CLASSROOM CAPSULES

EDITOR

Michael K. Kinyon
Indiana University South Bend
South Bend, IN 46634

Classroom Capsules consists primarily of short notes (1–3 pages) that convey new mathematical insights and effective teaching strategies for college mathematics instruction. Please submit manuscripts prepared according to the guidelines on the inside front cover to the new Editors Ricardo Alfaro and Steven Althoen, University of Michigan-Flint, Flint, MI 48502.

Distortion of average class size: The Lake Wobegon effect

Allen Schwenk (schwenk@wmich.edu), Western Michigan University, Kalamazoo, MI 49008

Universities love to boast about their presumably small average class size. And how often are parents stunned to hear how large their freshman's average class size seems to be? And since they are using the same words, we can safely assume that they are reporting the same thing. Can't we?

Actually, no! There is a distorting phenomenon that occurs when computing average class sizes that needs to be widely circulated. Namely, the average reported by the university, which is the same as the average experienced by the faculty, is significantly different from the average experienced by the students. Rather than try to explain this at the abstract level, let me present a hypothetical example.

A certain (remarkably small) university has precisely 200 students. All 200 are taking the same five courses, say English, Mathematics, Economics, History, and Psychology. Suppose that History and Psychology are taught in large lectures of 200, but all the other classes are taught in small sections of 20 students each. What is the average class size?

Course	Number of students in each section	Number of Sections	Total number of students in this course
English	20	10	200
Mathematics	20	10	200
Economics	20	10	200
History	200	1	200
Psychology	200	1	200
Totals		32	1000

The average class size is $\frac{1000}{32} = 31.25$. This is the figure an administrator would report. It is also the figure that the faculty experience when reporting their teaching loads, since each class is reported by one instructor. It accurately describes the experience of the university in reporting how large, on the average, are our classes. However, it is a serious error to presume that a typical student experiences the same "average." In our example we have 200 students all having the same class size experience. Each one of them has five classes, three with 20 students and two with 200. Thus, they all compute their personal average class size as $\frac{460}{5} = 92$. Observe that every single student in our example experiences an average class size nearly three times as large as the university claims! How can this be? How can every student have classes that are so much larger than average? Shades of Lake Wobegon, where the children are all above average!

The answer is that when the institution computes average class size it counts each class exactly once, but when students compute their own personal average class size the large lectures get experienced and reported 200 times while each small class gets reported only 20 times. This shifts the average considerably toward the higher end. Neither the university nor the student is wrong or deceitful, but the averages they report are very different animals. The average experience of the faculty and the administration is truthfully represented by the university average; the typical student-experienced class size is not. In all fairness to our students, we must understand this difference, and we must avoid the temptation to compute the average class size in the customary manner and then shape our opinions and policies on the premise that this average will somehow reflect the typical student experience.

My example is highly contrived to keep the arithmetic simple, but I can prove mathematically that the student-experienced average is always greater than the institution average. How much greater depends upon the disparity between the largest and the smallest classes. If classes are nearly uniform, the difference is insignificant, perhaps only a tenth of a point. But when there is a tenfold disparity, the student experience can be more than double, as my example illustrates. Of course in the real world we don't have all students taking precisely the same courses, but the effects are similar. If you don't trust my selection of an example, construct your own. I guarantee a similar discrepancy, though possibly less severe.

Before giving a rigorous comparison, we need to specify how we compute the average class size experienced by students. This is not a trivial issue. Consider two contrived cases.

Case 1. Ann is enrolled in one class that has 20 students, while Bob has a schedule of five classes, each with 200 students. Clearly Ann's average is 20 and Bob's is 200. Together they average 110.

Case 2. Carl has a single class of size 200, while Diane is in five classes of size 20. Again, the average size for both students is 110.

Something doesn't seem quite right here. We get the same average in each case, but doesn't Case 1 suggest mostly large classes while Case 2 has mostly small? The problem is that we are treating each student's experience as having equal weight when we compute these averages. But the one-class experience of Ann and Carl shouldn't count as heavily as the 5-course experience of Bob and Diane. And if we have a student who has taken 40 classes, that experience should influence the average more strongly than any of these. To reflect this observation, we choose to use a weighted average by weighting each student's personal average by the number of courses he or she has taken. Thus in Case 1 we get $\frac{1\cdot 20+5\cdot 200}{1+5} = 170$, while Case 2 produces $\frac{1\cdot 200+5\cdot 20}{1+5} = 50$. In effect, the personal average disappears from the formula and is replaced by a sum of all the student class sizes observed, then divided by the total number of classes. The weighted average seems to do a better job of reflecting the reality of student experience.

We are now ready to introduce a bit of notation to allow a rigorous comparison. Assume that the number of classes (or sections) of size i is given by c_i, and the maximum class size is m. Notice that each size i appears c_i times, and the total number of classes is the sum of the c_i's. The university average class size $\bar{c}$ is therefore

$$\bar{c} = \frac{\sum_{i=1}^{m} i \cdot c_i}{\sum_{i=1}^{m} c_i}.$$

To compute the student-experienced average, observe that each class of size i gets reported by i students, for a total of ic_i reports. This makes the weighted average of all the student computed average class sizes

$$\bar{s} = \frac{\sum_{i=1}^{m} i \cdot ic_i}{\sum_{i=1}^{m} ic_i}.$$

In the student average $\bar{s}$ the larger classes have higher weights. The university average $\bar{c}$ weighs each class equally. Consequently, for any distribution of class sizes whatsoever, the student average will always exceed the university average. The only way for equality to occur is to have all classes with equal sizes.

If that explanation is not to your liking, we can use statistics. Since the variance of $\bar{c}$ is the average of the squares of the deviations from the mean, it is certainly nonnegative [**2**, p. 98], so in our problem we find

$$\mathrm{Var}(\bar{c}) = \frac{\sum_{i=1}^{m} i^2 c_i}{\sum_{i=1}^{m} c_i} - (\bar{c})^2 = \frac{\sum_{i=1}^{m} i^2 c_i}{\sum_{i=1}^{m} c_i} - \left[\frac{\sum_{i=1}^{m} i \cdot c_i}{\sum_{i=1}^{m} c_i}\right]^2 \geq 0$$

Upon dividing by $\bar{c}$ we find

$$\left[\frac{\sum_{i=1}^{m} i^2 c_i}{\sum_{i=1}^{m} c_i}\right]\left[\frac{\sum_{i=1}^{m} c_i}{\sum_{i=1}^{m} i \cdot c_i}\right] - \left[\frac{\sum_{i=1}^{m} i \cdot c_i}{\sum_{i=1}^{m} c_i}\right] \geq 0$$

Canceling the $\sum c_i$ in the first term yields $\bar{s} - \bar{c} \geq 0$. So certainly the student-experienced weighted average is always greater than the university average, with equality only if the variance is zero; that is, only if all classes are precisely the same size.

Finally, if statistics is not your cup of tea, perhaps you are a fan of Cauchy and Schwarz. Define two vectors of dimension m, the maximum class size:

$$\vec{v} = [\sqrt{c_1}, \sqrt{c_2}, \sqrt{c_3}, \ldots, \sqrt{c_m}]$$
$$\vec{w} = [1\sqrt{c_1}, 2\sqrt{c_2}, 3\sqrt{c_3}, \ldots, m\sqrt{c_m}]$$

According to the Cauchy-Schwarz inequality [**1**, p. 195],

$$(\vec{v} \cdot \vec{v})(\vec{w} \cdot \vec{w}) \geq (\vec{v} \cdot \vec{w})^2$$
$$\left[\sum_{i=1}^{m} c_i\right]\left[\sum_{i=1}^{m} i^2 c_i\right] \geq \left[\sum_{i=1}^{m} i \cdot c_i\right]^2.$$

Upon dividing by $\left[\sum_{i=1}^{m} c_i\right]\left[\sum_{i=1}^{m} i \cdot c_i\right]$, we get, yet again, $\bar{s} \geq \bar{c}$. And when might these be equal? Why only if one vectors is a scalar multiple of the other. This happens only if precisely one term is nonzero. That is, if all classes are the same size.

In the real world this never happens. Consequently, the student average $\bar{s}$ *always exceeds* $\bar{c}$, the university reported average. Sometimes this can be by a small amount; sometimes, as in my hypothetical example, the effect is large. At a typical university I would guess that 90% of the students experience an average class size larger than what the university reports.

References

1. O. Bretscher, *Linear Algebra with Appications* (3rd ed.), Pearson Prentice Hall, 2005.
2. H. J. Larson, *Introduction to Probability Theory* (2nd ed.), Wiley, 1974.

Exhaustive sampling and related binomial identities

Jim Ridenhour (ridenhourj@apsu.edu) and David Grimmett (grimmettd@apsu.edu), Austin Peay State University, Clarksville, TN 37044

There are many situations that involve repeated sampling from the same set of observations. For example, suppose a professor has a test bank of 100 questions for a particular course and randomly chooses 25 of these questions for the final exam each semester. A persistent but not very talented student repeats the course several times. Obviously, the student has no chance of having seen all the questions before taking the course four times. What is the probability that the student will have seen all the questions after k repetitions? That is, what is the probability that the entire test bank will have been exhausted after k repetitions?

A more practical example involves drug testing. Suppose, for example that a bicycle race has 100 contestants and consists of several stages where random samples of 20 contestants are taken at each stage and screened for banned substances. If the race has 10 stages, what is the probability that each contestant will be tested for banned substances at least once?

Probability of exhaustion. We will assume that we are selecting k samples of size n from a population containing N members. We want to find the probability of the event E that the population has been exhausted in the k samples. That is, every member of the population has been included in at least one sample. Denote the members of the population by $x_1, x_2, \ldots, x_N$. We will calculate the probability of the complementary event E^C. Let E_i be the event that x_i has not been included in any of the k samples. Then $E^C = E_1 \cup E_2 \cup \cdots \cup E_N$. By the addition law of probability and the method of inclusion and exclusion,

$$P(E^C) = \sum_{i=1}^{N} P(E_i) - \sum_{i<j} P(E_i \cap E_j) + \sum_{p<q<r} P(E_p \cap E_q \cap E_r) - \cdots, \quad (1)$$

where the last sum consists of all terms with $N - n$ intersections. That is because each sample has n distinct elements and so the greatest number of elements that cannot be included is $N - n$. For a particular x_i, let B_{ij} be the event that x_i is not included in the jth sample. Then $E_i = B_{i1} \cap B_{i2} \cap \cdots \cap B_{ik}$. Moreover,

$$P(B_{ij}) = \frac{\binom{N-1}{n}}{\binom{N}{n}}$$

since we must choose a sample from N elements without including x_i. Since the samples are chosen with replacement, the events $B_{i1}, B_{i2}, \ldots, B_{ik}$ are independent,

