Time Examined Class Notes – March 12, 2019 (md)

1. We have drawn a distinction between *accuracy* and *precision.* Precision, or *stability* has to do with the reliability of repeated measurements giving the same value. Accuracy has to do with the measurements being close, on average, to the true value.

For example, looking at the collection of measurements of X in the figures **A** & **B**,

***** X	* * *X * * *
A: High precision, Low Accuracy	**B: Low Precision, High Accuracy**

A: Values are closely clustered together, but their average is far from X.

B: Values are very spread out, but their average is close to X.

2. To quote J & F-R (p. 47) "Some resonators have good stability, others have good accuracy; the best, for clockmakers, must have both." Fortunately, High-Q resonator systems do have both good stability *and* good accuracy. (see pp. 47 ff of J & F_R))

Stability: A High-Q system will only oscillate at f very near f_0, and it will maintain its oscillation for many cycles, so it continues to give the same f_0 repeatedly.

Accuracy: A High-Q system can only be stimulated by an energy source with f very near f_0, and a feedback system keeps adjusting the energy frequency to keep the power transfer at a maximum, so that the system stays at f very near f_0.

2. Problems in precision are usually due to statistical uncertainties in the measurement process, factors that might vary from measurement to measurement. Problems in accuracy are usually due to systematic uncertainties, in which an instrument or a part of the method are not working the way you think they should. These two types of uncertainties require different responses as the experimenter tackles them.

Systematic uncertainties are dealt with by examining the instruments and methods with a critical eye and refining the experimental process. We have about 150 years of experience at dealing with statistical variation. Let's look at an example.

4. Imagine a measurable quantity has a particular value, but with some statistical spread to it. For instance, "what is the average height of all 1800 Bard students" has an answer, but it would be hard to find. So, we might take a sample of students and find their average height. How good an estimate of the true average would that be?

We did a simulation in class. I had a set of 300 values of a variable x. As the frequency histogram shows, the values are spread out between x = 40 and x = 160, with a bell-shaped curve peaking around x = 100. We call this shape a "normal distribution", and it is a common data pattern in science measurements. This **population** has an **average** value μ (the center of the distribution) and a **standard deviation** σ (the spread of the distribution). By the way, the Full Width at Half Maximum of the normal distribution is FWHM = 2.2 σ.

I gave each student a sample of 20 values from that population, and I had already summed up the average of the first 5 entries on the sheet, and the average of all 20 entries. We then looked at the spread of these averages.

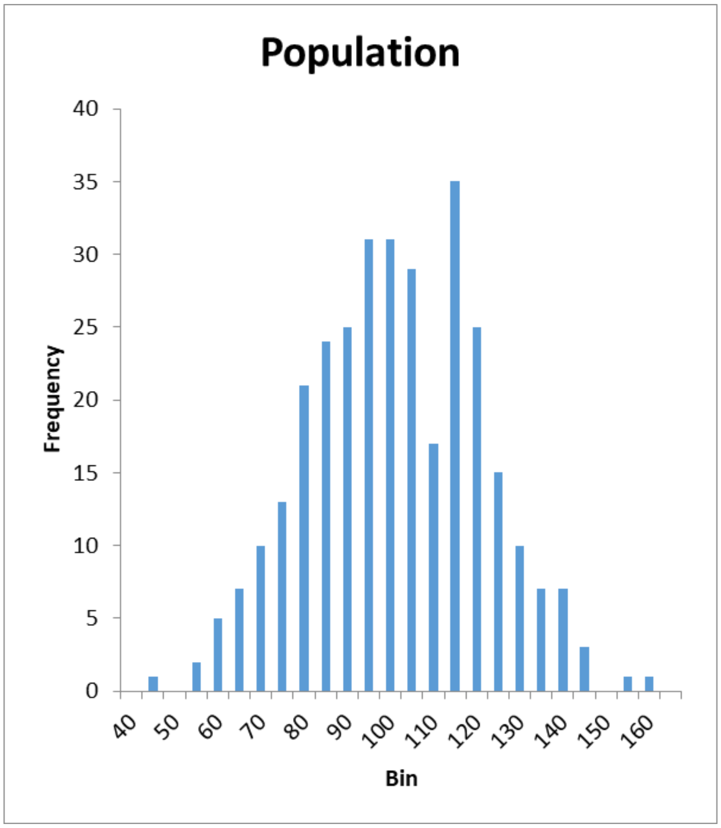

5-Avg		20-Avg
103.8		90.9
100.8		103.9
99.8		93.9
105.0		100.6
100.6		99.0
100.4		101.1
101.0		99.3
94.8		99.0
89.4		93.7
102.4		103.8
107.8		101.2
97.4		98.7
95.2		98.6
87.4		99.3
119.4		104.1
99.8		101.4
100.3	**Average**	**99.3**
7.35	**St. Dev.**	**3.74**

5. In each case, the averages clustered around x = 100, as expected, but the averages of 20 values were more tightly clustered than the averages of 5 values:

5-value averages	Average = 100.3	Standard Deviation = 7.4
20-value averages	Average = 99.3	Standard Deviation = 3.7

This shows a trend that we will see in most sampling situations. If we take a sample of size N, the average we find is our best estimate of the actual average value of the population, the number we wanted to measure. Our uncertainty in that answer gets better if we take a larger sample, but it does so by a factor of the square root of N. In this case, going from a sample of 5 to a sample of 20 increased N by a factor of 4, and so it reduced s by a factor of 2. I can't promise that the statistics will always work out so well, but the general effect is true: Increasing sample size improves our estimates, but only proportional to the square root of the increase.

6. We then moved on to a discussion of the Cesium Atomic Fountain Clock (F1) used by the National Institute of Standards and Technology (NIST) as the provider of the time standard for all of the U.S. (and many other countries). The detailed notes are in the handout I culled from the NIST website:

https://www.nist.gov/pml/time-and-frequency-division/primary-standard-nist-f1

and some of the details of the processes are in Chapters 5 & 7 of J & F-R.

Highlights of the Clock Details:

- Cesium-133 is an atom with a single unpaired electron orbiting the nucleus. The nucleus and electron each have a spin, acting like a small magnets.
- If the two magnets align parallel (both with North up), this gives a slightly higher energy than if they align anti-parallel (North up for one, South up for the other).
- A transition between the upper and lower energy states of this *hyperfine* energy difference corresponds to a photon of frequency 9,192,631,770 Hz.
- The Cesium atoms are cooled down to microKelvin temperatures by lasers, like the optical molasses described in J & F-R, Chapter 7. (nice animation online)
- Other lasers make the fountain effect, kicking the atoms up vertically and letting them drop down, passing twice through the microwave cavity that acts as the energy source for the resonator, over a period of time of about 1 second.
- The emissions from the excited atoms are then detected. The entire system has a feedback loop that continually tunes the absorption/emission frequencies to maximize the power absorbed by the stimulated atoms.
- The cold atoms and the feedback system tuning give a Q-value of 3×10^{15}.