Term Information

Effective Term	Spring 2019

General Information

Course Bulletin Listing/Subject Area	Statistics
Fiscal Unit/Academic Org	Statistics - D0694
College/Academic Group	Arts and Sciences
Level/Career	Undergraduate
Course Number/Catalog	1550
Course Title	Introduction to Statistical Reasoning
Transcript Abbreviation	Intr Stat Reason
Course Description	Introduction to statistical reasoning through data and application examples, including an introduction to coding in the R software; intended for students considering the Statistics major.
Semester Credit Hours/Units	Fixed: 3

Offering Information

Length Of Course	14 Week
Flexibly Scheduled Course	Never
Does any section of this course have a distance education component?	No
Grading Basis	Letter Grade
Repeatable	No
Course Components	Lecture
Grade Roster Component	Lecture
Credit Available by Exam	No
Admission Condition Course	No
Off Campus	Never
Campus of Offering	Columbus

Prerequisites and Exclusions

Prerequisites/Corequisites	Prereq or Coreq: Math 1152; or permission of the instructor.
Exclusions	
Electronically Enforced	Yes

Cross-Listings

Cross-Listings

Subject/CIP Code

Subject/CIP Code	27.0501
Subsidy Level	Baccalaureate Course
Intended Rank	Freshman, Sophomore

Requirement/Elective Designation

Required for this unit's degrees, majors, and/or minors
General Education course:
Data Analysis

Course Details

Course goals or learning objectives/outcomes

- Define and calculate basic probabilistic and statistical quantities (expectation, variance, hypothesis tests, confidence intervals, p-values).
- Critically assess statistical arguments.
- Explore statistical ideas through graphics and simulation in R.

Content Topic List

- Principles of data collection, including use of surveys
- Relationships among two or more variables
- Introduction to concepts of probability
- Introduction to statistical distributions
- Central limit theorem
- Concepts in hypothesis testing
- Frequentist vs. Bayesian inference
- Introduction to decision theory
- Correlation vs. causation
- Basics of statistical analyses in the R software

Sought Concurrence

No

Attachments

- GE_Rationale_1550.docx: GE Data Analysis Rationale
  (Other Supporting Documentation. Owner: Lee,Yoonkyung)
- GE_Assessment_Plan_1550.docx
  (GEC Course Assessment Plan. Owner: Lee,Yoonkyung)
- STAT1550_v2.pdf
  (Syllabus. Owner: Lee,Yoonkyung)
- Assessment_Question_Autumn_2016_JAS.pdf: Sample Assessment Questions
  (Other Supporting Documentation. Owner: Lee,Yoonkyung)

Comments

- [Nov 7] The GE goals and a grading scale have been added. The percents for grading components have been slightly adjusted, and attendance and participation has been clarified in the revised syllabus. Sample questions have been included in GE assessment plan in a separate file.

  [Sep 4] This course proposal is to be considered concurrently with the program proposal for the undergraduate Statistics major. *(by Lee,Yoonkyung on 11/07/2017 12:43 PM)*
- See 10-12-17 email to Y Lee. *(by Vankeerbergen,Bernadette Chantal on 10/12/2017 10:58 AM)*
- 09/05: Data Analysis ELOs and assessment plan requested. *(by Haddad,Deborah Moore on 09/05/2017 10:54 AM)*

Workflow Information

Status	User(s)	Date/Time	Step
Submitted	Lee,Yoonkyung	09/04/2017 09:14 PM	Submitted for Approval
Approved	Lee,Yoonkyung	09/04/2017 09:27 PM	Unit Approval
Revision Requested	Haddad,Deborah Moore	09/05/2017 10:54 AM	College Approval
Submitted	Lee,Yoonkyung	09/12/2017 01:36 PM	Submitted for Approval
Approved	Lee,Yoonkyung	09/12/2017 01:37 PM	Unit Approval
Approved	Haddad,Deborah Moore	09/12/2017 02:10 PM	College Approval
Revision Requested	Vankeerbergen,Bernadette Chantal	10/12/2017 10:58 AM	ASCCAO Approval
Submitted	Lee,Yoonkyung	11/07/2017 12:44 PM	Submitted for Approval
Approved	Lee,Yoonkyung	11/07/2017 12:44 PM	Unit Approval
Approved	Haddad,Deborah Moore	11/07/2017 01:38 PM	College Approval
Pending Approval	Nolen,Dawn Vankeerbergen,Bernadette Chantal Oldroyd,Shelby Quinn Hanlin,Deborah Kay Jenkins,Mary Ellen Bigler	11/07/2017 01:38 PM	ASCCAO Approval

Syllabus for Stat 1550: Introduction to Statistical Reasoning

Instructor:
Office:
Office Hours:
Office Phone:
E-mail:

Format: Three credit hours; three 55-minute in-class meetings per week

Prerequisites: Prerequisite or concurrent: Math 1152; or permission of the instructor.

Required Text: *How Not to be Wrong*, by Jordan Ellenberg, Penguin Press, 2014.

Course Description and Learning Outcomes: This course is intended for students (considering) majoring in Statistics. The goal of the class is to introduce statistical reasoning primarily through data and application examples, while attempting to give a flavor of what more rigorous statistical study is like. Students will learn about statistical arguments and statistical reasoning through example. They will also have some exposure to the basic mathematical arguments underlying statistical arguments, and some experience with coding in the R software.

Upon successful completion of the course, students will be able to

1. Define and calculate basic probabilistic and statistical quantities (expectation, variance, hypothesis tests, confidence intervals, p-values)
2. Critically assess statistical arguments
3. Explore statistical ideas through graphics and simulation in R.

This course satisfies the General Education (GE) requirement in *Data Analysis*.
Specifically, the **goal** is: Students develop skills in drawing conclusions and critically evaluating results based on data.
The **expected learning outcomes** are: Students understand basic concepts of statistics and probability, comprehend methods needed to analyze and critically evaluate statistical arguments, and recognize the importance of statistical ideas.

Assignments: Homework with some statistical and mathematical calculations; some conceptual questions; and some programming will be assigned approximately weekly. Some assignments will ask students to collect data and analyze it through a web interface (e.g., an R Shiny App) and discuss their findings. Students will be assessed on two exams (a midterm and a final). Class attendance and participation will also count toward the final grade.

Final Grade:
The final course grade will be based on the following weighting of assessment components:

Attendance and Participation – 10%
Homework – 30%
Midterm – 30%
Final – 30%

Attendance and Participation will be assessed in the following way:

- On 6 days selected at random throughout the semester, attendance will be taken; students will be expected to be present on at least 5 of these days, which will count toward 5% of the 10%.
- Students will be assigned in small groups to give a short presentation to the class, which will count toward the remaining 5%. These presentations will happen throughout the semester. The topic will be provided to the students, and will typically involve:
  - presenting their homework assignment that week (e.g., if the assignment involved collecting and analyzing data)
  - presenting on an example from the required text (or potentially an online article)

Grading Scale: The following grading scale will be used:

>90%	A/A-
80-90%	B-/B/B+
70-80%	C-/C/C+
60-70%	D/D+
<60%	E

Tentative Course Schedule:

Week	Topic
1	**Introduction.** Different types of sampling and their impact on inference. **Book Example:** (Introduction) Survivorship bias in bullet hole locations of military planes. **R:** Basic usage. **Assignment:** Conduct a survey; describe the type of sampling and scope of inference.
2	**Linearity and Nonlinearity.** Different type of relationships that might exist between variables. **Book Example:** (Ch. 1) the Laffer curve in economics. **R:** Basic graphing in R – plotting lines and curves. **Assignment:** R lab with plotting commands.
3	**Interpolation and Extrapolation.** Local linearity; predictions within and outside range of data. **Book Example:** (Ch. 3) the prediction that *everyone* will be obese by 2048. **Other Examples:** The Challenger explosion. **R:** Scatterplots in R; overlaying lines, curves, the lowess smoother. **Assignment:** R lab with data – a time series example, with questions about prediction.

4 **Basic Distribution Theory:** The Binomial and Normal Distributions.
Basic probability; definitions; properties; real-life applications; population vs. sample.
R: Normal curves in R; random sample generation.
Assignment: Mathematical calculations; collect data from a random mechanism (coins, dice).

5 **The Law of Large Numbers and the Central Limit Theorem.**
Book Example: (Ch. 4) coin flips; rare disease rates in large and small states; NBA shooting rates.
Other Examples: Students' collected data.
R: Illustration of these theorems by simulation.
Assignment: Simulation-based illustration of these theorems – with a different distribution.

6 **Midterm.**
The probability of improbable events.
Book Example: (Ch. 6) the Baltimore stockbroker (i.e. the Swami scam); the Bible code.
Other Examples: The birthday problem; "garden of forking paths" examples.

7-8 **Hypothesis Testing.** The standard statistical set-up.
p-values; statistical and clinical significance; multiple testing; publication bias; confidence intervals.
Book Example: (Ch. 7, 9) fMRI study of a dead fish; oral contraceptives and blood clots
The international journal of haruspicy.
Other Examples: studies of ESP.
Assignment. Mathematical calculations and conceptual exercises.

9 **Frequentist And Bayesian Inference.** Priors and posteriors. Rare event prediction.
Book Example: (Ch. 10) Does Facebook know you're a terrorist?
Other Examples: Rare disease screening – sensitivity vs. positive predictive value.
Assignment: Some calculations with Bayes Theorem.

10 **Expectation.**
Book Example: (Ch. 11) The lottery; betting; annuities; Buffon's needle
Assignment: Mathematical calculations.

11 **Loss / Utility Functions.**
Book Example: (Ch. 12) Missing planes; Pascal's wager; the St. Petersburg paradox
Assignment: Mathematical calculations.

12 **Variance.**
Book Example: (Ch. 13) More betting questions
Assignment: Mathematical calculations.

13-14 **Correlation vs. Causation.** Sampling and Experimental design.
Book Example: (Ch. 16) Relationships between smoking, marital status, and lung cancer.
Berkson's fallacy.
Assignment: Article critique (for one of several choices of articles) – correlation and causation.

Academic Misconduct: It is the responsibility of the Committee on Academic Misconduct to investigate or establish procedures for the investigation of all reported cases of student academic misconduct. The term "academic misconduct" includes all forms of student academic misconduct wherever committed; illustrated by, but not limited to, cases of plagiarism and dishonest practices in connection with examinations. Instructors shall report all instances of alleged academic misconduct to the committee (Faculty Rule 3335-5-487). For additional information, see the Code of Student Conduct `http://studentlife.osu.edu/csc/`.

Special Accommodations: Students with disabilities (including mental health, chronic or temporary medical conditions) that have been certified by the Office of Student Life Disability Services will be appropriately accommodated and should inform the instructor as soon as possible of their needs. The Office of Student Life Disability Services is located in 098 Baker Hall, 113 W. 12th Avenue; telephone 614-292-3307, slds@osu.edu; `slds.osu.edu`.

STAT 1550 GE Rationale

The goal of STAT 1550 is to introduce statistical reasoning through data and application examples, while attempting to give a flavor of what more rigorous statistical study is like. Students will learn about statistical reasoning, the basic mathematical arguments underlying statistical arguments, and coding in the R software.

The requirements for a course to be considered a GE Data Analysis course are reproduced below and numbered. Then, the topics taught in STAT 1550 are presented (as provided in the syllabus), together with an indication of which requirements these meet and expected number of instructional hours.

Core GE Data Analysis Requirements (at least 4 instructional hours spent on each)

- **C1.** Notions of probability. The axioms of probability, and basic probability calculations. Random variables, and probability calculations using random variables. Expected values.
- **C2.** Basics of statistical inference. Moving from a sample to a population. Bias and variance. Understanding the margin of error and confidence. The logic of statistical testing. The misuse of statistics.

Additional requirements (At least two out of four, with at least 3 instructional hours spent on each)

- **A1.** Understanding where data come from. Data sources. Discriminating between observational and experimental studies. (Random) sampling.
- **A2.** Summarizing data graphically and numerically. Discriminating between good and bad summaries. Understanding the advantages and disadvantages of a given summary.
- **A3.** Methods of statistical inference. Statistical testing. Constructing confidence intervals. Making quantitative statistical arguments using data. Understanding and verifying assumptions underlying a given inference.
- **A4.** Statistical modeling (e.g., regression models, analysis of variance). Interpreting the parameters underlying statistical models. Model assessment.

Presented below are the topics covered each week. This course focuses on statistics and data analysis, so generally all material each week falls into one of the six categories listed above. We have allocated instructional hours proportionally to the relevant categories, assuming 3 course hours per week. Even with some adjustment (such as holidays) to this allocation, we will meet the minimum number of instructional hours for the two core requirements, as well as for additional requirements A1 and A3.

Week	Main Topics	C1	C2	A1	A2	A3	A4
1	Different types of sampling and their impact on inference.			3			
2	Linearity and Nonlinearity. Different type of relationships that might exist between variables. Basic plotting in R.				2		1
3	Interpolation and Extrapolation. Local linearity; predictions within and outside range of data.		1		1		1
4	Basic Distribution Theory: The Binomial and Normal Distributions. Basic probability; definitions; properties; real-life applications; population vs. sample.	3					
5	The Law of Large Numbers and the Central Limit Theorem.	3					
6	Midterm. The probability of improbable events.	1	1				
7-8	Hypothesis Testing. The standard statistical set-up. p-values; statistical and clinical significance; multiple testing; publication bias; confidence intervals.		3			3	
9	Frequentist And Bayesian Inference. Priors and posteriors. Rare event prediction.	3					
10	Expectation.	2	1				
11	Loss / Utility Functions.		1			1	1
12	Variance.	2	1				
13-14	Correlation vs. Causation. Sampling and Experimental design.		2	3		1	
Total		14	10	6	3	5	3

Assessment of Data Analysis Learning Outcomes

Expected learning outcomes for GE in Data Analysis

ELO1	Students understand basic concepts of statistics and probability.
ELO2	Students comprehend methods needed to analyze and critically evaluate statistical arguments.
ELO3	Students recognize the importance of statistical ideas.

Method of assessment

Direct assessment of how well STAT 1550 achieves the three learning outcomes will be made through the use of embedded questions in the final. The instructor will be asked to indicate each student's level of achievement for each of the three ELOs based on the student's answers to the embedded questions. Table 1 provides general guidelines for selection of assessment questions for each of the ELOs.

Table 1. Guidelines for Assessment Questions

Outcome	Type of assessment questions
ELO1: Students understand basic concepts of statistics and probability.	Fundamental concept in statistics and/or probability, such as test-statistics, p-values, confidence intervals. Evaluate whether the student understands: - the definition of a given concept - how to calculate the quantity related to the concept - how to interpret the quantity in the context of the problem
ELO2: Students comprehend methods needed to analyze and critically evaluate statistical arguments.	One or several statistical methods, such as t-tests, analysis of variance, linear regression. Evaluate whether the student understands: - how the design of the data impacts the choice of an appropriate method - what are the core assumptions of a statistical method - how to evaluate the assumptions of the chosen method with the data

ELO3: Students recognize the importance of statistical ideas.	Understanding of principles, impacts and limitations of statistical analyses. Evaluate whether the student understands: - issues about using statistics in practice (e.g., using statistics for decision making, statistical vs. practical significance) - issues of statistical integrity (e.g., dealing with outliers, planned vs. unplanned tests, searching for p-values) - scope of inference (e.g., relating the sample to the population, causation vs. correlation, interpolation vs. extrapolation)

Level of student learning expected for ELOs

Table 2 shows a scoring rubric designed to help instructors and members of relevant committees assess how well students are meeting the Expected Learning Outcomes (ELO's) in a GE Data Analysis course. ELOs are deemed successful if at least 60% of students are ranking as Good or above in the rubric on an aggregate basis.

Table 2. A Scoring Rubric

Outcome	Excellent	Good	Fair	Poor
ELO1: Students understand basic concepts of statistics and probability.	Articulates clearly and concisely the definition and interpretation, and executes the calculation correctly.	Provides an imprecise but generally correct definition or interpretation with minor errors.	Demonstrates understanding of the material, but with some major errors.	Demonstrates little to no understanding of the material.
ELO2: Students comprehend methods needed to analyze and critically evaluate statistical arguments.	Demonstrates good understanding of how and when to apply a statistical method and how to evaluate its assumptions.	Demonstrates good but imprecise understanding of a statistical method or its assumptions with minor errors.	Demonstrates understanding, but with some major errors.	Demonstrates little to no understanding of the material.

ELO3: Students recognize the importance of statistical ideas.	Exhibits good understanding of statistical principles in data analysis.	Exhibits generally good but partial understanding of the principles.	Demonstrates some understanding, but with some major errors.	Demonstrates little to no understanding of the material.

Review process

After each offering of the course, the assessment results will be compiled by the instructor and reported to the undergraduate major committee. Information from the results will be used to provide guidance to subsequent instructors and will be shared with the Statistics undergraduate study chair and the co-director of the data analytics major. Revision of the assessment method or recalibration of the success criteria for achieving ELOs will be considered as more data are collected over time.

Sample questions for assessment

To provide examples of assessment questions, we attach a copy of the embedded questions that were used for assessment in the final exam of the Autumn 2016 offering of STAT 5301 (a GE Data Analysis course). In the exam, the first ELO was assessed by asking students to calculate a p-value and interpret its meaning with Question 3(a,b). The second ELO was assessed by asking students to identify the assumptions underlying two possible analyses of the same data, and to judge the appropriateness of each analysis based on whether these assumptions are met with Question 3(c,d). The third ELO was assessed by asking whether a causal relationship could be inferred from a randomized experiment, and about the applicability of the results to a different setting with Question 3(e,f).

3. (13 points) The "misery is not miserly" phenomenon refers to a person's spending judgment going haywire when sad. In a recent study, 40 young adults were each given \$10 and randomly assigned to either a Sad or Neutral group. The participants in the Sad group each watched a video about the death of a boy's mentor (from *The Champ*) and those in the Neutral group each watched a video on the Great Barrier Reef. After the video, each participant (separately) was offered the chance to trade some of the \$10 (in \$0.50 increments) for an insulated water bottle. The amount of money that each student was willing to trade for the water bottle was recorded, and these values are presented below:

Group	Purchase Price, in \$																			
Sad	0.0	3.0	1.0	0.0	0.0	0.0	4.0	4.0	4.0	4.0	0.0	3.0	1.0	1.0	0.0	3.5	1.5	2.5	2.0	3.5
Neutral	0.5	0.0	2.0	0.5	2.0	0.0	2.0	0.0	2.5	1.0	0.0	0.5	0.0	1.0	1.5	0.0	1.0	0.0	0.5	0.0

Let μ_1 be the mean of the Sad group and μ_2 be the mean of the Neutral group. Suppose that researchers wanted to test whether there is *any difference* in the purchase price between the two groups – i.e., to test:

$$H_0 : \mu_1 - \mu_2 = 0 \text{ vs. } H_A : \mu_1 - \mu_2 \neq 0$$

Two statisticians (Trevor and Edith) analyze the data. Some details of their analyses are presented on the next two pages; questions about these analyses begin on page 9.

Trevor's Analysis

Trevor decides to use the standard two-sample t-test (i.e., using the pooled estimate of the variance). He calculates the following summary statistics:

Group	N	Sample Mean	Sample Standard Deviation
1_Sad	20	1.90	1.62
2_Neutral	20	0.75	0.84

He looks at qq-plots of the values in each group and runs the function `t.test` in R; the plots and the output are printed below. **Note** that the p-value has been removed from the `t.test` output.

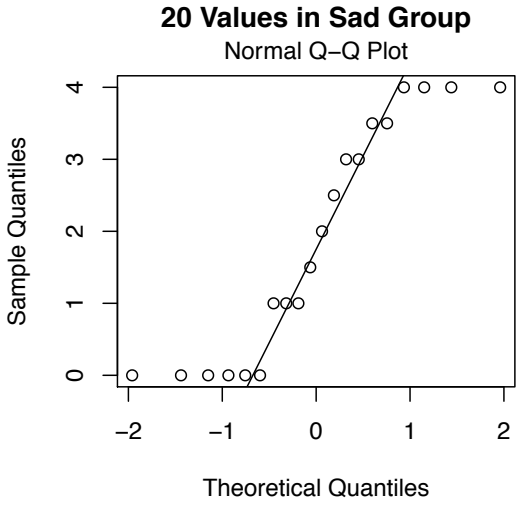

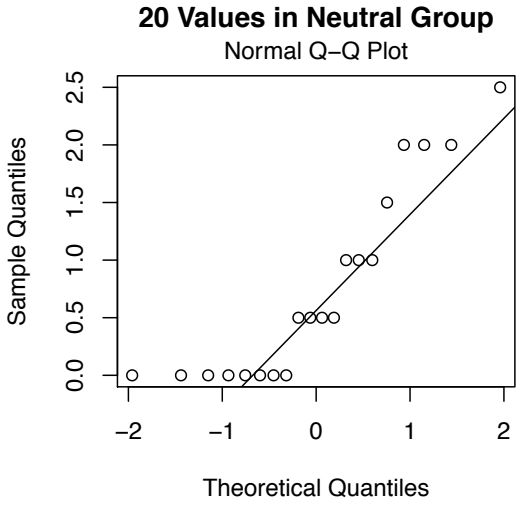

Trevor's R Output

```
> t.test(sad, neutral, var.equal=TRUE)

        Two Sample t-test

data:  sad and neutral
t = 2.8232, df = 38, p-value = XXXXX
alternative hypothesis: true difference in means is not equal to 0
95 percent confidence interval:
 0.3253953 1.9746047
sample estimates:
mean of x mean of y
     1.90      0.75
```

Edith's Analysis

Edith begins by taking pairwise differences, defining a new variable `diff = sad - neutral`. She decides to use the one-sample t-test on this new variable. She calculates the following summary statisitcs:

Variable	N	Sample Mean	Sample Standard Deviation
Diff	20	1.15	1.81

She looks at the qq-plot of the values of this new variable, and runs the function `t.test` in R; the plots and the output are printed below.

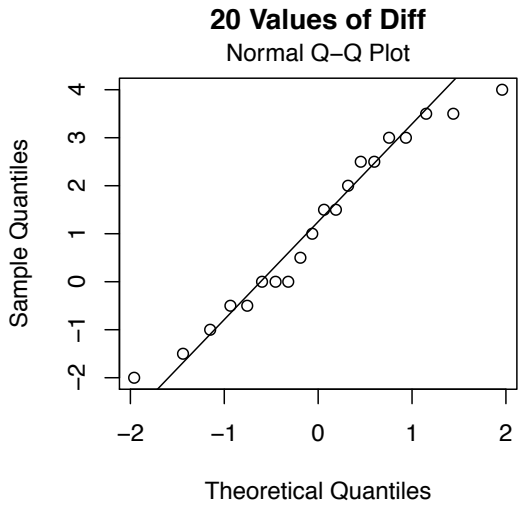

Edith's R Output

```
> t.test(diff)

        One Sample t-test

data:  diff
t = 2.8459, df = 19, p-value = 0.01033
alternative hypothesis: true mean is not equal to 0
95 percent confidence interval:
 0.304228 1.995772
sample estimates:
mean of x
     1.15
```

(a) (2 points) Let P denote the p-value in Trevor's analysis. Use the t-distribution table to calculate this p-value. Show all work. (You may express your answer as two values a and b so that $a \leq P \leq b$.)

(b) (2 points) To summarize his result, Trevor writes "The probability that H_0 is true is P." What is wrong with this statement? Using the definition of a p-value, correct the statement and describe how to interpret the p-value calculated in (a).

(c) (3 points) What are the main assumptions of Trevor's analysis? Is Trevor's analysis trustworthy? Why or why not?

(d) (3 points) What are the main assumptions of Edith's analysis? Is Edith's analysis trustworthy? Why or why not? Do you prefer Trevor's analysis or Edith's analysis?

(e) (2 points) Trevor concludes from his analysis that the differences he observed between the two groups are *due to* the video that the students in each group watched. Is Trevor's claim of a *causal* relationship between the video and the purchase price reasonable? Why or why not?

(f) (1 point) A water bottle manufacturer becomes interested in the results of this study as a potential approach for increasing revenue by *raising the prices* of their water bottles. To advertise on websites, the company pays a flat fee plus a commission for each water bottle sold. Current ads are still images of nature, and the company pays \$0.01 per resulting water bottle purchase. Based on the study, the company is considering replacing its ads with *sad video ads*; for these new video ads, the company would pay \$3.00 per resulting water bottle purchase. Would you advise the company to test out this approach? Why or why not?