30 ➛ The Base Rate

We just talked about a principle of probability—the product rule—that is easy to grasp, though sometimes also easy to forget. There is one other such principle that we need to talk about. It's usually called the *base rate.*

After much study, the police determine—never mind how—that for every 100 people who ride on an airplane, one of them, on average, is carrying drugs. They bring a police dog named Merlin to the airport to help find the culprits. Merlin sniffs every passenger who exits a plane, and he barks when he smells drugs. Merlin is highly skilled. He *never* misses drugs when they're present. He occasionally is prone to giving false alarms, but not often; if no drugs are present, he is 90 percent likely to keep quiet. Smithers steps off the plane, and Merlin barks at him. How likely is it that Smithers is carrying drugs?

I pose this question because many people get it wrong; you should very much want to get it right, because it comes up in different ways all the time. It seems at first that Smithers probably has drugs. Merlin makes mistakes only 10 percent of the time; if he barks, Smithers is 90 percent likely to be carrying drugs, isn't he? No. The chance that Smithers has drugs is about 1 in 11 (or 9 percent, if you prefer). The trick is that you have to remember not just Merlin's chance of being right but also the background chance that Smithers has drugs in the first place, which is low. Think of it this way: for every 100 passengers that come off the plane, only one of them will be carrying drugs. Merlin will always bark in those cases. But then when the other 99 passengers come off the plane without drugs, Merlin also will bark some of the time—10 percent of the time, to be exact. So out of this batch of 100 passengers, maybe Smithers is the one with the drugs, but chances are better that he's one of the 10 Merlin will accuse wrongly. Here is another way to express the point: when you hear of a 10 percent rate of error, you need to carefully ask, 10 percent of what? It might seem to mean that 1 out of every 10 people Merlin barks at will be innocent, but in fact it means that 1 out of 10 people who are innocent get barked at anyway—which is not the same thing. So to make sense out of Merlin's barks, you need to know how likely each person in the pool is to be guilty in the first place. You need to know their background likelihood of guilt; you need to know the base rate.

To make the point even more intuitive, try it like this: imagine that *one person in the world*—call him Waldo—is carrying some rare drug. In search of him, Merlin starts to sniff every person on the planet—all 6,415,000,000 of them. Merlin will bark if he finds Waldo; he also will bark wrongly 10 percent of the time. Merlin barks at you. How likely is it that he just found Waldo? Not very. On his way to finding the real Waldo he is going to set off about 641,500,000 false alarms. You're probably in that group. Hey, wait: if he barks wrongly only 10 percent of the time, doesn't that mean you're 90 percent likely to be the right man? No, and now it's easy to see why: he has a 10 percent error rate in screening people, *almost all of whom are innocent.* That background likelihood of innocence is the dominant fact here, not Merlin's rate of error.

Our airport example is similar to the case of the one man in the world. It just involves less overwhelming background numbers: a less imposing base rate, as it is called, or less imposing prior odds. The first big lesson of this discussion is simply to remember base rates and not just the probabilities in the foreground of a problem. But our second task is to figure out how to get the actual right answer in the airport problem and others like it, and without doing messy mathematical calculations that will cause most readers to run screaming for the exits. I will do my best to explain it in words. It's important, and it isn't that hard.

Basically you want to build a fraction to show how likely Merlin is to be right—1 chance in 11, or something like that. *On top you put the number of times (in 100) that the dog will bark and be right. On the bottom you put the number of times (in 100) that the dog will bark at all—rightly or wrongly.* That's what you really want to know: how the number of correct barks compares with the number of total barks. Sometimes instead of using the number of times these things would happen in 100 tries, you might use the number of times in 1,000. It doesn't matter, so long as you use the same number (say, 100 or 1,000) in all of your thinking. We'll see an example in a minute. But to stick with the airport case, the number on top is the number of times in 100 that Merlin will bark at someone who does have drugs. That's easy here; we know that one person in 100 carries them, and that Merlin always finds him. So the number on top is 1. On the bottom we put the number of times Merlin will bark, rightly or wrongly, in the course of 100 sniffs. He will bark rightly 1 time in 100, as we just saw. He will bark wrongly in 10 percent of the remaining 99 cases—which we can round off to 10. That's a total of 11 alarms by Merlin. So the odds that his barking will lead to a find are about 1/11.

Let's do another. When he applies for a green card so that he can work in the United States, Smithers is required to take a test to see whether he has AIDS. The background chance that someone from his country has the disease is 1 in 250. The AIDS test he takes catches the disease in every person who has it. It also produces false positives 4 percent of the time. (Tests for AIDS nowadays are better than that, but bear with the example.) The test on Smithers comes up positive. What are the chances that he has AIDS? At first the chances look great: the test yields a false positive only 4 percent of the time, so that means he's 96 percent likely to have the disease, right? Wrong, of course, as you now understand. We have to account for the base rate—the background unlikelihood that he has the disease. And we know how to do this. We build a fraction. On top goes the number of times the dog will bark and be right, or, here, the number of times the test will be positive and will rightly find AIDS. Since 1 person in 250 has AIDS, and since the test catches all of them, that's 4 in 1,000 (we'll use thousands here so that we can work with whole numbers). So we can put a 4 on the top of the fraction. On the bottom goes the number of times the dog will bark, rightly or wrongly—or, here, the number of times the AIDS test will be positive, rightly or wrongly, in 1,000 tries. First there are the 4 *accurate* positives that we just mentioned. Then there are the *erroneous* positives, which happen 4 percent of the time. That's 4 times in 100, or about 40 times in 1,000. ("About" is the right word because we should have removed the 4 of those 1,000 cases that do have AIDS. But we're close enough.) The number on the bottom of the fraction thus is 44: 4 correct positives and 40 bad ones in every 1,000 tests. Now we can see the odds that Smithers has AIDS when his test says so: 4/44, or 1/11, or about 9 percent—roughly the same figure as in our airport case, as it happens.[57]

In the appendix to this chapter there are some more problems like these in case you are interested in playing with them; they will show a few other complications that can arise. In the meantime, you get the idea. You can go a long way with the technique of putting the correct barks on top and *all* the barks on the bottom. That won't be enough to solve complicated problems, but it will help with the simple ones, which otherwise leave most people flummoxed. Notice a key move in all the problems: turning the percentages into stories where you imagine a hundred or a thousand run-throughs, with some number of them coming out one way and the rest coming out the other. This technique—thinking in frequencies, as it sometimes is called—helps make problems involving percentages easier to understand.[58] You have to be careful with it, because a statement of a probability may not actually amount to a claim about how often the thing has happened or will happen. When we speak of a 60 percent chance that the plaintiff was run over by one of the defendant's taxicabs, we might imagine 100 such accidents, and suppose that the defendant's cabs were responsible in 60 of them—but there may really be no other accidents quite like the one in the plaintiff's case. When we speak of probabilities in a case like that, we just are expressing our subjective sense of confidence that the thing did happen—perhaps the odds we would use in placing a bet on the matter. But thinking in frequencies still is a useful tool to make numbers easier to picture in the mind's eye.

The general family of ideas we are considering here goes by the name of Bayesian methods, after Sir Thomas Bayes, a British mathematician who lived in the 1700s. He was the creator of Bayes' theorem, which is a way of getting from prior odds to so-called posterior odds; the methods of calculation offered above are essentially shortcut versions of what Bayes offered, and take you to the same place.

Now we have some useful tools for thinking about the likelihood that something is true. But when can you use them in court? Start by thinking of a civil trial where the jurors hear some piece of testimony. They are supposed to decide whether the plaintiff's claim more likely than not is true—whether, to say it numerically, the plaintiff's story is more than 50 percent likely to be right. Maybe the new testimony the jurors hear gets them over the 50 percent bar, or maybe it doesn't; but where do they *begin*? What should they assume are the background odds that the plaintiff is right? Jurors who hear new evidence are like the police who hear the dog bark, or the immigrant whose AIDS test comes back positive: none of them can know what to make of this new information unless they know the base rate, or the prior odds of whatever it is they are trying to decide. Another way to put the point is that when we get new information on a question—a barking dog, an AIDS test, or testimony at a trial—we don't hear it against a backdrop of no opinion at all. It always causes us to revise (or *update,* as Bayesians say) our earlier sense, however weak, of what the answer to the question probably was. So when a piece of evidence is shown at a trial, what is the base rate from which the jury should operate—the benchmark comparable to the one chance in a hundred that the passenger on the airplane had drugs?

It seems natural to reply that the jurors should start with even odds: a belief that the plaintiff is as likely to be right as wrong. This sounds consistent with the preponderance of the evidence standard and the usual view that it amounts to confidence better than 50 percent. But it would lead to embarrassing problems.[59] It would mean that any evidence at all in the plaintiff's favor would get him to 51 percent and to victory. Suppose the plaintiff was attacked by someone wearing a ski mask; all he could see was that the attacker had brown hair. Sixty percent of the people in the region have brown hair. The plaintiff picks one of them at random and sues him for damages. If the jurors start with a belief that the defendant is 50 percent likely to be the right person, then all the plaintiff might have to do is point out that the defendant has brown hair; for that makes it slightly *more* likely than 50 percent that the defendant was indeed the attacker. But that seems crazy. The plaintiff just picked any brown-haired person at random and sued him; he probably is innocent. The analysis of this problem is explained in the appendix to this chapter, but the basic reason for the bizarre result is that we *began* by assuming—in our usual fanciful fashion—that if the plaintiff were attacked 100 times, the defendant would be the attacker in 50 of them. There was just no reason to think that.

We could go through a similar example on the criminal side. A starting assumption that the defendant is 50 percent likely to be guilty might seem fine, since the prosecutor still has to get well to the north of that number to gain a conviction under the "reasonable doubt" standard. But that's still ungenerous to the defendant, since he is supposed to be presumed innocent. That doesn't mean the jurors should assume it's equally likely that he's guilty or innocent; it means they should start with the assumption that he *isn't* guilty. But just how strong should the assumption be? What we now see is that the presumption of innocence, or any other statement of where the burden of proof lies, can be viewed as the selection of a base rate—a background likelihood of guilt or innocence that will then be revised by the evidence in the case. Should the jury be told to start by assuming that the defendant is no more likely to be guilty of the crime than any random person on the street? So it has been suggested;[60] but then should the strength of the assumption depend on how many people there *are* on the street? Isn't a defendant in a small town more likely to be the right person than a defendant in New York City? Well, nobody says so.[61]

There aren't any settled answers to these questions. We have seen earlier that jurors aren't instructed to turn standards of proof into numerical tests. By the same token they certainly aren't told to think in the probabilistic ways we have been developing here, nor do judges talk this way. Jurors, like anyone, all start with priors about how likely the defendant is to be liable (in a civil case) or guilty (in a criminal one), but they are told only who has the burden of proof and what standards they have to meet. The details are left to the imagination.

So much for quantifying the burden of proof. Now what happens if a plaintiff wants to use a base rate—a statistical generalization—as his *only* evidence to meet it? The classic example involves a plaintiff run over by a bus. He remembers nothing. All we know this time is that 80 percent of the buses in town are owned by the defendant. Should this evidence alone make the plaintiff a winner? It does seem more likely than not that the defendant owned the bus, and that's the question in a civil case. Yet nobody thinks the plaintiff should win such a lawsuit without other evidence.[62] Various reasons have been put forward to explain the intuition. As explained earlier, in a case like this the real question is how confident the jurors are that the plaintiff is right; and if the statistics are the plaintiff's only evidence, the jurors might not have much confidence despite the strength of the numbers. They might reasonably wonder *why* that is the only evidence the plaintiff has put forward. Was the other possible evidence unhelpful to his case? Was he too cheap to do a more thorough investigation? The point can be turned into a claim about incentives: forbidding a case to go forward just on the strength of statistical statements forces plaintiffs to try to find better evidence than that, which we want them to do. Occasionally it might be impossible, but it's hard to tell when that's true and when the plaintiff merely is shirking—so we assume the latter. Apart from all this, allowing the plaintiff to win with that sort of evidence might cause the defendant to lose every such case despite being responsible in only, say, 80 percent of them. That's a hard result, and one that gives the wrong incentive both to the defendant and to the other bus company in town (which never gets held liable).

Another possibility is that the law won't stand for liability on these facts because it would violate the principle that guilt or liability ought to follow from individualized findings of wrongdoing, not mere statements of probabilities. But that principle has been known to bend. In the 1970s many women developed tumors because their mothers took a drug called DES during pregnancy. Unfortunately, by the time the disease appeared the women couldn't figure out which companies had filled the prescriptions their mothers took many years earlier. Many courts have allowed those plaintiffs to win anyway; otherwise these cases would involve recurring misses of the kind discussed in chapter 28. The interesting point was how the courts structured the right to recover. They decided to let each plaintiff collect an amount from each maker of DES that corresponded to the maker's share of the market for the drug.[63] This isn't quite like the bus case because here we know that each maker of DES really did injure some large number of plaintiffs. The base rate here, in other words, isn't just a statement of subjective confidence. It's a true statement of frequencies. If Squibb supplied 30 percent of the DES to the market, the point isn't that we're 30 percent sure Squibb caused the plaintiff's injuries; it's that we're 100 percent sure Squibb caused 30 percent of *all* plaintiffs' injuries. When companies are held liable this way (it isn't done often), there is less violence done to usual notions of fairness and efficiency than we might find in the bus case.

The DES cases are exceptional; base rates alone generally aren't enough to win a case. But then when should they be allowed into a case at all? Let's consider an example of a case where Bayesian reasoning might be helpful. A woman is stabbed to death, and a partial palm print is found on the knife. Only one person in 1,000 has a print that matches. Smithers is one of them, so evidently there's only 1 chance in 1,000 that he didn't do the crime, right? No, that's wrong, as we saw in chapter 29; it's called the prosecutor's fallacy. But in that discussion we were assuming there was no other evidence in the case. Suppose there *is* some other evidence against Smithers; now what are we to make of the handprint? Bayes' theorem provides a way to think about this question that is precise—but not necessarily desirable.[64] Try treating the matching palm print as a new piece of evidence that arrives, like the barking dog, against the backdrop of some given likelihood that the defendant is guilty. (We could have treated the palm print as establishing the base rate and the other evidence as the new information that calls for updated odds; the decision about which information to treat as "in the background" is arbitrary.) In the case of the palm print, the base rate is just the jurors' confidence that Smithers is guilty, based on whatever other evidence they have encountered *before* hearing about the palm print. Perhaps they have already learned that Smithers had a grudge against the victim, and as a result they think that he is 25 percent likely to have committed the crime. Or perhaps they think it's only 1 percent likely. What is the effect of the palm print, again assuming that it would match 1 person in 1,000 in the population at large?

Once more the math will be put in the appendix (problem two); here let's just consider the result. If the jurors thought the defendant's guilt was 25 percent likely before hearing of the handprint, their confidence after hearing about it—or rather the posterior odds of his guilt, strictly from a mathematical standpoint—should become 99 percent. Perhaps more interestingly, suppose the jurors started out thinking there was only a 1 percent chance that Smithers did it. Add the handprint and the odds then go from 1 percent to 90 percent. (We're leaving out some complications—that maybe the handprint does belong to Smithers but he was framed, for example; those could be accounted for with no great change in the outcomes.) This might seem to contradict some lessons from the chapter on the product rule. Wasn't the whole point that you shouldn't jump to conclusions just because the person who committed the crime has some rare trait that the defendant also has? Yes; but this example shows how quickly that reasoning changes if you have something else—even a little something else—to connect the defendant to the crime.

The points just explained aren't very intuitive, and they may not match the actual effect of such evidence on a typical juror's felt sense of confidence—for in practice jurors aren't told how Thomas Bayes would have them update their earlier sense of guilt in view of new evidence, and they aren't invited to think in the way just described. The interesting question is why. Maybe the reason is that jurors would make a botch of it, especially if we tried to add in those complications we left out—that Smithers might have been framed and so forth. Maybe asking jurors to think in terms of numbers does a disservice to the defendant, or to the ritual of the trial, or to our values concerning the meaning of guilt and liability. We want to avoid finding someone guilty because of their membership in a statistical class.

But then comes a last puzzle: sometimes math *does* play a significant role in a case; sometimes courts will allow statistical generalizations as evidence, and sometimes they won't, though it's hard to say quite when and why. In one case a woman was hurt by ice that slid off a passing truck. The truck had a Hertz logo on the side, so she sued Hertz; the court allowed evidence that Hertz owns 90 percent of the trucks with that logo, and as a result was willing to presume it owned this one.[65] In another case a child had been abused and the accused defendant was his father. A doctor told the jury that "eighty to eighty-five percent of child sexual abuse is committed by a relative close to the child"; this was held objectionable, the court saying it was difficult to understand "how statistical information would assist a trier of fact in reaching a determination as to guilt in an individual case."[66] The point can't be (or can't *just* be) that the second case was a criminal prosecution, for prosecutors sometimes are allowed to bring in statistical base rates, too. In a rape case where the victim couldn't identify anything about her attacker except his race, for example, a court allowed the jury to convict the defendant solely on the basis of testimony that the DNA found at the crime would match only 1 in 420 billion black people—and that it matched him.[67]

So sometimes courts do allow statistical generalizations into a case and sometimes they don't—but when? The best efforts to answer the question turn up no rules but a few noticeable tendencies.[68] Such evidence is most likely to be let in when the case has a clearly defined statistical structure (some employment discrimination cases are like this), or when one side wants to rebut the other's claim that some bad event happened by chance (as when two test-takers accused of cheating claim that their similar answers were just a coincidence), or when it will be hard for a plaintiff to come up with any evidence *other* than these sorts of generalities (think again of the DES cases; if the makers hadn't been held liable on a market-share theory, they might routinely have gotten off with no liability).

But the more general point to observe is the pattern of deep ambivalence the law displays about quantification. In a sense all knowledge is based on likelihoods, is it not? Yet courts avoid the explicit use of likelihoods habitually and allow it unpredictably, leaving many of their standards and conclusions unspoken, vague, and perhaps arbitrary because these are thought to be lesser evils than letting numbers play too significant a role in legal judgments. This state of affairs can be seen as part of a more general set of problems we have considered at other points, too: When do we want rigorous quantification in law, and when and why do we avoid it?

APPENDIX ➻ More Fun with Base Rates

1. *The brown-haired defendant.* In the main part of the chapter we considered a case where an attacker has brown hair, as does 60 percent of the population; the plaintiff sues one of the brown-haired people at random. The point of this silly-sounding case is to see what happens—and how the plaintiff wins—when we start with an assumption of even odds that the defendant is the culprit. Our goal is to build a fraction where the top

number is the number of times in 100 that the plaintiff would see brown hair and it would be the defendant's. *If* we assume prior odds of 50/50, then that means in half the cases—50 percent—the defendant *is* the attacker; in all of those the plaintiff would correctly observe brown hair. So the top number in the fraction is 50. The bottom number is 50 plus however many times the plaintiff would see brown hair and it *wouldn't* be the defendant's (like the dog barking when there are no drugs). Here we find that by taking the 50 times when the attacker isn't the defendant and multiplying it by 60 percent (the number of brunettes). That gets us 30. So our final fraction is 5/8, which means the defendant—a random brunette!—is 63 percent likely to be the attacker. But of course that's wrong; as you can see, the incorrect 50 percent base rate is what makes this train of thought end as it does.

2. *The palm print.* In the text we also saw a case where jurors start out thinking the defendant has a 1 percent chance of being guilty, then are shown that his palm print matches the one on the murder weapon—and there was 1 chance in 1,000 of such a match occurring at random. So what is the updated chance that the defendant did the crime? Build the fraction. On top goes the number of times that the dog would bark and be right—or, here, the number of times the palm print would match and be right; this time let's do it out of 10,000 imaginary run-throughs. We started with a 1 percent chance that Smithers was guilty. That means that in every 100 cases "like this," Smithers did it; and his handprint will match, we can suppose, in all of them. So that's 100 times (out of 10,000) that the handprint will match and correctly identify him as the killer. Now the bottom of the fraction consists of all the times the print will match, period—whether rightly or wrongly. It will happen 100 times, as we just saw—plus the times out of 10,000 when the handprint will match but Smithers *didn't* do it, which we need to figure out next. That will be a rare event. One person in 1,000 has the match, so we could expect a bad or mistaken match about 10 times in 10,000 trials. Okay; now we have the bottom of the fraction, too: 110 (the 10 mistakes plus the 100 correct hits). The result looks like this: 100/110. Do the division and you get about a 90 percent chance that Smithers did it.

3. *The O. J. Simpson case.* When O. J. Simpson was put on trial for the murder of his ex-wife, the prosecution attempted to bring in evidence that he had beaten her during their marriage. One of Simpson's lawyers, Alan Dershowitz, famously argued that only a tiny percentage of husbands who beat their wives—"certainly fewer than 1 of 2,500"—go on to murder them; for every 3 million incidents of abuse, there were only

about 1,500 murders (we have rounded off the numbers). How relevant are these statistics? Not very. You can see the problem clearly by structuring it like our other exercises in this chapter. Simpson's beating his wife is like the dog barking because he thinks he smells drugs. It *may* indicate that some underlying fact is present (that Simpson murdered his wife, or the passenger in fact is carrying drugs). To figure out the likelihood, we build a fraction. On top goes the number of times the foreground fact will be present (wife beating) and will correctly suggest guilt. This time let's do it out of 100,000 tries. The background population consists of battered women. If there are 100,000 of them, then some number will be murdered by their husbands—let's say 1 in 2,500, which is the same as 4 in 10,000, or 40 in 100,000. So 40 is the number that goes on top of our fraction. On the bottom of it goes the number of times that a battered wife is murdered by *anyone,* husband or not. Of course we begin with the 40 murders committed by husbands, then we add the number of times such women are murdered by other people. Ah, but this is a very small number: about .05 percent, or 5 out of 100,000. So we add the 5 to the 40, and now we have our fraction: 40/45—or about 89 percent. That is the rough likelihood that if a battered wife is found murdered, she was killed by her husband, at least using the raw numbers borrowed here. I don't vouch for their accuracy; but even if you change them around in plausible ways, it's still clear that the evidence that Simpson beat his wife was far more important than Dershowitz's argument caused it to seem.[69]

Can you summarize why the initial claim was misleading? It was answering the wrong question. The important issue isn't the rate at which battered women are murdered by their husbands. It's how that number *compares* to the rate at which battered women are murdered by anyone—just as we asked how the number of correct barks compared to the number of total barks by the drug-sniffing dog. Merlin's barking wasn't as impressive a piece of evidence as it seemed, because he had a large error rate and most people on a plane aren't carrying drugs. Simpson's abuse was a more impressive piece of evidence than his lawyer made it seem, though, because *once we know a battered woman has been murdered,* her husband's guilt is quite likely—much more likely than possession of drugs by some random passenger on a plane. To turn the point around, the background rates of innocence are different. The rate at which innocent passengers are barked at by the dog is large compared to the rate at which he finds drugs. But the rate at which battered wives are murdered by people other than their husbands is small compared to the rate at which their husbands are guilty.

4. *The taxicab problem.* Here is another classic puzzle. It's a little more complicated than the others.[70] Smithers is run over by a taxi. It was owned either by the Green Cab company or the Blue Cab company. Eighty-five percent of the cabs in the town are green; but Smithers is 80 percent sure the cab that hit him was blue. (Or instead of saying he is "80 percent sure," perhaps imagine that his judgments about these things have been shown 80 percent likely to be right.) Assuming there is no other evidence, what is the chance that he was hit by a blue cab? Remember the method. We build a fraction. On top goes the number of times the dog will bark and be right—or, here, the number of times (in a hundred imaginary accidents) that Smithers will say the cab was blue and be correct. This time it takes an extra minute to figure. The drug-sniffing dog caught the drugs every time they were there, and the AIDS test caught the AIDS every time it was there, but Smithers isn't like that. Even when the cab really was blue, he'll see it right just 80 percent of the time. But don't be distracted: we still want to know the same thing, which is just how many times out of a hundred incidents Smithers will see blue and be right. Think in frequencies. Here 15 cabs out of our 100 imaginary incidents will be blue (since 85 percent of them are green). Smithers will spot 80 percent of them. 80 percent of 15 is 12. That's the number on top of our fraction. On the bottom goes *all* the times Smithers will see blue. That will be the 12 we just counted, plus all the times when he says the cab was blue but is wrong. How many times will that happen? Well, the cab will be green in 85 of the 100 tries we are imagining, and he will get 20 percent of those *wrong* (that's his error rate). Twenty percent of 85 is 17. So now we have our fraction:

$$\frac{\text{12 (the number of times in 100 Smithers will say the cab was blue and be right)}}{\text{29 (the number of times he will say it was blue, whether it was or not—12 + 17)}}$$

This problem was harder than the others because we had to worry about false negatives as well as false positives. We had to take the times the car was blue and multiply it by the chance that he would see it rightly; and we had to take the times the car was green and multiply it by the chance he would see it wrongly. Those numbers together (12 and 17) give us the total number of times in a hundred he will say the car was blue. They are the only numbers we need here. The first one goes on top and both of them get added together on the bottom. (The 12/29 fraction we end with can itself be made a percentage, of course: there's a 41 percent chance the cab was blue.)

SUGGESTIONS FOR FURTHER READING. Michael O. Finkelstein and William B. Fairley, *A Bayesian Approach to Identification Evidence,* 83 Harv. L. Rev. 489 (1970); Laurence H. Tribe, *Trial by Mathematics: Precision and Ritual in the Legal Process,* 84 Harv. L. Rev. 1329 (1971); Richard A. Posner, *An Economic Approach to the Law of Evidence,* 51 Stan. L. Rev. 1477 (1999); Gerd Gigerenzer and Ulrich Hoffrage, *How to Improve Bayesian Reasoning without Instruction: Frequency Formats,* 102 Psychol. Rev. 684 (1995); Richard D. Friedman, *A Presumption of Innocence, Not Even Odds,* 52 Stan. L. Rev. 873 (2000); Jonathan J. Koehler, *When Do Courts Think Base Rate Statistics Are Relevant?* 42 Jurimetrics J. 373 (2002); Peter Tillers and Eric D. Green eds., Probability and Inference in the Law of Evidence: The Uses and Limits of Bayesianism (1988).