

Experimental Design Basics

Floyd Bullard

Thursday, 16 November 2006

22 May 2006

The Essential
Ingredients

Hermit crab study

Control

Replication

Randomization

Summary

Blocking

The non-block and
the block design

Crab example

Summary

Bugaboos, or
FAQs

We'll consider a study on hermit crabs. The study described in this presentation was motivated by an actual study performed by a Duke University undergraduate at Duke's Marine Laboratory, but the experimental design is not the same and the data given here are made up.

The Essential
Ingredients

Hermit crab study

Control

Replication

Randomization

Summary

Blocking

The non-block and
the block design

Crab example

Summary

Bugaboos, or
FAQs

purpose of the study

The Essential Ingredients

Hermit crab study

Control

Replication

Randomization

Summary

Blocking

The non-block and
the block design

Crab example

Summary

Bugaboos, or FAQs

Does the “snugness” of a hermit crab’s (secondary) shell affect the crab’s insulation from temperature changes in the surrounding water?

essential question

Given shells of two different “snuggness” levels (let’s just call them “snug” and “loose”), do their temperature differentials (interior minus exterior) in warm exterior water differ from one another?

The Essential Ingredients

Hermit crab study

Control

Replication

Randomization

Summary

Blocking

The non-block and
the block design

Crab example

Summary

Bugaboos, or FAQs

an ideal world

In an ideal world, all hermit crabs would be exactly the same in every way, down to their very atoms.




Figure: An ideal world

an ideal experiment

Recall that we want to see whether a snug shell and a loose shell will lead to different temperature differentials for hermit crabs. In our ideal world, how would we design our experiment?

The Essential Ingredients

Hermit crab study

Control

Replication

Randomization

Summary

Blocking

The non-block and
the block design

Crab example

Summary

Bugaboos, or FAQs

an ideal experiment

Recall that we want to see whether a snug shell and a loose shell will lead to different temperature differentials for hermit crabs. In our ideal world, how would we design our experiment?



snug shell loose shell

Figure: Our experiment requires only two crabs

an ideal experiment

Recall that we want to see whether a snug shell and a loose shell will lead to different temperature differentials for hermit crabs. In our ideal world, how would we design our experiment?

$$\Delta T_1 = 1.4^\circ \quad \Delta T_2 = 2.1^\circ$$



snug shell



loose shell

Figure: Our experiment requires only two crabs

the real world

But alas: all crabs are not the same.



Figure: An non-ideal world

The Essential Ingredients

Hermit crab study

Control
Replication
Randomization
Summary

Blocking

The non-block and
the block design
Crab example
Summary

Bugaboos, or FAQs

Statistics exists because there is **variability** in all things.

The Essential
Ingredients

Hermit crab study

Control

Replication

Randomization

Summary

Blocking

The non-block and
the block design

Crab example

Summary

Bugaboos, or
FAQs

Statistics exists because there is **variability** in all things.

Trying to deal with variability is what statistics is largely about.

The Essential
Ingredients

Hermit crab study

Control

Replication

Randomization

Summary

Blocking

The non-block and
the block design

Crab example

Summary

Bugaboos, or
FAQs

essential ingredients

A common mantra about the essential ingredients in a good experiment is

The Essential Ingredients

Hermit crab study

Control

Replication

Randomization

Summary

Blocking

The non-block and
the block design

Crab example

Summary

Bugaboos, or FAQs

A common mantra about the essential ingredients in a good experiment is *randomization*,

The Essential Ingredients

Hermit crab study

Control

Replication

Randomization

Summary

Blocking

The non-block and
the block design

Crab example

Summary

Bugaboos, or FAQs

A common mantra about the essential ingredients in a good experiment is *randomization, replication,*

The Essential
Ingredients

Hermit crab study

Control

Replication

Randomization

Summary

Blocking

The non-block and
the block design

Crab example

Summary

Bugaboos, or
FAQs

A common mantra about the essential ingredients in a good experiment is *randomization, replication, and control*.

The Essential
Ingredients

Hermit crab study

Control

Replication

Randomization

Summary

Blocking

The non-block and
the block design

Crab example

Summary

Bugaboos, or
FAQs

A common mantra about the essential ingredients in a good experiment is *randomization, replication, and control*.

What do these things really mean? Let's look at them one by one.

The Essential Ingredients

Hermit crab study

Control

Replication

Randomization

Summary

Blocking

The non-block and
the block design

Crab example

Summary

Bugaboos, or FAQs

control = minimizing variability

As we have seen, in the ideal experiment, everything would be the same in the treatment groups except for the treatment itself.

The Essential Ingredients

Hermit crab study

Control

Replication

Randomization

Summary

Blocking

The non-block and
the block design

Crab example

Summary

Bugaboos, or FAQs

control = minimizing variability

As we have seen, in the ideal experiment, everything would be the same in the treatment groups except for the treatment itself.

Control means trying as hard as you can to make that true: minimizing variability.

The Essential
Ingredients

Hermit crab study

Control

Replication

Randomization

Summary

Blocking

The non-block and
the block design

Crab example

Summary

Bugaboos, or
FAQs

control = minimizing variability

As we have seen, in the ideal experiment, everything would be the same in the treatment groups except for the treatment itself.

Control means trying as hard as you can to make that true: minimizing variability. That means treating all crabs in exactly the same way, as much as possible. They get the same food, the same salinity in their water, *etc.*

The Essential
Ingredients

Hermit crab study

Control

Replication

Randomization

Summary

Blocking

The non-block and
the block design

Crab example

Summary

Bugaboos, or
FAQs

control = minimizing variability

As we have seen, in the ideal experiment, everything would be the same in the treatment groups except for the treatment itself.

Control means trying as hard as you can to make that true: minimizing variability. That means treating all crabs in exactly the same way, as much as possible. They get the same food, the same salinity in their water, *etc.*

(Control in this context does *not* refer to having a “control group”. That is not an essential ingredient to a good experiment.)

The Essential Ingredients

Hermit crab study

Control

Replication

Randomization

Summary

Blocking

The non-block and
the block design

Crab example

Summary

Bugaboos, or FAQs

even if our control were perfect

Let's suppose it is possible for us to exercise *perfect* control.
That is, suppose we are able to treat crabs identically in every single way except for the treatment we impose.

The Essential
Ingredients

Hermit crab study

Control

Replication

Randomization

Summary

Blocking

The non-block and
the block design

Crab example

Summary

Bugaboos, or
FAQs

even if our control were perfect

Let's suppose it is possible for us to exercise *perfect* control.
That is, suppose we are able to treat crabs identically in
every single way except for the treatment we impose.

$$\Delta T_1 = 1.4^\circ \quad \Delta T_2 = 2.1^\circ$$



snug shell



loose shell

Figure: What can we conclude?

treatment? crabs?

With only two crabs in our study—and knowing that they aren't perfectly identical—we can't know whether a response difference is due to the different treatments we imposed or to the inherent differences in crabs.

The Essential
Ingredients

Hermit crab study

Control

Replication

Randomization

Summary

Blocking

The non-block and
the block design

Crab example

Summary

Bugaboos, or
FAQs

treatment? crabs?

With only two crabs in our study—and knowing that they aren't perfectly identical—we can't know whether a response difference is due to the different treatments we imposed or to the inherent differences in crabs.

(Or due to the different ways we handled the crabs. After all, we can't *really* exercise “perfect control” over all external variables.)

The Essential Ingredients

Hermit crab study

Control

Replication

Randomization

Summary

Blocking

The non-block and
the block design

Crab example

Summary

Bugaboos, or FAQs

some variability we can't control

So we control as much variability as we can, trying to ensure that the only differences between the way our crabs are treated is the snugness of the shell.

But some variability (in particular, crabs) we just can't control.

The Essential
Ingredients

Hermit crab study

Control

Replication

Randomization

Summary

Blocking

The non-block and
the block design

Crab example

Summary

Bugaboos, or
FAQs

With only one crab receiving each treatment, there's just no way we'll ever know for sure that a different response is due to the different treatments we impose, or to inherent differences between the crabs.

The Essential Ingredients

Hermit crab study
Control

Replication

Randomization
Summary

Blocking

The non-block and
the block design
Crab example
Summary

Bugaboos, or FAQs

With only one crab receiving each treatment, there's just no way we'll ever know for sure that a different response is due to the different treatments we impose, or to inherent differences between the crabs.

But what if we gave our different treatments to multiple crabs, and the responses were consistent within a group, but different from group to group?

The Essential Ingredients

Hermit crab study
Control

Replication

Randomization
Summary

Blocking

The non-block and
the block design
Crab example
Summary

Bugaboos, or FAQs

this would be nice

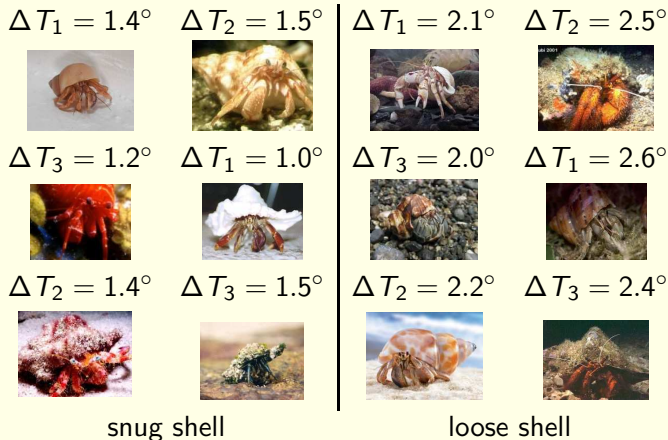


Figure: What can we conclude?

The Essential
Ingredients

Hermit crab study

Control

Replication

Randomization

Summary

Blocking

The non-block and
the block design

Crab example

Summary

Bugaboos, or
FAQs

Replication refers to having multiple **independent** experimental units in each treatment group.

The Essential Ingredients

Hermit crab study

Control

Replication

Randomization

Summary

Blocking

The non-block and
the block design

Crab example

Summary

Bugaboos, or FAQs

Replication refers to having multiple **independent** experimental units in each treatment group.

Why is **independent** in boldface? Because it's a “bugaboo”. We'll come back to that later in the talk.

The Essential Ingredients

Hermit crab study

Control

Replication

Randomization

Summary

Blocking

The non-block and
the block design

Crab example

Summary

Bugaboos, or FAQs

Replication refers to having multiple **independent** experimental units in each treatment group.

Why is **independent** in boldface? Because it's a “bugaboo”. We'll come back to that later in the talk.

(Replication in this context does *not* refer to replicating the entire experiment.)

The Essential
Ingredients

Hermit crab study
Control

Replication

Randomization
Summary

Blocking

The non-block and
the block design
Crab example
Summary

Bugaboos, or
FAQs

one problem remains

The Essential
Ingredients

Hermit crab study

Control

Replication

Randomization

Summary

Blocking

The non-block and
the block design

Crab example

Summary

Bugaboos, or
FAQs

What if there are *systematic* differences between the two groups?

systematic group differences are problematic

We want to be sure that the only difference between our treatment groups is the treatments themselves. It's hard to guarantee this when there may be invisible crab traits that could end up congregating in one group.

The Essential Ingredients

Hermit crab study

Control

Replication

Randomization

Summary

Blocking

The non-block and
the block design

Crab example

Summary

Bugaboos, or FAQs

By *randomly* allocating crabs to treatment groups, we guarantee* that the groups are not systematically different. Even traits that are invisible to us should be equalized between the two groups.

The Essential Ingredients

Hermit crab study

Control

Replication

Randomization

Summary

Blocking

The non-block and
the block design

Crab example

Summary

Bugaboos, or FAQs

By *randomly* allocating crabs to treatment groups, we guarantee* that the groups are not systematically different. Even traits that are invisible to us should be equalized between the two groups.

(*It is of course still *possible* for there to be, by chance, systematic group differences even after randomly allocating treatments. But that's where probability enters into the inference picture. And we're not going to get into that tonight.)

The Essential
Ingredients

Hermit crab study

Control

Replication

Randomization

Summary

Blocking

The non-block and
the block design

Crab example

Summary

Bugaboos, or
FAQs

randomization

So randomization means randomly allocating experimental units to treatment groups. (Or, equivalently, randomly allocating treatments to experimental units.)

The Essential
Ingredients

Hermit crab study

Control

Replication

Randomization

Summary

Blocking

The non-block and
the block design

Crab example

Summary

Bugaboos, or
FAQs

So randomization means randomly allocating experimental units to treatment groups. (Or, equivalently, randomly allocating treatments to experimental units.)

(Randomization in this context does *not* refer to randomly sampling crabs from a crab population! In fact, you should not call this *sampling* at all. If a student writes, “from our 12 crabs we will take a SRS of 6 to be assigned to snug shells...”, then you need to correct that student. Associating random sampling with random treatment allocation can only lead to confusion.)

The Essential
Ingredients

Hermit crab study

Control

Replication

Randomization

Summary

Blocking

The non-block and
the block design

Crab example

Summary

Bugaboos, or
FAQs

We want our treatment groups to be alike in every possible way except for our treatments so that a clear difference in responses can be attributable to the different treatments we imposed.

The Essential Ingredients

Hermit crab study

Control

Replication

Randomization

Summary

Blocking

The non-block and
the block design

Crab example

Summary

Bugaboos, or FAQs

We want our treatment groups to be alike in every possible way except for our treatments so that a clear difference in responses can be attributable to the different treatments we imposed.

- ▶ We *control* as much variability as we can.

The Essential
Ingredients

Hermit crab study

Control

Replication

Randomization

Summary

Blocking

The non-block and
the block design

Crab example

Summary

Bugaboos, or
FAQs

We want our treatment groups to be alike in every possible way except for our treatments so that a clear difference in responses can be attributable to the different treatments we imposed.

- ▶ We *control* as much variability as we can.
- ▶ We have *multiple units within treatment groups* because there will always be differences between any two experimental units.

We want our treatment groups to be alike in every possible way except for our treatments so that a clear difference in responses can be attributable to the different treatments we imposed.

- ▶ We *control* as much variability as we can.
- ▶ We have *multiple units within treatment groups* because there will always be differences between any two experimental units.
- ▶ And we *randomly allocate* units to treatment groups so that any invisible sources of variability will be equalized between the treatment groups.

The Essential
Ingredients

Hermit crab study

Control

Replication

Randomization

Summary

Blocking

The non-block and
the block design

Crab example

Summary

Bugaboos, or
FAQs

what we look for

If we've done all that—control, replication, and randomization—and we see a difference in responses that is “too large to be attributable to chance alone”, then it must be due to some difference between the two groups. And since we've made sure that the only difference between the groups is the treatments we imposed, then we can conclude that the different responses were caused by the different treatments.

The Essential Ingredients

Hermit crab study

Control

Replication

Randomization

Summary

Blocking

The non-block and
the block design

Crab example

Summary

Bugaboos, or FAQs

what we look for

If we've done all that—control, replication, and randomization—and we see a difference in responses that is “too large to be attributable to chance alone”, then it must be due to some difference between the two groups. And since we've made sure that the only difference between the groups is the treatments we imposed, then we can conclude that the different responses were caused by the different treatments.

(Even the best scientists and statisticians can inadvertently allow in confounding variables sometimes, usually something done to the treatment groups after they've been randomly allocated—perhaps in the process of applying the treatments. But the above summary of a good controlled experiment is at least true in principle.)

The Essential
Ingredients

Hermit crab study

Control

Replication

Randomization

Summary

Blocking

The non-block and
the block design

Crab example

Summary

Bugaboos, or
FAQs

Non-block model:

$$Y_{t,i} = \mu_t + e_{t,i}$$
$$e_{t,i} \stackrel{\text{iid}}{\sim} N(0, \sigma_t),$$

where

- ▶ $t = 1, 2$ indexes the treatment groups,
- ▶ $i = 1, \dots, n$ indexes the units within a treatment group,
- ▶ $Y_{t,i}$ is the response measured on the i^{th} unit in group t ,
- ▶ μ_t is the overall, underlying mean response of units in group t , and
- ▶ $e_{t,i}$ is the “error” (a lousy word choice) associated with the i^{th} unit in group t .

The Essential
Ingredients

Hermit crab study
Control
Replication
Randomization
Summary

Blocking

The non-block and
the block design
Crab example
Summary

Bugaboos, or
FAQs

$$Y_{b,t,i} = \tau_b + \mu_t + e_{b,t,i}$$

$$e_{b,t,i} \stackrel{\text{iid}}{\sim} N(0, \sigma_t),$$

where

- ▶ $b = 1, \dots, m$ indexes *blocks*, and
- ▶ τ_b is the underlying mean response associated with units in block b .

The Essential
Ingredients

Hermit crab study

Control

Replication

Randomization

Summary

Blocking

**The non-block and
the block design**

Crab example

Summary

Bugaboos, or
FAQs

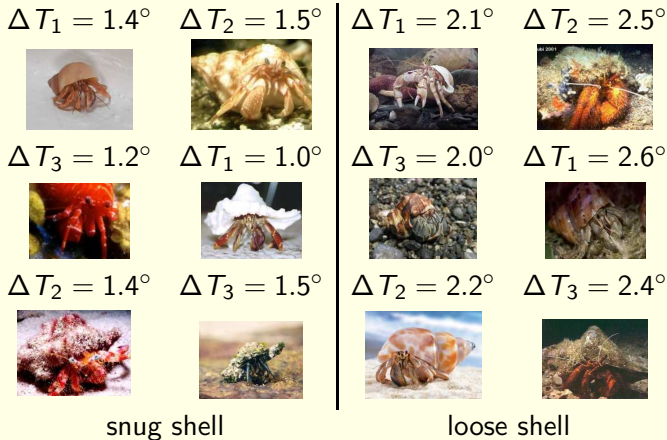


Figure: An obvious difference

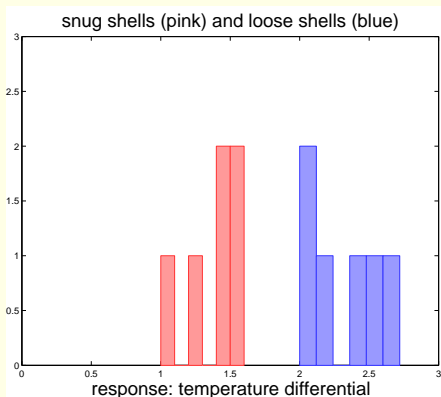


Figure: This is always nice to see.

The Essential Ingredients

Hermit crab study

Control

Replication

Randomization

Summary

Blocking

The non-block and
the block design

Crab example

Summary

Bugaboos, or FAQs

what about this?

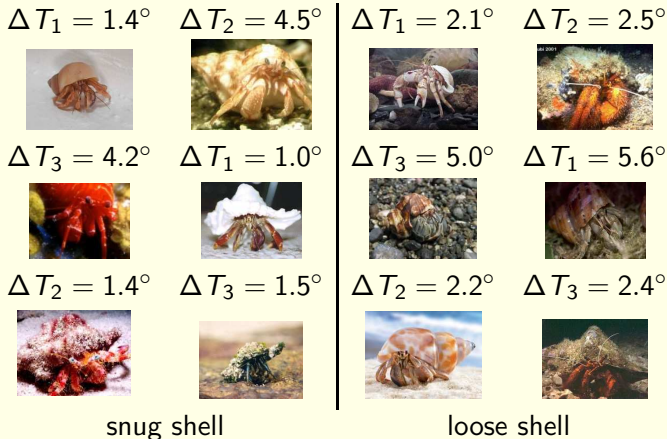


Figure: An obvious difference?

The Essential Ingredients

Hermit crab study
Control
Replication
Randomization
Summary

Blocking

The non-block and
the block design
Crab example
Summary

Bugaboos, or FAQs

not as nice




Figure: Is there still a significant treatment effect?

The Essential
Ingredients

Hermit crab study

Control

Replication

Randomization

Summary

Blocking

The non-block and
the block design

Crab example

Summary

Bugaboos, or
FAQs

Whoops!

I forgot to mention that in the previous two slides, four of the crabs had been observed, *a priori*, to have recently molted. It was thought that this might have an effect on the response variable, so the design was actually *blocked*:

- ▶ The four recently molted crabs were randomly allocated to the two treatment groups, two to each group,
- ▶ And the eight non-recently-molted crabs were randomly allocated to the two treatment groups, four to each group.

Can you guess which of the four crabs had recently molted?

The Essential
Ingredients

Hermit crab study
Control
Replication
Randomization
Summary

Blocking

The non-block and
the block design
Crab example
Summary

Bugaboos, or
FAQs

a blocked design

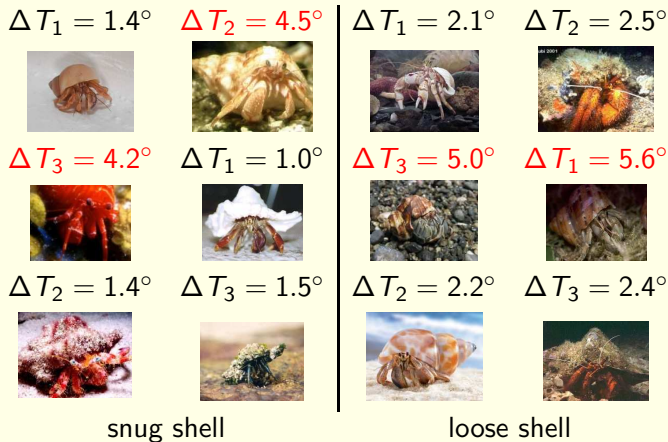


Figure: Responses in red correspond to crabs that had recently molted.

estimating block means τ_b

The four recently-molted crabs had responses of 4.5, 4.2, 5.0, and 5.6. These have a mean of $\hat{\tau}_{molt} \cong 4.825$.

estimating block means τ_b

The four recently-molted crabs had responses of 4.5, 4.2, 5.0, and 5.6. These have a mean of $\hat{\tau}_{molt} \cong 4.825$.

The eight non-recently-molted crabs has responses of 1.4, 1.0, 1.4, 1.5, 2.1, 2.5, 2.2, and 2.4. These have a mean of $\hat{\tau}_{nonmolt} \cong 1.8125$.

The Essential Ingredients

Hermit crab study

Control

Replication

Randomization

Summary

Blocking

The non-block and
the block design

Crab example

Summary

Bugaboos, or FAQs

estimating block means τ_b

The four recently-molted crabs had responses of 4.5, 4.2, 5.0, and 5.6. These have a mean of $\hat{\tau}_{molt} \cong 4.825$.

The eight non-recently-molted crabs has responses of 1.4, 1.0, 1.4, 1.5, 2.1, 2.5, 2.2, and 2.4. These have a mean of $\hat{\tau}_{nonmolt} \cong 1.8125$.

Let's go back and look at the responses again after we *subtract out the estimated block effects*.

The Essential Ingredients

Hermit crab study

Control

Replication

Randomization

Summary

Blocking

The non-block and
the block design

Crab example

Summary

Bugaboos, or FAQs

estimated block effects have been subtracted out

$$\Delta T_1 = -0.41^\circ$$



$$\Delta T_2 = -0.33^\circ$$



$$\Delta T_1 = 0.29^\circ$$



$$\Delta T_2 = 0.69^\circ$$



$$\Delta T_3 = -0.63^\circ$$



$$\Delta T_1 = -0.81^\circ$$



$$\Delta T_3 = 0.18^\circ$$



$$\Delta T_1 = 0.78^\circ$$



$$\Delta T_2 = -0.41^\circ$$



$$\Delta T_3 = -0.31^\circ$$



$$\Delta T_2 = 0.39^\circ$$



$$\Delta T_3 = 0.59^\circ$$



snug shell

loose shell

Figure: Responses in red correspond to crabs that had recently molted.

The Essential Ingredients

Hermit crab study

Control

Replication

Randomization

Summary

Blocking

The non-block and the block design

Crab example

Summary

Bugaboos, or FAQs

block effects removed

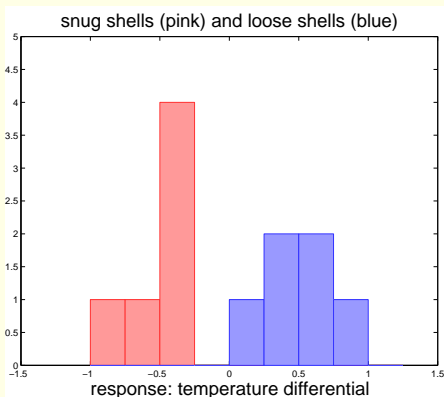


Figure: Is there a significant treatment effect?

The Essential Ingredients

Hermit crab study

Control

Replication

Randomization

Summary

Blocking

The non-block and
the block design

Crab example

Summary

Bugaboos, or FAQs

blocking: the upshot

Blocking is used when there are observable differences in the experimental units that are believed, a priori, to have an influence on the response variable.

The Essential Ingredients

Hermit crab study
Control
Replication
Randomization
Summary

Blocking

The non-block and
the block design
Crab example
Summary

Bugaboos, or FAQs

blocking: the upshot

Blocking is used when there are observable differences in the experimental units that are believed, a priori, to have an influence on the response variable.

Experimental units are randomly allocated to treatment groups *within blocks*.

The Essential
Ingredients

Hermit crab study
Control
Replication
Randomization
Summary

Blocking

The non-block and
the block design
Crab example
Summary

Bugaboos, or
FAQs

blocking: the upshot

Blocking is used when there are observable differences in the experimental units that are believed, a priori, to have an influence on the response variable.

Experimental units are randomly allocated to treatment groups *within blocks*.

The block effect is effectively removed by subtraction so that the treatment effect will show up more clearly.

The Essential Ingredients

Hermit crab study
Control
Replication
Randomization
Summary

Blocking

The non-block and
the block design
Crab example
Summary

Bugaboos, or FAQs

blocking: the upshot

Blocking is used when there are observable differences in the experimental units that are believed, a priori, to have an influence on the response variable.

Experimental units are randomly allocated to treatment groups *within blocks*.

The block effect is effectively removed by subtraction so that the treatment effect will show up more clearly.

(It's actually a little bit more complicated than this, but not much. You have to worry about degrees of freedom because you have estimated τ_b from the data.)

The Essential Ingredients

Hermit crab study
Control
Replication
Randomization
Summary

Blocking

The non-block and
the block design
Crab example
Summary

Bugaboos, or FAQs

a comment

The word “block” is not universally agreed upon by statisticians to mean the same thing. The block model I’ve shown here assumes that blocks do not interact with treatments. That is, the treatment effect is the same for all units regardless of what block they’re in.

The Essential Ingredients

Hermit crab study
Control
Replication
Randomization
Summary

Blocking

The non-block and
the block design
Crab example
Summary

Bugaboos, or FAQs

a comment

The word “block” is not universally agreed upon by statisticians to mean the same thing. The block model I’ve shown here assumes that blocks do not interact with treatments. That is, the treatment effect is the same for all units regardless of what block they’re in.

Under such a model, it would be unwise to perform, say, a drug study and “block on gender”, when men and women may react differently to a drug treatment.

The Essential Ingredients

Hermit crab study
Control
Replication
Randomization
Summary

Blocking

The non-block and
the block design
Crab example
Summary

Bugaboos, or FAQs

The word “block” is not universally agreed upon by statisticians to mean the same thing. The block model I’ve shown here assumes that blocks do not interact with treatments. That is, the treatment effect is the same for all units regardless of what block they’re in.

Under such a model, it would be unwise to perform, say, a drug study and “block on gender”, when men and women may react differently to a drug treatment. (Indeed, Linda Young has said before that you should *never* “block on gender” .)

The Essential
Ingredients

Hermit crab study

Control

Replication

Randomization

Summary

Blocking

The non-block and
the block design

Crab example

Summary

Bugaboos, or
FAQs

The word “block” is not universally agreed upon by statisticians to mean the same thing. The block model I’ve shown here assumes that blocks do not interact with treatments. That is, the treatment effect is the same for all units regardless of what block they’re in.

Under such a model, it would be unwise to perform, say, a drug study and “block on gender”, when men and women may react differently to a drug treatment. (Indeed, Linda Young has said before that you should *never* “block on gender”.)

But other statisticians will use the word “block” even though their model allows for treatment-block interactions. The Linda Young camp would say, “those aren’t *blocks*. They’re additional *factors*.”

The Essential Ingredients

- Hermit crab study
- Control
- Replication
- Randomization
- Summary

Blocking

- The non-block and the block design
- Crab example
- Summary

Bugaboos, or FAQs

a comment (continued)

The difference isn't very important for AP statistics students, because they don't need, at this level, to analyze a blocked design. They do need to know how to construct a blocked design, when it might be appropriate, and why it is an improvement over a non-blocked design.

The Essential
Ingredients

Hermit crab study

Control

Replication

Randomization

Summary

Blocking

The non-block and
the block design

Crab example

Summary

Bugaboos, or
FAQs

A matched-pairs design is a special case of a blocked design in which the blocks are all of size two. The removal of the block effects occurs when you subtract one response in a pair from the other, leaving only the treatment effect.

The Essential Ingredients

Hermit crab study

Control

Replication

Randomization

Summary

Blocking

The non-block and
the block design

Crab example

Summary

Bugaboos, or FAQs

In (more or less) Dick Schaeffer's words:

- ▶ We *control* what variability we can;
- ▶ We *block* to deal with the variability we can observe but not control;
- ▶ We *randomize* to deal with the variability we can't observe.

The Essential
Ingredients

Hermit crab study

Control

Replication

Randomization

Summary

Blocking

The non-block and
the block design

Crab example

Summary

Bugaboos, or
FAQs

random sampling versus random allocation of treatments

- ▶ Random sampling allows you to generalize to a larger population.
- ▶ Random allocation of treatments allows you to draw a cause-and-effect conclusion.

The Essential Ingredients

Hermit crab study

Control

Replication

Randomization

Summary

Blocking

The non-block and
the block design

Crab example

Summary

Bugaboos, or FAQs

surveys, observational studies, experiments

- ▶ Surveys use random sampling but do not involve any treatments.
- ▶ Experiments use random allocation of treatments but do not involve random sampling.
- ▶ Observational studies use neither random allocation of treatments nor random sampling. (But they're not useless!)

The Essential Ingredients

Hermit crab study
Control
Replication
Randomization
Summary

Blocking

The non-block and
the block design
Crab example
Summary

Bugaboos, or FAQs

independence of experimental units

Shrimp in a tank are one unit, not many.

Flowers in a vase are one unit, not many.

The Essential Ingredients

Hermit crab study

Control

Replication

Randomization

Summary

Blocking

The non-block and
the block design

Crab example

Summary

Bugaboos, or FAQs

independence of experimental units

Shrimp in a tank are one unit, not many.

Flowers in a vase are one unit, not many.

The issue is independence. Experimental units *must* have responses that are independent of one another.

The Essential Ingredients

Hermit crab study

Control

Replication

Randomization

Summary

Blocking

The non-block and
the block design

Crab example

Summary

Bugaboos, or FAQs

lurking variables

A quote from Paul Velleman: A lurking variable doesn't just affect the apparent relationship between two variables. An interaction term might do that but wouldn't be lurking. A lurking variable directly affects both X and Y and thereby makes it appear that X and Y are directly related to each other when, without the lurking variable, they would not be or would not be to that extent or in that direction.

My favorite is the strong positive association between the number of firefighters at a fire and the amount of damage. Perhaps you shouldn't call the fire department. The lurking variable is the size of the blaze, which "causes" both damage and fire fighters.

The Essential Ingredients

- Hermit crab study
- Control
- Replication
- Randomization
- Summary

Blocking

- The non-block and the block design
- Crab example
- Summary

Bugaboos, or FAQs

Also from Paul Velleman: *Confounded variables vary together so that one cannot tease apart which is responsible for any observed effect. But only predictors (or factors in an experiment) are said to be confounded. An external variable that is correlated with our response, but not associated with our factors is not a confounder because we will still be able to observe the effect of the factors on the response. Confounding can occur due to poor design in an experiment (offer both a low interest rate and low fee to one group of customers and a higher interest rate and higher fee to another; you'll never be able to tell whether customers were more motivated by the difference in interest rate or the difference in fee.)*

The Essential Ingredients

- Hermit crab study
- Control
- Replication
- Randomization
- Summary

Blocking

- The non-block and the block design
- Crab example
- Summary

Bugaboos, or FAQs

There are two kinds of blinding:

- ▶ If human subjects are involved in a study, they should be blinded (if possible) to which treatment group they are in.
- ▶ If the measurement of the response variable is in any way subjective, then the person doing the measurement should not know what treatment groups the experimental units belong to.

The Essential
Ingredients

Hermit crab study

Control

Replication

Randomization

Summary

Blocking

The non-block and
the block design

Crab example

Summary

Bugaboos, or
FAQs

Other bugaboos?

I have tried to think of common “bugaboos”—difficulties that students and teachers have with experimental design. The ones I’ve presented here show up often on the apstat listserv.

Are there others you’d like to talk about?

The Essential Ingredients

Hermit crab study

Control

Replication

Randomization

Summary

Blocking

The non-block and
the block design

Crab example

Summary

Bugaboos, or FAQs

Thank you so very much for inviting me to speak with you today.

I am always happy to correspond with other teachers about statistical or teaching issues.

bullard@ncssm.edu



The Essential
Ingredients

Hermit crab study
Control
Replication
Randomization
Summary

Blocking

The non-block and
the block design
Crab example
Summary

Bugaboos, or
FAQs