Introductory Statistics Explained
List of Supporting Videos

©Jeremy Balka

Last updated: March 29, 2019.

This is a complete list of available videos and descriptions. Some videos are on material not covered in 2040. At the start of each set of lecture outlines I give a list of videos for the chapter and state which ones are not required in 2040.

1 Descriptive Statistics

1.1 Measures of Central Tendency (8:31)

Measures of Central Tendency (8:31) (http://youtu.be/NM_iOLUwZFA)

A brief introduction to measures of central tendency. The mean, median, and mode are introduced and calculated for a simple example. The relationship between the mean and median for different shapes of distributions is then discussed.

The guinea pig survival time data can be found in:
Doksum, K. (1974). Empirical probability plots and statistical inference for nonlinear models in the two-sample case. Annals of Statistics, 2:267–277.

1.2 Measures of Variability (12:12)

Measures of Variability (12:12) (http://youtu.be/Cx2tGUze60s)

An introduction to measures of variability. I discuss the range, mean absolute deviation, variance, and standard deviation, and work through a simple example of calculating these quantities. I then discuss interpreting the standard deviation, including a brief discussion of the empirical rule.

The birth weight data is from random sample of 1000 males drawn from Table 7-2 (Live births, by birth weight and geography – Males) of the Statistics Canada publication 84F0210X, available at http://www.statcan.gc.ca/pub/84f0210x/2009000/t011-eng.htm.

1.3 The Sample Variance: Why Divide by n-1? (6:53)

The Sample Variance: Why Divide by n-1? (6:53) (http://youtu.be/9ONRMymR2Eg)

An informal discussion of why we divide by $n - 1$ in the sample variance formula. I give some motivation for why we should divide by something less than n, and (casually) discuss the concept of degrees of freedom (in the context of the sample variance). I have another video with a mathematical proof that dividing by $n - 1$ results in an unbiased estimator of the population variance, available at http://youtu.be/D1hgiAla3KI.

1.4 Z-Scores (As a Descriptive Measure of Relative Standing) (8:08)

Z-Scores (As a Descriptive Measure of Relative Standing) (8:08) (http://youtu.be/EhUvGRddC4M)

An introduction to z-scores as a descriptive measure of relative standing. (I don't do any probability calculations in this video.) I do a simple calculation example, discuss the empirical rule in the context of z-scores, and illustrate what the z-score tells us about how large or small an observation is.

The mean and standard deviation of the heights of American males is based on information from:

McDowell MA, Fryar CD, Ogden CL, Flegal KM. Anthropometric reference data for children and adults: United States, 2003-2006. National health statistics reports; no 10. Hyattsville, MD: National Center for Health Statistics. 2008.

(The mean is taken directly from this document, and the standard deviation is estimated based on the given percentiles.)

2 Probability

2.1 Basics of Probability: Unions, Intersections, Complements

Basics of Probability: Unions, Intersections, Complements (10:45) (https://youtu.be/B1v9OeCTlu0)

An introductory discussion of unions, intersections, and complements in the context of basic probability. (Pitched at a level appropriate for a typical introductory statistics course.) I include a discussion of mutually exclusive events, as well as the addition rule. I work through a simple die rolling example, and also an example of a scenario where we randomly sample from a population of people, where some have diabetes and some have hypertension. The various scenarios are illustrated with Venn and Euler diagrams. (The percentages given in the video are loosely based on statistics for 50 year-old Canadian males.)

2.2 An Introduction to Conditional Probability (12:01)

An Introduction to Conditional Probability (12:01) (https://youtu.be/bgCMjHzXTXs)

An introduction to conditional probability, pitched at a level appropriate for a typical introductory statistics course. I work through some simple examples in this introductory video, and briefly touch on the concept of independence in an example at the end of the video. I work through some harder examples in (Conditional Probability Example Problems (16:39) (https://youtu.be/ES9HFNDu4Bs)).

2.3 Independent Events (Basic Probability: Independence of Two Events) (21:25)

Independent Events (Basic Probability: Independence of Two Events) (https://youtu.be/1wuRV5zOPPE)

An introduction to the concept of independent events, pitched at a level appropriate for the probability section of a typical introductory statistics course. I give the definition of independence, work through some simple examples, and attempt to illustrate the meaning of independence in various ways. (Note: I use the phrase "not independent" rather than "dependent" almost exclusively. There is nothing wrong with calling events dependent when they are not independent, but I prefer to use "not independent" for a couple of reasons.)

(I'm on a bit of a probability run, but looking forward to getting back to statistics videos in the near future.)

2.4 Conditional Probability Example Problems (16:39)

Conditional Probability Example Problems (16:39) (`https://youtu.be/ES9HFNDu4Bs`)

Conditional probability example problems, pitched at a level appropriate for a typical introductory statistics course. I assume that viewers have already been introduced to the concepts of conditional probability and independence, but I do review the concepts along the way. I work through some problems with the conditional probability formula explicitly, and some using the reduced sample space argument.

The sudden death data is slightly modified from:

Naneix et al. (2015). Sudden adult death: An autopsy series of 534 cases with gender and control comparison. Journal of Forensic and Legal Medicine, 32:10-15.

The data was pulled from their Figure 3, and I pooled the Abdominal/pelvian and undetermined groups into "other", to make the example work better visually and have it be easier to follow. I took some slight liberties here, as "undetermined" is not the same as "other". Conscious choice, y'all.

2.5 What Does Independence Look Like on a Venn Diagram? (10:45)

What Does Independence Look Like on a Venn Diagram? (10:45) (`http://youtu.be/pV3nZAsJxl0`)

Usually, Venn diagrams are not very useful for illustrating independence, as the sizes of the circles and their intersections have no meaning. It can help to illustrate independence if we force the area of each region to be equal to its probability of occurring. Independence is even easier to see if we represent the events with rectangles instead of circles. I illustrate these concepts in this video.

All plots were created in R. The appropriate diameters of the circles and distance between the centres of the circles were calculated in R using the package venneuler.

Not all plots in this video are officially called Venn diagrams.

2.6 Are Mutually Exclusive Events Independent? (5:05)

Are Mutually Exclusive Events Independent? (5:04) (https://youtu.be/U0fsad9WWwk)

Are mutually exclusive events independent? I get asked variants of this question frequently, so it's evident that some students confuse these two concepts. The one sentence summary: If A and B are mutually exclusive events, then they are independent if and only if $\mathrm{P(A)} = 0$ or $\mathrm{P(B)} = 0$. I take a more detailed look in this video.

2.7 De Morgan's Laws (in a probability context) (7:40)

De Morgan's Laws (in a probability context) (7:40) (https://youtu.be/LBGbwQDhceg)

A discussion of De Morgan's laws, in the context of basic probability. I illustrate De Morgan's laws using Venn diagrams, describe their meaning in a worded example, and show how they might be useful in a probability calculation.

2.8 Proof that if two events are independent, so are their complements (4:45)

Proof that if two events are independent, so are their complements(4:45) (https://youtu.be/bnDpZNlVZ3k)

Here I prove that if events A and B are independent, so are Ac and Bc. I make use of De Morgan's Laws, without offering a formal proof of that part (but I do provide a brief Venn diagram justification of the needed bit).

2.9 Proof that if events A and B are independent, so are Ac and B (and A and Bc) (4:23)

Proof that if events A and B are independent, so are Ac and B (and A and Bc)(4:23) (https://youtu.be/8qyAth0T2rI)

Here I prove that if events A and B are independent, so are A complement and B. (And A and B complement, of course, as which event we call A and which we call B is arbitrary.)

Looking for a proof that if A and B are independent, so are their complements? That's here: https://youtu.be/bnDpZNlVZ3k

2.10 $P(A) = P(A \cap B) + P(A \cap Bc)$ (6:03)

P(A) = P(A n B) + P(A n Bc) (6:03) (https://youtu.be/Pu33eESSczU)

A quick video to illustrate that P(A) = P(A n B) + P(A n Bc), and work through a simple conditional probability example that makes use of this identity.

My Law of Total Probability video is here:

2.11 The Law of Total Probability (10:21)

The Law of Total Probability (10:21) (https://youtu.be/7t9jyikrG7w)

I discuss the Law of Total Probability. I begin with some motivating plots, then move on to a statement of the law, then work through two examples.

3 Discrete Probability Distributions

3.1 An Introduction to Discrete Random Variables and Discrete Probability Distributions (14:12)

An Introduction to Discrete Random Variables and Discrete Probability Distributions (14:12) (http://youtu.be/oHcrna8Fk18)

An introduction to discrete random variables and discrete probability distributions. A few examples of discrete and continuous random variables are discussed.

This is an updated and revised version of an earlier video. Those looking for my original Intro to Discrete Random Variables video can find it at: http://youtu.be/0P5WRKihQ4E

3.2 The Expected Value and Variance of Discrete Random Variables (11:20)

The Expected Value and Variance of Discrete Random Variables (11:20) (http://youtu.be/Vyk8HQOckIE)

An introduction to the expected value and variance of discrete random variables. The formulas are introduced, explained, and an example is worked through.

This is an updated and refined version of an earlier video. Those looking for the original version can find it at http://youtu.be/OvTEhNL96v0.

3.3 Introduction to the Bernoulli Distribution (5:02)

Introduction to the Bernoulli Distribution (5:02) (http://youtu.be/bT1p5tJwn_0)

An introduction to the Bernoulli distribution, a common discrete probability distribution.

3.4 The Bernoulli Distribution: Deriving the Mean and Variance (3:12)

The Bernoulli Distribution: Deriving the Mean and Variance (3:12) (http://youtu.be/bC6WIpRgMuc)

I derive the mean and variance of the Bernoulli distribution.

3.5 An Introduction to the Binomial Distribution (14:11)

An Introduction to the Binomial Distribution (14:11) (http://youtu.be/qIzC1-9PwQo)

An introduction to the binomial distribution. I discuss the conditions required for a random variable to have a binomial distribution, discuss the binomial probability mass function and the mean and variance, and look at two examples involving probability calculations.

The estimated probability of a 90 year old Canadian male surviving for one year was taken from Statistics Canada life tables, which can be found at (http://www.statcan.gc.ca/pub/84-537-x/4064441-eng.htm) . The probability given in the table is the estimated probability that a randomly selected Canadian male, given survival to his 90th birthday, survives until his 91st. I simplified this explanation a little in the example in the video.

3.6 Binomial/Not Binomial: Some Examples (8:08)

Binomial/Not Binomial: Some Examples (8:08) (http://youtu.be/UJFIZY0xx_s)

I work through a few word problems. For some the random variable has a binomial distribution, for others it does not. I hope to illustrate when the binomial distribution is appropriate, and when it is not. This video is inspired by the many times I've been asked something along the lines of "I don't understand when to use the binomial distribution".

3.7 The Binomial Distribution: Mathematically Deriving the Mean and Variance (13:54)

The Binomial Distribution: Mathematically Deriving the Mean and Variance (13:54) (http://youtu.be/8fqkQRjcR1M)

I derive the mean and variance of the binomial distribution. I do this in two ways. First, I assume that we know the mean and variance of the Bernoulli distribution, and that a binomial random variable is the sum of n independent Bernoulli random variables. I then take the more difficult approach, where we do not lie on this relationship and derive the mean and variance from scratch.

3.8 An Introduction to the Hypergeometric Distribution (15:35)

An Introduction to the Hypergeometric Distribution (15:35) (http://youtu.be/L2KMttDm3aY)

An introduction to the hypergeometric distribution. I briefly discuss the difference between sampling with replacement and sampling without replacement. I describe the conditions required for the hypergeometric distribution to hold, discuss the formula, and work through 2 simple examples.

I also discuss the relationship between the binomial distribution and the hypergeometric distribution, and a rough guideline for when the binomial distribution can be used as a reasonable approximation to the hypergeometric. I finish with a brief example involving the multivariate hypergeometric distribution.

3.9 An Introduction to the Poisson Distribution (9:03)

An Introduction to the Poisson Distribution (9:03) (http://youtu.be/jmqZG6roVqU)

An introduction to the Poisson distribution. I discuss the conditions required for a random variable to have a Poisson distribution. work through a simple calculation example, and briefly discuss of the relationship between the binomial distribution and the Poisson distributions.

3.10 Poisson or Not? (When does a random variable have a Poisson distribution?) (14:40)

Poisson or Not? (When does a random variable have a Poisson distribution?) (14:40) (http://youtu.be/sv_KXSiorFk)

A not-too-technical look at the conditions required for a random variable to have a Poisson distribution. It can be difficult to determine whether a random variable actually has a Poisson distribution, so here I look at a few examples and some visual illustrations that may help. There are no probability calculations carried out in this video. I assume that the viewer has already been introduced to the Poisson distribution, but I do a brief review at the start.

3.11 The Poisson Distribution: Mathematically Deriving the Mean and Variance (9:17)

The Poisson Distribution: Mathematically Deriving the Mean and Variance (9:17) (http://youtu.be/65n_v92JZeE)

I derive the mean and variance of the Poisson distribution.

3.12 Discrete Probability Distributions: Example Problems (Binomial, Poisson, Hypergeometric, Geometric) (14:51)

Discrete Probability Distributions: Some Examples (Binomial, Poisson, Hypergeometric, Geometric)(14:51) (http://youtu.be/Jm_Ch-iESBg)

I work through a few probability examples based on some common discrete probability distributions (binomial, Poisson, hypergeometric, geometric – but not necessarily in this order). I assume that you've been previously introduced to these distributions (although this isn't necessary for the geometric problem, as the probability is easily calculated from basic probability rules). Students sometimes have difficulty determining the appropriate distribution to use, so this video may give some help with the proper thought process.

3.13 The Relationship Between the Binomial and Poisson Distributions (5:24)

The Relationship Between the Binomial and Poisson Distributions (5:24) (http://youtu.be/eexQyHj6hEA)

A look at the relationship between the binomial and Poisson distributions (roughly, that the Poisson distribution approximates the binomial for large n and small p). I work through some calculations in an example, showing that the approximate probability from the Poisson can be quite close to the exact probability from the binomial distribution.

(The example used involves albinism. Albinism affects all races, but the rates of albinism vary a little around the world. In Europe and North America, roughly 1 in 20,000 people have some form of albinism).

3.14 Proof that the Binomial Distribution tends to the Poisson Distribution (8:25)

Proof that the Binomial Distribution tends to the Poisson Distribution (8:25) (http://youtu.be/ceOwlHnVCqo)

A proof that as n tends to infinity and p tends to 0 while np remains constant, the binomial distribution tends to the Poisson distribution.

3.15 Introduction to the Geometric Distribution (10:48)

Introduction to the Geometric Distribution (10:48) (`http://youtu.be/zq9Oz82iHf0`)

An introduction to the geometric distribution. I discuss the underlying assumptions that result in a geometric distribution, the formula, and the mean and variance of the distribution. I work through an example of the calculations and then discuss the cumulative distribution function.

3.16 Introduction to the Negative Binomial Distribution (7:33)

Introduction to the Negative Binomial Distribution (7:33) (http://youtu.be/BPlmjp2ymxw)

An introduction to the negative binomial distribution, a common discrete probability distribution. In this video I define the negative binomial distribution to be the distribution of the number of *trials* needed to obtain r successes in repeated independent Bernoulli trials. Different sources define it in different ways (the distribution of the number of *failures* before obtaining r successes, for example.)

3.17 Introduction to the Multinomial Distribution (11:15)

Introduction to the Multinomial Distribution (11:15) (http://youtu.be/syVW7DgvUaY)

An introduction to the multinomial distribution, a common discrete probability distribution. I discuss the basics of the multinomial distribution and work through two examples of probability calculations. For comparison purposes, I finish off with a quick example of a multivariate hypergeometric probability calculation.

3.18 Overview of Some Discrete Probability Distributions (6:21) (Bernoulli, Binomial, Geometric, Negative Binomial, Hypergeometric, Poisson)

Overview of Some Discrete Probability Distributions (6:21) (Bernoulli, Binomial, Geometric, Negative Binomial, Hypergeometric, Poisson) (http://youtu.be/UrOXRvG9oYE)

A brief overview of some common discrete probability distributions (Bernoulli, Binomial, Geometric, Negative Binomial, Hypergeometric, Poisson). I discuss when these distributions arise and the relationships between them. I do not do any calculations in this video, or discuss the probability mass functions or other characteristics of the distributions. This video is simply an overview of the distributions that can either be used as a summary recap, or a quick introduction.

4 Continuous Random Variables and Continuous Probability Distributions

4.1 An Introduction to Continuous Probability Distributions (5:52)

An Introduction to Continuous Probability Distributions (5:52) (http://youtu.be/OWSOhpS00_s)

An introduction to continuos random variables and continuous probability distributions. I briefly discuss the probability density function (pdf), the properties that all pdfs share, and the notion that for continuous random variables probabilities are areas under the curve.

4.2 Finding Probabilities and Percentiles for a Continuous Probability Distribution (11:59)

Finding Probabilities and Percentiles for a Continuous Probability Distribution

(11:59) (http://youtu.be/EPm7FdajBvc)

I work through an example of finding the median, areas under the curve, and the cumulative distribution function for a continuous probability distribution. I assume a basic knowledge of integral calculus. I derive the mean and variance of this distribution in another video: http://youtu.be/Ro7dayHU5DQ.

4.3 Deriving the Mean and Variance of a Continuous Probability Distribution (7:22)

Deriving the Mean and Variance of a Continuous Probability Distribution (7:22) (http://youtu.be/Ro7dayHU5DQ)

I work through an example of deriving the mean and variance of a continuous probability distribution. I assume a basic knowledge of integral calculus.

4.4 Introduction to the Continuous Uniform Distribution (7:03)

Introduction to the Continuous Uniform Distribution (7:03) (http://youtu.be/izE1dXrH5JA)

A brief introduction to the (continuous) uniform distribution. I discuss its pdf, median, mean, and variance. I also work through an example of finding a probability and a percentile. I don't do any integration in this video.

4.5 An Introduction to the Normal Distribution (5:27)

An Introduction to the Normal Distribution (5:27) (http://youtu.be/iYiOVISWXS4)

An introduction to the normal distribution, often called the Gaussian distribution. The normal distribution is an extremely important continuous probability distribution that arises very frequently in probability and statistics.

4.6 Standardizing Normally Distributed Random Variables (10:28)

Standardizing Normally Distributed Random Variables (10:28) (http://youtu.be/4R8xm19DmPM)

I discuss standardizing normally distributed random variables (turning variables with a normal distribution into something that has a standard normal distribution).

I work through an example of a probability calculation, and an example of finding a percentile of the distribution. It is assumed that you can find values from the standard normal distribution, using either a table or a computer.

The mean and variance of adult female heights in the US is estimated from statistics found in the National Health Statistics Reports:

McDowell MA, Fryar CD, Ogden CL, Flegal KM. Anthropometric reference data for children and adults: United States, 2003-2006. National health statistics reports; no 10. Hyattsville, MD: National Center for Health Statistics. 2008.

4.7 The Normal Approximation to the Binomial Distribution (14:10)

The Normal Approximation to the Binomial Distribution (14:10) (http://youtu.be/CCqWkJ_pqNU)

An introduction to the normal approximation to the binomial distribution. I discuss a guideline for when the normal approximation is reasonable, and the continuity correction. I work through examples of greater than, greater than or equal to, less than, and less than or equal to, and include plots that may help students to visualize the concepts. I assume that viewers are already familiar with finding values under the standard normal curve.

4.8 Normal Quantile-Quantile Plots (12:09)

Normal Quantile-Quantile Plots (12:09) (http://youtu.be/X9_ISJ0YpGw)

An introduction to normal quantile-quantile (QQ) plots (a graphical method for assessing whether a set of observations is approximately normally distributed). I discuss the motivation for the plot, the construction of the plot, then look at several examples. In the examples I look at what a normal quantile-quantile plot looks like when sampling from various other distributions. I then illustrate what normal QQ plots look like when sampling from a normal distribution by simulating several samples, for two different sample sizes.

4.9 An Introduction to the Chi-Square Distribution (5:29)

An Introduction to the Chi-Square Distribution (5:29) (http://youtu.be/hcDb12fsbBU)

A brief introduction to the chi-square distribution. I discuss how the chi-square distribution arises, its pdf, mean, variance, and shape.

4.10 An Introduction to the t Distribution (6:10)

An Introduction to the t Distribution (6:10) (http://youtu.be/T0xRanwAIiI)

An introduction to the t distribution, a common continuous probability distribution. I discuss how the t distribution arises, its pdf, its mean and variance, and its relationship to the standard normal distribution. I illustrate the relationship between the t distribution and the standard normal distribution through a series of plots.

This video contains some mathematical details regarding the t distribution (the pdf, mean, variance, how the distribution arises). I also have a video with a less technical introduction to the t distribution, available at http://youtu.be/Uv6nGIgZMVw.

4.11 An Introduction to the F Distribution (4:05)

An Introduction to the F Distribution (4:05) (http://youtu.be/G_RDxAZJ-ug)

A brief introduction to the F distribution, an important continuous probability distribution that frequently arises in statistical inference. I discuss how the F distribution arises, its pdf, mean, median, and shape.

5 Using Tables to Find Areas and Percentiles (Z, t, χ^2, F

5.1 Finding Areas Using the Standard Normal Table (for tables that give the area to left of z) (6:16)

Finding Areas Using the Standard Normal Table (for tables that give the area to left of z) (6:16) (http://youtu.be/-UljIcq_rfc)

I work through some examples of finding areas under the standard normal curve using the standard normal table. If you fully understand how to find values in the standard normal table, this video will not be of much use to you. The table used is one that gives areas to the left of z (the cumulative distribution function).

5.2 Finding Percentiles Using the Standard Normal Table (for tables that give the area to left of z) (7:33)

Finding Percentiles Using the Standard Normal Table (for tables that give the area to left of z) (7:33) (http://youtu.be/9KOJtiHAavE)

I work through some examples of finding the z value corresponding to a given area under the standard normal curve, using the standard normal table. If you fully understand how to find values in the standard normal table, this video will not be of much use to you. The table used is one that gives areas to the left of z (the cumulative distribution function).

5.3 Finding Areas Using the Standard Normal Table (for tables that give the area between 0 and z) (8:55)

Finding Areas Using the Standard Normal Table (for tables that give the area between 0 and z) (8:55) (http://youtu.be/dwRVY0FjpQM)

I work through some examples of finding areas under the standard normal curve using the standard normal table. If you fully understand how to find values in the standard normal table, this video will not be of much use to you. The table used is one that gives areas between 0 and z.

5.4 Finding Percentiles Using the Standard Normal Table (for tables that give the area between 0 and z) (9:42)

Finding Percentiles Using the Standard Normal Table (for tables that give the area between 0 and z) (9:42) (http://youtu.be/0qQEz74FfTI)

I work through some examples of finding the z value corresponding to a given area under the standard normal curve, using the standard normal table. If you fully understand how to find values in the standard normal table, this video will not be of much use to you. The table used is one that gives areas between 0 and z.

5.5 Using the t Table to Find Areas and Percentiles (7:56)

Using the t Table to Find Areas and Percentiles (7:56) (http://youtu.be/qH7QZoMbB80)

I show how to find areas and percentiles for the t distribution using the t table.

5.6 R Basics: Finding Percentiles and Areas for the t Distribution (3:00)

R Basics: Finding Percentiles and Areas for the t Distribution (3:00) (http://youtu.be/ZnIFX4PTAOo)

A quick intro to finding percentiles and areas for the t distribution using R (a free statistical software package), using the commands qt and pt.

5.7 Using the F Table to Find Areas and Percentiles (9:05)

Using the F Table to Find Areas and Percentiles (9:05) (http://youtu.be/mSn55vREkIw)

I show how to find percentiles and areas for the F distribution using the F table. In the first part of the video I work through examples finding areas and percentiles in the right tail of the distribution. (These are often of more importance to us than the left tail values.) In the second part of the video I show the mathematical "trick" that allows us to find left tail values from a table that gives only right tail values, and work through examples.

5.8 Finding Percentiles and Areas for the F Distribution Using R (3:11)

Finding Percentiles and Areas for the F Distribution Using R (3:11) (http://youtu.be/PZiVe5DMJWA)

I show how to find percentiles and areas for the F distribution using R (a free statistical software package), using the commands qf and pf.

5.9 Using the Chi-square Table to Find Areas and Percentiles (5:44)

Using the Chi-square Table to Find Areas and Percentiles (5:44) (http://youtu.be/C-0uN1inmcc)

I show how to find percentiles and areas for the chi-square distribution using the chi-square table.

5.10 Using R to Find Chi-square Areas and Percentiles (3:56)

Using R to Find Chi-square Areas and Percentiles (3:56) (http://youtu.be/SuAHkMtudbo)

I show how to find percentiles and areas for the chi-square distribution using R (a free statistical software package), using the commands qchisq and pchisq.

6 Sampling Distributions

6.1 Sampling Distributions: Introduction to the Concept (7:52)

Sampling Distributions: Introduction to the Concept (7:52) (http://youtu.be/Zbw-YvELsaM)

I discuss the concept of sampling distributions (an important statistical concept that underlies much of statistical inference), and illustrate the sampling distribution of the sample mean in a simple example.

6.2 The Sampling Distribution of the Sample Mean (11:39)

The Sampling Distribution of the Sample Mean (11:39) (http://youtu.be/q50GpTdFYyI)

I discuss the sampling distribution of the sample mean, and work through an example of a probability calculation. (I only briefly mention the central limit theorem here, but discuss it in more detail in another video).

The mean and standard deviation of the amount of protein in a quarter pound patty of lean beef was found in the USDA nutrient database.

6.3 Introduction to the Central Limit Theorem (13:14)

Introduction to the Central Limit Theorem (13:14) (http://youtu.be/Pujol1yC1_A)

I discuss the central limit theorem, a very important concept in the world of statistics. I illustrate the concept by sampling from two different distributions, and for both distributions plot the sampling distribution of the sample mean for various sample sizes. I also discuss why the central limit theorem is important in statistics, and work through a probability calculation. (For the most part this is a non-technical treatment, and simply illustrates the important implications of the central limit theorem.)

6.4 Deriving the Mean and Variance of the Sample Mean (5:07)

Deriving the Mean and Variance of the Sample Mean (5:07)
(http://youtu.be/7mYDHbrLEQo)

I derive the mean and variance of the sampling distribution of the sample mean. I have another video where I discuss the sampling distribution of the sample mean and work through some example probability calculations. It's called "The Sampling Distribution of the Sample Mean", and it's available at: http://youtu.be/q50GpTdFYyI.

6.5 Proof that the Sample Variance is an Unbiased Estimator of the Population Variance (6:58)

Proof that the Sample Variance is an Unbiased Estimator of the Population Variance (6:58)
(`http://youtu.be/D1hgiAla3KI`)

A proof that the sample variance (with n-1 in the denominator) is an unbiased estimator of the population variance.

7 Confidence Intervals

7.1 Introduction to Confidence Intervals (6:42)

Introduction to Confidence Intervals (6:42) (http://youtu.be/27iSnzss2wM)

An introduction to confidence intervals.

7.2 Deriving a Confidence Interval for the Mean (6:40)

Deriving a Confidence Interval for the Mean (6:40) (http://youtu.be/-iYDu8flFXQ)

I derive the appropriate formula for a confidence interval for the population mean mu, when sampling from a normally distributed population with a known value of sigma. (The purpose of this video is to show the rationale behind the confidence interval formula for mu.) I do not work through an example in this video; I work through an example in http://youtu.be/7I-HCHg2eTE.

7.3 Intro to Confidence Intervals for One Mean (Sigma Known)(10:37)

Intro to Confidence Intervals for One Mean: (Sigma Known) (10:37)
(http://youtu.be/KG921rfbTDw)

An introduction to confidence intervals for the population mean mu. These methods are appropriate when we are sampling from a normally distributed population, where the population standard deviation sigma is known. When the population standard deviation is not known, as is usually the case, we need to use a slightly different method (a method based on the t distribution).

Data reference The 2D:4D ratio data is simulated data with the same summary statistics as found in:

Stevenson et al. (2007). Attention Deficit/Hyperactivity Disorder (ADHD) Symptoms and Digit Ratios in a College Sample. American Journal of Human Biology. 19:41–50.

7.4 Finding the Appropriate z Value for the Confidence Interval Formula (5:37)

Finding the Appropriate z Value for the Confidence Interval Formula (5:37)
(http://youtu.be/grodoLzThy4)

I show how to find the appropriate z value (using the standard normal table) when calculating a confidence interval.

7.5 Confidence Intervals for One Mean: Interpreting the Interval (6:02)

Confidence Intervals for One Mean: Interpreting the Interval (6:02)
(http://youtu.be/JYP6gc–sGQ)

Interpreting a confidence interval for the population mean mu. Some reasonable interpretations are discussed, as are some common misconceptions. Many of these concepts hold for confidence intervals for other parameters.

7.6 What Factors Affect the Margin of Error? (4:05)

What Factors Affect the Margin of Error? (4:05) (http://youtu.be/NQtcGOhUWB4)

A look at the effect of various factors (the standard deviation, the sample size, and the confidence level) on the margin of error of a confidence interval for the population mean.

7.7 Confidence Intervals for One Mean: Sigma Not Known (t Method) (9:46)

Confidence Intervals for One Mean: Sigma Not Known (t Method) (9:46) (http://youtu.be/bFefxSE5bmo)

Introduction to confidence intervals for mu based on the t distribution. These methods are appropriate when we are sampling from a normally distributed population and the population standard deviation (sigma) is not known.

The cereal data used in this video is real data from a sample of 7 cereal boxes I purchased one day. The boxes had a stated weight of 750 grams. (I've changed the story slightly in the video, but it is real data.)

7.8 Intro to the t Distribution (non-technical) (8:55)

Intro to the t Distribution (non-technical) (8:55) (http://youtu.be/Uv6nGIgZMVw)

A brief non-technical introduction to the t distribution, how it relates to the standard normal distribution, and how it is used in inference for the mean.

7.9 Confidence Intervals for One Mean: Determining the Required Sample Size (5:15)

Confidence Intervals for One Mean: Determining the Required Sample Size (5:15) (http://youtu.be/7zcbVaVz_P8)

I find the sample size required to obtain a given margin of error in a confidence interval for mu. I discuss the appropriate formula and work through an example.

7.10 Confidence Intervals for One Mean: Investigating the Normality Assumption (8:57)

Confidence Intervals for One Mean: Assumptions (8:57)

(http://youtu.be/mE5vH2wDoIs)

A discussion of the assumptions when using the t procedure to construct a confidence interval for the population mean mu. The assumptions are discussed, and the effect of different violations of the normality assumption is investigated through simulation.

8 Hypothesis Testing

8.1 An Introduction to Hypothesis Testing (9:54)

An Introduction to Hypothesis Testing (9:54) (http://youtu.be/tTeMYuS87oU)

A first look at hypothesis testing.

8.2 Z Tests for One Mean: Introduction (11:13)

Z Tests for One Mean: Introduction (11:13) (http://youtu.be/pGv13jvnjKc)

An introduction to Z tests for a population mean mu. These tests are appropriate when sampling from a normally distributed population where sigma is known. I discuss the hypotheses and the underlying logic, and work through an example. I do not discuss the rejection region or p-value in an in-depth way in this video – I discuss them in other videos, and provide links at the end of this one.

8.3 Z Tests for One Mean: The Rejection Region Approach (10:24)

Z Tests for One Mean: The Rejection Region Approach (10:24) (http://youtu.be/60x86lYtWI4)

An introduction to the rejection region approach to reaching a conclusion in a Z test. This video discusses the rejection region in the context of a Z test for one mean, but the same logic holds for other Z tests.

8.4 Z Tests for One Mean: The p-value (10:02)

Z Tests for One Mean: The p-value (10:02) (http://youtu.be/m6sGjWz2CPg)

An introduction to the concept of the p-value in a Z test. This video discusses the p-value in the context of a Z test for one mean, but the same logic holds for other Z tests.

8.5 Z Tests for One Mean: An Example (6:26)

Z Tests for One Mean: An Example (6:26) (http://youtu.be/Xi33dGcZCA0)

An example of a Z test on the population mean mu. This test is appropriate when sampling from a normally distributed population, when the population standard deviation sigma is known.

8.6 What is a p-value?

Original (fast) version: What is a p-value? (5:44) (http://youtu.be/HTZ8YKgD0MI)

A brief intro to the concept of the p-value, in the context of one-sample Z tests for the population mean. Much of the underlying logic holds for other tests as well.

New version:

What is a p-value? (Updated and Extended Version) (10:51) (http://youtu.be/UsU-O2Z1rAs)

An introduction to the concept of the p-value, in the context of one-sample Z tests for the population mean. Much of the underlying logic holds for other tests as well. I discuss what a p-value is, then using simulation I illustrate its distribution when the null hypothesis is true and when the null hypothesis is false. I then give a rough guideline of what different p-values mean in terms of evidence against the null hypothesis.

8.7 Type I Errors, Type II Errors, and the Power of the Test (8:11)

Type I Errors, Type II Errors, and the Power of the Test (8:11) (http://youtu.be/7mE-K_w1v90)

A discussion of Type I errors, Type II errors, their probabilities of occurring (alpha and beta), and the power of a hypothesis test.

8.8 One-Sided Test or Two-Sided Test? (9:25)

One-Sided Test or Two-Sided Test? (9:25) (http://youtu.be/VP1bhopNP74)

A discussion of when to use a one-sided alternative hypothesis and when to use a two-sided alternative hypothesis in hypothesis testing. I assume that the viewer has already had a brief introduction to the notion of one-sided and two-sided tests.

8.9 Statistical Significance versus Practical Significance (4:47)

Statistical Significance versus Practical Significance (4:47) (http://youtu.be/_k1MQTUCXmU)

A brief discussion of the meaning of statistical significance, and how it is strongly related to the sample size.

8.10 The Relationship Between Confidence Intervals and Hypothesis Tests (5:36)

The Relationship Between Confidence Intervals and Hypothesis Tests (5:36) (http://youtu.be/k1at8VukIbw)

I discuss the relationship between a (two-sided) confidence interval and a two-sided hypothesis test. I discuss the relationship in terms of inference for one mean, but the same concept holds in many other settings as well.

8.11 Calculating Power and the Probability of a Type II Error (A One-Tailed Example) (11:32)

Calculating Power and the Probability of a Type II Error (A One-Tailed Example) (11:32)
(http://youtu.be/BJZpx7Mdde4)

An example of calculating power and the probability of a Type II error (beta), in the context of a one-tailed Z test for one mean. Much of the underlying logic holds for other types of tests as well.

8.12 Calculating Power and the Probability of a Type II Error (A Two-Tailed Example) (13:40)

Calculating Power and the Probability of a Type II Error (A Two-tailed Example) (13:40) (http://youtu.be/NbeHZp23ubs)

An example of calculating power and the probability of a Type II error (beta), in the context of a two-tailed Z test for one mean. Much of the underlying logic holds for other types of tests as well.

8.13 What Factors Affect the Power of a Z Test? (12:25)

What Factors Affect the Power of a Z Test? (12:25) (http://youtu.be/K6tado8Xcug)

A look at what factors influence the power of a test. This discussion is in the setting of a one-sample Z test on the population mean, but the concepts hold for many other types of test as well. I discuss what factors affect power, and illustrate the concepts visually using various plots. There are no power calculations carried out in this video; I have another video on calculating power: http://youtu.be/BJZpx7Mdde4

8.14 Hypothesis Testing in 17 Seconds (0:17)

Hypothesis Testing in 17 Seconds (0:17) (http://youtu.be/wyTwHmxs4ug)

An ultra-brief summary of hypothesis testing. (Just for fun.)

8.15 t Tests for One Mean: Introduction (13:46)

t Tests for One Mean: Introduction (13:46) (http://youtu.be/T9nI6vhTU1Y)

An introduction to t tests for one population mean. I briefly discuss when we use the test, and when we would use a z test instead. I also briefly discuss the hypotheses of the test, and the p-value for different alternatives. I then work through an example. (If you are comfortable with the basics of hypothesis testing, and understand the difference between t and z procedures from confidence intervals, then much of this video may be review.) If you are just looking for an example, it starts at 7:00.

The reaction time data is simulated data with the same summary statistics as found in:

Armstrong et al. (2012). Mild Dehydration Affects Mood in Healthy Young Women. Journal of Nutrition, 142: 382–388.

8.16 t Tests for One Mean: An Example (9:43)

t Tests for One Mean: An Example (9:43) (http://youtu.be/kQ4xcx6N0o4)

An example of a t test on one population mean. I first work through an example of a t test, then briefly investigate the influence of 3 outliers on the conclusions.

The sleep misperception index data is simulated data with the same summary statistics as found in:

Manconi et al. (2010). Measuring the error in sleep estimation in normal subjects and in patients with insomnia. Journal of Sleep Research, 19:478–486.

8.17 t Tests for One Mean: Investigating the Normality Assumption (7:55)

t Tests for One Mean: Investigating the Normality Assumption (7:55) (http://youtu.be/U1O4ZFKKD1k)

A discussion of the assumptions of the t test on one mean. (The assumptions are the same as those of the t confidence intervals for one mean.) The assumptions are discussed, and the effect of different violations of the normality assumption is investigated through simulation.

8.18 Hypothesis tests on one mean: t test or z test? (6:58)

Hypothesis tests on one mean: t test or z test? (6:58) (http://youtu.be/vw2IPZ2aD-c)

A look at at what influences the choice of the t test or z test in one-sample hypothesis tests on the population mean mu. I work through an example of a t test, and compare the p-value of the t test to would have been found had we (incorrectly) used a z test.

8.19 Don't watch this! (A t test example where nearly everything I say is wrong) (5:37)

Don't watch this! (A t test example where nearly everything I say is wrong) (5:37) (https://youtu.be/Y2hL5o8k12g)

(Recorded in 2013, but misplaced and never released.)

I work through an example of a t test, and (intentionally) make many false statements. Some of them might sound pretty reasonable. The lesson is: Get your statistics help from a reputable source!

8.20 Using the *t* Table to Find the P-value in One-Sample *t* Tests (7:11)

Using the *t* Table to Find the P-value in One-Sample *t* Tests (7:11) (http://youtu.be/tI6mdx3s0zk)

I work through examples of finding the p-value for a one-sample t test using the t table. (It's impossible to find the exact p-value using the t table. Here I illustrate how to find the appropriate interval of values in which the p-value must lie.)

8.21 Finding Areas Under the t Distribution (6:02)

Finding Areas Under the t Distribution (6:02) (http://youtu.be/2d9i2QeDJIY)

Here I work through two examples of finding areas under the t distribution, using both R and the t table.

9 Inference for Two Means

9.1 Inference for Two Means: Introduction (6:21)

Inference for Two Means: Introduction (6:21) (http://youtu.be/86ss6qOTfts)

I introduce inference procedures for the difference between two means in the case where the population standard deviations are known. I discuss the sampling distribution of the difference in the sample means, and discuss the confidence interval formula and the hypothesis test of the equality of population means.

The empathy/young offender data is simulated data with the same summary statistics as found in:

Owen, T., Fox, S. (2011). Experiences of shame and empathy in violent and non-violent young offenders. The Journal of Forensic Psychiatry & Psychology, 22(4):551-563.

9.2 The Sampling Distribution of the Difference in Sample Means $(\bar{X}_1 - \bar{X}_2)$ (10:08)

The Sampling Distribution of the Difference in Sample Means $(\bar{X}_1 - \bar{X}_2)$ (10:08) (http://youtu.be/4HB-FL529ag)

I discuss the characteristics of the sampling distribution of the difference in sample means $(\bar{X}_1 - \bar{X}_2)$. I then work through an example of a probability calculation that involves these concepts.

The values for male and female heights are based on information in the 2009-2011 Canadian Health Measures Survey.

9.3 Pooled-Variance t Tests and Confidence Intervals: Introduction (11:04)

Pooled-Variance t Tests and Confidence Intervals: Introduction (11:04) (http://youtu.be/NaZBdj0nCzQ)

An introduction to pooled-variance t tests and confidence intervals (in the setting of inference for two means).

The shame/young offender data is simulated data with the same summary statistics as found in:

Owen, T., Fox, S. (2011). Experiences of shame and empathy in violent and non-violent young offenders. The Journal of Forensic Psychiatry & Psychology, 22(4):551-563.

9.4 Pooled-Variance t Tests and Confidence Intervals: An Example (12:41)

Pooled-Variance t Tests and Confidence Intervals: An Example (12:41) (http://youtu.be/Q526z1mz4Sc)

I work through an example of a pooled-variance t test and confidence interval (in the setting of inference for two means).

The MENT/PTSD data is simulated data with the same summary statistics as found in:

Geraerts et al. (2009). Detecting deception of war-related posttraumatic stress disorder. The Journal of Forensic Psychiatry & Psychology, 20(2):278-285.

9.5 Welch (Unpooled Variance) t Tests and Confidence Intervals: Introduction (9:43)

Welch (Unpooled Variance) t Tests and Confidence Intervals: Introduction (9:43) (http://youtu.be/2-ecXltt2vI)

An introduction to Welch (unpooled variance) t tests and confidence intervals. (Inference for two means.)

The Cairo traffic police officer data is simulated data with the same summary statistics as found in:

Kamal, A., Eldamaty, S., and Faris, R. (1991). Blood level of Cairo traffic policemen. Science of the Total Environment, 105:165-170.

9.6 Welch (Unpooled Variance) t Tests and Confidence Intervals: An Example (10:13)

Welch (Unpooled Variance) t Tests and Confidence Intervals: An Example (10:13) (http://youtu.be/gzrmHpA54Sc)

I work through an example of a Welch (unpooled variance) t test and confidence interval (in the setting of inference for two means).

The antivenom/swelling data is found in:

Offerman et al. (2009). Subcutaneous crotaline Fab antivenom for the treatment of rattlesnake envenomation in a porcine model. Clinical Toxicology, 47: 61-68.

9.7 Pooled or Unpooled Variance t Tests and Confidence Intervals? (To Pool or not to Pool?) (11:52)

Pooled or Unpooled Variance t Tests and Confidence Intervals? (To Pool or not to Pool?) (11:52) (http://youtu.be/7GXnzQ2CX58)

I discuss the assumptions of both the pooled-variance and Welch (unpooled variance) t tests and confidence intervals, and their advantages and disadvantages. I illustrate some of pros and cons using the results of two simulations.

9.8 An Introduction to Paired-Difference Procedures (8:34)

An Introduction to Paired-Difference Procedures (8:34) (http://youtu.be/tZZt8f8URKg)

An introduction to paired-difference procedures. I briefly discuss how paired-difference scenarios arise, and briefly outline how we can use the paired-difference t procedure to construct confidence intervals and carry out hypothesis tests. (Paired-difference procedures are sometimes referred to as matched-pairs procedures, depending on the setting.)

The alcohol/reaction time data is loosely based on information in:

Anderson et al. (2011). Functional Imaging of Cognitive Control During Acute Alcohol Intoxication. Alcoholism: Clinical and Experimental Research, 35(1):156-165.

The identical twin-schizophrenia data is from:

Suddath et al. (1990). Anatomical abnormalities in the brains of monozygotic twins discordant for schizophrenia. New England Journal of Medicine, 322:789-794. Values used in this video are simulated values based on the summary statistics found in the paper.

9.9 An Example of a Paired-Difference t Test and Confidence Interval (12:06)

An Example of a Paired-Difference t Test and Confidence Interval (12:06) (http://youtu.be/upc4zN_-YFM)

An example of a paired-difference t test and confidence interval.

The data used in this video is from:

Penetar et al. (2012). The isoflavone puerarin reduces alcohol intake in heavy drinkers: A pilot study. Drug and Alcohol Dependence, 126:256-261. Values used in this video are simulated values based on the summary statistics found in the paper. (The summary statistics, test statistic, p-value, and overall conclusions are the same.)

9.10 Pooled-Variance t Procedures: Investigating the Normality Assumption (9:59)

Pooled-Variance t Procedures: Investigating the Normality Assumption (9:59) (http://youtu.be/zoJ5jK1V7Sc)

A discussion of the assumptions of pooled-variance t tests and confidence intervals for the difference in means. The assumptions are briefly discussed, and the effects of different violations of the normality assumption are investigated through simulation. The quick summary: Pooled-variance t procedures are more robust to violations of the normality assumption than their one-sample counterparts.

10 Inference for proportions

10.1 An Introduction to Inference for a Proportion (10:27)

An Introduction to Inference for a Proportion (10:27) (`http://youtu.be/owYtDtmrCoE`)

An introduction to inference procedures for a single proportion. I discuss confidence intervals and hypothesis testing methods for a single proportion, based on the normal approximation.

This video does not do an example of the calculations – I work through an example in another video.

10.2 The Sampling Distribution of the Sample Proportion $\hat{p}$ (9:49)

The Sampling Distribution of the Sample Proportion $\hat{p}$ (9:49) (`http://youtu.be/fuGwbG9_W1c`)

A discussion of the sampling distribution of the sample proportion. I discuss how the distribution of the sample proportion is related to the binomial distribution, discuss its mean and variance, and illustrate that the sample proportion is approximately normally distributed for large sample sizes.

10.3 Inference for a Proportion: An Example of a Confidence Interval and a Hypothesis Test (8:40)

Inference for a Proportion: An Example of a Confidence Interval and a Hypothesis Test (8:40) (http://youtu.be/M7fUzmSbXWI)

I work through an example of a confidence interval and a hypothesis test for a single proportion, using normal approximation methods (Z test and confidence interval).

The male birth rate data is from:
Koshy et al. (2010). Parental smoking and increased likelihood of female births. Annals of Human Biology, 37(6): 789–800.

10.4 Confidence Intervals for a Proportion: Determining the Minimum Sample Size (11:22)

Confidence Intervals for a Proportion: Determining the Minimum Sample Size (11:22) (http://youtu.be/mmgZI2G6ibI)

I discuss determining the minimum sample size required to achieve a given margin of error when estimating a population proportion. (For intervals based on the normal approximation.)

10.5 An Introduction to Inference for Two Proportions (15:10)

An Introduction to Inference for Two Proportions (15:10) (http://youtu.be/g0at6LpYvHc)

I discuss the sampling distribution of the difference in sample proportions, and confidence intervals and hypothesis tests for the difference in population proportions (using large sample Z procedures based on the normal approximation).

The male birth rate data is from:
Koshy et al. (2010). Parental smoking and increased likelihood of female births. Annals of Human Biology, 37(6): 789–800.

10.6 Inference for Two Proportions: An Example of a Confidence Interval and a Hypothesis Test (13:23)

Inference for Two Proportions: An Example of a Confidence Interval and a Hypothesis Test (13:23) (http://youtu.be/OIYkOiQX3fk)

I work through an example of a confidence interval and a hypothesis test for the difference in population proportions (based on the normal approximation).

The example involves an experiment investigating a possible effect of magnetic pulse on the ability of homing pigeons to navigate. The data used in this video is taken from:

Holland et al. (2013). A magnetic pulse does not affect homing pigeon navigation: a GPS tracking experiment. The Journal of Experimental Biology, 216: 2192-2200.

See the article for further details on the experiment.

11 Chi-square tests

11.1 Chi-square Tests for One-way tables (9:07)

Chi-square Tests for One-way Tables (9:07) (http://youtu.be/gkgyg-eR0TQ)

I introduce the chi-square test for one-way tables (sometimes called a goodness-of-fit test), and work through an example. The data used here is from a classic 1905 genetics experiment by William Bateson and Reginald Punnet.

11.2 Chi-square tests: Goodness of Fit for the Binomial Distribution (14:21)

Chi-square tests: Goodness of Fit for the Binomial Distribution (14:21) (http://youtu.be/O7wy6iBFdE8)

I work through an example of testing the null hypothesis that the data comes from a binomial distribution. I do this for two tests, one in which the probability of success is specified in the null hypothesis, and one where it is estimated from the data.

The Larry Bird free throw data based on information in:

Wardrop, R.L. (1995). Simpson's paradox and the hot hand in basketball. The American Statistician, 49 (1), 24-28.

11.3 Chi-square Tests for Two-way Tables (Chi-square Tests of Independence) (9:54)

Chi-square Tests of Independence (Chi-square Tests for Two-way Tables) (9:54) (http://youtu.be/L1QPBGoDmT0)

I introduce the chi-square test of independence and work through an example.

The binge drinking data is from:

Wechsler H, Lee JE, Kuo M, Lee H. (2000). College Binge Drinking in the 1990s: A Continuing Problem Ñ Results of the Harvard School of Public Health 1999 College Alcohol Study. Journal of American College Health. 2000. 48 (10): 199-210.

11.4 Chi-square tests for count data: Finding the p-value (5:14)

Chi-square tests for count data: Finding the p-value (updated and revised version) (5:14) (http://youtu.be/HwD7ekD5l0g)

I work through an example of finding the p-value for a chi-square test, using both the table and R.

12 Variances

12.1 An Introduction to Inference for One Variance (Assuming a Normally Distributed Population) (13:35)

An Introduction to Inference for One Variance (Assuming a Normally Distributed Population) (13:35) (`http://youtu.be/lyd4V8DFCjM`)

I discuss confidence intervals and hypothesis tests for a variance when sampling from a normally distributed population. I discuss the logic behind the procedures, discuss some characteristics of the sampling distribution of the sample variance, and give the appropriate formulas. I briefly discuss the results for an example problem, but I don't work through any of the calculations in this video. (I have another video in which I work through a complete example.)

The fat content in deep fried chicken sandwiches is based on information from the USDA Nutrient Database.

12.2 Inference for One Variance: An Example of a Confidence Interval and a Hypothesis Test (12:05)

Inference for One Variance: An Example of a Confidence Interval and a Hypothesis Test (12:05) (http://youtu.be/tsLGbpu_NPk)

I work through an example of a confidence interval and a hypothesis test for a variance (using chi-square methods that are appropriate when sampling from a normally distributed population).

The cereal data is from a sample of 15 bags of cereal I collected and weighed.

I have a video introduction to these procedures available at: http://youtu.be/lyd4V8DFCjM

12.3 Deriving a Confidence Interval for a Variance (Assuming a Normally Distributed Population) (4:17)

Deriving a Confidence Interval for a Variance (Assuming a Normally Distributed Population) (4:17) (http://youtu.be/q-cHZyOs5DQ)

I derive the appropriate formula for a confidence interval for a population variance (when we are sampling from a normally distributed population). I do not do any calculations or look at any examples in this video, I simply derive the appropriate confidence interval formula.

I work through an example at http://youtu.be/tsLGbpu_NPk

12.4 The Sampling Distribution of the Sample Variance (12:00)

The Sampling Distribution of the Sample Variance (12:00) (http://youtu.be/V4Rm4UQHij0)

A discussion of the sampling distribution of the sample variance. I begin by discussing the sampling distribution of the sample variance when sampling from a normally distributed population, and then illustrate, through simulation, the sampling distribution of the sample variance for a few other distributions.

12.5 Inference for a Variance: How Robust are These Procedures? (10:43)

Inference for a Variance: How Robust are These Procedures? (10:43)

(http://youtu.be/_7N35-RelI8)

A discussion of the effect of violations of the normality assumption on confidence intervals for a variance. The effects of different violations of the normality assumption are investigated through simulation. The quick summary: These procedures are very sensitive to violations of the normality assumption, and often perform very poorly when the normality assumption is violated.

12.6 An Introduction to Inference for the Ratio of Two Variances (16:00)

An Introduction to Inference for the Ratio of Two Variances (16:00) (http://youtu.be/kEnPOogexVY)

An introduction to confidence intervals and hypothesis tests for the ratio of two population variances (using F procedures based on the assumption of normally distributed populations). I introduce the methods, and take a quick look at an example and discuss the results.

I don't work through any calculations in this video, but I have another video in which I work through a complete example.

The summary statistics and distribution of the 2D:4D ratios were estimated from Figures 1, 2, and 3 in:

Özener et al. (2014). Inbreeding Is Associated with Lower 2D:4D Digit Ratio. American Journal of Human Biology. 26:183-188.

12.7 Inference for Two Variances: An Example of a Confidence Interval and a Hypothesis Test (10:01)

Inference for Two Variances: An Example of a Confidence Interval and a Hypothesis Test (10:01) (http://youtu.be/uJ8pLnGf-9Y)

I work through an example of a confidence interval and a hypothesis test for the ratio of population variances, using F procedures that are based on the assumption of normally distributed populations.

The example in this video involves tail lengths of male and female lizards. The summary statistics and distributions of the tail lengths were estimated from Table 1 and Figure 2 in:

Qu et al. (2011). Sexual dimorphism and female reproduction in two sympatric toad-headed lizards, *Phrynocephalus frontalis* and *P. versicolor* (Agamidae). Animal Biology. 61:139-151.

12.8 Deriving a Confidence Interval for the Ratio of Two Variances (4:29)

Deriving a Confidence Interval for the Ratio of Two Variances (4:29) (http://youtu.be/dx6-_d9CQcM)

I derive the appropriate formula for a confidence interval for the ratio of two population variances (when we are sampling from normally distributed populations). I do not do any calculations or look at any examples in this video, I simply derive the appropriate confidence interval formula.

12.9 The Sampling Distribution of the Ratio of Sample Variances (12:00)

The Sampling Distribution of the Ratio of Sample Variances (12:00) (http://youtu.be/lOez56i-yRk)

A discussion of the sampling distribution of the ratio of sample variances. I begin by discussing the sampling distribution of the ratio of sample variances when sampling from normally distributed populations, and then illustrate, through simulation, the sampling distribution of the ratio of sample variances for two other distributions.

12.10 Inference for the Ratio of Variances: How Robust are These Procedures? (9:34)

Inference for the Ratio of Variances: How Robust are These Procedures? (9:34) (http://youtu.be/4Hr56qUkohM)

A discussion of the effect of violations of the normality assumption on confidence intervals for the ratio of variances. The effects of different violations of the normality assumption are investigated through simulation. The quick summary: These procedures are very sensitive to violations of the normality assumption, and often perform very poorly when the normality assumption is violated.

This video is very similar in content and results to my video that investigates the effect of violations of the normality assumption on inference procedures for a single variance.

13 ANOVA

13.1 One-Way ANOVA: Introduction (5:44)

Introduction to One-Way ANOVA (5:44)
(http://youtu.be/QUQ6YppWCeg)

A brief introduction to one-way Analysis of Variance (ANOVA). I discuss the null and alternative hypotheses and conclusions of the test. I also illustrate the difference between and within group variance using a visual example.

In other related videos, I discuss the ANOVA formulas in detail and work through a real-world example.

13.2 One-Way ANOVA: The Formulas (9:06)

One-Way ANOVA: The Formulas (9:06)
(http://youtu.be/fFnOD7KBSbw)

I look at the calculation formulas and the meaning of the terms in the one-way ANOVA table.

In other related videos, I have a brief introduction to one-way ANOVA, and work through a real world example.

13.3 One-Way ANOVA: An Example (5:26)

A One-Way ANOVA Example (5:26)
(http://youtu.be/WUoVftXvjiQ)

A one-way ANOVA example.

The example and data used here is based on information and summary statistics from:

Grattan-Miscio, K. and Vogel-Sprott, M. (2005). Alcohol, intentional control, and inappropriate behavior: Regulation by caffeine or an incentive. Experimental and Clinical Psychopharmacology, 13:48-55.

13.4 One-Way ANOVA: LSD confidence intervals (8:38)

One-Way ANOVA: LSD confidence intervals (Updated and Revised Version) (8:38)
(http://youtu.be/kO8t_q-AXHE)

Finding significance in a one-way ANOVA F test means there is strong evidence that not all population means are equal. Which are different? We can explore that with a multiple comparison procedure. Here we look at the simplest one: Fisher's LSD confidence intervals.

13.5 One-Way ANOVA: Finding the p-value (4:52)

Finding the P-value in One-Way ANOVA (Updated and Revised Version) (4:52)
(http://youtu.be/XdZ7BRqznSA)

An example of finding the p-value in one-way ANOVA. The p-value (the area to the right of the F test statistic) is found using both the F table and the statistical software R.

14 Regression

14.1 Introduction to Simple Linear Regression (8:10)

Introduction to Simple Linear Regression(8:10)
(http://youtu.be/KsVBBJRb9TE)

An introduction to simple linear regression.

The pain-empathy data is estimated from a figure given in:

Singer et al. (2004). Empathy for pain involves the affective but not sensory components of pain. Science, 303:1157–1162.

The Janka hardness-density data is found in:

Hand, D.J., Daly, F. , Lunn, A.D., McConway, K., and Ostrowski, E., editors (1994). The Handbook of Small Data Sets. Chapman & Hall, London.

Original source: Williams, E.J. (1959). Regression Analysis. John Wiley & Sons, New York. Page 43, Table 3.7.

14.2 Simple Linear Regression: The Least Squares Regression Line (7:24)

Simple Linear Regression: The Least Squares Regression Line (7:24) (http://youtu.be/coQAAN4eY5s)

An introduction to the least squares regression line in simple linear regression.

The pain-empathy data is estimated from a figure given in:

Singer et al. (2004). Empathy for pain involves the affective but not sensory components of pain. Science, 303:1157-1162.

The Janka hardness-density data is found in:

Hand, D.J., Daly, F. , Lunn, A.D., McConway, K., and Ostrowski, E., editors (1994). The Handbook of Small Data Sets. Chapman & Hall, London.

Original source: Williams, E.J. (1959). Regression Analysis. John Wiley & Sons, New York. Page 43, Table 3.7.

14.3 Deriving the least squares estimators of the slope and intercept (simple linear regression) (12:13)

Deriving the least squares estimators of the slope and intercept (simple linear regression) (12:13) (https://youtu.be/ewnc1cXJmGA)

I derive the least squares estimators of the slope and intercept in simple linear regression (Using summation notation, and no matrices.) I assume that the viewer has already been introduced to the linear regression model, but I do provide a brief review in the few minutes. I assume that you have a basic knowledge of differential calculus, including the power rule and the chain rule.

If you are already familiar with the problem, and you are just looking for help with the mathematics of the derivation, the derivation starts at 3:26.

14.4 Simple Linear Regression: Interpreting Model Parameters (5:05)

Simple Linear Regression: Interpreting Model Parameters (Updated and Revised Version) (5:05) (http://youtu.be/I8Dr1OGUdZQ)

A look at the interpretation of the parameters $(\beta_0, \beta_1, \sigma^2)$, and their corresponding estimates in simple linear regression.

14.5 Simple Linear Regression: Assumptions (3:05)

Simple Linear Regression: Assumptions (Updated and Revised Version)(3:05)
(http://youtu.be/gHMTzdbpQTw)

A look at the assumptions on the epsilon term in our simple linear regression model.

14.6 Checking Assumptions with Residual Plots (8:04)

Simple Linear Regression: Checking Assumptions with Residual Plots (8:04)
(http://youtu.be/iMdtTCX2Q70)

An investigation of the normality, constant variance, and linearity assumptions of the simple linear regression model through residual plots.

The pain-empathy data is estimated from a figure given in:

Singer et al. (2004). Empathy for pain involves the affective but not sensory components of pain. Science, 303:1157–1162.

The Janka hardness-density data is found in:

Hand, D.J., Daly, F. , Lunn, A.D., McConway, K., and Ostrowski, E., editors (1994). The Handbook of Small Data Sets. Chapman & Hall, London.

Original source: Williams, E.J. (1959). Regression Analysis. John Wiley & Sons, New York. Page 43, Table 3.7.

14.7 Inference on the slope (the formulas) (6:57)

Inference on the Slope (The Formulas) (6:57)
(http://youtu.be/THzckPB7E8Q)

I show the formulas (and work through some of the underlying logic) for confidence intervals and hypothesis tests on the slope parameter in simple linear regression.

14.8 Inference on the Slope (An Example) (7:01)

Inference on the Slope (An Example) (7:01)
(http://youtu.be/nk_0RcHI-vo)

I work through an example, finding a confidence interval and carrying out a hypothesis test on the slope parameter.

The data used is estimated from a figure in:

Singer et al. (2004). Empathy for pain involves the affective but not sensory components of pain. Science, 303:1157–1162.

14.9 The Correlation Coefficient and Coefficient of Determination (11:08)

The Correlation Coefficient and the Coefficient of Determination (11:08) (http://youtu.be/cTolF3G5a1I)

A look at the Pearson correlation coefficient (r), the coefficient of determination (r 2), some of their properties and a few examples. (This is discussed in the context of measuring the strength of the linear relationship between 2 variables.)

One of the data sets used is estimated from a figure given in:

Singer et al. (2004). Empathy for pain involves the affective but not sensory components of pain. Science, 303:1157-1162.

14.10 Simple Linear Regression: An Example (9:51)

Simple Linear Regression: An Example (9:51)
(http://youtu.be/xIDjj6ZyFuw)

I work through an example relating eggshell thickness to DDT concentration, fitting the least squares line, using the line for prediction, interpreting the coefficient of determination, checking the residual plots and carrying out a test on the slope.

14.11 Simple Linear Regression: Always Plot Your Data! (5:25)

Simple Linear Regression: Always Plot Your Data! (5:25) (http://youtu.be/sfH43temzQY)

A quick look at a few cautions about simple linear regression.

Anscombe's quartet can be found in:

Anscombe, F. (1973). Graphs in statistical analysis. American Statistician, 27:17–21.

The fuel use vs speed data is found in:

West, B.H., McGill,R.N., Hodgson, J.W., Sluder, S.S. , and Smith D.E., Development and Verification of Light-Duty Modal Emissions and Fuel Consumption Values for Traffic Models, FHWA-RD-99-068, U.S. Department of Transportation, Federal Highway Administration, Washington, DC, March 1999.

The life expectancy vs television data is a random sample from the 2008 Economist Pocket World in Figures.

14.12 Simple Linear Regression: Transformations (7:27)

Simple Linear Regression: Transformations (7:27) (http://youtu.be/HIcqQhn3vSM)

A look at transformations in the context of simple linear regression. I look at two examples where taking a transformation (applying a function to the response and/or explanatory variables) can help to satisfy the assumptions of the simple linear regression model.

The brain and body weight data is from:

Sacher, G.A. Staffeldt, E. (1974). Relation of gestation time to brain weight for placental mammals: Implications for the theory of vertebrate growth. The American Naturalist, 108:593–615. The 99 observations represent one pair of measurements for each of the 99 species that had a brain and body weight measurement.

The hardness-density data is found in:

Hand, D.J., Daly, F. , Lunn, A.D., McConway, K., and Ostrowski, E., editors (1994). The Handbook of Small Data Sets. Chapman & Hall, London.

Original source: Williams, E.J. (1959). Regression Analysis. John Wiley & Sons, New York. Page 43, Table 3.7.

14.13 Intervals (for the Mean Response and a Single Response) in Simple Linear Regression (12:27)

Estimation and Prediction of the Response Variable in Simple Linear Regression (12:27) (http://youtu.be/V-sReSM887I)

I discuss confidence intervals for the mean of Y and prediction intervals for a single value of Y for a given value of X in simple linear regression. I discuss the form of the confidence interval for the mean and prediction interval for a single value and their standard errors.

The data used is estimated from a figure in:

Singer et al. (2004). Empathy for pain involves the affective but not sensory components of pain. Science, 303:1157–1162.

14.14 Leverage and Influential Points in Simple Linear Regression (7:15)

Leverage and Influential Points in Simple Linear Regression (7:15) (http://youtu.be/xc_X9GFVuVU)

A brief introduction to leverage and influence in simple linear regression. This video is about the basic concepts, and only briefly mentions numerical measures of leverage and influence at the end.

14.15 The Pooled-Variance t Test as a Regression (5:46)

The Pooled-Variance t Test as a Regression (5:46) (http://youtu.be/Gn78epv3jpo)

A look at the relationship between the pooled-variance t test and simple linear regression. I illustrate how the pooled-variance t test can be done as a regression, by declaring an appropriate indicator variable.

The Cairo traffic police officer data is simulated data with the same summary statistics as found in:

Kamal, A., Eldamaty, S., and Faris, R. (1991). Blood level of Cairo traffic policemen. Science of the Total Environment, 105:165Ð170.

15 Jimmy and Mr Snoothouse

15.1 Jimmy and Mr. S. discuss the t versus z issue in confidence intervals (1:03)

Jimmy and Mr. S. discuss the t versus z issue in confidence intervals (1:03) (http://youtu.be/QxcYJKETvD8)

15.2 Jimmy and Mr. Snoothouse discuss a hypothesis testing example (1:40)

Jimmy and Mr. Snoothouse discuss a hypothesis testing example (1:40) (http://youtu.be/6099XhoZgJ4)

15.3 Jimmy and Mr. Snoothouse discuss the interpretation of a confidence interval (2:09)

Jimmy and Mr. S discuss the interpretation of a confidence interval (2:09) (http://youtu.be/CF4Vdqiwhaw)

16 More advanced topics (beyond introductory statistics)

16.1 Deriving the mean and variance of the least squares slope estimator (beta_1 hat) in simple linear regression (10:54)

Deriving the mean and variance of the least squares slope estimator (beta_1 hat) in simple linear regression (10:54) (`https://youtu.be/rODUBTRUV0U`)

I derive the mean and variance of the sampling distribution of the slope estimator (beta_1 hat) in simple linear regression (in the fixed X case). I discuss the typical model assumptions, and discuss precisely where we use them as I carry out the derivations.

At the end, I briefly discuss the normality assumption, and how that leads to beta_1 hat being normally distributed. While I do discuss the real deal there, I go over it fairly quickly, as the main point of the video is deriving E(beta_1 hat) and Var(beta_1 hat).

Note that any time I use "errors" or "error terms" in this video, I am referring to the theoretical error terms (the epsilons) and not observed residuals from sample data.

Time stamps:

0:00 Brief discussion the simple linear regression model, assumptions, and some tools we will use. 2:58 Deriving E(beta_1 hat) 5:06 Deriving Var(beta_1 hat) 8:49 Discussion of normality of beta_1 hat.

16.2 In repeated sampling, on average what proportion of sample means would a randomly selected 95% CI for mu capture? (11:18)

In repeated sampling, on average what proportion of sample means would a randomly selected 95% CI for mu capture? (11:18) (https://youtu.be/GFT-2tY_6I0)

This is a bit of an unusual video for me, and if you're just looking for help with specific topics in a statistics course, you may not find it helpful. But I think it's well worth a watch.

Here I address what might seem at first like bit of a strange or uninformative question: In repeated sampling from a normally distributed population, on average what proportion of sample means would a randomly selected 95% CI for mu capture? I work through the calculations, then I discuss how this notion relates to reproducibility* studies.

This was inspired by a bad confidence interval interpretation that I heard a number of years ago (and have heard variants of ever since), where, when interpreting a 95% confidence interval for mu, the individual stated:

"If you repeat the same study a million times, then the mean of each one of those samples should fall in the interval 95% of the time." [Edited slightly to improve the readability.]

This is a poor interpretation of the interval, and simply untrue. It's just not the case. So, then, what is the probability a randomly selected 95% confidence interval for mu captures the mean of another sample of the same size from the same population?

This has applications in reproducibility* studies, and I briefly discuss that after working through the calculations. My discussion is not intended to be a complete discussion of issues in reproducibility, just a brief discussion of how the question I answer relates.

*In this video I use "reproducible" and "replicable" interchangeably. I know there has been much discussion in some circles of differences between those terms. Apologies f you find my casual use of these terms problematic or misleading.

Reference for the paper I bring up:

Open Science Collaboration. (2015). Estimating the reproducibility of psychological science. *Science*, 349(6251), 1–8.

Reference for a different discussion about how there is extra variability when comparing two statistics, and why the results found in the reference above might not be as bad as they appear at first blush:

Patil P., Peng R. D., Leek J. T. (2016). What should researchers expect when they replicate studies? A statistical view of replicability in psychological science. Perspect. Psychol. Sci. 11 539–544. 10.1177/1745691616646366

17 Retired Videos

The following videos have been replaced by upgraded and updated versions. They are not recommended, but they are still available on youtube and some people like them. They are included here for completeness.

Introduction to Discrete Random Variables and Discrete Probability Distributions (11:46) (http://youtu.be/0P5WRKihQ4E)

Expected Value and Variance of Discrete Random Variables (7:57) (http://youtu.be/OvTEhNL96v0)

Introduction to the Binomial Distribution (fast version) (8:02) (http://youtu.be/eSJ6ufTSJNk)

Introduction to the Poisson Distribution (fast version) (8:23) (http://youtu.be/8x3pnyYCBto)

The Hypergeometric Distribution: Introduction (fast version) (9:31) (http://youtu.be/BCeFgnh6A1U)

Discrete Probability Distributions: Some Examples (Binomial, Poisson, Hypergeometric, Geometric)(9:30) (http://youtu.be/Iu25wy7icok)

Standardizing Normally Distributed Random Variables (fast version) (6:38) (http://youtu.be/BN-2XOMnoCs)

The Sampling Distribution of the Sample Mean (fast version) (7:25) (http://youtu.be/0zqNGDVNKgA)

Deriving the Mean and Variance of the Sample Mean (fast version) (4:46) (http://youtu.be/JLmD0sJId1M)

An Introduction to the Chi-Square Distribution (4:10) (http://youtu.be/6Z4MBxye5VA)

An Introduction to the t Distribution (5:09) (http://youtu.be/TAQ8r1xUw0M)

Inference for One Proportion: Introduction (8:55) (http://youtu.be/G-tPLjTiSew)

Inference for One Proportion: An Example (5:50) (http://youtu.be/fVrSVYbrRzc)

Inference for One Proportion: Determining the Required Sample Size (6:47)
(http://youtu.be/PCL3g1RLUNU)

Inference for Two Proportions: Introduction (7:57) (http://youtu.be/hWYWHuH_zIw)

Inference for Two Proportions: An Example (9:15) (http://youtu.be/MTiZLnJU7FY)

Confidence Intervals for One Population Variance (9:57) (http://youtu.be/qwqB5a7_W44)

Hypothesis Tests for One Population Variance (8:52) (http://youtu.be/PweabcpqzYI)

Confidence Intervals for the Ratio of Population Variances (8:39) (http://youtu.be/64hFiLSq3Fg)

Hypothesis Tests for Equality of Two Variances (11:40) (http://youtu.be/7GJkxSYsX70)

One-Way ANOVA: Introduction (http://youtu.be/3MUxDrXyocc)

One-Way ANOVA: An Example)
(http://youtu.be/UjyVbuBRmY4)

Introduction to Simple Linear Regression
(http://youtu.be/LLeOVxODSls)

Simple Linear Regression: The Least Squares Regression Line
(http://youtu.be/VvQ-QkU8DPs)

Simple Linear Regression: Interpreting Model Parameters
(http://youtu.be/hhHz2xIVz4w)

Simple Linear Regression: Assumptions)
(http://youtu.be/UrD4n5cCL-M)

Inference on the slope (the formulas)
(http://youtu.be/8FwO89CXpsQ)

Simple Linear Regression: Checking assumptions with residual plots
(http://youtu.be/_NkWG1M69CI)

Inference on the slope (an example)
(http://youtu.be/TFkdFV_MgXs)

The Correlation Coefficient and the Coefficient of Determination (In Simple Linear Regression)
(http://youtu.be/dCgIavyFWIo)

Simple Linear Regression: An Example
(http://youtu.be/gJCFk-s80iI)

Simple Linear Regression: Some cautions
(http://youtu.be/aVsLT-CZZE0)

One-Way ANOVA: The Calculations
(http://youtu.be/AM6l5G20wGM)