Answer Key:

1. C
2. B
3. A
4. D
5. B

6. D
7. D
8. A
9. C
10. D

11. B
12. B
13. C
14. C
15. D

16. A
17. C
18. A
19. B
20. D

21. D
22. D
23. B
24. D
25. B

26. C
27. A
28. A
29. C
30. E

Solutions:

The first thing all students should do is carefully enter the three lists of numbers into a set of three lists in their calculator. For the sake of simplicity, let's enter the *Pretest Scores* into L_1, the *Posttest Scores* into L_2, and the set of differences in score (*Posttest – Pretest*) into L_3 which will aid in answering the first 20 questions.

1. C: Summing the set of 16 differences in score and then dividing the sum by 16 or equivalently by running the *1-Variable Statistics* program on L_3 gives us a sample mean for the set of differences in score of $\bar{X}_D = 15$.

2. B: The same output screen for answering question #1 also gives us a sum of squares of $\sum x_D = 3680$ and so the sample variance is $s_D = \frac{16}{15}\left(\frac{3680}{16} - 15^2\right) = \frac{16}{3}$ as an exact value.

3. A: Running any one of the *Linear Regression* programs with *X-list:* L_1 and *Y-list:* L_2 gives us a sample correlation coefficient of $r = 0.993$ when rounded to the thousandths place.

4. D: When a bivariate set of data is plotted in either a scatterplot or a residual plot, either is useful for establishing whether or not a linear relationship exists between the explanatory variable and the response variable of the bivariate set of data. Here they are for the data set in our scenario:

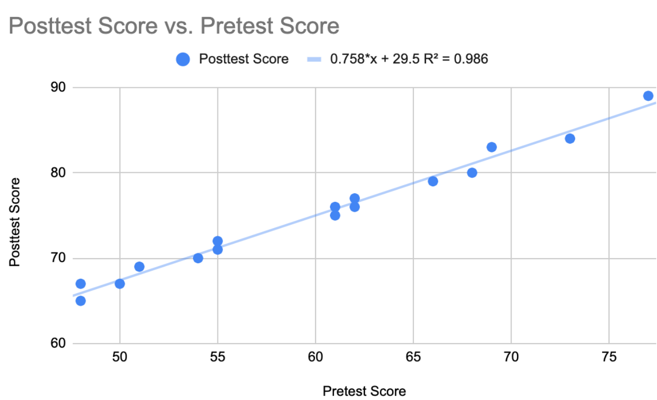

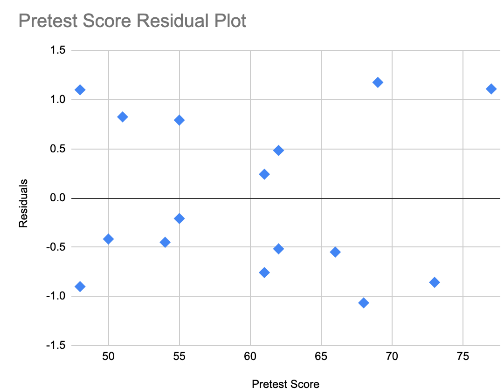

5. B: The running of the *Linear Regression* program with *X-list:* L_1 and *Y-list:* L_2 also gives us the following least-squares regression line equation in contextual form with the coefficients rounded to four decimal places: $\widehat{Posttest\ Score} = 29.5017 + 0.7583(Pretest\ Score)$.

6. D: Hopefully, students know how to utilize the "*RESID*" feature in their calculator to quickly and easily calculate the set of 16 residuals and save them in an empty list, say L_4. Then, one could either draw the residual plot above with their calculator's graphing features of simple skim through the list of residuals and note that there are 7 positive residuals and 9 negative residuals. Otherwise, one could simply calculate each of the 16 residuals by using the definition of the residual: $e = y - \hat{y} = y - (a + bx)$. This would take a little more time, but it is doable. The easiest way to do this is to do the following: $L_4 = 29.5017 + 0.7583(L_1)$. Here is a list of the 16 residuals rounded to the hundredths place for convenience. Clearly, 7 of them are positive and 9 of them are negative.

Observation	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
Residual	1.11	-0.86	1.18	-1.07	-0.55	0.48	-0.52	0.24	-0.76	0.79	-0.21	-0.45	0.82	-0.42	1.10	-0.90

7. D: Of all the *Linear Regression* programs available in their calculator, hopefully students choose to use the *Linear Regression T-Test* since it will help answer several of the questions in this section of the test. In particular, for this question, the estimated variance about the LSRL is the square of the value $s = 0.830960869$ which appears on the output screen because it represents the standard error of the residuals; and hence, the estimated standard deviation of the response variable about the LSRL. Thus, the answer we seek is $s^2 = 0.830960869^2 \approx 0.6905$ when rounded to four decimal places.

8. A: The calculator's *Linear Regression T-Test* output screen will also give us the answer to this question because the test statistic is $t = 31.66481546 \approx 31.6648$ when rounded, the *p-value* is $p = 9.894516 \times 10^{-15} \approx 0.0000$ when rounded, and the degrees of freedom are $df = 14$. Their sum is $31.6648 + 0.0000 + 14 = 45.6648$.

9. C: This time, we simply need to run the *Linear Regression T-Interval* program with *X-list:* L_1 and *Y-list:* L_2 which gives us the following 98% confidence interval for β with the limits appropriately rounded: $(0.6955, 0.8212)$

10. D: Since the hypothesis test of $H_0: \beta = 0\ vs.\ H_A: \beta > 0$ in question #8 rejects the null hypothesis and supports the alternative hypothesis at the 1% level of significance (which was stated an the required level of significance on the first page) and the 98% confidence interval constructed in question #9 does not capture any of the values specified for β in choices A, B, or C but it DOES capture the value of $\beta = 0.75$ in choice D; then the value of $\beta = 0.75$ is the only one of the claimed values of β that is NOT REJECTED as a plausible value for β.

11. B: Using the given information and formulas, we have $\mu_b = \beta = 0.76\ and\ \sigma_b = \frac{\sigma_e}{\sigma_x\sqrt{n}} = \frac{0.81}{9\sqrt{16}} = 0.0225$. Now, it is just a normal distribution calculation: $P(0.75 < b < 0.80) = normalcdf(0.75, 0.80, 0.76, 0.0225) \approx 0.6339$.

12. B: Here is a screenshot of the exact definition of an **unbiased estimator** from page 476 of *The Practice of Statistics Updated Sixth Edition* by Starnes and Tabor which is the premier AP Statistics textbook:

> **DEFINITION Unbiased estimator**
>
> A statistic used to estimate a parameter is an **unbiased estimator** if the mean of its sampling distribution is equal to the value of the parameter being estimated.

13. C: Running the *T-Test* program on the set of differences in score we entered earlier into L_3 and using the *Data* option with $\mu_0 = 12$ and the right-tail alternative selected gives us a rounded test statistic of $t = 5.196$ and a rounded *p-value* of $p = 0.000$. The degrees of freedom are $df = n - 1 = 16 - 1 = 15$ and since the p-value is less than the 0.01 level of significance we can reject the null hypothesis and support the alternative hypothesis. This makes the final answer $5.196 + 0.000 + 15 + 0.01 = 20.206$.

14. C: Confidence intervals, by definition, precisely give us a range of plausible values for the parameter of interest and so essentially any possible claimed value for the parameter of interest that is <u>not</u> captured by the interval is <u>rejected</u> as a plausible value for the parameter and any value that <u>is</u> captured by the interval is <u>not rejected</u> as a plausible value for the parameter.

15. D: Running the *T-Interval* program with the data in L_3 gives us the following 98% confidence interval for μ_G with the limits appropriately rounded: $(13.5, 16.5)$.

16. A: Since we have an SRS of 16 independent MAΘ Calculus student competitors, then the best unbiased point estimate for the population proportion of those who would potentially gain over 12 points on these parallel form tests is the number observed who did do just that out of our sample of 16: $\hat{p} = \frac{13}{16} = 0.8125$.

17. C: Clearly, the sample size of only 16 falls far short of the requisite criterion for using the normal approximation to the binomial distribution for constructing a confidence interval for an unknown population proportion using the usual formula of $\hat{p} \pm Z^*\sqrt{\frac{\hat{p}(1-\hat{p})}{n}}$. Depending on the textbook one uses, this requisite criterion is that we have observed at least 10 (or 5) successes as well as at least 10 (or 5) failures in our sample. Well, our sample has only 3 observed successes and 13 observed failures; thus, failing the criterion often referred to as the "large counts condition" or the "sufficient sample size condition" by most AP Statistics textbooks.

18. A: Since the matched-pairs t-test of $H_0: \mu_G = 12$ vs. $H_A: \mu_G > 12$ in question 13 resulted in rejecting the null hypothesis and supporting the alternative hypothesis and the 98% confidence interval for μ_G in question 15 landed completely above $\mu_G = 12$, there is statistically significant and convincing evidence that the mean gain in score in this scenario is greater than 12; and so, Señor Nieve should indeed market and sell Dr. Rampal's secret sleep audio program to MAΘ Calculus competition coaches so that they can apply it to their student competitors in hopes of becoming more competitive at Calculus. Thus, YES and YES! Oh, and by the way, they really are great guys as well!

19. B: Under the given assumptions and conditions, the sample mean gain in score required to reject the null hypothesis of $H_0: \mu_G = 12$ and support the alternative hypothesis of $H_A: \mu_G > 12$ at the 0.01 level of significance is $\bar{X}_D = invNorm\left(0.99, 12, \frac{2.4}{\sqrt{16}}\right) \approx 13.3958$. Thus, the power of the test is $normalcdf\left(13.3958, 999999, 14, \frac{2.4}{\sqrt{16}}\right) \approx$ 0.843 when rounded.

20. D: Increasing the level of significance of a hypothesis test increases the probability of committing a Type-I error while simultaneously increasing the power of the test; and as a consequence, decreases the probability of committing a Type-II error.

21. D: As noted in the stem of the question, all four of the listed choices are indeed potential issues with the pooling procedure described; however, the issues described in choices A through C should have relatively minimal impact on the results since it is likely that only a few individuals should end up in any of those three categories described. By contrast, the screening procedure specifically <u>fails</u> to exclude non-registered voters by asking a question as to the voting eligibility of the potential interviewee and so it will most certainly include persons over 18 who are specifically <u>not</u> registered to vote. This clearly fails to meet the specifications of the intended population of interest (namely, <u>registered</u> voters) which most certainly also has the potential to significantly bias the results of the poll.

22. D: The formula for calculating the number of derangements of the *n* items in a list is $!n = n! \sum_{k=0}^{n} \frac{(-1)^k}{k!}$. Note that the summation converges to e^{-1} as $n \to \infty$. As it turns out, the convergence is rather fast and so the number of derangements of an *n-item* list simplifies to $!n = \frac{n!}{e}$ rounded to the nearest integer. This makes the probability of a derangement of an *n-item* list converge to e^{-1} as $n \to \infty$ and for our 13 playing cards it is $\frac{2{,}290{,}792{,}932}{6{,}227{,}020{,}800} \approx 0.36787894$ which is indeed very close to $e^{-1} \approx 0.3678794412 \ldots$. This makes the probability we seek $\frac{0.3678794412}{4} \approx 0.092$ when rounded because Bernie first randomly selected one of the four suites to derange.

23. B: Since the two sample sizes for the two-sample t-test described in the problem are equal, the formula for the test statistic becomes $t = \frac{\bar{X}_1 - \bar{X}_2 - 0}{\sqrt{\frac{s_1^2}{n} + \frac{s_2^2}{n}}}$. Solving this for n and substituting in the given values we obtain the following:

$n = (s_1^2 + s_2^2)\left(\frac{t}{\bar{X}_1 - \bar{X}_2}\right)^2 = (2.3^2 + 1.9^2)\left(\frac{4.24}{25-23}\right)^2 \approx 40.00016$ which rounds to a sample size of *n = 40* and 4 + 0 = 4.

24. D: Random variable X_1 is a discrete uniform distribution on the set of positive integers {1, 2, 3, 4, 5, 6, 7 ,8 ,9} for which the subset of primes consists of {2, 3, 5, 7} which makes $P(X_1 \text{ } is \text{ } a \text{ } prime) = \frac{4}{9}$. NOTE: 1 is neither prime nor composite, by definition.

25. B: Random variable X_2 is a continuous uniform distribution on the closed interval $[0, 9]$ which makes the requested probability of $P(3 \leq X_2 < 6) = \frac{6-3}{9-0} = \frac{1}{3}$. NOTE: Since X_2 is continuous, the inclusion or the exclusion of the endpoints of the interval do not affect the calculation.

26. C: Random variable X_3 is a binomial distribution with $n = 12$ trials and probability of success $p_3 = 0.75$ which makes $\mu_3 = np_3 = 12(0.75) = 9$ and $\sigma_3 = \sqrt{np_3(1-p_3)} = \sqrt{12(0.75)(0.25)} = 1.5$. This makes the requested probability of $P(\mu_3 - 2\sigma_3 \le X_3 < \mu_3 + 2\sigma_3) = P(9 - 2 \times 1.5 \le X_3 < 9 + 2 \times 1.5) = P(6 \le X_3 < 12)$ which is equivalent to $P(6 \le X_3 \le 11) = binomcdf(12, 0.75, 11) - binomcdf(12, 0.75, 5) \approx 0.954$ when rounded.

27. A: Since X_1 is a discrete uniform distribution defined on the sample space {1, 2, 3,...9}, its mean is $\mu_1 = \frac{9+1}{2} = 5$ and since X_4 is a geometric distribution with probability of success $p_4 = 0.25$, its mean is $\mu_4 = \frac{1}{p_4} = \frac{1}{0.25} = 4$ and its variance is $\sigma_4^2 = \frac{1-p_4}{p_4^2} = \frac{0.75}{0.25^2} = 12$. Thus, we are looking for $P\left[(X_1 = \mu_1) \cup \left(X_4 = \frac{\sigma_4^2}{\mu_4}\right)\right] = P[(X_1 = 5) \cup (X_4 = 3)] =$ $P(X_1 = 5) + P(X_4 = 3) - P(X_1 = 5) \times P(X_4 = 3) = \frac{1}{9} + \left(\frac{3}{4}\right)^2\left(\frac{1}{4}\right) - \left(\frac{1}{9}\right)\left(\frac{3}{4}\right)^2\left(\frac{1}{4}\right) = \frac{17}{72}$ by the rule for the probability of the union of two independent events.

28. A: Random variable X_5 is a standard normal distribution so $\mu_5 = 0$ and $\sigma_5 = 1$. Random variable X_6 is a standard Student's t-distribution so $\mu_6 = 0$ and its variance is $\sigma_6^2 = \frac{df}{df-2} = \frac{12}{10} = 1.2$ (which was provided), so we have $\sigma_6 = \sqrt{1.2}$. Random variable X_7 is a χ_8^2 distribution so $\mu_7 = df = 8$ and $\sigma_7^2 = 2(df) = 16$ (which were provided); and so, we have $\sigma_7 = 4$. This makes the requested probability become the following:
$P[(\mu_5 - 2\sigma_5 \le X_5 < \mu_5 + 2\sigma_5) \cap (\mu_6 - 2\sigma_6 \le X_6 < \mu_6 + 2\sigma_6) \cap (\mu_7 - 2\sigma_7 \le X_7 < \mu_7 + 2\sigma_7)] =$
$P\left[(0 - 2(1) \le X_5 < 0 + 2(1)) \cap (0 - 2\sqrt{1.2} \le X_6 < 0 + 2\sqrt{1.2}) \cap (8 - 2(4) \le X_7 < 8 + 2(4))\right] =$
$P\left[(-2 \le X_5 < 2) \cap (-2\sqrt{1.2} \le X_6 < 2\sqrt{1.2}) \cap (0 \le X_7 < 16)\right] =$
$P(-2 \le X_5 < 2) \times P\left(-2\sqrt{1.2} \le X_6 < 2\sqrt{1.2}\right) \times P(0 \le X_7 < 16) =$
$normalcdf(-2, 2, 0, 1) \times tcdf\left(-2\sqrt{1.2}, 2\sqrt{1.2}, 12\right) \times \chi^2cdf(0, 16, 8) =$
$(0.954499876)(0.9510702741)(0.957619888) \approx 0.8693$ when rounded to the ten-thousandths place.
NOTE: We can multiply the probabilities above because the random variables are independent and the inclusion / exclusion of the endpoints of the intervals is irrelevant because all three random variables are continuous.

29. C: Let us first calculate the one remaining variance for random variable X_1 that we have not already calculated nor is it already provided. Entering 1, 2, 3, … 9 into a list in the calculator and running the *1-Variable Statistics* program gives us $\sigma_1^2 = \frac{285}{9} - 5^2 = \frac{20}{3}$. Now, since the 10 given random variable are all independent of each other, the mean of their sum, Y, is the sum of their means and the variance of their sum is the sum of their variances. Here is a table summarizing all the means and variances of the 10 independent random variables and their sum, Y:

	X_1	X_2	X_3	X_4	X_5	X_6	X_7	X_8	X_9	X_{10}	$Y = \sum_{k=1}^{10} X_k$
Mean	5	4.5	9	4	0	0	8	2	0.5	10	43
Variance	$\frac{20}{3}$	6.75	2.25	12	1	1.2	16	4	0.05	100	$\frac{1799}{12}$

Now, since we are taking an SRS of $n = 100$ from random variable Y, the sampling distribution of the sample mean for random variable $\bar{Y}$ is normally distributed according the Central Limit Theorem with a mean of $\mu_{\bar{Y}} = \mu_Y = 43$ and a standard deviation of $\sigma_{\bar{Y}} = \frac{\sigma_Y}{\sqrt{n}} = \sqrt{\frac{1799}{1200}}$. This makes the requested probability of $P(41.1 \le \bar{Y} \le 44.4) =$ $normalcdf\left(41.1, 44.4, 43, \sqrt{1799/1200}\right) \approx 0.8132$ when rounded to the ten-thousandths place.

30. E: Since at least one of the 10 random variables is continuous, we have that $P(X_k = \mu_k) = 0$ for at least one value of k. This makes the requested probability of $P(X_k = \mu_k \; for \; all \; k) = 0$. Alternatively, one could argue that there is insufficient information to answer this question which would lead to choice E anyway.