Statistical Methods for Quality Control

I. Acceptance Sampling

In general:

Acceptance sampling is used to decide whether to accept or reject shipments ("lots") of raw materials, parts, or finished goods from suppliers. The basic idea is that the recevier (us) has a procedure for 1.) selecting a sample of items from the shipment for testing 2.) testing the items selected 3.) making a decision on whether the lot is good enough to accept or bad enough to reject (& send back to the supplier). Acceptance sampling is concerned with setting the standards for the decision in step 3. We are looking at *attribute sampling* - the items tested are rated as *acceptable* or *defective* (no further measurement) and we are interested in the proportion of defective items.

Specification and use of a sampling plan:

A *sampling plan* (in our situation) consists of two values: a *sample size* (n) and an *acceptance criterion* (c) . Sample size should be clear. The acceptance criterion is the largest number of defective items allowed in the sample while allowing it to be accepted. The text shows calculations for an $n = 15, c = 0$ plan. With (for example) an $n = 20, c = 3$ plan, we would take a sample of 20 items and would accept the lot if 3 or fewer defectives were found, reject the lot if 4 or more defectives were found.

Each sampling plan has an *Operating Characteristic Curve*- x represents the (unknown) actual percent of defectives in the lot (that is, p) , and y is the probability that the sampling plan would accept such a lot [that is $P(\#\text{defectives } \leq c) = P(\#\text{defectives } = 0) + P(\#\text{defectives } = 1) + \ldots + P(\#\text{defectives } = c)$] . If the lots are large, we can use the binomial probability distribution to find the height of the curve because $P(\#\text{defectives } = m) = C(n, m)p^m(1 - p)^{(n-m)}$. The text shows this curve for $n = 15, c = 0$ on one graph in figure 20.12 on p. 869 and a comparative graph for this plan and four others in figure 20.13 on p. 870. Working by hand, we could get probabilities for several p values and sketch a smooth curve through them)

Setting up a sampling plan:

Setting the standards involves decisions about four values - two are measurements of quality (percent defective items) and two are measurements of risk (probability of wrong decisions - probability of type I & type II errors, in classical hypothesis testing). There is an "always acceptable" level - the text refers to this as p_0 , it is often called the "Acceptable Quality Level (AQL)" - any lot that contains a this percentage *or less* of defective items should be accepted.

There is also an "always rejectable level" - the text refers toi this as p_1 - called "Lot Tolerance Percent Defective (LTPD)" or "Rejectable Quality Level (RQL)" - any lot that contains this percentage *or more* of defective items should be rejected.

The risk levels of a plan are defined in terms of these quality levels: The *producers risk* (α) is the probability that a lot which is of acceptable quality (percent defectives $\leq p_0$) will be (wrongly) rejected, causing loss of sales (and probably shipping costs) to the producer/seller.

The *consumer's risk* (β) is the risk that a lot of unacceptable quality (percent defectives $\geq p_1$) will be (wrongly) accepted, causing excessive production or sales losses to the consumer/receiver.

All four of these values are involved in setting a sampling plan. Other considerations (which we will not factor in directly) are cost of testing (smaller samples are cheaper to test), cost of shipping/reshipping, desirability of keeping customers/suppliers happy, actual cost (to the "consumer") of using defective items, etc.

To read the characteristics of a plan: p_0 gives a point on the x-axis. The *Producer's Risk* is the distance down from the top (probability 1.00) to the point on the OC curve at this value. (Probability of rejecting a lot with proportion p_0 of defectives – an acceptable lot – is 1 − (probability of accepting a lot with p_0 defectives)). p_1 also gives a point on the x-axis. The *Consumer's Risk* is the height of the curve above this point (probability of accepting a lot with p_1 defectives - an unacceptable lot). That is $\alpha = 1 -$(height at p_0) and $\beta =$ (height at p_1)

In general, selection of a plan involves looking at OC curves for *several* plans to find one that gives the desired combination of p_0 with α and the desired combination of p_1 with β . As can be seen from the graphs on page 870 , sample size does not tell the whole story, though larger sample sizes tend to give us more chance to match the desired values. The figure on p871 shows reading the α and β values for the $n = 15, c = 0$ OC when $p_0 = .03$ and $p_1 = .15$