PSYC2001 – Research Methods 2

LECTURE 1 – INTRODUCTORY STATISTICS

Why do I have to learn Statistics?

- Psychology is a Science
- Science is a data-driven or evidence-based
- Statistics help you know what evidence is, and what is noise

Statistics is needed because measurement is noisy. We need to know if there is any signal among this noise.

Psychology is particularly hard. For example, if we want to know the effect of eating food on weight. Psychology is like measuring this with an individual eating some amount of food and then measure the displacement (measure how much the water rises in an active hot tub). The lack of direct observation demonstrates how hard psychology is.

Statistics involves determining whether there is any signal amongst different data. If there is a significant difference between different conditions.

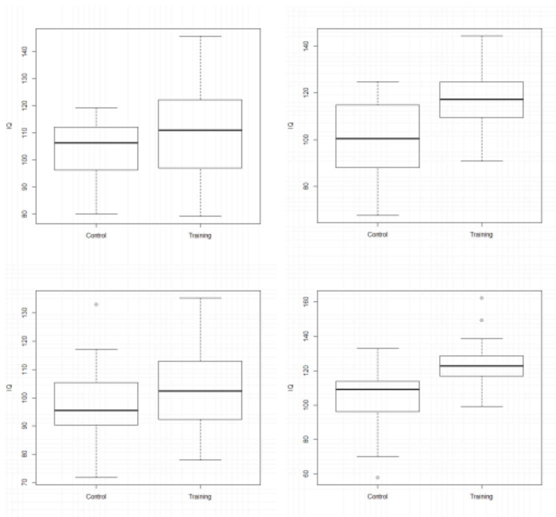

Revision of Methodology

- ***Independent Variable IV:*** Variable that is systematically varied (manipulated) by researcher
  - E.g. depression study: IV – type of treatment.
    Levels: Group 1 New Treatment
    Group 2 Standard Treatment
    Group 3: Waiting-list control
- ***Dependent Variable:*** variable that is measured or observed
  - E.g. depression study: DV – score on depression scale
- Question being addresses: does the IV affect (cause) the DV?
  - Depression study: does the treatment affect patient outcome (is one treatment better than another?)
  - Is the amygdala necessary for fear learning?
  - DO tinted glasses improve reading?

- ***Extraneous Variables***: any variable that is not the IV, that may affect the DV. E.g. depression study: personality of participants, demand characteristics, recent life events
  - Extraneous Variables are a threat to valid inference about the effect of IV on the DV ("internal validity")
  - Two negative effects of extraneous variables:
    - Make it harder to see the effects of the IV by adding variability ("error" or "noise")
    - If correlated with the IV ("confounding" variables), they make it impossible to determine the true effect of the IV on the DV
  - ways to control extraneous variables:
    - eliminate – e.g. outside noise
    - hold constant – e.g. temperature
    - measure and include in analysis – e.g. personality
    - randomise allocation to groups

Random Allocation

- In a between-group design, random allocation of participants to groups is the only method that controls for the thousands of extraneous variables affecting participants that the researcher doesn't know about or measure
- Without random allocation, the researcher can never be sure whether differences observed on the DV might have been due to a confounding "individual difference" variable – i.e., something about the participants rather than the IV.
- E.g. Depression Study:
  - Participants allowed to choose treatments: those who choose "easy" treatment might be less motivated
  - patients from one clinic given the standard treatment, patients from another clinic given the new treatment: clinics might differ in terms of patient factors such as severity, socio-economic status, diet etc. etc.
  - these are examples of confounding individual difference variables
- Random allocation "unconfounds" these variables by ensuring there are no systematic differences between groups prior to treatment
- It doesn't eliminate extraneous variables altogether – they still add noise to the measurement of the DV, but his effect can be handled statistically

Types of Research Designs

Research design	Causal inference
True experiment Controlled manipulation of IV Random allocation of participants to levels of IV	Strong cause-and-effect statement
Quasi-experiment Controlled manipulation of IV No random allocation	Moderate/weak
Correlational/observational study No manipulation of variables No random allocation Naturally occurring variation	Weak
Descriptive research **(surveys, descriptive studies, case-studies)** No manipulation of variables Intention is to observe and describe	None

REVISION OF STATISTICS

- ***Descriptive Statistics*** – facts and figure that describe populations or samples. Its purpose is to organise and summarise observations.
- ***Inferential Statistics*** – statistical procedures and their outcomes. Its purpose is to draw an inference (conclusion) about characteristics of a *population* from the characteristics of a sample.
- Example: what percentage of the general population is left handed?
  1. Select a representative sample from a population and ask each participant about handedness
  2. Descriptive statistics: calculate percentage left-handed in sample
  3. Inferential statistics: estimate population percentage (+/- margin of error)

Descriptive Statistics

	Population (parameters)	Sample (statistics)
Mean	μ	M
Variance	σ^2	S^2
Standard deviation	σ	S

Mean	$M = \frac{\sum X}{n}$
Variance	$S^2 = \frac{\sum (X-M)^2}{n}$
Standard deviation	$S = \sqrt{S^2} = \sqrt{\frac{\sum (X-M)^2}{n}}$

Comparing individual performance:

Standard score	$Z = \frac{X-M}{S}$

Normal Distribution

Many Variables have a symmetrical, bell-shaped distribution called a normal distribution. Area under curve corresponds to proportion or probability.

- Example: IQ
  $\mu=100, \sigma=15$

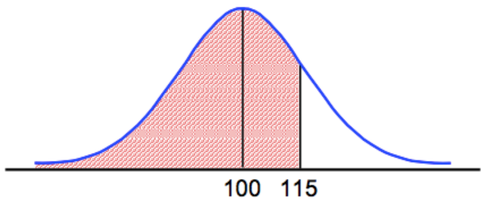

- Total area under curve = 1
- Tables of area under the normal distribution allow us to ask questions like:
  - what proportion of the population have an IQ < 80?
  - what is the probability that a person chosen at random has an IQ > 115?

- Revise first year methodology and statistics if necessary

LECTURE 2 – INFERENTIAL STATISTICS

Population and Sample

- The ***population*** is a large group of observations (people, objects) about which the researcher wants to draw conclusions
- Populations can by hypothetical. E.g. populations of all patients (past, present and future) who receive a particular treatment.
- Populations are usually large and unobtainable, so parameters cannot be directly measured.
- The ***Sample*** is a subset (small group) of the population. It should be representative of the population – this can be done from random sampling of subjects from the population.

Statistics is a way of making generalised conclusions which apply to a population from sample data. Statistics is needed to understand how much you can claim about truth given what your sample can tell you.

Parameters and Statistics

- Inferential *statistical* procedures use sample statistics to make inferences about population *parameters.*
- The sample must be representative of the population
- The can be achieved with random sampling. Random sampling means that anyone has an equal chance of being part of the sample. This can be contrasted with biased sampling.
  - Selection of any observation from the population is independent of the selection of any other observation
  - Every element in the population has an equal chance of being selected
  - Contrasted with *biased* sampling, e.g. only including people who have a telephone

$$\overline{x} = \frac{1}{n}\sum_{i=1}^{n} x_i$$

$$s = \sqrt{\frac{\sum (X - M)^2}{n - 1}}$$

Inferential Statistics

- Aim is to draw an inference about a population based on information obtained from a sample drawn from that population.
- Inferential statistics are based on descriptive statistics and can be used to:
  - Estimate populations parameters (find the true value of something, properties or ways of describing the population); or ESTIMATION
  - Decide whether to accept or reject a statement about a parameter. E.g. is the difference in the average effectiveness zero that is, is the effectiveness equal? You can test these types of statements regarding population parameters. HYPOTHESIS TESTING
- Inferential statistical procedures provide a way of determining whether a sample outcome (e.g. the results of an experiment) has or has not occurred by chance.

Estimation and Hypothesis Testing

Two inferential procedures include:

- Estimation – estimation of a population parameter (e.g. mean) through construction of *confidence interval.*
- Hypothesis testing – decision-making process regarding a statement about a population parameters.

Example: How effective is a particular treatment for anxiety?

- run an experiment – randomly allocate participants to either Group 1 (Treatment) or Group 2 (Control)
- measure all participants on DV (anxiety score)
- suppose we find: M1 = 20 and M2 = 22 ie, M1 – M2 = -2. We can see that in Group 1 the treatment was 2 points more effective than the control group.

What can we say about this result?

- Option 1: Treatment is effective i.e. difference of 2 points represents a *real* difference
- Option 2: Treatment is not effective - difference of 2 points has occurred by *chance*.

We need a way of distinguishing between these two possibilities i.e., What is the chance of obtaining a difference as large as 2 points if the treatment is not effective? To answer this, we need to look at *sampling variability*

Sampling Variability

- The value of a statistic will *vary* from sample to sample, due to chance
- E.g. large population where population mean $\mu = 30$
  - Randomly select a sample of 10 cases (n=10) and calculate the sample mean M

- The value of M is likely to be somewhere near 30, but is unlikely to be exactly 30. E.g., it might be 28 or 33. Occasionally, it might be further away, e.g. 41 or 12

- The same principle applies to other statistics (e.g. *s*) that estimate a population parameter (σ)

How can we describe this variability from sample to sample? E.g. if we are using M to estimate μ, how can we quantify the margin of error in our estimate?

Imagine a population that has an average of 30. We know the true value of the mean of the population. We randomly draw 10 numbers from this sample and get an average of 28.7. every time we do this we get a slightly different number. We are going to call this the sampling distribution of the mean.

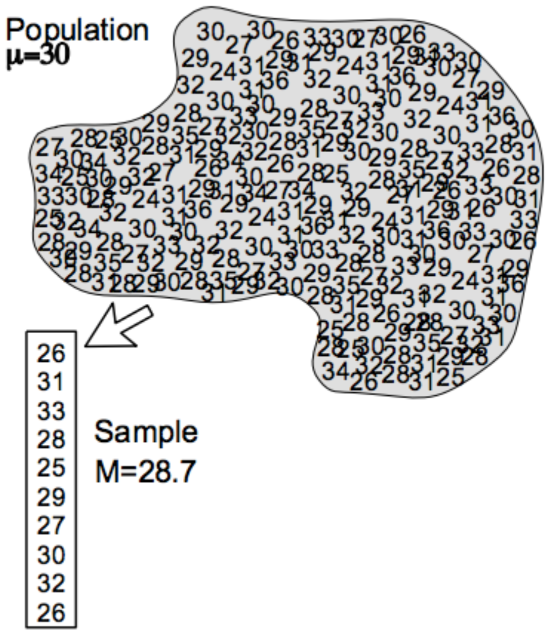

The sampling distribution is a hypothetical distribution of values of a particular sample statistic, formed by repeatedly drawing samples of n observations from a population and calculating the value of the statistic for each sample
in this course, we will focus on the ***sampling distribution of the mean***

Imagine drawing a very large number of samples of size n=10, and calculating the sample mean M for each one. You will end up with a long list of values of M. Now imagine forming a frequency distribution from this list.
The sampling distribution is the frequency distribution you would get if you used an infinite number of samples (i.e it is hypothetical).
You will always get a bell-shaped distribution. This has important characteristics:

- Its average is important as it tells us what we should expect to get out of each sample of 10
- The variance: it tells us how good our experiments are. The larger the experiment the more worse the experiment and vice versa.

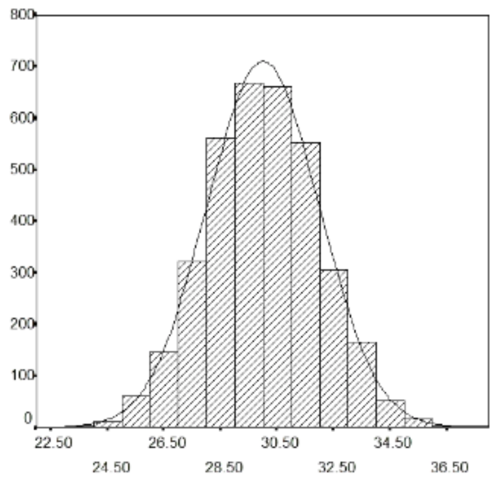

Figure 1: Frequency Distribution of Sample Means

If we get a wide distribution, it tells us that our experiments are going to be very bad. They aren't going to be an accurate and precise. Vice Versa.

- *mean* is μ (ie M is an unbiased estimator of μ)
- *variance* is $\sigma_M^2 = \frac{\sigma^2}{n}$
- *standard deviation* is $\sigma_M = \frac{\sigma}{\sqrt{n}}$

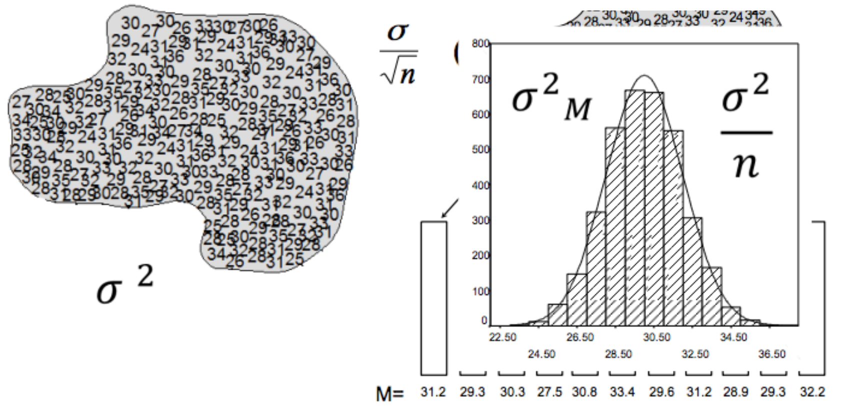

The average of this distribution is the true population μ. What this means your average your experiment is true. If you keep on getting samples from the population, on average the sample mean you get is the unbiased truth. The sample mean is an unbiased estimator of μ.

The variance (how wides is the sample distribution) an indicator of how good our experiment is. It is based on sigma squared and the number of samples you draw. The basic idea is that how variable the population is influences how variable the sample means are. As the size if your sample increases, variance decreases. The average of more things will be more true or closer to the truth compared to the average of les things.

If the *sample distribution of the mean (variance)* is large we are less likely to believe that our result is true. That s, it is more likely due to chance. If the sample distribution of the mean is small, it is more accurate and precise we are more likely to believe that our result was true.

Standard deviation (standard error of the mean).

Example: Effect of reducing standard deviation (sigma) we have a population of normal distribution. Mean of 100 and standard deviation of 20. We are going to take samples of 25 and calculate average. Repeat this again and again. Standard deviations of the mean is smaller than the distribution of individual scores.

As the variability (sigma = 10, sigma = 5) decreases the standard deviation of the mean variance also decreases.

Effect of σ on σ_M

μ=100
σ=20
n=25

μ=100
σ=10
n=25

μ=100
σ=5
n=25

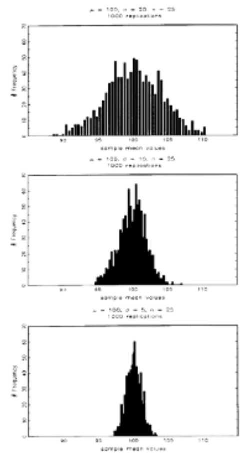

Example 2: Effect of reducing the number of participants in each sample (n). Run experiments with larger samples to get values closer to the true population mean.

Effect of n on σ_M

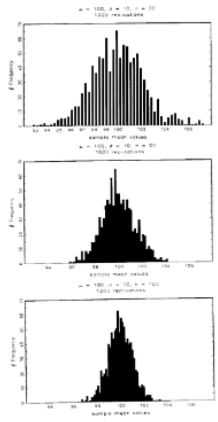

Shape of the Sampling Distribution

- If the population of scores is normally distributed, the sampling distribution of the mean will also be normally distributed

What if the population is non-normal?
Central Limit Theorem: sampling distribution of mean tends towards a normal distribution as n increases, regardless of the shape of the population distribution

- As long as n is reasonably large, we can assume the sampling distribution is normal

Critical information, since we know a lot about the normal distribution