Part IV- Chapter 16	**Random Variables**
Random variable Discrete random variable Continuous random variable	A variable, denoted by a capital letter (X, Y, Z etc.), whose value is a numerical outcome of a random event. The theoretical data (possible outcomes) of a probability model. Has a finite number of possible outcomes. Takes all values in an interval of numbers (infinite or bounded).
Probability model	A function that associates a probability P with each value of a discrete random variable X, denoted $P(X = x)$, or with any interval of values of a continuous random variable.
Probability histogram	Pictures the probability distribution of a discrete random variable. (a relative frequency histogram for a very large number of trials)
Density curve	Pictures the probability distribution of a continuous random variable (normal distributions are 1 type)
Expected value of a random variable.	The mean over the long run of a random variable. If the random variable is discrete, multiply each possible value by the probability that it occurs, and find the sum: $\mu_x = E(X) = \Sigma x_i p_i$
Variance of a random variable.	The expected value of the squared deviation from the mean $\sigma^2_x = \mathrm{Var}(X) = \Sigma(x_i - \mu_x)^2\, p_i$
Standard deviation of a random variable	Describes the spread of the model $\sigma_x = \mathrm{SD}(X) = \sqrt{\mathrm{Var}(X)}$
$\mu_{a+bX} =$ ______ $\sigma_{a+bX} =$ ______	$a + b\mu_X$ (a and b are constants) $b\sigma_X$
$\mu_{X+Y} =$ ______ $\mu_{X-Y} =$ ______ $\sigma_{X\pm Y} =$ ______	$\mu_X + \mu_Y$ $\mu_X - \mu_Y$ $\sqrt{\sigma^2_X + \sigma^2_Y}$, if X and Y are independent. (Pythagorean Theorem of Statistics)
$X_1 + X_2 \neq$ _____ $\mu_{X1+X2} =$ ______ $\mu_{X1-X2} =$ ______ $\sigma_{X1\pm X2} =$ ______	2X , (X_1 & X_2 are distinct random variables with the same μ and σ. They aren't like terms) $\mu_{X1} + \mu_{X2} = 2\mu_X$ $\mu_{X1} - \mu_{X2} = 0$ $\sqrt{\sigma^2_{X1} + \sigma^2_{X2}} = \sqrt{2\sigma^2_X} = \sigma_X\sqrt{2}$
If two independent continuous random variables have Normal models,	So does their sum or difference.
Part IV- Chapter 17	**Probability Models**
Bernoulli trial	1. two possible outcomes ("success" & "failure") 2. probability of success is constant p $q = 1 - p$ 3. trials are independent (or sample < 10% of population)
If number, X, of Bernoulli trials until next success [measuring until success]	Then Geometric probability model, Geom(p): $P(X = x) = q^{x-1}p$ (Expected # of trials until success) $\mu = \frac{1}{p}$ $\sigma = \sqrt{\frac{q}{p^2}}$
If number of successes, X, in n Bernoulli trials [number of successes, no *when*]	Then Binomial probability model, Binom(n, p): $P(X = k) = \binom{n}{k} p^k q^{n-k}, \; where \binom{n}{k} = \frac{n!}{k!(n-k)!}$

	(Expected # of successes) $\mu = np$ $\sigma = \sqrt{npq}$
Assumptions	Theoretical mathematical requirements (independence, large sample, etc.)
Conditions	Practical guidelines that confirm (or sometimes override) assumptions.
When using the Geometric or Binomial probability models check that you have _________	the 3 requirements of Bernoulli trials.
The Binomial probability model becomes difficult/impossible for _________. Fortunately it can be approximated by _________ as long as we meet the ________ Condition that ______________	large *n*. a Normal probability model Success/Failure we expect at least 10 successes and 10 failures: $np \geq 10$ and $nq \geq 10$
On the AP Exam students are required to ______, not just _____ the conditions. This means __________________	check state using the values given in the question to show your work!
Part V- Chapter 18	**Sampling Distribution Models (SDMs)**
Proportion	Ratio of: $\dfrac{\text{number of successes}}{\text{total}}$ for categorical data. [think percent]
We want to know the true population proportion (mean),__, but are often forced to work/estimate with a sample proportion (mean), ___ .	p (μ) $\hat{p}$ ($\bar{x}$) 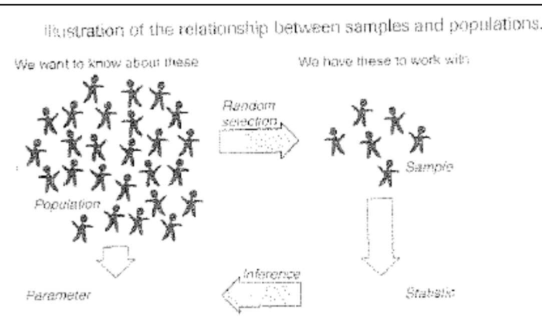
Sampling variability (sampling error)	No sample fully and exactly describes the population; sample proportions and means will vary from sample to sample. It is not just unavoidable – it's predictable! (with SDMs)
Sampling Distribution Model (SDM)	Shows how a statistic (sample proportion or mean) would vary in repeated (think infinite) samples of size *n*. We used to focus on the data, and derive the statistics from it. Now we focus on the statistic itself. The sample proportion (or mean) becomes our datum, and in our imaginations we compare that statistic to all other values we might have obtained from all the other samples of size n we might have taken.
The sample proportion, $\hat{p}$, does not have a binomial distribution because it is not the _________. But the SDM for a proportion appears to be _____	number of successes unimodal

and _____. When certain conditions are met, the _______ is a good SDM for a proportion.	roughly symmetric Normal model
Assumptions / Conditions for using a Normal model as the SDM for a proportion:	Assumptions: 1. Independent - sampled values must be independent of each other. Conditions: a) Randomization – SRS or at least representative and not biased. b) 10% Condition – If sampling w/o replacement Then n ≤ 10% of the population. 2. Sample Size - *n*, must be large enough. Conditions: a) Success/Failure - $np \geq 10$ and $nq \geq 10$.
Since the number of successes in the sample, X, is _________, we can obtain the mean and SD of the sample proportion by multiplying the mean and SD of the Binomial by the constant 1/n to get:	a Binomial random variable (n trials, probability p) $\mu(\hat{p}) = p$ $\sigma(\hat{p}) = SD(\hat{p}) = \sqrt{\frac{pq}{n}}$ $N\left(p, \sqrt{\frac{pq}{n}}\right)$
When we can understand and predict the variability of our estimates with SDMs, ______	we've taken the essential step toward seeing past that variability, so we can understand the world.
Means summarize _____________ data	quantitative
As long as the observations are ______, even if we sample from a skewed or bimodal population the ________________ tells us that the means (or proportions) of repeated random samples will tend to follow __________ as _______	independent Central Limit Theorem a Normal model the sample size grows.
Central Limit Theorem (CLT) [the fundamental theorem of statistics]	The sampling distribution model of the sample mean (and proportion) is approximately Normal for large *n*, regardless of the distribution of the population, as long as the observations are independent.
Assumptions / Conditions for using a Normal model as the SDM for a mean:	Assumptions: 1. Independent - sampled values must be independent of each other. Conditions: a) Randomization – SRS or at least representative and not biased. b) 10% Condition – If sampling w/o replacement Then n ≤ 10% of the population. 2. Sample Size - *n*, must be large enough. (More on this later) Conditions: a) For now, *Think* about your sample size in the context of what you know about the population, and then *Tell* whether the

	Large Enough Sample Condition has been met.
Unlike proportions, if we know the true population mean, μ, we don't automatically know the __	standard deviation of the population, σ.
For means the sampling distribution is centered at ________ and its standard deviation declines with the ________. So the Normal Model representing the SDM for a mean is _____	the true population mean $\mu(\bar{x}) = \mu$ square root of the sample size $\sigma(\bar{x}) = SD(\bar{x}) = \frac{\sigma}{\sqrt{n}}$ $N\left(\mu, \frac{\sigma}{\sqrt{n}}\right)$
Law of Diminishing Returns	Larger *n* yields smaller $\sigma(\bar{x})$ therefore $\bar{x}$ can tell us more about μ Unfortunately *n* only decreases $\sigma(\bar{x})$ at a rate of $\frac{1}{\sqrt{n}}$
Standard Error	If we don't know *p* or σ, then we must estimate the standard deviation of a sampling distribution with $\hat{p}$ or *s*. $SE(\hat{p}) = \sqrt{\frac{\hat{p}\hat{q}}{n}}$ $SE(\bar{x}) = \frac{s}{\sqrt{n}}$
Part V- Chapter 19	**Confidence Intervals for Proportions**
SDM for a proportion when we don't know *p*.	We don't know where to center our model and the best we can do for $\sigma(\hat{p})$ is the $SE(\hat{p})$ The resulting model is: $N\left(p, \sqrt{\frac{\hat{p}\hat{q}}{n}}\right)$ However, this still doesn't show us the value of *p*. The best we can do is to reach out with the $SE(\hat{p})$ on either side of $\hat{p}$ to create a confidence interval in an attempt to capture *p*.
Statistical inference	To use the sample we have at hand to say something about the world at large. In this case, we utilize the SDM of $\hat{p}$ to express our confidence in the results of any one sample.
Confidence interval [*p*-trap]	offers a range of plausible values for a model's parameter. For example: $\hat{p} \pm 2 \times SE(\hat{p})$
One-proportion *z*-interval [Official Name give to this type of confidence interval]	$\hat{p} \pm z^* \times SE(\hat{p})$
Margin of error *(ME)*	How far the confidence interval reaches out from $\hat{p}$ $\hat{p} \pm \overbrace{z^* \times SE(\hat{p})}$
z^*	Critical value – the number of standard errors to move away from the mean of the sampling distribution to correspond to the specified level of confidence.
To calculate z* for a particular level of confidence….	$z^* = \left\lvert invNorm\left(\frac{1 - confidence\ level}{2}\right)\right\rvert$
Assumptions / Conditions to	(See your inference guide)

check before creating (and believing) a confidence interval about a proportion:	
The more confident we want to be . . .	the larger the margin of error must be.
Every confidence interval is a balance between ____ and ____	certainty and precision.
The time to think about your margin of error, to see whether it's small enough to be useful, is	when you design your study or experiment and decide on *n*.
To get a narrower interval (decrease the *ME*) without giving up confidence,	You need to have less variability in your sample proportion, $\hat{p}$, by choosing a larger sample, *n*.
Law of Diminishing Returns	The larger the sample size, *n*, we have the narrower our confidence interval can be (at the rate of $\frac{1}{\sqrt{n}}$)
To calculate the sample size, *n*, necessary to reach conclusions that have a desired margin of error (degree of precision) and level of confidence:	Solve for *n* in: $ME = z^* \sqrt{\frac{\hat{p}\hat{q}}{n}}$ by substituting: *ME* = desired margin of error (as a decimal) z^* = critical value for desired level of confidence $\hat{p}$ = estimate based on experience or 0.5 (most cautious) $\hat{q} = 1 - \hat{p}$
Part V- Chapter 20	**Testing Hypotheses About Proportions**
Are the data consistent with the hypothesized SDM for a proportion?	We hypothesize a value, p_0, to construct a model for the unknown true population proportion, *p*. $N\left(p_0, \sqrt{\frac{p_0 q_0}{n}}\right)$ Then we test the sample proportion, $\hat{p}$, to see if it lends support to the hypothesis or casts doubt on the viability of the model. First find how many standard deviations $\hat{p}$ is from p_0 (you do remember the z-score from Unit I-F don't you?) $z = \frac{(\hat{p} - p_0)}{SD(\hat{p})}$ where $SD(\hat{p}) = \sqrt{\frac{p_0 q_0}{n}}$ Second use our standard normal model to change z-scores into percents like we did back in Unit I-F. These percents/probabilities are now called P-values and give the probability of observing the sample proportion, $\hat{p}$, (or one more extreme) given the original model is true.
Null hypothesis, H_0 [origi<u>null</u> skeptical hypothesis] [the normal chance outcome]	Proposes a parameter, p_0, and hypothesized value for an original population model that nothing interesting happened, or nothing has changed. H_0: $p = p_0$ *(hypothesized value)*

Alternative hypothesis, H_A [actual hypothesis] [that there is a real effect]	Represents the change or difference that we are interested in (what you want to show), usually a range of other possible values. The position we will have to take if the results are so unusual as to make the null hypothesis untenable. However, even when we reject the null hypothesis, we won't know the true value of the population parameter. (that is why we follow up with confidence intervals)
Two-sided alternative hypothesis	H_A: $p \neq p_0$ We are interested in deviations in *either* direction away from the hypothesized parameter value.
One-sided alternative hypothesis	H_A: $p > p_0$ or H_A: $p < p_0$ We are interested in deviations in *only one* direction away from the hypothesized parameter value.
Hypothesis are about __ not __	parameters not statistics (so no hats)
Hypothesis tests and confidence intervals share many of the same concepts.	Both rely on sampling distribution models, and because the models are the same and require the same assumptions, both check the same conditions.
Assumptions / Conditions for testing hypotheses about a proportion:	(See your inference guide)
One-proportion *z*-test	A test of the null hypothesis by referring the statistic $z = \frac{(\hat{p} - p_0)}{SD(\hat{p})}$ where $SD(\hat{p}) = \sqrt{\frac{p_0 q_0}{n}}$ to a standard normal model to find a P-value.
P-value [Probability-value] [% in tail(s) for a z-score]	The probability of observing a result at least as extreme as ours if the null hypothesis were true. A small value indicates either that the observation is improbable or that the probability calculation was based on incorrect assumptions. The assumed truth of the null hypothesis is the assumption under suspicion.
How low a P-value do we need?	Traditional: adopt a level of significance (alpha) of 10%,5%,1% etc Modern: think about what it says about the situation under consideration, and then make a decision.
A low P-value can never confirm that ________________, but it can convince us _________	the model is correct (beyond a reasonable doubt) that it is wrong.
Follow up a rejection of a hypothesis with _____________	a confidence interval that estimates the true value of the parameter
Am I surprised? How surprised am I? What would not surprise me?	Should I reject the null hypothesis? What's the P-value? Write a confidence interval for the parameter.
4-steps needed for inference problems: (based on the College Board's rubrics for the AP Exam)	*(See your inference guide)*
Part V- Chapter 21	**More about Tests**
Alpha level, α	The threshold P-value selected in advance that determines when we reject a null hypothesis, H_0. If we observe a statistic ($\hat{p}$) whose P-value based on the null

	hypothesis is less than α, we reject that null hypothesis.
Statistically significant	When the P-value falls below the alpha level, we say that the test is "statistically significant" at that alpha level. (But this doesn't necessarily have any practical importance.)
Significance level	The alpha level is also called the significance level, most often in a phrase such as a conclusion: "we reject the null hypothesis at the 5% significance level."
Don't just reject/fail to reject __ at an _______ level. Report the _________ as an indication of the strength of the evidence.	H_0 Alpha/significance P-value
When we perform a hypothesis test we can make mistakes in two ways: The more serious mistake is ___	Type I error – the null hypothesis is true, but we mistakenly reject it. Type II error – the null hypothesis is false, but we fail to reject it. depends on the situation.
Type I error, α	The error of rejecting a null hypothesis, H_0, when in fact it is true (also called a "false positive"). The probability of a Type I error is α, the chosen alpha level. (It happens when H_0 is true but we've had the bad luck of drawing an unusual sample.)
Type II error, β	The error of failing to reject a null hypothesis, H_0, when in fact it is false (also called a "false negative"). The probability of a Type II error is β. It is difficult to calculate because when H_0 is false, we don't know what value the parameter, p, really is.
Power	$1 - \beta$ The probability of correctly rejecting a false null hypothesis, H_0.
Reducing α to lower Type __ error will move _____________ and have the effect of increasing the probability of a Type __ error, __, and correspondingly reducing ___________	I the critical value, p^*, II β the power.
Effect size	$p - p_0$ How far the truth, p, lies from the null hypothesis, p_0.
The larger the effect size, the _______ the chance of making a Type __ error and the greater the _______ of the test.	smaller II power
Whenever a study fails to reject its null hypothesis, _________. H_0 may be false but our test is ..	the test's power comes into question. too weak to tell.
If we reduce Type I error, we automatically must _________ Type II error. But there is a way to reduce both:	increase we need to make both SDM curves narrower → by decreasing the spread (SD) → by increasing n (However the benefits are muted by the Law of Diminishing Returns)
The _____________ gives us the answer to a decision about a	hypothesis test

parameter; the ____________ tells us the plausible values of that parameter.	confidence interval
You can approximate a ______ by examining the confidence interval. Specifically, a confidence level of C% corresponds to ______________	hypothesis a two-sided hypothesis test with an α level of 100 – C% a one-sided hypothesis test with an α level of $\frac{1}{2}(100 - C\%)$
Part V- Chapter 22	**Comparing Two Proportions**
The sampling distribution of $\hat{p}_1 - \hat{p}_2$ is, under appropriate assumptions, modeled by …	A Normal model with: $\mu = p_1 - p_2$ $\quad SD(\hat{p}_1 - \hat{p}_2) = \sqrt{\frac{p_1 q_1}{n_1} + \frac{p_2 q_2}{n_2}}$
Assumptions / Conditions for using a Normal model as the SDM for a difference between two proportions: *(Also confidence intervals and testing hypotheses)*	(See your inference guide)
Two-proportion *z*-interval (confidence interval for $p_1 - p_2$)	$(\hat{p}_1 - \hat{p}_2) \pm z^* \times SE(\hat{p}_1 - \hat{p}_2)$ where $SE(\hat{p}_1 - \hat{p}_2) = \sqrt{\frac{\hat{p}_1 \hat{q}_1}{n_1} + \frac{\hat{p}_2 \hat{q}_2}{n_2}}$
Two-proportion *z*-test	H_0: $p_1 - p_2 = 0$. Because we hypothesize that the proportions are equal, we pool the groups to find an overall proportion: $\hat{p}_{pooled} = \frac{\#Success_1 + \#Success_2}{n_1 + n_2}$ and use that pooled value to estimate the standard error: $SE_{pooled}(\hat{p}_1 - \hat{p}_2) = \sqrt{\frac{\hat{p}_{pooled}\hat{q}_{pooled}}{n_1} + \frac{\hat{p}_{pooled}\hat{q}_{pooled}}{n_2}}$ Now refer the statistic $z = \frac{(\hat{p}_1 - \hat{p}_2) - 0}{SE_{pooled}(\hat{p}_1 - \hat{p}_2)}$ to a standard normal model to find a P-value.