UCLA-ALIGNED COURSE FIELD NOTES

STATS 100B

Introduction to Mathematical Statistics

Sampling, estimation, likelihood, confidence intervals, tests, comparisons, and statistical decision making

How these notes are written. I wrote each chapter the way I wish the material had been explained the first time: the real idea, the point where I usually get stuck, the simplest mental model that still stays honest, and a worked decision instead of a polished answer appearing from nowhere. Every chapter closes with practice and a fully written solution. If a sentence sounds impressive but does not help me solve or explain something, it does not belong here.

Daniel Mastick

Curriculum map, working notes, practice, and solutions

Course map

Week	Core thread	What should be solid by the end
1	**Samples, statistics, and sampling distributions**	Reconstruct the model, work an application, test its assumptions, and explain the result in ordinary language.
2	**Method of moments and maximum likelihood**	Reconstruct the model, work an application, test its assumptions, and explain the result in ordinary language.
3	**Estimator properties and information**	Reconstruct the model, work an application, test its assumptions, and explain the result in ordinary language.
4	**Confidence intervals and pivots**	Reconstruct the model, work an application, test its assumptions, and explain the result in ordinary language.
5	**Hypothesis testing and error tradeoffs**	Reconstruct the model, work an application, test its assumptions, and explain the result in ordinary language.
6	**One- and two-sample inference**	Reconstruct the model, work an application, test its assumptions, and explain the result in ordinary language.
7	**Analysis of variance and simple regression**	Reconstruct the model, work an application, test its assumptions, and explain the result in ordinary language.
8	**Survey sampling, resampling, and robust judgment**	Reconstruct the model, work an application, test its assumptions, and explain the result in ordinary language.
9	**Integration studio**	Combine several chapters in one case, preserve their assumptions, and reconcile the result across representations.
10	**Cumulative review**	Retrieve the full course map from memory and solve mixed problems without chapter labels.

Scope anchor: https://catalog.registrar.ucla.edu/course/2025/STATS100B. The sequence is aligned to the official catalog description and expanded into a practical ten-week study plan.

How I would use this packet

Treat each numbered section as one chapter. Read the formal explanation, pause at the student interpretation, and see whether the easy version still says the same thing. Rework the example without looking, then cover the solution in the chapter practice box. The cumulative set at the end is deliberately mixed: knowledge becomes durable when I have to choose the tool instead of being told which chapter I am in.

Contents

1 Samples, statistics, and sampling distributions **5**

1.1 Rebuild the chapter from first principles . . . 5

1.2 Details I would refuse to skip . . . 5

1.3 How this chapter connects . . . 6

1.4 The full working model . . . 6

2 Method of moments and maximum likelihood **7**

2.1 Rebuild the chapter from first principles . . . 7

2.2 Details I would refuse to skip . . . 8

2.3 How this chapter connects . . . 8

2.4 The full working model . . . 8

3 Estimator properties and information **10**

3.1 Rebuild the chapter from first principles . . . 10

3.2 Details I would refuse to skip . . . 10

3.3 How this chapter connects . . . 11

3.4 The full working model . . . 11

4 Confidence intervals and pivots **12**

4.1 Rebuild the chapter from first principles . . . 12

4.2 Details I would refuse to skip . . . 13

4.3 How this chapter connects . . . 13

4.4 The full working model . . . 13

5 Hypothesis testing and error tradeoffs **15**

5.1 Rebuild the chapter from first principles . . . 15

5.2 Details I would refuse to skip . . . 15

5.3 How this chapter connects . . . 16

5.4 The full working model . . . 16

6 One- and two-sample inference **17**

6.1 Rebuild the chapter from first principles . . . 17

6.2 Details I would refuse to skip . . . 18

6.3 How this chapter connects . . . 18

6.4 The full working model 18

7 Analysis of variance and simple regression **20**

7.1 Rebuild the chapter from first principles 20

7.2 Details I would refuse to skip 20

7.3 How this chapter connects 21

7.4 The full working model 21

8 Survey sampling, resampling, and robust judgment **22**

8.1 Rebuild the chapter from first principles 22

8.2 Details I would refuse to skip 23

8.3 How this chapter connects 23

8.4 The full working model 24

CHAPTER 1 Samples, statistics, and sampling distributions

The idea

A statistic is a function of the sample; its sampling distribution describes repeated-sample behavior under a population model. Pivots, sufficiency, ancillarity, and asymptotics separate data reduction from inference.

Rebuild the chapter from first principles

The claim I need to own is this: A statistic is a function of the sample; its sampling distribution describes repeated-sample behavior under a population model. Pivots, sufficiency, ancillarity, and asymptotics separate data reduction from inference. I do not count that as learned just because the sentence looks familiar. I should be able to name the objects, state what changes and what stays fixed, recover the rule from the definitions, and explain why the result is allowed.

The move

My working sequence

1. **Translate.** Rewrite the prompt in the language of samples, statistics, and sampling distributions. List the known quantities, the unknown, the units or types, and any hidden constraint.
2. **Choose.** State the theorem, model, algorithm, accounting rule, or physical law before substituting values. This is where I make the assumption visible.
3. **Execute.** Work one justified step at a time and keep intermediate state readable enough that I can audit it later.
4. **Verify.** Check a boundary case, sign, unit, invariant, order of magnitude, or alternative derivation. Then translate the result back into an ordinary sentence.

Details I would refuse to skip

- **Definitions:** I should be able to define every technical noun in the chapter without using that same noun in the definition.
- **Assumptions:** I write the condition that makes the method legal beside the method itself. A correct procedure under the wrong assumptions is still a wrong answer.
- **Representation:** I move freely among words, equations or code, a diagram, and a small numerical example. If one representation disagrees with another, that disagreement is useful evidence.
- **Failure modes:** I keep the nearest counterexample in mind: the boundary where the rule changes, the invariant breaks, or the model stops matching reality.
- **Communication:** My final answer names what the result means and what it does not establish. I do not let notation do the explanatory work for me.

Sanity check

Closed-note check. Can I teach samples, statistics, and sampling distributions on a blank page, derive or reconstruct its main rule, work a fresh example, and diagnose a deliberately broken solution? If I can only recognize the finished work, I am not done.

How this chapter connects

I read this chapter as the bridge from the course foundations to method of moments and maximum likelihood. The earlier material supplies the language and constraints; this chapter adds a new reasoning move; the later material asks me to use that move without being told its name. That connection matters because exams and real projects mix chapters on purpose.

The full working model

- **Core claim.** A statistic is a function of the sample; its sampling distribution describes repeated-sample behavior under a population model.
- **What moves.** Pivots, sufficiency, ancillarity, and asymptotics separate data reduction from inference.
- **What keeps the answer honest.** The model becomes useful only when I can apply samples, statistics, and sampling distributions to an unfamiliar case and defend the assumptions.
- **Mastery standard.** I can define the objects, rebuild the rule, work a fresh case, compare alternatives, and diagnose a broken solution without relying on recognition.

How I actually think about it

I imagine repeating the data collection, not the calculation, when I interpret a standard error.

The easy version

A statistic is a function of the sample; its sampling distribution describes repeated-sample behavior under a population model.

Worked example

Question. For iid normal data with variance 9 and n=36, what is the standard deviation of the sample mean?
Walkthrough. It is 3 divided by 6, or 0.5.

Common miss

The fastest way to get this chapter wrong is to use the visible procedure without naming its assumption, units, time period, reference group, or boundary. I put that condition beside the work and verify the result independently.

Solved chapter practice

Practice 1. For iid normal data with variance 9 and n=36, what is the standard deviation of the sample mean?
Solution. It is 3 divided by 6, or 0.5.

Practice 2. Explain samples, statistics, and sampling distributions to a classmate using one definition, one concrete example, one assumption, and one failure mode.
Solution. Begin with the core claim above, use the worked case as the concrete example, state the condition highlighted by the warning, and choose a failure where that condition does not hold. A complete explanation connects all four pieces instead of listing them.

Mastery practice A. Reconstruct samples, statistics, and sampling distributions from a blank page. Include the central definition, the governing rule, one required assumption, and one boundary case.
Solution. Start from the blue idea box, but rewrite it in my own words. Identify the objects and state, name the rule before using it, copy the relevant assumption from the details audit, and choose a boundary where that assumption becomes equality or fails. A complete solution explains why each part belongs; a list of vocabulary is not enough.

Mastery practice B. A classmate gets a polished-looking answer to a samples, statistics, and sampling distributions problem but never states a model or checks the result. Write the shortest useful review of that solution.
Solution. I would ask four concrete questions: What exactly are the inputs and desired output? Which rule licenses the main step? Which assumption could invalidate it? Which independent check supports the conclusion? If the answer cannot survive those four questions, I treat the numerical or verbal result as unverified.

CHAPTER 2 Method of moments and maximum likelihood

The idea

Method of moments matches sample and theoretical moments. Maximum likelihood chooses parameters making observed data most plausible under the model; invariance, score, information, and numerical optimization support its use.

Rebuild the chapter from first principles

The claim I need to own is this: Method of moments matches sample and theoretical moments. Maximum likelihood chooses parameters making observed data most plausible under the model; invariance, score, information, and numerical optimization support its use. I do not count that as learned just because the sentence looks familiar. I should be able to name the objects, state what changes and what stays fixed, recover the rule from the definitions, and explain why the result is allowed.

The move

My working sequence

1. **Translate.** Rewrite the prompt in the language of method of moments and maximum likelihood. List the known quantities, the unknown, the units or types, and any hidden constraint.
2. **Choose.** State the theorem, model, algorithm, accounting rule, or physical law before substituting values. This is where I make the assumption visible.
3. **Execute.** Work one justified step at a time and keep intermediate state readable enough that I can audit it later.
4. **Verify.** Check a boundary case, sign, unit, invariant, order of magnitude, or alternative derivation. Then translate the result back into an ordinary sentence.

Details I would refuse to skip

- **Definitions:** I should be able to define every technical noun in the chapter without using that same noun in the definition.
- **Assumptions:** I write the condition that makes the method legal beside the method itself. A correct procedure under the wrong assumptions is still a wrong answer.
- **Representation:** I move freely among words, equations or code, a diagram, and a small numerical example. If one representation disagrees with another, that disagreement is useful evidence.
- **Failure modes:** I keep the nearest counterexample in mind: the boundary where the rule changes, the invariant breaks, or the model stops matching reality.
- **Communication:** My final answer names what the result means and what it does not establish. I do not let notation do the explanatory work for me.

Sanity check

Closed-note check. Can I teach method of moments and maximum likelihood on a blank page, derive or reconstruct its main rule, work a fresh example, and diagnose a deliberately broken solution? If I can only recognize the finished work, I am not done.

How this chapter connects

I read this chapter as the bridge from samples, statistics, and sampling distributions to estimator properties and information. The earlier material supplies the language and constraints; this chapter adds a new reasoning move; the later material asks me to use that move without being told its name. That connection matters because exams and real projects mix chapters on purpose.

The full working model

- **Core claim.** Method of moments matches sample and theoretical moments.
- **What moves.** Maximum likelihood chooses parameters making observed data most plausible under the model; invariance, score, information, and numerical optimization support its use.

- **What keeps the answer honest.** The model becomes useful only when I can apply method of moments and maximum likelihood to an unfamiliar case and defend the assumptions.
- **Mastery standard.** I can define the objects, rebuild the rule, work a fresh case, compare alternatives, and diagnose a broken solution without relying on recognition.

How I actually think about it

My shorthand for this chapter is: my shorthand for this chapter is: likelihood is a function of the parameter with data fixed, not the probability the parameter is true.

The easy version

Method of moments matches sample and theoretical moments.

Worked example

Question. For Bernoulli observations with seven successes in ten trials, find the MLE of p.
Walkthrough. The log-likelihood is maximized at the sample proportion, p-hat=0.7.

Common miss

The fastest way to get this chapter wrong is to use the visible procedure without naming its assumption, units, time period, reference group, or boundary. I put that condition beside the work and verify the result independently.

Solved chapter practice

Practice 1. For Bernoulli observations with seven successes in ten trials, find the MLE of p.
Solution. The log-likelihood is maximized at the sample proportion, p-hat=0.7.

Practice 2. Explain method of moments and maximum likelihood to a classmate using one definition, one concrete example, one assumption, and one failure mode.
Solution. Begin with the core claim above, use the worked case as the concrete example, state the condition highlighted by the warning, and choose a failure where that condition does not hold. A complete explanation connects all four pieces instead of listing them.

Mastery practice A. Reconstruct method of moments and maximum likelihood from a blank page. Include the central definition, the governing rule, one required assumption, and one boundary case.
Solution. Start from the blue idea box, but rewrite it in my own words. Identify the objects and state, name the rule before using it, copy the relevant assumption from the details audit, and choose a boundary where that assumption becomes equality or fails. A complete solution explains why each part belongs; a list of vocabulary is not enough.

Mastery practice B. A classmate gets a polished-looking answer to a method of moments and maximum likelihood problem but never states a model or checks the result. Write the shortest useful review of that solution.
Solution. I would ask four concrete questions: What exactly are the inputs and desired output? Which rule licenses the main step? Which assumption could invalidate it? Which independent check supports the conclusion? If the answer cannot survive those four questions, I treat the

numerical or verbal result as unverified.

CHAPTER 3 Estimator properties and information

The idea

Bias, variance, mean squared error, consistency, efficiency, robustness, and sufficiency judge estimators differently. Rao-Blackwell improvement and Cramer-Rao bounds show how information and conditioning can reduce variance under assumptions.

Rebuild the chapter from first principles

The claim I need to own is this: Bias, variance, mean squared error, consistency, efficiency, robustness, and sufficiency judge estimators differently. Rao-Blackwell improvement and Cramer-Rao bounds show how information and conditioning can reduce variance under assumptions. I do not count that as learned just because the sentence looks familiar. I should be able to name the objects, state what changes and what stays fixed, recover the rule from the definitions, and explain why the result is allowed.

The move

My working sequence

1. **Translate.** Rewrite the prompt in the language of estimator properties and information. List the known quantities, the unknown, the units or types, and any hidden constraint.
2. **Choose.** State the theorem, model, algorithm, accounting rule, or physical law before substituting values. This is where I make the assumption visible.
3. **Execute.** Work one justified step at a time and keep intermediate state readable enough that I can audit it later.
4. **Verify.** Check a boundary case, sign, unit, invariant, order of magnitude, or alternative derivation. Then translate the result back into an ordinary sentence.

Details I would refuse to skip

- **Definitions:** I should be able to define every technical noun in the chapter without using that same noun in the definition.
- **Assumptions:** I write the condition that makes the method legal beside the method itself. A correct procedure under the wrong assumptions is still a wrong answer.
- **Representation:** I move freely among words, equations or code, a diagram, and a small numerical example. If one representation disagrees with another, that disagreement is useful evidence.
- **Failure modes:** I keep the nearest counterexample in mind: the boundary where the rule changes, the invariant breaks, or the model stops matching reality.

- **Communication:** My final answer names what the result means and what it does not establish. I do not let notation do the explanatory work for me.

Sanity check

Closed-note check. Can I teach estimator properties and information on a blank page, derive or reconstruct its main rule, work a fresh example, and diagnose a deliberately broken solution? If I can only recognize the finished work, I am not done.

How this chapter connects

I read this chapter as the bridge from method of moments and maximum likelihood to confidence intervals and pivots. The earlier material supplies the language and constraints; this chapter adds a new reasoning move; the later material asks me to use that move without being told its name. That connection matters because exams and real projects mix chapters on purpose.

The full working model

- **Core claim.** Bias, variance, mean squared error, consistency, efficiency, robustness, and sufficiency judge estimators differently.
- **What moves.** Rao-Blackwell improvement and Cramer-Rao bounds show how information and conditioning can reduce variance under assumptions.
- **What keeps the answer honest.** The model becomes useful only when I can apply estimator properties and information to an unfamiliar case and defend the assumptions.
- **Mastery standard.** I can define the objects, rebuild the rule, work a fresh case, compare alternatives, and diagnose a broken solution without relying on recognition.

How I actually think about it

My shorthand for this chapter is: my shorthand for this chapter is: an unbiased estimator can still be terrible if its variance is enormous.

The easy version

Bias, variance, mean squared error, consistency, efficiency, robustness, and sufficiency judge estimators differently.

Worked example

Question. Estimator A has bias 1 and variance 4; B is unbiased with variance 6. Compare MSE. **Walkthrough.** A has MSE 4+1 squared=5; B has MSE 6. Under squared error, A is better despite bias.

Common miss

The fastest way to get this chapter wrong is to use the visible procedure without naming its assumption, units, time period, reference group, or boundary. I put that condition beside the work and verify the result independently.

Solved chapter practice

Practice 1. Estimator A has bias 1 and variance 4; B is unbiased with variance 6. Compare MSE.
Solution. A has MSE 4+1 squared=5; B has MSE 6. Under squared error, A is better despite bias.

Practice 2. Explain estimator properties and information to a classmate using one definition, one concrete example, one assumption, and one failure mode.
Solution. Begin with the core claim above, use the worked case as the concrete example, state the condition highlighted by the warning, and choose a failure where that condition does not hold. A complete explanation connects all four pieces instead of listing them.

Mastery practice A. Reconstruct estimator properties and information from a blank page. Include the central definition, the governing rule, one required assumption, and one boundary case.
Solution. Start from the blue idea box, but rewrite it in my own words. Identify the objects and state, name the rule before using it, copy the relevant assumption from the details audit, and choose a boundary where that assumption becomes equality or fails. A complete solution explains why each part belongs; a list of vocabulary is not enough.

Mastery practice B. A classmate gets a polished-looking answer to a estimator properties and information problem but never states a model or checks the result. Write the shortest useful review of that solution.
Solution. I would ask four concrete questions: What exactly are the inputs and desired output? Which rule licenses the main step? Which assumption could invalidate it? Which independent check supports the conclusion? If the answer cannot survive those four questions, I treat the numerical or verbal result as unverified.

CHAPTER 4 Confidence intervals and pivots

The idea

A confidence procedure has repeated-sample coverage. Normal, t, chi-square, and likelihood-based pivots create intervals whose form depends on known parameters and sampling assumptions.

Rebuild the chapter from first principles

The claim I need to own is this: A confidence procedure has repeated-sample coverage. Normal, t, chi-square, and likelihood-based pivots create intervals whose form depends on known parameters and sampling assumptions. I do not count that as learned just because the sentence looks familiar. I should be able to name the objects, state what changes and what stays fixed, recover the rule from the definitions, and explain why the result is allowed.

The move

My working sequence

1. **Translate.** Rewrite the prompt in the language of confidence intervals and pivots. List the known quantities, the unknown, the units or types, and any hidden constraint.
2. **Choose.** State the theorem, model, algorithm, accounting rule, or physical law before substituting values. This is where I make the assumption visible.
3. **Execute.** Work one justified step at a time and keep intermediate state readable enough that I can audit it later.
4. **Verify.** Check a boundary case, sign, unit, invariant, order of magnitude, or alternative derivation. Then translate the result back into an ordinary sentence.

Details I would refuse to skip

- **Definitions:** I should be able to define every technical noun in the chapter without using that same noun in the definition.
- **Assumptions:** I write the condition that makes the method legal beside the method itself. A correct procedure under the wrong assumptions is still a wrong answer.
- **Representation:** I move freely among words, equations or code, a diagram, and a small numerical example. If one representation disagrees with another, that disagreement is useful evidence.
- **Failure modes:** I keep the nearest counterexample in mind: the boundary where the rule changes, the invariant breaks, or the model stops matching reality.
- **Communication:** My final answer names what the result means and what it does not establish. I do not let notation do the explanatory work for me.

Sanity check

Closed-note check. Can I teach confidence intervals and pivots on a blank page, derive or reconstruct its main rule, work a fresh example, and diagnose a deliberately broken solution? If I can only recognize the finished work, I am not done.

How this chapter connects

I read this chapter as the bridge from estimator properties and information to hypothesis testing and error tradeoffs. The earlier material supplies the language and constraints; this chapter adds a new reasoning move; the later material asks me to use that move without being told its name. That connection matters because exams and real projects mix chapters on purpose.

The full working model

- **Core claim.** A confidence procedure has repeated-sample coverage.
- **What moves.** Normal, t, chi-square, and likelihood-based pivots create intervals whose form depends

on known parameters and sampling assumptions.

- **What keeps the answer honest.** The model becomes useful only when I can apply confidence intervals and pivots to an unfamiliar case and defend the assumptions.
- **Mastery standard.** I can define the objects, rebuild the rule, work a fresh case, compare alternatives, and diagnose a broken solution without relying on recognition.

How I actually think about it

I interpret the procedure's coverage, not a post-data probability that the fixed parameter moves around.

The easy version

A confidence procedure has repeated-sample coverage.

Worked example

Question. A sample mean is 20 with standard error 2. Give a large-sample 95 percent interval.
Walkthrough. Use 20 plus or minus 1.96 times 2, giving approximately 16.08 to 23.92.

Common miss

The fastest way to get this chapter wrong is to use the visible procedure without naming its assumption, units, time period, reference group, or boundary. I put that condition beside the work and verify the result independently.

Solved chapter practice

Practice 1. A sample mean is 20 with standard error 2. Give a large-sample 95 percent interval.
Solution. Use 20 plus or minus 1.96 times 2, giving approximately 16.08 to 23.92.

Practice 2. Explain confidence intervals and pivots to a classmate using one definition, one concrete example, one assumption, and one failure mode.
Solution. Begin with the core claim above, use the worked case as the concrete example, state the condition highlighted by the warning, and choose a failure where that condition does not hold. A complete explanation connects all four pieces instead of listing them.

Mastery practice A. Reconstruct confidence intervals and pivots from a blank page. Include the central definition, the governing rule, one required assumption, and one boundary case.
Solution. Start from the blue idea box, but rewrite it in my own words. Identify the objects and state, name the rule before using it, copy the relevant assumption from the details audit, and choose a boundary where that assumption becomes equality or fails. A complete solution explains why each part belongs; a list of vocabulary is not enough.

Mastery practice B. A classmate gets a polished-looking answer to a confidence intervals and pivots problem but never states a model or checks the result. Write the shortest useful review of that solution.
Solution. I would ask four concrete questions: What exactly are the inputs and desired output? Which rule licenses the main step? Which assumption could invalidate it? Which independent

check supports the conclusion? If the answer cannot survive those four questions, I treat the numerical or verbal result as unverified.

CHAPTER 5 Hypothesis testing and error tradeoffs

The idea

A test specifies null and alternative, statistic, rejection rule, type I error, type II error, power, and effect size. P-values measure compatibility with the null model, not the probability the null is true.

Rebuild the chapter from first principles

The claim I need to own is this: A test specifies null and alternative, statistic, rejection rule, type I error, type II error, power, and effect size. P-values measure compatibility with the null model, not the probability the null is true. I do not count that as learned just because the sentence looks familiar. I should be able to name the objects, state what changes and what stays fixed, recover the rule from the definitions, and explain why the result is allowed.

The move

My working sequence

1. **Translate.** Rewrite the prompt in the language of hypothesis testing and error tradeoffs. List the known quantities, the unknown, the units or types, and any hidden constraint.
2. **Choose.** State the theorem, model, algorithm, accounting rule, or physical law before substituting values. This is where I make the assumption visible.
3. **Execute.** Work one justified step at a time and keep intermediate state readable enough that I can audit it later.
4. **Verify.** Check a boundary case, sign, unit, invariant, order of magnitude, or alternative derivation. Then translate the result back into an ordinary sentence.

Details I would refuse to skip

- **Definitions:** I should be able to define every technical noun in the chapter without using that same noun in the definition.
- **Assumptions:** I write the condition that makes the method legal beside the method itself. A correct procedure under the wrong assumptions is still a wrong answer.
- **Representation:** I move freely among words, equations or code, a diagram, and a small numerical example. If one representation disagrees with another, that disagreement is useful evidence.
- **Failure modes:** I keep the nearest counterexample in mind: the boundary where the rule changes, the invariant breaks, or the model stops matching reality.

- **Communication:** My final answer names what the result means and what it does not establish. I do not let notation do the explanatory work for me.

Sanity check

Closed-note check. Can I teach hypothesis testing and error tradeoffs on a blank page, derive or reconstruct its main rule, work a fresh example, and diagnose a deliberately broken solution? If I can only recognize the finished work, I am not done.

How this chapter connects

I read this chapter as the bridge from confidence intervals and pivots to one- and two-sample inference. The earlier material supplies the language and constraints; this chapter adds a new reasoning move; the later material asks me to use that move without being told its name. That connection matters because exams and real projects mix chapters on purpose.

The full working model

- **Core claim.** A test specifies null and alternative, statistic, rejection rule, type I error, type II error, power, and effect size.
- **What moves.** P-values measure compatibility with the null model, not the probability the null is true.
- **What keeps the answer honest.** The model becomes useful only when I can apply hypothesis testing and error tradeoffs to an unfamiliar case and defend the assumptions.
- **Mastery standard.** I can define the objects, rebuild the rule, work a fresh case, compare alternatives, and diagnose a broken solution without relying on recognition.

How I actually think about it

Before computing, I decide what mistake matters and what effect would be meaningful.

The easy version

A test specifies null and alternative, statistic, rejection rule, type I error, type II error, power, and effect size.

Worked example

Question. A two-sided z statistic is 2.3. What is the approximate p-value?
Walkthrough. It is about 0.021, twice the upper normal tail beyond 2.3, so a 5 percent test rejects but a 1 percent test does not.

Common miss

The fastest way to get this chapter wrong is to use the visible procedure without naming its assumption, units, time period, reference group, or boundary. I put that condition beside the work and verify the result independently.

Solved chapter practice

Practice 1. A two-sided z statistic is 2.3. What is the approximate p-value?
Solution. It is about 0.021, twice the upper normal tail beyond 2.3, so a 5 percent test rejects but a 1 percent test does not.

Practice 2. Explain hypothesis testing and error tradeoffs to a classmate using one definition, one concrete example, one assumption, and one failure mode.
Solution. Begin with the core claim above, use the worked case as the concrete example, state the condition highlighted by the warning, and choose a failure where that condition does not hold. A complete explanation connects all four pieces instead of listing them.

Mastery practice A. Reconstruct hypothesis testing and error tradeoffs from a blank page. Include the central definition, the governing rule, one required assumption, and one boundary case.
Solution. Start from the blue idea box, but rewrite it in my own words. Identify the objects and state, name the rule before using it, copy the relevant assumption from the details audit, and choose a boundary where that assumption becomes equality or fails. A complete solution explains why each part belongs; a list of vocabulary is not enough.

Mastery practice B. A classmate gets a polished-looking answer to a hypothesis testing and error tradeoffs problem but never states a model or checks the result. Write the shortest useful review of that solution.
Solution. I would ask four concrete questions: What exactly are the inputs and desired output? Which rule licenses the main step? Which assumption could invalidate it? Which independent check supports the conclusion? If the answer cannot survive those four questions, I treat the numerical or verbal result as unverified.

CHAPTER 6 One- and two-sample inference

The idea

Paired and independent designs have different sampling structures. Mean, proportion, and variance comparisons require appropriate standard errors, degrees of freedom, pooling assumptions, and confidence or test framing.

Rebuild the chapter from first principles

The claim I need to own is this: Paired and independent designs have different sampling structures. Mean, proportion, and variance comparisons require appropriate standard errors, degrees of freedom, pooling assumptions, and confidence or test framing. I do not count that as learned just because the sentence looks familiar. I should be able to name the objects, state what changes and what stays fixed, recover the rule from the definitions, and explain why the result is allowed.

The move

My working sequence

1. **Translate.** Rewrite the prompt in the language of one- and two-sample inference. List the known quantities, the unknown, the units or types, and any hidden constraint.
2. **Choose.** State the theorem, model, algorithm, accounting rule, or physical law before substituting values. This is where I make the assumption visible.
3. **Execute.** Work one justified step at a time and keep intermediate state readable enough that I can audit it later.
4. **Verify.** Check a boundary case, sign, unit, invariant, order of magnitude, or alternative derivation. Then translate the result back into an ordinary sentence.

Details I would refuse to skip

- **Definitions:** I should be able to define every technical noun in the chapter without using that same noun in the definition.
- **Assumptions:** I write the condition that makes the method legal beside the method itself. A correct procedure under the wrong assumptions is still a wrong answer.
- **Representation:** I move freely among words, equations or code, a diagram, and a small numerical example. If one representation disagrees with another, that disagreement is useful evidence.
- **Failure modes:** I keep the nearest counterexample in mind: the boundary where the rule changes, the invariant breaks, or the model stops matching reality.
- **Communication:** My final answer names what the result means and what it does not establish. I do not let notation do the explanatory work for me.

Sanity check

Closed-note check. Can I teach one- and two-sample inference on a blank page, derive or reconstruct its main rule, work a fresh example, and diagnose a deliberately broken solution? If I can only recognize the finished work, I am not done.

How this chapter connects

I read this chapter as the bridge from hypothesis testing and error tradeoffs to analysis of variance and simple regression. The earlier material supplies the language and constraints; this chapter adds a new reasoning move; the later material asks me to use that move without being told its name. That connection matters because exams and real projects mix chapters on purpose.

The full working model

- **Core claim.** Paired and independent designs have different sampling structures.
- **What moves.** Mean, proportion, and variance comparisons require appropriate standard errors, degrees of freedom, pooling assumptions, and confidence or test framing.

- **What keeps the answer honest.** The model becomes useful only when I can apply one- and two-sample inference to an unfamiliar case and defend the assumptions.
- **Mastery standard.** I can define the objects, rebuild the rule, work a fresh case, compare alternatives, and diagnose a broken solution without relying on recognition.

How I actually think about it

My shorthand for this chapter is: my shorthand for this chapter is: a paired design can remove person-to-person noise, but only if pairs are real and preserved in analysis.

The easy version

Paired and independent designs have different sampling structures.

Worked example

Question. Why not use an independent two-sample test for before-and-after measurements on the same people?
Walkthrough. The observations are correlated within person. Analyze differences so that pairing enters the variance and the estimand matches within-person change.

Common miss

The fastest way to get this chapter wrong is to use the visible procedure without naming its assumption, units, time period, reference group, or boundary. I put that condition beside the work and verify the result independently.

Solved chapter practice

Practice 1. Why not use an independent two-sample test for before-and-after measurements on the same people?
Solution. The observations are correlated within person. Analyze differences so that pairing enters the variance and the estimand matches within-person change.

Practice 2. Explain one- and two-sample inference to a classmate using one definition, one concrete example, one assumption, and one failure mode.
Solution. Begin with the core claim above, use the worked case as the concrete example, state the condition highlighted by the warning, and choose a failure where that condition does not hold. A complete explanation connects all four pieces instead of listing them.

Mastery practice A. Reconstruct one- and two-sample inference from a blank page. Include the central definition, the governing rule, one required assumption, and one boundary case.
Solution. Start from the blue idea box, but rewrite it in my own words. Identify the objects and state, name the rule before using it, copy the relevant assumption from the details audit, and choose a boundary where that assumption becomes equality or fails. A complete solution explains why each part belongs; a list of vocabulary is not enough.

Mastery practice B. A classmate gets a polished-looking answer to a one- and two-sample inference problem but never states a model or checks the result. Write the shortest useful review of that solution.

Solution. I would ask four concrete questions: What exactly are the inputs and desired output? Which rule licenses the main step? Which assumption could invalidate it? Which independent check supports the conclusion? If the answer cannot survive those four questions, I treat the numerical or verbal result as unverified.

CHAPTER 7 Analysis of variance and simple regression

The idea

ANOVA partitions variability into explained and residual components and tests equality of several means. Simple regression expresses the same projection logic with a quantitative predictor; F and t tests connect under special cases.

Rebuild the chapter from first principles

The claim I need to own is this: ANOVA partitions variability into explained and residual components and tests equality of several means. Simple regression expresses the same projection logic with a quantitative predictor; F and t tests connect under special cases. I do not count that as learned just because the sentence looks familiar. I should be able to name the objects, state what changes and what stays fixed, recover the rule from the definitions, and explain why the result is allowed.

The move

My working sequence

1. **Translate.** Rewrite the prompt in the language of analysis of variance and simple regression. List the known quantities, the unknown, the units or types, and any hidden constraint.
2. **Choose.** State the theorem, model, algorithm, accounting rule, or physical law before substituting values. This is where I make the assumption visible.
3. **Execute.** Work one justified step at a time and keep intermediate state readable enough that I can audit it later.
4. **Verify.** Check a boundary case, sign, unit, invariant, order of magnitude, or alternative derivation. Then translate the result back into an ordinary sentence.

Details I would refuse to skip

- **Definitions:** I should be able to define every technical noun in the chapter without using that same noun in the definition.
- **Assumptions:** I write the condition that makes the method legal beside the method itself. A correct procedure under the wrong assumptions is still a wrong answer.
- **Representation:** I move freely among words, equations or code, a diagram, and a small numerical example. If one representation disagrees with another, that disagreement is useful evidence.

- **Failure modes:** I keep the nearest counterexample in mind: the boundary where the rule changes, the invariant breaks, or the model stops matching reality.
- **Communication:** My final answer names what the result means and what it does not establish. I do not let notation do the explanatory work for me.

Sanity check

Closed-note check. Can I teach analysis of variance and simple regression on a blank page, derive or reconstruct its main rule, work a fresh example, and diagnose a deliberately broken solution? If I can only recognize the finished work, I am not done.

How this chapter connects

I read this chapter as the bridge from one- and two-sample inference to survey sampling, resampling, and robust judgment. The earlier material supplies the language and constraints; this chapter adds a new reasoning move; the later material asks me to use that move without being told its name. That connection matters because exams and real projects mix chapters on purpose.

The full working model

- **Core claim.** ANOVA partitions variability into explained and residual components and tests equality of several means.
- **What moves.** Simple regression expresses the same projection logic with a quantitative predictor; F and t tests connect under special cases.
- **What keeps the answer honest.** The model becomes useful only when I can apply analysis of variance and simple regression to an unfamiliar case and defend the assumptions.
- **Mastery standard.** I can define the objects, rebuild the rule, work a fresh case, compare alternatives, and diagnose a broken solution without relying on recognition.

How I actually think about it

ANOVA tells me whether some group difference exists; planned contrasts tell me which scientific comparison I care about.

The easy version

ANOVA partitions variability into explained and residual components and tests equality of several means.

Worked example

Question. What does a significant one-way ANOVA establish?
Walkthrough. It rejects equality of all group means. It does not identify which means differ or whether the differences are practically important; follow with planned or multiplicity-adjusted comparisons.

Common miss

The fastest way to get this chapter wrong is to use the visible procedure without naming its assumption, units, time period, reference group, or boundary. I put that condition beside the work and verify the result independently.

Solved chapter practice

Practice 1. What does a significant one-way ANOVA establish?
Solution. It rejects equality of all group means. It does not identify which means differ or whether the differences are practically important; follow with planned or multiplicity-adjusted comparisons.

Practice 2. Explain analysis of variance and simple regression to a classmate using one definition, one concrete example, one assumption, and one failure mode.
Solution. Begin with the core claim above, use the worked case as the concrete example, state the condition highlighted by the warning, and choose a failure where that condition does not hold. A complete explanation connects all four pieces instead of listing them.

Mastery practice A. Reconstruct analysis of variance and simple regression from a blank page. Include the central definition, the governing rule, one required assumption, and one boundary case.
Solution. Start from the blue idea box, but rewrite it in my own words. Identify the objects and state, name the rule before using it, copy the relevant assumption from the details audit, and choose a boundary where that assumption becomes equality or fails. A complete solution explains why each part belongs; a list of vocabulary is not enough.

Mastery practice B. A classmate gets a polished-looking answer to a analysis of variance and simple regression problem but never states a model or checks the result. Write the shortest useful review of that solution.
Solution. I would ask four concrete questions: What exactly are the inputs and desired output? Which rule licenses the main step? Which assumption could invalidate it? Which independent check supports the conclusion? If the answer cannot survive those four questions, I treat the numerical or verbal result as unverified.

CHAPTER 8 Survey sampling, resampling, and robust judgment

The idea

Simple random, stratified, cluster, and unequal-probability designs determine estimators and variance. Bootstrap and permutation methods approximate sampling or null distributions, while nonresponse and frame bias remain design problems.

Rebuild the chapter from first principles

The claim I need to own is this: Simple random, stratified, cluster, and unequal-probability designs determine estimators and variance. Bootstrap and permutation methods approximate sampling or null distributions, while nonresponse and frame bias remain design problems. I do not count that as learned just because the sentence looks familiar. I should be able to name the objects, state what changes and what stays fixed, recover the rule from the definitions, and explain why the result is allowed.

The move

My working sequence

1. **Translate.** Rewrite the prompt in the language of survey sampling, resampling, and robust judgment. List the known quantities, the unknown, the units or types, and any hidden constraint.
2. **Choose.** State the theorem, model, algorithm, accounting rule, or physical law before substituting values. This is where I make the assumption visible.
3. **Execute.** Work one justified step at a time and keep intermediate state readable enough that I can audit it later.
4. **Verify.** Check a boundary case, sign, unit, invariant, order of magnitude, or alternative derivation. Then translate the result back into an ordinary sentence.

Details I would refuse to skip

- **Definitions:** I should be able to define every technical noun in the chapter without using that same noun in the definition.
- **Assumptions:** I write the condition that makes the method legal beside the method itself. A correct procedure under the wrong assumptions is still a wrong answer.
- **Representation:** I move freely among words, equations or code, a diagram, and a small numerical example. If one representation disagrees with another, that disagreement is useful evidence.
- **Failure modes:** I keep the nearest counterexample in mind: the boundary where the rule changes, the invariant breaks, or the model stops matching reality.
- **Communication:** My final answer names what the result means and what it does not establish. I do not let notation do the explanatory work for me.

Sanity check

Closed-note check. Can I teach survey sampling, resampling, and robust judgment on a blank page, derive or reconstruct its main rule, work a fresh example, and diagnose a deliberately broken solution? If I can only recognize the finished work, I am not done.

How this chapter connects

I read this chapter as the bridge from analysis of variance and simple regression to the cumulative course model. The earlier material supplies the language and constraints; this chapter adds a new reasoning move; the later material asks me to use that move without being told its name. That connection matters because exams and real projects mix chapters on purpose.

The full working model

- **Core claim.** Simple random, stratified, cluster, and unequal-probability designs determine estimators and variance.
- **What moves.** Bootstrap and permutation methods approximate sampling or null distributions, while nonresponse and frame bias remain design problems.
- **What keeps the answer honest.** The model becomes useful only when I can apply survey sampling, resampling, and robust judgment to an unfamiliar case and defend the assumptions.
- **Mastery standard.** I can define the objects, rebuild the rule, work a fresh case, compare alternatives, and diagnose a broken solution without relying on recognition.

How I actually think about it

Resampling cannot repair a sample that systematically misses the population I claim to study.

The easy version

Simple random, stratified, cluster, and unequal-probability designs determine estimators and variance.

Worked example

Question. Why can stratification improve precision?
Walkthrough. If units are homogeneous within strata and different across strata, sampling each stratum controls representation and removes between-stratum variation from parts of the estimator variance.

Common miss

The fastest way to get this chapter wrong is to use the visible procedure without naming its assumption, units, time period, reference group, or boundary. I put that condition beside the work and verify the result independently.

Solved chapter practice

Practice 1. Why can stratification improve precision?
Solution. If units are homogeneous within strata and different across strata, sampling each stratum controls representation and removes between-stratum variation from parts of the estimator variance.

Practice 2. Explain survey sampling, resampling, and robust judgment to a classmate using one definition, one concrete example, one assumption, and one failure mode.
Solution. Begin with the core claim above, use the worked case as the concrete example, state the condition highlighted by the warning, and choose a failure where that condition does not hold. A complete explanation connects all four pieces instead of listing them.

Mastery practice A. Reconstruct survey sampling, resampling, and robust judgment from a blank page. Include the central definition, the governing rule, one required assumption, and one boundary case.

Solution. Start from the blue idea box, but rewrite it in my own words. Identify the objects and state, name the rule before using it, copy the relevant assumption from the details audit, and choose a boundary where that assumption becomes equality or fails. A complete solution explains why each part belongs; a list of vocabulary is not enough.

Mastery practice B. A classmate gets a polished-looking answer to a survey sampling, resampling, and robust judgment problem but never states a model or checks the result. Write the shortest useful review of that solution.

Solution. I would ask four concrete questions: What exactly are the inputs and desired output? Which rule licenses the main step? Which assumption could invalidate it? Which independent check supports the conclusion? If the answer cannot survive those four questions, I treat the numerical or verbal result as unverified.

Cumulative study plan

I use this packet in three passes. First I reconstruct each model and work its examples without looking. Second I mix chapters so the tool is no longer named for me. Third I explain a complete case aloud, including assumptions, units, uncertainty, failure modes, and what evidence would change my conclusion.