One-Factor Repeated Measures Design

James H. Steiger

Department of Psychology and Human Development
Vanderbilt University

P311, 2012

One-Factor Repeated Measures Design

One-Factor Repeated Measures Design

Introduction

- In the *one-way repeated measures* design, n randomly sampled subjects are measured repeatedly on a occasions. This is sometimes called the "subjects by trials" design for that reason.
- In the previous module, we discussed the *randomized blocks* design, which is in fact a special case of a mixed-effects two-way ANOVA, with 1 observation per cell, or a units per block.
- I pointed out that, when the "block size" is a, the a observations can be either different individuals matched by blocking, or repeated measurements on the same observational unit or subject.
- The one-way repeated measures design that employs the "nonadditive model" and assumes that the blocking factor (i.e., "Subjects") is a random effect, is identical to the randomized blocks model with a fixed effect treatment and one observation per cell that we examined earlier.
- The "additive model" discussed by MWS in section 14.2 of RDASA3 seems unrealistic. It assumes that the effect over repeated measures is the same for all subjects. I will not discuss it further here.

One-Factor Repeated Measures Design
Introduction

- In the *one-way repeated measures* design, n randomly sampled subjects are measured repeatedly on a occasions. This is sometimes called the "subjects by trials" design for that reason.
- In the previous module, we discussed the *randomized blocks* design, which is in fact a special case of a mixed-effects two-way ANOVA, with 1 observation per cell, or a units per block.
- I pointed out that, when the "block size" is a, the a observations can be either different individuals matched by blocking, or repeated measurements on the same observational unit or subject.
- The one-way repeated measures design that employs the "nonadditive model" and assumes that the blocking factor (i.e., "Subjects") is a random effect, is identical to the randomized blocks model with a fixed effect treatment and one observation per cell that we examined earlier.
- The "additive model" discussed by MWS in section 14.2 of RDASA3 seems unrealistic. It assumes that the effect over repeated measures is the same for all subjects. I will not discuss it further here.

One-Factor Repeated Measures Design

Introduction

- In the *one-way repeated measures* design, n randomly sampled subjects are measured repeatedly on a occasions. This is sometimes called the "subjects by trials" design for that reason.
- In the previous module, we discussed the *randomized blocks* design, which is in fact a special case of a mixed-effects two-way ANOVA, with 1 observation per cell, or a units per block.
- I pointed out that, when the "block size" is a, the a observations can be either different individuals matched by blocking, or repeated measurements on the same observational unit or subject.
- The one-way repeated measures design that employs the "nonadditive model" and assumes that the blocking factor (i.e., "Subjects") is a random effect, is identical to the randomized blocks model with a fixed effect treatment and one observation per cell that we examined earlier.
- The "additive model" discussed by MWS in section 14.2 of RDASA3 seems unrealistic. It assumes that the effect over repeated measures is the same for all subjects. I will not discuss it further here.

One-Factor Repeated Measures Design

Introduction

- In the *one-way repeated measures* design, n randomly sampled subjects are measured repeatedly on a occasions. This is sometimes called the "subjects by trials" design for that reason.
- In the previous module, we discussed the *randomized blocks* design, which is in fact a special case of a mixed-effects two-way ANOVA, with 1 observation per cell, or a units per block.
- I pointed out that, when the "block size" is a, the a observations can be either different individuals matched by blocking, or repeated measurements on the same observational unit or subject.
- The one-way repeated measures design that employs the "nonadditive model" and assumes that the blocking factor (i.e., "Subjects") is a random effect, is identical to the randomized blocks model with a fixed effect treatment and one observation per cell that we examined earlier.
- The "additive model" discussed by MWS in section 14.2 of RDASA3 seems unrealistic. It assumes that the effect over repeated measures is the same for all subjects. I will not discuss it further here.

One-Factor Repeated Measures Design

Introduction

- In the *one-way repeated measures* design, n randomly sampled subjects are measured repeatedly on a occasions. This is sometimes called the "subjects by trials" design for that reason.
- In the previous module, we discussed the *randomized blocks* design, which is in fact a special case of a mixed-effects two-way ANOVA, with 1 observation per cell, or a units per block.
- I pointed out that, when the "block size" is a, the a observations can be either different individuals matched by blocking, or repeated measurements on the same observational unit or subject.
- The one-way repeated measures design that employs the "nonadditive model" and assumes that the blocking factor (i.e., "Subjects") is a random effect, is identical to the randomized blocks model with a fixed effect treatment and one observation per cell that we examined earlier.
- The "additive model" discussed by MWS in section 14.2 of RDASA3 seems unrealistic. It assumes that the effect over repeated measures is the same for all subjects. I will not discuss it further here.

The Nonadditive Model for the $S \times A$ Design

- The additive model for the $S \times A$ design is essentially the same as for the randomized blocks RB-p design, which is in turn a mixed model ANOVA.
- The "Subjects" (S) factor is considered random, and the "Trials" factor (A) is fixed.
- You will note that the description of the model in RDASA3 Section 14.5.1 is incomplete in several respects.
- The Expected Mean Squares and F ratios for the non-additive model are given on the next slide. You will recognize them as equivalent to the values given for a 2-way mixed model randomized factorial design with one factor fixed and the other random.

The Nonadditive Model for the $S \times A$ Design

- The additive model for the $S \times A$ design is essentially the same as for the randomized blocks RB-p design, which is in turn a mixed model ANOVA.
- The "Subjects" (S) factor is considered random, and the "Trials" factor (A) is fixed.
- You will note that the description of the model in RDASA3 Section 14.5.1 is incomplete in several respects.
- The Expected Mean Squares and F ratios for the non-additive model are given on the next slide. You will recognize them as equivalent to the values given for a 2-way mixed model randomized factorial design with one factor fixed and the other random.

The Nonadditive Model for the $S \times A$ Design

- The additive model for the $S \times A$ design is essentially the same as for the randomized blocks RB-p design, which is in turn a mixed model ANOVA.
- The "Subjects" (S) factor is considered random, and the "Trials" factor (A) is fixed.
- You will note that the description of the model in RDASA3 Section 14.5.1 is incomplete in several respects.
- The Expected Mean Squares and F ratios for the non-additive model are given on the next slide. You will recognize them as equivalent to the values given for a 2-way mixed model randomized factorial design with one factor fixed and the other random.

The Nonadditive Model for the $S \times A$ Design

- The additive model for the $S \times A$ design is essentially the same as for the randomized blocks RB-p design, which is in turn a mixed model ANOVA.
- The "Subjects" (S) factor is considered random, and the "Trials" factor (A) is fixed.
- You will note that the description of the model in RDASA3 Section 14.5.1 is incomplete in several respects.
- The Expected Mean Squares and F ratios for the non-additive model are given on the next slide. You will recognize them as equivalent to the values given for a 2-way mixed model randomized factorial design with one factor fixed and the other random.

The Nonadditive Model for the $S \times A$ Design

Expected Mean Squares
1-Way Repeated Measures
Design

Effect	*A fixed* *S random*
A	$\sigma_e^2 + bn\theta_A^2 + n\sigma_{SA}^2$
S	$\sigma_e^2 + an\sigma_S^2$
SA	$\sigma_e^2 + n\sigma_{SA}^2$

- From the above, you can see that the F test for A is MS_A/MS_{SA}.
- However, when the null hypothesis of no S effect is true, $E(MS_S) = \sigma_e^2$, and no term in the table of Expected Mean Squares matches that.
- In this case, we cannot produce a test for an S effect in the standard manner.
- However, the notion that there is no difference between subjects is not taken seriously in most contexts, and so there is generally no interest in performing this test.
- In a similar vein, there is a Tukey Test for Nonadditivity. Using the non-additive model instead of the additive model costs some power, but since the model is the realistic one, I advocate using it.

The Nonadditive Model for the $S \times A$ Design

Expected Mean Squares
1-Way Repeated Measures Design

Effect	*A fixed* *S random*
A	$\sigma_e^2 + bn\theta_A^2 + n\sigma_{SA}^2$
S	$\sigma_e^2 + an\sigma_S^2$
SA	$\sigma_e^2 + n\sigma_{SA}^2$

- From the above, you can see that the F test for A is MS_A/MS_{SA}.
- However, when the null hypothesis of no S effect is true, $E(MS_S) = \sigma_e^2$, and no term in the table of Expected Mean Squares matches that.
- In this case, we cannot produce a test for an S effect in the standard manner.
- However, the notion that there is no difference between subjects is not taken seriously in most contexts, and so there is generally no interest in performing this test.
- In a similar vein, there is a Tukey Test for Nonadditivity. Using the non-additive model instead of the additive model costs some power, but since the model is the realistic one, I advocate using it.

The Nonadditive Model for the $S \times A$ Design

Expected Mean Squares
1-Way Repeated Measures
Design

Effect	*A fixed* *S random*
A	$\sigma_e^2 + bn\theta_A^2 + n\sigma_{SA}^2$
S	$\sigma_e^2 + an\sigma_S^2$
SA	$\sigma_e^2 + n\sigma_{SA}^2$

- From the above, you can see that the F test for A is MS_A/MS_{SA}.
- However, when the null hypothesis of no S effect is true, $E(MS_S) = \sigma_e^2$, and no term in the table of Expected Mean Squares matches that.
- In this case, we cannot produce a test for an S effect in the standard manner.
- However, the notion that there is no difference between subjects is not taken seriously in most contexts, and so there is generally no interest in performing this test.
- In a similar vein, there is a Tukey Test for Nonadditivity. Using the non-additive model instead of the additive model costs some power, but since the model is the realistic one, I advocate using it.

The Nonadditive Model for the $S \times A$ Design

Expected Mean Squares
1-Way Repeated Measures
Design

Effect	*A fixed* *S random*
A	$\sigma_e^2 + bn\theta_A^2 + n\sigma_{SA}^2$
S	$\sigma_e^2 + an\sigma_S^2$
SA	$\sigma_e^2 + n\sigma_{SA}^2$

- From the above, you can see that the F test for A is MS_A/MS_{SA}.
- However, when the null hypothesis of no S effect is true, $E(MS_S) = \sigma_e^2$, and no term in the table of Expected Mean Squares matches that.
- In this case, we cannot produce a test for an S effect in the standard manner.
- However, the notion that there is no difference between subjects is not taken seriously in most contexts, and so there is generally no interest in performing this test.
- In a similar vein, there is a Tukey Test for Nonadditivity. Using the non-additive model instead of the additive model costs some power, but since the model is the realistic one, I advocate using it.

The Nonadditive Model for the $S \times A$ Design

Expected Mean Squares
1-Way Repeated Measures
Design

Effect	*A fixed* *S random*
A	$\sigma_e^2 + bn\theta_A^2 + n\sigma_{SA}^2$
S	$\sigma_e^2 + an\sigma_S^2$
SA	$\sigma_e^2 + n\sigma_{SA}^2$

- From the above, you can see that the F test for A is MS_A/MS_{SA}.
- However, when the null hypothesis of no S effect is true, $E(MS_S) = \sigma_e^2$, and no term in the table of Expected Mean Squares matches that.
- In this case, we cannot produce a test for an S effect in the standard manner.
- However, the notion that there is no difference between subjects is not taken seriously in most contexts, and so there is generally no interest in performing this test.
- In a similar vein, there is a Tukey Test for Nonadditivity. Using the non-additive model instead of the additive model costs some power, but since the model is the realistic one, I advocate using it.

The Sphericity Assumption

- We recall from our earlier discussions that ANOVA involves an assumption of equal variances.
- In repeated measures designs, an additional assumption of *sphericity* is required.
- This assumption states that, if we consider all levels of the random effects variable, and compute all pairwise differences between these levels, they must have equal variance.
- Note that, if there are a levels of the treatment, then there are $\binom{a}{2} = a(a-1)/2$ differences to be calculated.

The Sphericity Assumption

- We recall from our earlier discussions that ANOVA involves an assumption of equal variances.
- In repeated measures designs, an additional assumption of *sphericity* is required.
- This assumption states that, if we consider all levels of the random effects variable, and compute all pairwise differences between these levels, they must have equal variance.
- Note that, if there are a levels of the treatment, then there are $\binom{a}{2} = a(a-1)/2$ differences to be calculated.

The Sphericity Assumption

- We recall from our earlier discussions that ANOVA involves an assumption of equal variances.
- In repeated measures designs, an additional assumption of *sphericity* is required.
- This assumption states that, if we consider all levels of the random effects variable, and compute all pairwise differences between these levels, they must have equal variance.
- Note that, if there are a levels of the treatment, then there are $\binom{a}{2} = a(a-1)/2$ differences to be calculated.

The Sphericity Assumption

- We recall from our earlier discussions that ANOVA involves an assumption of equal variances.
- In repeated measures designs, an additional assumption of *sphericity* is required.
- This assumption states that, if we consider all levels of the random effects variable, and compute all pairwise differences between these levels, they must have equal variance.
- Note that, if there are a levels of the treatment, then there are $\binom{a}{2} = a(a-1)/2$ differences to be calculated.

Effect of Violating Sphericity

- If sphericity is violated, the Type I Error rate for the F test of the A effect will be inflated above the nominal level.
- There is a test for sphericity, Mauchley's Test, that can be performed. However, it is sensitive to non-normality.
- If all population variances are identical for the levels of A, *and* all covariances are also equal, then the covariance matrix for A meets the assumption of *compound symmetry* and also meets the assumption of sphericity. However data can exhibit sphericity without exhibiting compound symmetry — compound symmetry is the stronger assumption.
- Of course, sphericity can never be violated if $a = 2$. Why not? (C.P.)

Effect of Violating Sphericity

- If sphericity is violated, the Type I Error rate for the F test of the A effect will be inflated above the nominal level.
- There is a test for sphericity, Mauchley's Test, that can be performed. However, it is sensitive to non-normality.
- If all population variances are identical for the levels of A, *and* all covariances are also equal, then the covariance matrix for A meets the assumption of *compound symmetry* and also meets the assumption of sphericity. However data can exhibit sphericity without exhibiting compound symmetry — compound symmetry is the stronger assumption.
- Of course, sphericity can never be violated if $a = 2$. Why not? (C.P.)

Effect of Violating Sphericity

- If sphericity is violated, the Type I Error rate for the F test of the A effect will be inflated above the nominal level.
- There is a test for sphericity, Mauchley's Test, that can be performed. However, it is sensitive to non-normality.
- If all population variances are identical for the levels of A, *and* all covariances are also equal, then the covariance matrix for A meets the assumption of *compound symmetry* and also meets the assumption of sphericity. However data can exhibit sphericity without exhibiting compound symmetry — compound symmetry is the stronger assumption.
- Of course, sphericity can never be violated if $a = 2$. Why not? (C.P.)

Effect of Violating Sphericity

- If sphericity is violated, the Type I Error rate for the F test of the A effect will be inflated above the nominal level.
- There is a test for sphericity, Mauchley's Test, that can be performed. However, it is sensitive to non-normality.
- If all population variances are identical for the levels of A, *and* all covariances are also equal, then the covariance matrix for A meets the assumption of *compound symmetry* and also meets the assumption of sphericity. However data can exhibit sphericity without exhibiting compound symmetry — compound symmetry is the stronger assumption.
- Of course, sphericity can never be violated if $a = 2$. Why not? (C.P.)

Dealing with Non-Sphericity

The Epsilon-Adjusted Test

- One can adjust the F statistic in the repeated measures design by multiplying both numerator and denominator degrees of freedom by a multiplier known as ϵ.
- In that case, $df_A = (a-1)\epsilon$ and $df_{SA} = (a-1)(n-1)\epsilon$.
- The multiplier ϵ is a number that varies from $1/(a-1)$ and 1. Thus, if maximum non-sphericity is observed, the degrees of freedom are adjusted to 1 and $(n-1)$ instead of $(a-1)$ and $(a-1)(n-1)$.
- Unfortunately we can only estimate ϵ. Two estimates are in common use. One is traceable back to Box(1954), but is commonly referred to as the Greenhouse-Geisser $\hat{\epsilon}$. The other estimator, the Huynh-Feldt $\tilde{\epsilon}$, tends to result in slightly more rejections.

Dealing with Non-Sphericity

The Epsilon-Adjusted Test

- One can adjust the F statistic in the repeated measures design by multiplying both numerator and denominator degrees of freedom by a multiplier known as ϵ.
- In that case, $df_A = (a-1)\epsilon$ and $df_{SA} = (a-1)(n-1)\epsilon$.
- The multiplier ϵ is a number that varies from $1/(a-1)$ and 1. Thus, if maximum non-sphericity is observed, the degrees of freedom are adjusted to 1 and $(n-1)$ instead of $(a-1)$ and $(a-1)(n-1)$.
- Unfortunately we can only estimate ϵ. Two estimates are in common use. One is traceable back to Box(1954), but is commonly referred to as the Greenhouse-Geisser $\hat{\epsilon}$. The other estimator, the Huynh-Feldt $\tilde{\epsilon}$, tends to result in slightly more rejections.

Dealing with Non-Sphericity
The Epsilon-Adjusted Test

- One can adjust the F statistic in the repeated measures design by multiplying both numerator and denominator degrees of freedom by a multiplier known as ϵ.
- In that case, $df_A = (a-1)\epsilon$ and $df_{SA} = (a-1)(n-1)\epsilon$.
- The multiplier ϵ is a number that varies from $1/(a-1)$ and 1. Thus, if maximum non-sphericity is observed, the degrees of freedom are adjusted to 1 and $(n-1)$ instead of $(a-1)$ and $(a-1)(n-1)$.
- Unfortunately we can only estimate ϵ. Two estimates are in common use. One is traceable back to Box(1954), but is commonly referred to as the Greenhouse-Geisser $\hat{\epsilon}$. The other estimator, the Huynh-Feldt $\tilde{\epsilon}$, tends to result in slightly more rejections.

Dealing with Non-Sphericity

The Epsilon-Adjusted Test

- One can adjust the F statistic in the repeated measures design by multiplying both numerator and denominator degrees of freedom by a multiplier known as ϵ.
- In that case, $df_A = (a-1)\epsilon$ and $df_{SA} = (a-1)(n-1)\epsilon$.
- The multiplier ϵ is a number that varies from $1/(a-1)$ and 1. Thus, if maximum non-sphericity is observed, the degrees of freedom are adjusted to 1 and $(n-1)$ instead of $(a-1)$ and $(a-1)(n-1)$.
- Unfortunately we can only estimate ϵ. Two estimates are in common use. One is traceable back to Box(1954), but is commonly referred to as the Greenhouse-Geisser $\hat{\epsilon}$. The other estimator, the Huynh-Feldt $\tilde{\epsilon}$, tends to result in slightly more rejections.

A Repeated Measures Analysis of Seasonal Depression

- Table 14.1 in RDASA3 presents Beck Depression scores for 14 subjects over 4 seasons (Winter, Spring, Summer, Fall).
- A full SPSS analysis of these data is shown in Table 14.6. Here, we reproduce this analysis in R.
- We could perform this analysis "from scratch," using the standard `lm()` function, after writing our own functions to perform the Mauchley test and calculate the Huynh-Feldt and Greenhouse-Geisser estimates. However, these tasks have been nicely automated for us in an R package `ez`.
- This package still has some bugs in it, and I am trying to determine how far we can trust it.

A Repeated Measures Analysis of Seasonal Depression

- To begin with, we load the data

```
> Table1401 <- read.csv("Table1401.csv")
> Table1401
   Subject Winter Spring Summer   Fall
1        1  7.500 11.554  1.000  1.208
2        2  7.000  9.000  5.000 15.000
3        3  1.000  1.000  0.000  0.000
4        4  0.000  0.000  0.000  0.000
5        5  1.059  0.000  1.097  4.000
6        6  1.000  2.500  0.000  2.000
7        7  2.500  0.000  0.000  2.000
8        8  4.500  1.060  2.000  2.000
9        9  5.000  2.000  3.000  5.000
10      10  2.000  3.000  4.208  3.000
11      11  7.000  7.354  5.877  9.000
12      12  2.500  2.000  0.009  2.000
13      13 11.000 16.000 13.000 13.000
14      14  8.000 10.500  1.000 11.000
```

A Repeated Measures Analysis of Seasonal Depression

Reshaping the Data

- This format seems natural for the data — one line for each subject.
- Unfortunately, this is not the format typically required by analysis of variance routines, because in this case, the subject is being treated as a level of the Subject (S) factor in the design, and each column is a repeated measure that is treated as a level of the A factor.
- In effect, then, we have a $S \times A$, 14×4 ANOVA with 1 observation per cell, and the data need to be recast that way. How?
- Well, we could do it by hand, or we could write a brief program in R to do it. This is a skill we'll need to master, especially if we do repeated measures ANOVAs often.
- There is another way. We can use a library, `reshape`, designed especially for this task.

A Repeated Measures Analysis of Seasonal Depression
Reshaping the Data

- We begin by loading in the `reshape` library.

```
> library(reshape)
```

- Next we `melt` the data to get it into the proper form, in which *Subject* and *Season* are treated as factors with 14 and 4 levels, respectively.
- R cannot guess what names we want to call the variables, so we change the column names.

```
> rm.data <- melt(Table1401,id=c("Subject"),measured=c("Depression"))
> colnames(rm.data)[2:3] <- c("Season","Depression")
> head(rm.data)

  Subject Season Depression
1       1 Winter      7.500
2       2 Winter      7.000
3       3 Winter      1.000
4       4 Winter      0.000
5       5 Winter      1.059
6       6 Winter      1.000

> tail(rm.data)

   Subject Season Depression
51       9   Fall          5
52      10   Fall          3
53      11   Fall          9
54      12   Fall          2
55      13   Fall         13
56      14   Fall         11
```

Analyzing with ezANOVA

- Next, we load the `ez` library, which is specially designed to handle repeated measures designs.
- We call the function `ezANOVA`, which requires specification of the dependent variable, subject ID variable, within-subject factors, and between-subject factors (factors that have different groups of subjects at different levels.
- There are no between-subject variables in this design.
- The results are on the next slide – with just a few exceptions, these match the results produced by SPSS.
- Since the upper bound for ϵ is 1.0, SPSS truncates the Huynh-Felt estimator at 1.00, while `ezANOVA` reports it at 1.04, but apparently truncates it internally, as the p-value is consistent with an ϵ of 1.0.

Analyzing with ezANOVA

```
> library(ez)
> results <-ezANOVA(data=rm.data,dv=.(Depression),
+ wid=.(Subject),within=.(Season),type="III",return_aov=TRUE)
> results

$ANOVA
  Effect DFn DFd       F          p p<.05        ges
2 Season   3  39 3.00115 0.04202195     * 0.04621462

$`Mauchly's Test for Sphericity`
  Effect        W         p p<.05
2 Season 0.648444 0.4078981

$`Sphericity Corrections`
  Effect       GGe      p[GG] p[GG]<.05      HFe      p[HF] p[HF]<.05
2 Season 0.8316087 0.05321143           1.044805 0.04202195         *

$aov

Call:
aov(formula = formula(aov_formula), data = data)

Grand Mean: 4.132607

Stratum 1: Subject

Terms:
                Residuals
Sum of Squares   779.0988
Deg. of Freedom        13

Residual standard error: 7.741491

Stratum 2: Subject:Season

Terms:
                   Season Residuals
Sum of Squares   47.77842 206.96047
Deg. of Freedom         3        39

Residual standard error: 2.303623
Estimated effects may be unbalanced
```

Analyzing with ezANOVA

Table 14.6 SPSS output for the data of Table 14.1: Mauchly's test of sphericity (*a*), ANOVA (*b*), and MANOVA (*c*)

(*a*) Mauchly's test of sphericity

Within-subjects effect	Mauchly's *W*	Approx. chi-square	*df*	Sig.	Epsilon[a] Greenhouse–Geisser	Huynh–Feldt	Lower bound
Season	.648	15.078	5	.408	.832	1.000	.333

[a] May be used to adjust the degrees of freedom for the averaged tests of significance. Corrected tests are displayed in the tests of within-subjects effects table.

(*b*) Tests of within-subjects effects

Source		Type III sum of squares	*df*	Mean square	*F*	Sig.
Seasons	Sphericity assumed	47.778	3	15.926	3.001	.042
	Greenhouse–Geisser	47.778	2.495	19.151	3.001	.053
	Huynh–Feldt	47.778	3.000	15.926	3.001	.042
	Lower bound	47.778	1.000	47.778	3.001	.107
Error (seasons)	Sphericity assumed	206.960	39	5.307		
	Greenhouse–Geisser	206.960	32.433	6.381		
	Huynh–Feldt	206.960	39.000	5.307		
	Lower bound	206.960	13.000	15.920		

(*c*) Multivariate tests

Effect		Value	*F*	Hypothesis *df*	Error *df*	Sig.
Seasons	Pillai's trace	.392	2.365[a]	3.000	11.000	.127
	Wilks' lambda	.608	2.365[a]	3.000	11.000	.127
	Hotelling's trace	.645	2.365[a]	3.000	11.000	.127
	Roy's largest root	.645	2.365[a]	3.000	11.000	.127

[a] Exact statistic.

A Multivariate Approach

- Strictly speaking, repeated measures data are *multivariate*, in that each subject (or unit of observation) produces several responses.
- The ANOVA approach to repeated measures requires the strong assumption of *sphericity* in order to reduce a multivariate problem to a univariate problem.
- An alternate approach is to use a truly multivariate approach, that allows the data to have any population covariance structure.
- This approach is actually a generalization of the univariate t test called *Hotelling's* T^2.

A Multivariate Approach

- Strictly speaking, repeated measures data are *multivariate*, in that each subject (or unit of observation) produces several responses.
- The ANOVA approach to repeated measures requires the strong assumption of *sphericity* in order to reduce a multivariate problem to a univariate problem.
- An alternate approach is to use a truly multivariate approach, that allows the data to have any population covariance structure.
- This approach is actually a generalization of the univariate t test called *Hotelling's* T^2.

A Multivariate Approach

- Strictly speaking, repeated measures data are *multivariate*, in that each subject (or unit of observation) produces several responses.
- The ANOVA approach to repeated measures requires the strong assumption of *sphericity* in order to reduce a multivariate problem to a univariate problem.
- An alternate approach is to use a truly multivariate approach, that allows the data to have any population covariance structure.
- This approach is actually a generalization of the univariate t test called *Hotelling's* T^2.

A Multivariate Approach

- Strictly speaking, repeated measures data are *multivariate*, in that each subject (or unit of observation) produces several responses.
- The ANOVA approach to repeated measures requires the strong assumption of *sphericity* in order to reduce a multivariate problem to a univariate problem.
- An alternate approach is to use a truly multivariate approach, that allows the data to have any population covariance structure.
- This approach is actually a generalization of the univariate t test called *Hotelling's* T^2.

Hotelling's T^2

- The ordinary one-sample t statistic tests whether a single mean μ is equal to some hypothesized value μ_0.
- The one-sample Hotelling's T^2 statistic tests whether, simultaneously, a *vector* (list) of means $\boldsymbol{\mu}$ is equal to a vector of hypothesized values $\boldsymbol{\mu_0}$.
- So, suppose we have 4 measured variables, as in the Depression data, and we wish to test whether the 4 repeated measures all have equal means.
- One way of testing this is to test whether the three pairwise differences between the means all have means of zero, using the one-sample Hotelling's test.
- This can be done in a couple of seconds, using R. Just load in the `ICSNP` library, and use the `HotellingsT2` function.

Hotelling's T^2

- First, we create the difference scores, then we pass them to the Hotelling T^2 function for analysis.
- We obtain an F statistic of 2.36, which is not significant.

```
> library(ICSNP)
> X <- with(Table1401,data.frame(cbind(
+ Winter-Spring, Spring-Summer,Summer-Fall)))
> HotellingsT2(X)

        Hotelling's one sample T2-test

data:  X
T.2 = 2.3654, df1 = 3, df2 = 11, p-value = 0.1268
alternative hypothesis: true location is not equal to c(0,0,0)
```