1 Review questions

(a) *There are 10 students in a class. The professor randomly assigns final exam grades, A through C, by giving 3 As, 5 Bs, and 2 Cs. What is the mean GPA of a student in the class? The variance?*

Solution: the mean GPA is

$$\mu_{\text{GPA}} = \frac{1}{10}\big(3(4) + 5(3) + 2(2)\big) = 3.1$$

Using the mean, we can compute the variance

$$\sigma^2_{\text{GPA}} = \frac{1}{10}\big(3(4-3.1)^2 + 5(3-3.1)^2 + 2(2-3.1)^2\big) = 0.49$$

This tells us that the standard deviation is

$$\sigma_{\text{GPA}} = 0.7$$

(We computed these values assuming that the 10 students were our population, as in the question statement; how would they change is the 10 students were a sample from a larger population?)

(b) *How many different ways can the professor assign the 4 As?*

Solution: order does not matter (think about why if it's not clear to you; an A is an A regardless of whether you get the first A or the fourth A), so we use

$${}_{10}C_3 = \frac{10!}{7!3!} = 120$$

(c) *How many different grade assignments are there?*

Solution: there are two ways of attacking this problem. The first method uses the multiplication principle, while the second applies a more intuitive approach.

Consider this: from the 10 students, 3 are chosen to receive As and the rest receive something else. There are ${}_{10}C_3$ ways of doing this. From the remaining 7 students, 5 are chosen to receive Bs and the remaining 2 get something else (in this case, a C). There are ${}_7C_5$ ways of doing this. By the multiplication principle, then, the total number of grade assignments is

$${}_{10}C_3 \; {}_7C_5 = \left(\frac{10!}{7!3!}\right)\left(\frac{7!}{5!2!}\right) = \frac{10!}{3!5!2!} = 2520$$

As an alternate solution, consider this question the way we thought about deriving combinations from permutations: we look at a permutation, and then realize we don't care about the order of the elements so we divide out by the number of permutations of the elements we selected. In this case, there are 10! total permutations of students; without loss of generality, we assume that from a given permutation we give the first 3 students As, students 4 through 8 Bs, and the remaining 2 students Cs.

However, we have the usual problem that there is more than one permutation which will yield the same grade assignment if we proceed in this manner. So we need to normalize by the number of orderings within each assigned grade. There are 3! ways of permuting the A students, 5! ways of permuting the B students, and 2! ways of permuting the C students. Then the total number of grade assignments is

$$\frac{10!}{3!5!2!} = 2520$$

(d) *To get their exam grades, students wait outside the professor's office and enter one-at-a-time to view their exam (to protect privacy). The first student in line enters the professor's office, and exits telling the remaining 9 students in line that he received a B on the exam. Should the remaining students be more relaxed about their own grades?*

Solution: this question is a little obtuse, but hopefully presents useful problem-translation practice. Here, a student will be more relaxed about her own grades if her expected GPA increases as a function of the news she received from the first student. This is not particularly like anything we've seen in class so far, but uses tools we're well accustomed to.

To begin, we compute the conditional PMF of the grade distribution. As an example, letting G_i be the grade of student i, we compute $\Pr(G_i = A|G_1 = B)$ explicitly and then state the answers for the other two grades. We know, by definition of conditional probability,

$$\Pr(G_i = A|G_1 = B) = \frac{\Pr(G_i = A, G_1 = B)}{\Pr(G_1 = B)}$$

The naïve probability $\Pr(G_1 = B)$ is the probability that any student's grade is a B, which is $\frac{5}{10}$. $\Pr(G_i = A, G_1 = B)$ is a little tougher to compute; if i receives an A while 1 receives a B, the remaining 8 students are left to split 2 As, 4 Bs, and 2 Cs. The total number of assignments for these 8 students (of the remaining grades) is

$$\frac{8!}{2!4!2!} = 420$$

Then we see $\Pr(G_i = A, G_1 = B) = \frac{420}{2520}$; this tells us

$$\Pr(G_i = A|G_1 = B) = \frac{1}{3}$$

Applying similar logic (although the algebra certainly changes), we have

$$\Pr(G_i = B|G_1 = B) = \frac{4}{9} \qquad \Pr(G_i = C|G_1 = B) = \frac{2}{9}$$

So we now have a new PMF, conditional on the information we received from the first student:

$$f_{\text{GPA}}(x) = \begin{cases} \frac{3}{9} & \text{if } x = 4 \\ \frac{4}{9} & \text{if } x = 3 \\ \frac{2}{9} & \text{if } x = 2 \end{cases}$$

From this we can compute the conditional expectation,

$$E(G_i|G_1 = B) = \frac{1}{9}\left(\frac{3}{9}(4) + \frac{4}{9}(3) + \frac{2}{9}(2)\right) = 3.11$$

Since i's expected GPA was 3.1 before the announcement of student 1's grade and is 3.11 after, she is right to feel a little bit better about her prospects (if just a little better).

(As an aside, notice that we could have gotten the same PMF using the somewhat hand-wavy argument that since one B was "eaten up" by the first student, there are (for example) 3 ways out of 9 for student i now to receive an A. This logic will obtain the correct PMF, and is a useful sanity check.)

(e) *Forget for a moment that we are talking about only 10 students, and suppose that instead we know that 30% of some number of students receive an A in the class. The professor has allocated $\frac{2}{3}$ of the As in the class to students who attend office hours; 40% of students attend office hours regularly. What is the probability of receiving an A, conditional on attending office hours?*

Solution: this is a straightforward application of Bayes' rule. We see

$$\Pr(G_i = A|\text{OH}) = \frac{\Pr(\text{OH}|G_i = A)\Pr(G_i = A)}{\Pr(\text{OH})} = \frac{\frac{2}{3}(0.3)}{0.4} = \frac{1}{2}$$

(To make this a little harder, what is the percent-attendance rate of office hours such that a student who attends office hours is no more likely to receive an A than a student who does not? We haven't provided enough information to solve an expected GPA question in this setup, but we do have enough to solve for the A threshold.)

(f) *Is it a good choice for a student to attend office hours, if they are interested in receiving an A in the class? Assume that the class is large enough that if one more student attends office hours, the attendance rate is still roughly* 40%.

Solution: if the students only care about receiving an A in the class, they will want to attend office hours if $\Pr(G_i = A|\text{OH}) > \Pr(G_i = A|\text{not OH})$. By the law of total probability,

$$\Pr(G_i = A) = \Pr(G_i = A|\text{OH})\Pr(\text{OH}) + \Pr(G_i = A|\text{not OH})\Pr(\text{not OH})$$

Using $\Pr(G_i = A|\text{OH}) = \frac{1}{2}$ from the previous question, we have

$$\Pr(G_i = A|\text{not OH}) = \frac{\Pr(G_i = A) - \Pr(G_i = A|\text{OH})\Pr(\text{OH})}{\Pr(\text{not OH})} = \frac{0.3 - 0.5(0.4)}{0.6} = \frac{1}{6}$$

Since $\frac{1}{6} < \frac{1}{2}$, a student interested in receiving an A should attend office hours. We didn't explicitly use the, "class is large enough," assumption anywhere, but if we were trying to model this as an economic choice we would need to know that the student's decision to attend office hours wouldn't significantly alter the probability of receiving an A, conditional on attending office hours (the wording is confusing, but think about it for a moment).

(g) *Suppose now that the* 10 *students are in a single section, which is part of a much larger class. For the entire class, the population standard deviation* $\sigma_{GPA} = 0.5$ *is known with population mean known to be* $\mu_{GPA} = 2.65$. *What is the minimum probability that a random student's grade is between a C and a B+?*

Solution: this question requires use of Chebyshev's inequality (the same as Chebyshev's theorem). We change notation slightly and let $\hat{G}_i$ be the grade of i, while G_i will be the corresponding GPA points.

$$\begin{aligned}
\Pr(\hat{G}_i \in \{C, C+, B-, B, B+\}) &= \Pr(2.0 \le G_i \le 3.3) \\
&= \Pr(|G_i - 2.65| \le 0.65) \\
&= \Pr\left(|G_i - 2.65| \le \frac{13\sigma_{\text{GPA}}}{10}\right) \\
&\le 1 - \left(\frac{10}{13}\right)^2 \\
&= 0.4083
\end{aligned}$$

Then at least 40% of the class should receive between a C and a B+.

(h) *Suppose there are* 320 *students in the class. What is a lower bound for the number we expect to receive between a C and a B+?*

Solution: this is a natural extension of the previous question. If the class has 320 students and $\Pr(2.0 \leq G_i \leq 3.3) \geq 0.4083$, the minimum bound for the number of students we expect to receive between a C and a B+ is

$$320(0.4083) = 130.6509$$

As an aside, this quantity is not particularly meaningful. Still, we can calculate it all the same.

(i) *Suppose now only that the population variance $\sigma^2_{GPA} = 0.25$ is known, but the class mean is not. Using the section's grades (as presumptively assigned in part (a)), compute a 95% confidence interval for μ. You are free to assume that 10 is "sufficiently large" to apply the standard tricks, even though you should know better.*

Solution: a standard confidence interval is

$$\mu \in \left[\overline{G} - E, \overline{G} + E\right]$$

From part (a), we have the section's mean is $\mu_{\text{GPA}} = 3.1$. Now, the section is a sample of the larger class, so the sample mean is $\overline{G} = 3.1$. The margin of error E is determined in the standard way; with a 95% confidence interval, $\alpha = 0.05$, so we have

$$E = z_{\alpha/2}\left(\frac{\sigma}{\sqrt{N}}\right) = 1.96\left(\frac{0.25}{\sqrt{10}}\right) = 0.3099$$

Then a 95% confidence interval for μ is

$$\mu \in [3.1 - 0.3099, 3.1 + 0.3099] = [2.7901, 3.4099]$$

(This confidence interval corresponds to a known standard deviation; how does it change if the standard deviation is estimated from the sample?)

(j) *Suppose now that the TA cannot see his section's average test score (MyUCLA will not allow it), but after talking to 10 students finds their mean exam GPA to be $\bar{G}_k = 3.1$. The professor announces that the class mean is $\mu_{GPA} = 2.7$. Can the TA say with 1% significance that his section outperformed the class? Again, you are free to assume that 10 is "significantly large."*

Solution: this is a one-sided hypothesis test. Let μ_k be the section mean; our hypotheses are then

$$H_0: \quad \mu_k = \mu_{\text{GPA}} = 2.7$$
$$H_1: \quad \mu_k > \mu_{\text{GPA}}$$

The critical value of this test is $z_{0.01} = 2.3263$, corresponding to 1% significance (the entire mass of which lies in the left-hand tail, *why?*). We then compute the z-statistic of our sample mean according to the null hypothesis,

$$\sqrt{N}\left(\frac{3.1 - 2.7}{\sigma}\right) = \sqrt{10}\left(\frac{0.4}{0.5}\right) = 2.5298$$

Since $2.5298 < 2.3263$, we can safely reject the null hypothesis; with 1% significance, the data supports the section outperforming the remainder of the class.

(How would this change if we needed to estimate the class variance with sample variance $S = 0.49$?)