

Chance and Uncertainty: Probability Theory

Formally, we begin with a set of *elementary events*, precisely one of which will eventually occur. Each elementary event has associated with it a *probability*, which represents the likelihood that a particular event will occur; the probabilities are all nonnegative, and sum to 1.

For example, if a five-card hand is dealt from a thoroughly-shuffled deck of 52 cards, there are 2,598,960 different elementary events (different hands that might be dealt), each of which has probability $1/2,598,960$ of occurring.

An *event* is a collection of elementary events, and its probability is the sum of all of the probabilities of the elementary events in it.

For example, there are 778,320 distinct five-card hands that contain precisely one ace. Therefore, the event $A =$ “the hand contains precisely one ace” has probability $778,320/2,598,960$ of occurring; this is written as $\Pr(A) = 0.299474$.

[In many applications, the set of elementary events is not made explicit. For example, we can consider $\Pr(\text{the Cubs win the World Series})$ without needing to detail every possible way they might do so.]

The shorthand $A \cup B$ is written to mean “at least one of A or B occurs” (more concisely, “A or B”; more verbosely, “the disjunction of A and B”), and $A \cap B$ is written to mean “both A and B occur” (more concisely, “A and B”; more verbosely, “the conjunction of A and B”). A sometimes-useful relationship is

$$\Pr(A) + \Pr(B) = \Pr(A \cup B) + \Pr(A \cap B) .$$

Two events A and B are *mutually exclusive* (or *disjoint*) if it is impossible for both to occur, i.e., if $A \cap B = \emptyset$. The *complement* of A, written A^c , is the event “A does *not* occur”. Obviously,

$$\Pr(A) + \Pr(A^c) = 1 .$$

Whenever $\Pr(A) > 0$, the *conditional probability* that event B occurs, given that A occurs (or has occurred, or will occur), is $\Pr(B | A) = \Pr(A \cap B)/\Pr(A)$.

Two events A and B are *independent* if $\Pr(A \cap B) = \Pr(A) \cdot \Pr(B)$. Note that when A and B are independent, $\Pr(B | A) = \Pr(B)$. Three events, A, B, and C, are *mutually independent* if each pair is independent, and furthermore $\Pr(A \cap B \cap C) = \Pr(A) \cdot \Pr(B) \cdot \Pr(C)$. Similarly, any number of events are mutually independent if the probability of every conjunction of events is simply the product of the event probabilities.

[The following example illustrates why we must be careful in our definition of mutual independence: Let our elementary events be the outcomes of two successive flips of a fair coin. Let $A =$ “first flip is Heads”, $B =$ “second flip is Heads”, and $C =$ “exactly one flip is Heads”. Then each pair of events is independent, but the three are not mutually independent.]

A collection of events A_1, \dots, A_k is *mutually exclusive and exhaustive* if each is disjoint from the others, and together they cover all possibilities. If A_1, \dots, A_k are mutually exclusive and exhaustive, then

$$\begin{aligned}\Pr(B) &= \Pr(B \cap A_1) + \dots + \Pr(B \cap A_k) \\ &= \Pr(B \mid A_1) \cdot \Pr(A_1) + \dots + \Pr(B \mid A_k) \cdot \Pr(A_k) .\end{aligned}$$

Often, one observes some consequence B of a chance event, and wishes to make inferences about the outcome of the original event A . In such cases, it is typically easier to compute $\Pr(B \mid A)$, so the following well-known “rule” is of use.

Bayes' Rule:
$$\Pr(A \mid B) = \Pr(A \cap B) / \Pr(B) = \Pr(B \mid A) \cdot \Pr(A) / \Pr(B) .$$

The result at the top of this page can be used to compute $\Pr(B)$.

Example: On the basis of a patient's medical history and current symptoms, a physician feels that there is a 20% chance that the patient is suffering from disease X . There is a standard blood test which can be conducted: The test has a false-positive rate of 25% and a false-negative rate of 10%. What will the physician learn from the result of the blood test?

$$\begin{aligned}\Pr(X \text{ present} \mid \text{test positive}) &= \Pr(\text{test positive} \mid X \text{ present}) \cdot \Pr(X \text{ present}) / \Pr(\text{test positive}) \\ &= 0.9 \cdot 0.2 / \Pr(\text{test positive}) , \text{ and}\end{aligned}$$

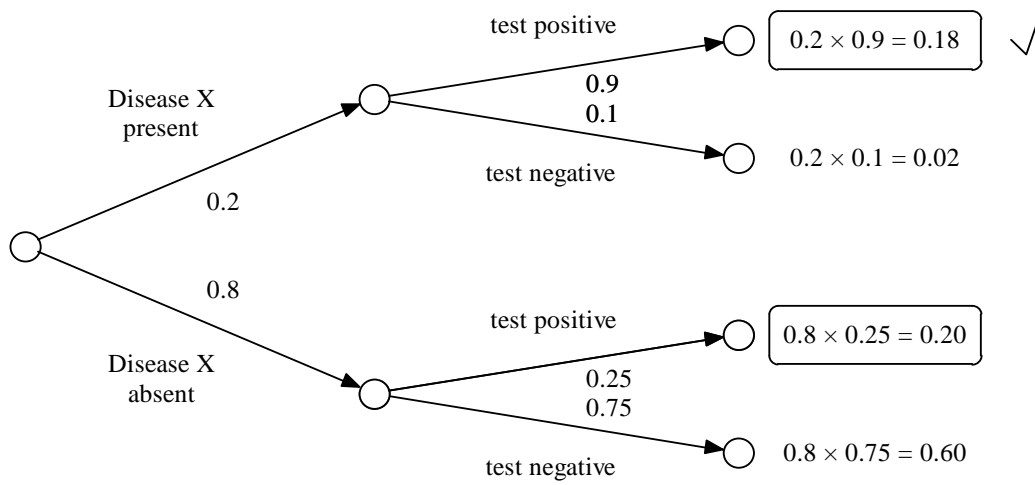
$$\begin{aligned}\Pr(\text{test positive}) &= \Pr(\text{test positive} \mid X \text{ present}) \cdot \Pr(X \text{ present}) \\ &\quad + \Pr(\text{test positive} \mid X \text{ not present}) \cdot \Pr(X \text{ not present}) \\ &= 0.9 \cdot 0.2 + 0.25 \cdot 0.8 = 0.38 , \text{ so}\end{aligned}$$

$$\Pr(X \text{ present} \mid \text{test positive}) = 0.18 / 0.38 = 0.474 .$$

Similarly,

$$\begin{aligned}\Pr(X \text{ present} \mid \text{test negative}) &= \Pr(\text{test negative} \mid X \text{ present}) \cdot \Pr(X \text{ present}) / \Pr(\text{test negative}) \\ &= 0.1 \cdot 0.2 / \Pr(\text{test negative}) = 0.02 / (1 - 0.38) = 0.032 .\end{aligned}$$

The preceding calculations are much simpler to visualize in “probability-tree” form:



From the tree we can immediately write

$$\Pr(X \text{ present} \mid \text{test positive}) = 0.18 / (0.18 + 0.20) .$$

What Are The “Odds”?

In both gambling and gaming settings, you’ll often hear probabilities referred to indirectly through a reporting of the “odds” associated with a wager or the outcome of a game. The traditional method for stating the so-called “odds against” is to compare the chance that one might lose, if things work out unfavorably, to the chance that one might win: If p is your probability of winning, then the odds you are facing are $1-p:p$.

Odds are scalable, and are frequently scaled to whole numbers. For example, if your probability of winning a gamble is $1/3$, then the odds can be reported as $2/3:1/3$, but are usually reported as $2:1$ (and read as “the odds are 2 to 1 against you”) instead. The original probability can be easily recaptured from an “odds against” statement: Odds of $a:b$ correspond to a probability (of winning) of $b/(a+b)$.

Parimutuel Wagering

When you gamble at a casino, you are playing against the “house.” However, when you bet at a racetrack, be it on thoroughbreds, trotters, or greyhounds – I’ll assume horses, for purposes of discussion – you are betting against the other bettors. The amounts of money bet on the various horses determine the “track odds,” which are stated in terms of monetary outcomes: If a horse goes off at 4:1 odds, this means that a \$1 bet will bring you either a \$1 loss, or a \$5 payback (for a \$4 profit).

“Win” betting, under a parimutuel system, is simple to understand. The track collects all the bets, takes out a fixed percentage (as its earnings – and tax obligations – on the race), and returns all the residual money to those who bet on the winning horse (in proportion to their bets). For example, assume you bet \$20 on “Raging Bob” to win, and he does indeed win the race. Also assume that the track takes 20% of the total of all “win” bets. (Actual track percentages in the U.S. are typically fixed by local law, at somewhere between 15% and 20% of the total betting pool. The track percentage on “exotic” bets, such as exactas and quinellas, might be even higher.) If a total of \$20,000 was wagered, of which \$2000 was bet on “Raging Bob,” then the track keeps \$4000, and the remaining \$16,000 goes to the winning bettors, who get back \$8 for each dollar bet: You receive \$160, for a profit of \$140 on your bet.

Before the race begins, if the current bets are as stated above, the track will announce the odds on your pick as 7:1 (i.e., 140:20, scaled down). If you believe that the actual probability that your pick will win the race is greater than $1/(7+1) = 12.5\%$, then you will have a positive expected return from placing a small wager on this horse.

Before you race out to the track, note that the “final” odds determine the payoff, and that empirical research has shown that much of the “smart” money gets bet in the last minute before the betting windows close. Since the “posted” odds change as more money is bet on the race, and tend to lag the actual odds by about a minute, you can never be sure, when placing a bet, on the actual payoff you’ll receive if your bet wins.

Also note that your bet will itself shift the odds. This effect is greatest when the odds on a horse are “long”, e.g., 50:1. A large bet at the last minute could substantially reduce the prospective payoff if your “long-shot” wins the race.

A “place” bet is a bet that your horse will finish either first or second. The track collects all the “place” bets into a pool, takes its percentage, pays back the winning bettors (on either of the top two finishers) the amount they bet, and then splits the remaining money into two equal shares. The bettors on the horse that

won proportionately split one share, and the bettors on the horse that finished second proportionately split the other share. “Show” bets – that your horse will finish in the top three – are handled similarly.

Can you make money at the track? Of course you can, *IF* you are much better able to predict outcomes than most of the other bettors. (It’s not enough to be just a *bit* better, since you must overcome the track’s takeout.) Much statistical analysis has gone into the development of horserace betting systems. A good system must combine the ability to (probabilistically) predict race outcomes with the ability to determine when the posted odds differ enough from the prediction to offer a profit. A common feature of any sensible system is that it will often recommend not betting at all (when, for example, the posted odds are in line with the predictions of the system).

One business application of parimutuel betting is to forecasting: If you wish to derive a consensus forecast of the likelihood of one or the other of several events occurring, you could endow each of your forecasters with a budget, and require that they all place bets on the outcome. The resulting odds could then be taken as their actual (combined) forecast. Goldman Sachs and Deutsche Bank have recently opened an online enterprise where bettors can wager on such economic derivatives as the inflation rate, housing starts, or new unemployment claims. The odds that result can be used as a composite (across the bettors) estimate of the economic future

“Game” Odds

A common marketing technique is to offer consumers participation in a “collection game.” McDonald’s, for example, runs an annual “Monopoly” game in conjunction with Hasbro (the publisher of the board game). Stickers representing Monopoly “properties” are attached to food-product containers, and a customer who collects some particular combination of properties over the duration of the game can claim a prize (usually a free McDonald’s or Hasbro product). McDonald’s benefits directly from heightened consumer interest, and Hasbro indirectly by building name recognition for its product line.

A legal requirement (in most states) of such games is that the odds of winning the various prizes must be posted. The standard language states, “If you collect k stamps, the odds of winning prize P are 1 in x .” This translates directly to “the odds against winning are $x-1:1$.”