Design and Analysis of Scientific Experiments

R. A. Bailey
University of St Andrews

SISCO
6 February 2022

Consultation or collaboration with scientist

If this is a new consultation, does it happen in good time, or the day before the experiment starts?

Example

Scientist telephones RAB, gives her lots of complicated details about his proposed experiment, and asks for her advice.

Consultation or collaboration with scientist

If this is a new consultation, does it happen in good time, or the day before the experiment starts?

Example

Scientist telephones RAB, gives her lots of complicated details about his proposed experiment, and asks for her advice.

RAB: I am afraid that that is rather complicated. I will need to think about it. When are you planning to start this?

Consultation or collaboration with scientist

If this is a new consultation, does it happen in good time, or the day before the experiment starts?

Example

Scientist telephones RAB, gives her lots of complicated details about his proposed experiment, and asks for her advice.

RAB: I am afraid that that is rather complicated. I will need to think about it. When are you planning to start this?

Scientist: Tomorrow.

Consultation or collaboration with scientist

If this is a new consultation, does it happen in good time, or the day before the experiment starts?

Example

Scientist telephones RAB, gives her lots of complicated details about his proposed experiment, and asks for her advice.

RAB: I am afraid that that is rather complicated. I will need to think about it. When are you planning to start this?

Scientist: Tomorrow.

RAB: I am afraid that I cannot do it by then.

Consultation or collaboration with scientist

If this is a new consultation, does it happen in good time, or the day before the experiment starts?

Example

Scientist telephones RAB, gives her lots of complicated details about his proposed experiment, and asks for her advice.

RAB: I am afraid that that is rather complicated. I will need to think about it. When are you planning to start this?

Scientist: Tomorrow.

RAB: I am afraid that I cannot do it by then.

Scientist: I do not need you to do anything; I just need to tick the box that says that I have consulted a statistician.

Different vocabulary and notation

Example

The horticulturalist describes five different species, all with Latin names.

Different vocabulary and notation

Example

The horticulturalist describes five different species, all with Latin names.

RAB calls them A, B, C, D and E.

Different vocabulary and notation

Example

The horticulturalist describes five different species,
all with Latin names.

RAB calls them A, B, C, D and E.

Of course, the real Latin names are used in the details of the experiment,
but RAB uses A, B, C, D and E in her thinking and planning.
So she has to make a dictionary for herself, linking the letters to the real names.

Admitting ignorance

Scientists may not like to admit that they do not know the meaning of some word or notation.

Admitting ignorance

Scientists may not like to admit that they do not know the meaning of some word or notation.

The statistician absolutely must not be afraid of admitting ignorance and asking for clarification.

Admitting ignorance

Scientists may not like to admit that they do not know the meaning of some word or notation.

The statistician absolutely must not be afraid of admitting ignorance and asking for clarification.

Example

A small group of scientists come to consult RAB. They describe their proposed experiment, and give her typed notes about it, including some equations for their assumed model.

Admitting ignorance

Scientists may not like to admit that they do not know the meaning of some word or notation.

The statistician absolutely must not be afraid of admitting ignorance and asking for clarification.

Example

A small group of scientists come to consult RAB. They describe their proposed experiment, and give her typed notes about it, including some equations for their assumed model.

RAB: What do you mean by σ^2?

Admitting ignorance

Scientists may not like to admit that they do not know the meaning of some word or notation.

The statistician absolutely must not be afraid of admitting ignorance and asking for clarification.

Example

A small group of scientists come to consult RAB. They describe their proposed experiment, and give her typed notes about it, including some equations for their assumed model.

RAB: What do you mean by σ^2?

Scientists: Everyone knows that σ^2 means variance,
so we do not need to explain that.

Admitting ignorance

Scientists may not like to admit that they do not know the meaning of some word or notation.

The statistician absolutely must not be afraid of admitting ignorance and asking for clarification.

Example

A small group of scientists come to consult RAB. They describe their proposed experiment, and give her typed notes about it, including some equations for their assumed model.

RAB: What do you mean by σ^2?

Scientists: Everyone knows that σ^2 means variance,
so we do not need to explain that.

RAB: But I think that your set-up will have two different variances, so we need different notation for each,
and we have to explain which is which.

Politeness or honesty?

One day I was visiting some mathematicians in the London School of Economics. We had lunch in the staff canteen, where we were joined by a few biologists. During lunch, we had the following conversation.

Politeness or honesty?

One day I was visiting some mathematicians in the London School of Economics. We had lunch in the staff canteen, where we were joined by a few biologists. During lunch, we had the following conversation.

One mathematician: I think that I have just proved the following interesting theorem.

Politeness or honesty?

One day I was visiting some mathematicians in the London School of Economics. We had lunch in the staff canteen, where we were joined by a few biologists. During lunch, we had the following conversation.

One mathematician: I think that I have just proved the following interesting theorem.

Three other mathematicians: No, you must be wrong because …

Politeness or honesty?

One day I was visiting some mathematicians in the London School of Economics. We had lunch in the staff canteen, where we were joined by a few biologists. During lunch, we had the following conversation.

One mathematician: I think that I have just proved the following interesting theorem.

Three other mathematicians: No, you must be wrong because …

All the biologists: How can you be so rude to your colleague?

Politeness or honesty?

One day I was visiting some mathematicians in the London School of Economics. We had lunch in the staff canteen, where we were joined by a few biologists. During lunch, we had the following conversation.

One mathematician: I think that I have just proved the following interesting theorem.

Three other mathematicians: No, you must be wrong because …

All the biologists: How can you be so rude to your colleague?

All the mathematicians, including the one who had made the claim: If I make a mistake, I prefer it if other people tell me as soon as possible, so that I do not waste any more time using a false assumption.

Who is involved?

If any part of the experimental procedure
(such as sowing seeds, injecting patients, collecting data)
involves more than one person,

Who is involved?

If any part of the experimental procedure
(such as sowing seeds, injecting patients, collecting data)
involves more than one person,
then the allocation should be part of the design, should be recorded, and should be used in the data analysis.

Who is involved?

If any part of the experimental procedure
(such as sowing seeds, injecting patients, collecting data)
involves more than one person,
then the allocation should be part of the design, should be recorded, and should be used in the data analysis.

Example

In an experiment on growing soft fruit, seasonal fruit-pickers were employed to harvest the ripe fruit. Each fruit-picker was allocated to a certain number of rows of fruit each day, so that any differences between fruit-pickers could be allowed for in the data analysis.

Who is involved?

If any part of the experimental procedure
(such as sowing seeds, injecting patients, collecting data)
involves more than one person,
then the allocation should be part of the design, should be recorded, and should be used in the data analysis.

Example

In an experiment on growing soft fruit, seasonal fruit-pickers were employed to harvest the ripe fruit. Each fruit-picker was allocated to a certain number of rows of fruit each day, so that any differences between fruit-pickers could be allowed for in the data analysis. One day, someone noticed that all the fruit-pickers were working in the same row.

Who is involved?

If any part of the experimental procedure
(such as sowing seeds, injecting patients, collecting data)
involves more than one person,
then the allocation should be part of the design, should be recorded, and should be used in the data analysis.

Example

In an experiment on growing soft fruit, seasonal fruit-pickers were employed to harvest the ripe fruit. Each fruit-picker was allocated to a certain number of rows of fruit each day, so that any differences between fruit-pickers could be allowed for in the data analysis. One day, someone noticed that all the fruit-pickers were working in the same row.
When queried, they said

> *Of course! We always do this. We do this job so that we have other people to talk to while we are doing it.*

Not just biology

A former undergraduate student of mine at St Andrews decided to do a final-year project investigating how undergraduates in Physics conducted their experiments.

Not just biology

A former undergraduate student of mine at St Andrews decided to do a final-year project investigating how undergraduates in Physics conducted their experiments.

She found that they were allocated to certain times in the lab, and certain pieces of equipment.

Not just biology

A former undergraduate student of mine at St Andrews decided to do a final-year project investigating how undergraduates in Physics conducted their experiments.

She found that they were allocated to certain times in the lab, and certain pieces of equipment.

But the students preferred to work with their friends, so they arranged to swap times and/or pieces of equipment.

Not just biology

A former undergraduate student of mine at St Andrews decided to do a final-year project investigating how undergraduates in Physics conducted their experiments.

She found that they were allocated to certain times in the lab, and certain pieces of equipment.

But the students preferred to work with their friends, so they arranged to swap times and/or pieces of equipment.

They did not want to get into trouble for this, so they reported their results as if they had done their experiment at the allocated time and on the allocated piece of equipment.

Data collection

I always advise the following.

- Specify units of measurement in advance.

Data collection

I always advise the following.

- Specify units of measurement in advance.
- Measure items in natural spatial or temporal order.

Data collection

I always advise the following.

- Specify units of measurement in advance.
- Measure items in natural spatial or temporal order.
- No copying of data by hand.

Data collection

I always advise the following.

- Specify units of measurement in advance.
- Measure items in natural spatial or temporal order.
- No copying of data by hand.
- No re-ordering of the data.

Data collection

I always advise the following.

- Specify units of measurement in advance.
- Measure items in natural spatial or temporal order.
- No copying of data by hand.
- No re-ordering of the data.
- No intermediate calculations.

Data collection

I always advise the following.

- Specify units of measurement in advance.
- Measure items in natural spatial or temporal order.
- No copying of data by hand.
- No re-ordering of the data.
- No intermediate calculations.
- No delegation to juniors.

Data sniffing

Look over data for obvious anomalies or outliers or bad practice (for example, change of measurement units).

Query dubious data while there is still time to investigate.

Data sniffing

Look over data for obvious anomalies or outliers or bad practice (for example, change of measurement units).
Query dubious data while there is still time to investigate.

Example

RAB: Why is one measurement more than 15 times as big as any of the others?

Data sniffing

Look over data for obvious anomalies or outliers or bad practice (for example, change of measurement units).
Query dubious data while there is still time to investigate.

Example

RAB: Why is one measurement more than 15 times as big as any of the others?

Scientists: Oh, we had not noticed that, we will check.

…

Data sniffing

Look over data for obvious anomalies or outliers or bad practice (for example, change of measurement units).
Query dubious data while there is still time to investigate.

Example

RAB: Why is one measurement more than 15 times as big as any of the others?

Scientists: Oh, we had not noticed that, we will check.

...

Scientists (a bit later): Sorry, we had made a copying mistake, we have corrected it now.

More examples of data scrutiny

Example

In an experiment at an agricultural research station in New Zealand, the hardness of kiwi fruit was measured. Preliminary data analysis made the statistician suspicious of the results. Then he noticed that the data had been recorded in two different handwritings. He re-analysed the data, including an unknown constant to multiply all the data in the second handwriting. The fitted value of the constant was 2.2.
What does this suggest?

More examples of data scrutiny

Example

In an experiment at an agricultural research station in New Zealand, the hardness of kiwi fruit was measured. Preliminary data analysis made the statistician suspicious of the results. Then he noticed that the data had been recorded in two different handwritings. He re-analysed the data, including an unknown constant to multiply all the data in the second handwriting. The fitted value of the constant was 2.2.
What does this suggest?

Example

In an experiment on wheat yields, I noticed that the numbers recorded for the last 12 plots out of 72 were noticably lower than the rest. I asked why. “It started to rain during harvest, when the harvester was about 12 plots away from the end.”
I was able to include a covariate in the data analysis to allow for this.

Data analysis

In many cases, data analysis begins with analysis of variance (ANOVA).

Data analysis

In many cases, data analysis begins with analysis of variance (ANOVA).

This produces a table,
with rows labelled by sources of variation, such as treatments, labs, error (everything left over, also called residual)
and columns headed by things like sum of squares, degrees of freedom, mean square, variance ratio (also called F).

Data analysis

In many cases, data analysis begins with analysis of variance (ANOVA).

This produces a table,

with rows labelled by sources of variation, such as treatments, labs, error (everything left over, also called residual)

and columns headed by things like sum of squares, degrees of freedom, mean square, variance ratio (also called F).

If MS(treatments) $>>$ MS(residual) then there are almost certainly differences between treatments.

Data analysis

In many cases, data analysis begins with analysis of variance (ANOVA).

This produces a table,

with rows labelled by sources of variation, such as treatments, labs, error (everything left over, also called residual)

and columns headed by things like sum of squares, degrees of freedom, mean square, variance ratio (also called F).

If MS(treatments) $>>$ MS(residual) then there are almost certainly differences between treatments.

If MS(treatments) $\approx$ MS(residual) then we do not have enough evidence that there are differences between treatments. (There may be, but these data do not give us evidence to claim that.)

Data analysis

In many cases, data analysis begins with analysis of variance (ANOVA).
This produces a table,
with rows labelled by sources of variation, such as treatments, labs, error (everything left over, also called residual)
and columns headed by things like sum of squares, degrees of freedom, mean square, variance ratio (also called F).
If MS(treatments) $>>$ MS(residual) then there are almost certainly differences between treatments.
If MS(treatments) $\approx$ MS(residual) then we do not have enough evidence that there are differences between treatments. (There may be, but these data do not give us evidence to claim that.)
If MS(treatments) $<<$ MS(residual) then you should query the data.

Data analysis

In many cases, data analysis begins with analysis of variance (ANOVA).
This produces a table,
with rows labelled by sources of variation, such as treatments, labs, error (everything left over, also called residual)
and columns headed by things like sum of squares, degrees of freedom, mean square, variance ratio (also called F).
If MS(treatments) $>>$ MS(residual) then there are almost certainly differences between treatments.
If MS(treatments) $\approx$ MS(residual) then we do not have enough evidence that there are differences between treatments. (There may be, but these data do not give us evidence to claim that.)
If MS(treatments) $<<$ MS(residual) then you should query the data.

Once when I did query this, I got the reply “Oh, we used a block design, because we know that we are supposed to, but we did not want to burden your brain with too many details.”

Rothamsted Experimental Station (Harpenden)

This was founded by Sir John Bennet Lawes in 1843.

Broadbalk

Rothamsted Experimental Station (Harpenden)

This was founded by Sir John Bennet Lawes in 1843.

trees →

Broadbalk

Rothamsted Experimental Station (Harpenden)

This was founded by Sir John Bennet Lawes in 1843.

trees →

Broadbalk

I worked in the Statistics Department there from 1981 to 1990.

Blocking

We have 6 varieties to compare in this field.
How do we avoid bias?

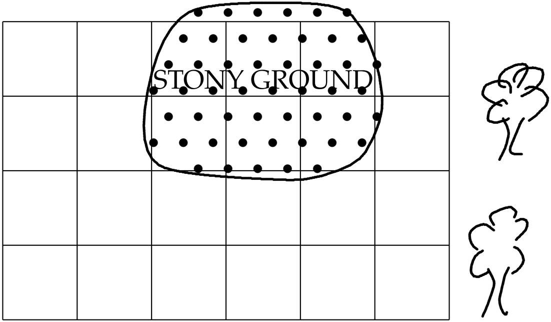

Blocking

We have 6 varieties to compare in this field.
How do we avoid bias?

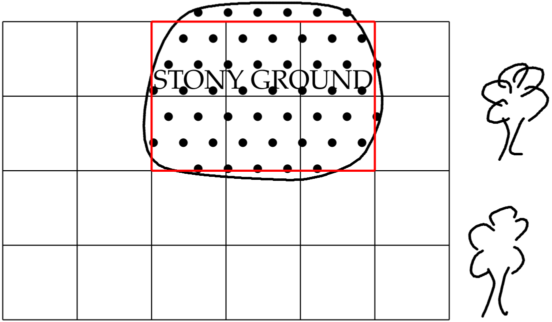

Blocking

We have 6 varieties to compare in this field.
How do we avoid bias?

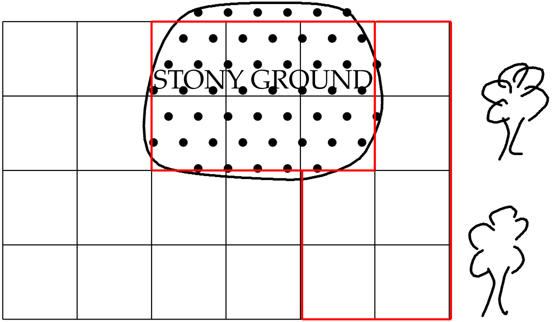

Blocking

We have 6 varieties to compare in this field.
How do we avoid bias?

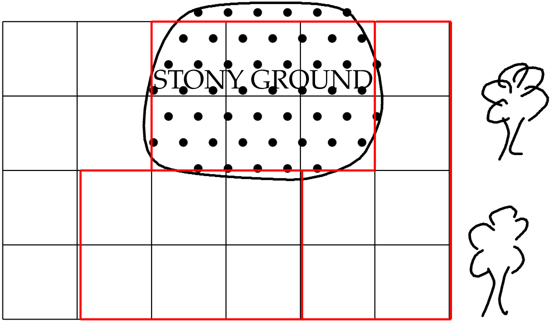

Blocking

We have 6 varieties to compare in this field.
How do we avoid bias?

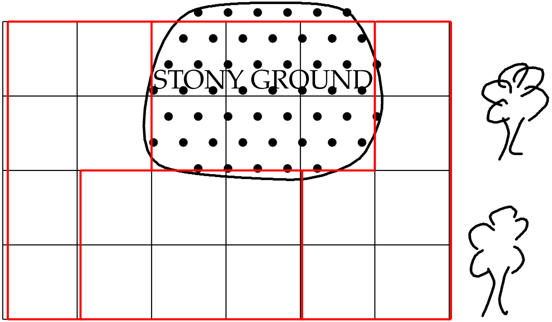

Blocking

We have 6 varieties to compare in this field.
How do we avoid bias?

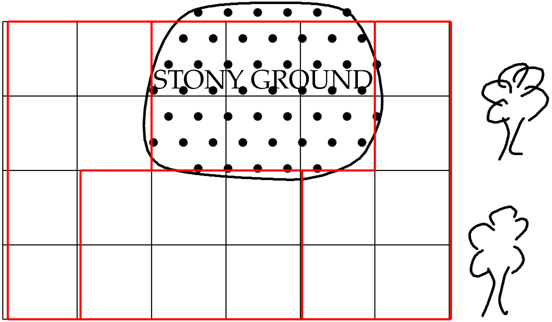

Partition the experimental units into homogeneous blocks and plant each variety in one plot in each block.

Latin squares

Suppose that there are n^2 experimental units
and that they can be considered as a $n \times n$ square array.

Latin squares

Suppose that there are n^2 experimental units
and that they can be considered as a $n \times n$ square array.

One possibility is rows and columns in a field.

Latin squares

Suppose that there are n^2 experimental units
and that they can be considered as a $n \times n$ square array.

One possibility is rows and columns in a field.

Another possibility is n patients in a trial for n months,
each able to change their drug each month.

Latin squares

Suppose that there are n^2 experimental units
and that they can be considered as a $n \times n$ square array.

One possibility is rows and columns in a field.

Another possibility is n patients in a trial for n months,
each able to change their drug each month.

The design is called a Latin square if every treatment occurs
exactly once in each row and once in each column.

An experiment on potatoes at Ely in 1932

E	B	F	A	C	D
B	C	D	E	F	A
A	E	C	B	D	F
F	D	E	C	A	B
D	A	B	F	E	C
C	F	A	D	B	E

Treatment	A	B	C	D	E	F
Extra nitrogen	0	0	0	1	1	1
Extra phosphate	0	1	2	0	1	2

Column-complete Latin squares

Definition

A Latin square is column-complete if each treatment is immediately followed, in the same column, by each other treatment exactly once.

0	1	2	3	4	5
1	2	3	4	5	0
5	0	1	2	3	4
2	3	4	5	0	1
4	5	0	1	2	3
3	4	5	0	1	2

Column-complete Latin squares

Definition

A Latin square is column-complete if each treatment is immediately followed, in the same column, by each other treatment exactly once.

0	1	2	3	4	5
1	2	3	4	5	0
5	0	1	2	3	4
2	3	4	5	0	1
4	5	0	1	2	3
3	4	5	0	1	2

Column-complete Latin squares

Definition

A Latin square is column-complete if each treatment is immediately followed, in the same column, by each other treatment exactly once.

0	1	2	3	4	5
1	2	3	4	5	0
5	0	1	2	3	4
2	3	4	5	0	1
4	5	0	1	2	3
3	4	5	0	1	2

These squares are widely used in tasting experiments and in trials of new drugs to alleviate symptoms of chronic conditions. (Rows represent time-periods; columns represent people.)

Complete Latin squares

A Latin square is row-complete if each treatment is immediately followed, in the same row, by each other treatment exactly once.

Complete Latin squares

A Latin square is row-complete if each treatment is immediately followed, in the same row, by each other treatment exactly once.

A Latin square is complete if it is both row-complete and column-complete.

An experiment at Rothamsted that I designed

An experiment at Rothamsted that I designed

This is a complete Latin square with six treatments.

Unintended consequences

The following year at Rothamsted,
I was flooded with requests from scientists who wanted to do an experiment in a complete Latin square of size 6×6.

Unintended consequences

The following year at Rothamsted,
I was flooded with requests from scientists who wanted to do an experiment in a complete Latin square of size 6×6.

RAB: There is nothing special about 6.
I can give you a Latin square like that for any even number.

Unintended consequences

The following year at Rothamsted,
I was flooded with requests from scientists who wanted to do an experiment in a complete Latin square of size 6×6.

RAB: There is nothing special about 6.
I can give you a Latin square like that for any even number.
For some experiments on the ground,
an East neighbour is as bad as a West neighbour,
and a South neighbour is as bad as a North neighbour.

Unintended consequences

The following year at Rothamsted,
I was flooded with requests from scientists who wanted to do an experiment in a complete Latin square of size 6×6.

RAB: There is nothing special about 6.
I can give you a Latin square like that for any even number.
For some experiments on the ground,
an East neighbour is as bad as a West neighbour,
and a South neighbour is as bad as a North neighbour.
In that case, I can give you an appropriate Latin square of any size.

There are more possible designs than you can imagine

The scientist's pre-conceived design from *Statistics for Science 101* may not be ideal.

There are more possible designs than you can imagine

The scientist's pre-conceived design from *Statistics for Science 101* may not be ideal.

They are often taught only three possibilities.

- A Latin square.

There are more possible designs than you can imagine

The scientist's pre-conceived design from *Statistics for Science 101* may not be ideal.

They are often taught only three possibilities.

- A Latin square.
- A block design,
  with each treatment occurring once in each block.

There are more possible designs than you can imagine

The scientist's pre-conceived design from *Statistics for Science 101* may not be ideal.

They are often taught only three possibilities.

- A Latin square.
- A block design,
  with each treatment occurring once in each block.
- A completely randomized design,
  which has no restrictions on which treatment can go where.

An experiment on pesticides

A famous company conducted an experiment to compare a new pesticide which they had developed, a standard pesticide, and "no treatment". They wanted to convince the Ministry of Agriculture, Fisheries and Foods that their new pesticide was effective but did not harm ladybirds.

An experiment on pesticides

A famous company conducted an experiment to compare a new pesticide which they had developed, a standard pesticide, and "no treatment". They wanted to convince the Ministry of Agriculture, Fisheries and Foods that their new pesticide was effective but did not harm ladybirds.

MAFF asked me to investigate the data from the experiment. I saw that the company had divided a field into three areas, applied their new pesticide to one area, used the standard pesticide on another area, and put nothing on the third area. Later they had measured the number of ladybirds on three samples from each area.

An experiment on pesticides

A famous company conducted an experiment to compare a new pesticide which they had developed, a standard pesticide, and "no treatment". They wanted to convince the Ministry of Agriculture, Fisheries and Foods that their new pesticide was effective but did not harm ladybirds.

MAFF asked me to investigate the data from the experiment. I saw that the company had divided a field into three areas, applied their new pesticide to one area, used the standard pesticide on another area, and put nothing on the third area. Later they had measured the number of ladybirds on three samples from each area.

Thus there is no way of distinguishing the differences between treatments from the differences between areas.

An experiment on pesticides

A famous company conducted an experiment to compare a new pesticide which they had developed, a standard pesticide, and "no treatment". They wanted to convince the Ministry of Agriculture, Fisheries and Foods that their new pesticide was effective but did not harm ladybirds.

MAFF asked me to investigate the data from the experiment. I saw that the company had divided a field into three areas, applied their new pesticide to one area, used the standard pesticide on another area, and put nothing on the third area. Later they had measured the number of ladybirds on three samples from each area.

Thus there is no way of distinguishing the differences between treatments from the differences between areas.

This is a common mistake, known as false replication.

How did they justify their design?

RAB: What design did you use?

How did they justify their design?

RAB: What design did you use?

Company: It was a completely randomized design.

How did they justify their design?

RAB: What design did you use?

Company: It was a completely randomized design.

RAB: I can see that it was not completely randomized, because all the samples for each treatment come from the same part of the field.

How did they justify their design?

RAB: What design did you use?

Company: It was a completely randomized design.

RAB: I can see that it was not completely randomized, because all the samples for each treatment come from the same part of the field.

Company: It was not a Latin square.
There were no blocks, so it was not a block design.
Therefore it was completely randomized (quotes a respectable textbook).

Another example of false replication

Germany company Bayer AG makes agricultural chemicals (among other things). They carried out a large experiment in Mecklenburg–West Pomerania. Two large study sites were chosen, each containing many farms. At one site, all farmers treated their oilseed rape seeds with clothianidin before sowing; at the other site, no farmers did this.

Another example of false replication

Germany company Bayer AG makes agricultural chemicals (among other things). They carried out a large experiment in Mecklenburg–West Pomerania. Two large study sites were chosen, each containing many farms. At one site, all farmers treated their oilseed rape seeds with clothianidin before sowing; at the other site, no farmers did this.

Three species of bees were used. For each species, within each site, six locations were chosen, three at the edges of fields growing oilseed rape, three at a fixed minimum distance from such fields. At each location, several bee colonies were placed. Many characteristics of bee performance were measured in each colony.

Another example of false replication

Germany company Bayer AG makes agricultural chemicals (among other things). They carried out a large experiment in Mecklenburg–West Pomerania. Two large study sites were chosen, each containing many farms. At one site, all farmers treated their oilseed rape seeds with clothianidin before sowing; at the other site, no farmers did this.

Three species of bees were used. For each species, within each site, six locations were chosen, three at the edges of fields growing oilseed rape, three at a fixed minimum distance from such fields. At each location, several bee colonies were placed. Many characteristics of bee performance were measured in each colony.

An entire issue of the journal *Ecotoxicology* in 2016 was devoted to conclusions from data from this experiment.

Another example of false replication

Germany company Bayer AG makes agricultural chemicals (among other things). They carried out a large experiment in Mecklenburg–West Pomerania. Two large study sites were chosen, each containing many farms. At one site, all farmers treated their oilseed rape seeds with clothianidin before sowing; at the other site, no farmers did this.

Three species of bees were used. For each species, within each site, six locations were chosen, three at the edges of fields growing oilseed rape, three at a fixed minimum distance from such fields. At each location, several bee colonies were placed. Many characteristics of bee performance were measured in each colony.

An entire issue of the journal *Ecotoxicology* in 2016 was devoted to conclusions from data from this experiment.

Jeremy Greenwood (CREEM, St Andrews) and RAB wrote a Letter to the Editor pointing out that there was no way of distinguishing any effect of clothianidin from any inherent differences between sites.

An experiment on zebrafish

Josh Paik took my module MT4614 *Design of Experiments* in 2020. He is now doing a PhD at Penn State University, USA.

An experiment on zebrafish

Josh Paik took my module MT4614 *Design of Experiments* in 2020. He is now doing a PhD at Penn State University, USA.
A post-doc in neuroscience asked for his help in an experiment.

An experiment on zebrafish

Josh Paik took my module MT4614 *Design of Experiments* in 2020. He is now doing a PhD at Penn State University, USA.

A post-doc in neuroscience asked for his help in an experiment.

Post-doc: I am going use ten zebrafish, each for one day.

An experiment on zebrafish

Josh Paik took my module MT4614 *Design of Experiments* in 2020. He is now doing a PhD at Penn State University, USA.

A post-doc in neuroscience asked for his help in an experiment.

Post-doc: I am going use ten zebrafish, each for one day. Starting at 9 am, I will fire a laser at a specific group of neurons for 3 milliseconds. Then I will turn off the laser, wait 10 seconds, then do it again. Repeat until 5 minutes are up, then record the neuron activity for the next 55 minutes.

An experiment on zebrafish

Josh Paik took my module MT4614 *Design of Experiments* in 2020. He is now doing a PhD at Penn State University, USA.

A post-doc in neuroscience asked for his help in an experiment.

Post-doc: I am going use ten zebrafish, each for one day.

Starting at 9 am, I will fire a laser at a specific group of neurons for 3 milliseconds. Then I will turn off the laser, wait 10 seconds, then do it again. Repeat until 5 minutes are up, then record the neuron activity for the next 55 minutes.

At 10 am, I do the same thing, firing the laser for 6 milliseconds.

An experiment on zebrafish

Josh Paik took my module MT4614 *Design of Experiments* in 2020. He is now doing a PhD at Penn State University, USA.

A post-doc in neuroscience asked for his help in an experiment.

Post-doc: I am going use ten zebrafish, each for one day. Starting at 9 am, I will fire a laser at a specific group of neurons for 3 milliseconds. Then I will turn off the laser, wait 10 seconds, then do it again. Repeat until 5 minutes are up, then record the neuron activity for the next 55 minutes. At 10 am, I do the same thing, firing the laser for 6 milliseconds. At 11 am, I do the same thing, firing the laser for 9 milliseconds.

An experiment on zebrafish

Josh Paik took my module MT4614 *Design of Experiments* in 2020. He is now doing a PhD at Penn State University, USA.

A post-doc in neuroscience asked for his help in an experiment.

Post-doc: I am going use ten zebrafish, each for one day.

Starting at 9 am, I will fire a laser at a specific group of neurons for 3 milliseconds. Then I will turn off the laser, wait 10 seconds, then do it again. Repeat until 5 minutes are up, then record the neuron activity for the next 55 minutes.

At 10 am, I do the same thing, firing the laser for 6 milliseconds.

At 11 am, I do the same thing, firing the laser for 9 milliseconds.

After that, the zebrafish will be euthanized, because it would be cruel to keep it alive.

An experiment on zebrafish

Josh Paik took my module MT4614 *Design of Experiments* in 2020. He is now doing a PhD at Penn State University, USA.

A post-doc in neuroscience asked for his help in an experiment.

Post-doc: I am going use ten zebrafish, each for one day. Starting at 9 am, I will fire a laser at a specific group of neurons for 3 milliseconds. Then I will turn off the laser, wait 10 seconds, then do it again. Repeat until 5 minutes are up, then record the neuron activity for the next 55 minutes.

At 10 am, I do the same thing, firing the laser for 6 milliseconds.

At 11 am, I do the same thing, firing the laser for 9 milliseconds.

After that, the zebrafish will be euthanized, because it would be cruel to keep it alive.

Josh: Can you do 3, 6 and 9 in different orders for some zebrafish? Different times of day may have different effects.

An experiment on zebrafish

Josh Paik took my module MT4614 *Design of Experiments* in 2020. He is now doing a PhD at Penn State University, USA.

A post-doc in neuroscience asked for his help in an experiment.

Post-doc: I am going use ten zebrafish, each for one day. Starting at 9 am, I will fire a laser at a specific group of neurons for 3 milliseconds. Then I will turn off the laser, wait 10 seconds, then do it again. Repeat until 5 minutes are up, then record the neuron activity for the next 55 minutes. At 10 am, I do the same thing, firing the laser for 6 milliseconds. At 11 am, I do the same thing, firing the laser for 9 milliseconds. After that, the zebrafish will be euthanized, because it would be cruel to keep it alive.

Josh: Can you do 3, 6 and 9 in different orders for some zebrafish? Different times of day may have different effects.

Post-doc: I suppose so.

An experiment on zebrafish

Josh Paik took my module MT4614 *Design of Experiments* in 2020. He is now doing a PhD at Penn State University, USA.

A post-doc in neuroscience asked for his help in an experiment.

Post-doc: I am going use ten zebrafish, each for one day. Starting at 9 am, I will fire a laser at a specific group of neurons for 3 milliseconds. Then I will turn off the laser, wait 10 seconds, then do it again. Repeat until 5 minutes are up, then record the neuron activity for the next 55 minutes. At 10 am, I do the same thing, firing the laser for 6 milliseconds. At 11 am, I do the same thing, firing the laser for 9 milliseconds. After that, the zebrafish will be euthanized, because it would be cruel to keep it alive.

Josh: Can you do 3, 6 and 9 in different orders for some zebrafish? Different times of day may have different effects.

Post-doc: I suppose so.

Josh: Why 10 fish? Can you do 12?

An experiment on zebrafish

Josh Paik took my module MT4614 *Design of Experiments* in 2020. He is now doing a PhD at Penn State University, USA.

A post-doc in neuroscience asked for his help in an experiment.

Post-doc: I am going use ten zebrafish, each for one day. Starting at 9 am, I will fire a laser at a specific group of neurons for 3 milliseconds. Then I will turn off the laser, wait 10 seconds, then do it again. Repeat until 5 minutes are up, then record the neuron activity for the next 55 minutes. At 10 am, I do the same thing, firing the laser for 6 milliseconds. At 11 am, I do the same thing, firing the laser for 9 milliseconds. After that, the zebrafish will be euthanized, because it would be cruel to keep it alive.

Josh: Can you do 3, 6 and 9 in different orders for some zebrafish? Different times of day may have different effects.

Post-doc: I suppose so.

Josh: Why 10 fish? Can you do 12?

Post-doc: Yes, but it will take me at least two extra days.

Design for the zebrafish experiment

day	time		
	9 am	10 am	11 am
1	3	6	9
2	6	9	3
3	3	6	9
4	9	3	6
5	6	9	3
6	3	9	6
7	9	6	3
8	6	3	9
9	9	3	6
10	9	6	3
11	6	3	9
12	3	9	6

More complicated structures

There are many experiments where we need more complicated structures than those I have shown so far.

More complicated structures

There are many experiments where we need more complicated structures than those I have shown so far.

The blocks may have a natural size which is smaller than the number of treatments.

More complicated structures

There are many experiments where we need more complicated structures than those I have shown so far.

The blocks may have a natural size which is smaller than the number of treatments.

The treatments may consist of all combinations of levels of two or more treatment factors.
(The experiment on potatoes in Ely was a simple example.)

More complicated structures

There are many experiments where we need more complicated structures than those I have shown so far.

The blocks may have a natural size which is smaller than the number of treatments.

The treatments may consist of all combinations of levels of two or more treatment factors.
(The experiment on potatoes in Ely was a simple example.)

The experiment may have two or more phases.

Incomplete-block designs

Suppose that there are b blocks of size k, and v treatments, where v divides bk and $k < v$.

Suppose that there are b blocks of size k, and v treatments, where v divides bk and $k < v$.

A design is optimal if it minimizes the sum of the variances V_{ij} of the estimators of the differences between all pairs of treatments.

Incomplete-block designs

Suppose that there are b blocks of size k, and v treatments, where v divides bk and $k < v$.

A design is optimal if it minimizes the sum of the variances V_{ij} of the estimators of the differences between all pairs of treatments.

The concurrence of two treatments is the number of blocks in which they both occur.

Incomplete-block designs

Suppose that there are b blocks of size k, and v treatments, where v divides bk and $k < v$.

A design is optimal if it minimizes the sum of the variances V_{ij} of the estimators of the differences between all pairs of treatments.

The concurrence of two treatments is the number of blocks in which they both occur.

A design is balanced if all concurrences are equal.

Incomplete-block designs

Suppose that there are b blocks of size k, and v treatments, where v divides bk and $k < v$.

A design is optimal if it minimizes the sum of the variances V_{ij} of the estimators of the differences between all pairs of treatments.

The concurrence of two treatments is the number of blocks in which they both occur.

A design is balanced if all concurrences are equal.

Theorem

If a balanced incomplete-block design (BIBD) exists for these values of b, k and v then it is optimal.

Suppose that there are b blocks of size k, and v treatments, where v divides bk and $k < v$.

A design is optimal if it minimizes the sum of the variances V_{ij} of the estimators of the differences between all pairs of treatments.

The concurrence of two treatments is the number of blocks in which they both occur.

A design is balanced if all concurrences are equal.

Theorem

If a balanced incomplete-block design (BIBD) exists for these values of b, k and v then it is optimal.

Combinatorialists know many ways of constructing balanced incomplete-block designs.

What should we do if there are no BIBDs for these parameters?

Suggestion: Use a BIBD with a smaller value of b?

What should we do if there are no BIBDs for these parameters?

Suggestion: Use a BIBD with a smaller value of b?
Response: No, because each extra block will decrease some variances.

What should we do if there are no BIBDs for these parameters?

Suggestion: Use a BIBD with a smaller value of b?
Response: No, because each extra block will decrease some variances.

Suggestion: Use combinatorial techniques such as graph theory, symmetry and association schemes to find families of block designs that are optimal.

What should we do if there are no BIBDs for these parameters?

Suggestion: Use a BIBD with a smaller value of b?
Response: No, because each extra block will decrease some variances.

Suggestion: Use combinatorial techniques such as graph theory, symmetry and association schemes to find families of block designs that are optimal.
Response: Yes, this has been very successful.

What should we do if there are no BIBDs for these parameters?

Suggestion: Use a BIBD with a smaller value of b?
Response: No, because each extra block will decrease some variances.

Suggestion: Use combinatorial techniques such as graph theory, symmetry and association schemes to find families of block designs that are optimal.
Response: Yes, this has been very successful.

Suggestion: Simply use computer search. We can simplify this by assuming that all treatments occur equally often.

What should we do if there are no BIBDs for these parameters?

Suggestion: Use a BIBD with a smaller value of b?
Response: No, because each extra block will decrease some variances.

Suggestion: Use combinatorial techniques such as graph theory, symmetry and association schemes to find families of block designs that are optimal.
Response: Yes, this has been very successful.

Suggestion: Simply use computer search. We can simplify this by assuming that all treatments occur equally often.
Response: Equal replication used to be part of the folklore. We now know that optimal designs do not always have this property.

Designs for $k = 2$ when $b = v$ (blocks shown as edges)

The only connected equireplicate design is the cycle.

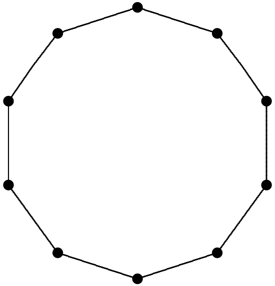

Designs for $k = 2$ when $b = v$ (blocks shown as edges)

The only connected equireplicate design is the cycle.

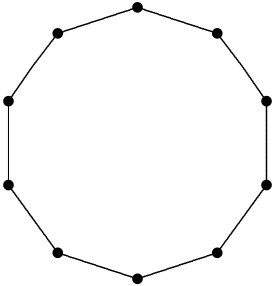

If the distance between i and j is w

$$V_{ij} = \frac{2w(v-w)}{v}\sigma^2,$$

Designs for $k = 2$ when $b = v$ (blocks shown as edges)

The only connected equireplicate design is the cycle.

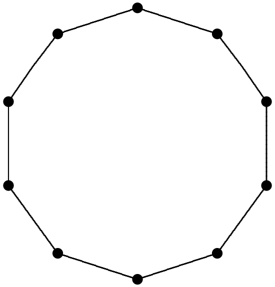

If the distance between i and j is w

$$V_{ij} = \frac{2w(v-w)}{v}\sigma^2,$$

$> 4\sigma^2$ if $v \geq 10$ and $3 \leq w$.

Designs for $k = 2$ when $b = v$ (blocks shown as edges)

The only connected equireplicate design is the cycle.

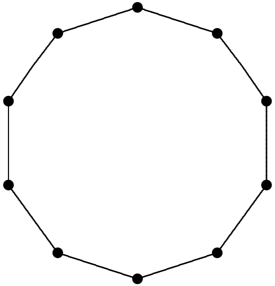

If the distance between i and j is w

$$V_{ij} = \frac{2w(v-w)}{v}\sigma^2,$$

$> 4\sigma^2$ if $v \geq 10$ and $3 \leq w$.

Here is an alternative design.

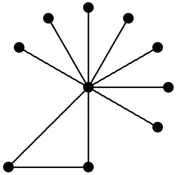

Designs for $k = 2$ when $b = v$ (blocks shown as edges)

The only connected equireplicate design is the cycle.

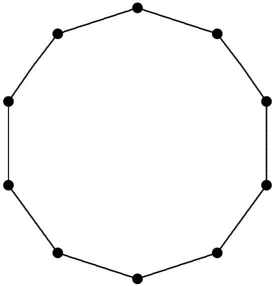

If the distance between i and j is w

$$V_{ij} = \frac{2w(v-w)}{v}\sigma^2,$$

$> 4\sigma^2$ if $v \geq 10$ and $3 \leq w$.

Here is an alternative design.

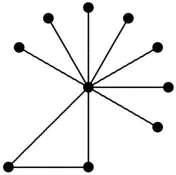

$V_{ij} \leq 4\sigma^2$ for all i, j.

Designs for $k = 2$ when $b = v$ (blocks shown as edges)

The only connected equireplicate design is the cycle.

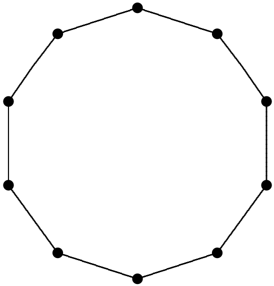

If the distance between i and j is w

$$V_{ij} = \frac{2w(v-w)}{v}\sigma^2,$$

$> 4\sigma^2$ if $v \geq 10$ and $3 \leq w$.

Here is an alternative design.

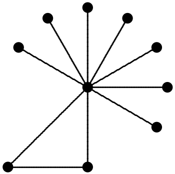

$V_{ij} \leq 4\sigma^2$ for all i, j.

A star attached to one vertex of a triangle is optimal for all $v \geq 12$.

Factorial designs

Example

In an experiment on sugar beet (in the UK, in 1935), the treatments were all combinations of levels of the following factors.

Sowing date	18 April	9 May	25 May
Spacing between rows	10 inches	15 inches	20 inches
Nitrogen fertilizer	none	0.3 cwt/acre	0.6 cwt/acre

There were three blocks of nine plots each.

Factorial designs

Example

In an experiment on sugar beet (in the UK, in 1935), the treatments were all combinations of levels of the following factors.

Sowing date	18 April	9 May	25 May
Spacing between rows	10 inches	15 inches	20 inches
Nitrogen fertilizer	none	0.3 cwt/acre	0.6 cwt/acre

There were three blocks of nine plots each.

Techniques from the theory of Abelian groups can be used to construct a suitable design.

An example of a two-phase experiment

The treatments are 10 varieties of common beans.

An example of a two-phase experiment

The treatments are 10 varieties of common beans.

In Phase I, these are grown in a field, in 10 blocks of size 6.

An example of a two-phase experiment

The treatments are 10 varieties of common beans.

In Phase I, these are grown in a field, in 10 blocks of size 6.

In Phase II, a sample of beans is taken from each plot.
Each sample is cooked in a special machine. The measured response is the time taken to properly cook the beans.

An example of a two-phase experiment

The treatments are 10 varieties of common beans.

In Phase I, these are grown in a field, in 10 blocks of size 6.

In Phase II, a sample of beans is taken from each plot.
Each sample is cooked in a special machine. The measured response is the time taken to properly cook the beans.

In Phase II, only four samples can be processed per day.
So we should treat days as 15 blocks of size 4.

An example of a two-phase experiment

The treatments are 10 varieties of common beans.

In Phase I, these are grown in a field, in 10 blocks of size 6.

In Phase II, a sample of beans is taken from each plot.
Each sample is cooked in a special machine. The measured response is the time taken to properly cook the beans.

In Phase II, only four samples can be processed per day.
So we should treat days as 15 blocks of size 4.

Now the design consists of
one function allocating bean varieties to plots in the field, and another function allocating each plot to a run of the cooking machine.

Computer search

At a conference on variety-testing in Słupia Wielka, Poland, in June 2018, Nha Vo-Thanh (Universität Hohenheim) gave a talk explaining his work with Hans-Peter Piepho on several different methods of computer search to find a good design for this experiment.

Computer search

At a conference on variety-testing in Słupia Wielka, Poland, in June 2018, Nha Vo-Thanh (Universität Hohenheim) gave a talk explaining his work with Hans-Peter Piepho on several different methods of computer search to find a good design for this experiment.

That evening, I got out some paper and a pen, and scribbled down some ideas, using my pattern approach. Very soon, I had a design with a smaller value of the average variance than he had found.

Computer search

At a conference on variety-testing in Słupia Wielka, Poland, in June 2018, Nha Vo-Thanh (Universität Hohenheim) gave a talk explaining his work with Hans-Peter Piepho on several different methods of computer search to find a good design for this experiment.

That evening, I got out some paper and a pen, and scribbled down some ideas, using my pattern approach. Very soon, I had a design with a smaller value of the average variance than he had found.

Since then, we have combined our approaches. Computer search may get stuck in a local optimum. Using a combinatorial approach to get a good starting design may overcome this.

Are statisticians just computers?

In the late twentieth century, several scientific research organizations got rid of all their statisticians, saying

We've all got computers now, so who needs statisticians?

Are statisticians just computers?

In the late twentieth century, several scientific research organizations got rid of all their statisticians, saying

We've all got computers now, so who needs statisticians?

I hope that my various stories convince you that statistical advice and collaboration,
in the design of the experiment, in the collection of data, and the data analysis, are worthwhile.

Ongoing collaboration

Ecologist Julia Reiss was working in a research group in the School of Biological Sciences at Queen Mary, University of London. They did an experiment about biodiversity, collected the data, then wondered about analysing it.

Ongoing collaboration

Ecologist Julia Reiss was working in a research group in the School of Biological Sciences at Queen Mary, University of London. They did an experiment about biodiversity, collected the data, then wondered about analysing it.

Julia: We need to consult a statistician

Ongoing collaboration

Ecologist Julia Reiss was working in a research group in the School of Biological Sciences at Queen Mary, University of London. They did an experiment about biodiversity, collected the data, then wondered about analysing it.

Julia: We need to consult a statistician
Group leader: We do not know any of them.

Ongoing collaboration

Ecologist Julia Reiss was working in a research group in the School of Biological Sciences at Queen Mary, University of London. They did an experiment about biodiversity, collected the data, then wondered about analysing it.

Julia: We need to consult a statistician
Group leader: We do not know any of them.

So Julia went away and searched on the web. She found my web page, where I advertised free statistical consultancy to QMUL researchers every Wednesday afternoon.

Ongoing collaboration

Ecologist Julia Reiss was working in a research group in the School of Biological Sciences at Queen Mary, University of London. They did an experiment about biodiversity, collected the data, then wondered about analysing it.

Julia: We need to consult a statistician
Group leader: We do not know any of them.

So Julia went away and searched on the web. She found my web page, where I advertised free statistical consultancy to QMUL researchers every Wednesday afternoon.

Also, the School of Mathematical Sciences was in a building just over the road from the School of Biological Sciences.

Ongoing collaboration

Ecologist Julia Reiss was working in a research group in the School of Biological Sciences at Queen Mary, University of London. They did an experiment about biodiversity, collected the data, then wondered about analysing it.

Julia: We need to consult a statistician
Group leader: We do not know any of them.

So Julia went away and searched on the web. She found my web page, where I advertised free statistical consultancy to QMUL researchers every Wednesday afternoon.

Also, the School of Mathematical Sciences was in a building just over the road from the School of Biological Sciences.

So she came to see me on the next Wednesday afternoon, and we have been collaborating ever since.

Biodiversity experiments

When we started, this seemed to be the received wisdom.

Treatments:	random set of species
Measured response Y:	some eco-desirable outcome
Conclusion:	the greater the number of different species, the better the outcome.

A more carefully controlled experiment

A, B, C, D, E, F — six types of freshwater "shrimp".

Put 12 shrimps in a jar containing stream water and alder leaf litter.

Measure how much leaf litter is eaten after 28 days.

A more carefully controlled experiment

A, B, C, D, E, F — six types of freshwater "shrimp".
Put 12 shrimps in a jar containing stream water and alder leaf litter.
Measure how much leaf litter is eaten after 28 days.

Experimental unit = jar.

A more carefully controlled experiment

A, B, C, D, E, F — six types of freshwater "shrimp".
Put 12 shrimps in a jar containing stream water and alder leaf litter.
Measure how much leaf litter is eaten after 28 days.

Experimental unit = jar.

Assemblage identity			Richness Level
A, …, F	monoculture	12 of type A	1
AB, …, EF	duoculture	6 of A, 6 of B	2
ABC, …, DEF	triculture	4 of A, 4 of B, 4 of C	3

A more carefully controlled experiment

A, B, C, D, E, F — six types of freshwater "shrimp".
Put 12 shrimps in a jar containing stream water and alder leaf litter.
Measure how much leaf litter is eaten after 28 days.

Experimental unit = jar.

	Assemblage identity			Richness Level
6	A, ..., F	monoculture	12 of type A	1
15	AB, ..., EF	duoculture	6 of A, 6 of B	2
20	ABC, ..., DEF	triculture	4 of A, 4 of B, 4 of C	3
41				

A more carefully controlled experiment

A, B, C, D, E, F — six types of freshwater "shrimp".
Put 12 shrimps in a jar containing stream water and alder leaf litter.
Measure how much leaf litter is eaten after 28 days.

Experimental unit = jar.

	Assemblage identity			Richness Level
6	A, …, F	monoculture	12 of type A	1
15	AB, …, EF	duoculture	6 of A, 6 of B	2
20	ABC, …, DEF	triculture	4 of A, 4 of B, 4 of C	3
41				

The experiment was carried out in 4 blocks of 41 jars.

A more carefully controlled experiment

A, B, C, D, E, F — six types of freshwater "shrimp".
Put 12 shrimps in a jar containing stream water and alder leaf litter.
Measure how much leaf litter is eaten after 28 days.

Experimental unit = jar.

	Assemblage identity			Richness Level
6	A, …, F	monoculture	12 of type A	1
15	AB, …, EF	duoculture	6 of A, 6 of B	2
20	ABC, …, DEF	triculture	4 of A, 4 of B, 4 of C	3
41				

The experiment was carried out in 4 blocks of 41 jars.
Actually 42 jars, because untreated jars were included,
but their data was so obviously different that it was excluded from further modelling.

Initial model fitting

The biologist fitted the model ‘Richness’ with 3 parameters,
one for each level of richness,
and found no evidence of any differences between the levels.

Initial model fitting

The biologist fitted the model 'Richness' with 3 parameters, one for each level of richness, and found no evidence of any differences between the levels.

This model for the response Y is

$$\mathbb{E}(Y) = \begin{cases} \alpha_1 & \text{on monocultures A}, \ldots, \text{F} \\ \alpha_2 & \text{on duocultures AB}, \ldots, \text{EF} \\ \alpha_3 & \text{on tricultures ABC}, \ldots, \text{DEF} \end{cases}$$

Initial model fitting

The biologist fitted the model 'Richness' with 3 parameters,
one for each level of richness,
and found no evidence of any differences between the levels.

This model for the response Y is

$$\mathbb{E}(Y) = \begin{cases} \alpha_1 & \text{on monocultures A}, \ldots, \text{F} \\ \alpha_2 & \text{on duocultures AB}, \ldots, \text{EF} \\ \alpha_3 & \text{on tricultures ABC}, \ldots, \text{DEF} \end{cases}$$

The data did not give any evidence against the null hypothesis that

$$\alpha_1 = \alpha_2 = \alpha_3 :$$

this is the 'Constant' model, or null model.

Call in a statistician

Assemblage identity			R	$x1$	$x2$	$x3$	$x4$	$x5$	$x6$
1	A	12 of type A	1	12	0	0	0	0	0
$\vdots$			$\vdots$						
6	F	12 of type F	1	0	0	0	0	0	12
7	AB	6 of A, 6 of B	2	6	6	0	0	0	0
$\vdots$			$\vdots$						
21	EF	6 of E, 6 of F	2	0	0	0	0	6	6
22	ABC	4 of A, 4 of B, 4 of C	3	4	4	4	0	0	0
$\vdots$			$\vdots$						
41	DEF	4 of D, 4 of E, 4 of F	3	0	0	0	4	4	4

Call in a statistician

	Assemblage identity		R	$x1$	$x2$	$x3$	$x4$	$x5$	$x6$
1	A	12 of type A	1	12	0	0	0	0	0
$\vdots$			$\vdots$						
6	F	12 of type F	1	0	0	0	0	0	12
7	AB	6 of A, 6 of B	2	6	6	0	0	0	0
$\vdots$			$\vdots$						
21	EF	6 of E, 6 of F	2	0	0	0	0	6	6
22	ABC	4 of A, 4 of B, 4 of C	3	4	4	4	0	0	0
$\vdots$			$\vdots$						
41	DEF	4 of D, 4 of E, 4 of F	3	0	0	0	4	4	4

I suggested the model ‘Type’ with 6 parameters $\beta_1, \ldots, \beta_6$:

$$\mathbb{E}(Y) = \sum_{i=1}^{6} \beta_i x_i$$

($\sum x_i = 12$ always, so no need for intercept.)

Family of expectation models (picture vs. formulae)

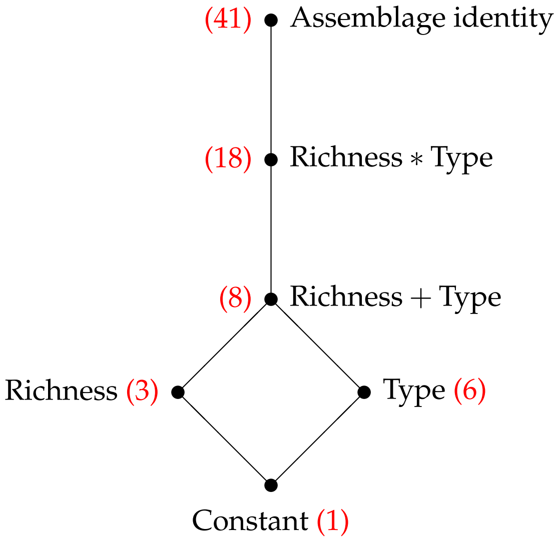

Family of expectation models (picture vs. formulae)

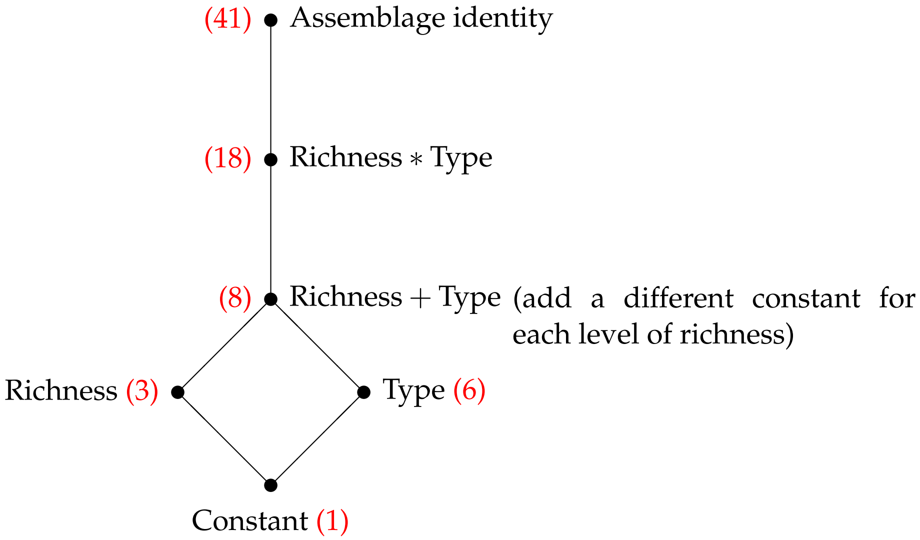

Family of expectation models (picture vs. formulae)

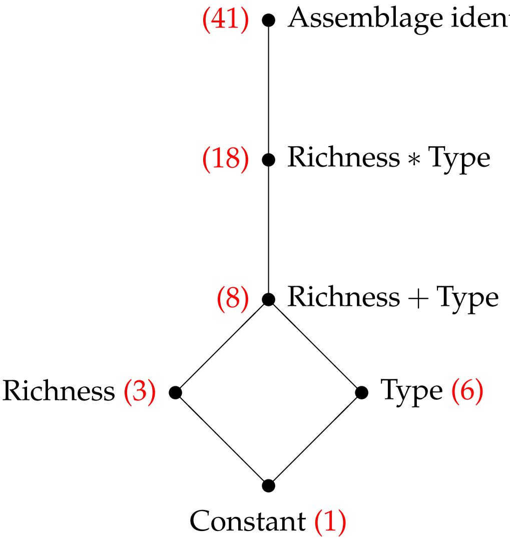

Richness $*$ Type: (β_i can change with each level of richness but does not depend on what else is present)

Richness + Type: (add a different constant for each level of richness)

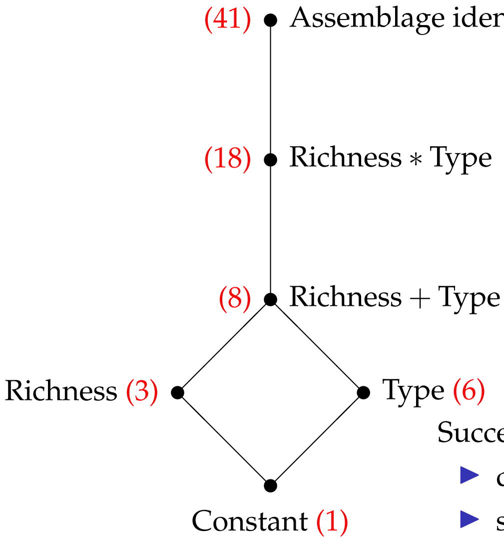

Richness ∗ Type (β_i can change with each level of richness but does not depend on what else is present)

Richness + Type (add a different constant for each level of richness)

Success: an ecology journal published

- diagram of family of models
- statement that each row of an ANOVA table is for a difference between models.

What the data showed: mean squares (picture/numbers)

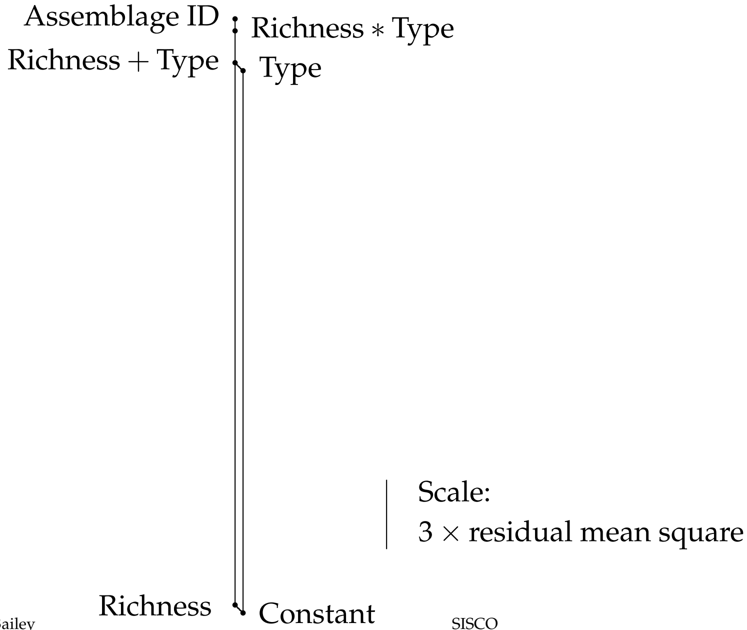

What the data showed: mean squares (picture/numbers)

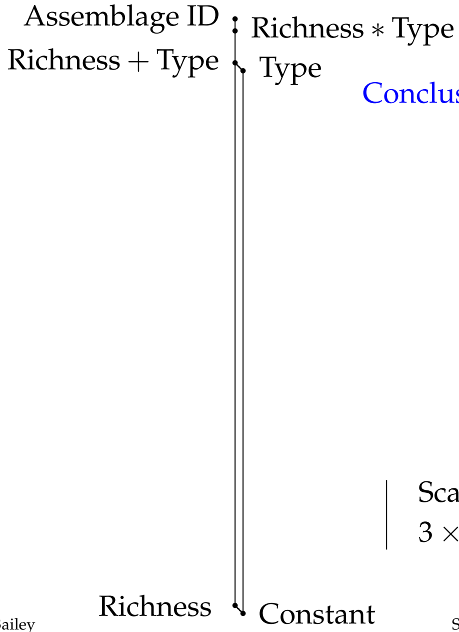

Conclusions:

Scale:
$3 \times$ residual mean square

What the data showed: mean squares (picture/numbers)

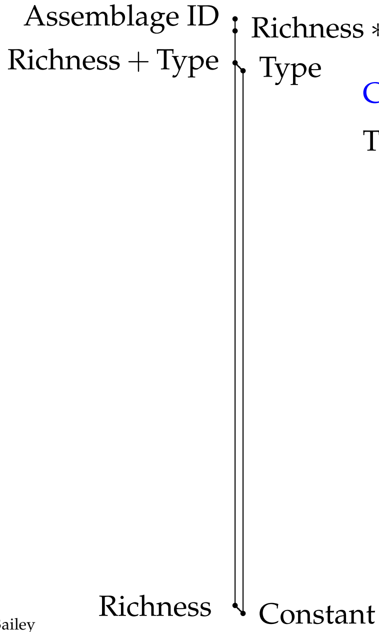

Conclusions:

The model Richness does not explain the data.

Scale:
3 × residual mean square

What the data showed: mean squares (picture/numbers)

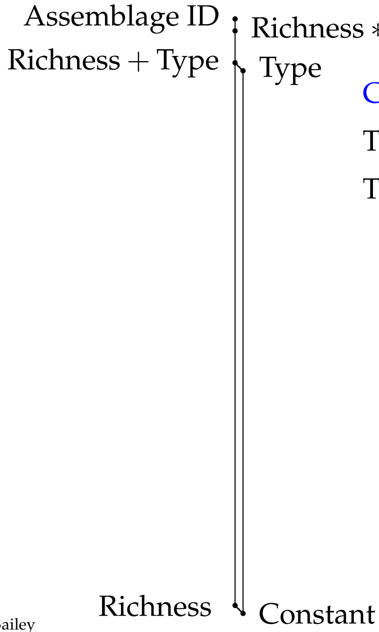

Conclusions:

The model Richness does not explain the data.

The model Type explains the data well.

Scale:
$3 \times$ residual mean square

What the data showed: mean squares (picture/numbers)

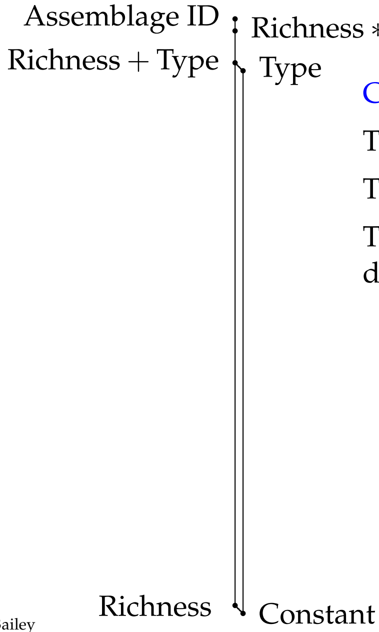

Conclusions:

The model Richness does not explain the data.

The model Type explains the data well.

There is no evidence that any larger model does any better.

Scale:
$3 \times$ residual mean square

What the data showed: mean squares (picture/numbers)

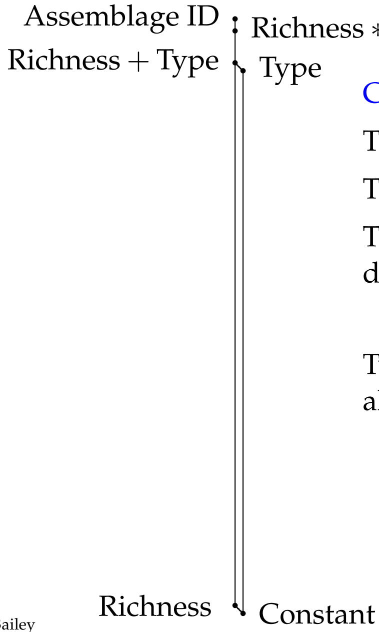

Conclusions:

The model Richness does not explain the data.

The model Type explains the data well.

There is no evidence that any larger model does any better.

Two experiments, with two responses each, all led to similar conclusions.

Scale:
$3 \times$ residual mean square

A new experiment on a different ecosystem (7 types)

Assemblage identity			Richness Level
A, ..., G	monoculture	12 of type A	1
AB, ..., FG	duoculture	6 of A, 6 of B	2
ABC, ..., EFG	triculture	4 of A, 4 of B, 4 of C	3

A new experiment on a different ecosystem (7 types)

	Assemblage identity			Richness Level
7	A, ..., G	monoculture	12 of type A	1
21	AB, ..., FG	duoculture	6 of A, 6 of B	2
35	ABC, ..., EFG	triculture	4 of A, 4 of B, 4 of C	3
63				

A new experiment on a different ecosystem (7 types)

	Assemblage identity			Richness Level
7	A, . . . , G	monoculture	12 of type A	1
21	AB, . . . , FG	duoculture	6 of A, 6 of B	2
35	ABC, . . . , EFG	triculture	4 of A, 4 of B, 4 of C	3
63				

"Do I really need all 35 tricultures?"

A new experiment on a different ecosystem (7 types)

	Assemblage identity			Richness Level
7	A, ..., G	monoculture	12 of type A	1
21	AB, ..., FG	duoculture	6 of A, 6 of B	2
35	ABC, ..., EFG	triculture	4 of A, 4 of B, 4 of C	3
63				

"Do I really need all 35 tricultures?"

"Use 7 tricultures making a balanced incomplete-block design."

A new experiment on a different ecosystem (7 types)

	Assemblage identity			Richness Level
7	A, …, G	monoculture	12 of type A	1
21	AB, …, FG	duoculture	6 of A, 6 of B	2
35	ABC, …, EFG	triculture	4 of A, 4 of B, 4 of C	3
63				

"Do I really need all 35 tricultures?"

"Use 7 tricultures making a balanced incomplete-block design."

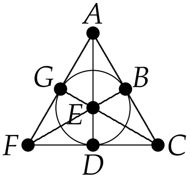

A new experiment on a different ecosystem (7 types)

	Assemblage identity			Richness Level
7	A, . . . , G	monoculture	12 of type A	1
21	AB, . . . , FG	duoculture	6 of A, 6 of B	2
35	ABC, . . . , EFG	triculture	4 of A, 4 of B, 4 of C	3
63				

"Do I really need all 35 tricultures?"

"Use 7 tricultures making a balanced incomplete-block design."

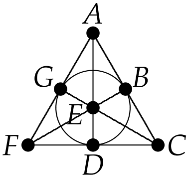

Another success: *Advances in Ecological Research* published this picture of the Fano plane.