Throughout the class, in addition to readings, I will give you frequent handouts to supplement our discussion. These will often contain additional informations (and occasional references) which are beyond the scope of in-class discussion.

The Binomial Distribution

In class, we discussed the example of flipping a fair coin. This means that each flip is independent of the previous flip, and that each flip has a probability, $p = 0.5$. Now, if all events are random, then it is easy to see that if I flip a fair coin twice, there are 4 ways that the results can come out:

HH HT TH TT

And that only 1 of those combinations produces two heads. Now, consider:

$$\text{Probability of an outcome} = \frac{\text{Number of ways to get the outcome}}{\text{Number of total outcomes}}$$

In other words, the probability of getting two heads in a row is 1/4=25%. The probability of getting one of each is 1/2=50%, and the probability of two tails in a row is 25%. We call this sort of result a *binomial* statistic, since in each case, the coin can either come up heads or tails, but it must come up one of them ("bi" is the Greek prefix meaning "two").

Now, consider the probabilities if we flip many more times. For example, if we flip 10 times, we have 1024 combinations! The reason is that each flip can be either heads or tails, and thus each flip doubles our total sequence. Or, for those of you who are mathematically minded:

$$\text{Number of total outcomes} = 2^N$$

where we'll use the variable, N to represent the number of times that we're flipping the coin.

Now, this gets very, very large, very quickly, and thus, since there is only 1 way to get, say, all heads, then it is very, very unlikely to get all heads over a long series of flips.

Beyond that, calculating the exact odds, of say, 4 heads in a series of 10 flips, is a bit tricky (and beyond the scope of the course). But we can see what the relative probabilities of the different outcomes might look like:

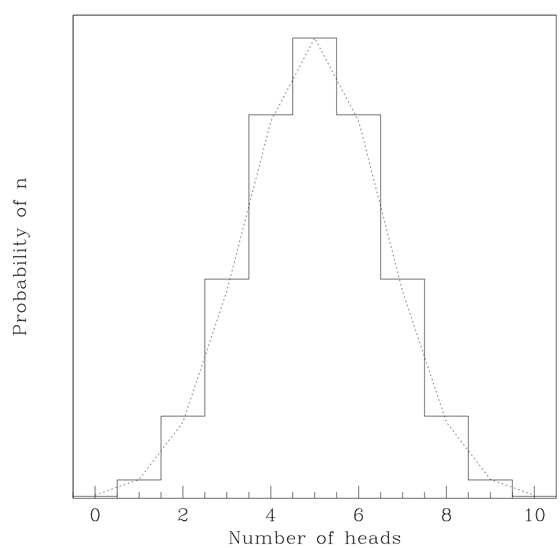

You will notice several things about this *distribution function*. First, the most likely scenario is that you will get 5 heads – exactly half. Secondly, it is still fairly likely that you well get 4 or 6 (or even 3 or 7) heads, but fairly unlikely that you will get 0,1,2,8,9 or 10 heads. But it *could* happen.

We describe the average (in this case, the most likely) outcome as the *mean*, and usually label it with the greek letter, μ (mu). We call the width of the distribution, the standard deviation, and usually give it the Greek letter, σ (sigma).

What does the error mean? Well, a $1-\sigma$ error means if we did the experiment a gazillion times, 68% of the times, the result would be in that range. Furthermore, 95% of the time the result will be within $2-\sigma$ of the average. Most scientists use a "$2-\sigma$" result as an acceptable standard for an experiment.

Now, you'll notice something else. In addition to a solid distribution, there is a smoother looking dotted line which roughly follows the same distribution. This is known as a normal distribution, and obeys the mathematical relation:

$$p(n) \propto e^{-\frac{(n-\mu)^2}{2\sigma^2}}$$

I know. It looks ugly. That's why you should just remember what μ and σ *mean*.

Now, the cool thing is that if you flip lots and lots of times, the binomial distribution looks like a normal distribution. If, instead of a fair coin, we have a probability p of coming up heads, we find:

$$\mu = Np \tag{1}$$

(which we've already found), and

$$\sigma = \sqrt{Np(1-p)} \tag{2}$$

which looks complicated unless you plug in $p = 0.5$, and you see:

$$\sigma = \sqrt{N}/2$$

How well does this work? Well, let's plug in for 50 and 100 coin flips:

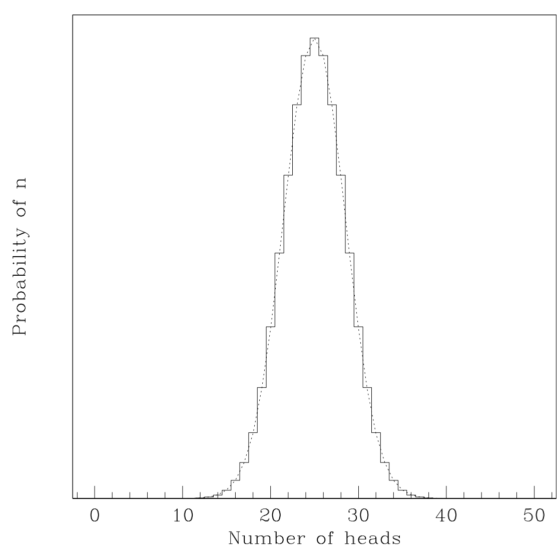

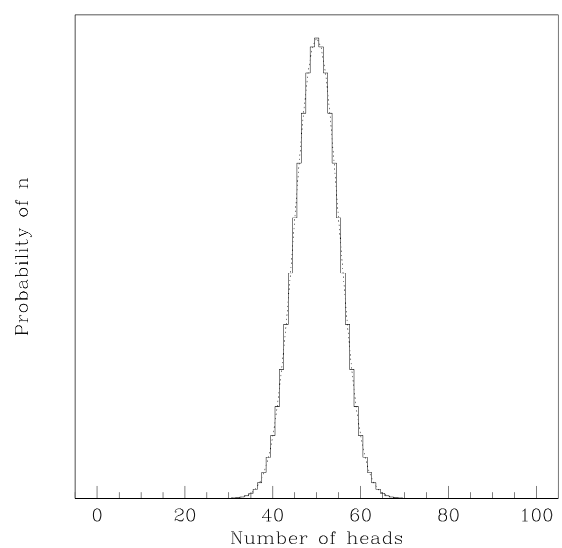

Pretty good, eh?