Clinician's Probability Primer

Stuart Spitalnic, MD

The lottery is a tax on people who are bad at math.
—Unattributed bumper sticker

Consider how often in clinical practice the following statements are made:

- "I doubt the patient is having a pulmonary embolism, because…"
- "The *P* value in that study was .06; thus…"
- "Combining these test results essentially rules in the diagnosis of…"

We use testing strategies to increase the certainty of our diagnostic reasoning. We select antibiotics that are most likely to treat an infection. When we read the medical literature, we consider what findings may be "real" and which ones might have been observed "by chance." For most of our clinical practice activities, phrases such as "ruled in," "ruled out," and "high likelihood" are adequate to describe the certainty of our conclusions. Because these phrases represent a range of numerical probabilities, however, the rules of probability can be applied to tell us with what frequency we can expect particular outcomes.

To understand test properties, including sensitivity, specificity, and predictive values, and to appreciate the meaning of *P* values and confidence intervals in the medical literature, it is essential to understand some basic concepts of probability. This article will introduce elements of probability typically referred to casually in the clinical arena. By developing a better understanding of probability, one can predict more accurately the outcomes of medical decisions, improve diagnostic test strategies, and better appreciate the real value of conclusions in the medical literature. Future articles in the *Primer in Literature Interpretation* series will explore topics in test properties, including sensitivity and specificity, predictive values, likelihood ratios, and Bayes' formula.

DEFINITIONS AND NOTATIONS

Probabilities range from 0.00 to 1.00 (0% to 100%). If something has a probability of 0.00, that means it is never observed; if the probability is 1.00, then its observation is a certainty. The probability of an event refers to its relative frequency. For example, if the field of a spinner is divided into 4 equal quarters labeled A, B, C, and D (**Figure 1**), we would expect the spinner to land in the "A" region 25% of the time, or 1 out of every 4 spins. The probability of the spinner landing in the "A" region (or, for that matter, in any of the other 3 letter fields) is .25. Note that in practice, if we spin the spinner 8 times, it might land in the "A" region the expected 2 times but also might land there never, once, 6 times, or any number of times between 0 and 8.

Probability values reflect the frequency that an outcome will be observed as the number of observations becomes very large. If we flipped a coin 6 times, we would not be surprised if it came up heads 4 times and tails twice; however, if we flipped a coin 600 times, we would become suspicious if it came up heads 400 times. In the 6-flips example, we still believe the coin is "fair," even after observing 2 of 3 flips coming up heads. In contrast, in the 600-flips example, we conclude that the coin is coming up head an excess number of times, even though in both examples the proportionate outcome is the same. Although the reasons behind our different conclusions in the 6-flips and 600-flips examples may not be immediately apparent now, once they are appreciated, significance testing as presented in every medical research paper will begin to become much clearer.

"*P*[event]" is the notation used to express the probability of an event. Therefore, the probability of the spinner landing in the "B" region can be expressed as *P*[B]. In addition, we might be interested in the probability of something not happening, as in the probability that a patient does not have a disease. Using the spinner example, we can ask, "What is the probability of the spinner not landing in the 'C' region?" If, for 3 out of 4 spins, the spinner is not expected to land in the "C" region, then the probability of it not landing there is .75. The notation used to express the probability of something not being observed (in this case, not

Dr. Spitalnic is an Assistant Professor of Medicine and an Assistant Residency Director of the Residency Program in Emergency Medicine, Brown University School of Medicine, Providence, RI; he is also a member of the Hospital Physician *Editorial Board.*

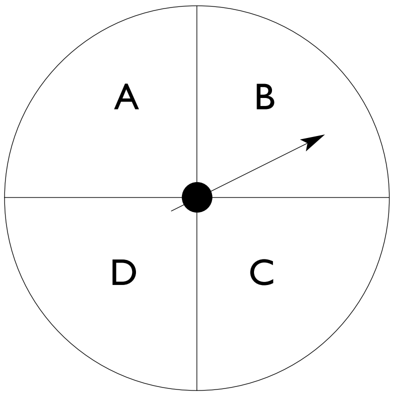

Figure 1. A spinner with 4 equally sized, labeled quadrants.

observing the spinner landing in the "C" region) is $P[\bar{C}]$; the bar over the "C" means "not."

A concept closely related to probability is odds. Whereas probability reflects the chance of something happening from 0.00 (impossible) to 1.00 (certainty), odds—which can range from 0 to infinity—reflect the ratio of the chance of something happening to the chance of it not happening. That is, if something has a probability of happening of .75, there is a probability of .25 that it will not happen. The odds of it happening are .75/.25 or 3 (often represented as 3:1), meaning it will occur 3 times for every 1 time that it does not. Before anyone says, "Wait a minute. A 3:1 horse doesn't win 75% of the time," note that odds in horse racing express the odds against the horse winning; the ratio would be the chance of the horse losing divided by the chance of it winning. (Note that horse-racing odds have to do with payout for a winning ticket, not the real chance of a horse winning, but that's another story.) When discussing odds, it is important to express whether we are concerned with the odds for or against something happening. A later article in this series will revisit odds and discuss a concept known as the *odds ratio.*

CONDITIONAL PROBABILITIES

When we talk about the probability of a given event, the implied question is, "What is the relative frequency we will observe a particular event, considering all possible outcomes?" When we restrict the number of possible outcomes, we are seeking a conditional probability. Consider the following two questions:

- What is the probability that someone who has 2 children has 2 girls?
- What is the probability that someone who has 2 children has 2 girls, given that at least 1 of the children is a girl?

In the first question, there are 4 equally possible ways to have 2 children: (GG, BG, GB, BB). Because the "GG" combination is considered a success and represents 1 of 4 possible outcomes, the probability of success is 1/4, or *P*[GG] = .25. In the second question, we have placed a condition on the possible outcomes, namely, that there is at least 1 girl. With this restriction, there are three equally likely arrangements (GG, BG, GB), and the probability is 1/3. The notation *P*[A|B] is used to express conditional probability, read as the chance of outcome A, given restriction B. In this example, we would write *P*[GG|at least 1 girl]. If we had asked the question, "What is the probability that someone who has 2 children has 2 girls, given that the oldest child is a girl?" then the probability would be 1/2 or 0.5, because then there would be only 2 possibilities (GG, GB), 1 of which is a success.

When the restriction has no effect on the outcome of interest, it is said to be independent. An example is contained in the question, "What is the probability of flipping a coin and having it come up heads, given that a 3 has been rolled on a die?" The probability is .5, because the die does not restrict the possible outcomes or the likelihood of any particular outcome of the coin flip.

Test properties are typically expressions of conditional probabilities. When we quote a sensitivity, what we are relating is the chance that a test will be positive, not in all people who undergo the test but in those who have a particular condition. The negative predictive value is the chance a patient will not have a condition, given a negative test result. When we are interested not just in the probability of a particular event but the probability of an event given specific restrictions, we are referring to a conditional probability. If we substituted "a negative test" for A, and "a patient free of disease" for B, *P*[A|B] would express the probability of a negative test, given a patient free of disease; this is the specificity of a test.

Conditional probabilities have many direct clinical applications and are essential for the proper incorporation of tests into diagnostic strategies. They also are commonly employed in the risk stratification of patients for conditions such as pulmonary embolism, appendicitis, and ischemic cardiac disease.

COMBINING PROBABILITIES

There are rules for determining the probability of observing a set of conditions when we know the probability of the individual components of the set. Considering 2 outcomes of interest (eg, "A" and "B"), we could ask 3 basic questions: (1) what is the probability of observing A and B; (2) what is the probability of observing A or B; and (3) what is the probability of observing A, given B? The negative questions (eg, what is the probability of observing neither A nor B) can be derived by remembering that the probability of observing neither A nor B—provided that there must be some outcome—is 1 minus the probability of observing either A or B.

ADDING PROBABILITIES

When we are interested in the probability of 1 of several mutually exclusive possibilities, the probability is the sum of the individual probabilities. In the spinner example, if we are interested in the probability of the spinner landing in the "A" or the "B" regions, given that these outcomes are mutually exclusive, the probability would be $.25 + .25 = .5$.

When we are interested in the probability of events that are not mutually exclusive, then the probability of finding any 1 of them is the sum of the individual probabilities, minus any overlap. Referring to **Figure 2**, if we are interested in the probability of selecting someone who has diabetes OR hypertension, the probability would be $P[H] + P[D] - P[H \& D]$. Subtracting $P[H \& D]$ prevents double counting the overlap area. Note that this is a more general form of the addition rule and that, if events are mutually exclusive, the overlap area (H & D) would be 0.

MULTIPLYING PROBABILITIES

If we are interested in the probability of observing the combined outcome of several independent events, this overall probability is the product of the individual probabilities. As a simple illustration, consider the spinner example in Figure 1 and the example of a fair coin toss. What is the probability of simultaneously having the spinner land in the "A" region and of having a coin come up heads on a toss? The probability of the first is .25, and the probability of the second is .50, so the combined probability is $.25 \times .50 = .125$; in other words, in 1 out of 8 tries you would expect to observe both an "A" and heads. Using another example, if you were to have 2 children, what is the chance of them both being girls? Assuming that each child has a .50 chance of being a girl, the chance of both children being girls is $.50 \times .50 = .25$.

What about instances when the combined outcomes are not independent? For example, what is the probability that someone is diabetic AND hypertensive? We are interested in determining the amount of the overlap between the 2 fields. For this, we will need to use conditional probability by resolving the question into 2 independent questions, the probabilities of which can then be multiplied. We can consider the question of who is diabetic AND hypertensive as the probability of someone being diabetic multiplied by

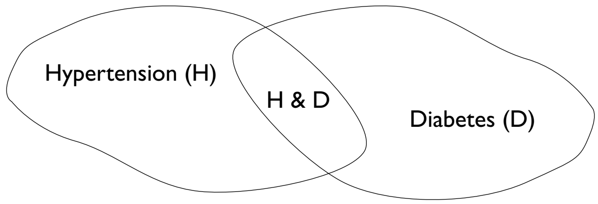

Figure 2. Drawing illustrating an overlap of those who have diabetes mellitus and hypertension.

the probability of someone being hypertensive, given that this person is diabetic. This can be represented as $P[H \& D] = P[D] \times P[H|D]$. By the same token, $P[H \& D] = P[H] \times P[D|H]$; thus, $P[H] \times P[D|H] = P[D] \times P[H|D]$. If the observation of interest is truly independent (ie, there is no overlap) then $P[H|D]$ would be 0 and, therefore, $P[H \& D]$ would be 0.

SUMMARY

We constantly (and appropriately) express uncertainty of clinical outcomes in terms of probability, but rarely are the fundamentals of thinking in terms of probability reviewed. In order to appreciate the true meaning of test results, to understand the actual risks our patients are exposed to with and without treatment, and to properly interpret the results of research papers, it is essential to develop at least a basic competence in reasoning with probabilities. It has been the purpose of this discussion to be a reminder of these basic ideas, so they can be built on in future articles on statistical concepts and literature interpretation. The next installment in this series will be on sensitivity, specificity, and predictive values, and appreciating the fundamentals of probability will greatly enhance one's understanding of these topics. Bayes' formula and likelihood ratios will likewise be discussed in future articles in this series; these topics will further develop principles of working with conditional probabilities. **HP**

EDITOR'S NOTE: Readers interested in further instruction regarding evidence-based medicine should consult the "Evidence-Based Practice" feature of *Seminars in Medical Practice*, available at http://www.turner-white.com/smp/smp_EBP.htm.

SUGGESTED READINGS

Freund JE. Introduction to probability. Encino (CA): Dickenson Publishing Co; 1973.

Langley R. Practical statistics simply explained. Rev ed. New York: Dover Publications; 1971.

Rosner BA. Fundamentals of biostatistics. 4th ed. Belmont (CA): Duxbury Press; 1995.