

Two-factor experiments

Chapter 12.3.1 and 12.3.2 (pp. 488–499). Randomized block designs are in 12.3.3.

The CRD (completely randomized design) for comparing insecticides varies the levels of one factor (insecticide), while controlling other factors that influence survival time. The inferences from the one-way ANOVA apply to beetles with a given age from the selected strain that might be given the selected concentration of the insecticides. Any generalization of the conclusions to other situations must be justified scientifically, typically through further experimentation.

There are several ways to broaden the scope of the study. For example, several strains of beetles or several concentrations of the insecticide might be used. For simplicity, consider a simple two-factor experiment where three concentrations (Low, Medium, and High) are applied with each of the four insecticides. This is a completely crossed **two-factor experiment** where each of the $4 \times 3 = 12$ combinations of the two factors (insecticide and dose) are included in the comparison of survival times. With this experiment, the scientist can compare insecticides, compare concentrations, and check for an interaction between dose and insecticide.

Assuming that 48 beetles are available, the scientist would randomly assign them to the 12 experimental groups, giving prespecified numbers of beetles to the 12 groups. For simplicity, assume that the experiment is **balanced**, that is, the same number of beetles (4) is assigned to each group ($12 \times 4 = 48$). This is a CRD with two factors.

A Simplified View of Two-Factor Analysis

The following R output gives an analysis of survival times of groups of four beetles randomly allocated to twelve treatment groups obtained by crossing the levels of four insecticides (A,B,C,D) at each of three concentrations of the insecticides (1=Low, 2=Medium, 3=High):

```
> d <- read.table("c:/tim/PubH7400/insect.txt")
> dose <- factor(d[,1]); type <- factor(d[,2]); time <- d[,3]
> fit <- aov(time~dose+type+dose*type)
> summary(fit)
```

	Df	Sum Sq	Mean Sq	F value	Pr(>F)	
dose	2	1.03301	0.51651	23.2217	3.331e-07	***
type	3	0.92121	0.30707	13.8056	3.777e-06	***
dose:type	6	0.25014	0.04169	1.8743	0.1123	
Residuals	36	0.80072	0.02224			

This is a balanced 4-by-3 **factorial** design (two-factor design) that is replicated four times. Three variables are needed to uniquely represent each response: dose (1-3), insecticide (1-4 for insecticides A-D), and the survival time (called time). The unit of measure for the survival times is 10 hours. That is, 0.3 is a survival time of 3 hours. The data are given below.

	Dose 1	Dose 2	Dose 3
Insecticide 1	0.31 0.45 0.46 0.43	0.36 0.29 0.40 0.23	0.22 0.21 0.18 0.23
Insecticide 2	0.82 1.10 0.88 0.72	0.92 0.61 0.49 1.24	0.30 0.37 0.38 0.29
Insecticide 3	0.43 0.45 0.63 0.76	0.44 0.35 0.31 0.40	0.23 0.25 0.24 0.22
Insecticide 4	0.45 0.71 0.66 0.62	0.56 1.02 0.71 0.38	0.30 0.36 0.31 0.33

The basic *interaction* model for the two-factor design, as applied to this experiment is that the

$$\text{Response} = \text{Grand Mean} + \text{Dose Effect} + \text{Insect Effect} + \text{Dose} * \text{Insect Interaction} + \text{Residual}.$$

The assumptions for the analysis of the model are identical to those for a one-way ANOVA on the 12 treatment combinations (dose and insecticide).

Where $i = 1, 2, 3$ are dose levels, $j = 1, 2, 3, 4$ are insecticide levels, and $k = 1, 2, \dots, 12$ are the twelve beetles in each of the $3 \times 4 = 12$ treatment combinations, we write the model

$$Y_{ijk} = \mu + \alpha_i + \beta_j + (\alpha\beta)_{ij} + e_{ijk},$$

with the restrictions that $\sum_{i=1}^3 \alpha_i = 0$, $\sum_{j=1}^4 \beta_j = 0$, $\sum_{i=1}^3 (\alpha\beta)_{ij} = 0$, and $\sum_{j=1}^4 (\alpha\beta)_{ij} = 0$. These restrictions enhance the interpretability of the model, but we will not discuss them further.

The parameter μ is interpreted as an overall, or *grand* mean. The α_1, α_2 , and α_3 are the three main dose effects beyond μ , while the $\beta_1, \beta_2, \beta_3$ and β_4 are the four main insecticide effects beyond μ . There are *twelve* interaction terms $(\alpha\beta)_{ij}$. It is of interest to test the hypothesis that $H_0 : (\alpha\beta)_{ij} = 0$ for all i and j . If there is no interaction present we then have the *additive* model:

$$Y_{ijk} = \mu + \alpha_i + \beta_j + e_{ijk},$$

and the dose and insecticide effects simply add to the overall mean μ . This simpler model is a bit easier to interpret than the full model that includes an interaction term.

For example, with no interaction, for a fixed insecticide, say j , the mean difference in dose level 2 and dose level 1 is given by

$$[\mu + \alpha_2 + \beta_j] - [\mu + \alpha_1 + \beta_j] = \alpha_2 - \alpha_1.$$

This latter term does not involve the insecticide level j implying that the difference in mean survival times for sets of two doses is *independent of the insecticide used*. When an interaction is present this difference *does depend* on the insecticide used.

The ANOVA table for this experimental design decomposes the total variation in the data, as measured by the Total SS, into components that measure the variation of marginal sample means of dose and insecticide separately (Dose SS and Insecticide SS), *if interaction is included* a component that measures the degree to which the factors interact (dose by insecticide SS), and a component that pools the sample variances across the 12 samples (Error SS).

Each SS has a df , given in the following skeleton ANOVA table for the interaction model. As usual, the MS for each source of variation is the corresponding SS divided by the df . The MS Error estimates the common population variance for the 12 treatments.

Source	df	SS	MS
Dose	$3 - 1 = 2$		
Insecticide	$4 - 1 = 3$		
Interaction	$2 \times 3 = 6$		
Error	$48 - 12 = 36$		
Total	$48 - 1 = 47$		

There are three usual tests of interest, a test for no interaction in the model that contains the interaction term, and tests for no main effects in the additive model:

1. The test of no interaction between dose and insecticide, $H_0 : (\alpha\beta)_{ij} = 0$, is based on the p -value for the F -statistic: $F_{obs} = \text{MS Interaction} / \text{MS Error}$. Here, the p -value is 0.11 and we accept that we can drop the interaction term. We can then refit the model without the interaction term using `fit <- aov(time~dose+type)`. This will increase the SS Residual as the variability in survival times explained by the interaction term will be added into the error, or unexplained variation.
2. The test of no dose effect, $H_0 : \alpha_i = 0$, in the *additive* model. The absence of a dose effect implies that each dose level has the same **population mean response when the means are averaged over levels of insecticide**. The test for no dose effect is based on the p -value for the F -statistic: $F_{obs} = \text{MS Dose} / \text{MS Error}$. This hypothesis is rejected when the marginal means for dose vary significantly relative to the error variation.

3. The test of no insecticide effect, $H_0 : \beta_j = 0$, in the *additive* model. The absence of an insecticide effect implies that each level of insecticide has the same **population mean** response **when the means are averaged over levels of** dose. The test for no insecticide effect is based on the p -value for the F -statistic: $F_{obs} = \text{MS Insect} / \text{MS Error}$. This hypothesis is rejected when the insecticide marginal means vary significantly relative to the within sample variation.

To understand interaction, consider the Dose*Insecticide profile plot given below. For each dose, we have a plot of the *sample* mean survival times $\bar{Y}_{ij\bullet}$ across insecticides, giving 3 profiles. There is no interaction in the data if these profiles are parallel. The formal test for no interaction is a check on whether the profile plots of the *population* means are perfectly parallel.




Figure 1: Plot from `interaction.plot(dose,type,time)`.




Figure 2: Plot from `interaction.plot(type,dose,time)`.

These lines are not parallel, but they have the same general trend. If we accept that the interaction is not significant, we may fit the additive model instead and base inferences on additive structure.

Analysis of the Survival Times of Beetles

The R implementation of the two-factor analysis is similar to a randomized block analysis, except, if desired, you need to specify an interaction effect between dose and insecticide (a crossed effect) in the `aov()` model statement. The output is also similar, except for the added information on the interaction effect. We re-examine the R output for the interaction model:

```
              Df  Sum Sq Mean Sq F value    Pr(>F)
dose           2  1.03301  0.51651  23.2217 3.331e-07 ***
type           3  0.92121  0.30707  13.8056 3.777e-06 ***
dose:type      6  0.25014  0.04169   1.8743  0.1123
Residuals    36  0.80072  0.02224
---
```

The table provides SS for dose, insecticide, and the dose by insecticide interaction. The Mean Squares, F -statistics and p -values for testing these effects are given. The only real inference of interest here, the F -test for no dose by insecticide interaction, is not significant at the 0.10 level (p -value=0.1123). However, if you decide that the interaction *is* important (and here the p -value here *is* fairly small), then the mean differences in dose effect *changes with the level of insecticide*.

For example, we can look at how the mean differences in survival times between doses 1 and 3 *change* across insecticides A, B, C, and D. This is carried out by looking at *contrasts* in the model parameters. Specifically, the four differences of interest are

$$\begin{aligned} & [\mu + \alpha_3 + \beta_1 + (\alpha\beta)_{31}] - [\mu + \alpha_1 + \beta_1 + (\alpha\beta)_{11}], \\ & [\mu + \alpha_3 + \beta_2 + (\alpha\beta)_{32}] - [\mu + \alpha_1 + \beta_2 + (\alpha\beta)_{12}], \\ & [\mu + \alpha_3 + \beta_3 + (\alpha\beta)_{33}] - [\mu + \alpha_1 + \beta_3 + (\alpha\beta)_{13}], \text{ and} \\ & [\mu + \alpha_3 + \beta_4 + (\alpha\beta)_{34}] - [\mu + \alpha_1 + \beta_4 + (\alpha\beta)_{14}]. \end{aligned}$$

Each one of these comparisons is of the form $\mu_{3j} - \mu_{1j}$ where μ_{ij} is the average survival time for dose i and poison j and is given by $\mu_{ij} = \mu + \alpha_i + \beta_j + (\alpha\beta)_{ij}$ in the interaction model. **Stata** output for computing estimates, confidence intervals, and p -values that each of these four differences is zero follows.

```
. lincom _b[dose[3]]+_b[poison[1]]+_b[dose[3]*poison[1]]-
      _b[dose[1]]-_b[poison[1]]-_b[dose[1]*poison[1]]
```

time	Coef.	Std. Err.	t	P> t	[95% Conf. Interval]	
(1)	-.2025	.105457	-1.92	0.063	-.4163767	.0113767

```
. lincom _b[dose[3]]+_b[poison[2]]+_b[dose[3]*poison[2]]-
      _b[dose[1]]-_b[poison[2]]-_b[dose[1]*poison[2]]
```

time	Coef.	Std. Err.	t	P> t	[95% Conf. Interval]	
(1)	-.545	.105457	-5.17	0.000	-.7588767	-.3311233

```
. lincom _b[dose[3]]+_b[poison[3]]+_b[dose[3]*poison[3]]-
      _b[dose[1]]-_b[poison[3]]-_b[dose[1]*poison[3]]
```

time	Coef.	Std. Err.	t	P> t	[95% Conf. Interval]	
(1)	-.3325	.105457	-3.15	0.003	-.5463767	-.1186233

```
. lincom _b[dose[3]]+_b[poison[4]]+_b[dose[3]*poison[4]]-
      _b[dose[1]]-_b[poison[4]]-_b[dose[1]*poison[4]]
```

time	Coef.	Std. Err.	t	P> t	[95% Conf. Interval]	
(1)	-.285	.105457	-2.70	0.010	-.4988767	-.0711233

Obtaining these contrasts in R is only slightly more complicated. One approach is to use `esticon()` from the package `doBy`:

```
> fit <- lm(time~dose+type+dose*type)
> fit
Call: lm(formula = time ~ dose + type + dose * type)

Coefficients:
(Intercept)      dose2      dose3      type2      type3      type4
    0.4125    -0.0925    -0.2025     0.4675     0.1550     0.1975
dose2:type2 dose3:type2 dose2:type3 dose3:type3 dose2:type4 dose3:type4
    0.0275    -0.3425    -0.1000    -0.1300     0.1500    -0.0825

> esticon(fit,c(0,0,1,0,0,0,0,0,0,0,0))
Confidence interval ( WALD ) level = 0.95
  beta0 Estimate Std.Error  t.value DF Pr(>|t|)  Lower.CI  Upper.CI
1     0  -0.2025  0.105457 -1.920214 36 0.0627808 -0.4163767 0.01137673
> esticon(fit,c(0,0,1,0,0,0,0,1,0,0,0))
Confidence interval ( WALD ) level = 0.95
  beta0 Estimate Std.Error  t.value DF Pr(>|t|)  Lower.CI  Upper.CI
1     0  -0.545  0.105457 -5.167983 36 9e-06 -0.7588767 -0.3311233
> esticon(fit,c(0,0,1,0,0,0,0,0,0,1,0))
Confidence interval ( WALD ) level = 0.95
  beta0 Estimate Std.Error  t.value DF Pr(>|t|)  Lower.CI  Upper.CI
1     0  -0.3325  0.105457 -3.152944 36 0.0032533 -0.5463767 -0.1186233
> esticon(fit,c(0,0,1,0,0,0,0,0,0,0,1))
Confidence interval ( WALD ) level = 0.95
  beta0 Estimate Std.Error  t.value DF Pr(>|t|)  Lower.CI  Upper.CI
1     0  -0.285  0.105457 -2.702523 36 0.0104323 -0.4988767 -0.07112327
```

We see that when testing mean differences *individually*, we have significant differences between doses 1 and 3 for insecticides B, C, and D, but not A (p -values are A: 0.063, B: 0.000, C: 0.003, D: 0.010). When we bound the family error rate (FER) at 0.05 using Bonferroni's method, we find significant differences for only B and C (p -values are all multiplied by $c = 4$ yielding A: 0.252, B: < 0.005 , C: 0.0012, D: 0.050). We conclude that there is a significant mean difference between doses 1 and 3 for insecticides B and C.

If we decide that the interaction is *not* important, we may fit the additive model. The additive model forces the mean differences in dose effects to be exactly the same regardless of the insecticide. We may thus look at difference in doses 3 and 1 and simultaneously draw conclusions for all four insecticides. Below we look at the estimate of $\alpha_3 - \alpha_1$, which is the difference between μ_{3j} and μ_{1j} when $\mu_{ij} = \mu + \alpha_i + \beta_j$, as in the additive model.

```
. lincom _b[dose[3]]-_b[dose[1]]
```

```
-----  
time |      Coef.   Std. Err.      t    P>|t|     [95% Conf. Interval]  
-----+-----  
(1) |   -0.34125   0.0559247    -6.10  0.000   -0.4541105   -0.2283895  
-----
```

OR

```
> fit <- lm(time~dose+type)
```

```
> coef(fit)
```

```
(Intercept)      dose2      dose3      type2      type3      type4  
0.45229167 -0.07312500 -0.34125000 0.36250000 0.07833333 0.22000000
```

```
> esticon(fit,c(0,0,1,0,0,0))
```

```
Confidence interval ( WALD ) level = 0.95
```

```
beta0 Estimate  Std.Error  t.value DF Pr(>|t|)  Lower.CI  Upper.CI  
1      0 -0.34125 0.05592465 -6.10196 42 3e-07 -0.4541105 -0.2283895
```

Checking assumptions

The $(ijk)^{th}$ fitted value is denoted $\widehat{E}(Y_{ijk})$, and estimates $E(Y_{ijk})$ under the model. For example, under an additive model it is given by $\widehat{E}(Y_{ijk}) = \hat{\mu} + \hat{\alpha}_i + \hat{\beta}_j$, under the interaction model it is $\widehat{E}(Y_{ijk}) = \hat{\mu} + \hat{\alpha}_i + \hat{\beta}_j + \widehat{(\alpha\beta)}_{ij}$. The $(ijk)^{th}$ residual is what we saw versus what is expected under the model, $r_{ijk} = y_{ijk} - \widehat{E}(Y_{ijk})$. These are obtained `fit$fit` and `fit$res` in R and estimate the unknown errors e_{ijk} .

We should check two things. The first is that the observed data variability is roughly constant across factor levels and/or predicted values. In either case there should be no pattern. The second is that the residuals are approximately normal; a histogram could be viewed, a normal probability plot (not discussed) assessed, or a formal test can be carried out.

A plot of the residuals versus the predicted values suggests that the variability in the survival times increases with increasing mean.

Transforming survival times to the reciprocal scale (which turns the response into a rate) is often suggested with these data, and in this case eliminates evidence of interaction and non-constant variance.

Analyzing beetles expiration rates will be a homework problem.




Figure 3: Plot from `plot(fitfit,fitres)`.

An unbalanced two-factor experiment and analysis: insulin levels in rats

The sample sizes are rarely equal for the different treatments in an experiment. This has no consequence on the specification of a model, and we proceed as in the balanced case.

The data below are the insulin levels in rats a certain length of time after a fixed dose of insulin was injected into their jugular or portal veins. This is a two-factor study with two vein types (jugular=1, portal=2) and three time levels (0 minutes = 1, 30 minutes = 2, and 60 minutes = 3). A feature of this experiment is that the rats used in the six vein and time combinations are distinct. I will fit a two-factor interaction model, which assumes that the responses are independent within and across treatments. The design is unbalanced, with sample sizes varying from 3 to 12.

Vein	Time	Insulin Levels
jugular	0	18 36 12 24 43
jugular	30	61 116 63 132 68 37
jugular	60	18 133 33
portal	0	96 72 34 41 98 77 120 49 92 111 99 94
portal	30	146 193 78 127 136 144 115 199 253 338
portal	60	132 110 141 204 69 152 196 195 84 105 71 83

An alternative experimental design might randomly assign rats to the two vein groups, and then measure the insulin levels of each rat at the three time points. Depending on the questions of interest, you could compare veins using a one-way MANOVA, or a repeated measures design that allows correlated responses within rats.

```
> d <- read.table("c:/tim/PubH7400/rats.txt")
```

```
> d
```

	V1	V2	V3
1	1	1	18
2	1	1	36
3	1	1	12
4	1	1	24
5	1	1	43
6	1	2	61
7	1	2	116
8	1	2	63
9	1	2	132
10	1	2	68

et cetera...

38	2	3	110
39	2	3	141
40	2	3	204
41	2	3	69
42	2	3	152
43	2	3	196
44	2	3	195
45	2	3	84
46	2	3	105
47	2	3	71
48	2	3	83

The interaction model is written

$$Y_{ijk} = \mu + \alpha_i + \beta_j + (\alpha\beta)_{ij} + \epsilon_{ijk}.$$

Here, $i = 1, 2$ denote the two vein types, $j = 1, 2, 3$ denote the three times, and $k = 1, 2, \dots, n_{ij}$ denotes the k^{th} rat out of n_{ij} in the group with vein i and time j . You should verify that $n_{11} = 5$, $n_{12} = 6$, $n_{13} = 3$, $n_{21} = 12$, $n_{22} = 10$, and $n_{23} = 12$.

The profile or *interaction* plot shows roughly parallel lines indicating that the interaction term may not be important.




Figure 4: Roughly parallel lines indicate lack of interaction effect.

The profile plot indicates that the insulin level is at its highest (of the three times considered) at 30 minutes for either vein considered alone. We fit the model that includes the vein by time interaction and find

```
> vein=factor(d[,1]); time=factor(d[,2]); insulin=d[,3]
> fit <- aov(insulin~vein+time+vein*time)
> summary(fit)
```

	Df	Sum Sq	Mean Sq	F value	Pr(>F)	
vein	1	46400	46400	18.9363	8.465e-05	***
time	2	50332	25166	10.2705	0.0002338	***
vein:time	2	2746	1373	0.5603	0.5752396	
Residuals	42	102914	2450			

The ANOVA table lists **Type I** SS. When the design is unbalanced there are two types of SS one could consider: Type I and Type III. `aov()` gives Type I. Either is fine for determining whether the interaction should be included. I will briefly describe the associated tests in class.

We see that the interaction is not significant (p -value=0.575). Recall that the jugular and portal profiles are reasonably parallel, which is consistent with a lack of interaction. Because of the lack of interaction, the difference in mean levels for the portal veins is reasonably consistent across times.

Since we accept that there is no interaction here, it makes sense to compare the overall main effects *vein* and *time* using pairwise comparisons from a fit of the additive model. First the Type I ANOVA table:

```
> fit <- aov(insulin~vein+time)
> summary(fit)
```

	Df	Sum Sq	Mean Sq	F value	Pr(>F)	
vein	1	46400	46400	19.323	6.883e-05	***
time	2	50332	25166	10.480	0.0001896	***
Residuals	44	105660	2401			

The test of $H_0 : \alpha_1 = \alpha_2 = 0$ in the model $Y_{ijk} = \mu + \alpha_i + e_{ijk}$ yields a p -value of 0.000068. The test of $H_0 : \beta_1 = \beta_2 = \beta_3 = 0$ in the model $Y_{ijk} = \mu + \alpha_i + \beta_k + e_{ijk}$ is 0.00019.

Pairwise comparisons:

```
> TukeyHSD(fit)
```

```
  Tukey multiple comparisons of means
```

```
 95% family-wise confidence level
```

```
$vein
```

	diff	lwr	upr	p adj
2-1	68.40336	37.04159	99.76513	6.88e-05

```
$time
```

	diff	lwr	upr	p adj
2-1	77.76057	36.3606232	119.160509	0.0001201
3-1	42.98164	0.8767947	85.086494	0.0445013
3-2	-34.77892	-77.4960892	7.938246	0.1304624

We see that, on average, the insulin level is between 37 and 100 higher for the portal versus the jugular veins *regardless of the time level*. Since there is no interaction here, the difference in insulin levels is the same at time = 0, time = 30, and time = 60 minutes. Similarly we may look at *differences* in insulin levels at the three times *independent of vein type*. The *p*-values for testing that there is no difference between (1) 60 minutes and 30 minutes, (2) 30 minutes and 0 minutes, and (3) 60 minutes and 0 minutes are (1) 0.130, (2) 0.044, and (3) 0.000 with an FER of 0.05. We accept the first null but reject the latter two.

What are your thoughts on the residual plot?




Figure 5: `plot(fitfit,fitres)`.

My thoughts are that the model doesn't fit as well as it could.

Again, we could try the reciprocal transformation. Another commonly-used transformation in such cases is the natural logarithm.

The model then becomes

$$\log Y_{ijk} = \mu + \alpha_i + \beta_j + (\alpha\beta)_{ij} + \epsilon_{ijk}.$$

Again, we see the interaction is not needed:

```
> fit <- aov(log(insulin)~vein+time+vein*time)
> summary(fit)
              Df Sum Sq Mean Sq F value    Pr(>F)
vein           1  8.6458   8.6458 28.2496 3.385e-06 ***
time           1  2.1735   2.1735  7.1019  0.01072 *
vein:time      1  0.1438   0.1438  0.4698  0.49665
Residuals    44 13.4662   0.3061

> fit <- aov(log(insulin)~vein+time) # fit additive model
> plot(fit$fit,fit$res) # residuals versus predicted values; okay
> hist(fit$res)        # histogram of residuals; looks okay
> boxplot(fit$res~vein) # variability versus vein type; looks okay
> boxplot(fit$res~time) # variability versus time; looks okay
```

The residual plot looks a lot better...




Figure 6: `plot(fitfit,fitres)`.

The final model is additive for log-insulin values:

$$\log Y_{ijk} = \mu + \alpha_i + \beta_j + \epsilon_{ijk}.$$

This implies

$$\log Y_{ijk} \sim N(\mu + \alpha_i + \beta_j, \sigma^2),$$

and so

$$P(\log Y_{ijk} \leq \mu + \alpha_i + \beta_j) = 0.5,$$

implying

$$P(Y_{ijk} \leq e^{\mu + \alpha_i + \beta_j}) = 0.5.$$

The *median* of insulin level Y_{ijk} is $e^{\mu + \alpha_i + \beta_j}$.

Therefore fixing a time level at j , comparing the ratio of medians across veins gives

$$\frac{e^{\mu+\alpha_2+\beta_j}}{e^{\mu+\alpha_1+\beta_j}} = e^{\alpha_2-\alpha_1}.$$

Similarly, the ratio of medians across time gives

$$e^{\beta_3-\beta_2}, e^{\beta_3-\beta_1}, \text{ and } e^{\beta_2-\beta_1}.$$

```
> TukeyHSD(fit)
  Tukey multiple comparisons of means
    95% family-wise confidence level

Fit: aov(formula = log(insulin) ~ vein + time)

$vein
      diff      lwr      upr p adj
2-1 0.9337261 0.6338878 1.233564 1e-07

$time
      diff      lwr      upr    p adj
2-1 0.8558846 0.46007501 1.25169427 0.0000126
3-1 0.4948461 0.09229708 0.89739505 0.0126916
3-2 -0.3610386 -0.76944171 0.04736456 0.0925955
```

These are 95% *simultaneous* CI's for time, and the “regular” CI for vein, for which there's only one comparison.

The 95% CI for $\alpha_2 - \alpha_1$ is (0.63, 1.23).

$$0.63 \leq \alpha_2 - \alpha_1 \leq 1.23 \Leftrightarrow e^{0.63} \leq e^{\alpha_2 - \alpha_1} \leq e^{1.23}.$$

That is

$$1.9 \leq \frac{\text{Median insulin level vein 2}}{\text{Median insulin level vein 1}} \leq 3.4.$$

The median insulin levels differ significantly across veins ($p \approx 10^{-7}$); specifically, median insulin from vein one is between 1.9 to 3.4 times greater than the median insulin level from vein two, for a fixed time.

What can you say about time for a fixed vein?