

**The second-ace puzzle** A deck has four cards: the ace and deuce of hearts, and the ace and deuce of spades. After a fair shuffle of the deck, two cards are dealt to Alice. It is easy to see that, at this point, there is a probability of  $1/6$  that Alice has both aces, a probability of  $5/6$  that Alice has at least one ace, a probability of  $1/2$  that Alice has the ace of spades, and a probability of  $1/2$  that Alice has the ace of hearts: of the six possible deals of two cards out of four, Alice has both aces in one of them, at least one ace in five of them, the ace of hearts in three of them, and the ace of spades in three of them. (For readers unfamiliar with probability, there is an introduction in Chapter 2.)

Alice then says, “I have an ace.” Conditioning on this information (by discarding the possibility that Alice was dealt no aces), Bob computes the probability that Alice holds both aces to be  $1/5$ . This seems reasonable. The probability, according to Bob, of Alice having two aces goes up if he learns that she has an ace. Next, Alice says, “I have the ace of spades.” Conditioning on this new information, Bob now computes the probability that Alice holds both aces to be  $1/3$ . Of the three deals in which Alice holds the ace of spades, she holds both aces in one of them. As a result of learning not only that Alice holds at least one ace, but that the ace is actually the ace of spades, the conditional probability that Alice holds both aces goes up from  $1/5$  to  $1/3$ . But suppose that Alice had instead said, “I have the ace of hearts.” It seems that a similar argument again shows that the conditional probability that Alice holds both aces is  $1/3$ .

Is this reasonable? When Bob learns that Alice has an ace, he knows that she must have either the ace of hearts or the ace of spades. Why should finding out which particular ace it is raise the conditional probability of Alice having two aces? Put another way, if this probability goes up from  $1/5$  to  $1/3$  whichever ace Alice says she has, and Bob knows that she has an ace, then why isn't it  $1/3$  all along?

**The Monty Hall puzzle** The Monty Hall puzzle is very similar to the second-ace puzzle. Suppose that you're on a game show and given a choice of three doors. Behind one is a car; behind the others are goats. You pick door 1. Before opening door 1, host Monty Hall (who knows what is behind each door) opens door 3, which has a goat. He then asks you if you still want to take what's behind door 1, or to take instead what's behind door 2. Should you switch? Assuming that, initially, the car was equally likely to be behind each of the doors, naive conditioning suggests that, given that it is not behind door 3, it is equally likely to be behind door 1 and door 2, so there is no reason to switch. On the other hand, the car is equally likely to be behind each of the doors. If it is behind door 1, then you clearly should not switch; but if it is not behind door 1, then it must be behind door 2 (since it is obviously not behind door 3), and you should switch to door 2. Since the probability that it is behind door 1 is  $1/3$ , it seems that, with probability  $2/3$ , you should switch. But if this reasoning is correct, then why exactly is the original argument incorrect?

The second-ace puzzle and the Monty Hall puzzle are the stuff of puzzle books. Nevertheless, understanding exactly why naive conditioning does not give reasonable answers in these cases turns out to have deep implications, not just for puzzles, but for important statistical problems.

An agent typically acquires new information all the time. How should the new information affect her beliefs? The standard way of incorporating new information in probability theory is by *conditioning*. This is what Bob used in the second-ace puzzle to incorporate the information he got from Alice, such as the fact that she

holds an ace or that she holds the ace of hearts. This puzzle already suggests that there are subtleties involved with conditioning. Things get even more complicated if uncertainty is not represented using probability, or if the new information does not come in a nice package that allows conditioning. (Consider, e.g., information like “people with jaundice typically have hepatitis.”) Chapter 3 examines conditioning in the

It turns out that in order to represent the puzzle formally, it is important to describe the *protocol* used by Alice. The protocol determines the set of *runs*, or possible sequences of events that might happen. The key question here is what Alice's protocol says to do after she has answered "yes" to Bob's question as to whether she has an ace. Roughly speaking, if her protocol is "if I have the ace of spades, then I will say that, otherwise I will say nothing," then  $1/3$  is indeed Bob's probability that Alice has both aces. This is the conditional probability of Alice having both aces given that she has the ace of spades. On the other hand, suppose that her protocol is "I will tell Bob which ace I have; if I have both, I will choose at random between the ace of hearts and the ace of spades." Then, in fact, Bob's conditional probability should not go up to  $1/3$  but should stay at  $1/5$ . The different protocols determine different possible runs and so result in different probability spaces. In general, it is critical to make the protocol explicit in examples such as this one.

4. Their captors have decided that two of three prisoners—Smith, Jones, and Fitch—will be executed tomorrow. The choice has been made at random, but the identity of the unfortunate selectees is to be kept from the prisoners until the final hour. The prisoners, who are held in separate cells, unable to communicate with each other, know this. Fitch asks a guard to tell the name of one of the other prisoners who will be executed. Regardless of whether Fitch was chosen or not, one of the others will be executed, so the guard reasons that he is not giving Fitch any illicit information by answering truthfully. He says: “Jones will be executed.” Fitch is heartened by the news for he reasons that his probability of being the one who escapes execution has risen from  $\frac{1}{3}$  to  $\frac{1}{2}$ . Has Fitch made a mistake? Has the guard? Use Bayes’ theorem to analyze the reasoning involved. (Hint: Calculate

the probability that Fitch will not be executed given that *the guard tells him that Jones will be executed*, not the probability that Fitch will not be executed given that Jones will be. What assumptions are possible about the probability that the guard tells Fitch that Jones will be executed given that Fitch escapes execution?)

4.  $F$  = “Fitch will not be executed”  
 $J$  = “Jones will not be executed”  
 $S$  = “Smith will not be executed”  
 $G$  = “Guard tells Fitch that Jones will be executed”

According to the problem,  $\Pr(F) = \Pr(J) = \Pr(S) = \frac{1}{3}$ . We need to calculate  $\Pr(F \text{ given } G)$ . Using Bayes’ theorem:

$$\Pr(F \text{ given } G) = \frac{\Pr(F) \cdot \Pr(G \text{ given } F)}{\Pr(F) \cdot \Pr(G \text{ given } F) + \Pr(J) \cdot \Pr(G \text{ given } J) + \Pr(S) \cdot \Pr(G \text{ given } S)}$$

The guard is said to be truthful, so  $\Pr(G \text{ given } J) = 0$ . What is  $\Pr(G \text{ given } S)$ ? If Smith will not be executed, then Fitch and Jones will both be executed. But the guard cannot tell Fitch that he will be executed, so in this case he must tell him that Jones will be executed. So  $\Pr(G \text{ given } S) = 1$ . So far:

$$\Pr(F \text{ given } G) = \frac{(\frac{1}{3}) \cdot \Pr(G \text{ given } F)}{(\frac{1}{3}) \cdot \Pr(G \text{ given } F) + (\frac{1}{3})}$$

Everything turns on  $\Pr(G \text{ given } F)$ . If Fitch will not be executed, Jones and Smith will be. Will the guard, in this case, say “Jones” or Smith?” The problem gives no reason why he should prefer one rather than another, so he might flip a fair coin, in which case  $\Pr(G \text{ given } F) = \frac{1}{2}$ . Then  $\Pr(F \text{ given } G) = (\frac{1}{6})/(\frac{3}{6}) = \frac{1}{3}$ , and Fitch has no better prospects than before.

On the other hand, you might imagine that the guard has special reasons to say “Jones” in this case, so that  $\Pr(F \text{ given } G) = 1$ . If so,



Fitch has good news, for then:  $\Pr(F \text{ given } G) = (\frac{1}{3})/(\frac{2}{3}) = \frac{1}{2}$ . But you might just as well imagine that the guard has special reasons to say “Smith” in this case, so that  $\Pr(F \text{ given } G) = 0$ . If so, Fitch has bad news, for then  $\Pr(F \text{ given } G) = 0$ !

now, I just consider how sets of probabilities can be used to deal with the *three-prisoners puzzle*.

**Example 3.3.1** The three-prisoners puzzle is an old chestnut that is somewhat similar in spirit to the second-ace puzzle discussed in Chapter 1, although it illustrates somewhat different issues.

Of three prisoners  $a$ ,  $b$ , and  $c$ , two are to be executed, but  $a$  does not know which. He therefore says to the jailer, “Since either  $b$  or  $c$  is certainly going to be executed, you will give me no information about my own chances if you give me the name of one man, either  $b$  or  $c$ , who is going to be executed.” Accepting this argument, the jailer truthfully replies, “ $b$  will be executed.” Thereupon  $a$  feels happier because before the jailer replied, his own chance of execution was  $2/3$ , but afterward there are only two people, himself and  $c$ , who could be the one not executed, and so his chance of execution is  $1/2$ .

Note that in order for  $a$  to believe that his own chance of execution was  $2/3$  before the jailer replied, he seems to be implicitly assuming the principle of indifference. A straightforward application of the principle of indifference also seems to lead to  $a$ 's believing that his chances of execution goes down to  $1/2$  after hearing the jailer's statement. Yet it seems that the jailer did not give him any new relevant information. Is  $a$  justified in believing that his chances of avoiding execution have improved? If so, it seems that  $a$  would be equally justified in believing that his chances of avoiding execution would have improved if the jailer had said “ $c$  will be executed.” It seems that  $a$ 's prospects improve no matter what the jailer says! That does not seem quite right.

The principle of indifference is implicitly being applied here to a space consisting of three worlds—say  $w_a$ ,  $w_b$ , and  $w_c$ —where in world  $w_x$ , prisoner  $x$  is pardoned. But this representation of a world does not take into account what the jailer says. Perhaps a better representation of a possible situation is as a pair  $(x, y)$ , where  $x, y \in \{a, b, c\}$ . Intuitively, a pair  $(x, y)$  represents a situation where  $x$  is pardoned and the jailer says that  $y$  will be executed in response to  $a$ 's question. Since the jailer answers truthfully,  $x \neq y$ ; since the jailer will never tell  $a$  directly that  $a$  will be executed,  $y \neq a$ . Thus, the set of possible worlds is  $\{(a, b), (a, c), (b, c), (c, b)\}$ . The event *lives -a*— $a$  lives—corresponds to the set  $\{(a, b), (a, c)\}$ . Similarly, the events *lives -b* and *lives -c* correspond to the sets  $\{(b, c)\}$  and  $\{(c, b)\}$ , respectively. Assume in accord with the principle of indifference that each prisoner is equally likely to be pardoned, so that each of these three events has probability  $1/3$ .

The event *says -b*—the jailer says  $b$ —corresponds to the set  $\{(a, b), (c, b)\}$ ; the story does not give a probability for this event. To do standard probabilistic condition-

ing, this set must be measurable and have a probability. The event  $\{(c, b)\}$  (*lives-c*) has probability  $1/3$ . But what is the probability of  $\{(a, b)\}$ ? That depends on the jailer's strategy in the one case where he has a choice, namely, when  $a$  lives. He gets to choose between saying  $b$  and  $c$  in that case. The probability of  $(a, b)$  depends on the probability that he says  $b$  if  $a$  lives; that is,  $\mu(\text{says-}b \mid \text{lives-}a)$ .

If the jailer applies the principle of indifference in choosing between saying  $b$  and  $c$  if  $a$  is pardoned, so that  $\mu(\text{says-}b \mid \text{lives-}a) = 1/2$ , then  $\mu(\{(a, b)\}) = \mu(\{(a, c)\}) = 1/6$ , and  $\mu(\text{says-}b) = 1/2$ . With this assumption,

$$\mu(\text{lives-}a \mid \text{says-}b) = \mu(\text{lives-}a \cap \text{says-}b) / \mu(\text{says-}b) = (1/6) / (1/2) = 1/3.$$

Thus, if  $\mu(\text{says-}b) = 1/2$ , the jailer's answer does not affect  $a$ 's probability.

Suppose more generally that  $\mu_\alpha$ ,  $0 \leq \alpha \leq 1$ , is the probability measure such that  $\mu_\alpha(\text{lives-}a) = \mu_\alpha(\text{lives-}b) = \mu_\alpha(\text{lives-}c) = 1/3$  and  $\mu_\alpha(\text{says-}b \mid \text{lives-}a) = \alpha$ . Then straightforward computations show that

$$\begin{aligned} \mu_\alpha(\{(a, b)\}) &= \mu_\alpha(\text{lives-}a) \times \mu_\alpha(\text{says-}b \mid \text{lives-}a) = \alpha/3, \\ \mu_\alpha(\text{says-}b) &= \mu_\alpha(\{(a, b)\}) + \mu_\alpha(\{(c, b)\}) = (\alpha + 1)/3, \text{ and} \\ \mu_\alpha(\text{lives-}a \mid \text{says-}b) &= \frac{\alpha/3}{(\alpha + 1)/3} = \alpha/(\alpha + 1). \end{aligned}$$

Thus,  $\mu_{1/2} = \mu$ . Moreover, if  $\alpha \neq 1/2$  (i.e., if the jailer had a particular preference for answering either  $b$  or  $c$  when  $a$  was the one pardoned), then  $a$ 's probability of being executed would change, depending on the answer. For example, if  $\alpha = 0$ , then if  $a$  is pardoned, the jailer will definitely say  $c$ . Thus, if the jailer actually says  $b$ , then  $a$  knows that he is definitely not pardoned, that is,  $\mu_0(\text{lives-}a \mid \text{says-}b) = 0$ . Similarly, if  $\alpha = 1$ , then  $a$  knows that if either he or  $c$  is pardoned, then the jailer will say  $b$ , while if  $b$  is pardoned the jailer will say  $c$ . Given that the jailer says  $b$ , from  $a$ 's point of view the one pardoned is equally likely to be him or  $c$ ; thus,  $\mu_1(\text{lives-}a \mid \text{says-}b) = 1/2$ . In fact, it is easy to see that if  $\mathcal{P}_J = \{\mu_\alpha : \alpha \in [0, 1]\}$ , then  $(\mathcal{P}_J \mid \text{says-}b)_*(\text{lives-}a) = 0$  and  $(\mathcal{P}_J \mid \text{says-}b)^*(\text{lives-}a) = 1/2$ .

To summarize, the intuitive answer—that the jailer's answer gives  $a$  no information—is correct if the jailer applies the principle of indifference in the one case where he has a choice in what to say, namely, when  $a$  is actually the one to live. If the jailer does not apply the principle of indifference in this case, then  $a$  may gain information. On the other hand, if  $a$  does not know what strategy the jailer is using to answer (and is not willing to place a probability on these strategies), then his prior point probability of  $1/3$  “diffuses” to an interval. ■