Ch12-13: Factorial ANOVA, Blocking ANOVA, and χ^2 Test

- ***Please download:*** ***11-ANOVA.xls***
- ***HW8 this week***
- ***Presentations in two weeks!***

27 Mar 2012

Dr. Sean Ho

busi275.seanho.com

Outline for today

- Factorial ANOVA (multiple nominal IVs)
  - Graphing
  - Main effects and interaction
- Randomized Complete Block ANOVA
  - Fixed vs. random effects
  - Post-hoc analysis: Fisher's LSD
- χ^2 goodness-of-fit test: observed vs. expected
  - Test against uniform distribution
    - CHIDIST(), CHITEST()
  - Test of normality
  - Test 2 nominal vars for independence
    - Marginal probabilities

Overview: regression / ANOVA

- Both: DV (outcome) quantitative
- Regression: IVs quant. (or use dummy coding)
  - ANOVA: IVs nominal
- Global F-test: do all IVs together impact DV?
  - R^2: fraction of variability in DV explained
- Main effects: each IV's contribution
  - Significance (p-val) vs. effect size (unique R^2)
  - Follow-up (nominal IVs): Tukey-Kramer
    - See “Delivery” in **11-ANOVA.xls**
- Interaction:
  - Multiple regr: include $X_1 * X_2$ as a predictor
  - Factorial ANOVA (2-way, 3-way, etc.)

2-way ANOVA

DV: Purch Amt

	IV1: Gender	
	$40 $20 $25 $32	$17 $21 $19 $22
IV2: Src	$30 $26 $25	$22 $19
	$40 $45 $52 $43 $39 $48	$50 $60 $55

- Equivalent to multiple regression
  - Except with nominal predictors
  - N-way ANOVA for N predictors
- IVs are “between-groups” factors:
  - Divide up sample into cells
  - Each participant in only 1 cell
- If your IVs are mixed continuous / nominal, try regression using dummy-coded variables
  - Although this may result in many IVs!
- You can also try ANCOVA:
  - Continuous “covariates” are first factored out then regular ANOVA is done on residuals

2-way: assumptions

DV: Purch Amt

	IV1: Gender	
IV2: Src	$40 $20 $25 $32	$17 $21 $19 $22
	$30 $26 $25	$22 $19
	$40 $45 $52 $43 $39 $48	$50 $60 $55

- Same as for regular ANOVA, per cell:
  - DV continuous
  - Independent observations and independent cells (groups)
  - DV normal within each cell
  - Variance of DV similar across all cells:
    - (largest SD) / (smallest SD) < 2
    - Better: Levene's test of homoscedasticity
- The last two are less important as long as:
  - Total sample size is reasonably large (>50)
  - Balanced design: all cells similar sample size
  - No rows/cols are completely empty

Graphing 2-way ANOVA data

- Dataset: “Eukanuba” in **11-ANOVA.xls**
- The DV has a different distribution in each cell
- One way to visualize: condense it down to the average of DV within each cell
- Pivot Table:
  - Formula (row)
  - Plant (col)
  - Average of Fat (data)
- Try a line chart:

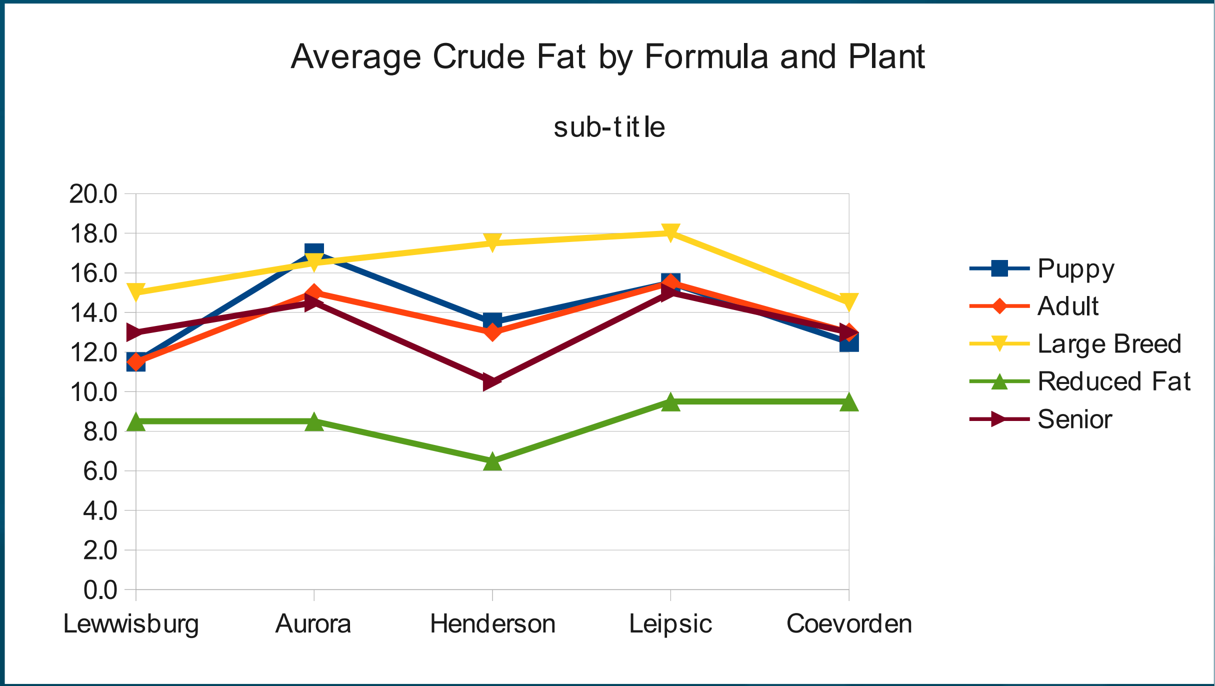

2-way ANOVA: model

- Main effects on each IV, plus interactions:
  - Fat = b_0 + (Formula effect) + (Plant effct) + (Formula*Plant effect) + ε
- Decomposition of variance:
  - $SS_{tot} = SS_{For} + SS_{Plt} + SS_{For*Plt} + SS_{resid}$
- Global F-test looks for any effect of IVs on DV
  - If not significant, check for violations of assumptions
- Effect size η^2 is akin to R^2: $1 - (SS_{resid} / SS_{tot})$

2-way ANOVA: calculating

-	IV_1 (a levels)	IV_2 (b levels)	IV_1*IV_2 (Interaction)
SS	$bn\sum_{i=1}^{a}(\bar{x}_i-\bar{x})^2$	$an\sum_{j=1}^{b}(\bar{x}_j-\bar{x})^2$	$n\sum_{i=1}^{a}\sum_{j=1}^{b}(\bar{x}_{ij}-\bar{x}_i-\bar{x}_j+\bar{x})^2$
df	a - 1	b - 1	(a - 1) * (b - 1)

- Also find SS_{tot} as before, and SS_{res}
  - df_{tot} = n – 1, and df_{res} = n – ab
  - The SS and df always add up:
    - Tot = IV_1 + IV_2 + (IV_1*IV_2) + Resid
- 3 F-tests: IV_1, IV_2, and interaction
  - e.g., main effect on IV_1: $F = MS_1 / MS_{res}$

Main effects

- A main effect is a one-way ANOVA on one IV, after controlling for the other predictors
  - Analogous to t-tests on slope for each IV in multiple regression
  - Here, the main effects are themselves F-tests
- E.g., do females spend more at your site, after accounting for source?
  - 2-way ANOVA on both Gender and Source, then look at main effect of Gender
- E.g., do different formulas have different fat content, across all plants?

Interactions

- When the effect of one IV on the DV changes, depending on the level of the moderator
- e.g., females spend more in response to print ads, but males spend more in response to web
- e.g., Henderson generally has lower fat than the other plants, except for Large Breed, where it has the second-highest fat:
- Plot means, note change in shape of the curves

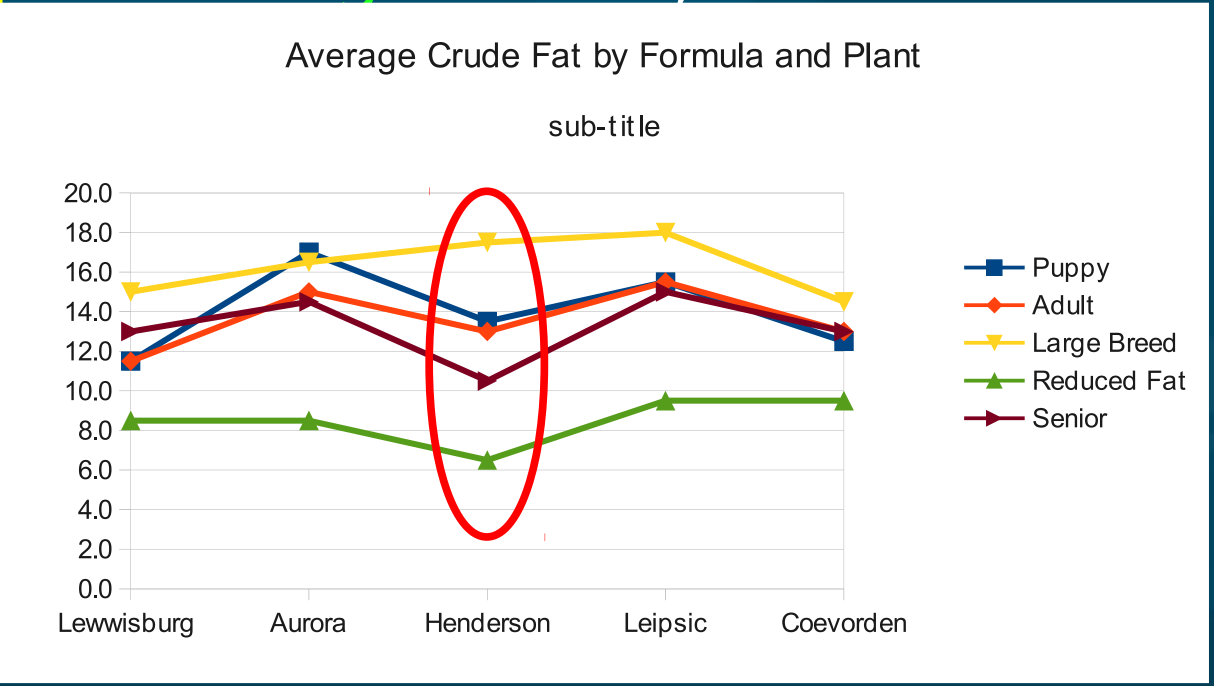

Outline for today

- Factorial ANOVA (multiple nominal IVs)
  - Graphing
  - Main effects and interaction
- Randomized Complete Block ANOVA
  - Fixed vs. random effects
  - Post-hoc analysis: Fisher's LSD
- χ^2 goodness-of-fit test: observed vs. expected
  - Test against uniform distribution
    - CHIDIST(), CHITEST()
  - Test of normality
  - Test 2 nominal vars for independence
    - Marginal probabilities

Randomized Complete Block

- Special case of 2-way ANOVA, where each cell has only 1 observation
- IV1 is the factor: typically a fixed effect
  - Fixed: levels are set in the hypothesis: e.g., gender, province, store branch, plant
- IV2 is the block: typically a random effect
  - Random: levels are sampled from a population: e.g., customer, truck, day
- e.g, “Applebees” in **11-ANOVA.xls**
  - Factor: Restaurant
  - Block: Week
  - DV: Revenue

	Restaurant		
Week	8.34	6.79	9.18
	10.7	10.0	12.8
	...	...	...
	...	...	...
	...	...	...

Randomized Block model

- A complete 2-way ANOVA on this data would have zero residual in each cell
  - So the interaction term serves as “residual”
  - Tot = Factor + Blocking + Residual
    - $df_{res} = (a - 1)(b - 1)$
- Factor effect (IV_1): $F = MS_1 / MS_{res}$
  - This is usually what we're most interested in
- Blocking effect (IV_2): $F = MS_2 / MS_{res}$
  - If non-sig, then blocking was not necessary and we could've just done a 1-way ANOVA

Post-hoc: Fisher's LSD test

- If the factor effect is significant, one post-hoc test we can use is Fisher's least sig. diff. test
  - Like Tukey-Kramer, but for equal-size cells
- Critical range:
  - t: 2-tails, use df_{res}
  - b: # blocks (IV_2)

$$LSD = t\sqrt{\frac{2\,MS_{res}}{b}}$$

- For all pairs of levels of the main factor, if the difference of means $|\bar{x}_i - \bar{x}_j|$ exceeds LSD, then those two groups differ significantly
  - Use the results to cluster the factor levels

Outline for today

- Factorial ANOVA (multiple nominal IVs)
  - Graphing
  - Main effects and interaction
- Randomized Complete Block ANOVA
  - Fixed vs. random effects
  - Post-hoc analysis: Fisher's LSD
- χ^2 goodness-of-fit test: observed vs. expected
  - Test against uniform distribution
    - CHIDIST(), CHITEST()
  - Test of normality
  - Test 2 nominal vars for independence
    - Marginal probabilities

Goodness of fit

- The χ^2 (chi-squared) test is one way to assess goodness of fit:
  - How well an observed distribution fits a hypothesized distribution
  - Hypothesized distribution can be uniform, normal, etc.

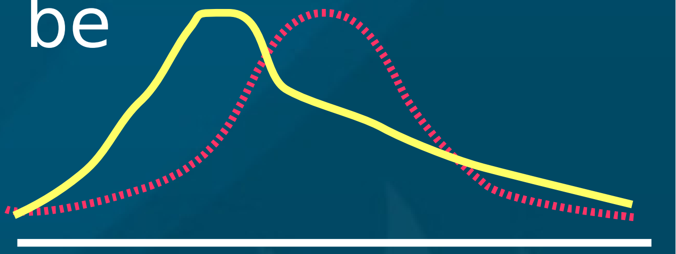

- χ^2 can also be applied to test if two nominal variables are independent
  - Compare pivot table (contingency table) with hypothesized results if vars independent
  - Analogous to correlation for continuous vars

χ^2 vs. uniform distribution

- e.g., are technical support calls evenly distributed across the weekdays?
  - H_0: evenly distributed, matches uniform dist.
- Expected # calls per day (uniform distribution):
  - Total observed calls (1300), divided by 5

Observed	**290**	**250**	**238**	**257**	**265**
Expected	**260**	**260**	**260**	**260**	**260**

- Test statistic:

$$\chi^2 = \sum \frac{(o_i - e_i)^2}{e_i}$$

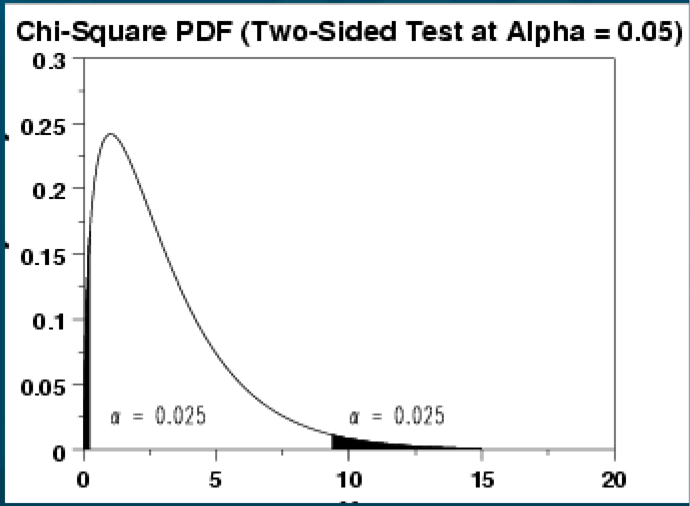

itl.nist.gov

- Use CHIDIST(χ^2, *#cells* – 1)
  - Or CHITEST(*obs*, *exp*)

χ^2 vs. normal distribution

- e.g., are student test scores normally distrib?
  - Other normality tests: Shapiro-Wilk, K-S
- Count frequency of test scores by bins
- How to find expected frequencies?
  - Find mean, SD of the data
  - Use NORMDIST() to find percentage of the data that would lie within each bin on the ideal normal:
    - NORMDIST(80, μ, σ, 1) – NORMDIST(75, μ, σ, 1)

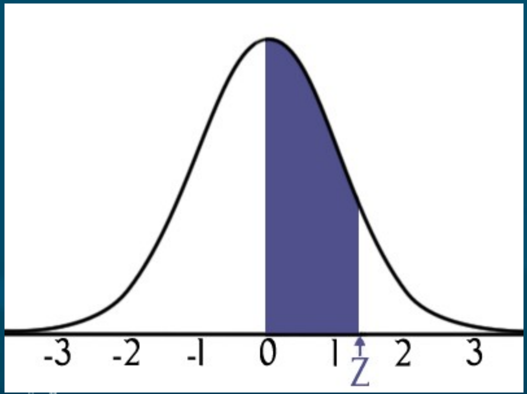

Bin	Freq	Norm Freq
40		
45	1	1.30
50	4	2.77
55	15	5.14
60	7	8.28
65	5	11.60
70	7	14.11
75	11	14.93
80	15	13.72
85	22	10.97
90	9	7.62
95	4	4.60
100	0	2.41

Outline for today

- Factorial ANOVA (multiple nominal IVs)
  - Graphing
  - Main effects and interaction
- Randomized Complete Block ANOVA
  - Fixed vs. random effects
  - Post-hoc analysis: Fisher's LSD
- χ^2 goodness-of-fit test: observed vs. expected
  - Test against uniform distribution
    - CHIDIST(), CHITEST()
  - Test of normality
  - Test 2 nominal vars for independence
    - Marginal probabilities

Contingency tables

- Joint freq. distribs for multiple nominal variables
  - Each cell of the table holds the # (frequency) of observations that match that combo
  - Pivot tables, with Count in the Data field
- E.g., Handedness vs. Gender
  - H_0: handedness is independent of gender: the probability of being left-handed stays the same, regardless of the gender
    - P(left | M) = P(left | F) = P(left)

Observed	**Left**	**Right**
Male	17	163
Female	33	167

χ^2 on 2-way contingency

- Expected values assume independence
- Calculate marginal probabilities:
  - P(female) = 200/380 ≈ 52.6%
  - P(left) = 50/380 ≈ 13.2%

	L	R	Tot
M	17	163	180
F	33	167	200
Tot	50	330	380

- Assuming independence,
  - P(F ∩ L) = P(F) * P(L) = (.526)(.132)
- Thus the expected count for (F ∩ L) is
  - P(F) * P(L) * (total) = (.526)(.132)(380)
- Calculate χ^2 summed over all cells
  - df = (#rows – 1) (#cols – 1)
    - = 1 in this case!

Summary on χ^2

- Test of goodness-of-fit: observed vs. expected
- May apply to a single nominal variable:
  - Expected distrib. may be uniform, normal, …
- May apply to two nominal variables:
  - Expected distrib. is that vars are independent
  - Akin to correlation on continuous variables
    - Large $\chi^2 \leftrightarrow |r| \approx 1$
- But only an approximation to the true distrib:
  - Results may be invalid if cell counts are <5
  - May need to combine levels of a var

TODO

- HW8 due Thu
- Projects: be pro-active and self-led
  - All groups have passed REB by now
  - Presentations on 10Apr (2 weeks from now!)
  - Remember your potential clients:
    what questions would they like answered?
  - Tell a story/narrative in your presentation
- Email me your preferences for presentation slot