```
> model1 <- lm(yield ~ Treatment + After1970, data = Yields)
>
> coef(model1)
      (Intercept) TreatmentFertilised     TreatmentManure    TreatmentStopped      After1970Before
        1.5145833           1.9616667           3.0122222           0.8094444           -0.9035417
> confint(model1)
                         2.5 %     97.5 %
(Intercept)          1.1006485  1.9285182
TreatmentFertilised  1.4836959  2.4396374
TreatmentManure      2.5342515  3.4901930
TreatmentStopped     0.3314737  1.2874152
After1970Before     -1.2620197 -0.5450636
```

Here you have the results of a linear model using the Rothamsted data from week's 8 and 9.

The model has a response of plant yield and explanatory variables of **fertiliser treatment** and **time** (before and after 1970). The data have been plotted to the right. The first column shows the effect of treatment and the second column shows the effect of time within each treatment. So, the first column shows pre-1970 only. The second column shows both times.

Using the output from the model above, try and draw the model line (defined by the beta values/coefficient estimates) for each effect.

Hint: think about how many coefficients you have so how many different slopes you expect.

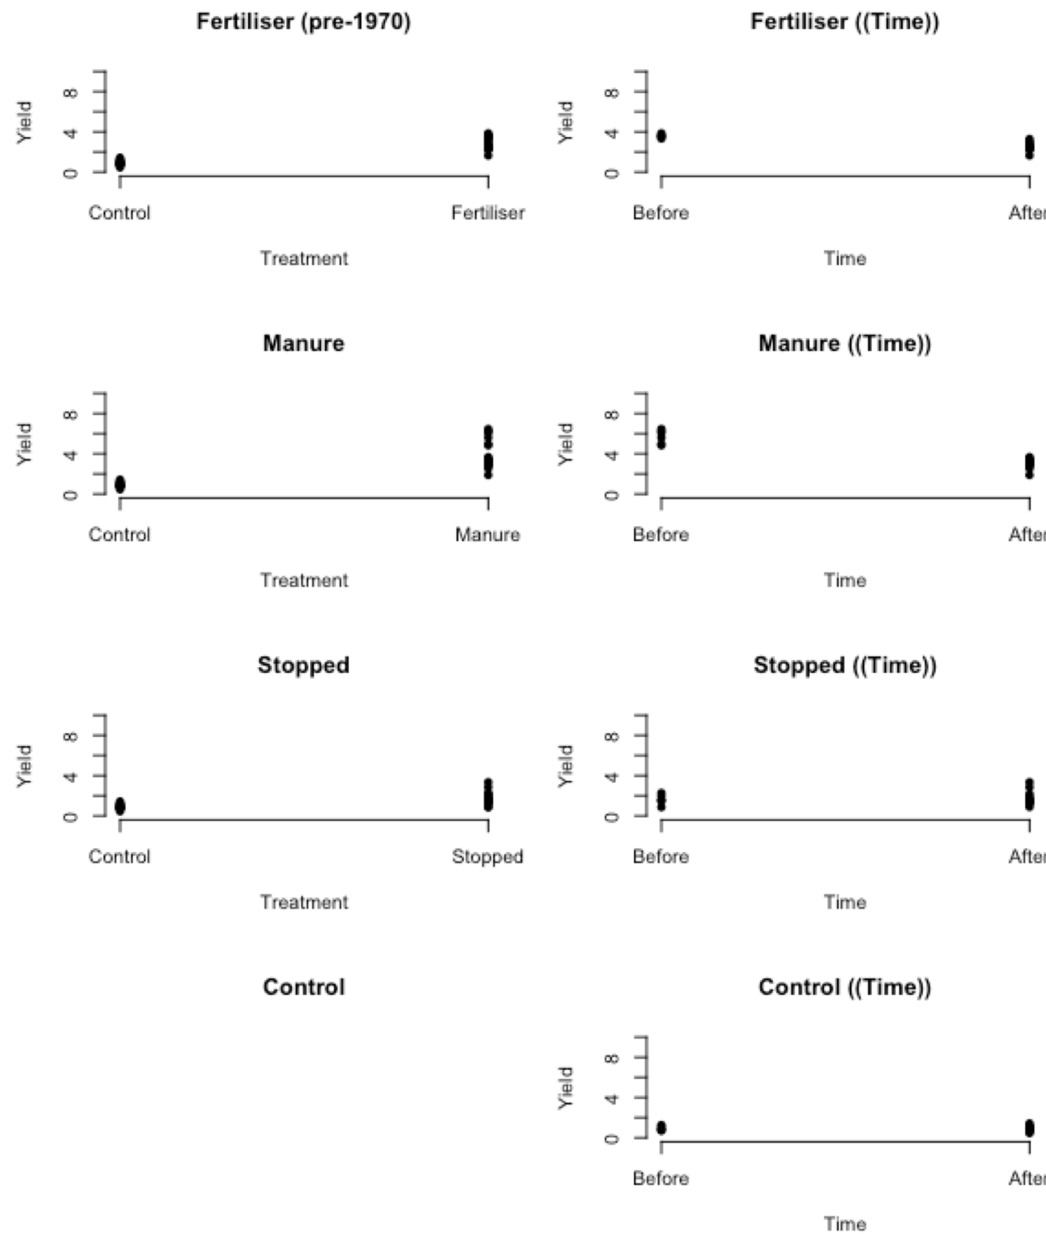

```
> model2 <- lm(yield ~ Treatment * After1970, data = Yields)
>
> coef(model2)
                     (Intercept)                TreatmentFertilised                      TreatmentManure
                      0.94166667                         2.64000000                           4.75500000
                TreatmentStopped                   After1970Before TreatmentFertilised:After1970Before
                      0.68000000                        -0.04416667                          -1.01750000
    TreatmentManure:After1970Before    TreatmentStopped:After1970Before
                     -2.61416667                         0.19416667
> confint(model2)
                                        2.5 %     97.5 %
(Intercept)                         0.5499153  1.3334180
TreatmentFertilised                 2.0859799  3.1940201
TreatmentManure                     4.2009799  5.3090201
TreatmentStopped                    0.1259799  1.2340201
After1970Before                    -0.5239621  0.4356288
TreatmentFertilised:After1970Before -1.6960332 -0.3389668
TreatmentManure:After1970Before    -3.2926999 -1.9356334
TreatmentStopped:After1970Before   -0.4843666  0.8726999
```

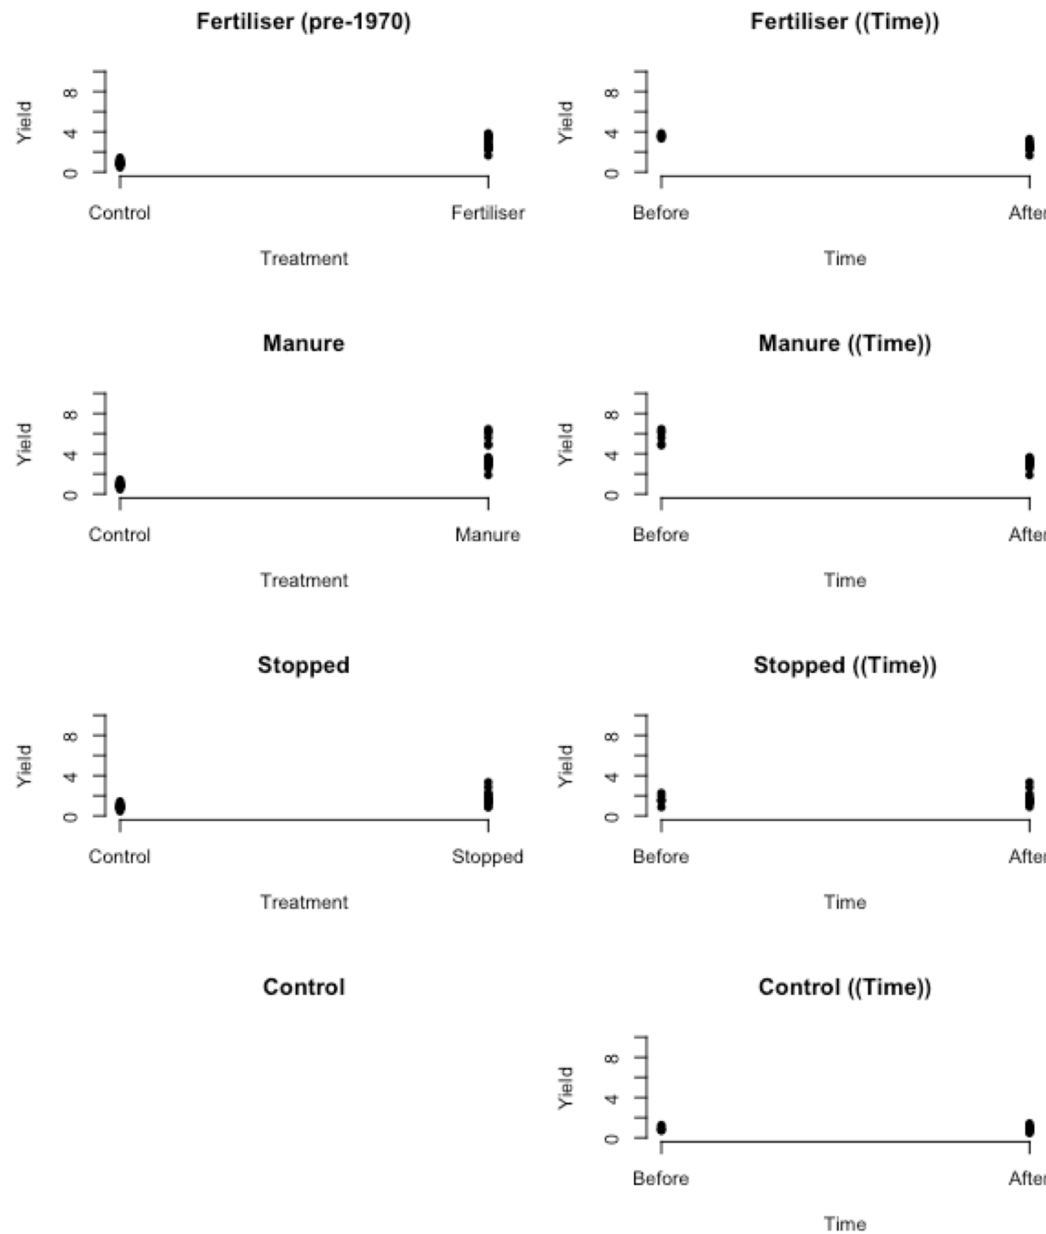

Now, you have an interaction model from the same variables. Repeat the exercise from before.

How many lines do you have this time? What has changed?

```
> model1 <- lm(yield ~ Treatment + After1970, data = Yields)
>
> coef(model1)
      (Intercept) TreatmentFertilised      TreatmentManure     TreatmentStopped       After1970Before
        1.5145833           1.9616667            3.0122222            0.8094444            -0.9035417
> confint(model1)
                         2.5 %     97.5 %
(Intercept)          1.1006485  1.9285182
TreatmentFertilised  1.4836959  2.4396374
TreatmentManure      2.5342515  3.4901930
TreatmentStopped     0.3314737  1.2874152
After1970Before     -1.2620197 -0.5450636
```

Here you have the results of a linear model using the Rothamsted data from week's 8 and 9.

The model has a response of plant yield and explanatory variables of **fertiliser treatment** and **time** (before and after 1970). The data have been plotted to the right. The first column shows the effect of treatment and the second column shows the effect of time within each treatment. So, the first column shows pre-1970 only. The second column shows both times.

Using the output from the model above, try and draw the model line (defined by the beta values/coefficient estimates) for each effect.

Hint: think about how many coefficients you have so how many different slopes you expect.

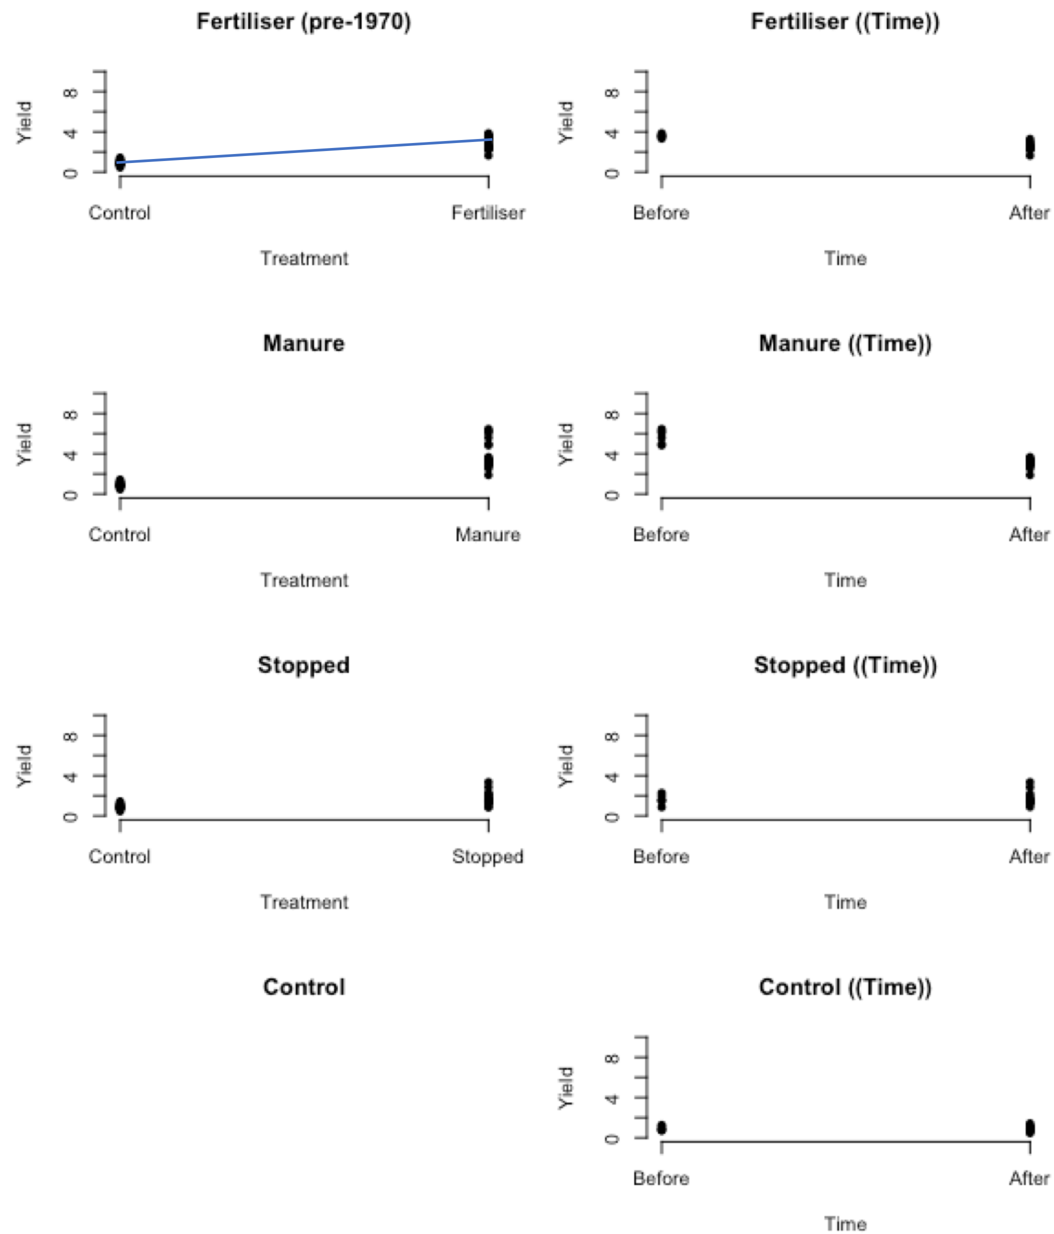

```
> model1 <- lm(yield ~ Treatment + After1970, data = Yields)
>
> coef(model1)
      (Intercept) TreatmentFertilised    TreatmentManure   TreatmentStopped      After1970Before
        1.5145833           1.9616667          3.0122222          0.8094444           -0.9035417
> confint(model1)
                        2.5 %     97.5 %
(Intercept)         1.1006485  1.9285182
TreatmentFertilised 1.4836959  2.4396374
TreatmentManure     2.5342515  3.4901930
TreatmentStopped    0.3314737  1.2874152
After1970Before    -1.2620197 -0.5450636
```

Here you have the results of a linear model using the Rothamsted data from week's 8 and 9.

The model has a response of plant yield and explanatory variables of **fertiliser treatment** and **time** (before and after 1970). The data have been plotted to the right. The first column shows the effect of treatment and the second column shows the effect of time within each treatment. So, the first column shows pre-1970 only. The second column shows both times.

Using the output from the model above, try and draw the model line (defined by the beta values/coefficient estimates) for each effect.

Hint: think about how many coefficients you have so how many different slopes you expect.

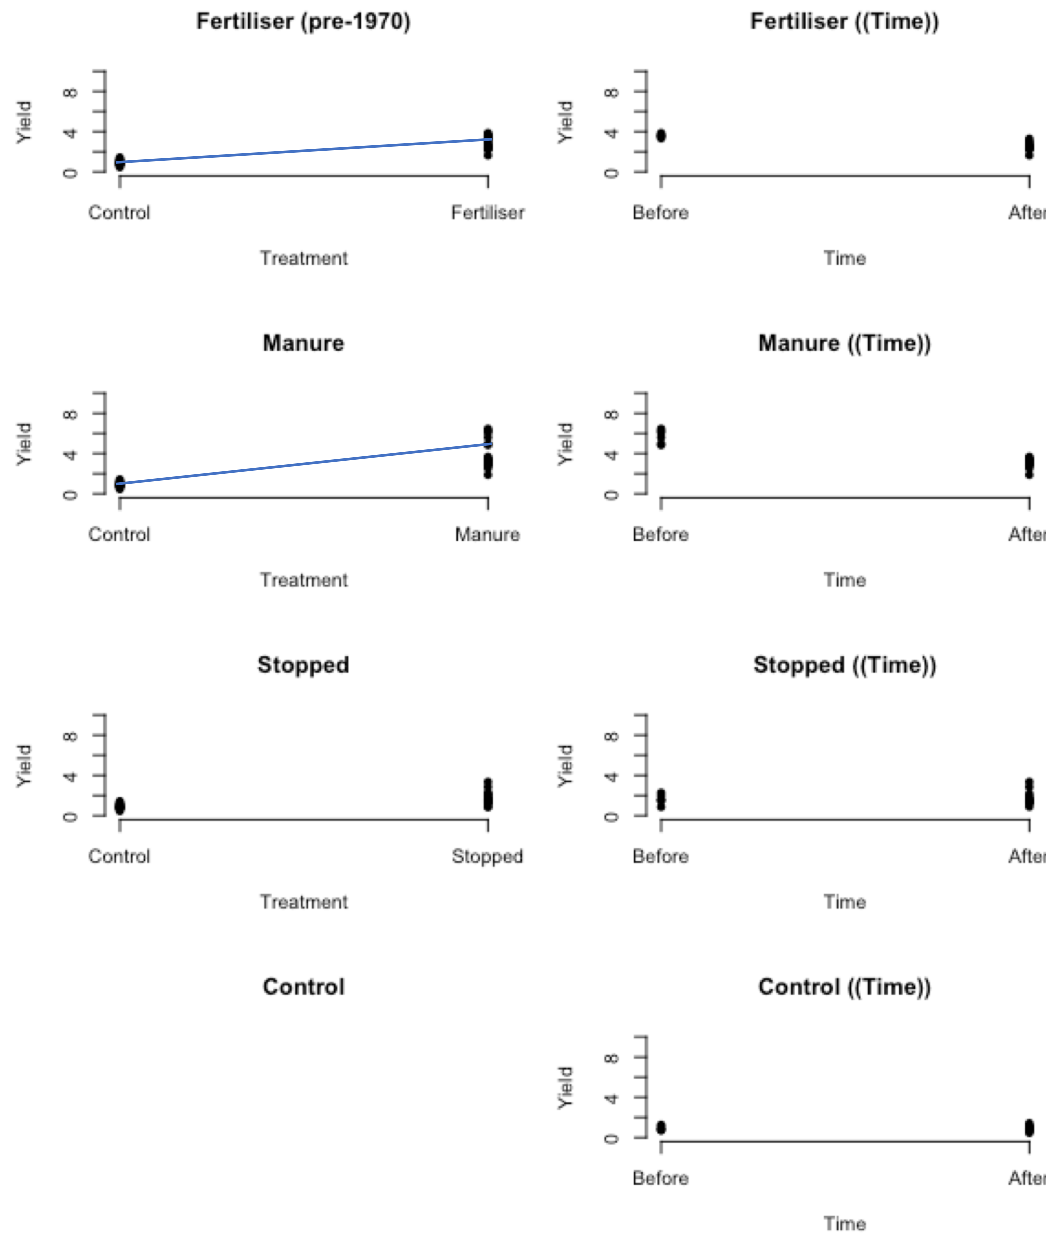

```
> model1 <- lm(yield ~ Treatment + After1970, data = Yields)
>
> coef(model1)
      (Intercept) TreatmentFertilised      TreatmentManure     TreatmentStopped      After1970Before
        1.5145833           1.9616667            3.0122222            0.8094444           -0.9035417
> confint(model1)
                         2.5 %     97.5 %
(Intercept)          1.1006485  1.9285182
TreatmentFertilised  1.4836959  2.4396374
TreatmentManure      2.5342515  3.4901930
TreatmentStopped     0.3314737  1.2874152
After1970Before     -1.2620197 -0.5450636
```

Here you have the results of a linear model using the Rothamsted data from week's 8 and 9.

The model has a response of plant yield and explanatory variables of **fertiliser treatment** and **time** (before and after 1970). The data have been plotted to the right. The first column shows the effect of treatment and the second column shows the effect of time within each treatment. So, the first column shows pre-1970 only. The second column shows both times.

Using the output from the model above, try and draw the model line (defined by the beta values/coefficient estimates) for each effect.

Hint: think about how many coefficients you have so how many different slopes you expect.

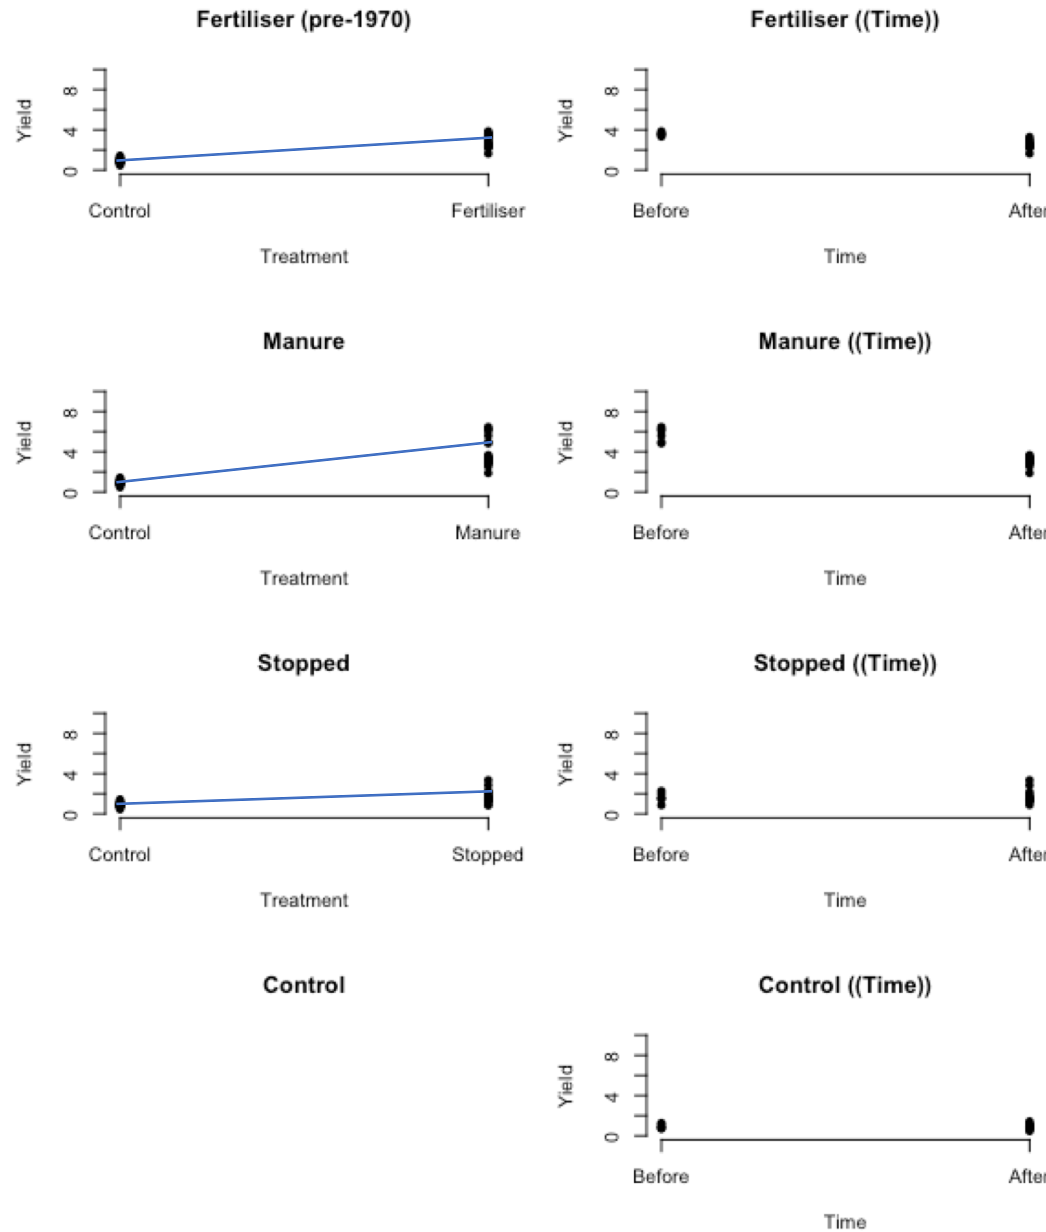

```
> model1 <- lm(yield ~ Treatment + After1970, data = Yields)
>
> coef(model1)
       (Intercept) TreatmentFertilised     TreatmentManure     TreatmentStopped     After1970Before
         1.5145833           1.9616667           3.0122222            0.8094444          -0.9035417
> confint(model1)
                          2.5 %     97.5 %
(Intercept)           1.1006485  1.9285182
TreatmentFertilised   1.4836959  2.4396374
TreatmentManure       2.5342515  3.4901930
TreatmentStopped      0.3314737  1.2874152
After1970Before      -1.2620197 -0.5450636
```

Only one estimate for time – so the same in each group. Similar to when we have continuous variables. The ideas are the same for both.

Here you have the results of a linear model using the Rothamsted data from week's 8 and 9.

The model has a response of plant yield and explanatory variables of **fertiliser treatment** and **time** (before and after 1970). The data have been plotted to the right. The first column shows the effect of treatment and the second column shows the effect of time within each treatment. So, the first column shows pre-1970 only. The second column shows both times.

Using the output from the model above, try and draw the model line (defined by the beta values/coefficient estimates) for each effect.

Hint: think about how many coefficients you have so how many different slopes you expect.

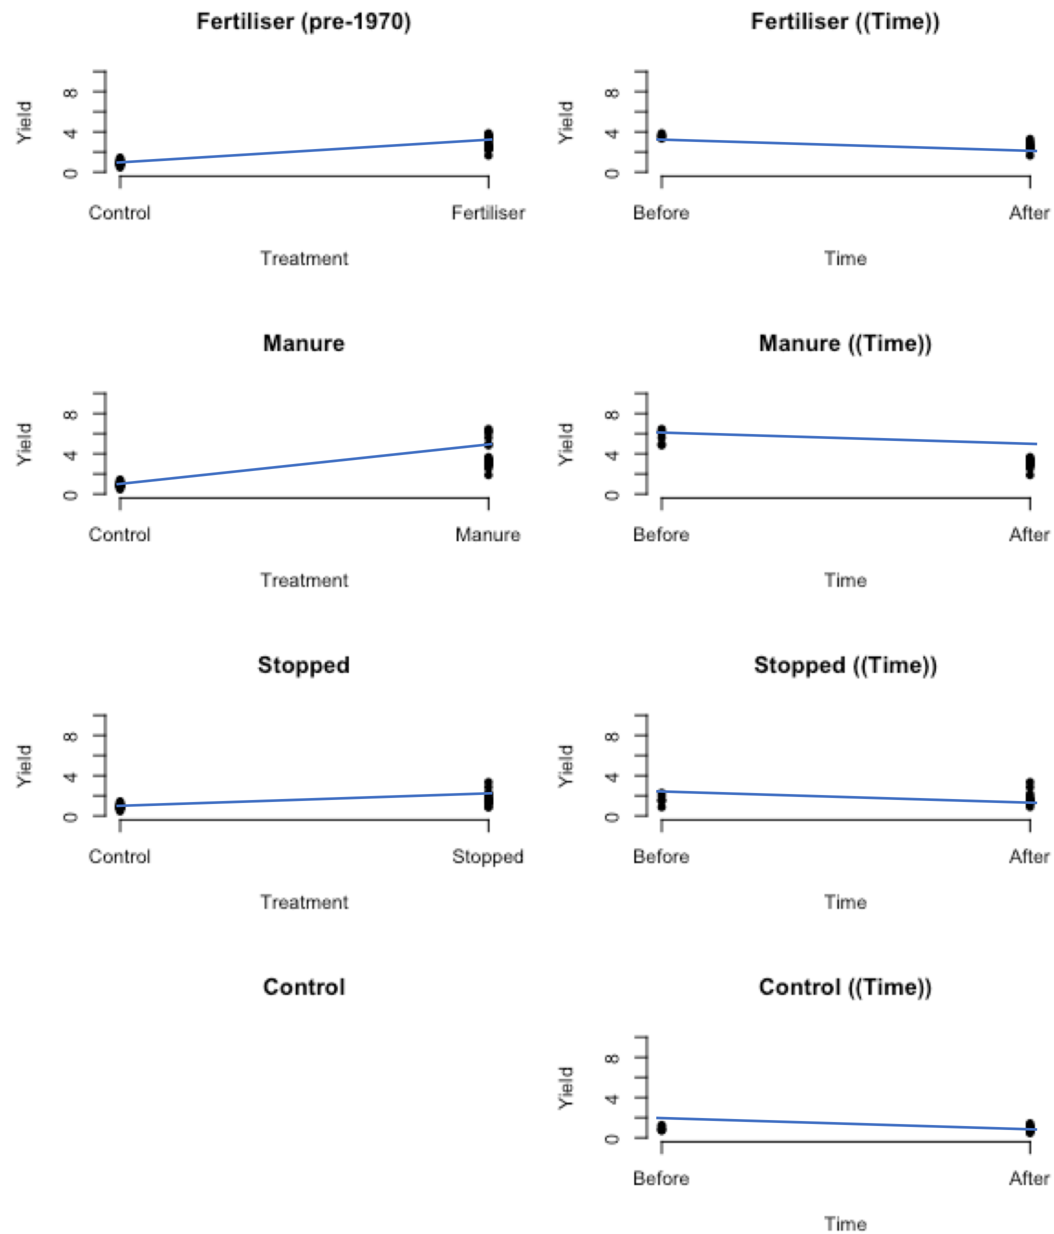

```
> model2 <- lm(yield ~ Treatment * After1970, data = Yields)
>
> coef(model2)
                    (Intercept)             TreatmentFertilised                      TreatmentManure
                     0.94166667                      2.64000000                           4.75500000
               TreatmentStopped                 After1970Before  TreatmentFertilised:After1970Before
                     0.68000000                     -0.04416667                          -1.01750000
TreatmentManure:After1970Before TreatmentStopped:After1970Before
                    -2.61416667                      0.19416667
> confint(model2)
                                         2.5 %     97.5 %
(Intercept)                          0.5499153  1.3334180
TreatmentFertilised                  2.0859799  3.1940201
TreatmentManure                      4.2009799  5.3090201
TreatmentStopped                     0.1259799  1.2340201
After1970Before                     -0.5239621  0.4356288
TreatmentFertilised:After1970Before -1.6960332 -0.3389668
TreatmentManure:After1970Before     -3.2926999 -1.9356334
TreatmentStopped:After1970Before    -0.4843666  0.8726999
```

Each slope is different!

Here we have interaction terms, they give us the difference in slope of time from the control effect to all other groups. So, we now have different slopes for time in each level of treatment.

E.g. the effect of time (**from after to before**) in the fertiliser group is -0.044 – (-1.0175)

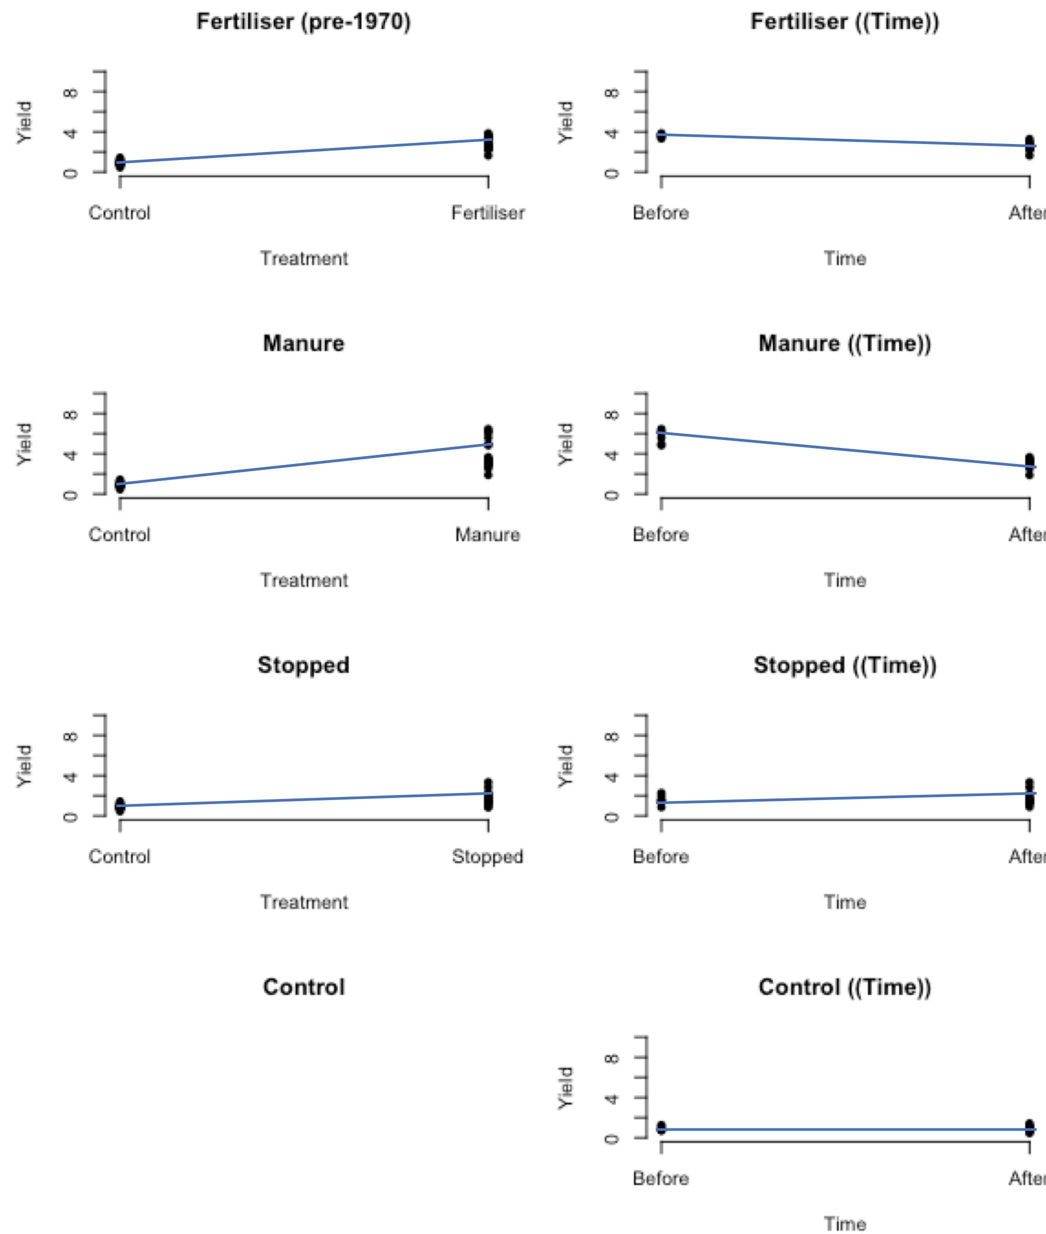

Now, you have an interaction model from the same variables. Repeat the exercise from before.

How many lines do you have this time? What has changed?