$p = ?$

<u>1-sided (single limit) cases:</u> (single lower limit L)

$$p = \Phi\left(\frac{L-\mu}{\sigma}\right)$$

Depending upon set-up p is either a function of μ only or a function of both μ <u>and</u> σ.

<u>2-sided (double limits) cases:</u>

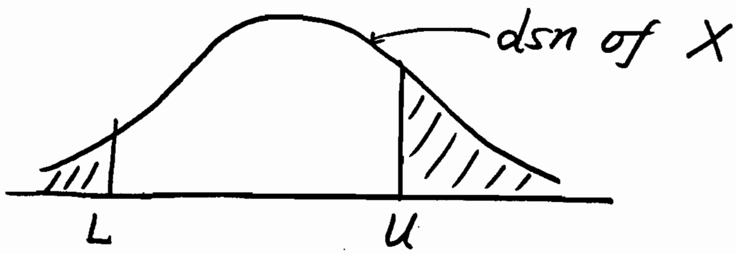

$$p = 1 - \left[\Phi\left(\frac{U-\mu}{\sigma}\right) - \Phi\left(\frac{L-\mu}{\sigma}\right)\right]$$

Again, depending upon scenario this is either a function of μ alone, or a function of μ and σ.

Single Limit "σ known" cases (lower limit version)

Small μ makes p big; small $\bar{x}$ suggests that μ is small.

So we should reject for small $\bar{x}$; i.e.,

reject lot if $\bar{x} < K$

accept lot if $\bar{x} > K$.

$p(\mu) = \Phi\left(\frac{L-\mu}{\sigma}\right)$ (a function of μ alone)

$P_a(\mu) = P[\bar{x} > K] = 1 - P[\bar{X} < K]$

$= 1 - \Phi\left(\frac{K-\mu}{\sigma/\sqrt{n}}\right)$ (a function of μ alone)

So as I vary μ

$(p(\mu), P_a(\mu))$

traces out a path in (p, P_a)-plane that I can use as an OC-curve.

See §8.2 of V&J for classical 2-point design problem (choosing n, K to get OC-curve to pass through 2 points of interest). (page 471)

Double Limits "σ known" cases:

μ far from $\frac{U+L}{2}$ makes p big.

Therefore one might reject a lot if $\bar{x}$ is too far from $\frac{U+L}{2}$; i.e.,

accept lot if $\frac{U+L}{2} - k < \bar{x} < \frac{U+L}{2} + k$

reject lot otherwise

$$p(\mu) = 1 - \left[\Phi\left(\frac{U-\mu}{\sigma}\right) - \Phi\left(\frac{L-\mu}{\sigma}\right)\right]$$

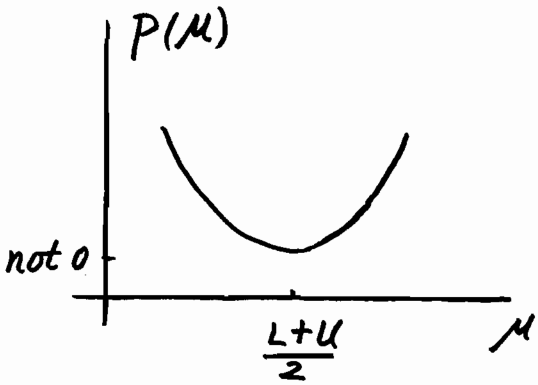

p is minimum at $\frac{L+U}{2}$ and there are 2 μ's giving each value of p (these are symmetric about $\frac{L+U}{2}$).

$$P_a(\mu) = P\left[\frac{U+L}{2} - k < \bar{x} < \frac{U+L}{2} + k\right]$$

$$= \Phi\left(\frac{\frac{U+L}{2} + k - \mu}{\sigma/\sqrt{n}}\right) - \Phi\left(\frac{\frac{U+L}{2} - k - \mu}{\sigma/\sqrt{n}}\right).$$

$P_a(\mu)$ is also symmetric about $\mu = \frac{L+U}{2}$.

Thus, as in the single limit case, I can let μ vary and

$$(p(\mu), P_a(\mu))$$

traces out a path in the (p, P_a)-plane that can serve as an OC-curve (both μ's giving a particular p give the same P_a).

—

See V&J for classical (2-point) design (choice of k and n). (page 471)

<u>Single Limit "σ unknown" Cases</u>: (lower limit version)

Note that p depends upon (μ, σ) through "z-value associated with L"

$$\frac{L-\mu}{\sigma}.$$

When this is small (large negatively), all is well, i.e., p is small.

So a plausible acceptance sampling criterion is to accept if the sample version of this is small; i.e.,

$$\text{accept if } \quad \frac{L-\bar{x}}{s} \quad \text{is small.}$$

That is, accept if $\frac{\bar{x}-L}{s}$ is large ($>k$);

i.e., $\boxed{\begin{array}{l} \text{accept if } \bar{x}-L > ks \\ \text{reject if } \bar{x}-L < ks \end{array}}$

$$P_a(\mu,\sigma) = P\left[\frac{\bar{x}-L}{s} > k\right]$$

$$= P\left[\frac{\frac{\bar{x}-\mu}{\sigma/\sqrt{n}} - \frac{L-\mu}{\sigma/\sqrt{n}}}{\frac{s}{\sigma/\sqrt{n}}} > k\right]$$

$$= P\left[\frac{\text{a } N(0,1) \text{ r.v.} - \frac{L-\mu}{\sigma/\sqrt{n}}}{\sqrt{\frac{\text{a } \chi^2_{n-1} \text{ r.v.}}{n-1}}} > \sqrt{n}\,k\right]$$

independent

$$= P\left[\begin{array}{c}\text{a noncentral} \\ t_{n-1}\left(\sqrt{n}\left(\frac{\mu-L}{\sigma}\right)\right) \text{ r.v.}\end{array} > \sqrt{n}\,k\right]$$

So both p and P_a depend on (μ,σ) only through $\underline{z = \frac{L-\mu}{\sigma}}$.

— All (μ,σ) with a given p have the same P_a and thus once again I can get an "OC curve" by plotting

$$(p(z), P_a(z))$$

for various z.

Note as an alternative to exact (noncentral t) calculations, it has been common in SQC to use <u>the Wallis approximation</u>:

$\bar{x} \pm ks$ is "approximately" normal with

mean $\mu \pm k\sigma$

variance $\sigma^2\left(\frac{1}{n} + \frac{k^2}{2n}\right)$.

Note that $Var(s) = \sigma^2 c_5^2(n)$

also $Var(s) \approx \left(\frac{1}{2\sqrt{ES^2}}\right)^2 Var(s^2)$

$$= \frac{1}{4\sigma^2} \cdot \frac{2\sigma^4}{(n-1)} = \frac{\sigma^2}{2(n-1)}$$

$$\approx \frac{\sigma^2}{2n}$$

$\Rightarrow Var(\bar{x} \pm ks) \approx \sigma^2\left(\frac{1}{n} + \frac{k^2}{2n}\right)$.

Using the Wallis approximation,

$$P_a(\mu,\sigma) = P[\bar{x} - L > ks]$$

$$\overset{\shortparallel}{P_a(z)} \quad = P[\bar{x} - ks > L]$$

$$\approx 1 - \Phi\left(\frac{L-(\mu-k\sigma)}{\sigma\sqrt{\frac{1}{n} + \frac{k^2}{2n}}}\right)$$

$$= 1 - \Phi\left(\frac{\frac{L-\mu}{\sigma} + k}{\sqrt{\frac{1}{n} + \frac{k^2}{2n}}}\right) \qquad z$$

Design / Choice of n, k: classical 2-point design problem is "solved" in V&J using the Wallis approximation. (page 472)

Double Limits "σ unknown" Cases

This is a messy problem without clean answers. It's not even obvious what kind of acceptance regions to try.

Note that a picture like this has motivated plans with similar acceptance regions (in the $(\bar{x}, s)$-plane).

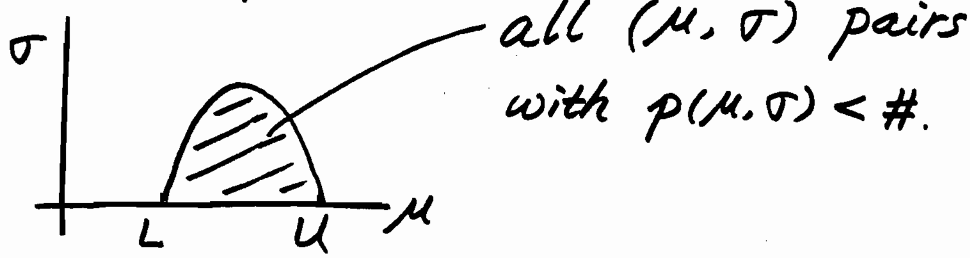

E.g., Duncan suggests

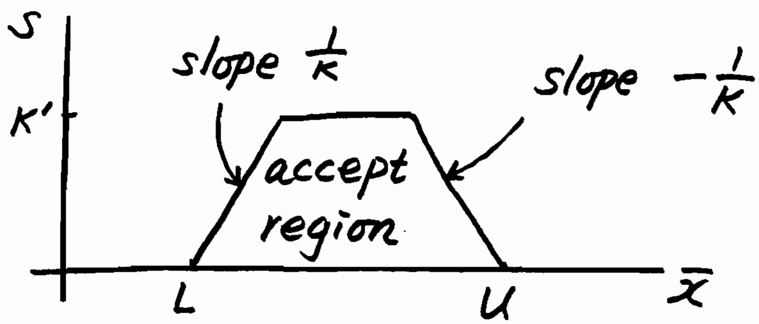

Others have studied plans that for some estimator of p, say $\hat{p}(\bar{x}, s)$,

accept lot if $\hat{p} < K$

reject lot if $\hat{p} > K$.

A common choice for $\hat{p}$ is the UMVUE of p.

— This is used, e.g. in the famous Military Standard 414 Variables Acceptance Sampling Plans.

— Such plans (based on $\hat{p}$) usually lead to pictures like

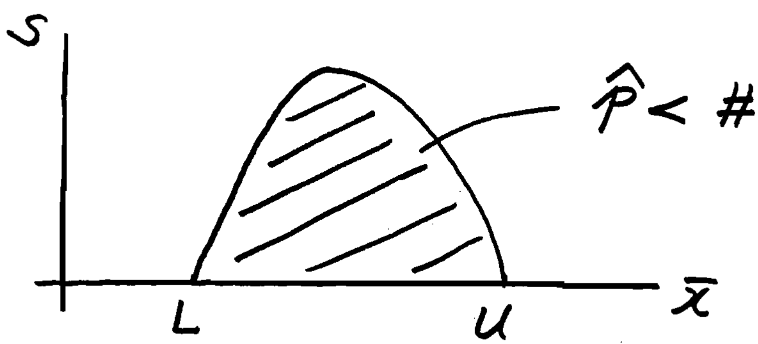

But it's not easy to describe them or their statistical properties.

Worse yet no one knows how to cook up a method for which every (μ, σ) pair with a given p has the same P_a; i.e., one has to deal with OC "bands."

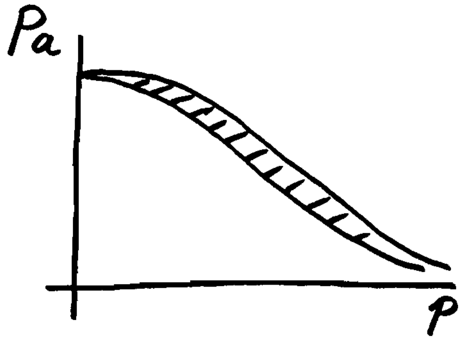

For the best of these this band may not be so wide — perhaps thin enough to ignore this issue.

Philosophy / Caveats

1) In comparison to corresponding attributes plans, variables acceptance sampling plans typically provide huge sample size savings.

   This is particularly true where my concern is small p.

2) But this stuff depends critically on the parametric model assumptions (particularly on the shape of the extreme tails of a distribution) that can only be checked with a huge sample size.

3) Besides, it's not so clear that treating at (μ, σ) with a given p equally is rational.

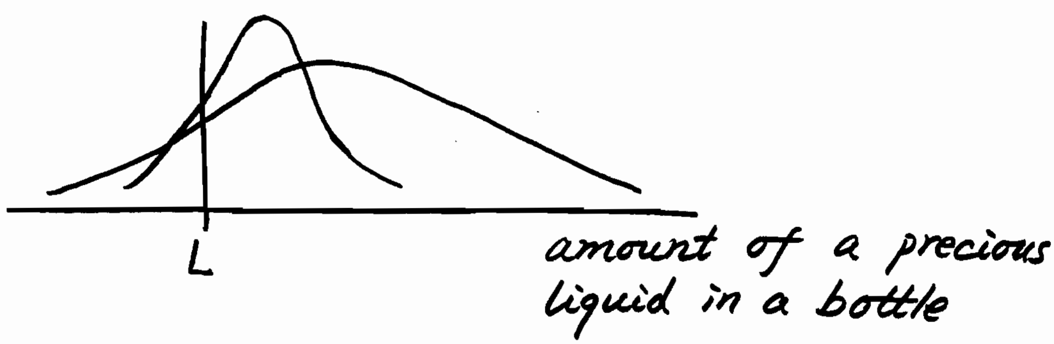

4) This whole enterprise (variables acceptance sampling) comes unglued if I admit any possibility of measurement error.

$X \sim N(\mu, \sigma^2)$ ← dsn of the "real" variation in population

If what I observe is

$$\underset{\nearrow}{y = x} + \underset{\nwarrow}{\varepsilon}$$

$N(\mu, \sigma^2)$ $N(0, r^2)$

independent

I end up with $\bar{y}$ and s_y (not $\bar{x}$, s).
As long as $r^2 > 0$ there will be σ^2's small enough so that population variability hides in measurement noise.

One winds up with very unpleasant "OC bands."

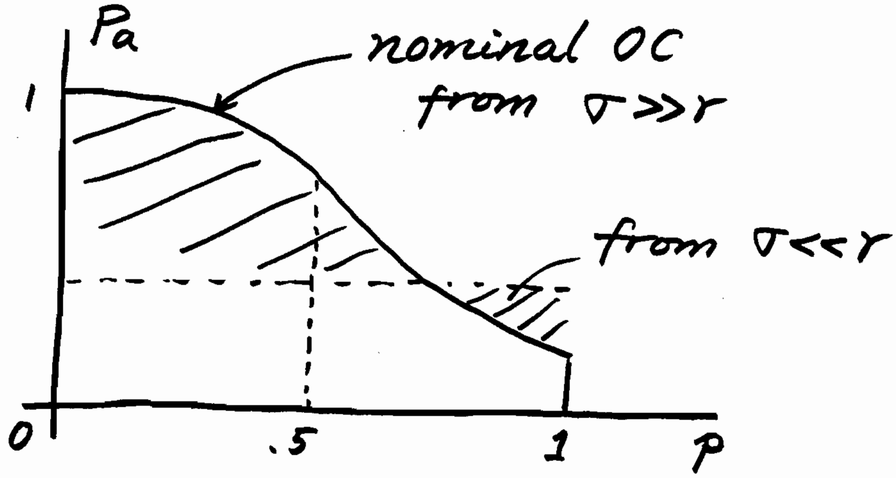

The rest of V&J is about experimental design and analysis (applied to process improvement).