MA206X — WPR2 STATISTICS REFERENCE CARD — AY 26-2

Statistical Inference (Lessons 17–25)

SAMPLING DISTRIBUTIONS & CLT (Lesson 17)

Concept	Formula / Property
Sample mean	$\bar{X} \sim N\left(\mu,\ \frac{\sigma^2}{n}\right)$ for large n (CLT)
Standard error	$SE(\bar{X}) = \frac{\sigma}{\sqrt{n}}$ (est. $\frac{s}{\sqrt{n}}$)
Sample prop.	$\hat{p} \approx N\left(p,\ \frac{p(1-p)}{n}\right)$ for large n
CLT conditions	Large n (rule of thumb: $n \geq 30$); for proportions: $np \geq 10$ and $n(1-p) \geq 10$

HYPOTHESIS TESTING FRAMEWORK (Lesson 20)

H_0	Null hypothesis (status quo, $=$)
H_a	Alternative ($\neq$, $>$, or $<$)
α	Significance level (prob of Type I error)
Test statistic	Standardized measure of distance from H_0
p-value	P(test stat as or more extreme $\mid H_0$)
Decision rule	Reject H_0 if p-value $< \alpha$
Type I error	Reject H_0 when H_0 is true; $P = \alpha$
Type II error	Fail to reject H_0 when H_a true; $P = \beta$
Power	$1 - \beta =$ prob. of correctly rejecting false H_0

CI–HT duality: If a two-sided $(1-\alpha)\%$ confidence interval a two-sided hypothesis test, at level α.

ONE-SAMPLE INFERENCE (Lessons 18–22)

Mean (σ unknown) — One-sample t	
Conditions	Random sample, normal pop (or $n \geq 30$ by CLT)
Test stat	$t = \frac{\bar{x} - \mu_0}{s/\sqrt{n}} \sim t_{n-1}$
CI	$\bar{x} \pm t_{\alpha/2,\,n-1} \cdot \frac{s}{\sqrt{n}}$
R	`t.test(x, mu=mu0, alternative=...)`
Mean (σ known) — Z-test	
Test stat	$z = \frac{\bar{x} - \mu_0}{\sigma/\sqrt{n}}$
CI	$\bar{x} \pm z_{\alpha/2} \cdot \frac{\sigma}{\sqrt{n}}$
Proportion — One-proportion z-test	
Conditions	$np_0 \geq 10$ and $n(1-p_0) \geq 10$
Test stat	$z = \frac{\hat{p} - p_0}{\sqrt{p_0(1-p_0)/n}}$
CI	$\hat{p} \pm z_{\alpha/2}\sqrt{\frac{\hat{p}(1-\hat{p})}{n}}$
R	`prop.test(x, n, p=p0)`
Sample Size	
For mean	$n = \left(\frac{z_{\alpha/2} \cdot \sigma}{E}\right)^2$ where $E =$ margin of error
For prop	$n = \left(\frac{z_{\alpha/2}}{E}\right)^2 \hat{p}(1-\hat{p})$; use $\hat{p} = 0.5$ if unknown

TWO-SAMPLE INFERENCE (Lessons 23–25)

Independent Means — Two-sample t	
Conditions	Two independent random samples; both normal (or large n)
Test stat	$t = \frac{(\bar{x}_1 - \bar{x}_2) - D_0}{\sqrt{s_1^2/n_1 + s_2^2/n_2}}$
CI	$(\bar{x}_1 - \bar{x}_2) \pm t_{\alpha/2,\,df}\sqrt{\frac{s_1^2}{n_1} + \frac{s_2^2}{n_2}}$
df (conserv.)	$df = \min(n_1 - 1,\ n_2 - 1)$
R	`t.test(x1, x2)`
Independent Means — Pooled t (equal variances)	
s_p^2	$\frac{(n_1-1)s_1^2 + (n_2-1)s_2^2}{n_1 + n_2 - 2}$
Test stat	$t = \frac{(\bar{x}_1 - \bar{x}_2) - D_0}{s_p\sqrt{1/n_1 + 1/n_2}}$ $df = n_1 + n_2 - 2$
R	`t.test(x1, x2, var.equal=TRUE)`
Paired Data — Paired t (Lesson 24)	
Setup	Compute $d_i = x_i - y_i$; then one-sample t on d_i's
Test stat	$t = \frac{\bar{d} - D_0}{s_d/\sqrt{n}} \sim t_{n-1}$
CI	$\bar{d} \pm t_{\alpha/2,\,n-1} \cdot \frac{s_d}{\sqrt{n}}$
When?	Same subjects measured twice; natural pairing
R	`t.test(x1, x2, paired=TRUE)`
Two Proportions (Lesson 25)	
Conditions	$n_i\hat{p}_i \geq 10$ and $n_i(1-\hat{p}_i) \geq 10$ for $i = 1, 2$
Test stat	$z = \frac{\hat{p}_1 - \hat{p}_2}{\sqrt{\hat{p}(1-\hat{p})(1/n_1 + 1/n_2)}}$ where $\hat{p} = \frac{x_1 + x_2}{n_1 + n_2}$ (pooled under H_0)
CI	$(\hat{p}_1 - \hat{p}_2) \pm z_{\alpha/2}\sqrt{\frac{\hat{p}_1(1-\hat{p}_1)}{n_1} + \frac{\hat{p}_2(1-\hat{p}_2)}{n_2}}$
R	`prop.test(c(x1,x2), c(n1,n2))`

R-LITE INTERPRETER COMMANDS

Function	What it returns
`qnorm(p)`	z such that $P(Z \leq z) = p$
`pnorm(z)`	$P(Z \leq z)$ for std normal
`qt(p, df)`	t such that $P(T \leq t) = p$
`pt(t, df)`	$P(T \leq t)$ for t_{df}

Examples:
$z_{0.025}$: `qnorm(0.975)` $= 1.960$.
$t_{0.025,9}$: `qt(0.975, 9)` $= 2.262$.
Upper tail: $P(T > 2.3)$: `1 - pt(2.3, df)`.

TEST SELECTION, INTERPRETATION & CRITICAL VALUES

WHICH TEST DO I USE?

Data	Parameter	Procedure
1 quant sample	μ	One-sample t-test
1 categorical	p	One-proportion z-test
2 indep quant	$\mu_1 - \mu_2$	Two-sample t-test
2 paired quant	μ_d	Paired t-test
2 indep categ	$p_1 - p_2$	Two-proportion z-test

Key question: Are the two samples *independent* or *paired*?
Paired = same subjects, before/after, matched pairs.

HYPOTHESIS TEST STEPS

1. Hypotheses: State H_0 and H_a in terms of parameter.
2. Conditions: Verify assumptions for chosen test.
3. Test statistic: Compute t or z from data.
4. p-value: Find tail probability using R-lite.
For t-tests:
H_a: $> \Rightarrow$ `1 - pt(t, df)`
H_a: $< \Rightarrow$ `pt(t, df)`
H_a: $\neq \Rightarrow$ `2*(1 - pt(abs(t), df))`
For z-tests (proportions):
H_a: $> \Rightarrow$ `1 - pnorm(z)`
H_a: $< \Rightarrow$ `pnorm(z)`
H_a: $\neq \Rightarrow$ `2*(1 - pnorm(abs(z)))`
5. Decision: Reject H_0 if $p < \alpha$; else fail to reject.
6. Conclusion: Interpret in context using plain language.

INTERPRETATION TEMPLATES

Confidence Interval:
"We are [confidence level]% confident that the true [parameter in context] is between [lower] and [upper]."

p-value / HT Conclusion:
"With a p-value of _____ which is less than α, we reject H_0, meaning that [context of H_a]."

"With a p-value of _____ which is greater than α, we FTR H_0, meaning that [insufficient evidence for context of H_a]."

Common mistakes to avoid:
× "95% probability the true mean is in the interval"
✓ "95% of such intervals contain the true mean"
× "Accept H_0" ✓ "Fail to reject H_0"
(absence of evidence ≠ evidence of absence)
× "The null is true" ✓ "Insufficient evidence against the null"

COMMON CRITICAL VALUES

CL	α	$z_{\alpha/2}$	R code	
90%	0.10	1.645	`qnorm(0.95)`	
95%	0.05	1.960	`qnorm(0.975)`	
99%	0.01	2.576	`qnorm(0.995)`	

For t critical values, use `qt(1-alpha/2, df)`.
Example: 95% CI with $df = 9$: `qt(0.975, 9)` = 2.262.

POOLED vs. UNPOOLED vs. PAIRED

Pooled t	Use when you **assume** $\sigma_1^2 = \sigma_2^2$. Uses s_p^2 and $df = n_1+n_2-2$. More powerful when assumption holds, but **not robust** if violated.
Unpooled t	Default for independent samples. Does **not** assume equal variances. Use $df = \min(n_1 - 1, n_2 - 1)$ for hand calculations.
Paired t	Use when observations are **naturally paired** (same subject, before/after, matched). Compute d_i; then one-sample t on differences. $df = n - 1$.

Rule of thumb: If $\max(s_1, s_2)/\min(s_1, s_2) < 2$, pooling is reasonable.

NOTE ON PROPORTIONS — HT vs. CI

Hypothesis test for $p_1 - p_2$: Use **pooled** $\hat{p} = \frac{x_1+x_2}{n_1+n_2}$ in the standard error (because H_0 assumes $p_1 = p_2$).

Confidence interval for $p_1 - p_2$: Use **unpooled** $\hat{p}_1$ and $\hat{p}_2$ separately in the standard error (no assumption of equal proportions).

Same idea for one proportion: HT uses p_0; CI uses $\hat{p}$.

ASSUMPTIONS & CONDITIONS CHECKLIST

All tests	Random sampling / random assignment
t-tests	Normal population or $n \geq 30$ (CLT); independent observations
Paired t	Differences d_i are approximately normal (check via histogram/QQ of differences)
Pooled t	Equal population variances ($\sigma_1^2 = \sigma_2^2$)
z-tests (prop)	$np_0 \geq 10$, $n(1 - p_0) \geq 10$ (one-prop) $n_i\hat{p}_i \geq 10$, $n_i(1 - \hat{p}_i) \geq 10$ (two-prop)

ERRORS, POWER & PRACTICAL SIGNIFICANCE

Type I (α)	Reject true H_0 — "false alarm"
Type II (β)	Fail to reject false H_0 — "missed detection"
Power ($1 - \beta$)	$P(\text{reject } H_0 \mid H_a \text{ true})$; increases with n, α, effect size
Stat vs Prac	Small $p \not\Rightarrow$ large/important effect. Always consider the magnitude of the estimated difference.

General Reminders: Check assumptions before testing • Always state hypotheses in terms of *population parameters* (μ, p), not statistics ($\bar{x}$, $\hat{p}$) • Use R-lite for critical values and p-values • Report conclusions in context • CI & HT are complementary — use both when asked • Stat sig ≠ practical sig