

Stat 13, Intro. to Statistical Methods for the Life and Health Sciences.

1. CIs for a single mean, used car example.
2. Factors affecting CI width, and CI meaning.
3. When to use which multiplier.
4. Statistical and practical significance, and longevity example.
5. Observational studies: association, confounding, and nightlights example.
6. Observational studies and experiments.
7. Experiments and aspirin example.
8. Random sampling and random assignment.
9. Blinding.
10. Portacaval shunt example.

<http://www.stat.ucla.edu/~frederic/13/sum24> .

Remember to submit hw2 by Aug21 at 10pm.

hw3 is due Wed Aug28, 10pm. 4.CE.10, 5.3.28, 6.1.17, and 6.3.14. In 5.3.28d, use the theory-based formula. You do not need to use an applet.

Midterm is Thu Aug29, 10am-11:50am.

Spanking and IQ

4.CE.10 Studies have shown that children in the U.S. who have been spanked have a significantly lower IQ score on average than children who have not been spanked.

- a. Is it legitimate to conclude from this study that spanking a child causes a lower IQ score? Explain why or why not.
- b. Explain why conducting a randomized experiment to investigate this issue (of whether spanking causes lower IQs) would be possible in principle but ethically objectionable.

5.3.28 Recall the data from the Physicians' Health Study: Of the 11,034 physicians who took the placebo, 138 developed ulcers during the study. Of the 11,037 physicians who took aspirin, 169 developed ulcers.

- Define the parameters of interest. Assign symbols to these parameters.
- State the appropriate null and alternative hypotheses in symbols.
- Explain why it would be okay to use the theory-based method (that is, normal distribution based method) to find a confidence interval for this study.
- Use an appropriate applet to find and report the theory-based 95% confidence interval.
- Does the 95% confidence interval contain 0? Were you expecting this? Explain your reasoning.
- Interpret the 95% confidence interval in the context of the study.
- Use the 95% confidence interval to state a conclusion about the strength of evidence in the context of the study.
- Relatively speaking, is the 95% confidence interval narrow or wide? Explain why that makes sense.

5.3.29 Recall the data from the Physicians' Health Study:


every week. Be sure to compare and contrast the shape, center, and spread for study hours' distributions for males and females.

6.1.16 Reconsider the data in the previous question about number of hours spent studying.

- Find the median number of study hours for both males and females. What do these numbers tell us about the two data sets?
- Find the inter-quartile range for the number of study hours for both males and females. What do these numbers tell us about the two data sets?
- Construct parallel boxplots by hand for the two data sets.

Gettysburg Address

6.1.17 The graph below displays the distribution of word lengths (number of letters) in the Gettysburg Address, which you explored in Exploration 2.1A.



- Describe the shape of this distribution.
- Based on this shape, do you expect the median to be less than the mean, greater than the mean, or very close to the mean? Explain.

The following table lists how often each of the word lengths appears for these 268 words.

Word length	1	2	3	4	5	6	7	8	9	10	11
Number of words	7	49	54	59	34	27	15	6	10	4	3

- Determine the median word length of these 268 words.
- The mean word length is 4.29 letters per word. Is the median greater than, less than, or very close to the mean? Does this confirm your answer to part (b)?
- Calculate the five-number summary of the word lengths.

College student bedtimes*

6.1.18 In a survey, 30 college students were asked what their usual bedtime was and the results are shown in the 6.1.18 dotplot in terms of hours after midnight. Negative responses are hours before midnight.

- Determine the five-number summary for the bed times.
- What is the inter-quartile range?
- The earliest bedtime is 11:30 PM (represented by -0.50 on the graph). If that person's usual bedtime is actually 9:00 PM and that change was made in the dotplot, does that change the inter-quartile range? Would it change the standard deviation?


Candy bars

6.1.19 Weights of 20 Mounds® candy bars and 20 PayDay® candy bars, in grams, are shown in the 6.1.19 dotplots.

- Describe how the distributions of weights of the two types of candy bars differ in both variability and center.
- Based on your answers to part (a), which set of candy bar weights has the lowest standard deviation? Which has the lowest mean?
- Would you say there is an association between the type of candy bar and the weight? Why or why not?



EXERCISE 6.1.18



EXERCISE 6.1.19

- h. Summarize your conclusions about the research question of the study. Be sure to comment on statistical significance, confidence/estimation, causation, and generalization.

Perceived wealth

6.3.13 Do people tend to spend money differently based on perceived changes in wealth? In a study conducted by Epley et al. (2006), 47 Harvard undergraduates were randomly assigned to receive either a "bonus" check of \$50 or a "rebate" check of \$50. A week later, each student was contacted and asked whether they had spent any of that money, and if yes, how much. In this exercise we will focus on how much money they recalled spending when contacted a week later. It turned out that those in the "bonus" group spent an average of about \$22, compared to \$10 in the "rebate" group.

- Identify the observational units.
- Identify the explanatory and response variables. Identify each as either categorical or quantitative.
- State the appropriate null and alternative hypotheses in the context of the study.
- In the article that appeared in the *Journal of Behavioral Decision Making*, the researchers reported neither the sample size nor the sample SD of each group. In this exercise you will explore whether and how the strength of evidence is impacted by the sample size and sample SD. Complete the following table by finding the t -statistic and a p -value for a theory-based test of significance comparing two means under each of the four different scenarios.
- Summarize what your analysis has revealed about the effects of the sample size breakdown and the sample standard deviations on the values of the t -statistic and p -value.

Nostril breathing and cognitive performance*

6.3.14 In an article titled "Unilateral Nostril Breathing Influences Lateralized Cognitive Performance" that appeared in *Brain and Cognition* (1989), researchers Block

et al. published results from an experiment involving assessments of spatial and verbal cognition when breathing through only the right versus left nostril.

The subjects were 30 male and 30 female right-handed introductory psychology students who volunteered to participate in exchange for course credit. Initial testing on spatial and verbal tests revealed the following summary statistics. Note that the scores on the spatial task can range from 0 to 40, whereas those on the verbal task can go from 0 to 20. The distributions are not strongly skewed on either scale or for males or females.

Sex	Spatial		Verbal	
	Mean	SD	Mean	SD
Male	10.20	2.70	10.90	3.00
Female	7.80	2.50	15.10	3.40

- Consider comparing males to females with regard to performance on the spatial assessment task. State the appropriate null and alternative hypotheses in the context of the study.
- Explain why it is valid to use the theory-based method for producing a p -value to test the hypotheses stated in part (a).
- Carry out the appropriate test to produce a p -value to test the hypotheses stated in part (a) and interpret the p -value.
- Find a 95% confidence interval for the difference in mean scores of males and females with regard to performance on spatial assessments. Interpret the interval.
- Based on your p -value, state a conclusion in the context of the study. Be sure to comment on statistical significance, estimation (confidence interval), causation, and generalization.
- Repeat the investigation comparing males and females, this time on verbal performance. Be sure to address the questions asked in parts (a)–(e).

Scenario		Sample sizes	Sample means	Sample SDs	t -statistic	p -value
1	Bonus	24	22	5		
	Rebate	23	10	5		
2	Bonus	24	22	10		
	Rebate	23	10	10		
3	Bonus	30	22	5		
	Rebate	17	10	5		
4	Bonus	30	22	10		
	Rebate	17	10	10		


EXERCISE 6.3.13

1. Used Cars

Example 3.3


Used Cars

The following histogram displays data for the selling price of 102 Honda Civics that were listed for sale on the Internet in July 2006.



Used Cars

- The average of this sample is $\bar{x} = \$13,292$ with a standard deviation of $s = \$4,535$.
- What can we say about μ , the average price of all used Honda Civics?



Used Cars

- While we should be cautious about our sample being representative of the population, let's treat it as such.
- μ might not equal \$13,292 (the sample mean), but it should be close.
- To determine how close, we can construct a confidence interval.

Confidence Intervals

- Remember the basic form of a confidence interval is:

$$\text{statistic} \pm \text{multiplier} \times (\text{SD of statistic})$$

SD of statistic is also called Standard Error (SE).

- In our case, the statistic is \bar{x} and for a 95% CI our multiplier is 1.96, so we write our 1.96SE confidence interval as:

$$\bar{x} \pm 1.96(\text{SE})$$

Confidence Intervals

- It is important to note that the SE, which is the SD of \bar{x} , is not the same as the SD of our sample, $s = \$4,535$.
- There is more variability in the data (the car-to-car variability) than in sample means.
- The SE is s/\sqrt{n} . Which means we can write a 1.96SE confidence interval as:


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

Summary Statistics

- In many circumstances we will use a multiplier that is based on a t -distribution, instead of 1.96. The t multiplier is dependent on the sample size and confidence level.
- For a theory-based confidence interval for a population mean (called a one-sample t -interval) to be valid, the observations should be approximately iid (independent and identically distributed), and either the population should be normal or n should be large. Check the sample distribution for skew and asymmetry.

Confidence Intervals

- We find our 95% CI for the mean price of all used Honda Civics is from \$12,401.20 to \$14,182.80.
- Notice that this is a much narrower range than the prices of all used Civics.
- For a 99% confidence interval, it would be wider. The multiplier would be 2.58 instead of 1.96.



Factors that Affect the Width of a Confidence Interval

Section 3.4

Factors Affecting Confidence Interval Widths


- **Level of confidence** (e.g., 90% vs. 95%)
 - As we increase the confidence level, we increase the width of the interval.
- **Sample size**
 - As sample size increases, variability decreases and hence the standard error will be smaller. This will result in a narrower interval.
- **Sample standard deviation**
 - A larger standard deviation, s , will yield a wider interval.
 - For sample proportions, wider intervals when \hat{p} is closer to 0.5. $s = \sqrt{[\hat{p} (1-\hat{p})]}$.

Level of Confidence


- If we have a wider interval, we should be more confident that we have captured the population proportion or population mean.
- We could see this with repeated tests of significance.
 - A higher confidence level corresponds to a lower significance level, and one must go farther to the left and farther to the right in our tables to get our confidence interval.

Sample Size


- We know as sample size increases, the variability (and thus standard deviation) in our null distribution decreases



$n = 90$ (SD = 0.054)



$n = 361$ (SD = 0.026)



$n = 1444$ (SD = 0.013)

Sample size	90	361	1444
SD of null distr.	0.053	0.027	0.013
Margin of error	$2 \times \text{SD} = 0.106$	$2 \times \text{SD} = 0.054$	$2 \times \text{SD} = 0.026$
Confidence interval	(0.091, 0.303)	(0.143, 0.251)	(0.171, 0.223)

Sample Size




- (With everything else staying the same) increasing the sample size will make a confidence interval narrower.

Notice:

- The observed sample proportion is the midpoint. (that won't change)
- The margin of error is a multiple of the standard deviation so as the standard deviation decreases, so will the margin of error.

Value of \hat{p} (or the value used for π under the null)

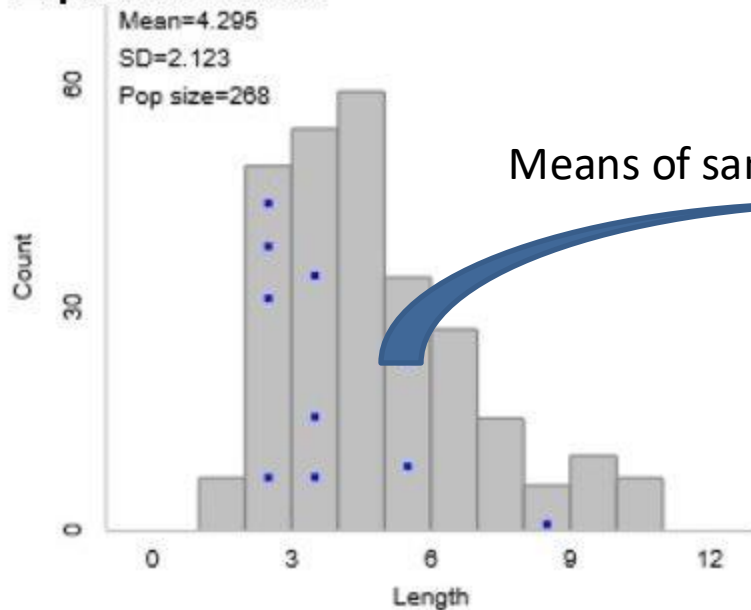
- As the value that is used under the null gets farther away from 0.5, the standard error decreases.
- When this standard error is used in the 1.96 SE method, the interval gets gradually narrower.




Standard Deviation

- Suppose we are taking repeated samples of a population.
- How do we estimate what the standard error (standard deviation of the null distribution) will be? s/\sqrt{n} .

Population data:



Means of samples of size 10.



Standard Deviation

- The SE (SD of the null distribution) is approximated by s/\sqrt{n} .
- Remember that $1.96 s/\sqrt{n}$ is the **margin** of error for a 95% confidence interval, so as the standard deviation of the sample data, s , increases so does the width of the confidence interval.
- Intuitively this should make sense, more variability in the data should be reflected by a wider confidence interval.

Formulas for Theory-Based Confidence Intervals

$$\hat{p} \pm multiplier \times \sqrt{\frac{\hat{p}(1-\hat{p})}{n}} \quad \bar{x} \pm multiplier \times \frac{s}{\sqrt{n}}$$

- The width of the confidence interval increases as level of confidence increases (multiplier)
- The width of the confidence interval decreases as the sample size increases
- The value \hat{p} also has a more subtle effect. The farther it is from 0.5 the smaller the width.
- The width of the confidence interval increases as the sample standard deviation increases.

What does 95% confidence mean?

- If we repeatedly sampled from a population and constructed 95% confidence intervals, 95% of our intervals will contain the population parameter.
- Notice the interval is the random event here.

What does 95% confidence mean?

- Suppose a 95% confidence interval for a mean is 2.5 to 4.3. We would say we are 95% confident that the population mean is between 2.5 and 4.3.
 - Does that mean that 95% of the data fall between 2.5 and 4.3?
 - No
 - Does that mean that in repeated sampling, 95% of the sample means will fall between 2.5 and 4.3?
 - No
 - Does that mean that there is a 95% chance the population mean is between 2.5 and 4.3?
 - Not quite but close.

What does 95% confidence mean?

- What does it mean when we say we are 95% confident that the population mean is between 2.5 and 4.3?
 - It means that if we repeated this process (taking random samples of the same size from the same population and computing 95% confidence intervals for the population mean) repeatedly, 95% of the confidence intervals we find would contain the population mean.
 - $P(\text{confidence interval contains } \mu) = 95\%$.

3. For CIs, when to use 1.96 from the normal,
& when to use a multiplier based on the t distribution.

iid = independent and identically distributed.

if the observations are iid. and n is large, then

$$P(\mu \text{ is in the range } \bar{x} \pm 1.96 \sigma/\sqrt{n}) \sim 95\%.$$

If the observations are iid and normal, and σ is known, then

$$P(\mu \text{ is in the range } \bar{x} \pm 1.96 \sigma/\sqrt{n}) \sim 95\%.$$

If the obs. are iid and normal and σ is unknown, then

$$P(\mu \text{ is in the range } \bar{x} \pm t_{\text{mult}} s/\sqrt{n}) \sim 95\%.$$

where t_{mult} is the multiplier from the t distribution.

This multiplier depends on n .

4. Statistical and Practical significance.

- *Statistically significant* means that the results are unlikely to happen by chance alone.
- *Practically important* means that the difference is large enough to matter in the real world.

Cautions

- Practical importance is context dependent and somewhat subjective.
- Well designed studies try to equate statistical significance with practical importance, but not always.
- Look at the sample size.
 - If very large, expect significant results.
 - If very small, don't expect significant results. (A lot of missed opportunities---type II errors.)

Longevity example.

According to data from the WHO (2014) and World Cancer Report (2014), the average number of cigarettes smoked per adult per day in the U.S. is 2.967, and in Latvia it is 2.853.

The sample sizes are huge, so even this little difference is stat. sig. (In the U.S., the National Health Interview Survey has $n > 87000$).

If you do not like cigarette smoke around you, should you move to Latvia?

The difference is statistically significant, but not practically significant for most purposes.

Causation.

Chapter 4

Big Idea of Chapter 4

- Previously research questions focused on **one** proportion
 - What proportion of the time did Marine choose the right bag?
- We will now start to focus on research questions comparing **two** groups.
 - Are smokers more likely than nonsmokers to have lung cancer?
 - Are children who used night lights as infants more likely to need glasses than those who didn't use night lights?

Big Idea of Chapter 4

- Typically we observe two groups and we also have two variables (like smoking and lung cancer).
- So with these comparisons, we will:
 - determine when there is an association between our two variables.
 - discuss when we can conclude the outcome of one variable causes a change in the other.

Observational studies and confounding.

Types of Variables

- When two variables are involved in a study, they are often classified as explanatory and response
- **Explanatory variable** (Independent, Predictor)
 - The variable we think may be causing or explaining or used to predict a change in the response variable. (Often this is the variable the researchers are manipulating.)
- **Response variable** (Dependent)
 - The variable we think may be being impacted or changed by the explanatory variable.
 - The one we are interested in predicting.

Roles of Variables

- Choose the explanatory and response variable:
 - Smoking and lung cancer
 - Heart disease and diet
 - Hair color and eye color
- Sometimes there is a clear distinction between explanatory and response variables and sometimes there isn't.

Observational Studies

- In observational studies, researchers *observe* and measure the explanatory variable but do not set its value for each subject.
- Examples:
 - A significantly higher proportion of individuals with lung cancer smoked compared to same-age individuals who don't have lung cancer.
 - College students who spend more time on Facebook tend to have lower GPAs.

Do these studies prove that smoking *causes* lung cancer or Facebook *causes* lower GPAs?

Night Lights and Nearsightedness

Example 4.1

Nightlights and Near-Sightedness

- Near-sightedness often develops in childhood
- Recent studies looked to see if there is an association between near-sightedness and night light use with infants
- Researchers interviewed parents of 479 children who were outpatients in a pediatric ophthalmology clinic
- Asked whether the child slept with the room light on, with a night light on, or in darkness before age 2
- Children were also separated into two groups: near-sighted or not near-sighted based on the child's recent eye examination

Night-lights and near-sightedness

	Darkness	Night Light	Room Light	Total
Near-sighted	18	78	41	137
Not near-sighted	154	154	34	342
Total	172	232	75	479

The largest group of near-sighted kids slept in rooms with night lights. It might be better to look at the data in terms of proportions.

Conditional proportions

$$18/172 \approx 0.105 \quad 78/232 \approx 0.336 \quad 41/75 \approx 0.547$$

Night lights and near-sightedness

	Darkness	Night Light	Room Light	Total
Near-sighted	10.5% 18/172	33.6% 78/232	54.7% 41/75	137
Not near-sighted	154	154	34	342
Total	172	232	75	479

- Notice that as the light level increases, the percentage of near-sighted children also increases.
- We say there is an **association** between near-sightedness and night lights.
- Two variables are **associated** if the values of one variable provide information (help you predict) the values of the other variable.

Night lights and near-sightedness

- While there is an association between the lighting condition and nearsightedness, can we claim that night lights and room lights *caused* the increase in near-sightedness?
- Might there be other reasons for this association?

Night lights and near-sightedness

- Could parents' eyesight be another explanation?
 - Maybe parents with poor eyesight tend to use more light to make it easier to navigate the room at night and parents with poor eyesight also tend to have children with poor eyesight.
 - Now we have a third variable of *parents' eyesight*
 - *Parents' eyesight* is considered a **confounding variable**.
 - Other possible confounders? Wealth? Books? Computers?

Confounding Variables

- A **confounding variable** is associated with both the explanatory variable and the response variable.
- We say it is confounding because its effects on the response cannot be separated from those of the explanatory variable.
- Because of this, we can't draw cause and effect conclusions when confounding variables are present.

Confounding Variables

- Since confounding variables can be present in observational studies, we can't conclude causation from these kinds of studies.
- This doesn't mean the explanatory variable isn't influencing the response variable. **Association may not imply causation, but can be a pretty big hint.**

Observational studies versus Experiments

Section 4.2

Observational Studies vs. Experiments

- In an **observational study**, the researchers do not set the level of the explanatory variable for each subject. Typically each subject herself decides her level of the explanatory variable. Sometimes nature decides.
- For example, the researchers didn't control which children slept with a night light on or not.
- Observational studies always have potential confounding variables present and these may prevent us from determining cause and effect.

Observational Studies vs. Experiments

- In an **experiment**, the researchers set the level of the explanatory variable for each subject.
- These levels may correspond to a treatment and control.
- Well designed experiments can control for confounding variables by making the treatment and control groups very similar except for what the experimenter manipulates.

Experiments and aspirin example.

Physicians' Health Study I (study aspirin's affect on reducing heart attacks.

- Started in 1982 with 22,071 male physicians.
- The physicians were **randomly assigned** into one of two groups.
 - Half took a 325mg aspirin every other day and half took a placebo.

Results

- Intended to go until 1995, the aspirin study was stopped in 1988 after finding significant results.
- 189 (1.7%) heart attacks occurred in the placebo group and 104 (0.9%) in the aspirin group. This is a 45% reduction in heart attacks for the aspirin group.
- What about confounding variables? Could the aspirin group be different than the placebo group in some other ways?
 - Did they have a better diet?
 - Did they exercise more?
 - Were they genetically less likely to have heart attacks?
 - Were they younger?

The Big Idea

- Confounding variables are often circumvented in experiments due to the **random assignment** of subjects to treatment groups.
- Randomly assigning people to groups tends to balance out all other variables between the groups.
- So confounding variables, including ones the researchers didn't anticipate, should be roughly equalized between the two groups and therefore should not be confounding.
- **Thus, cause and effect conclusions are sometimes possible in experiments through random assignment.** It must be a well run experiment though.

Random sampling and random assignment.

- With observational studies, **random sampling** is often done. This possibly allows us to make inferences from the sample to the population where the sample was drawn.
- With experiments, **random assignment** is done. This might allows us to conclude causation.

- The Physician's Health Study used random assignment. Did it also use random sampling?
- No, hardly any experiments use random sampling, but get their subjects in other ways.
- The Physician's Health Study sent out invitation letters and questionnaires to all 261,248 male physicians between 40 and 84 years of age who lived in the United States.
- Of the 59,285 who were willing to participate in the trial, 26,062 were told they could not because of some medical condition or current medical treatment.

- So to what group can we generalize the results that taking aspirin can reduce heart attacks?
 - Just physicians in the study?
 - All male physicians between 40-84 years old?
 - All males physicians?
 - All males between 40-84 years olds?
 - All males?
 - Everyone between 40-84 years old?
 - Everyone?

Article Baseline Demographics After Random Assignment

Parameter	Placebo (n=129)	Uceris (n=128)
Mean age, years (range)	39.9 (12–68)	37.6 (13–66)
Men	77 (59.7)	70 (54.7)
Women	52 (40.3)	58 (45.3)
Mean disease duration (yrs)	6.3	5.5
Duration ≤1 year, n (%)	23 (17.8)	28 (21.9)
Duration >5 years, n (%)	51 (39.5)	44 (34.4)
Proctosigmoiditis	64 (49.6)	58 (45.3)
Left-sided colitis	44 (34.1)	37 (28.9)
Mean baseline UCDAI score	6.2	6.5
Mean baseline EI score	6.6	6.5
Prior mesalazine use	75 (58.1)	66 (51.6)
Prior sulfasalazine use	28 (21.7)	33 (25.8)

Sandborn WJ, Travis S, Moro L, Jones R, Gaultier T, Bagin R, Huang M, Yeung P, Ballard ED 2nd Once-daily budesonide MMX[®] extended-release tablets induce remission in patients with mild to moderate ulcerative colitis: results from the CORE I study. *Gastroenterology* 2012 Nov;143(5):1218-26

Blocking and Random Assignment

- The goal in random assignment is to make the two groups as similar as possible in all ways other than the treatment.
- Sometime there are known confounders and you can block on (control for) these variables.
- For example, if our subjects consist of 60% females and 40% males, we can force each group to be 60% female and 40% male, using a matched pair design.
- Blocking makes sense when there are known confounders you want to control for. But randomly assigning subjects to groups makes them as similar as possible in terms of unknown confounders.

Blinding.

Even in experiments, the treatment and control groups can be different in ways other than the explanatory variable. This is especially true when the response variable is somewhat subjective.

Pain is an example. One study found that 1/4 of patients suffering from post-operative pain, when given a placebo (just a pill of sugar and water) claimed they experienced "significant prompt pain relief".

Blinding.

People might not be able to judge their own levels of pain very well, and may be influenced by the belief that they have taken an effective treatment.

Thus in an experiment with such a response variable, researchers should ensure the subject does not know whether he or she received the treatment or the control. This is called blinding.

In a *double-blind* experiment, neither the subject nor the researcher recording the response variable knows the level of the explanatory variable for each subject, i.e. treatment or control.

Portacaval shunt example.

The following example shows the importance of doing a randomized controlled experiment.

The portacaval shunt is a medical procedure aimed at curbing bleeding to death in patients with cirrhosis of the liver.

The following table summarizes 51 studies on the portacaval shunt. The poorly designed studies were very enthusiastic about the surgery, while the carefully designed studies prove that the surgery is largely ineffective.

Design	Degree of enthusiasm		
	High	Moderate	None
No controls	24	7	1
Controls, but not randomized	10	3	2
Randomized controlled	0	1	3

Portacaval shunt example.

Why did the poorly designed studies come to the wrong conclusion?

A likely explanation is that in the studies where patients were not randomly assigned to the treatment or control group, by and large the healthier patients were given the surgery.

This alone could explain why the treatment group outlived the control group in these studies.

Design	Degree of enthusiasm		
	High	Moderate	None
No controls	24	7	1
Controls, but not randomized	10	3	2
Randomized controlled	0	1	3