

Re: quality control

Source: <http://sci.tech--archive.net/Archive/sci.stat.math/2006-11/msg00384.html>

- *From:* David Winsemius <doe_snot@xxxxxxxxxxxx>
 - *Date:* Tue, 14 Nov 2006 10:00:53 -0600
-

Richard Ulrich <Rich.Ulrich@xxxxxxxxxxxx> wrote in
news:6fki12560tr7t2uac2burppvem3eohol2t@xxxxxxxx:

On Mon, 13 Nov 2006 10:04:03 -0600, David Winsemius
<doe_snot@xxxxxxxxxxxx> wrote:

Richard Ulrich <Rich.Ulrich@xxxxxxxxxxxx> wrote in
news:vnsfl2d685snrulj63m9iel784mhgnjop8@xxxxxxxx:

On 12 Nov 2006 17:00:04 -0800, "Frank"
<deps_bear@xxxxxxxxxxxx> wrote:

If I know a product fails .01% of the time
and I have 1500 items
I'm running through a process. How many
items do I need to check
with, say, 99% confidence that all the items
are built correctly.

How many failures do you expect? Almost always, zero.
This is dealing with exact probabilities. For a higher failure
rate, you might want to look at the p of success, and raise
to a power, e.g., $(.9999)^n$. For the tiny p of 0.01%,
the figuring can be pretty much additive

You want to have only so many items *unchecked* that
there
will be, on the average, only 1 bad item in 100 samplings ---
so that 99 times out 100, there will be none.

You expect 1 failure in 10,000. One hundred samplings
that each fail to test 100 items will meet that condition.
So you need to check 1400 of each 1500.

Re: quality control

How did you go from 100 samplings of size 100 to the number 1400?

What I got was 100 samplings of 100, with a long-term average of 1 defect per 10,000. The 1500 and subsequently 1400 are largely irrelevant to the problem, as I construed it. I agree that it is not necessarily a good way to devise a real-life test of defects. As you point out, it is even more difficult to assure that there are *no* defects, than what you get with this strong assumption. Basically, you might as well go with 100% testing.

If the long term rate is 1 per 10,000, then any sample of 100 has a (very close to) 99% chance of being clean. That's easy arithmetic, following stern logic.

However you did it, you are then claiming that after examining 1400 items that you are 99% confident that there are no defective items in the remaining 100 items, when the past experience indicated that the failure rate was 0.0001 (so it was quite unlikely that you would have observed any even if all 1500 were examined? Why not stop at 1300? Seems to me that you might have incorrectly inverted the problem to one of reducing the size of the population at risk.

The best answer to a binomial acceptance sampling problem in a finite population that I have seen recently was given by Ted Harding in Medstats:

<http://groups.google.com/group/MedStats/msg/65e62e78dfe9b3f0?dmode=source&hl=en>

When I apply his method I get a confidence level of .9333 for an expected rate of 1/1500 and an observed of zero after 1400 were sampled and that is a much higher predicted rate (at least on a ratio scale) than the OP specified. (code in the R system.)

I took the expected rate of 1/10,000 as guaranteed --
I don't see how you fit that into a 2x2 table, which is what the hypergeometric describes, e.g., Fisher's Exact Test.

I used a higher event rate to set an upper bound on the confidence associated with an observed value of zero defects.

Do the entries of `phyper()` represent a 2x2 table in some fashion? For problems of finite samples, with larger numbers of errors, the counts can be construed as numbers (okay-seen, defects-seen; yet-to-see-okay,

Re: quality control

defects–yet–to–see) — for a table with cells (A,B; C,D).

Phyper(.) is a cumulative probability rather than a single table probability. In R the convention is for "p" functions to be cumulative probabilities and "d" functions to be densities. In this case x is zero and we have one table, but x could be set higher. The table is: (k, x ; N–k–M, M–x) so it is (1400, 0, 99, 1).

If 20 defects in a batch of 1500 were the highest acceptable to "upper management" and we wanted to know the sample size needed to assure that the rate in a single sample was below 20/1500 at a 95% level, you would set M=20 and adjust k to find the sample size that would give the desired level of confidence with zero defects.

That is a problem with the sample providing the rates.

```
N<-1500; M<-1; x<-0;
k<-1400; 1-phyper(x,M,N-M,k)
```

```
[1] 0.9333333
```

So even with a higher prior than specified, you cannot get close to 99%. If you assume a prior of 2/1500, you can get to 0.9956 after sampling 1400.

Huh? A higher prior *what*? I don't get this. A higher defect rate should not have a higher confidence of more zeros

I wouldn't use exactly that language. A higher event rate decreases the probability that zero will be an observed value. Apologies if "prior" was confusing. A higher assumed (or "acceptable". or "specified") defect rate raises the probable number of defects in a sample when k is of reasonable size. It changes the expected distribution of counts in a sample of fixed size. The question is finding a sample size that makes an observed value of zero become meaningful and to specify the nature of that meaning. If the defect rate is higher, then the probability of a sample of any size having _zero_ defects will become smaller. The question is: How to increase the "value" of zero_observed after a sample of (variable) size from a population of fixed size?

As you noted, it is not possible to use this to precisely address the question of probabilities in the original question, because you cannot set up an expected number of events properly. I was using it to set an upper bound. I still think (as I originally replied to Frank) that it

Re: quality control

Re: quality control

requires no math to appreciate that one cannot use an acceptance sampling method on a single sample when the sought-after, not-bigger-than-X, defect rate is materially lower than $1/\text{sample size}$.

I will admit that in reading the literature on exact and approximate binomial problems, it is not unusual to see multiple formulations with different answers. We do agree (and the OP apparently agrees) that the problem posed was not solved to his satisfaction. He has left the consulting room following yet another disappointing encounter with a sample size question. Another victim of unreasonably high expectations regarding the "power of statistics."

--

David Winsemius

.