جامعة نيويورك أبوظبي

NYU | ABU DHABI

PSYCH-UH 1004Q: Statistics for Psychology

Class 1: Why do we need statistics?

Prof. Jon Sprouse
Psychology

My name, your name

My full title and name is Prof. Jon Sprouse. But, we will be working closely together this semester, so you can call me "Jon". If that feels too informal, I am also ok with "Prof. Sprouse" or "Dr. Sprouse". But I don't require the title. I know that I am a professor — I have the 15 years of memories of it.

You all have been around for a bit, so you know this, but I want to say it out loud. I carry a bunch of privilege as a middle aged white professor - nobody ever questions my credentials. This makes it easier for me to be ok with using my first name. Other professors or administrators with PhDs may not be so lucky, so I recommend always using Prof. Lastname or Dr. Lastname when you first meet people at a university. They will tell you if it is ok to use their first name.

My pronouns are he/him. For some reason, the online systems (Brightspace and Albert) don't allow professors to add their pronouns. If I could upload them, I would. I think this is a terrific new feature of these systems.

Finally, the system also allows you to upload a pronunciation of your name. That is great too! If you do it, I will try my very best to learn them.

Why do we do science?

It is about asking questions, getting answers, and changing your mind!

We all have beliefs about how the universe works.

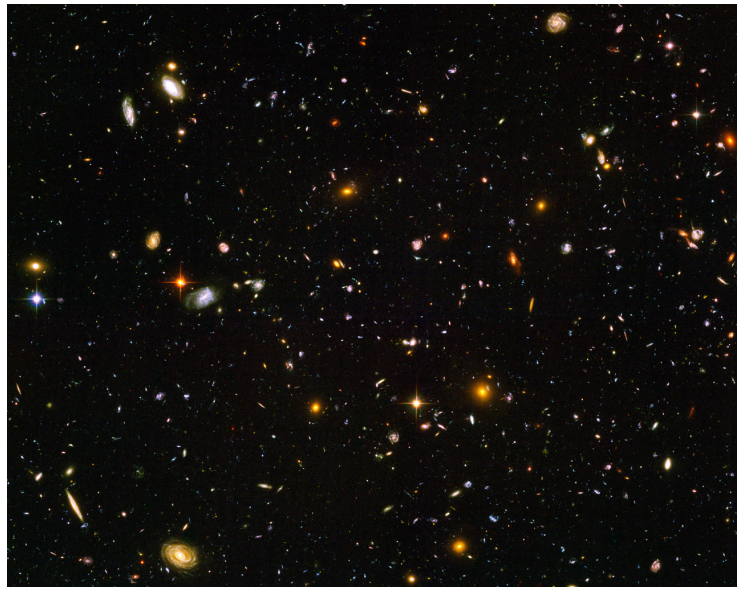

Science gives you a set of rules for asking questions about the universe, gathering data to answer that question, and most importantly, for **changing your mind** when you encounter new evidence.

You've probably seen diagrams like this before - something showing the "scientific method". This course is provides one key part to this process - how to interpret the outcome of a test!

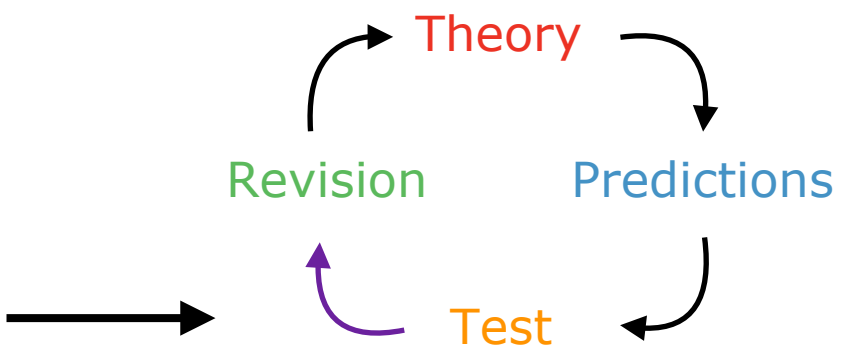

Science has never been more important

Scientific debates are becoming more and more relevant to the world we live in. As information becomes easier and easier to access, it is critical that we understand how to use evidence to prove/disprove theories.

The Impact of Vaccines in the United States

Disease	Baseline 20th Century Pre-Vaccine Annual Cases	2008 Cases*	Percent Decrease
Measles	503,282	55	99.9%
Diphtheria	175,885	0	100%
Mumps	152,209	454	95.7%
Pertussis	147,271	10,735	92.7%
Smallpox	48,164	0	100%
Rubella	47,745	11	99.9%
Haemophilus influenzae type b, invasive	20,000	30	99.9%
Polio	16,316	0	100%
Tetanus	1,314	19	98.6%

*Provisional. Widespread use of vaccines in the United States has eliminated or almost eliminated infectious diseases that were once terrifying household names. Credit: Morbidity and Mortality Weekly Report, Centers for Disease Control and Prevention, 4/2/99, 12/25/09, 3/12/10

In short, there are a number of debates in society that depend upon an understanding of what it means to use evidence.

Science has never been more important

Scientific debates are becoming more and more relevant to the world we live in. As information becomes easier and easier to access, it is critical that we understand how to use evidence to prove/disprove theories.

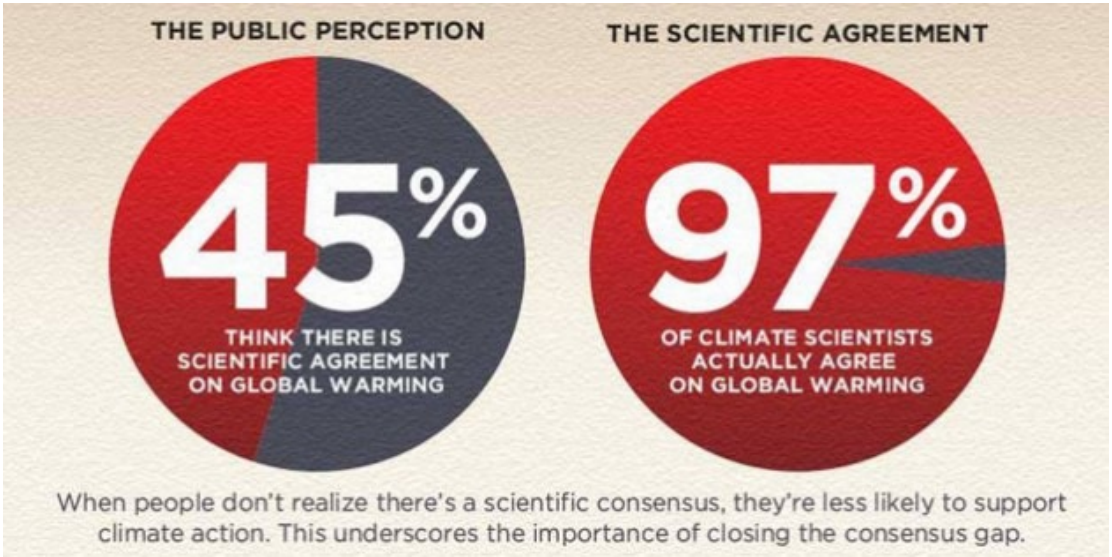

In short, there are a number of debates in society that depend upon an understanding of what it means to use evidence.

Why do we care about the scientific method?

The bottom line is that science is incredibly successful, and appears to be substantially more successful than other knowledge-gathering methods.

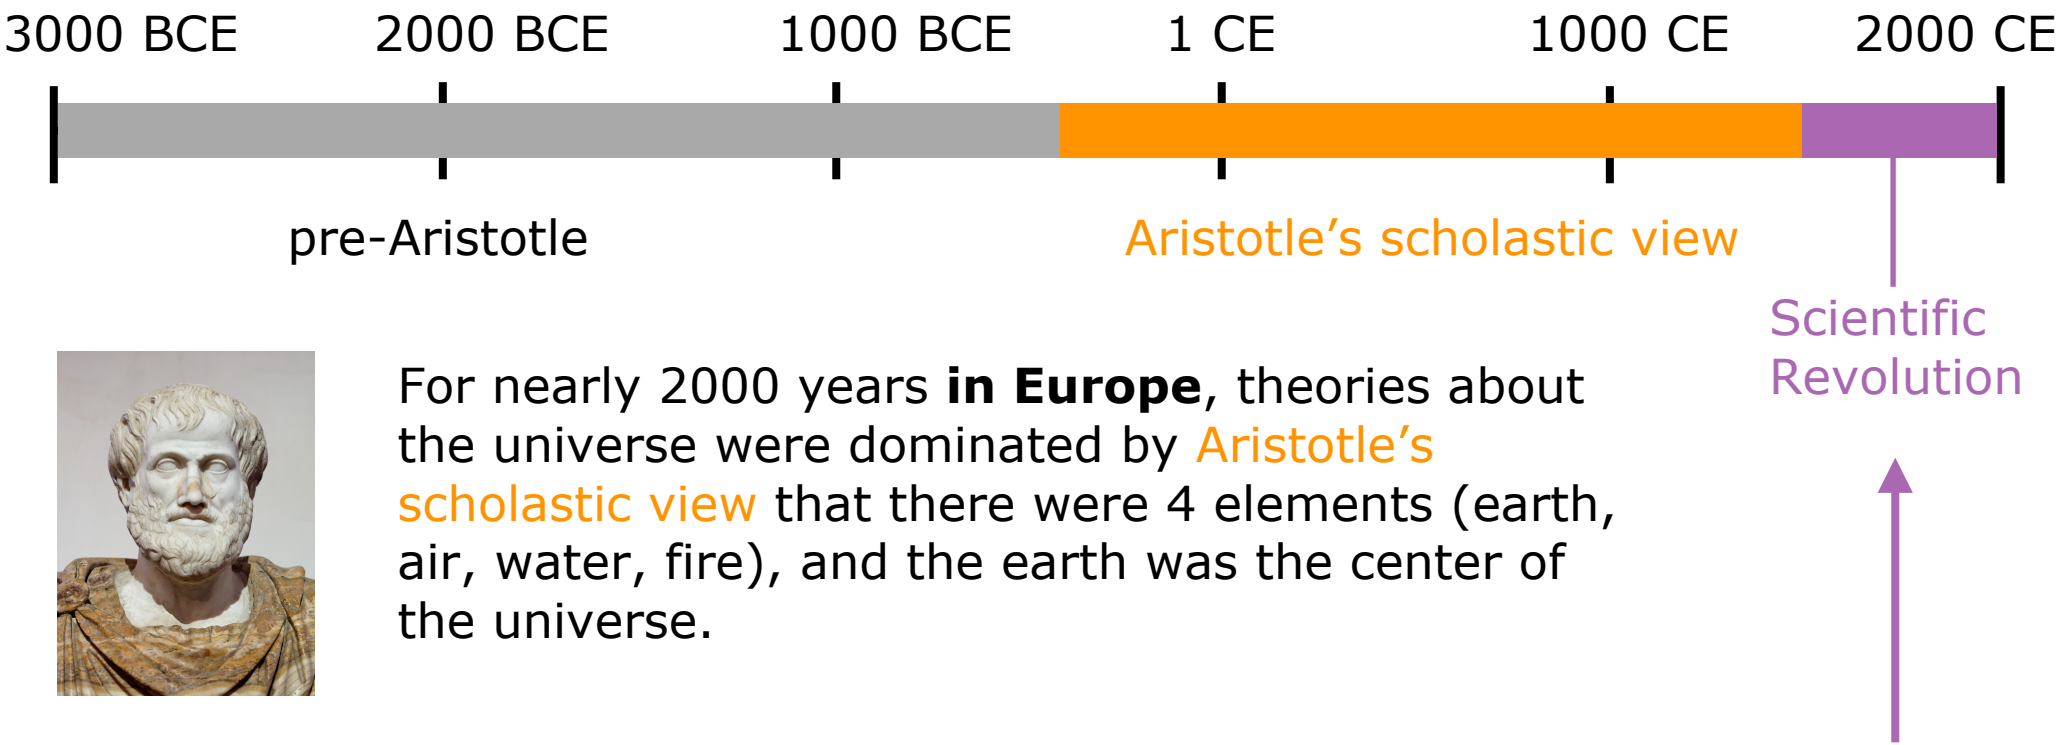

For nearly 2000 years **in Europe**, theories about the universe were dominated by Aristotle's scholastic view that there were 4 elements (earth, air, water, fire), and the earth was the center of the universe.

Then from 1550-1700 prominent thinkers started to question those theories, and even question how to build a theory. In Europe's history, we call this the age of enlightenment or the scientific revolution. From that point forward, the expansion of human knowledge has been dramatic!

Why do we need statistics?

The Scientific Method and Statistics

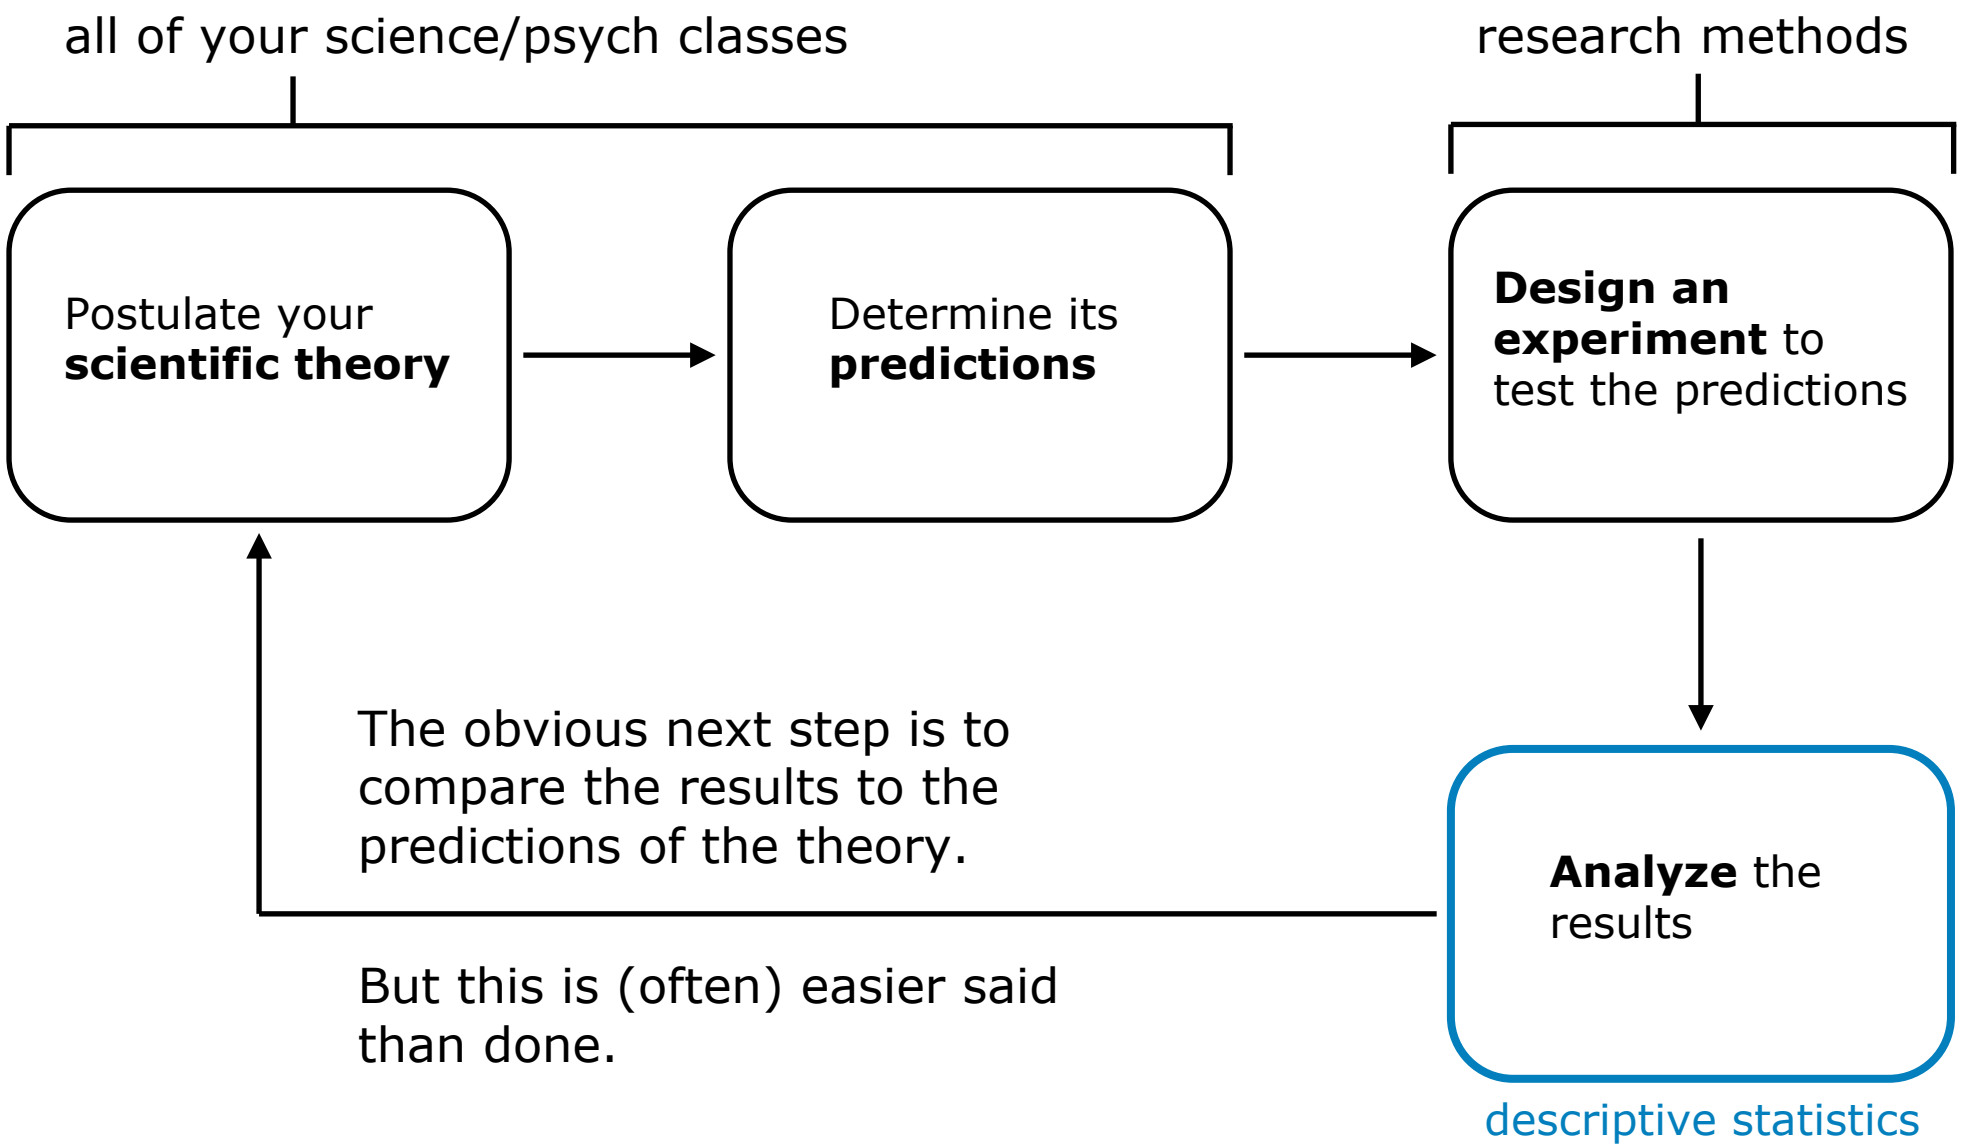

Without variability, we wouldn't need statistics

If the scientific question is about something without variability - where it either happens or not, and could not just happen by chance - then you don't need statistics.

These are the kinds of questions that people give as examples of the scientific method.

So we go out and look at what kinds of ravens we can find. Perhaps most of them are black, but we happen to come across a white raven too. Then we would know that our theory is false. Case closed.

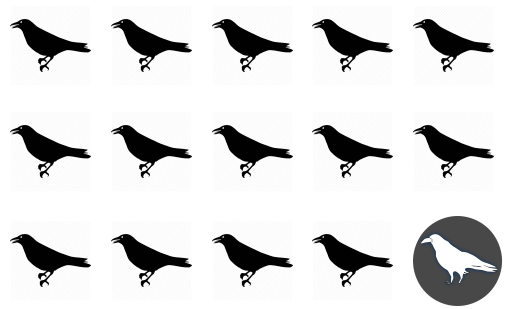

I know this seems like a silly example, but the general form does happen in real science. A number of questions in natural sciences have this form - e.g., do we see a particle of a certain type or not?

The challenge in these cases is usually making the measurement - it can be very difficult to measure things that are very small or very weak or very far away! But that challenge is not solved by statistics.

We need statistics because of variability

Here is a famous example from the early history of statistics.

Two biologists were having tea in the early 1920s: Ronald Fisher and Muriel Bristol. Bristol told Fisher that she could taste the difference between the order that the tea and milk were poured into the cup: tea-1st/milk-2nd versus milk-1st/tea-2nd.

So they designed an experiment to test it. Fisher poured 8 cups of tea, 4 with the tea first, and 4 with the milk first. He presented them to Bristol in a random order. And she made a guess for each one.

Here's the question: **How many does she have to get right to prove that she can taste the difference?**

We need statistics because of variability

The issue, as you probably immediately thought, is that even if Bristol could **not** taste the difference, guessing randomly could lead to 0, 2, 6, or 8 correct guesses! (She knew it was 4 and 4, so the correct guesses go by 2s.)

Intuitively, you probably feel the following:

4 correct (4 incorrect) feels like the most likely outcome if one is guessing.

8 correct (0 incorrect) feels very unlikely, as if Bristol can taste the difference.

But what about 6 correct (2 incorrect)? How do you feel about that?

Another way to put this is that we expect the results of experiments to vary from one experiment to another even when nothing scientifically interesting is happening. The questions we want to ask are how much they will vary when nothing interesting is happening, and how the actual results we observed compare to that. Did we see something extreme enough to conclude that something scientifically interesting happened?

What actually happened?

Just FYI, **Bristol got all 8 correct!** Fisher calculated the probability of that outcome by just guessing (chance)- there is about a **1.4%** probability of doing that by guessing, so he concluded that he was wrong and she was right.

People definitely have opinions about the order. Here is a poll from the UK about it.

This apparently has a chemical explanation based on how the milk molecules get cooked, or not, based on the order. I am not a chemist, so this is beyond me.

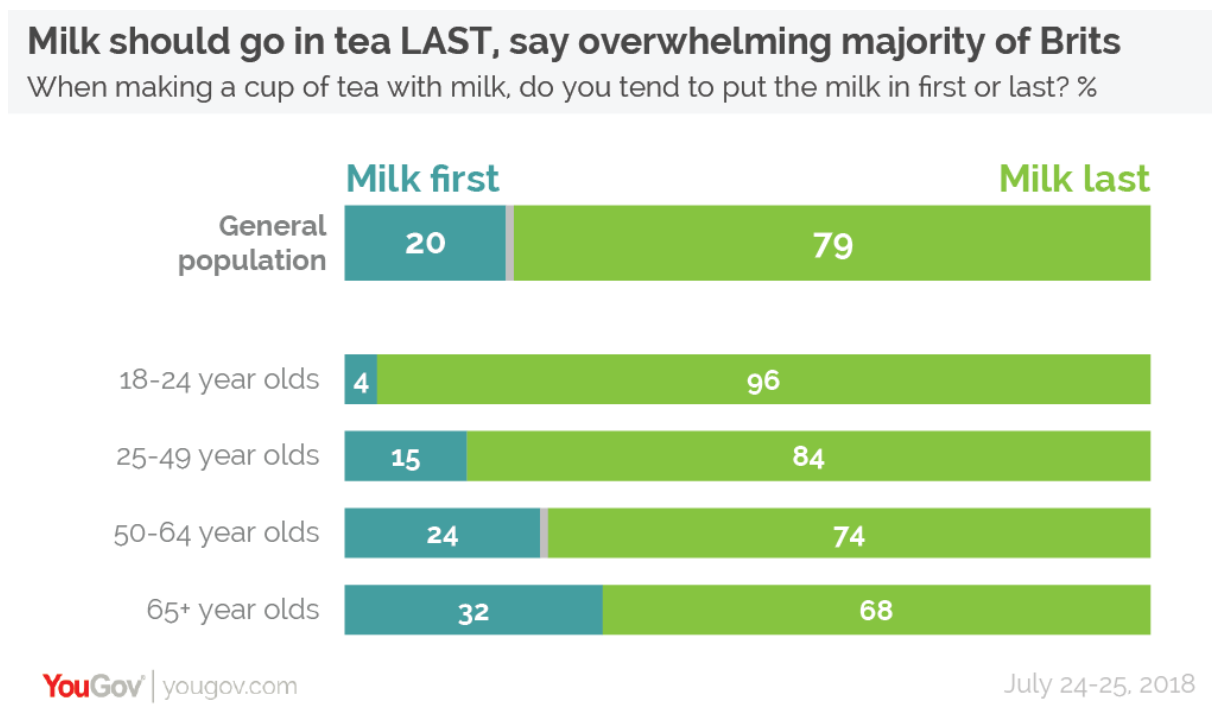

We need statistics because of variability

In the early 1900s, Guinness (the beer brewery in Ireland) hired a chemist and amateur mathematician William Sealy Gosset to help them use scientific methods to determine the best farming techniques for the ingredients of their beer (barley, hops, etc).

The issue is that farming is difficult and expensive, especially in the early 1900s. So if farmers wanted to try a new technique or a new strain of crop, they could really only risk a few acres at a time.

With small test sizes, variability becomes an even bigger problem. We expect crops to vary in their qualities for any number of reasons (soil, sun, rain, etc). If a farmer found that their 4 acres of experimental crop was better (or worse) than the rest, could they conclude that it was due to the new technique, or due to the natural variation that we expect in farming?

Gosset invented the *t*-test to deal with the small numbers of plots that farmers could test. It is a foundational test in statistics. We'll see it later in this course!

We need statistics because of variability

And just to be clear, we still need statistics today. Climate change is a great example. Temperatures vary. The big question underlying the political debate about climate change (in countries where there is still a debate) is what is causing the variability - human factors (gas emissions) or natural factors.

There are multiple methods of investigating this, including massive computer simulations.

But one set of methods are statistical - to compare the expected variability from natural factors to the temperature changes that we observe.

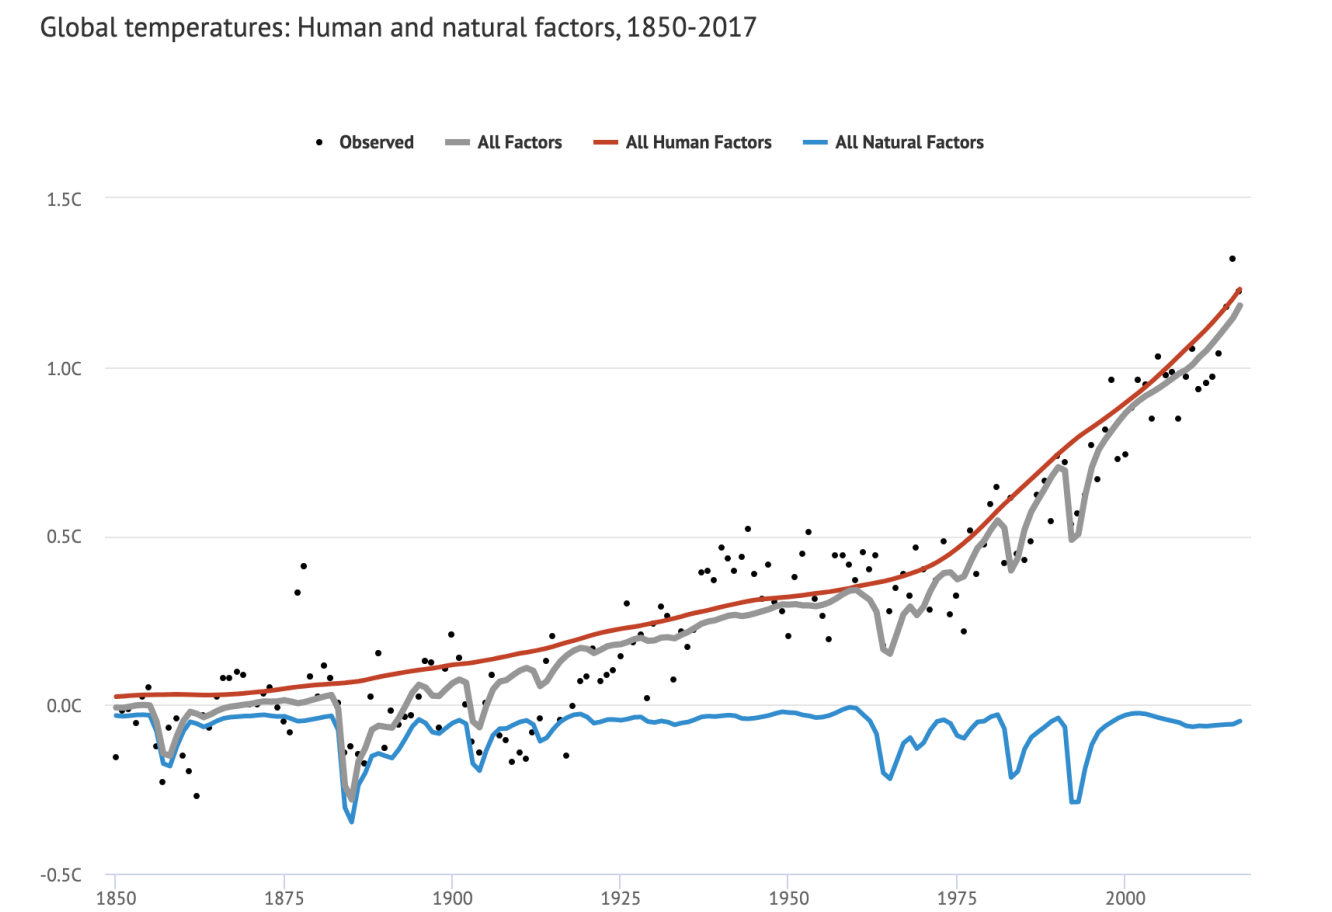

The Scientific Method and Statistics

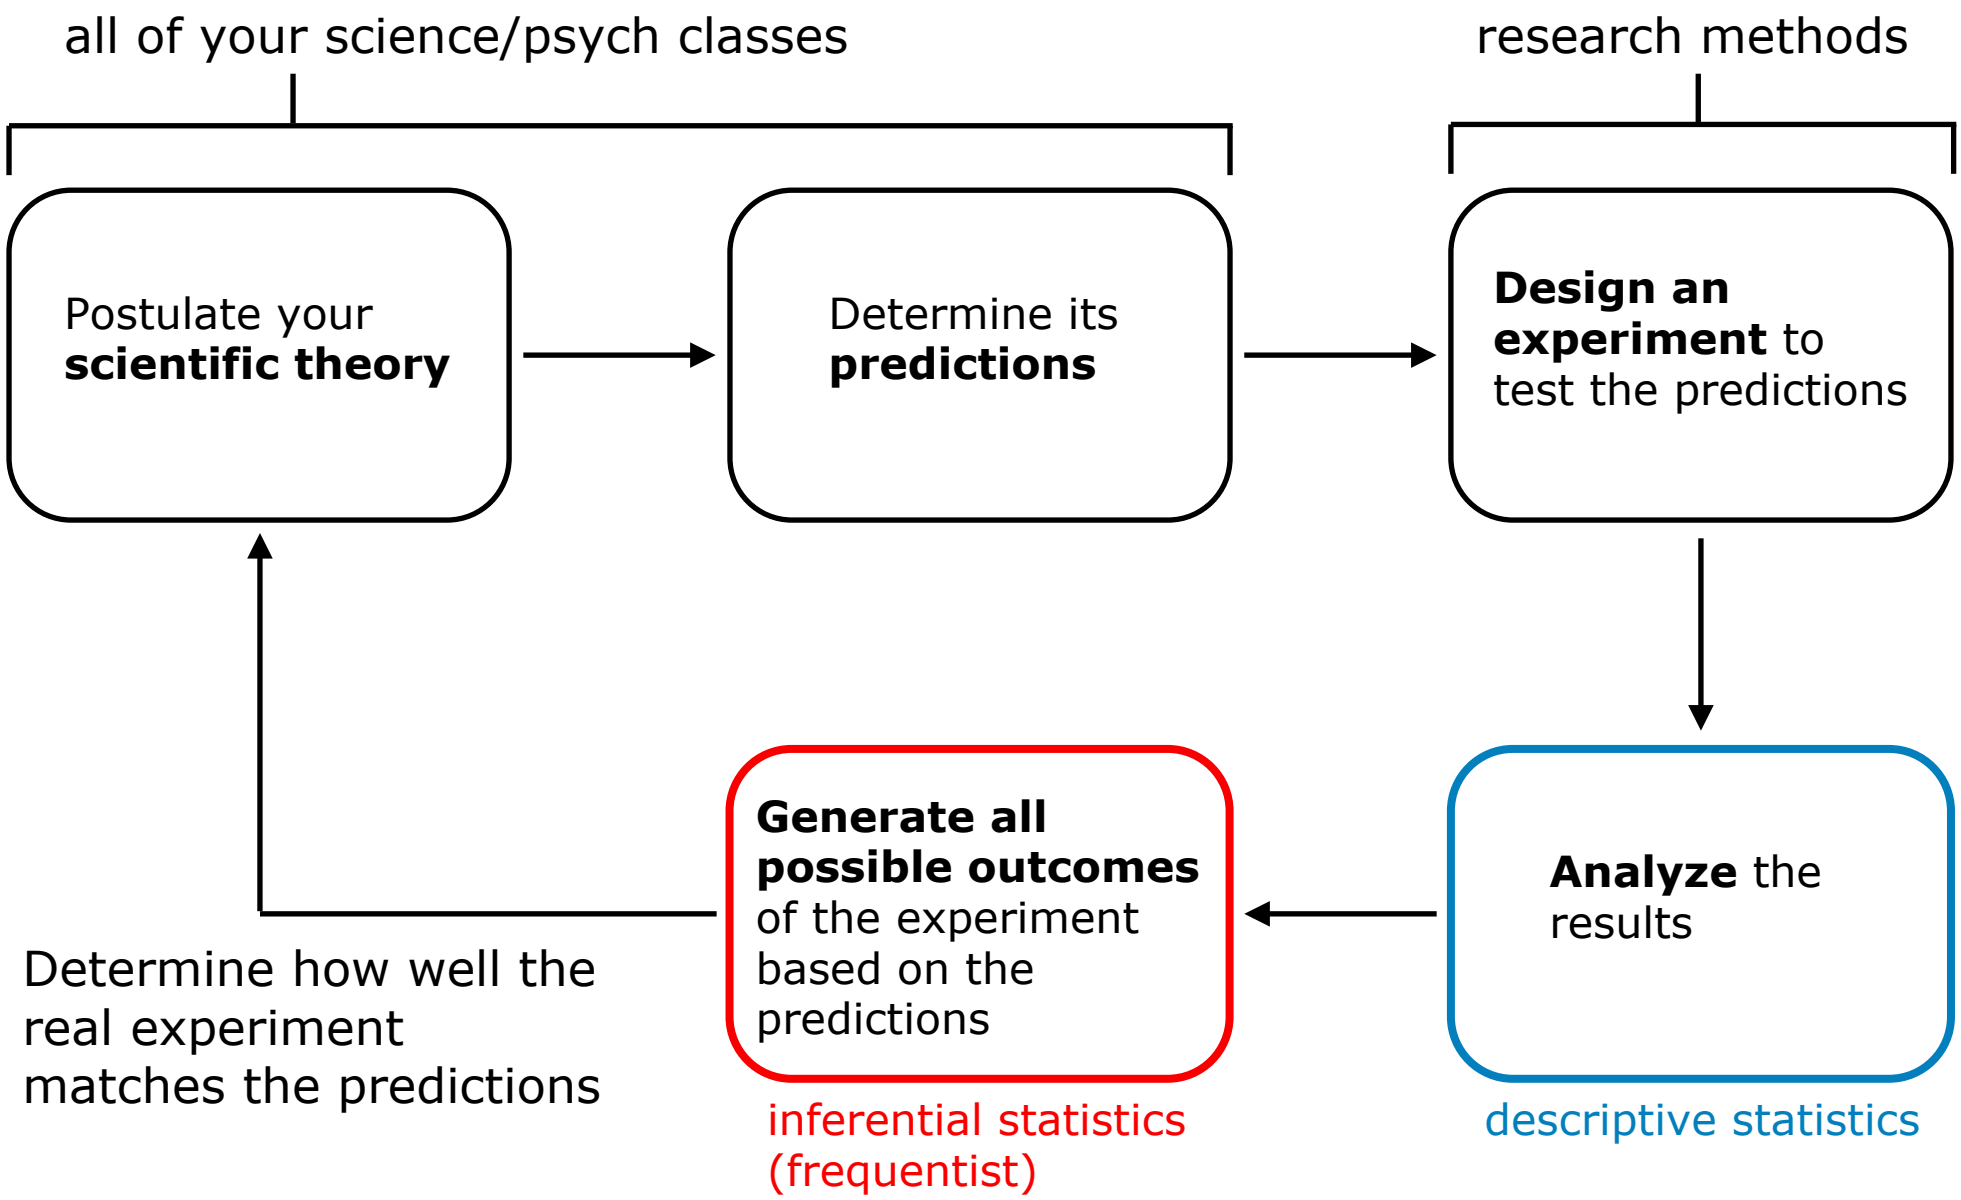

But will you really use it?

YES

Statistics is one of those courses that sticks with you after college. You may not be running statistical tests yourself after college, but you will be trying to answer questions using data. And statistics gives you a foundation for interpreting data rigorously in the face of inherent variability.

Psych majors: You will definitely use statistics in your own research.

Science majors: You will likely use statistics in your research.

Data careers: You will likely use statistics.

Daily life: You will see statistics everywhere — in news articles, in decisions about medicine, in product comparisons, in political debates, etc. This course will allow you to evaluate those statistics directly!

The components of the course

The components of this course all in one place

Lectures - Monday/Wednesday at 12:45-2:00pm

The lectures will focus on the conceptual part (the theory) of statistics.

Recitations - Friday at 10:50am-12:05pm **(8 of them)**

In the recitations you will learn R - a free, open source, cross-platform programming language specifically for statistics (and plotting). I will discuss this in more detail in a couple of slides. They will be both in-person and online.

Homework problems edited from the textbook (8 of them)

These will be due Sundays at 11:59pm as indicated on the schedule.

Take-home exams (2 of them)

These will be assigned at the beginning of each unit of the course so you can work on them as you learn. I'll say more in the next slides.

Wk	Date		Topic	Reading	Assignments due
Unit I: Descriptive Statistics and Basic Inferential Statistics					
1	8/30	W	Introduction: Why do we need statistics?		
	9/01	F	No recitation		
2	9/04	M	Science is our goal	Ch 1	
	9/06	W	Frequencies	Ch 2	
	9/08	F	Recitation 1		HW1
	9/10	Su	HW 1 due at 11:59pm		
3	9/11	M	Central tendency and variability	Ch 3	
	9/13	W	z-scores and the sampling distribution of the mean	Ch 4	HW2
	9/15	F	Recitation 2		
	9/17	Su	HW 2 due at 11:59pm		
4	9/18	M	Putting it in practice		
	9/20	W	Foundations of Null Hypothesis Testing		
	9/22	F	Recitation 3		
	9/24	Su	No homework (you're welcome)		
5	9/25	M	The logic of Null Hypothesis Testing 1	Ch 5	
	9/27	W	The logic of Null Hypothesis Testing 2		HW 3
	9/29	F	Recitation 4		
	10/01	Su	HW 3 due at 11:59pm		
6	10/02	M	t-test: one sample	Ch 6	
	10/04	W	Confidence intervals		HW4
	10/06	F	Recitation 5		
	10/08	Su	HW 4 due at 11:59pm		

The standard rhythm of the week is as follows:

Lectures on M and W.

Recitation on Friday.

Homework is due on Sundays at 11:59pm

There are times when there is no homework, like here in week 4.

And there are times when there is no recitation - **like this week**!

7	10/09	M	t-test: two independent samples	Ch 7	
	10/11	W	Statistical power and effect size	Ch 8	(practice problems)
	10/13	F	No Recitation – Spring Break		
	10/15	Su	No homework – Spring Break		
	10/16	M	No class – exam equivalent		
			Spring Break		
8	10/30	M	No class – Exam 1 due at 11:59pm		
Unit II: Advanced Inferential Statistics					
8	11/01	W	Statistical power and effect size	Ch 8	
	11/03	F	No recitation		(practice problems)
	11/05	Su	No homework		
9	11/06	M	Linear correlation	Ch 9	
	11/08	W	Linear regression	Ch 10	HW5
	11/10	F	Recitation 6		
	11/12	Su	HW 5 due at 11:59pm		
10	11/13	M	t-test: matched/paired (and linear regression)	Ch 11	
	11/15	W	ANOVA: the logic	Ch 12	HW6
	11/17	F	No recitation		
	11/19	Su	HW 6 due at 11:59pm		

In the week leading up to spring break, there is no recitation and no homework.

Even though there is no homework, I have given you practice problems if you want to practice the content of that week before doing the exam.

There is also no class on the day the exam is due - the monday after spring break. Again, you can use that day to finish the exam.

Also, there is no class on the Monday before spring break. You can use that day to work on the exam or generally finish up midterms.

These also make sure I don't overload the content of the course.

7	10/09	M	t-test: two independent samples	Ch 7	
	10/11	W	Statistical power and effect size	Ch 8	(practice problems)
	10/13	F	No Recitation – Spring Break		
	10/15	Su	No homework – Spring Break		
	10/16	M	No class – exam equivalent		
			Spring Break		
8	10/30	M	No class – Exam 1 due at 11:59pm		
Unit II: Advanced Inferential Statistics					
8	11/01	W	Statistical power and effect size	Ch 8	
	11/03	F	No recitation		(practice problems)
	11/05	Su	No homework		
9	11/06	M	Linear correlation	Ch 9	
	11/08	W	Linear regression	Ch 10	HW5
	11/10	F	Recitation 6		
	11/12	Su	HW 5 due at 11:59pm		
10	11/13	M	t-test: matched/paired (and linear regression)	Ch 11	
	11/15	W	ANOVA: the logic	Ch 12	HW6
	11/17	F	No recitation		
	11/19	Su	HW 6 due at 11:59pm		

In the week the exam is due, there is no recitation or homework. You will have done enough.

Again, I have created practice problems if you want to practice that week's content.

No recitation on 11/17. This is just because we want to get the full ANOVA lesson finished before the next recitation.

11	11/20	M	ANOVA: one-way, independent samples		
	11/22	W	Multiple comparisons: the logic	Ch 13	
	11/24	F	Recitation 7		HW7
	11/26	Su	HW 7 due at 11:59pm		
12	11/27	M	Multiple comparisons: the corrections	Ch 13	
	11/29	W	No Class – Legislative day = Friday (No Recitation)		
	12/01	F	No Recitation – National Day		(Practice problems)
	12/03	Su	No homework – National Day		
13	12/04	M	ANOVA: two-way, independent samples	Ch 14	
	12/06	W	ANOVA: two-way, independent samples		
	12/08	F	Recitation 8		
	12/10	Su	HW 8 due at 11:59pm		HW8
14	12/11	M	Interactions	Ch 14	
	12/13	W	Wrap-up		
	12/15	F	No recitation (finals period)		
	12/17	Su	No homework (finals period)		
	12/20	W	Exam 2 due at 11:59pm		

The week of National Day, there is no class on Wednesday because of a legislative day; there is no recitation and no homework.

Again, I have created practice problems if you want to practice that week's content.

11	11/20	M	ANOVA: one-way, independent samples		
	11/22	W	Multiple comparisons: the logic	Ch 13	
	11/24	F	Recitation 7		HW7
	11/26	Su	HW 7 due at 11:59pm		
12	11/27	M	Multiple comparisons: the corrections	Ch 13	
	11/29	W	No Class – Legislative day = Friday (No Recitation)		
	12/01	F	No Recitation – National Day		(Practice problems)
	12/03	Su	No homework – National Day		
13	12/04	M	ANOVA: two-way, independent samples	Ch 14	
	12/06	W	ANOVA: two-way, independent samples		
	12/08	F	Recitation 8		
	12/10	Su	HW 8 due at 11:59pm		HW8
14	12/11	M	Interactions	Ch 14	
	12/13	W	Wrap-up		
	12/15	F	No recitation (finals period)		
	12/17	Su	No homework (finals period)		
	12/20	W	Exam 2 due at 11:59pm		

There is no homework or recitation during the last week of the semester.

The second exam is due during finals period!

The homework problems

The homework problems all come from the textbook.

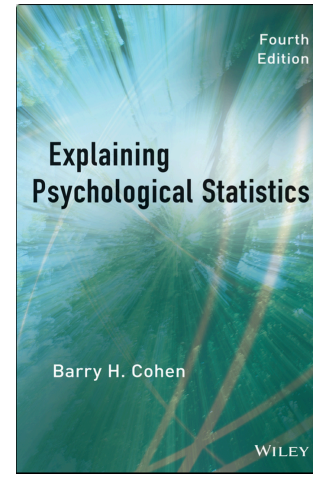

But sometimes the problems are a little unclear, so we have created word documents where we have edited the problems a bit for clarity. We do tell you which problem it is so you can see them in the book too, but please answer the problems as written in the word document!

The homework problems are fairly math-y. Even though we will almost always use a computer to do actual statistical analyses, it is important to work through the math a little bit because the math will help to reinforce the concepts. But always think of it in this direction: concept first, then math. The math is there to make the concept precise. Not the other way around.

We will typically assign only 2-6 problems per assignment. So these should not be a burden. (That said, they can have multiple parts, sorry, statistics has multiple steps!) If you want do extra problems for practice, the book is full of them! The ones with asterisks have answers at the back of the book!

The take-home exams

The take-home exams will always have two components: a conceptual part that literally asks questions about the statistical theory we have been learning, and a practical part that asks you to do an analysis (with R).

These exams are open note and open book. My goal is for you to learn the concepts. You will use these for the rest of your life as a scientist. My goal is not to force you to do some math calculations under time pressure.

There is (obviously) no time limit when you work on the exams. You can work on them in parallel with the weekly lessons or wait until the end of each unit. Your choice.

However, you must work alone. You cannot ask your classmates for help on the exams. I will take you on your honor that you worked alone on the exams.

I have placed three "exam days" in the schedule just as if we were taking them in class. But we are not. Those are simply days without class. If I didn't do this, we'd end up covering more material, and that defeats the purpose of take-home exams (which is to make the course more reasonable)!

There is no perfect textbook in statistics

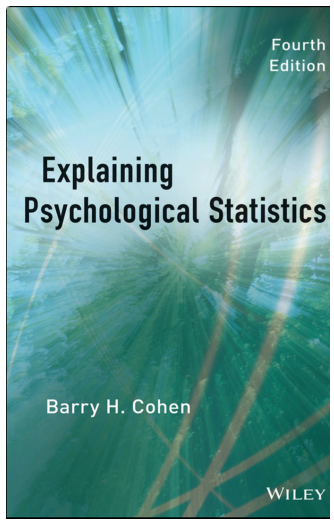

The textbook we use is fairly old school. It is a classic presentation of frequentist statistics. It is also fairly math-centric (but the math is only arithmetic, no calculus or anything like that). It was written by an NYU-NY professor!

We use it because we want you to own a classic book. It is the kind of book you can use as a reference in the future if you ever need to look up a classic concept.

That said, we know that there is no perfect textbook for statistics. Most people, including me, learn statistics in waves. So we end up reading several different books. The question really just is which one to start with. After that, you can go on and read different ones to get different approaches. I will recommend several others (free and not) at the end of the semester for you to continue learning if you want!

We (the psych program) think this is a good starting book, but if it doesn't speak to you, don't worry too much about it. It is just one small part of the course. I will also be teaching these concepts to you in class in my own way, so you will get several different presentations of the material!

What do you need to read?

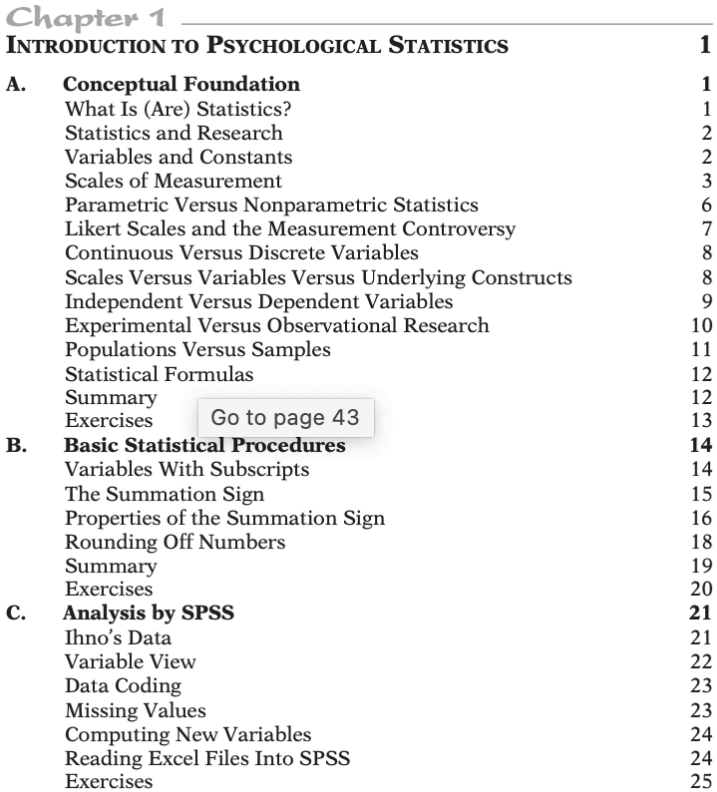

Chapter 1

INTRODUCTION TO PSYCHOLOGICAL STATISTICS 1

Go to page 43

Each chapter of the book is split into 3 sections. You only have to read two:

A: Conceptual Foundation

B: Basic Statistical Procedures

~~C: Analysis by SPSS~~

Do not read part C in any chapter. We won't use SPSS in this course. If, later, you want to learn SPSS, you can go back to this book. But I don't think you will want to!

Part C is always about SPSS. SPSS is a piece of software to do statistics. It is frequently used in classes because it is "point and click", which makes it easy to use when you are first starting out with statistics.

But that ease of use can turn into limitations when you actually start to do your own research (for your capstone, etc), and later in your career. So we want to use a tool that will serve you better at NYUAD and in life. That is R.

Why are we choosing R over SPSS?

R is a programming language specifically designed for statistics. It is a bit harder to learn than SPSS… but it is worth it!

Benefits of R for our course:

It is free, open-source, and cross-platform.

It has a giant user community. Anything you want to do has probably been done before, so there are pre-built packages and internet help groups galore.

It allows you to do three things that you absolutely need to do as a data scientist: (i) manipulate data files, (ii) analyze your data, and (iii) create publication-quality figures. It is simply the right tool for the job.

Benefits of R after this course:

R is a popular tool in data science more generally. It is a nice skill to add to your resumé. We will build a foundation in R so that if you decide to pursue a career that is in any way related to data, you will already have some experience with it.

Recitations

There are 8 recitations spread throughout the semester. These will be where you learn R!

The recitations will be taught by two TAs in the course: Dr. Pietro Cerrone and Dr. Pasha Koval.

They are both postdoctoral researchers here at NYUAD. They both have PhDs in Linguistics, and use R extensively in their own research. They also both work with me on language research, so we make a pretty good team!

The recitations will be held simultaneously in person and by zoom. The link can be found on Brightspace. We do this for multiple reasons. First, R is coding, so you will be sitting in front of a computer. If you want to be remote with that, why not? Second, the recitations are Friday mornings. It may be easier to join from home than to travel to class. And we want to make it easy for you to attend! That said, if you prefer in-person, it will also be in-person. The best of both worlds!

Grades

The grade weights are as follows:

Activity Detail	Percentage	Due
Homework	33.33%	Weekly
Exam 1	33.33%	10/30
Exam 2	33.33%	12/20

And the grade distribution is:

Letter	Min	Max
A	93	100
A-	90	92
B+	87	89
B	83	86
B-	80	82
C+	77	79
C	73	76
C-	70	72
D+	67	69
D	63	66
D-	60	62
F	0	59

That is it!

Remember that **there is no recitation this week**.