Handout

The setup

An *experiment* is a well-defined process with an uncertain outcome (for example, toss three fair coins).

The *sample space* is the set of possible outcomes (e. g. {HHH,HHT,HTH,THH,HTT,THT,TTH,TTT}).

An *event* is a set of outcomes (a subset of the sample space) (e. g. , A = {exactly one head} = {HTT,THT,TTH}).

An event is said to have occurred if one of the outcome it contains occurs.

Basic rules of probability

$0 \leq$ Pr(A) ≤ 1, for any event A.

Pr($\mathcal{S}$) = 1, for the sample space $\mathcal{S}$.

If A and B are *mutually exclusive*, Pr(A or B) = Pr(A) + Pr(B).

More rules

Pr(not A) = 1 – Pr(A)

Pr(A or B) = Pr(A) + Pr(B) – Pr(A and B)

Conditional probability

Pr(A | B) = "probability of A given B" = Pr(A and B) / Pr(B), provided Pr(B) > 0.

Independence

Events A and B are *independent* if Pr(A and B) = Pr(A) Pr(B), or equivalently, if Pr(A | B) = Pr(A) and Pr(B | A) = Pr(B).

Still more rules

Pr(A and B) = Pr(B) Pr(A | B) = Pr(A) Pr(B | A)

Pr(A) = Pr(A and B) + Pr(A and not B) = Pr(B) Pr(A | B) + Pr(not B) Pr(A | not B)

Bayes rule

Pr(A | B) = Pr(A) Pr(B | A) / Pr(B) = Pr(A) Pr(B | A) / { Pr(A) Pr(B | A) + Pr(not A) Pr(B | not A) }

An example

Suppose a test for HIV correctly gives a positive result, if a person is infected, with probability 99.5%, and correctly gives a negative result, if a person is not infected, with probability 98%.

(a) Suppose that 0.1% of a population are infected with HIV. Consider drawing a person at random and testing him or her for HIV infection. Calculate Pr(infected | test is positive).

Let I = { the person is infected } and P = { the person tests positive }.

From the information above, on the *sensitivity* and *specificity* of the test, we have
Pr(P | I) = 0.995 and Pr(not P | not I) = 0.98.

From this information,
Pr(not P | I) = 1 – Pr(P | I) = 0.005 and Pr(P | not I) = 1 – Pr(not P | not I) = 0.02.
Note further that Pr(I) = 0.001, and so Pr(not I) = 0.999.

We seek to calculate Pr(I | P). We use Bayes's rule. {Why use Bayes's rule here? Because we want to "turn around the conditioning." We want to write Pr(I | P) in terms of things like Pr(P | I).}

Pr(I | P) = Pr(I) Pr(P | I) / [Pr(I) Pr(P | I) + Pr(not I) Pr(P | not I)]
= 0.001×0.995 / ($0.001 \times 0.995 + 0.999 \times 0.02$) $\approx$ 4.7%

(b) Consider a person drawn from a high-risk group, so that they have, *a priori*, probability 30% of being infected. Calculate Pr(infected | test is positive).

In this case, we have Pr(I) = 0.3 and Pr(not I) = 1 – Pr(I) = 0.7.

Thus, Pr(I | P) = 0.3×0.995 / ($0.3 \times 0.995 + 0.7 \times 0.02$) $\approx$ 96%

Mutually Exclusive vs Independent

1. Suppose that events A and B are *mutually exclusive*. Calculate, in terms of Pr(A) and Pr(B),

   (a) Pr(A or B) :
   $\longrightarrow$ Pr(A or B) = Pr(A) + Pr(B).

   (b) Pr(A and B) :
   $\longrightarrow$ Pr(A and B) = 0, since they can't both happen.

2. Suppose that A and B are *independent*. Calculate, in terms of Pr(A) and Pr(B),

   (a) Pr(A and B) :
   $\longrightarrow$ Pr(A and B) = Pr(A) $\times$ Pr(B).

   (b) Pr(A or B) :
   $\longrightarrow$ Pr(A or B) = Pr(A) + Pr(B) – Pr(A) $\times$ Pr(B).

3. Suppose that A and B are both *mutually exclusive and independent*. What can we say about Pr(A or B), Pr(A and B), Pr(A), and Pr(B)?

   Since A and B are independent, Pr(A and B) = Pr(A) $\times$ Pr(B). But since A and B are mutually exclusive, Pr(A and B) = 0. Thus Pr(A) $\times$ Pr(B) = 0. And so either Pr(A) = 0, or Pr(B) = 0, or both.

   In other words, either A or B (or both) *cannot happen*!

   The point: Generally when we talk about independent events, they are *not* mutually exclusive, and vice versa.