

# **1 Introduction**

Statistics and Bioinformatics  
University of Zadar  
Winter, 2010-2011

## Contact Information

- Instructor: doc. dr. sci. Stewart T. Schultz
  - email: [sschultz@unizd.hr](mailto:sschultz@unizd.hr)
  - phone: 200-653
- Teaching Assistant: Tomislav Sarić
  - email: [tosaric@unizd.hr](mailto:tosaric@unizd.hr)
  - phone: 095 904 6339
- Course website: [www.stewartschultz.com/statistics/course](http://www.stewartschultz.com/statistics/course)
  - Username: `zadar77`
  - Password: `zadar77`

If you have any questions, feel free to contact either one of us by email or phone.

## Class schedule

- Lectures
  - Tuesday: 14:00 to 16:00 (two hours), SK-INF
  - Thursday: 15:00 to 16:00 (one hour), SK-INF
- Computer laboratory
  - Tuesday: 16:00 to 18:00 (two hours), SK-INF

Laboratory exercises will be using R, a programming language and an environment for statistical computing and graphics. See <http://www.r-project.org/>.

## Course requirements

- Problem sets (15)
- 3 midterm exams
- Optional final exams

You learn statistics by doing statistics, therefore we have weekly problem sets, which will be handed out on Thursday at the laboratory session, and returned the following Thursday.

## What is statistics?

- Set of tools
- to convert *data*
- to *probability*.

In this course, our data examples for the most part will come from agricultural research.

## What is bioinformatics?

- Application of statistical tools
- to convert *molecular data*
- to *probability*.

Bioinformatics originally referred to information contained in any biological system, from genomes to ecosystems, but since 1980 it has referred exclusively to molecular data, primarily the analysis of sequences of DNA and amino acids.

## 2 Probability

### What is probability?

- Axiomatic
  1.  $0 < P(A) < 1$
  2. Impossible event:  $P(A) = 0$
  3. Certain event:  $P(A) = 1$
  4. Complement of  $A$ ,  $A'$ :  $P(A') = 1 - P(A)$
  5. Two incompatible events  $A$  and  $B$ :  $P(A \cup B) = P(A) + P(B)$
  6. Any two events  $A$  and  $B$ :  $P(A \cup B) = P(A) + P(B) - P(A \cap B)$

- Relative frequencies

– Repeat the experiment  $n$  times, tally the frequency of event  $A$ ,  $n_A$ .

–

$$P(A) = \lim_{n \rightarrow \infty} \frac{n_A}{n}$$

Both definitions of probability are useful, the axiomatic definition for simple symmetric systems and cases of random sampling, and the relative frequency definition within any experimental context. These axioms are easily visualized in a Venn Diagram or probability tree.

### Conditional probability

**Definition 1.** The probability of  $A$  given  $B \equiv P(A/B)$ .

$$P(A/B) = \frac{P(A \cap B)}{P(B)}$$

Thus,

$$P(A \cap B) = P(A/B)P(B)$$

Conditional probability simply recognizes the intuitive fact that if you know something about the context of an event, then its probability might change.

### Independent events

**Definition 2.** Event  $A$  is independent of event  $B$  if and only if

$$P(A/B) = P(A)$$

Thus,

$$P(A \cap B) = P(A/B) \times P(B) = P(A) \times P(B)$$

“Product rule for independent events.”

Event  $A$  is independent of event  $B$  if the probability of  $A$  is the same whether or not  $B$  occurs.

### Bayes' Theorem

**Theorem 3.**

$$\begin{aligned} P(A/B) &= \frac{P(B/A)P(A)}{P(B)} \\ &= \frac{P(B \cap A)}{P(B \cap A) + P(B \cap A')} \\ &= \frac{P(B/A)P(A)}{P(B/A)P(A) + P(B/A')P(A')} \end{aligned}$$

Bayes' Theorem allows calculation of the probability of event  $A$  conditioned on  $B$ , if you know the probability of event  $B$  conditioned on  $A$ . This is easily illustrated by a Venn Diagram or by a probability tree.

## Screening test

**Definition 4.** A screening test is a preliminary test given to a sample of individuals to determine if they are likely to have a condition such as a disease, genotype, or contamination. Two events are analyzed:  $C$ , the presence of the condition; and  $T$ , a positive test result.

### Screening test properties

Fundamental properties of a screening test  $T$  for a condition  $C$ :

- Positive predictive value: probability of having the condition given a positive test result;  $PPV = P(C/T)$ .
- Negative predictive value: probability of not having the condition given a negative test result;  $NPV = P(C'/T')$ .
- False positive rate: the probability of a positive test result given that the individual does not have the condition;  $\alpha = P(T/C') = 1 - \text{specificity}$ .
- False negative rate: the probability of a negative test result given that the individual has the condition;  $\beta = P(T'/C) = 1 - \text{sensitivity}$ .

- $$\text{Power} = 1 - \beta = \text{sensitivity}$$

Illustrate with a probability tree.

## Fagan plot

**Definition 5.** An intuitive plot that relates the prior probability of a condition  $P(C)$  to the posterior probability of a condition  $P(C/T)$ , based on the ratio of true positive rate to false positive rate, or power/ $\alpha$ .

## Lecture 2: Review of basic concepts of probability

### Another textbook in Croatian

Franji, Iva. (Bio)Statistika. Zagreb.

[www.pbf.hr/hr/content/download/11976/54038/version/1/.../SKRIPTA.pdf](http://www.pbf.hr/hr/content/download/11976/54038/version/1/.../SKRIPTA.pdf)

### Basic concepts of probability

- Definition
- Independence
- Conditional probability
- Screening test
- Bayes Theorem
- Venn Diagram, Probability Tree

## 3 Counting Rules

### Lecture 3: Counting rules

## Schedule for today

- Lecture 14:00 to 16:00
- Computer lab 1 16:00 to 17:00
- Computer lab 2 17:00 to 18:00

## Basic concepts of probability

- Definition
- Independence
- Conditional probability
- Screening test
- Bayes Theorem
- Venn Diagram, Probability Tree

## What is statistics?

**Definition 6.**     • Set of tools

- to convert *data*
- to *probability*.

## For example

Bayes' Theorem, derived from probability axioms.

Review the example from yesterday, draw the diagram again. Prevalence = 0.1, power = 0.8, and false positive rate = 0.044.

Emphasize: Definition of prevalence. Prevalence is just  $P(C)$ , the probability of the condition in the population that you are interested. So in English we refer to the prevalence of AIDS, or the prevalence of breast cancer, or the prevalence of drug use, the prevalence of obesity, prevalence of alcoholism. All these are the overall probability that a random person drawn from the population has the condition.

The positive predictive value of the test is the proportion of all tests that are true positives.

The negative predictive value of the test is the proportion of all tests that are true negatives.

The power of the test is the proportion of all positive people who test positive.

The false positive rate of the test is the proportion of all negative people who test positive.

## Screening test in R

In R, we can construct a screening test table using:

- package TeachingDemos
- SensSpec.demo()

There are three ways to visualize conditional probability: probability trees, Venn diagrams, and  $2 \times 2$  tables. The R demonstration constructs a full screening test table from three inputs: the sensitivity of the test, the specificity of the test, and the prevalence of the condition (or disease) in the population.

Illustrate it with the EMIT test, Enzyme Multiplied Immunoassay Test, for cocaine:  $\alpha = 0.05$ , power = 0.99, prevalence = 0.01.

False positive rate for opiates is 0.45 False positive rate for amphetamines is 0.8.

Example: mammography

power = 0.85 false positive rate = 0.05 prevalence in the population of breast cancer, for 40 yo women = 0.0064

## 4 Random variables

### Random variables

**Definition 7.** A random variable is a quantity whose value is not fixed, but which can take on different values according to a probability distribution.

Example: the number of boys in a family of three. There are four possibilities, 0, 1, 2, and 3, and each of these possibilities can be assigned a probability. In order to calculate these probabilities, first we need to tally up all the possible outcomes for families of three. Here are the possible outcomes: BBB, GGG, BGG, GBG, GGB, GBB, BGB, BBG. So there are eight possible outcomes. Now if we assume that each of these outcomes is equally likely, then we can calculate the probabilities.

$$P(0) = 1/8 \quad P(1) = 3/8 \quad P(2) = 3/8 \quad P(3) = 1/8$$

So: we have a random variable, let's call it  $x$ , the number of boys in a family of three children. There are four possible values of this variable. Each of these values has an associated probability.

This result can be illustrated with the binomial expansion.

$$(p + q)^3 = p^3 + 3p^2q + 3pq^2 + q^3.$$

### Axiomatic definition of probability

**Definition 8.**

$$P(A) = \frac{\text{number of outcomes giving event A}}{\text{total number of equally likely outcomes}}$$

### Fundamental counting principle

**Definition 9.** If a first task can be done in any one of  $m$  different ways and, after that task is completed, a second task can be done in any one of  $n$  different ways, then both tasks can be done, in the order given, in  $m \times n$  different ways.

This can be illustrated with a probability tree diagram.

### Permutations

**Definition 10.** A permutation is an ordered arrangement of a set of distinctly different items. For example, two permutations of the numbers one through five are 12345 and 53142.

### Permutations

**Theorem 11.** *Through a direct application of the Fundamental Counting Principle, the number of permutations of  $n$  distinct items is given by*

$${}_nP_n = n! = n(n-1)(n-2) \dots (2)(1)$$

*The number of permutations of  $r$  items selected from  $n$  distinctly different items is*

$${}_nP_r = \frac{n!}{(n-r)!}$$

Example: 6 people seated in six chairs along a straight line. Example: 4 people out of 6 seated in four chairs along a straight line.

### Combinations

**Definition 12.** A combination is an arrangement of a set of distinct items in which different orderings of the same items are considered the same. In other words, a combination is just a subset of a set of distinct elements.

**Theorem 13.** *The number of combinations of  $r$  items selected from  $n$  distinct items is*

$${}_nC_r = \frac{n!}{r!(n-r)!} = \binom{n}{r}$$

Example: you go to the store to buy two shirts. The store has 3 different kinds of shirts. How many combinations of two shirts out of 3 can you buy? You can buy shirts 1 and 2, shirts 1 and 3, or shirts 2 and 3. So there are three possible combinations of two shirts selected out of a set of three shirts. Note that shirts 1 and 2 is the same combination as shirts 2 and 1, because order does not matter in this problem.

To understand the equation, we can arrange the possibilities under the general equation:

(number of combinations)  $\times$  (number of permutations per combination) = total number of permutations

The total number of permutations is six: 12, 21, 13, 31, 23, 32. The number of permutations per combination is 2 (12 and 21, etc.). So the number of combinations is three.

We can then write down each of these: the number of permutations is  $3 \times 2$ . In terms of factorials, we can write this as  $3!/1!$ . The number of permutations per combination is  $2 \times 1$ , which is  $2!$ . So if we solve this equation for the number of combinations, we have

$$\text{number of combinations} = \frac{\text{total number of permutations}}{\text{number of permutations per combination}} = \frac{n!}{r!(n-r)!}$$

### Laboratory exercises in R

- TeachingDemos: a package that illustrates statistical relationships
- scan() for entering small quantities of data
- sample() for random sampling from datasets
- SensSpec.demo() for constructing a screening test table
- fagan.plot() for illustrating screening test

## 5 Review of counting rules and binomial distribution

### Lecture 4: Review of counting rules and binomial distribution

#### What is statistics?

**Definition 14.**     • Set of tools

- to convert *data*
- to *probability*.

#### Our tools so far:

- Product rule for independence
- Bayes' theorem for the probability of a condition
- Binomial distribution for probability of "successes" in a binomial experiment

#### Binomial distribution is derived from

- Fundamental counting rule
- Number of permutations of  $n$  items taken  $r$  at a time
- Number of combinations of  $n$  items taken  $r$  at a time
- We assume independence!

## Fundamental counting principle

**Definition 15.** If a first task can be done in any one of  $m$  different ways and, after that task is completed, a second task can be done in any one of  $n$  different ways, then both tasks can be done, in the order given, in  $m \times n$  different ways.

Examples: Number of possible phone numbers beginning with a given prefix, say 200. Number of possible passwords, assuming either a letter or a digit, assuming say 26 letters and 10 digits, and letters can be either upper or lower case. Monkeys at a keyboard typing letters at random, what is the probability that a monkey types out Shakespeare's Hamlet? There are 199749 characters in Shakespeare's Hamlet.

Answer: around one in  $37^{199749} = 10^{313246.7}$ .

## Permutations

**Definition 16.** A permutation is an ordered arrangement of a set of distinctly different items. For example, two permutations of the numbers one through five are 12345 and 53142.

## Permutations

**Theorem 17.** *Through a direct application of the Fundamental Counting Principle, the number of permutations of  $n$  distinct items is given by*

$${}_nP_n = n! = n(n-1)(n-2)\dots(2)(1)$$

*The number of permutations of  $r$  items selected from  $n$  distinctly different items is*

$${}_nP_r = \frac{n!}{(n-r)!}$$

Examples: how many ways is it possible to seat 25 people in a straight line of chairs? Answer: 25!. How many distinct ways is it possible to arrange the 10 letters of the word *statistics*? Let's say we want to elect a president, vice-president, and secretary of our class. How many different ways can we do this?

## Combinations

**Definition 18.** A combination is an arrangement of a set of distinct items in which different orderings of the same items are considered the same. In other words, a combination is just a subset of a set of distinct elements.

**Theorem 19.** *The number of combinations of  $r$  items selected from  $n$  distinct items is*

$${}_nC_r = \frac{n!}{r!(n-r)!} = \binom{n}{r}$$

Examples: say we go shopping for shirts. At the store there are 10 shirts for sale and we want to buy 3. How many different combinations of 3 shirts out of 10 are possible? Let's say we want to buy two pairs of pants out of 15 that are available. How many combinations are possible? Now how many possible combinations of both pants and shirts are possible if we have 3 shirts and 2 pants, out of 10 and 15?

# 6 Random variables: Binomial and Poisson

## Lecture 5: Random variables

Today's schedule:

- Lecture on random variables, binomial, poisson 14:00 to 16:00
- Laboratory exercise Group 1: 16:00 to 17:00



- Laboratory exercise Group 2: 17:00 to 18:00
- Problem Set 2 is due today at the end of the laboratory session

Announcement: notes from graphics pad will be uploaded to website as a JPEG.  
Questions?

## What is statistics?

**Definition 20.** • Set of tools

- to convert *data*
- to *probability*.

## Random variables

**Definition 21.** A random variable is a quantity whose value is not fixed, but which can take on different values according to a **probability distribution**. A probability distribution consists of a chart, table, graph, or formula that specifies all the values that the random variable can assume along with the probability of those values occurring.

Examples: we take a random sample of students and measure height or weight. A random sample of plants and measure seed set, fruit set, harvest yield, growth rate, survivorship from seed to harvest. Or the number of offspring in a family with a trait or condition of interest, such as disease, or disease resistance, or growth rate, or some trait we're interested in.

Example: flipping a coin once: the probability distribution is 0.5 on heads and 0.5 on tails. Rolling a die once: the probability distribution is 1/6 on each of the six possibilities. The complication here comes when you flip a coin more than once. Now how do we distribute probabilities over 1, 2, 3, 4, etc. heads?

Draw a graph of a probability distribution to visualize the distribution of probability across possible values of the variable.

There are two types of probability distribution the probability mass distribution for discrete variables, and the probability density distribution for continuous variables

## Probability distributions

Two kinds:

- Probability mass function for a discrete variable
- Probability density function for a continuous variable

## Probability distributions

Two fundamental properties:

- Location
- Spread

Draw pictures of probability distributions varying in location and spread.

## Measures of location, or “central tendency”

- Mode
- Median
- Mean, or “Expected value” of  $x$ ,  $E(x) = \mu$

### Mean of a distribution

**Definition 22.** The mean of a discrete probability distribution is

$$E(x) = \mu = \sum_i x_i P(x_i).$$

The mean of a continuous probability distribution is

$$E(x) = \mu = \int_{-\infty}^{+\infty} x f(x) dx.$$

In both cases, the mean is the “balance point” of the distribution.

### Measures of variability

- Variance =  $Var(x) = \sigma_x^2$
- Standard deviation =  $\sqrt{Var(x)} = \sigma_x$

### Variance of a distribution

**Definition 23.** The variance of a discrete probability distribution is

$$Var(x) = E(x - \mu)^2 = \sum_i (x_i - \mu)^2 P(x_i).$$

The variance of a continuous probability distribution is

$$Var(x) = E(x - \mu)^2 = \int_{-\infty}^{+\infty} (x - \mu)^2 f(x) dx.$$

### Properties of the mean

- $E(x + c) = E(x) + c$
- $E(ax) = aE(x)$
- $E(x + y) = E(x) + E(y)$ , even if  $x$  and  $y$  are dependent (correlated).
- $E(xy) \neq E(x)E(y)$
- $E(x^2) \neq [E(x)]^2$

### Properties of the variance

- $Var(ax) = a^2 Var(x)$
- $Var(x + y) = Var(x) + Var(y)$  only if  $x$  and  $y$  are independent (not correlated)

### Let's apply to the binomial distribution

### Axiomatic definition of probability

**Definition 24.**

$$P(A) = \frac{\text{number of outcomes giving event A}}{\text{total number of equally likely outcomes}}$$

$$P(A) = (\text{number of possible outcomes giving event A}) \\ \times (\text{probability of one outcome giving event A})$$

## Binomial distribution

**Definition 25.** Consider  $n$  independent trials of an experiment where at each trial there are exactly two possible outcomes: success with probability  $p$  that is constant from trial to trial and failure with probability  $1 - p$ . Let  $X$  equal the number of successes out of the  $n$  repeated trials. The random variable  $X$  is called a **binomial random variable**, and its probability distribution is called a **binomial distribution**.

## Mean and variance of the binomial distribution

**Theorem 26.** If  $x$  is a binomial random variable with  $n$  trials and probability of success  $p$ , then the mean of this variable is  $np$  and its variance is  $np(1 - p) = npq$ .

## Binomial distribution

**Theorem 27.** Let the random variable  $X$  denote the number of successes in  $n$  trials of a binomial experiment. Then  $X$  has a binomial distribution and the formula for the probability of observing exactly  $x$  successes in  $n$  trials is given by

$$P(X = x) = \binom{n}{x} p^x q^{n-x} \text{ for } x = 0, 1, 2, \dots, n, \text{ where}$$

$x$  = number of trials,  $x$  = number of successes in  $n$  trials,  $p$  = probability of success in a single trial,  $q = 1 - p$ , and  $\binom{n}{x} = {}_n C_x = \frac{n!}{x!(n-x)!}$

## Poisson distribution

The Poisson distribution can be thought of a special case of the binomial distribution for very large  $n$  and very small  $p$ .

In that case, the mean  $np$  and variance  $np(1 - p)$  are essentially equal to each other.

## Poisson distribution

The probability of exactly  $k$  events during some interval or across some spatial region, given that the mean number of events is  $\lambda$ , is equal to

$$P(k, \lambda) = \frac{\lambda^k e^{-\lambda}}{k!}$$

## Binomial distribution in R

- `dbinom(0:10, 10, 0.5)` gives the probability of each of zero to 10 successes in a sample of 10, where the probability of one success is 0.5
- `pbinom(0:10, 10, 0.5)` gives the cumulative probability of each of zero to zero to 10 successes in a sample of 10, where the probability of one success is 0.5.
- `plot(0:10, dbinom(0:10, 10, 0.5))` plots the above on a graph, which gives a visual depiction of the binomial probability distribution for those values.

## Poisson distribution in R

- `dpois(0:10, 2)` gives the probability of each of zero to 10 events across some spatial or temporal sampling area, where mean number of events in that area is 2
- `ppois(0:10, 2)` gives the cumulative probability of each of zero to 10 events across some spatial or temporal sampling area, where mean number of events in that area is 2.
- `plot(0:10, dpois(0:10, 2))` plots the above on a graph, which gives a visual depiction of the poisson probability distribution for those values.

## 7 Random variables: Review


### Lecture 6: Random variables: review

What is statistics?

**Definition 28.**     • Set of tools


- to convert *data*
- to *probability*.

Probability distributions



**Probability distributions**

In this probability distribution, the probability that  $x$  is greater than one equals one-half:



### Probability distributions

Two fundamental properties:

- Location
- Spread

### Measures of location, or “central tendency”

- Mode
- Median
- Mean, or “Expected value” of  $x$ ,  $E(x) = \mu$

### Mean of a distribution

**Definition 29.** The mean of a discrete probability distribution is

$$E(x) = \mu = \sum_i x_i P(x_i).$$


The mean of a continuous probability distribution is

$$E(x) = \mu = \int_{-\infty}^{+\infty} x f(x) dx.$$

In both cases, the mean is the “balance point” of the distribution.

### Mean of a distribution

The mean is the balancing point of the distribution, if we treat the area under the curve as physical mass.



### Measures of variability

- Variance =  $Var(x) = \sigma_x^2$
- Standard deviation =  $\sqrt{Var(x)} = \sigma_x$

### Variance of a distribution

**Definition 30.** The variance of a discrete probability distribution is

$$Var(x) = E(x - \mu)^2 = \sum_i (x_i - \mu)^2 P(x_i).$$

The variance of a continuous probability distribution is

$$Var(x) = E(x - \mu)^2 = \int_{-\infty}^{+\infty} (x - \mu)^2 f(x) dx.$$

### Properties of the mean

- $E(c) = c$ , where  $c$  is a constant.
- $E(x + c) = E(x) + c$ , where  $c$  is a constant.
- $E(cx) = cE(x)$ , where  $c$  is a constant.
- $E(x + y) = E(x) + E(y)$ , even if  $x$  and  $y$  are dependent (correlated).
- $E[E(x), E(y)] = E(x, y)$ ,
- $E(xy) \neq E(x)E(y)$
- $E(x^2) \neq [E(x)]^2$

## Binomial distribution

- We have  $n$  trials, and the probability of success on each trial is  $p$ .
- $p$  is a constant.
- Each trial is independent of every other trial (one trial does not affect another trial).
- Mean of the binomial distribution  $E(x) = np$ .
- Variance of the binomial distribution  $Var(x) = np(1 - p) = npq$ .

## Binomial distribution

Probability of  $r$  successes in  $n$  trials:

$$\begin{aligned} P(r, n) &= \text{number of ways } r \text{ successes} \times \text{probability of one way} \\ &= \text{number of ways to get } r \text{ successes} \times p^r q^{n-r} \\ &= \text{combinations of } n \text{ items taken } r \text{ at a time} \times p^r q^{n-r} \\ &= \frac{n!}{r!(n-r)!} \times p^r q^{n-r} \end{aligned}$$

## Statistical test

We want to know if some observation can be explained by chance variation, or if some additional cause is present.

1. We suppose that we have a known probability distribution.
2. We suppose that our observation was drawn from that probability distribution.
3. We ask: what is the probability that we would have obtained a value this extreme (or more) from this distribution?
4. If this probability is less than 0.05, then we conclude that this observation could not (likely) have occurred by chance; i.e. it is “unusual” or “unexpected.” Therefore, there must have been some cause outside the variation in the given probability distribution.
5. If this probability is greater than or equal to 0.05, then we conclude that this observation could have occurred by chance.

## Binomial distribution in R

- `dbinom(0:10, 10, 0.5)` gives the probability of each of zero to 10 successes in a sample of 10, where the probability of one success is 0.5
- `pbinom(0:10, 10, 0.5)` gives the cumulative probability of each of zero to zero to 10 successes in a sample of 10, where the probability of one success is 0.5.
- `plot(0:10, dbinom(0:10, 10, 0.5))` plots the above on a graph, which gives a visual depiction of the binomial probability distribution for those values.

Examples: number of lymphocytes in a sample of blood cells. Number of a particular type of offspring. Number of individuals dying over a given time interval (also might be a Poisson distribution), number of disease infected individuals in a population.

## Poisson distribution

- Poisson distribution is the binomial distribution with large  $n$  and small  $p$ .
- If  $n > 100$  and  $p < 0.01$ , and all the assumptions of the binomial are met, then we have the Poisson distribution.
- Then the mean is equal to the variance.
- $\sigma^2 = \text{Var}(x) = np(1 - p) \cong np = \text{mean}(x) = \mu$ .
- Poisson distribution is appropriate for events in space or time, where one event is independent of other events.
- Examples: genetic mutation, dispersion of individuals in space, morphology, weather events (e.g. storms), accidents, number of deaths in a region or across a time interval, number of parasites found on a host, radioactive decay.

## Poisson distribution in R

- `dpois(0:10, 2)` gives the probability of each of zero to 10 events across some spatial or temporal sampling area, where mean number of events in that area is 2
- `ppois(0:10, 2)` gives the cumulative probability of each of zero to 10 events across some spatial or temporal sampling area, where mean number of events in that area is 2.
- `plot(0:10, dpois(0:10, 2))` plots the above on a graph, which gives a visual depiction of the poisson probability distribution for those values.


Examples from the binomial distribution: can a person predict the outcome of a coin flip? Do this in class. Another example: lymphocytes in a sample of 100 blood cells. Probability of HIV positive is 0.008. What if 500 babies are screened, and 5 are found positive? Number of people getting cancer over a particular time interval. Let's say the probability that someone dies between age 20 and age 40 is 0.002. Let's say we have 100 students age 20 in a class, and they have their 20-year reunion, and 3 people died. Is that an unusual class?

Examples from the Poisson distribution: let's say there is an expected average of 1 new mutation per seed. We screen a seed for mutations, and we find 3 new mutations. Is this expected purely by chance? Number of buras per winter is 10. This winter there were 20. Is this expected by chance, or could some other factor be involved, maybe climate change? Let's say the average number of offspring of a particular bird is 3. We take a particular bird we're interested in and find that it produces 6 offspring. Could this have happened by chance, or is some other factor involved (perhaps an unusual diet)?

## Normal distribution

Fundamental rule: approximately 95% of all observations lie within 2 standard deviations of the mean!





## 8 Normal distribution

### Lecture 7: Normal distribution and the Central Limit Theorem


Reading: pp. 11 - 44 in Vasilj.

#### What is statistics?

**Definition 31.** • Set of tools

- to convert *data*
- to *probability*.

#### Probability distributions




```
plot(0:5, dbinom(0:5, 5, .5), type = "h", xlab = "Successes", ylab = "Probability", main = "Binomial, n = 5, p = 0.5")
```

### Binomial distribution

- We have  $n$  trials, and the probability of success on each trial is  $p$ .
- $p$  is a constant.
- Each trial is independent of every other trial (one trial does not affect another trial).
- Mean of the binomial distribution  $E(x) = np$ .
- Variance of the binomial distribution  $Var(x) = np(1 - p) = npq$ .

### Binomial distribution



$$\mu = E(x) = np = 5 \times 0.5 = 2.5.$$

$$\sigma^2 = Var(x) = E(x - \mu)^2 = np(1 - p) = 5 \times 0.5 \times 0.5 = 1.25.$$

$$\sigma = \sqrt{\sigma^2} = \sqrt{1.25} = 1.12.$$

### Binomial distribution: Example

- Given: probability that a blood cell is a lymphocyte is 0.22 in a healthy human population.
- In one patient's blood test, 100 cells were counted and 32 were lymphocytes.
- Is this normal?
- Need to calculate  $P(X = 32)$  or more from a binomial distribution with  $n = 100$  and  $p = 0.22$ .
- $P(X = r) = \frac{n!}{r!(n-r)!} \times p^r q^{n-r}$

### Statistical test

We want to know if some observation can be explained by chance variation, or if some additional cause is present.

1. We suppose that we have a known probability distribution.
2. We suppose that our observation was drawn from that probability distribution.

3. We ask: what is the probability that we would have obtained a value this extreme (or more) from this distribution?
4. If this probability is less than 0.05, then we conclude that this observation could not (likely) have occurred by chance; i.e. it is “unusual” or “unexpected.” Therefore, there must have been some cause outside the variation in the given probability distribution.
5. If this probability is greater than or equal to 0.05, then we conclude that this observation could have occurred by chance.

### Binomial distribution: Example 2

- Given: probability that a coin lands heads is 0.5.
- I flip a coin 10 times, and a student guesses correctly on 4 flips.
- Can the student predict a coin flip?
- Need to calculate  $P(X = 4)$  or better from a binomial distribution with  $n = 10$  and  $p = 0.5$ .

### Poisson distribution


- Poisson distribution is the binomial distribution with large  $n$  and small  $p$ .
- If  $n > 100$  and  $p < 0.01$ , and all the assumptions of the binomial are met, then we have the Poisson distribution.
- Then the mean is equal to the variance.
- $\sigma^2 = Var(x) = np(1 - p) \cong np = mean(x) = \mu$ .
- Poisson distribution is appropriate for events in space or time, where one event is independent of other events.
- Examples: genetic mutation, dispersion of individuals in space, morphology, weather events (e.g. storms), accidents, number of deaths in a region or across a time interval, number of parasites found on a host, radioactive decay.

### Poisson distribution: Example 1

- Given: in a normal plant, the mean number of mutations per genome each generation is 1.0.
- I expose a plant to a chemical, then measure 3.0 mutations in one seed.
- Is the chemical a mutagen?
- Need to calculate  $P(X = 3)$  or more from a Poisson distribution with  $\lambda = 1.0$ .
- $P(X = k) = \frac{\lambda^k e^{-\lambda}}{k!}$ .


### Normal distribution

- The normal distribution is bell-shaped.
- The normal distribution is symmetrical.
- Approximately 95% of all observations lie within 2 standard deviations of the mean.




### Normal distribution: Example 1

- The mean height of an adult Arabidopsis plant is 15 cm, and the standard deviation of height is 2 cm.
- I randomly sample a seed from a mutant variety of Arabidopsis. I grow it to adult, and its height is 10 cm.
- Does the mutant gene cause stunted growth?
- Need to calculate  $P(X = 10)$  or lower, assuming a normal distribution.




## Normal distribution

- The normal distribution is bell-shaped.
- The normal distribution is symmetrical.
- Approximately 95% of all observations lie within 2 standard deviations of the mean.



## Normal distribution

$$f(x) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{(x-\mu)^2}{2\sigma^2}}$$



Where did this equation come from?

## Central Limit Theorem

$$f(x) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{(x-\mu)^2}{2\sigma^2}}$$

**Theorem 32.** The sum of  $n$  identically-distributed variables call  $S_n$ . Then  $S_n$  approaches a normal distribution as  $n$  approaches infinity.

## Central Limit Theorem

The central limit theorem can be demonstrated in a Galton board, or a Quincunx machine:

[http://www.youtube.com/watch?v=xDIyA0Ba\\_yU](http://www.youtube.com/watch?v=xDIyA0Ba_yU)

<http://www.mathsisfun.com/data/quincunx.html>

## Normal distribution

The binomial distribution becomes the normal distribution if  $np > 5$ . Compare the two when  $n = 16$  and  $p = 0.5$ :




Fig. 7-2 Plot of binomial for  $N = 16$  and  $p = 0.5$  and normal curve with mean = 8 and standard deviation = 2.

Similarly, the Poisson distribution becomes a normal distribution when  $\lambda > 5$ .

## 9 Central Limit Theorem continued

### Lecture 8: Central limit theorem, normal distribution continued

Announcements:

- Reading for today's subject: pp. 48 - 52 in Vasilj.
- First midterm exam is Tuesday, November 9, at 4PM, here. Will cover all material in the course through today.
- Answer keys to the first three problem sets are uploaded to the course webpage, [www.stewartschultz.com/statistics/course/](http://www.stewartschultz.com/statistics/course/).
- Pick up your corrected problem sets after class today from me. Study for the test by checking your answers against the key. See me or Tomislav if you have any questions.

## Central Limit Theorem

The central limit theorem can be demonstrated in a Galton board, or a Quincunx machine:

[http://www.youtube.com/watch?v=xDIyA0Ba\\_yU](http://www.youtube.com/watch?v=xDIyA0Ba_yU)

<http://www.mathsisfun.com/data/quincunx.html>

Any repeated random process eventually generates the normal distribution. In this case the repeated random process is a ball bouncing against pegs, left to right with probability 0.5. In other words a binomial variable with  $n$  equal to the number of rows on the board, and  $p$  equal to any constant.

### Central Limit Theorem

$$f(x) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{(x-\mu)^2}{2\sigma^2}}$$

**Theorem 33.** *The sum of  $n$  identically-distributed variables call  $S_n$ . Then  $S_n$  approaches a normal distribution as  $n$  approaches infinity.*

Question: What if the variables are not identically-distributed? See the Quincunx machine to see what happens.

Answer: we can calculate the sample mean!

The sample mean is calculated  $\bar{x} = \frac{\sum_i^n x_i}{n}$ .

Question: what is the shape of the distribution of the sample mean?

Answer: We know that the sample mean should approach a normal distribution, by the Central Limit Theorem.

### Central Limit Theorem

$$f(x) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{(x-\mu)^2}{2\sigma^2}}$$

**Theorem 34.** *The sum of  $n$  identically-distributed variables call  $S_n$ . Then  $S_n$  approaches a normal distribution as  $n$  approaches infinity.*

*Corollary: the sample mean of these  $n$  identically-distributed variables is normally distributed.*

*Question: what is the variance of the sample mean?*

*Corollary: the variance of the sample mean of the  $n$  variables is their variance divided by  $n$ . The standard deviation of the sample mean of the  $n$  variables is the standard deviation divided by  $\sqrt{n}$ .*

Note that the central limit theorem requires that all the variables that we sum have the same "identical" distribution, identical mean and variance. If they don't, then the sum is no longer normally distributed.

This can be demonstrated on the Quincunx machine on the mathisfun website, by changing the  $p$ -value. Let it run with  $p=0.5$  for awhile, then switch to  $p=.1$  or something, and see what happens. You get a bimodal distribution.

Are we then out of luck? Can we convert this to a normal distribution somehow? Yes we can, by taking a random sample from it. Every time we take a sample from it, we have a sample from the same "identical" distribution. The distribution does not change between samples. So if we take a random sample from that distribution, then get the sum of that sample, the sum is normally distributed, by the central limit theorem.

But, if the sum is normally distributed, then the sample mean also is. The sample mean is just  $\bar{x} = \frac{\sum_i^n x_i}{n}$ , where  $n$  is the sample size. Notice that the numerator is the sum of many identically distributed variables (we know they are all identically distributed because we have sampled randomly from the same distribution), so the numerator is likely to be normally distributed if  $n$  is large enough. Dividing by  $n$  is just dividing by a constant, so that has no effect on the distribution, so the sum divided by  $n$  is also normally distributed. So the sample mean is normally distributed.

This is a very useful result. By the central limit theorem, we know that the sample mean will be approximately normally distributed, and it will be closer to the normal distribution as  $n$ , the sample size, increases. Generally, if we are sampling from a distribution that is already close to a normal distribution, such as height or mass of individuals, then we need exceed only a small  $n$  in order for the mean to be normally distributed. Usually  $n > 20$  or so is enough sample size to assume the mean is normal.

Draw a normal distribution, show its mean, now imagine that we take random samples of  $n = 100$  or some other number from this population, and for each sample we calculate the mean and mark it on the graph. Will our estimates of the mean be as spread out as the original distribution? No, their

distribution will be narrower. In other words, the distribution of the sample mean has a lower variance than the distribution of the population you are sampling from. How much lower? We can calculate this from the properties of the variance that we demonstrated last week.

### Properties of the variance

- $Var(ax) = a^2 Var(x)$
- $Var(x + y) = Var(x) + Var(y)$  only if  $x$  and  $y$  are independent (not correlated)

### Central Limit Theorem

#### Theorem 35. [...]

*Corollary: the variance of the sample mean of the  $n$  variables is their variance divided by  $n$ . The standard deviation of the sample mean of the  $n$  variables is the standard deviation divided by  $\sqrt{n}$ .*

### Central Limit Theorem

*Proof.* We know that  $Var(x + y) = Var(x) + Var(y)$ , if  $x$  and  $y$  are independent. Also,  $Var(ax) = a^2 Var(x)$ . Then  $Var(x_1 + x_2) = Var(x_1) + Var(x_2) = 2Var(x)$ . Then  $Var(x_1 + x_2 + x_3 + \dots + x_n) = nVar(x)$ . Or, equivalently,  $Var(\sum_i^n x_i) = nVar(x)$ . But we want to know  $Var(\sum_i^n x_i/n)$ . This is just the sum divided by the constant. Therefore the variance is the sum of the variances divided by the constant squared,  $Var(\bar{x}) = Var(\sum_i^n x_i/n) = nVar(x)/n^2 = Var(x)/n$ . Then the standard deviation is the square root of the variance:  $SD(\bar{x}) = \sqrt{Var(\bar{x})} = \sqrt{Var(x)/n} = SD(x)/\sqrt{n} \equiv \text{S.E.M.}$   $\square$

### Central Limit Theorem

$$f(x) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{(x-\mu)^2}{2\sigma^2}}$$

**Theorem 36.** *The sum of  $n$  identically-distributed variables call  $S_n$ . Then  $S_n$  approaches a normal distribution as  $n$  approaches infinity.*

*Corollary: the sample mean of these  $n$  identically-distributed variables is normally distributed.*

*Corollary: the variance of the sample mean of the  $n$  variables is their variance divided by  $n$ . The standard deviation of the sample mean of the  $n$  variables is the standard deviation divided by  $\sqrt{n}$ .*

### Central Limit Theorem

General procedure if we know the true standard deviation:

- Given we know the true mean  $\mu$  and true standard deviation  $\sigma$  from the distribution.
- We have an observation  $x_i$ .
- We calculate the  $Z$  score of our observation:

$$Z = \frac{x_i - \mu}{\sigma}.$$

- This is the distance from the mean to the observation in units of standard deviation.
- If this distance is in the tail of the distribution (e.g. if  $|Z| > 1.96$ ) then we conclude that the observation is unusual.
- Or, we can calculate the probability that the observation is at least as extreme as observed.



Example. Let's imagine that a hotel owner has a hotel with 9 rooms. He has a boiler that can give 90 minutes of hot water for the 9 showers. He wants to know what is the probability that if he has all 9 guests showering at the same time, they will use up all the 90 minutes of water? Well that would require that the average shower length for those 9 guests is 10 minutes. So, what's the probability that the average (sample mean) shower length for 9 random guests is 10 minutes? To calculate this we need to know the mean and standard deviation. Let's say he knows (he has been keeping track over several months) that the mean shower length of people in his hotel is 8 minutes, and the standard deviation is 3 minutes. What is the probability that the mean shower length of 9 people would be higher than 10 minutes? The standard deviation of the mean, or S.E.M. is 3 divided by the square root of 9, or  $3/3 = 1$ . The mean is 8 minutes. So 10 minutes is 2 standard deviations higher than 8, and 2 standard deviations is 95% of the normal distribution. So our Z-score is 2.0, and the probability that 9 people all showering at the same time would use up all the water is approximately 0.025, or 2.5%.

NOTE our assumptions here. We assume the sample mean is a normal distribution. Can we assume this? The reason we assume it is that we have a sample of 9. Is 9 large enough of a sample to assume normality? It might be, if the variable of shower time is close to normal, as you can see from the R TeachingDemo `clt.examp()`. Just adding together 5 or so uniform random variates is enough to give a distribution close to normal. But in practice, we might be a little concerned that the distribution is normal if we have a sample size only of 9. We also assume that we already know the mean and standard deviation of shower time of the hotel guests, and that the 9 people are a random sample from all potential guests of this hotel. If they are not a random sample, then we can't make any conclusions. For example, the mean and standard deviation may have changed recently because people changed their showering habits, etc., so now we have a sample from a different distribution than that assumed.

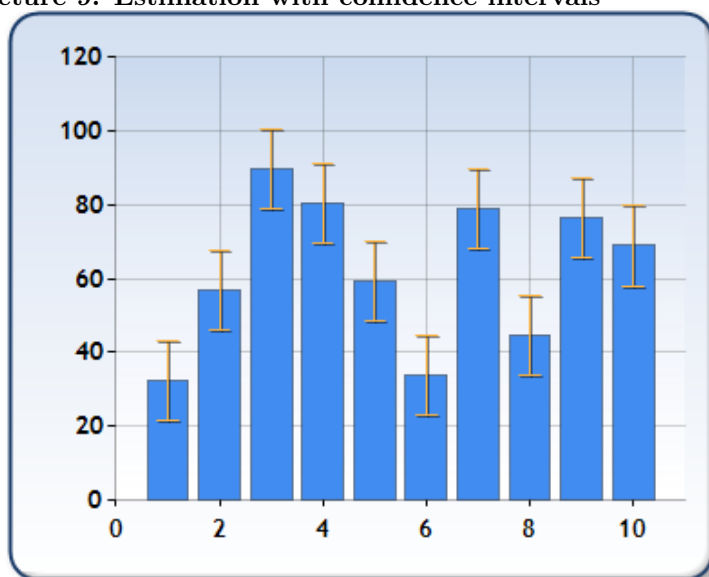
## 10 Estimation

### Lecture 9: Estimation with confidence intervals

Announcements:

- Reading for today's subject: pp. 52 - 54 in Vasilj.
- No lab exercises today.
- However, Problem Set 4 is due today, please turn in at the end of the lecture.
- First midterm exam is today, at 16:00, here.
- Today: lecture on Z scores, estimation, and confidence intervals (14:00 to 16:00).

### Lecture 9: Estimation with confidence intervals



## Z scale


General procedure if we know the true standard deviation:

- Given we know the true mean  $\mu$  and true standard deviation  $\sigma$  from the distribution.
- We have an observation  $x_i$ .
- We calculate the  $Z$  score of our observation:

$$Z_i = \frac{x_i - \mu}{\sigma}.$$

- This is the distance from the mean to the observation in units of standard deviation.
- If this distance is in the tail of the distribution (e.g. if  $|Z_i| > 1.96$ ) then we conclude that the observation is unusual.
- $Z_{0.025} = 1.96$ .

## Z distribution




$$f(x) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{(x-\mu)^2}{2\sigma^2}}$$

$$f(x) = \frac{1}{\sqrt{2\pi}} e^{-\frac{x^2}{2}}$$

## Z distribution in R: pnorm()

For any  $Z_i$  score, the probability that  $Z$  is less than or equal to  $Z_i$ ,


$$P(Z \leq Z_i) = \text{pnorm}(Z_i).$$



You should know (memorize) that the Z variable is normally distributed with mean 0 and standard deviation 1. The Z distribution is called the “standard normal” distribution, i.e. the distribution with mean 0 and standard deviation 1. All normal distributions can be converted to the Z distribution; in fact it makes them easier to work with. It’s just changing the scale.

Example: assume the mean height of men is 175 cm, and the standard deviation is 8 cm, and height is normally distributed. Then my height of 183 cm has a Z score of 1, it is one standard deviation higher than the mean. In R, the probability that a random man is shorter than I am is `pnorm(1)`. The probability that a random man is taller is `1-pnorm(1)`. `pnorm()` assumes that the mean is zero and the standard deviation is 1 unless you tell it otherwise. So, for example, `pnorm(1)` is the same as `pnorm(183, mean=175, sd = 8)` since 183 has a Z score of 1 if the mean is 175 and the sd is 8. So in general, for any Z score, `pnorm(Z)` is the probability of Z or smaller.

## Z distribution



## Confidence interval

**Definition 37.** A confidence interval is any interval calculated around an estimator that has a particular probability of containing the true value of the corresponding parameter of the population. The parameter of interest is usually the population mean.

### Calculating the 95% confidence interval

$$P(-1.96 \leq Z \leq 1.96) = 0.95.$$

$$P\left(-1.96 \leq \frac{\bar{x} - \mu}{\frac{\sigma}{\sqrt{n}}} \leq 1.96\right) = 0.95.$$

$$P\left(-1.96 \frac{\sigma}{\sqrt{n}} \leq \bar{x} - \mu \leq 1.96 \frac{\sigma}{\sqrt{n}}\right) = 0.95.$$

$$P\left(-\bar{x} - 1.96 \frac{\sigma}{\sqrt{n}} \leq -\mu \leq -\bar{x} + 1.96 \frac{\sigma}{\sqrt{n}}\right) = 0.95.$$

$$P\left(\bar{x} - 1.96 \frac{\sigma}{\sqrt{n}} \leq \mu \leq \bar{x} + 1.96 \frac{\sigma}{\sqrt{n}}\right) = 0.95.$$

### Calculating the 95% confidence interval

$$P\left(\bar{x} - 1.96 \frac{\sigma}{\sqrt{n}} \leq \mu \leq \bar{x} + 1.96 \frac{\sigma}{\sqrt{n}}\right) = 0.95.$$

$$\bar{x} \pm 1.96 \frac{\sigma}{\sqrt{n}}$$

This is the 95% confidence interval for the sample mean.

We are 95% confident that the true mean is contained in this interval.

$$Z_{0.025} = 1.96.$$

$$\bar{x} \pm Z_{0.025} \frac{\sigma}{\sqrt{n}}.$$

For example, we want to estimate the mean height of Croatian men. We take a random sample of 5 men, measure their height. Let's say the mean is 169. What is the 95% confidence interval? It is (162, 176). Let's say we take another sample of 20 men, and find that the sample mean is 176. Then the 95% confidence interval is (172.5, 179.5). Let's say we take another sample of 100 men, and find the sample mean is 174. Then the 95% confidence interval is (172.4, 175.6). Let's say we take a random sample of 1000 men, and get a mean of 175.3. Now the 95% confidence interval is (174.8, 175.8). You should do these calculations yourself to check your knowledge. As you can see, as the sample size increases, the width of the confidence interval gets smaller. In other words, our estimate becomes closer and closer to the true mean, and our 95% confidence of the true mean extends over a smaller and smaller segment of possible values.

We don't have to calculate the 95% confidence interval, we could also calculate the 90% confidence interval, or the 99% confidence interval, or some other interval. It depends on how certain we would like to be about the true value of the population mean. The 99% confidence interval is wider than the 95% and 90% confidence intervals.


This procedure can be followed for any sample mean from any distribution, as long as  $n$ , the sample size is large enough that we can assume the normal distribution. The calculation can be done for the binomial distribution and for the Poisson distribution. The mean of a binomial variable (the number of successes out of  $n$  trials, with probability of success  $p$ ), is just the observed proportion of successes,  $n_{\text{successes}}/n$ .

### Lecture 10: Estimation with confidence intervals, continued

Announcements:

- Reading for this week's subject: pp. 54 - 79 in Vasilj.
- However, Problem Set 4 is past due, and Problem set 5 is due today, please turn in at the end of the lecture.
- Exams will be returned today.

## Lecture 10: Estimation with confidence intervals



### Population versus sample

**Definition 38.** A **population** is the set of all persons or things of interest. A **sample** is a subset of the population. Convention: Greek letters are used for population parameters ( $\mu$  for mean,  $\sigma$  for standard deviation). Latin letters are used for the sample ( $\bar{x}$  for sample mean,  $s$  for sample standard deviation). We say  $\bar{x}$  is an **estimate** of  $\mu$ , and  $s$  is an estimate of  $\sigma$ .

Show a diagram of population and sample, with corresponding probability distributions. Show the definition of mean of the probability distribution, and the definition of mean of the sample.

We can illustrate the effects of sample size on the distribution of the mean with R, using the `dnorm()` function. Say for heights of men, the distribution will have a mean of 175 cm and standard deviation of 8 cm. To make a plot, we type `xv = seq(140, 215, .1)` to get the x values to plot, and then plot the distribution with `plot(xv, dnorm(xv, mean=175, sd = 8), type="l")`. We can show the effect of sample size with `points(xv, dnorm(xv, mean=175, sd = 8/sqrt(n)), where  $n$  is the sample size.`

For example: IQ scores for adults are normally distributed, with mean  $\mu = 100$  and standard deviation  $\sigma = 15$ . If a random sample of 16 adults is selected, find the probability that the average IQ score in the sample will be at least 92.5. Here the standard deviation of the mean is  $15/4 = 3.75$ , and  $92.5$  is  $(-7.5)/3.75 = -2$  on the  $Z$  scale. So the probability that the mean is at least  $-2$  on the  $Z$  scale is  $1 - \text{pnorm}(-2) = 0.977$ .

We can use the standard error of the mean (SEM) to construct the margin of error around the estimate of the true mean.

Draw another picture of the normal distribution, and show that 95% of the distribution is between plus or minus 1.96 standard deviations.

As I showed last time, the 95% confidence interval is easily calculated:

### Calculating the 95% confidence interval

$$P(-1.96 \leq Z \leq 1.96) = 0.95.$$

$$P\left(-1.96 \leq \frac{\bar{x} - \mu}{\frac{\sigma}{\sqrt{n}}} \leq 1.96\right) = 0.95.$$

$$\text{Therefore, } P\left(\bar{x} - 1.96 \frac{\sigma}{\sqrt{n}} \leq \mu \leq \bar{x} + 1.96 \frac{\sigma}{\sqrt{n}}\right) = 0.95.$$

$$\bar{x} \pm 1.96 \frac{\sigma}{\sqrt{n}}$$

This is the 95% confidence interval for the sample mean.

$$Z_{0.025} = 1.96.$$

$$\bar{x} \pm Z_{0.025} \frac{\sigma}{\sqrt{n}}.$$

**Calculating the 95% confidence interval of the sample mean  $\bar{x}$**

$$\bar{x} \pm Z_{\alpha/2} \frac{\sigma}{\sqrt{n}} = (\bar{x} - E, \bar{x} + E)$$

where

$1 - \alpha$  = confidence coefficient  $(1 - \alpha)100\%$  = level of confidence  $\alpha$  = the error probability  $\bar{x}$  = the mean of the sample  $\sigma$  = the standard deviation of the population  $n$  = the sample size  $Z_{\alpha/2}$  = the critical value of  $Z$ , or the  $Z$  value that cuts off a right-tail area of  $\alpha/2$  under the standard normal curve; that is,  $1 - P(Z < Z_{\alpha/2}) = \alpha/2$   $Z_{\alpha/2} \frac{\sigma}{\sqrt{n}} = E$ , the margin of error.

**Calculating the 95% confidence interval of the sample mean  $\bar{x}$**

$$\bar{x} \pm Z_{\alpha/2} \frac{\sigma}{\sqrt{n}} = (\bar{x} - E, \bar{x} + E)$$

**Table 10.1** Commonly Used Critical Values  $Z_{\alpha/2}$

Level of Confidence $(1 - \alpha)100\%$	Critical Value $Z_{\alpha/2}$
90%	1.645
95%	1.96
99%	2.576

As an example, let's say there is a new diet for people who want to lose weight. We sample 100 people who use the diet for two months, and we keep track of their weight gain during that time. The sample mean weight gain is -1 kg, and the standard deviation weight gain is 7 kg. Does the diet work? The standard error of the mean is 7/10 kg. Thus the mean, -1, is less than 2 standard errors from zero, so it is not unexpected assuming the true mean is zero. So we cannot conclude that the diet works.

**The sample standard deviation**

**Definition 39.** The sample variance is written  $s^2$  and is defined as:

$$s^2 = \frac{\sum_1^n (x_i - \bar{x})^2}{n - 1}.$$

$n - 1$  is called the "degrees of freedom" in this calculation.

## 11 Z tests and t tests

### Lecture 11: Z tests and t tests

Announcements:

- Reading for this week's subject: pp. 54 - 79 in Vasilj.

- Problem Set 4 is past due, Problem set 5 is due, please turn in at the end of the lecture.
- Exams will be returned today.
- If you have any questions about the exams, please see me or Tomislav.
- Next exam will include all questions in this exam that less than 50% of people answered correctly. These are questions 9, 24, 25, 21, 16, and 5. So make sure you understand those questions!

### IMPORTANT ANNOUNCEMENT

Your final grade in this course will be based 75% on exams (midterm and/or final), and 25% on problem sets (and class attendance for regular students, participation, and behavior). **EACH WEEKLY PROBLEM SET MUST BE RETURNED THE FOLLOWING TUESDAY IN CLASS!** If a problem set is late, you receive no credit (because answers will be posted that day).

Also, Problem sets 4 through 6 are all due IN CLASS ON TUESDAY, NOVEMBER 23.

### Population versus sample

**Definition 40.** A **population** is the set of all persons or things of interest. A **sample** is a subset of the population. Convention: Greek letters are used for population parameters ( $\mu$  for mean,  $\sigma$  for standard deviation. Latin letters are used for the sample ( $\bar{x}$  for sample mean,  $s$  for sample standard deviation). We say  $\bar{x}$  is an **estimate** of  $\mu$ , and  $s$  is an estimate of  $\sigma$ .

Show a drawing: we have several possible populations, and a sample. Which population did the sample come from? To answer that question we ask "is the sample unusual relative to the reference population?"

### Z test in general

$$Z = \frac{\text{observed statistic} - \text{expected}}{\text{true standard deviation of observed statistic}}$$

If observed statistic is the mean of a sample of  $n$ , then its standard deviation is  $\sigma/\sqrt{n}$ .

If observed statistic is a binomial variate, its standard deviation is  $\sqrt{np(1-p)}$ .

If observed statistic is a Poisson variate, its standard deviation is  $\sqrt{\lambda}$ .

Show a drawing of the normal distribution with tails, on the Z scale. This represents the reference population. We can plot our sample on this reference population. if the sample is in one of the tails, then we conclude the sample is unusual and it was not drawn from the reference population.

### Z test in general

$$Z = \frac{\text{observed statistic} - \text{expected}}{\text{true standard deviation of observed statistic}}$$

If observed statistic is the mean of a sample of  $n$ , then its standard deviation is  $\sigma/\sqrt{n}$ .

If observed statistic is a binomial variate, its standard deviation is  $\sqrt{np(1-p)}$ .

If observed statistic is a Poisson variate, its standard deviation is  $\sqrt{\lambda}$ .

First we have to convert to the Z scale. To do so, we need to know the standard deviation, the denominator of the Z expression. If the observation is the mean of the sample, the standard deviation is the SEM, or the standard deviation of the distribution divided by the square root of the sample size. If the sample size is one, then we have just one observation, and the standard deviation in the denominator is just the standard deviation of the distribution.

Take some examples from Problem sets 5 and 6. The orchard question, question 2 on PS 5. Here we have the sample mean, so we use the SEM in the denominator.

Then Mendel's pea experiment. Here we have the binomial distribution, so we can use the binomial variance to calculate the Z. The denominator is just the square root of the binomial variance. We can assume a normal distribution because  $np > 5$ .

Then try a Poisson example, with the Poisson we also know what the true standard deviation is, if we know the true mean. Let's say we know the true mutation rate is 1 mutation per seed. We screen 100

seeds and find a total of 118 mutations. Is that unusual if the normal mutation rate is 1 per seed? The Z value would be  $(212-100)/(\sqrt{100}) = 1.8$ , which is extremely unusual. We can assume normality because we added up 100 seeds, and the sum should be close to normal according to the Central Limit Theorem.

### What is the variance of the difference between two sample means?

Properties of the variance:

- $Var(ax) = a^2 Var(x)$
- $Var(x + y) = Var(x) + Var(y)$  only if  $x$  and  $y$  are independent (not correlated)
- $Var(\bar{x}_1 - \bar{x}_2) = Var(\bar{x}_1 + -\bar{x}_2) = Var(\bar{x}_1) + Var(-\bar{x}_2)$
- $Var(\bar{x}_1) + (-1)^2 Var(\bar{x}_2) = Var(\bar{x}_1) + Var(\bar{x}_2)$ , only if  $\bar{x}_1$  and  $\bar{x}_2$  are independent.
- If  $\bar{x}_1$  and  $\bar{x}_2$  are sample means, then their variances are respectively  $Var(\bar{x}_1) = \sigma_{x1}^2/n_{x1}$  and  $Var(\bar{x}_2) = \sigma_{x2}^2/n_{x2}$ .
- So  $Var(\bar{x}_1 - \bar{x}_2) = \sigma_{x1}^2/n_{x1} + \sigma_{x2}^2/n_{x2}$ .

We can also use a Z test to test for a difference between two samples. If we have two samples: were they drawn from the same distribution with true mean and standard deviation, or were they drawn from different distributions? Our observed statistic is the difference between the two sample means. Our expected value of the difference is (usually) zero. And in the denominator, the standard deviation of the difference is the square root of the sum of the variances of the two sample means.

### Z test in general

$$Z = \frac{\text{observed statistic} - \text{expected}}{\text{true standard deviation of observed statistic}}$$

If observed statistic is the difference between two sample means,  $\bar{x}_1 - \bar{x}_2$ , then the expected is usually zero, and

the standard deviation of the difference is

$$\text{standard deviation of difference} = \sqrt{\sigma_{x1}^2/n_{x1} + \sigma_{x2}^2/n_{x2}}$$

So

$$Z = \frac{\bar{x}_1 - \bar{x}_2}{\sqrt{\sigma_{x1}^2/n_{x1} + \sigma_{x2}^2/n_{x2}}}$$

Up until now we have assumed we know the true value of the standard deviation. We needed to know that because we needed to assume it was a constant, so we could assume that Z was normal. If you divide a normal variate with a constant, then the resulting variate is also normal. But if we don't know the true value of the standard deviation, we can calculate it from the sample. But then, it is no longer a constant, it is a variable. Every time we take a sample, the calculated standard deviation is slightly different. In that case, the distribution of the Z will no longer be normal. It will be more spread out than normal. Because it is no longer normal, it is no longer called a Z variate or Z score or Z distribution. Rather it is called a t distribution. The shape of a t distribution was derived by William Sealey Gossett (1876 - 1937) a British statistician who worked for the Guinness brewery. He called himself "Student", his pen-name, and the t distribution is also called "Student's t distribution".

### Z versus t


$$Z = \frac{\bar{x} - \text{expected}}{\frac{\sigma}{\sqrt{n}}}$$

$$t = \frac{\bar{x} - \text{expected}}{\frac{s}{\sqrt{n}}}$$



$$s = \text{sample standard deviation} = \sqrt{\frac{\sum_1^n (x - \bar{x})^2}{n - 1}}$$

$Z$  versus  $t$



## Lecture 12: $Z$ tests and $t$ tests, continued

Announcements:

- Reading for this week's subject: pp. 347 - 353 in Horvat.
- Your final grade in this course will be based 75% on exams (midterm and/or final), and 25% on problem sets (and class attendance for regular students, participation, and behavior).
- EACH WEEKLY PROBLEM SET MUST BE RETURNED THE FOLLOWING TUESDAY IN CLASS! If a problem set is late, you receive no credit (because answers will be posted that day).
- Also, Problem sets 4 through 6 are all due IN CLASS ON TUESDAY, NOVEMBER 23.

### More announcements

- If you have any questions about the exams, please see me or Tomislav.
- Next exam will include all questions in this exam that less than 50% of people answered correctly. These are questions 9, 24, 25, 21, 16, and 5. So make sure you understand those questions!

### $Z$ test in general

$$Z = \frac{\text{observed statistic} - \text{expected}}{\text{true standard deviation of observed statistic}}$$

If observed statistic is the mean of a sample of  $n$ , then its standard deviation is  $\sigma/\sqrt{n}$ .

If observed statistic is a binomial variate, its standard deviation is  $\sqrt{np(1-p)}$ .

If observed statistic is a Poisson variate, its standard deviation is  $\sqrt{\lambda}$ .

If observed statistic is the difference between two sample means, its standard deviation is  $\sqrt{\sigma_1^2/n_1 + \sigma_2^2/n_2}$ .

**Example: poll results Z test**

Let's say a poll of 1000 random Croatians found that 47.5% Were in favor of entering the EU. Is this unusual if the true proportion was 51.0% in favor?

The  $Z$  value for this calculation is  $-35/15.80823 = -2.214037$ . The associated probability is 0.01341312, so this is an unusual result, and we conclude that Croatians as a whole are not in favor of entering the EU according to this poll.

**Chi square test**

$$\chi^2 = \frac{\sum_1^n (\text{count of observed} - \text{count of expected})^2}{\text{count of expected}}$$

Degrees of freedom is the number of counts that are free to vary in this calculation.

**Example: poll results chi square test**

Let's say a poll of 1000 random Croatians found that 47.5% Were in favor of entering the EU. Is this unusual if the true proportion was 51.0% in favor?

We can redo this problem as a chi-square test, which is mathematically equivalent to a  $Z$  test. The chi square statistic is  $\chi^2 = 35^2/510 + 35^2/490 = 4.901961 = -2.214037^2 = Z^2$ . So a question involving a binomial variable can be answered either by a  $Z$  test or a chi-square test, assuming that the normality assumption of the  $Z$  test is met, namely that the binomial variate is normally distributed, which is generally assumed to be true if the mean of the binomial variable is greater than 5 ( $np > 5$ ).

If you have a choice between a  $Z$  test and a chi-square test, which should you choose? Because they are mathematically identical, it makes little difference. However, with the  $Z$  test it is possible for the  $Z$  value to be positive or negative. It is not possible for the chi-square value to be negative. So if you use a chi-square test, you need to keep in mind the direction of the difference. In the example above 475 Croatians were in favor and 525 were opposed to EU entry. If these numbers were reversed, they would have given the same chi-square statistic. While the chi-square test is insensitive to the direction of any result, the  $Z$  test is sensitive.

**Z test versus t test**

What if we don't know  $\sigma$ ? We calculate it from the sample as the sample standard deviation:

$$s = \sqrt{\frac{\sum_1^n (x_i - \bar{x})^2}{n - 1}}$$

$n - 1$  is called  $v$ , or the "degrees of freedom" used in the calculation of  $s$ .

**Z test versus t test**

$$Z = \frac{\text{observed} - \text{expected}}{\text{true standard deviation of observed}}$$

$$Z = \frac{\text{observed} - \text{expected}}{\sigma_{\text{obs}}}.$$

$$t = \frac{\text{observed} - \text{expected}}{\text{estimated standard deviation of observed}}$$

$$t = \frac{\text{observed} - \text{expected}}{s_{\text{obs}}}.$$


**Z test versus t test**

$Z$  is normally distributed, with mean of 0 and standard deviation of 1.

$t$  is very close to normally distributed, symmetrical, with mean of 0 and standard deviation slightly higher than 1.

But  $t$  is practically indistinguishable from a normal distribution if  $n > 30$ .

$t$  distribution for various sample sizes (or degrees of freedom)



**Example: a sample of heights**

Say we have a random sample of 15 women, their sample mean height is  $\bar{x} = 173$  cm, and sample standard deviation of height is  $s = 9.2$  cm. Is this unusual if the population mean is  $\mu = 169$  cm?

**Example: exam scores**

In our Exam 1, the mean score of women was 72.7 and the mean score for men was 70.3. The standard deviation for women was 7.7 and for men was 8.3. There were 14 women and 13 men. Is the sex difference in means unusual, if we expect the difference to be zero?

## 12 Hypothesis testing

### Lecture 13: Hypothesis testing

Announcements:

- Reading for this week's subject: pp. 61 - 79 in Vasilj.
- Problem sets 4 through 6 are due today in class.
- If you have any questions about the exams, please see me or Tomislav.
- Next exam will include all questions in this exam that less than 50% of people answered correctly. These are questions 9, 24, 25, 21, 16, and 5. So make sure you understand those questions!

### Lecture 13: Hypothesis testing

- We have a question about a population.
- We take a sample, calculate a statistic.
- We plot that statistic on the expected probability distribution.
- If the statistic is within the tail, we conclude that the sample is "unusual" relative to the reference population.

## Hypothesis testing

### Hypothesis testing: the $\alpha$ level, or significance level, of a test

### Hypothesis testing: procedure

1. We ask a yes/no question about a **population**.
2. We answer the question yes, and answer the question no, using symbols for the population means.
3. We label one answer the **null hypothesis** and the other answer the **alternative hypothesis**.
4. We decide the criterion for rejecting the null hypothesis.
5. The test is one of: **two-tailed**, **right-tailed**, or **left-tailed**.
6. We take a sample, and calculate our **test statistic** ( $Z$  or  $t$  for now)
7. We find if the observed test statistic is in the rejection region (**critical region** or **tail**) of the distribution.
8. If the statistic is in the rejection region, we **reject** the null hypothesis and **accept** the alternative hypothesis.
9. If the statistic is not in the rejection region, we **retain** the null hypothesis, and do not accept the alternative hypothesis.

### Examples

I flip a coin  $n$  times and get  $x$  heads. Is this a fair coin?

First we define the true means of the population we're interested in. The population is all possible coin flips of this particular coin. The population mean could be defined as  $p$ , the true proportion heads in an infinite number of flips. So first we answer the question "yes": yes, it is a fair coin would be equivalent to writing the equation  $p = 0.5$ . If the true proportion heads in an infinite number of flips is 0.5, then the coin is fair. Then we answer the question "no": if the coin is not fair, then the probability of landing heads is something besides 0.5. So the equation for this is  $p \neq 0.5$ . Now we decide: which of these equations is the null hypothesis and which is the alternative hypothesis? The null hypothesis is the hypothesis containing the equals sign, which is the first equation above. So the null hypothesis  $H_0$  is  $p = 0.5$ . So we can write  $H_0 : p = 0.5$ . This means the other equation is the alternative hypothesis:  $H_A : p \neq 0.5$ . Now we decide whether this is a 2-tailed, left-tailed, or right-tailed hypothesis. Because the alternative hypothesis contains the  $\neq$  sign, we know that this is a two-tailed test. This means the null hypothesis can be rejected in either direction, right or left. To the right, the coin's chances of landing heads is greater than 0.5. To the left, its chances of landing heads are less than 0.5. Next we write down the equation for the statistic that tests the null hypothesis. We have been using either  $t$  or  $Z$  so far. In this case we can use a  $Z$  test if  $n$  is sufficiently large (then we have a normal distribution).  $Z = \frac{x/n - 0.5}{\sqrt{p(1-p)/n}}$ . What is the criterion for rejection of the null hypothesis? The criterion is  $|Z| > Z_{crit}$ , where  $Z_{crit}$  is the value of  $Z$  that defines what the tail is of the distribution. Since this is a two tailed test, and the tails by definition contain  $\alpha$ , and  $\alpha$  is conventionally set to 0.5, then the  $Z_{crit}$  must be the  $Z$  value that cuts off a tail of 0.025. This  $Z$  value is 1.96. So  $|Z| > 1.96$  is the criterion for rejecting the null hypothesis that the coin is fair.

See the solutions to the examples below in the tablet file for this lecture.

### Examples

I create a new genotype of olive tree. Do its olives produce more oil than the old genotype?

## Rules

- If the alternative hypothesis contains the  $\neq$  sign, the test is two-tailed.
- If the alternative hypothesis contains the  $<$  sign, the test is left-tailed.
- If the alternative hypothesis contains the  $>$  sign, the test is right-tailed.

## Examples

I have a new chemical pesticide. Does it leave less residue on apples than the old kind?

## Examples

I take a poll. Are Croatians in favor of entering the EU?

## Examples

Is the sex ratio of males to females 1:1 in some population?

## Terminology

- The **critical value** of a test statistic is the value that defines the rejection region of the distribution.
- $t_{0.05/2}$  is the critical  $t$ -value for  $\alpha = 0.05$  and a 2-tailed test. Each tail here contains 0.025 of the probability.
- $t_{0.05/1}$  is the critical  $t$ -value for  $\alpha = 0.05$  and a 1-tailed test. The single tail contains 0.05 of the probability.
- $t_{0.01/1}$  is the critical  $t$ -value for  $\alpha = 0.01$  and a 1-tailed test. The single tail contains 0.01 of the probability.

## Examples

Is the level of arsenic in a well greater than some minimum value  $x$ ?

## Example using a single-sample $t$ -test

Does a particular fertilizer increase the potassium content of leaves? 16 plants were grown, and the potassium content before and after fertilizer was measured, and the difference reported. The mean difference was 5, and the standard deviation of the difference was 10.

## Example using a two-sample $t$ -test

Our first exam was taken by 14 women and 13 men. The mean female score was 72.7, and the mean male score was 70.3. The standard deviation in score was 7.7 for females and 8.3 for males. Is the mean for the female population the same as the mean for the male population?

**Equal variance  $t$ -test: use the pooled variance to calculate the sample standard deviation**

$$\begin{aligned}\text{pooled variance} &= \frac{(n_1 - 1)s_1^2 + (n_2 - 1)s_2^2}{n_1 + n_2 - 2} \\ s &= \sqrt{\text{pooled variance}} = \sqrt{\frac{(n_1 - 1)s_1^2 + (n_2 - 1)s_2^2}{n_1 + n_2 - 2}} \\ t &= \frac{\bar{x}_1 - \bar{x}_2}{\sqrt{s^2/n_1 + s^2/n_2}} = \frac{\bar{x}_1 - \bar{x}_2}{s\sqrt{1/n_1 + 1/n_2}}\end{aligned}$$

### Unequal variance $t$ -test

Calculate separate variances for the two samples:

$$t = \frac{\bar{x}_1 - \bar{x}_2}{\sqrt{s_1^2/n_1 + s_2^2/n_2}}$$

What is the degrees of freedom? Either use the smallest of  $n_1 - 1$  and  $n_2 - 1$ , or use R to calculate the following degrees of freedom:

$$\text{d.f.} = \frac{(s_1^2/n_1 + s_2^2/n_2)^2}{(s_1^2/n_1)^2/(n_1 - 1) + (s_2^2/n_2)^2/(n_2 - 1)}.$$

beginframe

### Example using a $Z$ test on proportions

1000 Slovenians and 1000 Croatians were polled. A total of 41% of Slovenians and 47% of Croatians stated in the poll that they believed that Croatia should enter the EU. Is the true proportion of people in favor of Croatia entering the EU the same in the two countries?

### Power of a two-sample $t$ -test

Our first exam was taken by 14 women and 13 men. The standard deviation in score was 7.7 for females and 8.3 for males. What is the power of a  $t$ -test on a difference in the means, if the true difference is 5 points, and  $\alpha = 0.05$  for a 2-tailed test?

### Power of a two-sample $t$ -test

The standard deviation in score of our first exam was 7.7 for females and 8.3 for males. What sample size would be necessary to detect a true difference of 1 point between the sexes in either direction? Assume  $\alpha = 0.05$ .

### Power of a two-sample $t$ -test

The standard deviation in score of our first exam was 7.7 for females and 8.3 for males. What sample size would be necessary to detect a true difference of 1 point between the sexes in either direction? Assume  $\alpha = 0.05$ .

### Power of a two-sample $t$ -test

The standard deviation in score of our first exam was 7.7 for females and 8.3 for males. What sample size would be necessary to detect a true difference of 1 point between the sexes in either direction? Assume  $\alpha = 0.05$ .

## Lecture 14: Hypothesis testing, continued

Announcements:

- Reading for this week's subject: pp. 61 - 79 in Vasilj.

### Hypothesis testing: example

We have a new variety of tomato. Does it have **the same** resistance to fusarium wilt as the old variety? To answer this question, we do an experiment. We expose 100 new plants to fusarium, and found that the mean growth rate was 35 with standard deviation 10. We know from previous experience that the mean growth rate of the old genotype is 33.

### Hypothesis testing: procedure

1. We ask a yes/no question about a **population**.
2. We answer the question yes, and answer the question no, using symbols for the population means.
3. We label one answer the **null hypothesis** and the other answer the **alternative hypothesis**.
4. We decide the criterion for rejecting the null hypothesis.
5. The test is one of: **two-tailed**, **right-tailed**, or **left-tailed**.
6. We take a sample, and calculate our **test statistic** ( $Z$  or  $t$  for now)
7. We find if the observed test statistic is in the rejection region (**critical region** or **tail**) of the distribution.
8. If the statistic is in the rejection region, we **reject** the null hypothesis and **accept** the alternative hypothesis.
9. If the statistic is not in the rejection region, we **retain** the null hypothesis, and do not accept the alternative hypothesis.

### Rules

- If the alternative hypothesis contains the  $\neq$  sign, the test is two-tailed.
- If the alternative hypothesis contains the  $<$  sign, the test is left-tailed.
- If the alternative hypothesis contains the  $>$  sign, the test is right-tailed.

### Hypothesis testing: example

We have a new variety of tomato. Does it have **the same** resistance to fusarium wilt as the old variety? To answer this question, we do an experiment. We expose 100 new plants to fusarium, and found that the mean growth rate was 35 with standard deviation 10. We know from previous experience that the mean growth rate of the old genotype is 33.

### Hypothesis testing: example

#### Hypothesis testing: example

We have a new variety of tomato. Does it have **lower** resistance to fusarium wilt than the old variety? To answer this question, we do an experiment. We expose 100 new plants to fusarium, and found that the mean growth rate was 35 with standard deviation 10. We know from previous experience that the mean growth rate of the old genotype is 33.

### Hypothesis testing: example

#### Hypothesis testing: example

We have a new variety of tomato. Does it have **higher** resistance to fusarium wilt than the old variety? To answer this question, we do an experiment. We expose 100 new plants to fusarium, and found that the mean growth rate was 35 with standard deviation 10. We know from previous experience that the mean growth rate of the old genotype is 33.

### Hypothesis testing: example

### Some terminology

- If we reject the null hypothesis, then we conclude our data are **significant**, or **significantly different** from the null.
- The **probability** of the test is the probability of a test result that is of equal or greater departure from the null hypothesis, if the null hypothesis were true.
- In a two-tailed test, the probability is twice the tail probability. In a one-tailed test, the probability is the one-tail probability.

### Some more terminology

- The least significant difference:

$$t = \frac{\text{observed} - \text{expected}}{s_{\text{observed}}}$$
$$t_{\text{crit}} = \frac{\text{observed} - \text{expected}}{s_{\text{observed}}}$$
$$\text{LSD} = t_{\text{crit}} s_{\text{observed}}$$

### Hypothesis testing: example

We have a new variety of tomato. Does it have **the same** resistance to fusarium wilt as the old variety? To answer this question, we do an experiment. We expose 100 new plants to fusarium, and found that the mean growth rate was 35 with standard deviation 10. We know from previous experience that the mean growth rate of the old genotype is 33.

### Hypothesis testing: example

### Hypothesis testing: example

### Hypothesis testing: example

We have a new variety of tomato. Does it have **the same** resistance to fusarium wilt as the old variety? To answer this question, we do an experiment. We expose 100 new plants and 100 old plants to fusarium, and found that the mean growth rate was 35 with standard deviation 10 for the new plants, and mean 33 and standard deviation 8 for the old plants.

### Hypothesis testing: example

### Unequal variance *t*-test

Calculate separate variances for the two samples:

$$t = \frac{\bar{x}_1 - \bar{x}_2}{\sqrt{s_1^2/n_1 + s_2^2/n_2}}$$

What is the degrees of freedom? Either use the smallest of  $n_1 - 1$  and  $n_2 - 1$ , or use R to calculate the following degrees of freedom:

$$\text{d.f.} = \frac{(s_1^2/n_1 + s_2^2/n_2)^2}{(s_1^2/n_1)^2/(n_1 - 1) + (s_2^2/n_2)^2/(n_2 - 1)}.$$

### Hypothesis testing: example



**Equal variance  $t$ -test:** use the pooled variance to calculate the sample standard deviation

$$\begin{aligned}\text{pooled variance} &= \frac{(n_1 - 1)s_1^2 + (n_2 - 1)s_2^2}{n_1 + n_2 - 2} \\ s &= \sqrt{\text{pooled variance}} = \sqrt{\frac{(n_1 - 1)s_1^2 + (n_2 - 1)s_2^2}{n_1 + n_2 - 2}} \\ t &= \frac{\bar{x}_1 - \bar{x}_2}{\sqrt{s^2/n_1 + s^2/n_2}} = \frac{\bar{x}_1 - \bar{x}_2}{s\sqrt{1/n_1 + 1/n_2}}\end{aligned}$$

### Hypothesis testing: example

We have a new variety of tomato. Does it have **the same** resistance to fusarium wilt as the old variety? To answer this question, we do an experiment. We expose 100 new plants and 100 old plants to fusarium, and found that the mean growth rate was 35 with standard deviation 10 for the new plants, and mean 33 and standard deviation 8 for the old plants.

### Hypothesis testing: example

### Hypothesis testing: one-sample test of proportion example

We have a new variety of tomato. Does it have **the same** resistance to fusarium wilt as the old variety? To answer this question, we do an experiment. We expose 100 new plants to fusarium, and found that 72 are resistant. We know in advance that 65% of the old plants are resistant.

### Hypothesis testing: example

### Hypothesis testing: two-sample test of proportions example

We have a new variety of tomato. Does it have **the same** resistance to fusarium wilt as the old variety? To answer this question, we do an experiment. We expose 100 new plants and 100 old plants to fusarium, and found that 72 new are resistant, and 65 old are resistant.

### Hypothesis testing: example

### Two sample test of proportions

Given  $x_1$  successes in a sample of  $n_1$ , and  $x_2$  successes in a sample of  $n_2$ . Therefore  $\hat{p}_1 = x_1/n_1$ , and  $\hat{p}_2 = x_2/n_2$ .

We define  $\bar{p} = \frac{x_1 + x_2}{n_1 + n_2}$ , and  $\bar{q} = 1 - \bar{p}$ .

Then our test statistic is:

$$Z = \frac{\hat{p}_1 - \hat{p}_2}{\sqrt{\frac{\bar{p}\bar{q}}{n_1} + \frac{\bar{p}\bar{q}}{n_2}}}$$

## 13 Statistical power

### Lecture 15: Type I, II errors, power of a hypothesis test

Announcements:

- Midterm exam 2 will be a week from today: Tuesday, 7 December, 14:00 - 16:00.
- Exam will cover problem sets 5, 6, 7, and 8, and selected problems from Problem Set 4.
- Answer keys for Problem Sets 5, 6, and 7 are posted on the web. You'll write down the answers to Problem Set 8 yourselves this week. Tomislav will help you with this on Thursday.
- Today: statistical power, and errors of a hypothesis test.
- Thursday: review session on all material for the test, especially this week's problem set.

### Problem

Statistical significance is not the same as biological “significance” (or meaning).

If you do not reject the null hypothesis, that could be because:

- 1) the null hypothesis is true, or
- 2) your test was not powerful enough: your sample size was too low.

What is the power of a test?

Let’s say you do a t-test to compare the mean weight of apples from grower 1 to grower 2, and find that you can’t reject the null hypothesis, and there is no “significant” difference. Does that mean there is really no true difference in weight of apples between the two growers? No. There could be a difference, but your test was not sensitive enough, i.e. its power was too low. How do you increase power? By increasing the sample size. If you increase the sample size, your estimate of the mean apple weight gets better and better, because the standard error of the mean gets smaller and smaller, so eventually, with high enough sample size, you can now reject the null hypothesis.

So this begs the question: what sample size do you need to increase the power of the test enough to reject the null hypothesis?

### Problem

As you recall, a screening test has four combinations of test outcome/condition.

In a screening test, a false negative occurs because power is not 100%. If you performed more than one test on a person, you might find that the first negative was a fluke, and that subsequent tests were positive. This might reduce the probability that you will falsely conclude that there is no condition (or disease etc.). So increasing sample size reduces the probability of false negative conclusions, and therefore increases the power of a test.

How do you calculate the power of a hypothesis test?

### Example

As you recall, a screening test has four combinations of test outcome/condition.

### Errors

Situation	Decision	
	Accept $H_0$	Reject $H_0$
$H_0$ true	$1 - \alpha$	$\alpha$
$H_0$ false	$\beta$	$1 - \beta$

### Terminology

- **Type I error** is  $\alpha$ .
- **Type II error** is  $\beta$ .
- **Sensitivity** of the test is the power, which is  $1 - \beta$ .
- **Specificity** of the test is  $1 - \alpha$ .

## Basic power equations, one-sample Z-test

$$Z_{pow} = |Z_{crit}| - \frac{\delta}{\frac{\sigma}{\sqrt{n}}} = |Z_{crit}| - \frac{\delta\sqrt{n}}{\sigma}$$
$$n = \sigma^2 \frac{(Z_{pow} - Z_{crit})^2}{\delta^2}$$
$$\delta = \sigma \frac{(Z_{crit} - Z_{pow})}{\sqrt{n}}$$

## Calculating power using R

$$\text{Power} = 1 - \text{pnorm}(Z_{pow})$$
$$Z_{pow} = \text{qnorm}(1 - \text{power})$$
$$Z_{crit} = \pm \text{qnorm}(1 - \alpha/T),$$

$T$  is the number of tails in the test. Note that the  $Z_{crit}$  is negative for a left-tailed test.

Also, the function `power.examp()` in `TeachingDemos` is a nice graphical demonstration of power. For example, you type `power.examp(diff = 2, alpha = 0.05, stdev = 10, n = 20)` to compute the power of detecting a true difference of 2, at  $\alpha = 0.05$ , for a standard deviation of 10 and sample size of 20.

### Example

A buyer wants to buy a large batch of apples from a grower. But he does not want to buy if the true mean weight is less than 150 g. He can't weigh all the apples, but he can weigh a sample of  $n$ , and calculate the sample mean. Let's say the true standard deviation is 30g. If he weighs 20 apples and does a  $Z$  test he may or may not reject the null hypothesis. If the **true** mean weight of apples is 150 g, what is the probability that he rejects the null hypothesis, and rejects the apples? Assume he is doing a one-tailed  $Z$ -test, with a significance level of 0.05 ( $\alpha = 0.05$ ). He rejects the apples only if he rejects the null hypothesis.

### Example

What if the true mean weight of the grower's apples is 149 g? Will he reject the null hypothesis? 145 g? 140 g?

### Example

What is the least significant difference (LSD) that the buyer must see before he rejects the grower's apples?

### Example

What is the probability that he will see the LSD if the true mean weight of the grower's apples is 149 g? 145 g? 140 g?

Power goes up with sample size.

### Example

If his sample size has to be 20, then what true difference ( $\delta$ ) will he be able to detect with his test 80% of the time? What if the sample size is 30? 50? 100? 500?

### Example

If he wants to detect, at a power of 80%, a true difference of only 2 g, what sample size will he need? 1 g? 0.5 g?

You need much larger sample sizes to detect smaller true differences.

### Example

If he wants to detect a true difference of 10 g, at a power of 80%, what should his sample size be? At a power of 90%? 95%? 99%?

You can detect smaller true differences with higher sample sizes.

### Example

If he wants to detect a true difference of 15 g, at a power of 80%, with a sample size of 20. What effect  $\alpha$  should he use?

If  $\alpha$  is smaller, power also decreases.

## 14 Testing for equality of variances: the $F$ -test

### Lecture 16: Testing for equality of variances: the $F$ -test

Announcements:

- Midterm exam 2 has been graded and results will be posted to the website.
- Problem set 8 due now.
- Problem set 9 due next Tuesday – Croatian translation will be uploaded to website today.
- Today: Testing for equality of variances with an  $F$  test.

Up until now, we have been using two samples for testing of the equality of the means of two populations. Another common question in biology is whether the variances of two populations are the same. Remember variance is a measure of the spread of a distribution, and the standard deviation is the square root of the variance.

Remember how variance is defined: differently for a continuous variable and a discrete variable.

### Measures of variability

- Variance =  $Var(x) = \sigma_x^2$
- Standard deviation =  $\sqrt{Var(x)} = \sigma_x$

### Variance of a distribution

**Definition 41.** The variance of a discrete probability distribution is

$$Var(x) = E(x - \mu)^2 = \sum_i (x_i - \mu)^2 P(x_i).$$

The variance of a continuous probability distribution is

$$Var(x) = E(x - \mu)^2 = \int_{-\infty}^{+\infty} (x - \mu)^2 f(x) dx.$$

These are the definitions of the absolute variance from a probability distribution.

We can visualize changes in the variance of a distribution using R, by plotting the distribution. For a normal distribution, we can do a plot with `dnorm()`. For example, set the `xvalues` as `xval = seq(-3, 3, .01)`. Then plot with `plot(xval, dnorm(xval), type = "l")`. Then you can see how changes in the variance affect changes in the distribution with `lines(xval, dnorm(xval, sd = 0.8), type = "l")`.

Now we imagine that we take a random sample from that distribution, and estimate the true variance from the sample.

## Calculation of the sample variance

$$\text{Sample variance} = s^2 = \frac{\sum_1^n (x_i - \bar{x})^2}{n - 1}$$

$$\text{Sample standard deviation} = s = \sqrt{s^2}$$

$n - 1$  is called  $v$ , or the “degrees of freedom” used in the calculation of  $s$ .

We can calculate variance of a sample  $x$  in R with `var(x)`. We can take a random sample from a normal distribution with the function `rnorm()`. So for example `x = rnorm(100)` creates 100 datapoints from a normal distribution with mean 0 and variance 1. We can calculate the sample variance of these data with `var(rnorm(100))`.

So, how do we test the hypothesis that two samples were drawn from the same population variance? We can't do a  $t$ -test, because the sample standard deviation is not normally distributed. Instead we do an  $F$ -test.

## $F$ -test steps

1. We want to know if the two population variances are equal. Null hypothesis is that  $\sigma_1^2 = \sigma_2^2$ .
2. Calculate the sample variance in the two samples,  $s_1^2$  and  $s_2^2$ .
3. Form the ratio of the larger over the smaller.
4. This ratio has an  $F$ -distribution with degrees of freedom for the numerator and denominator.
5. Compare the observed  $F$  with  $F_{crit}$ , which is `qf(1 - alpha/2, dfn, dfd)`. Note that we divide alpha by 2 for a two-tailed test.
6. Calculate the probability of the observed  $F$  as `2(1 - pf(F, dfn, dfd))`. Note that we multiply by 2 if this is a two-tailed test.

You can look up the critical values of  $F$ ,  $F_{crit}$  in Tablica C1 of the Vasilj textbook, page 299.

A good example of the use of the  $F$ -test is the Bumpus sparrow data. In Providence, Rhode Island, USA, in the 1890s there was a severe winter storm. Somebody collected 136 stunned sparrows and took them to Bumpus. Some of them revived, the rest died. He measured various morphological characteristics of the survivors and dead, and concluded that natural selection had eliminated individuals that were far from the average, which he considered the ideal, or the value of maximum fitness. If natural selection favors individuals closer to the mean, then the variance of the survivors should be lower than the variance of the non-survivors. We can test this hypothesis with an  $F$ -test.

## Bumpus data

In females, the variance of the humerus length of survivors was 0.176, and the variance for the non-survivors was 0.434. The number of survivors was 21, and the number of non-survivors was 28. Question: is the population variance of survivors different from the population variance of non-survivors?

## Exam scores, men and women

In our first midterm exam, the sample variance for women was 59.2, and the sample variance for men was 68.9. There were 14 women and 13 men taking the exam. Is the true population variance for women different from that for men?

## Confidence interval for the population variance

$$\chi^2 = \frac{vs^2}{\sigma^2},$$

$v$  is the degrees of freedom; so:

$$P(\chi_{0.025}^2 < \frac{vs^2}{\sigma^2} < \chi_{0.975}^2) = 0.95$$

$$P\left(\frac{vs^2}{\chi_{0.975}^2} < \sigma^2 < \frac{vs^2}{\chi_{0.025}^2}\right) = 0.95$$

High end of 95% confidence interval:  $\frac{vs^2}{\chi_{0.025}^2}$

Low end of 95% confidence interval:  $\frac{vs^2}{\chi_{0.975}^2}$

Note that this method assumes that the variable itself is normally distributed, and it is sensitive to small departures from normality. So you should use it with care!

## 15 Analysis of variance: ANOVA

### Lecture 17: Analysis of variance: ANOVA


Announcements:

- Exam has been graded. See website for results.
- Reading: Vasilj pp. 83-97.
- Problem set 9 due in class on Thursday, December 16.
- Problem set 10 due next Tuesday – Croatian translation will be uploaded to website.
- Today: Analysis of Variance (ANOVA).


**F-test purpose: to compare two variances**

**F-test purpose: to compare two variances**

**F distribucija** je teoretska distribucija vjerojatnosti, krivulja je asimetrična i različito izgleda, ovisno o  $n_1$  i  $n_2$  (slika 19).



**F-test purpose: to compare two variances**




The example we gave last time was the bampus sparrow data, where the humerus length in survivors had a significantly lower variance than the humerus length in non-survivors. The  $F$  value was 2.47, with a  $p = 0.04$ .


#### **$F$ -test steps**

1. We want to know if the two population variances are equal. Null hypothesis is that  $\sigma_1^2 = \sigma_2^2$ .
2. Calculate the sample variance in the two samples,  $s_1^2$  and  $s_2^2$ .
3. Form the ratio of the larger over the smaller.
4. This ratio has an  $F$ -distribution with degrees of freedom for the numerator and denominator.
5. Compare the observed  $F$  with  $F_{crit}$ , which is  $qf(1 - \alpha/2, dfn, dfd)$ . Note that we divide alpha by 2 for a two-tailed test.
6. Calculate the probability of the observed  $F$  as  $2(1 - pf(F, dfn, dfd))$ . Note that we multiply by 2 if this is a two-tailed test.

**$F$ -test purpose: to compare two variances**



*F*-test purpose: to compare two variances





### ANOVA problem


We want to test for the equality of means in more than two samples.

If we have two samples:  $t$ -test or  $Z$ -test.

If we have more than two samples: ANOVA.


In fact, ANOVA for two samples is mathematically the same as a  $t$ -test.

### ANOVA problem



Use the function `chist()` to demonstrate sampling multiple samples from the same population.

### ANOVA problem



### Single-factor ANOVA, $g$ samples (groups), same $n$ in all groups

1. We want to know if all population means are equal to each other. The null hypothesis is  $\mu_1 = \mu_2 = \dots = \mu_k$ .
2. Calculate the **observed** variance among the group means.  $\text{Var}(\bar{x}_{obs}) = \frac{\sum (\bar{x}_k - \bar{\bar{x}})^2}{g-1}$ .
3. Calculate the **expected** variance among group means,  $\text{Var}(\bar{x}_{exp}) = \frac{s^2}{n}$ .
4.  $s^2$  is the pooled sample variance across all groups:  $s^2 = \sum s_k^2 / g$ .
5. Form the ratio of observed to expected:

$$F = \frac{\text{Var}(\bar{x}_{obs})}{\text{Var}(\bar{x}_{exp})} = \frac{\frac{\sum (\bar{x}_k - \bar{\bar{x}})^2}{g-1}}{\frac{s^2}{n}}.$$

### Single-factor ANOVA, $g$ samples (groups), same $n$ in all groups

1. Compare the observed  $F$  with  $F_{crit}$ , which is  $\text{qf}(1 - \alpha/2, \text{dfn}, \text{dfd})$ . Note that we divide alpha by 2 for a two-tailed test. Numerator degrees of freedom is  $g - 1$ . Denominator degrees of freedom is  $\sum_k n_k - g$ .
2. Calculate the probability of the observed  $F$  as  $2(1 - \text{pf}(F, \text{dfn}, \text{dfd}))$ . Note that we multiply by 2 if this is a two-tailed test.

### Single-factor ANOVA, $g$ samples (groups), same $n$ in all groups

- 1.

$$F = \frac{\text{Var}(\bar{x}_{obs})}{\text{Var}(\bar{x}_{exp})} = \frac{\frac{\sum (\bar{x}_k - \bar{\bar{x}})^2}{g-1}}{\frac{s^2}{n}} = \frac{n \sum (\bar{x}_k - \bar{\bar{x}})^2 / (g-1)}{s^2}.$$

$$F = \frac{\text{SS}_g / \text{df}_g}{s^2}.$$

2.

$$s^2 = \frac{(n-1) \sum (x_1 - \bar{x}_1)^2 + \dots + (n-1) \sum (x_g - \bar{x}_g)^2}{\sum n - g}.$$

3.

$$F = \frac{SS_g/\text{df}_g}{SS_w/\text{df}_w} = \frac{MS_g}{MS_w}.$$

**Single-factor ANOVA,  $g$  samples (groups), different  $n_k$  in all groups**

1.

$$F = \frac{\sum n_k (\bar{x}_k - \bar{\bar{x}})^2 / (g-1)}{s^2}.$$

$$F = \frac{SS_g/\text{df}_g}{s^2}.$$

2.

$$s^2 = \frac{(n_1-1) \sum (x_1 - \bar{x}_1)^2 + \dots + (n_g-1) \sum (x_g - \bar{x}_g)^2}{\sum n_k - g}.$$

3.

$$F = \frac{SS_g/\text{df}_g}{SS_w/\text{df}_w} = \frac{MS_g}{MS_w}.$$

## 16 Multiple comparisons and 2-factor ANOVA

### Lecture 18: 2-factor ANOVA


Announcements:

- Exam has been graded. See website for results.
- Reading: Vasilj pp. 187-225.
- Problem set 10 due Thursday.
- Problem set 11 due on January 4, 2011.
- Today: Multiple comparisons, and 2-factor analysis of variance (ANOVA).

**$F$ -test purpose: to compare two variances**

**$F$ -test purpose: to compare two variances**

**F distribucija** je teoretska distribucija vjerojatnosti, krivulja je asimetrična i različito izgleda, ovisno o  $n_1$  i  $n_2$  (slika 19).



The  $F$ -test is used to test for the equality of population variances, using the ratio of sample variances, which follows an  $F$  distribution. The  $F$ -test is used in the analysis of variance (ANOVA), to test for the equality of means, not variances, by comparing the observed variance among sample means to the predicted variance among sample means if all samples were drawn from the same population.

The single-factor ANOVA can be thought of as a multi-sample  $t$ -test. If there are only two samples, then the ANOVA is mathematically equivalent to a two-sample  $t$ -test assuming equal variances in the two populations.

#### Single-factor ANOVA, $g$ samples (groups), same $n$ in all groups

1. We want to know if all population means are equal to each other. The null hypothesis is  $\mu_1 = \mu_2 = \dots = \mu_k$ .
2. Calculate the **observed** variance among the group means.  $\text{Var}(\bar{x}_{obs}) = \frac{\sum(\bar{x}_k - \bar{\bar{x}})^2}{g-1}$ .
3. Calculate the **expected** variance among group means,  $\text{Var}(\bar{x}_{exp}) = \frac{s^2}{n}$ .
4.  $s^2$  is the pooled sample variance across all groups:  $s^2 = \sum s_k^2 / g$ .
5. Form the ratio of observed to expected:

$$F = \frac{\text{Var}(\bar{x}_{obs})}{\text{Var}(\bar{x}_{exp})} = \frac{\frac{\sum(\bar{x}_k - \bar{\bar{x}})^2}{g-1}}{\frac{s^2}{n}}.$$

#### Single-factor ANOVA, $g$ samples (groups), same $n$ in all groups

1. Compare the observed  $F$  with  $F_{crit}$ , which is  $\text{qf}(1 - \alpha, \text{dfn}, \text{dfd})$ . Note that an ANOVA is a 1-tailed test, because it's not possible for the true ratio to be less than 1. Numerator degrees of freedom is  $g - 1$ . Denominator degrees of freedom is  $\sum_k n_k - g$ .
2. Calculate the probability of the observed  $F$  as  $1 - \text{pf}(F, \text{dfn}, \text{dfd})$ .

**Single-factor ANOVA,  $g$  samples (groups), same  $n$  in all groups**

1.

$$F = \frac{\text{Var}(\bar{x}_{obs})}{\text{Var}(\bar{x}_{exp})} = \frac{\frac{\sum(\bar{x}_k - \bar{\bar{x}})^2}{g-1}}{\frac{s^2}{n}} = \frac{n \sum(\bar{x}_k - \bar{\bar{x}})^2 / (g-1)}{s^2}.$$

$$F = \frac{SS_g / df_g}{s^2}.$$

2.

$$s^2 = \frac{(n-1) \sum(x_1 - \bar{x}_1)^2 + \dots + (n-1) \sum(x_g - \bar{x}_g)^2}{\sum n - g}.$$

3.

$$F = \frac{SS_g / df_g}{SS_w / df_w} = \frac{MS_g}{MS_w}.$$

**Single-factor ANOVA,  $g$  samples (groups), different  $n_k$  in all groups**

1.

$$F = \frac{\sum n_k (\bar{x}_k - \bar{\bar{x}})^2 / (g-1)}{s^2}.$$

$$F = \frac{SS_g / df_g}{s^2}.$$

2.

$$s^2 = \frac{(n_1-1) \sum(x_1 - \bar{x}_1)^2 + \dots + (n_g-1) \sum(x_g - \bar{x}_g)^2}{\sum n_k - g}.$$

3.

$$F = \frac{SS_g / df_g}{SS_w / df_w} = \frac{MS_g}{MS_w}.$$

In the single-factor ANOVA, we have just one response variable and one predictor variable. The response variable is a continuous variable (like height or weight), and the predictor variable is a factor (like sex or species or treatment).

In an ANOVA, the result is an analysis of variance table, which always has the same components.

**Single-factor ANOVA**

### **Analysis of Variance Table**

**Response: Prevalence**

	Df	Sum Sq	Mean Sq	F value	Pr(>F)
Alcohol	3	3.4510	1.15032	12.617	6.914e-07 ***
Residuals	84	7.6584	0.09117		

**Single-factor ANOVA**

1.

$$F = \frac{MS_g}{MS_w}.$$

2.

$$MS_g = \frac{SS_g}{df_g}$$

3.

$$SS_g = \sum n_k(\bar{x}_k - \bar{\bar{x}})^2$$

4.

$$df_g = g - 1$$

5.

$$MS_w = \text{pooled variance} = s^2 = \frac{SS_w}{df_w}$$

### Single-factor ANOVA

1.

$$MS_w = \text{pooled variance} = s^2 = \frac{SS_w}{df_w}$$

2.

$$SS_w = (n_1 - 1) \sum (x_1 - \bar{x}_1)^2 + \dots + (n_g - 1) \sum (x_g - \bar{x}_g)^2$$

3.

$$df_w = n - g$$

For example, let's take the dataset `data(chickwts)` in R. This is the weight of chicks after some length of time fed on different diets. We want to know if there is any difference among the diets in the weight of chicks. We answer this question with a single-factor ANOVA: `anova(lm(weight ~ feed))`. We see in the ANOVA that there is a significant difference. The next question is: which diets are actually different from each other? We can visualize the differences among diets with the command `boxplot(weight ~ feed)`, but the plot doesn't answer the question. We can answer this question by doing all possible *t*-tests for all possible pairs of diets. But for each *t*-test, the false positive rate is 0.05. If we do many *t*-tests, the probability that at least one of the tests gives a false positive is greater than 0.05. So doing a large number of tests, looking for a significant result, is a "fishing expedition" that greatly increases the false positive rate. We can correct this problem by adjusting the probability of the test upwards based on the number of tests performed, and still using the 0.05 criterion for significance. This adjusted probability is used in the Tukey Honestly Significant Difference method for performing multiple *t*-tests.

### Multiple comparisons problem

- Multiple *t*-tests inflate the false positive rate!
- Therefore, we need to adjust for the number of comparisons.
- The Tukey Honestly Significant Difference tests makes this adjustment.
- In R, the function is `TukeyHSD()`.
- First do `a = aov(response ~ predictor)`, then `TukeyHSD(a, "predictor")`.

For the chick weight data, we do:

```
data(chickwts)
attach(chickwts)
a = aov(weight ~string~ feed)
Tukey(a, "feed")

Tukey multiple comparisons of means
 95% family-wise confidence level

Fit: aov(formula = weight ~string~ feed)

$feed
```

	diff	lwr	upr	p adj
horsebean-casein	-163.383333	-232.346876	-94.41979	0.0000000
linseed-casein	-104.833333	-170.587491	-39.07918	0.0002100
meatmeal-casein	-46.674242	-113.906207	20.55772	0.3324584
soybean-casein	-77.154762	-140.517054	-13.79247	0.0083653
sunflower-casein	5.333333	-60.420825	71.08749	0.9998902
linseed-horsebean	58.550000	-10.413543	127.51354	0.1413329
meatmeal-horsebean	116.709091	46.335105	187.08308	0.0001062
soybean-horsebean	86.228571	19.541684	152.91546	0.0042167
sunflower-horsebean	168.716667	99.753124	237.68021	0.0000000
meatmeal-linseed	58.159091	-9.072873	125.39106	0.1276965
soybean-linseed	27.678571	-35.683721	91.04086	0.7932853
sunflower-linseed	110.166667	44.412509	175.92082	0.0000884
soybean-meatmeal	-30.480519	-95.375109	34.41407	0.7391356
sunflower-meatmeal	52.007576	-15.224388	119.23954	0.2206962
sunflower-soybean	82.488095	19.125803	145.85039	0.0038845

This shows us the difference in mean weight between each feed type, the 95% confidence interval for the difference (adjusted for the multiple comparisons) and the probability in the *t*-test (again, adjusted for the multiple comparisons). Here we find out that the largest significant difference in weight is between sunflower and horsebean diets. Some diets are not significantly different from each other, as for example sunflower and casein.

## 2-factor ANOVA: calculations

We do three hypothesis tests in 2-factor ANOVA. We test for:

- the main effect of factor 1.
- the main effect of factor 2.
- the interaction between the two factors.

Take, for example, two factors, very commonly used: sex, and drugs. Often when you take a medication, there are side effects. For example, if you take steroids. The side effects are different in men and women. For example, women taking steroids often develop a deeper voice.

## 2-factor ANOVA: calculations

- The mean squares for each effect are the same as for the single-factor ANOVA.
- The within mean squares is different from the single-factor ANOVA: it is the pooled variance within each cell, rather than within each effect.
- The interaction is calculated as a difference.

## 2-factor ANOVA: calculations

- The mean squares for each effect are the same as for the single-factor ANOVA.
- The within mean squares is different from the single-factor ANOVA: it is the pooled variance within each cell, rather than within each effect.
- The interaction is calculated as a difference.

As an example, let's do a 2-factor ANOVA for voice deepness, with factors steroid use and sex.

Here are the data: no steroids, female: mean is 3, variance is 3, sample size is 20. yes steroids, female: mean is 6, variance is 3, sample size is 20. no steroids, male: mean is 7, variance is 2, sample size is 20. yes steroids, male: mean is 8, variance is 2, sample size is 20.

How do we calculate the ANOVA table for these data, assuming we're testing for a drug effect, a sex effect, and the interaction between the two? See the tablet file for the calculations.

Some additional calculations are generally interesting: the coefficient of determination, which is the proportion of the total variation that is explained by the effect. This is just the sum of squares for the effect divided by the total sum of squares. The total sum of squares is the sum of all the sums of squares (factor1, factor2, interaction, within). Also, when we have a 2-by-2 table, we can actually calculate the interaction. This is effect of factor 1 within the first level of factor two, minus the effect of factor 1 within the second level of factor 2. In the steroid example the sample or observed interaction is  $(3-6) - (7-8) = (-3) - (-1) = -2$ .

What are the reasons for doing a 2-factor ANOVA, rather than two single-factor ANOVAs? First, if there is an interaction, then you explain more of the variation. Second, if there is an interaction, then your interpretation of the main effect is different, it becomes conditional. And third, whether there is an interaction or not, then the mean squares within is smaller if you do a 2-factor ANOVA, which results in a higher power of the test.

**2-factor ANOVA: coefficient of determination,  $R^2$**

### Analysis of Variance Table

**Response: Prevalence**

	Df	Sum Sq	Mean Sq	F value	Pr(>F)
Alcohol	3	3.4510	1.15032	12.617	6.914e-07 ***
Residuals	84	7.6584	0.09117		

In the above ANOVA table, the coefficient of determination is  $3.45/(3.45 + 7.66) = 0.31$ . In other words, 31% of the variation in the response variable is explained by the predictor. Variation is synonymous with "sum of squares". To calculate the coefficient of determination ( $R^2$ ) for a factor, we divide its sum of squares with the total sum of squares. The total sum of squares is the factor sum of squares plus the within sum of squares.

**2-factor ANOVA: rationale**

Why do a 2-factor ANOVA rather than two 1-factor ANOVAs?

- If an interaction exists, then the interpretation of the main effect is different.
- If an interaction exists, then we can demonstrate it and increase the proportion of variation that we have explained (i.e. increase the coefficient of determination,  $R^2$ ).
- If there is a large effect of both of the factors, then including both in the same analysis greatly reduces the within mean squares, which increases the power of the test for demonstrating both main effects.

## 17 Covariance, correlation, and linear regression

**Lecture 19: Covariance, correlation, and linear regression**

Announcements:

- Reading: Chapter 7 in Vasilj.
- Problem set 11 due today.
- Today: Covariance, correlation, and linear regression



**Problem. Two continuous variables: are they related?**

Examples

- Parents versus offspring.
- Growth curves (organs, organisms, populations).
- Allometric relationships.
- Physiological relationships.
- Ecological relationships.

**Correlation versus causation**



Storks

**Correlation versus causation**




Babies

**Correlation versus causation**




Storks and babies?

Correlation versus causation



Correlation versus causation



### Covariance

1.  $Cov(x, y) = E[(x - \bar{x})(y - \bar{y})]$

- 2.

$$Cov(x, y) = \frac{\sum_i [(x_i - \bar{x})(y_i - \bar{y})]}{n - 1}$$

### Properties of the covariance

1.  $Cov(x, a) = 0$ , where  $a$  is any constant.
2.  $Cov(x, x) = Var(x)$ , for any variable  $x$ .
3.  $Cov(x, y) = Cov(y, x)$ , for any variables  $x$  and  $y$ .
4.  $Cov(x, bx) = bVar(x)$ , for any variables  $x$  and  $y$  and constant  $b$ .
5.  $Cov(ax, by) = abCov(x, y)$ , for any constants  $a$  and  $b$ .
6.  $Cov(x + a, y + b) = Cov(x, y)$ , for any constants  $a$  and  $b$ .
7. If  $x$  and  $y$  are independent, then  $Cov(x, y) = 0$ .

### Calculating the slope and intercept of the best-fit line

The best fit line is the line that minimizes the residual sum of squares.

- 1.

$$\text{Slope} = b = \frac{Cov(x, y)}{Var(x)}.$$

2. Intercept =  $a = \bar{y} - b\bar{x}$ .

### Correlation versus covariance

Correlation coefficient is the scaled covariance.

1. The maximum covariance between  $x$  and  $y$  is  $s_x s_y$ .

2.

$$\text{Cor}(x, y) = r = \frac{\text{Cov}(x, y)}{s_x s_y}.$$

3. The correlation coefficient lies between -1 and +1.

### Strength of the correlation coefficient

korelacijski koeficijent ( $r$ )	jačina korelacije
0.00 – 0.10	nema
0.10 – 0.25	vrlo slaba
0.25 – 0.40	slaba
0.40 – 0.50	srednja
0.50 – 0.75	jaka
0.75 – 0.90	vrlo jaka
0.90 – 1.00	potpuna

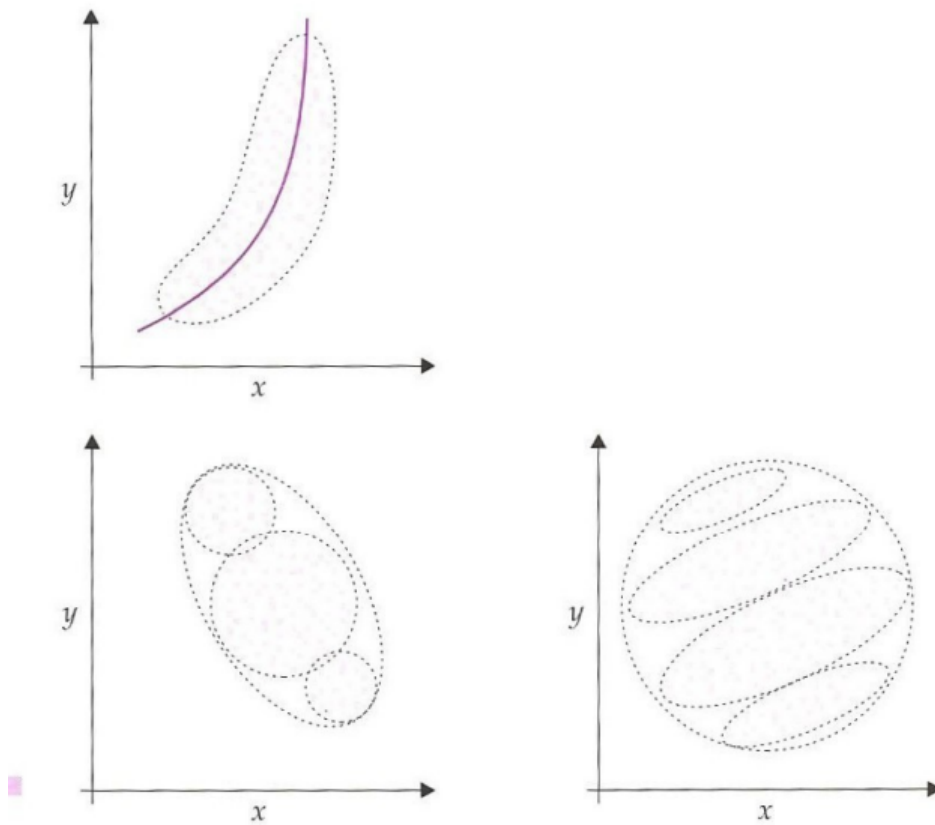
### Simple linear regression

1. Calculate the best-fit straight line  $\hat{y} = a + bx$ .
2. Calculate the residual sum of squares.  $SS_{res} = \sum (y - \hat{y})^2$ .
3. Calculate the line sum of squares  $SS_x = \sum (\hat{y} - \bar{y})^2$ .
4. Calculate the corresponding residual and line mean squares:  $MS_{res} = SS_{res}/df_{res}$ ; and  $MS_x = SS_x/df_x$   $df_{res} = n - 2$ ;  $df_x = 1$
5. Form the  $F$  ratio and calculate its tail probability  $F = MS_{res}/MS_x$ .
6. Or let R do it all with `anova(lm(y ~ x))`.

### Assumptions of linear regression

1. The  $x$  values are measured without any error
2. The relationship between  $x$  and  $y$  is linear
3. The residuals are normally distributed
4. The variance is constant (regardless of the value of  $x$ )

## Interpretation of correlation



## Lecture 20: Regression

Announcements:


- Reading: Chapter 7 in Vasilj.
- Problem sets 12 and 13 due Thursday, January 13 at 15:00.
- Third colloquium exam on Tuesday, January 18, 14:00 to 18:00.
- Course anketa on Thursday, January 20 at 15:00.
- Today: Regression review, nonlinear regression

### Problem. Two continuous variables: are they related?


Examples

- Parents versus offspring.
- Growth curves (organs, organisms, populations).
- Allometric relationships.
- Regression of traits on fitness.
- Physiological relationships.
- Ecological relationships.

### Regression of traits on fitness



### Ecological regressions can be messy



### Covariance

1.  $Cov(x, y) = E[(x - \bar{x})(y - \bar{y})]$

- 2.

$$Cov(x, y) = \frac{\sum_i [(x_i - \bar{x})(y_i - \bar{y})]}{n - 1}$$

### Properties of the covariance

1.  $Cov(x, a) = 0$ , where  $a$  is any constant.
2.  $Cov(x, x) = Var(x)$ , for any variable  $x$ .
3.  $Cov(x, y) = Cov(y, x)$ , for any variables  $x$  and  $y$ .
4.  $Cov(x, bx) = bVar(x)$ , for any variables  $x$  and  $y$  and constant  $b$ .
5.  $Cov(ax, by) = abCov(x, y)$ , for any constants  $a$  and  $b$ .
6.  $Cov(x + a, y + b) = Cov(x, y)$ , for any constants  $a$  and  $b$ .
7. If  $x$  and  $y$  are independent, then  $Cov(x, y) = 0$ .

### Simple linear regression

1. Calculate the best-fit straight line  $\hat{y} = a + bx$ .
2. Calculate the residual sum of squares.  $SS_{res} = \sum (y - \hat{y})^2$ .
3. Calculate the line sum of squares  $SS_x = \sum (\hat{y} - \bar{y})^2$ .
4. Calculate the corresponding residual and line mean squares:  $MS_{res} = SS_{res}/df_{res}$ ; and  $MS_x = SS_x/df_x$   $df_{res} = n - 2$ ;  $df_x = 1$
5. Form the  $F$  ratio and calculate its tail probability  $F = MS_{res}/MS_x$ .
6. Or let R do it all with `anova(lm(y ~ x))`.

### Nonlinear regression

1. Approach 1: linearize if possible, and then do a simple linear regression.
2. Approach 2: use the raw data and find the best-fit line by a computer search.


### Nonlinear regression

1. Calculate the best-fit curve in R using `nls()`.
2. Plot the best fit curve over the scatter diagram using `curve()`.
3. Estimate the standard error of each parameter in the equation.
4. Form the  $t$  value and test for the significance of each parameter in the equation.
5. Or let R do it all with `summary(nls())`.


### Assumptions of linear regression

1. The  $x$  values are measured without any error
2. The relationship between  $x$  and  $y$  is linear
3. The residuals are normally distributed
4. The variance is constant (regardless of the value of  $x$ )

### Correlation versus causation

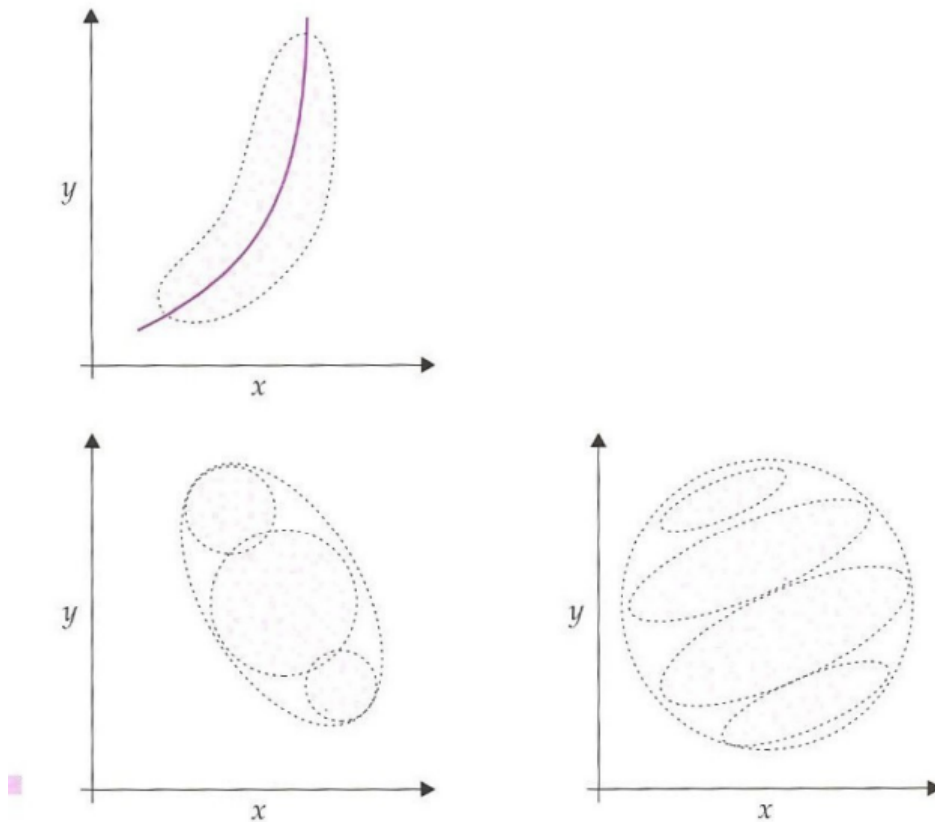


### Correlation versus causation



### Interpretation of correlation





## Lecture 21: Non-parametric statistics, course review

Announcements:

- Reading: Chapter 7 in Vasilj.
- Problem sets 12 and 13 due now.
- Third colloquium exam on Tuesday, January 18, 14:00 to 18:00.
- Course anketa on Thursday, January 20 at 15:00.
- Today: Non-parametric statistics, review

### Assumptions of “parametric” statistical tests

1. Observations are independent of each other.
2. Variables are measured without bias.
3. Predictor variables are measured exactly.
4. Residuals are normally distributed.
5. The true variance is independent of the predictor variable.

### We can test these assumptions

1. Test for normality with Shapiro-Wilkes test, `shapiro.test()` in R.
2. You can do `plot(lm(y ~ x))` to look at both normality and assumption of constant variance.

**If the data do not meet the assumptions**

Two choices:

1. Transform the response variable (for example, log, square root) and retest the data.
2. Use a non-parametric test instead.

**Non-parametric tests**

1. Non-parametric single factor ANOVA: Kruskal-Wallis test, `kruskal.test()`
2. Non-parametric multiple comparisons: Pairwise Wilcoxon test, `pairwise.wilcox.test()`
3. Or, just rank transform and do a parametric test.

**What is statistics?**

- Set of tools
- to convert *data*
- to *probability*.

**What is statistics?**

- Set of tools
- to convert *data*
- to *probability*.