or coloured hair, tattoos, and multiple piercings are more likely to be smuggling illegal drugs at border crossings than conservative-looking older women wearing business suits. But a particular young man might show no more signs of illegal drug use than a particular older woman: both might have normal speech, pupil dilation, sense of time, and so on. What should the border guards believe? So long as they have good reason to believe their information about base rates (and in this case they do), and so long as they do not have any strong particular evidence to the contrary, the border guards should believe that the younger man is more likely to be smuggling illegal drugs than the older woman. This is not to say that the border guards should be confident that the younger man is smuggling illegal drugs. Far from it! Actual base rates for smuggling illegal drugs are well below 50% even for young men who use illegal drugs, much less for young men who have an “alternative” appearance. But it *is* to say that if forced to bet, the young man would be a better bet than the older woman for being a person smuggling illegal drugs. Now, what should the border guards *do*, given this belief? That is another question altogether.

What border guards in fact do is stop and search many more young alternative men than older conservative women, but they let many of the young alternative men through without searching them and stop the occasional older conservative woman for a search too. That is, they stop people and search them based in part on base rate information. Of course, border guards also stop people and search them based in part on specific information about individuals: people who seem especially nervous, people who seem intoxicated, and so on are likely to be singled out for searches. But even when none of these signs are present, people are still searched, and these searches are more likely to be performed on the people whose appearance correlates better with illegal border activities. Is it *morally right* for border guards to do this? This is a difficult question. It seems in part to require us to think about the harm caused by smuggling, the harm done to people who are searched even though they are innocent, and about the unfairness of searching people based on superficial statistical evidence rather than more substantial evidence particular to that individual. I will not try to answer this question here, but take a moment to consider it for yourself. Clearly *some* actions are morally wrong

when guided only by base rate information. (Arresting every long-haired man to cross the border would clearly be wrong, for example, even if the base rates for long-haired men being criminals was higher than the average for the population.) Is this *particular* policy wrong? What do you think?

Bonus #2: Truth

A final topic worth considering is that of truth. This section is not a required one for credit, but it addresses a topic that many undergraduate students find confusing.

What is truth? This is a very old philosophical question, but there is an equally ancient answer that has seemed correct to most people, philosophers and non-philosophers alike. For a statement or belief to be true is for what is stated or believed to be the case. For the statement “Halifax is the capital of Nova Scotia” to be true is for it to be the case that Halifax is the capital of Nova Scotia, and not Sydney or Yarmouth, for example. For the belief that Ben likes hot chocolate to be true is for it to be the case that Ben really does like hot chocolate. In general, beliefs and statements say that the world is a certain way, and for them to be true is for the world to be the way they say it is; if the world is not “as advertised” by the statement or belief, then the statement or belief is false. This is called the *correspondence theory* of truth. It is quite commonsensical, and although there are philosophical arguments against it, there are philosophical arguments against almost every interesting claim, and we will not worry about them in this manual.

Statements and beliefs (or, according to some philosophers, propositions) are the *only* things which are either truth or false. If you say “I am hungry” or think “Elle a soif,” you have said or thought something which may be true or which may be false. On the other hand, if you say “Let’s go” or wonder “Can I really see myself staying at this job for the next ten years?,” your utterance or state of mind is neither true nor false, because commands and questions are just not the sorts of things that can be true or false. This is pretty obvious.

Interestingly, it seems that not all statements (or beliefs) are true or false. Being a statement or belief is *necessary*, if the concepts of truth and falsity are going to be relevant, but being a statement or belief does not seem to *guarantee* that the concepts of truth and falsity are relevant. Consider the following two examples:

1. When Hitler conquered the Soviet Union, he had Stalin executed.
2. These contraptions called “backless chairs” are really a kind of stool, because they don’t have backs.

Both are puzzling. In the case of (1), just try saying that it is either true or false. If you say that it is true, then that sounds like Hitler had Stalin executed, which is obviously not correct. But if you say that (1) is false, that makes it sound like Hitler conquered the Soviet Union but spared Stalin’s life, which is also wrong. Statement (1) seems to have a *false presupposition*—that Hitler conquered the Soviet Union. It does not seem that it can be true or false without that presupposition being true, but the presupposition itself is false, and so statement (1) can be neither true nor false. In the case of (2), the problem is different. The problem there seems to be that “chair” and “stool” are vague terms, and do not clearly include or exclude backless chairs. Most of our terms (and the corresponding concepts) are vague in some respect—what shade marks the precise border between red and orange, for

instance, or which hair is it that makes a truly hairy armpit hairy? There seems to be no answer. When one tries to use an imprecise term in a way that would only make sense if the term were more precise, then one is likely to say something that is neither true nor false.

After taking into account the few sorts of exceptions there are, however, it seems safe to say that all other statements, thoughts, beliefs, claims, factual pronouncements, and the like (but not commands, questions, desires, or their ilk) are either true or false.

An important thing to understand about truth is that truth is not relative, in the usual sense of the claim “truth is relative.” If some statement is true, then it is true. It is not true for me but false for you, because that would mean that the world both would and would not be “as advertised,” and that is impossible. Think of the statement “A meteor of over 1kg in mass is about to strike the building we are in.” If this statement is true, then we are about to die; if it is false, we can probably relax. A person who insists “that’s not true for me!” would seem not to understand that, if the statement is true, he is going to die, regardless of what he considers to be “true for him.” The world of causes and effects does not care about what is “true for me” or “false for you,” but only about what is true or false for everybody—about what is really the case.

Now, it is true that sometimes people say things that suggest that truth is relative, and such statements are not entirely crazy. So, what could people mean by such statements? What they usually mean is not that truth is relative, but one of a variety of other things that sound similar but have importantly different meanings. The slogan of the truth-is-relative crowd could be put “that’s true for you but not for me.” What follows are various interpretations of this slogan which give it non-crazy meanings.

Perhaps most commonly, people say “that’s true for you but not for me” and mean “that’s true in the sense you intend it, but false in the sense I intend.”

E.g.: Doug says “David Suzuki is incredibly famous” and Arnold Schwarzenegger replies “that might be true as far as you’re concerned, but not where I’m from.” Here, Schwarzenegger might be fairly interpreted as saying Suzuki has fame in Doug’s circle of acquaintances, or in Canada, but is not famous in Hollywood or internationally.

Other times, people mean something like “if I said that, it wouldn’t be true.”

E.g.: Grigori says “I like pistachio ice cream” and Tatiana replies that “that’s not true for me,” meaning “I wouldn’t say that, because I don’t like pistachio ice cream.”

Perhaps a person saying “that’s true for you but not for me” means something like “I have a different opinion, but I respect your opinion also.”

E.g.: Irit says “the museum wasted its money buying a Jackson Pollock” and Yonatan replies “that may be true for you, but not for me,” meaning “you may think the museum wasted its money, and I respect your thinking so, but I think the money was well spent.”

In all of these cases, it is clear that there is no one statement that one person says is both true and false. Instead, one statement is held to be true, and a slightly different statement is held to be false. If this is kept in mind, many confusions about truth can be avoided.

One more potential confusion should also be eliminated. Sometimes it is hard (or even impossible) to determine whether or not some particular statement or belief is true or false. For example, statements such as “The first Ming emperor had a cancerous mole on his left wrist when he died” are probably completely uncheckable. No one can now know whether that statement is true or false. Does that show that the statement is neither true nor false? Not at all! What we *know* is not the same as *what is the case*. If I flip a coin so that no one can see it, it must be showing heads or tails (or be standing on its end), but no one knows which. Nevertheless, clearly, either heads is showing or tails. If I guess “Heads!” what I believe might be true, or it might be false. I do not know which until someone looks, and I might never know which (maybe a meteor will hit the coin before anyone looks), but my belief is either true or false nonetheless. Likewise with long-dead emperors and the details of their healths, and with other uncheckable claims and beliefs: just because one cannot *check*, it does not follow that no *fact* exists. This is just common sense. How you look at the world does not do much to change what is out there to see.

Epistemology and Metaphysics

The distinction between what one knows or believes and what is true is generally known as the distinction between epistemology and metaphysics. There are various traditional branches of philosophy, and *epistemology*, the study of knowledge and belief justification, and *metaphysics*, the study of what exists and how things are related (at the most abstract level), are two of them. Students (and professors, too) sometimes forget the distinction between the two branches in their work, since the two fields are very closely connected (the only things you can study directly are things that exist, after all), so it will be worth saying just a little more about it.

Suppose I tell you that Phillip is not a human being, though he looks it. You wonder why I think so, and I say he is non-human because he has a Neanderthal ancestor, while true humans are descended solely from the beings that, in Europe at least, displaced the Neanderthals. You might object to this claim along two different lines. You might say “that doesn’t matter—we call people human if they look like us, think like us, and so on, and who their 1000-greats grandparents were makes no difference.” This is a metaphysical response. It claims that I do not understand the essence of being human, the qualities that make an animal a human being. But you might also object “how can you possibly know that about Phillip? I’ve heard that the Neanderthal were so like modern humans that, after all these generations, it would be impossible to find out something like that from genetic tests, and there aren’t any genealogical records.” This is an epistemological response. It asks how I could possibly know that my claim is true.

In the above example, should you argue with me on epistemological or metaphysical grounds? It all depends on what you believe. Perhaps you believe that *Phillip* is human but there is nothing in principle to prevent someone else from being inhuman by reason of Neanderthal ancestors. Then you would want to make

the epistemological argument—“why do you think Phillip is any more likely to be descended from Neanderthals than you?” On the other hand, perhaps you think it is ridiculous to worry about distant ancestry in this way. Perhaps you even think (from his thick eyebrows and low forehead) that Phillip might have Neanderthal ancestors, but you do not think that is relevant. Then you would want to make the metaphysical argument—“who his ancient ancestors are makes no difference; most of Phillip’s ancestors are obviously human and that makes him a human too.”

Often, the difference between the facts and our beliefs about them seems uncontroversial. Be warned, though, that the distinction is not always so clear-cut. For example, some philosophers known as “logical positivists” held that if there were no *possible* observation that could be evidence against a claim, then that claim would be meaningless (or have only “poetic” meaning), and thus could not be true. So the claim “there is a completely undetectable elf in my pocket who never does anything,” which looks at first glance like a perfectly comprehensible metaphysical claim (of a rather odd sort) would be taken to be neither true nor false by the positivists. They were well aware that they were mixing epistemology (evidence) with metaphysics (the facts) and were ready to defend their view. If you decide that an epistemological claim is relevant to an apparently metaphysical claim—that what we believe is relevant to the truth of what we say—be sure that you too are ready to justify such a bold move.