UCLA-ALIGNED COURSE FIELD NOTES

STATS 101B

Introduction to Design and Analysis of Experiment

Randomization, blocking, factorial structure, ANOVA, contrasts, power, and reproducible experimental reasoning

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	**Experimental units, treatments, and causal estimands**	Reconstruct the model, work an application, test its assumptions, and explain the result in ordinary language.
2	**Randomization, replication, and control**	Reconstruct the model, work an application, test its assumptions, and explain the result in ordinary language.
3	**Completely randomized designs and one-way ANOVA**	Reconstruct the model, work an application, test its assumptions, and explain the result in ordinary language.
4	**Contrasts, multiple comparisons, and familywise error**	Reconstruct the model, work an application, test its assumptions, and explain the result in ordinary language.
5	**Blocking, matching, and randomized block designs**	Reconstruct the model, work an application, test its assumptions, and explain the result in ordinary language.
6	**Factorial experiments and interaction**	Reconstruct the model, work an application, test its assumptions, and explain the result in ordinary language.
7	**Power, effect size, and sample-size planning**	Reconstruct the model, work an application, test its assumptions, and explain the result in ordinary language.
8	**Diagnostics, missing outcomes, and experimental reporting**	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/STATS101B. 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

CHAPTER 1 Experimental units, treatments, and causal estimands

The idea

An experiment defines units, treatments, outcomes, assignment, and the causal contrast before data collection. Potential outcomes clarify that an individual treatment effect is unobservable and that assignment creates comparable groups in expectation.

Rebuild the chapter from first principles

The claim I need to own is this: An experiment defines units, treatments, outcomes, assignment, and the causal contrast before data collection. Potential outcomes clarify that an individual treatment effect is unobservable and that assignment creates comparable groups in expectation. 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 experimental units, treatments, and causal estimands. 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 experimental units, treatments, and causal estimands 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 randomization, replication, and control. 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.** An experiment defines units, treatments, outcomes, assignment, and the causal contrast before data collection.
- **What moves.** Potential outcomes clarify that an individual treatment effect is unobservable and that assignment creates comparable groups in expectation.
- **What keeps the answer honest.** The model becomes useful only when I can apply experimental units, treatments, and causal estimands 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 label the experimental unit before discussing sample size; repeated measurements are not automatically new independent units.

The easy version

An experiment defines units, treatments, outcomes, assignment, and the causal contrast before data collection.

Worked example

Question. Twenty classrooms receive one teaching method each and 600 students are measured. What is the assignment unit?
Walkthrough. The classroom is the assignment unit. Student outcomes provide information, but analysis must account for clustering within the 20 randomized classrooms.

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. Twenty classrooms receive one teaching method each and 600 students are measured. What is the assignment unit?
Solution. The classroom is the assignment unit. Student outcomes provide information, but analysis must account for clustering within the 20 randomized classrooms.

Practice 2. Explain experimental units, treatments, and causal estimands 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 experimental units, treatments, and causal estimands 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 experimental units, treatments, and causal estimands 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 Randomization, replication, and control

The idea

Random assignment protects against systematic confounding; replication estimates variability; controls define the counterfactual; blinding limits behavior and measurement bias. Randomization inference connects the assignment mechanism to a null distribution.

Rebuild the chapter from first principles

The claim I need to own is this: Random assignment protects against systematic confounding; replication estimates variability; controls define the counterfactual; blinding limits behavior and measurement bias. Randomization inference connects the assignment mechanism to a null distribution. 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 randomization, replication, and control. 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 randomization, replication, and control 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 experimental units, treatments, and causal estimands to completely randomized designs and one-way anova. 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.** Random assignment protects against systematic confounding; replication estimates variability; controls define the counterfactual; blinding limits behavior and measurement bias.
- **What moves.** Randomization inference connects the assignment mechanism to a null distribution.

- **What keeps the answer honest.** The model becomes useful only when I can apply randomization, replication, and control 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

The randomization procedure is part of the analysis, so I preserve exactly how treatment was assigned.

The easy version

Random assignment protects against systematic confounding; replication estimates variability; controls define the counterfactual; blinding limits behavior and measurement bias.

Worked example

Question. Why is alternating treatment assignment not equivalent to random assignment?
Walkthrough. The sequence is predictable, allowing time trends or investigator choices to align with treatment. A genuine random mechanism prevents that systematic alignment.

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 is alternating treatment assignment not equivalent to random assignment?
Solution. The sequence is predictable, allowing time trends or investigator choices to align with treatment. A genuine random mechanism prevents that systematic alignment.

Practice 2. Explain randomization, replication, and control 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 randomization, replication, and control 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 randomization, replication, and control 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 Completely randomized designs and one-way ANOVA

The idea

A completely randomized design assigns treatments across comparable units. One-way ANOVA partitions total variation into between-treatment and within-treatment components, and its F statistic compares signal with residual noise.

Rebuild the chapter from first principles

The claim I need to own is this: A completely randomized design assigns treatments across comparable units. One-way ANOVA partitions total variation into between-treatment and within-treatment components, and its F statistic compares signal with residual noise. 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 completely randomized designs and one-way anova. 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 completely randomized designs and one-way anova 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 randomization, replication, and control to contrasts, multiple comparisons, and familywise error. 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 completely randomized design assigns treatments across comparable units.
- **What moves.** One-way ANOVA partitions total variation into between-treatment and within-treatment components, and its F statistic compares signal with residual noise.
- **What keeps the answer honest.** The model becomes useful only when I can apply completely randomized designs and one-way anova 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

The omnibus F test answers whether some mean differs; it does not tell me which scientific contrast matters.

The easy version

A completely randomized design assigns treatments across comparable units.

Worked example

Question. What does rejecting the one-way ANOVA null establish?
Walkthrough. It provides evidence that not all treatment means are equal under the model. Planned or multiplicity-adjusted contrasts are needed to locate meaningful differences.

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 rejecting the one-way ANOVA null establish?
Solution. It provides evidence that not all treatment means are equal under the model. Planned or multiplicity-adjusted contrasts are needed to locate meaningful differences.

Practice 2. Explain completely randomized designs and one-way anova 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 completely randomized designs and one-way anova 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 completely randomized designs and one-way anova 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 Contrasts, multiple comparisons, and familywise error

The idea

Contrasts encode targeted linear comparisons among treatment means. Orthogonality can separate questions, while Tukey, Bonferroni, Holm, and false-discovery procedures manage different error goals when many claims are tested.

Rebuild the chapter from first principles

The claim I need to own is this: Contrasts encode targeted linear comparisons among treatment means. Orthogonality can separate questions, while Tukey, Bonferroni, Holm, and false-discovery procedures manage different error goals when many claims are tested. 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 contrasts, multiple comparisons, and familywise error. 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 contrasts, multiple comparisons, and familywise error 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 completely randomized designs and one-way anova to blocking, matching, and randomized block designs. 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.** Contrasts encode targeted linear comparisons among treatment means.
- **What moves.** Orthogonality can separate questions, while Tukey, Bonferroni, Holm, and false-discovery procedures manage different error goals when many claims are tested.
- **What keeps the answer honest.** The model becomes useful only when I can apply contrasts, multiple comparisons, and familywise error 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 write planned contrasts from the research question before looking at means; otherwise the data quietly choose the hypothesis.

The easy version

Contrasts encode targeted linear comparisons among treatment means.

Worked example

Question. Why does testing ten pairs at 5 percent inflate false-positive risk?
Walkthrough. Each test offers another chance to reject under a global null. Familywise procedures adjust thresholds or p-values to control the probability of at least one false rejection.

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 does testing ten pairs at 5 percent inflate false-positive risk?
Solution. Each test offers another chance to reject under a global null. Familywise procedures adjust thresholds or p-values to control the probability of at least one false rejection.

Practice 2. Explain contrasts, multiple comparisons, and familywise error 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 contrasts, multiple comparisons, and familywise error 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 contrasts, multiple comparisons, and familywise error 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 Blocking, matching, and randomized block designs

The idea

Blocking groups similar units before randomization so nuisance variation is removed from error. Matched pairs are blocks of size two; treatment comparisons occur within blocks, and treatment-by-block interaction determines how well additivity works.

Rebuild the chapter from first principles

The claim I need to own is this: Blocking groups similar units before randomization so nuisance variation is removed from error. Matched pairs are blocks of size two; treatment comparisons occur within blocks, and treatment-by-block interaction determines how well additivity works. 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 blocking, matching, and randomized block designs. 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 blocking, matching, and randomized block designs 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 contrasts, multiple comparisons, and familywise error to factorial experiments and interaction. 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.** Blocking groups similar units before randomization so nuisance variation is removed from error.
- **What moves.** Matched pairs are blocks of size two; treatment comparisons occur within blocks, and treatment-by-block interaction determines how well additivity works.
- **What keeps the answer honest.** The model becomes useful only when I can apply blocking, matching, and randomized block designs 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

Blocking spends design effort on a source of variation I do not care about so treatment differences become easier to see.

The easy version

Blocking groups similar units before randomization so nuisance variation is removed from error.

Worked example

Question. When is blocking likely to improve precision?
Walkthrough. When units within a block are similar on outcome-relevant factors and blocks differ meaningfully. A weak or irrelevant blocking variable may add complexity without reducing residual 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. When is blocking likely to improve precision?
Solution. When units within a block are similar on outcome-relevant factors and blocks differ meaningfully. A weak or irrelevant blocking variable may add complexity without reducing residual variance.

Practice 2. Explain blocking, matching, and randomized block designs 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 blocking, matching, and randomized block designs 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 blocking, matching, and randomized block designs 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 Factorial experiments and interaction

The idea

Factorial designs vary multiple factors together, estimating main effects and interactions efficiently. An interaction means one factor's effect depends on another; hierarchy guides model construction and plots make conditional effects visible.

Rebuild the chapter from first principles

The claim I need to own is this: Factorial designs vary multiple factors together, estimating main effects and interactions efficiently. An interaction means one factor's effect depends on another; hierarchy guides model construction and plots make conditional effects visible. 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 factorial experiments and interaction. 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 factorial experiments and interaction 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 blocking, matching, and randomized block designs to power, effect size, and sample-size planning. 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.** Factorial designs vary multiple factors together, estimating main effects and interactions efficiently.
- **What moves.** An interaction means one factor's effect depends on another; hierarchy guides model construction and plots make conditional effects visible.
- **What keeps the answer honest.** The model becomes useful only when I can apply factorial experiments and interaction 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

If interaction is real, I stop telling one average main-effect story and compare the factor at each meaningful level.

The easy version

Factorial designs vary multiple factors together, estimating main effects and interactions efficiently.

Worked example

Question. A fertilizer helps only under high irrigation. What pattern is this?
Walkthrough. It is a fertilizer-by-irrigation interaction. The fertilizer effect must be reported separately by irrigation level rather than as one universal average.

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 fertilizer helps only under high irrigation. What pattern is this?
Solution. It is a fertilizer-by-irrigation interaction. The fertilizer effect must be reported separately by irrigation level rather than as one universal average.

Practice 2. Explain factorial experiments and interaction 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 factorial experiments and interaction 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 factorial experiments and interaction 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 Power, effect size, and sample-size planning

The idea

Power is the probability of detecting a specified effect under a design and analysis. It depends on effect size, variance, sample size, allocation, clustering, multiplicity, and significance level; planning requires a scientifically meaningful target.

Rebuild the chapter from first principles

The claim I need to own is this: Power is the probability of detecting a specified effect under a design and analysis. It depends on effect size, variance, sample size, allocation, clustering, multiplicity, and significance level; planning requires a scientifically meaningful target. 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 power, effect size, and sample-size planning. 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 power, effect size, and sample-size planning 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 factorial experiments and interaction to diagnostics, missing outcomes, and experimental reporting. 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.** Power is the probability of detecting a specified effect under a design and analysis.
- **What moves.** It depends on effect size, variance, sample size, allocation, clustering, multiplicity, and significance level; planning requires a scientifically meaningful target.

- **What keeps the answer honest.** The model becomes useful only when I can apply power, effect size, and sample-size planning 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 choose the smallest effect worth acting on before running a power calculation, not after seeing what sample size is convenient.

The easy version

Power is the probability of detecting a specified effect under a design and analysis.

Worked example

Question. What happens to power when residual variance falls while all else stays fixed?
Walkthrough. Power rises because the treatment signal is larger relative to noise. Blocking or better measurement can therefore increase power without adding units.

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 happens to power when residual variance falls while all else stays fixed?
Solution. Power rises because the treatment signal is larger relative to noise. Blocking or better measurement can therefore increase power without adding units.

Practice 2. Explain power, effect size, and sample-size planning 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 power, effect size, and sample-size planning 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 power, effect size, and sample-size planning 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 Diagnostics, missing outcomes, and experimental reporting

The idea

ANOVA assumptions concern independent errors, variance structure, and model form. Residual diagnostics, transformation, robust alternatives, attrition analysis, protocol deviations, effect intervals, and preregistered reporting protect the causal story.

Rebuild the chapter from first principles

The claim I need to own is this: ANOVA assumptions concern independent errors, variance structure, and model form. Residual diagnostics, transformation, robust alternatives, attrition analysis, protocol deviations, effect intervals, and preregistered reporting protect the causal story. 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 diagnostics, missing outcomes, and experimental reporting. 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 diagnostics, missing outcomes, and experimental reporting 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 power, effect size, and sample-size planning 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.** ANOVA assumptions concern independent errors, variance structure, and model form.
- **What moves.** Residual diagnostics, transformation, robust alternatives, attrition analysis, protocol deviations, effect intervals, and preregistered reporting protect the causal story.
- **What keeps the answer honest.** The model becomes useful only when I can apply diagnostics, missing outcomes, and experimental reporting 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: randomization begins credibility; transparent handling of attrition and deviations has to carry it through the final table.

The easy version

ANOVA assumptions concern independent errors, variance structure, and model form.

Worked example

Question. Why can treatment-dependent dropout bias an experiment?
Walkthrough. Observed outcomes may no longer represent the randomized groups. Compare attrition, investigate reasons, use defensible missing-data methods, and report sensitivity to plausible unseen outcomes.

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 treatment-dependent dropout bias an experiment?
Solution. Observed outcomes may no longer represent the randomized groups. Compare attrition, investigate reasons, use defensible missing-data methods, and report sensitivity to plausible unseen outcomes.

Practice 2. Explain diagnostics, missing outcomes, and experimental reporting 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 diagnostics, missing outcomes, and experimental reporting 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 diagnostics, missing outcomes, and experimental reporting 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.