

Description	Quick start
Menu	Syntax
Options for multiple-sample tests	Options with one-sample tests
Remarks and examples	Stored results
Methods and formulas	References
Also see	

Description

`mvtest means` performs one-sample and multiple-sample multivariate tests on means. These tests assume multivariate normality.

See [\[MV\] mvtest](#) for other multivariate tests.

Quick start

Test that the means of v_1 , v_2 , and v_3 are equal

```
mvtest means v1 v2 v3
```

Test that the means of v_1 , v_2 , and v_3 are zero

```
mvtest means v1 v2 v3, zero
```

Test that the means vector equals the hypothesized vector `myvec`

```
mvtest means v1 v2 v3, equals(myvec)
```

Test that the means of v_1 , v_2 , and v_3 are the same across groups defined by `catvar` assuming homogeneous covariance matrices

```
mvtest means v1 v2 v3, by(catvar)
```

Same as above, but allow heterogeneous covariance matrices

```
mvtest means v1 v2 v3, by(catvar) heterogeneous
```

Same as above, but use a likelihood-ratio test instead of a Wald test

```
mvtest means v1 v2 v3, by(catvar) lr
```

Menu

Statistics > Multivariate analysis > MANOVA, multivariate regression, and related > Multivariate test of means, covariances, and normality

Syntax

Multiple-sample tests

```
mvtest means varlist [if] [in] [weight], by(groupvars) [multisample_options]
```

One-sample tests

```
mvtest means varlist [if] [in] [weight] [, one-sample_options]
```

<i>multisample_options</i>	Description
Model	
* <i>by</i> (<i>groupvars</i>)	compare subsamples with same values in <i>groupvars</i>
<u>missing</u>	treat missing values in <i>groupvars</i> as ordinary values
Options	
<u>homogeneous</u>	test for equal means with homogeneous covariance matrices across by-groups; the default
<u>heterogeneous</u>	James's test for equal means, allowing heterogeneous covariance matrices across by-groups
<u>lr</u>	likelihood-ratio test for equal means, allowing heterogeneous covariance matrices across by-groups
<u>protect</u> (<i>spec</i>)	run protection as a safeguard against local minimum with the group means as initial values; use only with <code>lr</code> option

* *by*(*groupvars*) is required.

<i>one-sample_options</i>	Description
Options	
<u>equal</u>	test that variables in <i>varlist</i> have equal means; the default
<u>zero</u>	test that means of <i>varlist</i> are all equal to 0
<u>equals</u> (<i>M</i>)	test that mean vector equals vector <i>M</i>
<u>linear</u> (<i>V</i>)	test that mean vector of <i>varlist</i> satisfies linear hypothesis described by matrix <i>V</i>

`bootstrap`, `by`, `collect`, `jackknife`, `rolling`, and `statsby` are allowed; see [U] 11.1.10 Prefix commands.

Weights are not allowed with the `bootstrap` prefix; see [R] `bootstrap`.

`aweight`s are not allowed with the `jackknife` prefix; see [R] `jackknife`.

`aweight`s and `fweight`s are allowed; see [U] 11.1.6 `weight`.

Options for multiple-sample tests

Model

by(*groupvars*) is required with the multiple-sample version of the test. Observations with the same values in *groupvars* form a sample. Observations with missing values in *groupvars* are ignored, unless the `missing` option is specified.

`missing` specifies that missing values in *groupvars* are treated like ordinary values.

Options

`homogeneous`, the default, specifies the hypothesis that the mean vectors are the same across the by-groups, assuming homogeneous covariance matrices across the by-groups. `homogeneous` produces the four standard tests of multivariate means (Wilks's lambda, Pillai's trace, Lawley–Hotelling trace, and Roy's largest root).

`heterogeneous` removes the assumption that the covariance matrices are the same across the by-groups. This is the multivariate Behrens–Fisher problem. With two groups, the MNV test, an affine-invariant modification by [Krishnamoorthy and Yu \(2004\)](#) of the Nel–Van der Merwe (1986) test, is displayed. With more than two groups, the Wald test, with p -values adjusted as proposed by [James \(1954\)](#), is displayed.

`lr` removes the assumption that the covariance matrices are the same across the by-groups and specifies that a likelihood-ratio test be presented. The associated estimation problem may have multiple local optima, though this seems rare with two groups.

`protect(spec)` is a technical option accompanying `lr`, specifying that the “common means” model is fit from different starting values to ascertain with some confidence whether a global optimum to the underlying estimation problem was reached. The initialization of the common means proposed in [Mardia, Kent, and Bibby \(1979\)](#) is always used as well. If the different trials do not converge to the same solution, the “best” one is used to obtain the test, and a warning message is displayed.

`protect(groups)` specifies to fit the common means model using each of the group means as starting values for the common means.

`protect(randobs, reps(#))` specifies to fit the common means model using $\#$ random observations as starting values for the common means.

`protect(#)` is a convenient shorthand for `protect(randobs, reps(#))`.

Options with one-sample tests

Options

`equal` performs Hotelling's test of the hypothesis that the means of all variables in `varlist` are equal.

`zero` performs Hotelling's test of the hypothesis that the means of all variables in `varlist` are 0.

`equals(M)` performs Hotelling's test that the vector of means of the k variables in `varlist` equals M . The matrix M must be a $k \times 1$ or $1 \times k$ vector. The row and column names of M are ignored.

`linear(V)` performs Hotelling's test that the means satisfy a user-specified set of linear constraints, represented by V . V must be a matrix vector with k or $k + 1$ columns, where k is the number of variables in `varlist`. Let A be a matrix of the first k columns of V . Let b be the last column of V if V has $k + 1$ columns and a column of 0s otherwise. The linear hypothesis test is that A times a column vector of the means of `varlist` equals b . `mvtest` ignores matrix row and column names.

Remarks and examples

Remarks are presented under the following headings:

One-sample tests for mean vectors

Multiple-sample tests for mean vectors

One-sample tests for mean vectors

One-sample and multiple-sample tests of means are available with the `mvtest means` command. One-sample tests include tests that the means of *varlist* are equal, the means of *varlist* equal a given vector, the means of *varlist* are zero, and linear combinations of the means of *varlist* equal a given vector.

We first explore the use of `mvtest means` for testing the one-sample hypothesis that the means of *varlist* are equal.

▷ Example 1

The cost on a per-mile basis of 25 gasoline trucks used for transporting milk are provided in three categories: fuel, repair, and capital (Johnson and Wichern 2007, 269).

```
. use https://www.stata-press.com/data/r19/milktruck
(Milk transportation costs for 25 gasoline trucks (Johnson and Wichern 2007))
```

```
. summarize
```

Variable	Obs	Mean	Std. dev.	Min	Max
fuel	25	12.56	5.382	4.24	29.11
repair	25	8.1612	4.631723	1.35	17.44
capital	25	10.5444	3.687688	3.28	17.59

Are the means of the three costs equal? The `equal` option of `mvtest means` provides a way of testing this hypothesis.

```
. mvtest means fuel repair capital, equal
```

```
Test that all means are the same
```

```
Hotelling T2 = 35.25
Hotelling F(2,23) = 16.89
Prob > F = 0.0000
```

We reject the null hypothesis of equal means for fuel, repair, and capital costs.

◀

Hotelling's T^2 statistic is a multivariate generalization of the univariate t statistic; see [R] `ttest`. A test of the bivariate hypothesis that the means of the repair and capital costs are equal could be obtained with

```
ttest repair == capital
```

or with

```
mvtest means repair capital, equal
```

The square of the t statistic from `ttest` equals the T^2 value from `mvtest means`. With `ttest`, you are limited to comparing the means of two variables; with `mvtest means`, you can simultaneously compare the means of two or more variables.

The `equals()` option of `mvtest means` provides Hotelling's T^2 statistic for the test that the mean vector for *varlist* equals a given vector. This provides a multivariate generalization of the univariate t statistic obtained using the `ttest varname == #` syntax of [R] `ttest`.

▷ Example 2

We compare the measurements of the available and exchangeable soil calcium (y_1 and y_2) and turnip-green calcium (y_3) at 10 locations in the South (Rencher and Christensen 2012, 66; Kramer and Jensen 1969) to the values 15.0, 6.0, and 2.85 respectively (Rencher and Christensen 2012, 133–134).

```
. use https://www.stata-press.com/data/r19/turnip
(Calcium in soil and turnip greens, Rencher and Christensen (2012))
. summarize y*

```

Variable	Obs	Mean	Std. dev.	Min	Max
y1	10	28.1	11.85514	6	40
y2	10	7.18	8.499908	1.6	30
y3	10	3.089	.5001211	2.7	4.38

```
. matrix Mstd = (15.0, 6.0, 2.85)
. mvtest means y* , equals(Mstd)
Test that means equal vector Mstd
    Hotelling T2 =    24.56
    Hotelling F(3,7) =    6.37
    Prob > F =    0.0207
```

The calcium measurements from these 10 locations in the South do not appear to match the hypothesized values.



The zero option of `mvtest means` tests the hypothesis that the means of *varlist* are zero. The same result could be obtained by creating a column or row vector of the appropriate length filled with zeros and supplying that to the `equals()` option.

```
mvtest means y1 y2 y3, zero
```

would give the same test as

```
matrix Zero = 0,0,0
mvtest means y1 y2 y3, equals(Zero)
```

This same test against a zero-mean vector can be obtained with the `hotelling` command; see [\[MV\] hotelling](#). For example,

```
hotelling y1 y2 y3
```

`mvtest means` also tests that linear combinations of the means of *varlist* equal a given vector.

▶ Example 3

The `linear()` option of `mvtest means` can be used to obtain the same result as in [example 1](#), testing that the fuel, repair, and capital costs are equal. We do this by constructing two appropriate linear combinations of our three variables and testing that the means of these two linear combinations are zero.

```
. use https://www.stata-press.com/data/r19/milktruck
(Milk transportation costs for 25 gasoline trucks (Johnson and Wichern 2007))
. matrix C = 1, -1, 0 \ 0, 1, -1
. matrix list C
C[2,3]
    c1  c2  c3
r1   1  -1   0
r2   0   1  -1
. mvtest means fuel repair capital, linear(C)
Test that mean vector satisfies linear hypothesis C
    Hotelling T2 =    35.25
    Hotelling F(2,23) =    16.89
    Prob > F =    0.0000
```

We formed a matrix `C` that contrasted fuel to repair (the first row of `C`) and repair to capital (the second row of `C`). Note that we need not set the matrix row and column names of `C`. By default, the linear contrast was tested equal to a vector of zeros.

We could explicitly append an extra column of zeros in our matrix and obtain the same result.

```
. matrix Czero = C, (0 \ 0)
. matrix list Czero
Czero[2,4]
   c1  c2  c3  c4
r1   1  -1   0   0
r2   0   1  -1   0
. mvtest means fuel repair capital, linear(Czero)
Test that mean vector satisfies linear hypothesis Czero
      Hotelling T2 =      35.25
      Hotelling F(2,23) =     16.89
      Prob > F =      0.0000
```

Values other than zeros could have been appended to `C` to test if the linear combinations equal those other values.

◀

[Rencher and Christensen \(2012, 152–154\)](#) discusses one-sample profile analysis. The `linear()` option of `mvtest means` allows you to do this and other one-sample comparisons of interest.

Multiple-sample tests for mean vectors

Multiple-sample tests of mean vectors are also supported by `mvtest means`. The groups defining the multiple samples are specified with the `by()` option. The test that is presented depends on whether homogeneity of variance is assumed and whether there are more than two groups. The `homogeneity` option, the default, provides four standard multivariate tests (Wilks's lambda, Pillai's trace, Lawley–Hotelling trace, and Roy's largest root) under the assumption that the group covariance matrices are equal. The remaining possibilities do not assume equal covariances for the groups. The `heterogeneous` option with two `by-groups` presents the affine-invariant modification by Krishnamoorthy and Yu (2004) of the test proposed by [Nel and Van der Merwe \(1986\)](#). The `heterogeneous` option with more than two `by-groups` presents a Wald test along with James's approximation to the p -value of the test. The `lr` option also removes the assumption of equal covariance matrices for the groups and produces a likelihood-ratio test for the equality of the group means.

▷ Example 4

In [example 2](#) of [\[MV\] manova](#), we introduce two variables measured on four groups of rabbits. The groups have differing sample sizes. `mvtest means` with the `by()` option can test the hypothesis that the means of the two variables are the same for the four groups of rabbits under the assumption that the groups have equal covariance matrices.

```
. use https://www.stata-press.com/data/r19/metabolic
(Table 4.5 Metabolic comparisons of rabbits -- Rencher (1998))
. mvtest means y1 y2, by(group)
```

Test for equality of 4 group means, assuming homogeneity

	Statistic	F(df1,	df2)	= F	Prob>F
Wilks' lambda	0.1596	6.0	32.0	8.02	0.0000 e
Pillai's trace	1.2004	6.0	34.0	8.51	0.0000 a
Lawley-Hotelling trace	3.0096	6.0	30.0	7.52	0.0001 a
Roy's largest root	1.5986	3.0	17.0	9.06	0.0008 u

e = exact, a = approximate, u = upper bound on F

We reject the null hypothesis and conclude that the means are likely different between the four groups of rabbits.

The statistics reported above are the same as reported by `manova y1 y2 = group` in [example 2](#) of [\[MV\] manova](#). `mvtest means y1 y2, by(group)` homogeneous would also have produced the same results because homogeneous is the default when `by()` is specified.

◀

▶ Example 5

Continuing with the rabbit data, restricting ourselves to examining only the first two groups of rabbits and continuing to assume equal covariance matrices, we obtain the following:

```
. mvtest means y1 y2 if group < 3, by(group)
```

Test for equality of 2 group means, assuming homogeneity

	Statistic	F(df1,	df2)	= F	Prob>F
Wilks' lambda	0.3536	2.0	11.0	10.05	0.0033 e
Pillai's trace	0.6464	2.0	11.0	10.05	0.0033 e
Lawley-Hotelling trace	1.8279	2.0	11.0	10.05	0.0033 e
Roy's largest root	1.8279	2.0	11.0	10.05	0.0033 e

e = exact, a = approximate, u = upper bound on F

We reject the null hypothesis of equal means for the two groups of rabbits.

With only two groups, the four multivariate tests above are equivalent. Because there were only two groups, we could have also produced this same F test with `hotelling y1 y2 if group < 3, by(group)`; see [\[MV\] hotelling](#).

◀

▶ Example 6

We now remove the assumption of equal covariance matrices for these two groups of rabbits and see if our conclusions change.

```
. mvtest means y1 y2 if group < 3, by(group) heterogeneous
```

Test for equality of 2 group means, allowing for heterogeneity

```
MNV F(2,9.5) = 9.92
Prob > F = 0.0047
```

Removing the assumption of equal covariance matrices still leads to rejection of the null hypothesis that the means for the two groups of rabbits are equal.

Because there were only two groups, an F statistic based on an affine-invariant modification by [Krishnamoorthy and Yu \(2004\)](#) of the test proposed by [Nel and Van der Merwe \(1986\)](#) was presented. ◀

▷ Example 7

If we attempt to test all four groups of rabbits while removing the assumption of equal covariance matrices,

```
. mvtest means y1 y2, by(group) heterogeneous
```

we receive an error message indicating that we have a singular covariance matrix. This is because there are only two observations for the fourth group of rabbits.

If we omit the fourth group, we obtain

```
. mvtest means y1 y2 if group < 4, by(group) heterogeneous
Test for equality of 3 group means, allowing for heterogeneity
Wald chi2(4) =      34.08
Prob > chi2 =    0.0000 (chi-squared approximation)
Prob > chi2 =    0.0017 (James' approximation)
```

With more than two groups, a Wald χ^2 statistic is presented along with two p -values. The first p -value is the approximate one based on the χ^2 distribution with four degrees of freedom. The second, more accurate, p -value is based on James's (1954) approximation. Both p -values lead to rejection of the null hypothesis; the three groups of rabbits are unlikely to have equal means. ◀

▷ Example 8

We can request a likelihood-ratio test instead of a Wald test by specifying the `lr` option. Like `heterogeneous`, the `lr` option does not assume that the group covariance matrices are equal.

```
. mvtest means y1 y2 if group < 4, by(group) lr
Test for equality of 3 group means, allowing for heterogeneity
LR chi2(4) =      21.32
Prob > chi2 =    0.0003
```

The likelihood-ratio test also leads us to reject the null hypothesis of equal means for the three groups of rabbits. ◀

The computation of the likelihood-ratio test requires fitting the multivariate normal distribution with common means while allowing for different covariance matrices. The iterative fitting process may converge to local solutions rather than to the global solution, invalidating the reported test ([Buot, Hoşten, and Richards 2007](#)). As a precaution, you may use the `protect()` option to request fitting from different starting values so that the test results are based on the best solution found, and you can check whether the likelihood surface has multiple optima.

Stored results

`mvtest` means without the `by()` option (that is, a one-sample means test) stores the following in `r()`:

Scalars

<code>r(T2)</code>	Hotelling T^2
<code>r(F)</code>	F statistic
<code>r(df_m)</code>	model degrees of freedom
<code>r(df_r)</code>	residual degrees of freedom
<code>r(p_F)</code>	p -value for F test

Macros

<code>r(Ftype)</code>	type of model F test
-----------------------	------------------------

`mvtest` means with `by()` but without the `1r` or `heterogeneous` options (that is, a multiple-sample means test, assuming homogeneity) stores the following in `r()`:

Scalars

<code>r(F)</code>	F statistic
<code>r(df_m)</code>	model degrees of freedom
<code>r(df_r)</code>	residual degrees of freedom
<code>r(p_F)</code>	p -value for F test

Macros

<code>r(Ftype)</code>	type of model F test
-----------------------	------------------------

Matrices

<code>r(stat_m)</code>	MANOVA model tests
------------------------	--------------------

`mvtest` means with `by()` defining two groups and with the `heterogeneous` option (that is, a two-sample test of means, allowing for heterogeneity) stores the following in `r()`:

Scalars

<code>r(F)</code>	F statistic
<code>r(df_m)</code>	model degrees of freedom
<code>r(df_r)</code>	residual degrees of freedom
<code>r(p_F)</code>	p -value for F test

Macros

<code>r(Ftype)</code>	type of model F test
-----------------------	------------------------

`mvtest` means with `by()` defining more than two groups and with the `heterogeneous` option (that is, a multiple-sample test of means, allowing for heterogeneity) stores the following in `r()`:

Scalars

<code>r(chi2)</code>	χ^2 statistic
<code>r(df)</code>	degrees of freedom for χ^2 test
<code>r(p_chi2)</code>	p -value for χ^2 test
<code>r(p_chi2_James)</code>	p -value for Wald test via James's approximation

Macros

<code>r(chi2type)</code>	type of model χ^2 test
--------------------------	-----------------------------

mvtest means with the `by()` and `lr` options (that is, a likelihood-ratio multiple-sample test of means, allowing for heterogeneity) stores the following in `r()`:

Scalars	
<code>r(chi2)</code>	χ^2 statistic
<code>r(df)</code>	degrees of freedom for χ^2 test
<code>r(p_chi2)</code>	p -value for χ^2 test
<code>r(rc)</code>	return code
<code>r(uniq)</code>	1/0 if protection runs yielded/did not yield same solution (<code>protect()</code> only)
<code>r(nprotect)</code>	number of protection runs (<code>protect()</code> only)
Macros	
<code>r(chi2type)</code>	type of model χ^2 test
Matrices	
<code>r(M)</code>	maximum likelihood estimate of means

Methods and formulas

When comparing the formulas in this section with those found in some multivariate texts, be aware of whether they define the sample covariance matrix with a divisor of N or $N - 1$. We use N . The formulas for several of the statistics are presented differently depending on your choice of divisor (but are still equivalent).

Methods and formulas are presented under the following headings:

One-sample tests for mean vectors
Multiple-sample tests for mean vectors

One-sample tests for mean vectors

Let the sample consist of N i.i.d. observations, $\mathbf{x}_i, i = 1, \dots, N$, from a k -variate multivariate normal distribution, $MVN_k(\boldsymbol{\mu}, \boldsymbol{\Sigma})$, with sample mean $\bar{\mathbf{x}} = 1/N \sum_{i=1}^N \mathbf{x}_i$ and sample covariance matrix $\mathbf{S} = 1/N \sum_{i=1}^N (\mathbf{x}_i - \bar{\mathbf{x}})(\mathbf{x}_i - \bar{\mathbf{x}})'$.

mvtest means with the `equals()` option tests that a mean vector equals a fixed vector, $H_0: \boldsymbol{\mu} = \boldsymbol{\mu}_0$, and produces a Hotelling T^2 statistic, which is equivalent to the likelihood-ratio test (Mardia, Kent, and Taylor 2024, 129–130)

$$T_{\text{omf}}^2 = (N - 1)(\bar{\mathbf{x}} - \boldsymbol{\mu}_0)' \mathbf{S}^{-1} (\bar{\mathbf{x}} - \boldsymbol{\mu}_0)$$

Under the null hypothesis,

$$F_{\text{omf}} = \frac{N - k}{(N - 1)k} T_{\text{omf}}^2$$

is distributed $F(k, N - k)$.

mvtest means with the `zero` option tests that the mean vector equals $\mathbf{0}$ and is obtained from T_{omf}^2 by setting $\boldsymbol{\mu}_0 = \mathbf{0}$. For this case, denote Hotelling's T^2 as T_{omz}^2 and the corresponding F statistic as F_{omz} .

mvtest means with the `linear()` option tests that the mean vector $\boldsymbol{\mu}$ satisfies a linear hypothesis, $H_0: \mathbf{C}\boldsymbol{\mu} = \mathbf{b}$, and produces a Hotelling T^2 test, which is equivalent to the likelihood-ratio test (Mardia, Kent, and Taylor 2024, 135–136)

$$T_{\text{omc}}^2 = (N - 1)(\mathbf{C}\bar{\mathbf{x}} - \mathbf{b})' (\mathbf{C}\mathbf{S}\mathbf{C}')^{-1} (\mathbf{C}\bar{\mathbf{x}} - \mathbf{b})$$

Under the null hypothesis,

$$F_{\text{omc}} = \frac{N - q}{(N - 1)q} T_{\text{omc}}^2$$

is distributed $F(q, N - q)$, where q is the rank of \mathbf{CSC}' , typically the number of rows of \mathbf{C} .

`mvtest means` with the `equal` option tests that all means are equal and is obtained from T_{omc}^2 by setting $\mathbf{C} = (\mathbf{I}, -\mathbf{1})$ and $\mathbf{b} = \mathbf{0}$. For this case, denote Hotelling's T^2 as T_{omc}^2 and the corresponding F statistic as F_{omc} .

Multiple-sample tests for mean vectors

Let there be $m \geq 2$ independent samples with the j th sample containing N_j i.i.d. observations, \mathbf{x}_{ji} , $i = 1, \dots, N_j$, from a k -variate multivariate normal distribution, $\text{MVN}_k(\boldsymbol{\mu}_j, \mathbf{\Sigma}_j)$. The observed j th sample mean is $\bar{\mathbf{x}}_j = 1/N_j \sum_{i=1}^{N_j} \mathbf{x}_{ji}$ and covariance is $\mathbf{S}_j = 1/N_j \sum_{i=1}^{N_j} (\mathbf{x}_{ji} - \bar{\mathbf{x}}_j)(\mathbf{x}_{ji} - \bar{\mathbf{x}}_j)'$. Let $N = \sum_{j=1}^m N_j$.

The tests for the hypothesis that the mean vector is the same across m independent samples, obtained from `mvtest means` with the `by()` option, come in four different flavors, depending on whether the additional assumption is made that the covariance matrix \mathbf{S}_j is the same over the m samples (the homogeneous and heterogeneous options) and on the number of samples (whether the `by()` option defines more than two groups). If equal covariance is not assumed, the problem is commonly referred to as the multivariate Behrens–Fisher problem.

When assuming equal covariance matrices for the `by()` groups (the homogeneous option, the default), `mvtest means` tests the equality of the group means by using `manova` to compute the classic quartet of test statistics: Wilks's lambda, Pillai's trace, Lawley–Hotelling trace, and Roy's largest root. See [\[MV\] manova](#) for details. For $m = 2$ samples, the four tests are equivalent.

`mvtest means` has the `heterogeneous` and `lr` options, which remove the assumption of equal covariance matrices for the `by()` groups. The statistic produced with the `heterogeneous` option depends on whether there are $m = 2$ groups or $m > 2$ groups.

With the `heterogeneous` option and $m = 2$ samples, the test for equal means is computed using the affine-invariant modification by [Krishnamoorthy and Yu \(2004\)](#) of the test proposed by [Nel and Van der Merwe \(1986\)](#):

$$T_{\text{mm2}}^2 = (\bar{\mathbf{x}}_1 - \bar{\mathbf{x}}_2)' \tilde{\mathbf{S}}^{-1} (\bar{\mathbf{x}}_1 - \bar{\mathbf{x}}_2)$$

where

$$F_{\text{mm2}} = \frac{v - k + 1}{vk} T_{\text{mm2}}^2$$

is approximately $F(k, v - k + 1)$ distributed, and where $\tilde{\mathbf{S}}_j = \mathbf{S}_j / (N_j - 1)$, $\tilde{\mathbf{S}} = \tilde{\mathbf{S}}_1 + \tilde{\mathbf{S}}_2$, $v = k(k + 1) / (a_1 + a_2)$, and

$$a_j = \left[\text{trace} \left\{ (\tilde{\mathbf{S}}_j \tilde{\mathbf{S}}^{-1})^2 \right\} + \text{trace} \left(\tilde{\mathbf{S}}_j \tilde{\mathbf{S}}^{-1} \right)^2 \right] / (N_j - 1)$$

With the heterogeneous option and $m > 2$ samples, `mvtest means` computes the Wald test for equal means (Seber 1984, 445–447)

$$T_{\text{mmw}} = \sum_{j=1}^m (\bar{\mathbf{x}}_j - \bar{\mathbf{x}})' \mathbf{W}_j (\bar{\mathbf{x}}_j - \bar{\mathbf{x}})$$

where $\mathbf{W}_j = \{\mathbf{S}_j / (N_j - 1)\}^{-1}$, $\mathbf{W} = \sum_{j=1}^m \mathbf{W}_j$, and $\bar{\mathbf{x}} = \mathbf{W}^{-1} \sum_{j=1}^m \mathbf{W}_j \bar{\mathbf{x}}_j$.

James (1954) showed that the upper α quantile of T_{mmw} under the null hypothesis can, to order N_i^{-1} , be obtained as $\chi_{r,\alpha}^2(a + b\chi_{r,\alpha}^2)$, where $\chi_{r,\alpha}^2$ is the upper α quantile of a χ^2 with $r = k(m - 1)$ degrees of freedom, and

$$a = 1 + \frac{1}{2r} \sum_{j=1}^m \frac{\{\text{trace}(\mathbf{I} - \mathbf{W}^{-1}\mathbf{W}_j)\}^2}{N_j - 1}$$

$$b = \frac{1}{r(r+2)} \left[\sum_{j=1}^m \frac{\text{trace}\{(\mathbf{I} - \mathbf{W}^{-1}\mathbf{W}_j)^2\}}{N_j - 1} + \frac{\{\text{trace}(\mathbf{I} - \mathbf{W}^{-1}\mathbf{W}_j)\}^2}{2(N_j - 1)} \right]$$

`mvtest` computes the p -value associated with the observed T_{mmw} by inverting James’s expansion with a scalar solver using Brent’s (1973) algorithm.

With the `lr` option, `mvtest means` provides a likelihood-ratio test for constant means across $m \geq 2$ samples, without the assumption of constant covariances. To fit the null model with a common mean, $\boldsymbol{\mu}_c$, an iterative procedure proposed by Mardia, Kent, and Bibby (1979, 142–143) is used. Let h denote the iteration. Let $\widehat{\boldsymbol{\Sigma}}_j^{(0)} = \mathbf{S}_j$, and $\widehat{\boldsymbol{\mu}}_c^{(0)}$ be obtained using the second formula below, then iterate these two formulas starting with $h = 1$ until convergence is achieved.

$$\widehat{\boldsymbol{\Sigma}}_j^{(h)} = \mathbf{S}_j + (\bar{\mathbf{x}}_j - \widehat{\boldsymbol{\mu}}_c^{(h-1)}) (\bar{\mathbf{x}}_j - \widehat{\boldsymbol{\mu}}_c^{(h-1)})'$$

$$\widehat{\boldsymbol{\mu}}_c^{(h)} = \left\{ \sum_{j=1}^m N_j (\widehat{\boldsymbol{\Sigma}}_j^{(h)})^{-1} \right\}^{-1} \left\{ \sum_{j=1}^m N_j (\widehat{\boldsymbol{\Sigma}}_j^{(h)})^{-1} \bar{\mathbf{x}}_j \right\}$$

v

The likelihood-ratio χ^2 statistic

$$\chi_{\text{mml}}^2 = \sum_{j=1}^m N_j \ln \left\{ 1 + (\bar{\mathbf{x}}_j - \widehat{\boldsymbol{\mu}}_c^{(h)})' (\widehat{\boldsymbol{\Sigma}}_j^{(h)})^{-1} (\bar{\mathbf{x}}_j - \widehat{\boldsymbol{\mu}}_c^{(h)}) \right\}$$

is approximately χ^2 distributed with $k(m - 1)$ degrees of freedom.

References

- Brent, R. P. 1973. *Algorithms for Minimization without Derivatives*. Englewood Cliffs, NJ: Prentice Hall. (Reprinted in paperback by Dover Publications, Mineola, NY, January 2002).
- Buot, M.-L. G., S. Hoşten, and D. S. P. Richards. 2007. Counting and locating the solutions of polynomial systems of maximum likelihood equations, II: The Behrens–Fisher problem. *Statistica Sinica* 17: 1343–1354.
- James, G. S. 1954. Tests of linear hypotheses in univariate and multivariate analysis when the ratios of the population variances are unknown. *Biometrika* 41: 19–43. <https://doi.org/10.2307/2333003>.

- Johnson, R. A., and D. W. Wichern. 2007. *Applied Multivariate Statistical Analysis*. 6th ed. Englewood Cliffs, NJ: Prentice Hall.
- Kramer, C. Y., and D. R. Jensen. 1969. Fundamentals of multivariate analysis, part I. Inference about means. *Journal of Quality Technology* 1: 120–133. <https://doi.org/10.1080/00224065.1969.11980360>.
- Krishnamoorthy, K., and J. Yu. 2004. Modified Nel and Van der Merwe test for the multivariate Behrens–Fisher problem. *Statistics and Probability Letters* 66: 161–169. <https://doi.org/10.1016/j.spl.2003.10.012>.
- Mardia, K. V., J. T. Kent, and J. M. Bibby. 1979. *Multivariate Analysis*. London: Academic Press.
- Mardia, K. V., J. T. Kent, and C. C. Taylor. 2024. *Multivariate Analysis*. 2nd ed. Hoboken, NJ: Wiley.
- Nel, D. G., and C. A. Van der Merwe. 1986. A solution to the multivariate Behrens–Fisher problem. *Communications in Statistics—Theory and Methods* 15: 3719–3735. <https://doi.org/10.1080/03610928608829342>.
- Rencher, A. C. 1998. *Multivariate Statistical Inference and Applications*. New York: Wiley.
- Rencher, A. C., and W. F. Christensen. 2012. *Methods of Multivariate Analysis*. 3rd ed. Hoboken, NJ: Wiley. <https://doi.org/10.1002/9781118391686>.
- Seber, G. A. F. 1984. *Multivariate Observations*. New York: Wiley. <https://doi.org/10.1002/9780470316641>.

Also see

- [MV] **hotelling** — Hotelling’s T^2 generalized means test
- [MV] **manova** — Multivariate analysis of variance and covariance
- [R] **mean** — Estimate means
- [R] **ttest** — t tests (mean-comparison tests)

