More about the Normal distribution (section 1.3 plus a bit)

We wait for Chapter 5 for the details of the fact that averages tend to have Normal distributions. We will see that a related phenomenon is that many measurements have a distribution that is well approximated by a normal distribution (that is, the measurements themselves may vary according to the Normal curve, not only the sample averages of measurements). For now, you will have to take on faith that the Normal distribution deserves special consideration.

Here is the thing that you need to be able to do with any random variable that can be assumed to have a normal distribution:

If X is Normal with mean μ and standard deviation σ, then the probability $P(a<X<b)$ can be computed from the $N(0,1)$ table in the following way:

First re-express $P(a<X<b)$ in terms of a RV Z that does have a $N(0,1)$ distribution.

Let $Z=(X-\mu)/\sigma$. Then if X is $N(\mu,\sigma)$, Z will be $N(0,1)$. So $P(a<X<b)$ is the same as $P((a-\mu)/\sigma <Z<(b-\mu)/\sigma)$. We can look this up in the $N(0,1)$ table p 740-741.

Example: X has mean 100 and sd 15. What is $P(85 < X < 115)$?

Compute

$(85-100)/15 = -1$

$(115-100)/15 = +1$

So $P(85<X<115) = P(-1<Z<+1) = .8413-.1587 = .6826$

Every Normal distribution has the property that

A proportion 0.6826 lies within 1 SD of its mean.
A proportion 0.9544 lies within 2 SDs of its mean
A proportion 0.9974 lies within 3 SDs of its mean

and as an approx rule of thumb:

68% within 1 SD
95% within 2 SD
100% within 3 SD

In Chapter 5 we will see that the Binomial RV can be understood to be proportional to a certain sample average, and since averages tend to have normal distributions, it is not surprising that the normal distribution should in some cases provide a good approximation to the Binomial distribution. This fact has some practical utility although it is this utility is fairly minor when one has access to statistical software. However, we will still pursue it for the theoretical importance.

The Normal Approximation to the Binomial Distribution

The normal approximation to any distribution is done by plugging in the actual mean and variance from the target distribution into the Z formula: Z=(X-μ)/σ

Now for the Binomial with number of trials n and probability of success p, the mean is np and the sd is $\sqrt{np(1-p)}$. So to compute P(X<a) for example, one computes

$$P(Z < \frac{(a-np)}{\sqrt{np(1-p)}})$$

For example, if I toss a fair coin 25 times, what is the chance that I get 10 or fewer heads?

P(X≤10) = $P(Z \le \frac{(10-25*0.5)}{\sqrt{25*0.5(1-0.5)}})$ = $P(Z \le \frac{-2.5}{\sqrt{6.25}}) = P(Z \le -1)$=.1587

In other words, I get 10 or fewer heads in almost 16% of such experiments.

Now we could have computed this directly from the Binomial Distribution:

See Table A.1 pp 736-738 – actually the part headed n=25 on p 738.

P(X≤10) = .212

Now this is not very close to .1587. But that is because we were a little sloppy in applying the normal approximation. Look at the picture on p 169 at the bottom. The area under the curve (Normal) is trying to approximate the area under the rectangles (Binomial). So to compute X≤10, we need to go to 10.5 on the normal curve to catch the comparable area for the binomial probabilities. So if we redo the normal calculation,

P(X≤10) = $P(Z \le \frac{(10.5-25*0.5)}{\sqrt{25*0.5(1-0.5)}})$ = $P(Z \le \frac{-2.0}{\sqrt{6.25}}) = P(Z \le -0.8)$=.0.2119

which is very accurate indeed!

When n is large, the .5 correction will make less of a difference.

The approximation is good as long as the binomial is not too skewed left or right. A rule of thumb for this is that both np and n(1-p) is at least 10. In our example, we had np=12.5 and n(1-p)= 12.5 , and we saw the approximation was excellent, once properly done.

PS: Have you heard of the MENSA organization? They admit only members who can prove their IQ is in the top 2% of the population. If IQ has X has mean 100 and sd 15, then you can check that an IQ of at least 131 is necessary for membership.