

Package ‘qvalue’

April 15, 2017

Type Package

Title Q-value estimation for false discovery rate control

Version 2.6.0

Date 2015-03-24

Author John D. Storey with contributions from Andrew J. Bass, Alan Dabney and David Robinson

Maintainer John D. Storey <jstorey@princeton.edu>, Andrew J. Bass <ajbass@princeton.edu>

biocViews MultipleComparisons

Description This package takes a list of p-values resulting from the simultaneous testing of many hypotheses and estimates their q-values and local FDR values. The q-value of a test measures the proportion of false positives incurred (called the false discovery rate) when that particular test is called significant. The local FDR measures the posterior probability the null hypothesis is true given the test's p-value. Various plots are automatically generated, allowing one to make sensible significance cut-offs. Several mathematical results have recently been shown on the conservative accuracy of the estimated q-values from this software. The software can be applied to problems in genomics, brain imaging, astrophysics, and data mining.

VignetteBuilder knitr

Imports splines, ggplot2, grid, reshape2

Suggests knitr

Depends R(>= 2.10)

URL <http://github.com/jdstorey/qvalue>

License LGPL

RoxygenNote 5.0.1

NeedsCompilation no

R topics documented:

empPvals	2
hedenfalk	3
hist.qvalue	4
lfdr	6
pi0est	7
plot.qvalue	9
qvalue	10
summary.qvalue	12
write.qvalue	14

Index	15
--------------	-----------

empPvals	<i>Calculate p-values from a set of observed test statistics and simulated null test statistics</i>
----------	---

Description

Calculates p-values from a set of observed test statistics and simulated null test statistics

Usage

```
empPvals(stat, stat0, pool = TRUE)
```

Arguments

stat	A vector of calculated test statistics.
stat0	A vector or matrix of simulated or data-resampled null test statistics.
pool	If FALSE, stat0 must be a matrix with the number of rows equal to the length of stat. Default is TRUE.

Details

The argument `stat` must be such that the larger the value is the more deviated (i.e., "more extreme") from the null hypothesis it is. Examples include an F-statistic or the absolute value of a t-statistic. The argument `stat0` should be calculated analogously on data that represents observations from the null hypothesis distribution. The p-values are calculated as the proportion of values from `stat0` that are greater than or equal to that from `stat`. If `pool=TRUE` is selected, then all of `stat0` is used in calculating the p-value for a given entry of `stat`. If `pool=FALSE`, then it is assumed that `stat0` is a matrix, where `stat0[i,]` is used to calculate the p-value for `stat[i]`. The function `empPvals` calculates "pooled" p-values faster than using a for-loop.

See page 18 of the Supporting Information in Storey et al. (2005) PNAS (<http://www.pnas.org/content/suppl/2005/08/26/0504609102.DC1/04609SuppAppendix.pdf>) for an explanation as to why calculating p-values from pooled empirical null statistics and then estimating FDR on these p-values is equivalent to directly thresholding the test statistics themselves and utilizing an analogous FDR estimator.

Value

A vector of p-values calculated as described above.

Author(s)

John D. Storey

References

Storey JD and Tibshirani R. (2003) Statistical significance for genome-wide experiments. *Proceedings of the National Academy of Sciences*, 100: 9440-9445.

<http://www.pnas.org/content/100/16/9440.full>

Storey JD, Xiao W, Leek JT, Tompkins RG, Davis RW. (2005) Significance analysis of time course microarray experiments. *Proceedings of the National Academy of Sciences*, 102 (36), 12837-12842.

<http://www.pnas.org/content/102/36/12837.full.pdf?with-ds=yes>

See Also

[qvalue](#)

Examples

```
# import data
data(hedenfalk)
stat <- hedenfalk$stat
stat0 <- hedenfalk$stat0 #vector from null distribution

# calculate p-values
p.pooled <- empPvals(stat=stat, stat0=stat0)
p.testspecific <- empPvals(stat=stat, stat0=stat0, pool=FALSE)

# compare pooled to test-specific p-values
qqplot(p.pooled, p.testspecific); abline(0,1)
```

hedenfalk

P-values and test-statistics from the Hedenfalk et al. (2001) gene expression dataset

Description

The data from the breast cancer gene expression study of Hedenfalk et al. (2001) were obtained and analyzed. A comparison was made between 3,226 genes of two mutation types, BRCA1 (7 arrays) and BRCA2 (8 arrays). The data included here are p-values, test-statistics, and permutation null test-statistics obtained from a two-sample t-test analysis on a set of 3170 genes, as described in Storey and Tibshirani (2003).

Usage

```
data(hedenfalk)
```

Value

A list called hedenfalk containing:

`p` Vector of 3,170 p-values of tests comparing BRCA1 to BRCA2.
`stat` Vector of 3,170 absolute two-sample t-statistics comparing BRCA1 to BRCA2.
`stat0` A 3,170 by 100 matrix of absolute two-sample t-statistics from 100 independent permutations of the BRCA1 and BRCA2 labels; the row `stat0[i,]` contains the permutation statistics corresponding to observed statistic `stat[i]`.

References

Hedenfalk I et al. (2001). Gene expression profiles in hereditary breast cancer. *New England Journal of Medicine*, 344: 539-548.

Storey JD and Tibshirani R. (2003). Statistical significance for genome-wide studies. *Proceedings of the National Academy of Sciences*, 100: 9440-9445.

<http://www.pnas.org/content/100/16/9440.full>

See Also

[qvalue](#), [empPvals](#)

Examples

```
# import data
data(hedenfalk)
stat <- hedenfalk$stat
stat0 <- hedenfalk$stat0 #vector from null distribution

p.pooled <- empPvals(stat=stat, stat0=stat0)
p.testspecific <- empPvals(stat=stat, stat0=stat0, pool=FALSE)

#compare pooled to test-specific p-values
qqplot(p.pooled, p.testspecific); abline(0,1)

# calculate q-values and view results
qobj <- qvalue(p.pooled)
summary(qobj)
hist(qobj)
plot(qobj)
```

hist.qvalue

Histogram of p-values

Description

Histogram of p-values

Usage

```
## S3 method for class 'qvalue'
hist(x, ...)
```

Arguments

x A q-value object.
... Additional arguments, currently unused.

Details

This function allows one to view a histogram of the p-values along with line plots of the q-values and local FDR values versus p-values. The π_0 estimate is also displayed.

Value

Nothing of interest.

Author(s)

Andrew J. Bass

References

Storey JD. (2002) A direct approach to false discovery rates. *Journal of the Royal Statistical Society, Series B*, 64: 479-498.

<http://onlinelibrary.wiley.com/doi/10.1111/1467-9868.00346/abstract>

Storey JD and Tibshirani R. (2003) Statistical significance for genome-wide experiments. *Proceedings of the National Academy of Sciences*, 100: 9440-9445.

<http://www.pnas.org/content/100/16/9440.full>

Storey JD. (2003) The positive false discovery rate: A Bayesian interpretation and the q-value. *Annals of Statistics*, 31: 2013-2035.

http://projecteuclid.org/DPubS/Repository/1.0/Disseminate?view=body&id=pdf_1&handle=euclid.aos/1074290335

conservative point estimation, and simultaneous conservative consistency of false discovery rates: A unified approach. *Journal of the Royal Statistical Society, Series B*, 66: 187-205.

<http://onlinelibrary.wiley.com/doi/10.1111/j.1467-9868.2004.00439.x/abstract>

Storey JD. (2011) False discovery rates. In *International Encyclopedia of Statistical Science*.

http://genomine.org/papers/Storey_FDR_2011.pdf

<http://www.springer.com/statistics/book/978-3-642-04897-5>

See Also

[qvalue](#), [plot.qvalue](#), [summary.qvalue](#)

Examples

```
# import data
data(hedenfalk)
p <- hedenfalk$p

# make histogram
qobj <- qvalue(p)
hist(qobj)
```

lfdr

*Estimate local False Discovery Rate (FDR)***Description**

Estimate the local FDR values from p-values.

Usage

```
lfdr(p, pi0 = NULL, trunc = TRUE, monotone = TRUE, transf = c("probit",
  "logit"), adj = 1.5, eps = 10^-8, ...)
```

Arguments

p	A vector of p-values (only necessary input).
pi0	Estimated proportion of true null p-values. If NULL, then <code>pi0est</code> is called.
trunc	If TRUE, local FDR values >1 are set to 1. Default is TRUE.
monotone	If TRUE, local FDR values are non-decreasing with increasing p-values. Default is TRUE; this is recommended.
transf	Either a "probit" or "logit" transformation is applied to the p-values so that a local FDR estimate can be formed that does not involve edge effects of the [0,1] interval in which the p-values lie.
adj	Numeric value that is applied as a multiple of the smoothing bandwidth used in the density estimation. Default is <code>adj=1.0</code> .
eps	Numeric value that is threshold for the tails of the empirical p-value distribution. Default is 10^{-8} .
...	Additional arguments, passed to <code>pi0est</code> .

Details

It is assumed that null p-values follow a Uniform(0,1) distribution. The estimated proportion of true null hypotheses $\hat{\pi}_0$ is either a user-provided value or the value calculated via `pi0est`. This function works by forming an estimate of the marginal density of the observed p-values, say $\hat{f}(p)$. Then the local FDR is estimated as $\text{lfdr}(p) = \hat{\pi}_0 / \hat{f}(p)$, with adjustments for monotonicity and to guarantee that $\text{lfdr}(p) \leq 1$. See the Storey (2011) reference below for a concise mathematical definition of local FDR.

Value

A vector of estimated local FDR values, with each entry corresponding to the entries of the input p-value vector p.

Author(s)

John D. Storey

References

Efron B, Tibshirani R, Storey JD, and Tisher V. (2001) Empirical Bayes analysis of a microarray experiment. *Journal of the American Statistical Association*, 96: 1151-1160.

<http://www.tandfonline.com/doi/abs/10.1198/016214501753382129>

Storey JD. (2003) The positive false discovery rate: A Bayesian interpretation and the q-value. *Annals of Statistics*, 31: 2013-2035.

http://projecteuclid.org/DPubS/Repository/1.0/Disseminate?view=body&id=pdf_1&handle=euclid.aos/1074290335

Storey JD. (2011) False discovery rates. In *International Encyclopedia of Statistical Science*.

http://genomine.org/papers/Storey_FDR_2011.pdf

<http://www.springer.com/statistics/book/978-3-642-04897-5>

See Also

[qvalue](#), [pi0est](#), [hist.qvalue](#)

Examples

```
# import data
data(hedenfalk)
p <- hedenfalk$p
lfdrVals <- lfdr(p)

# plot local FDR values
qobj = qvalue(p)
hist(qobj)
```

pi0est

Proportion of true null p-values

Description

Estimates the proportion of true null p-values, i.e., those following the Uniform(0,1) distribution.

Usage

```
pi0est(p, lambda = seq(0.05, 0.95, 0.05), pi0.method = c("smoother",
  "bootstrap"), smooth.df = 3, smooth.log.pi0 = FALSE, ...)
```

Arguments

p	A vector of p-values (only necessary input).
lambda	The value of the tuning parameter to estimate π_0 . Must be in [0,1). Optional, see Storey (2002).
pi0.method	Either "smoother" or "bootstrap"; the method for automatically choosing tuning parameter in the estimation of π_0 , the proportion of true null hypotheses.
smooth.df	Number of degrees-of-freedom to use when estimating π_0 with a smoother. Optional.

`smooth.log.pi0` If TRUE and `pi0.method = "smoother"`, π_0 will be estimated by applying a smoother to a scatterplot of $\log(\pi_0)$ estimates against the tuning parameter λ . Optional.

... Arguments passed from `qvalue` function.

Details

If no options are selected, then the method used to estimate π_0 is the smoother method described in Storey and Tibshirani (2003). The bootstrap method is described in Storey, Taylor & Siegmund (2004). A closed form solution of the bootstrap method is used in the package and is significantly faster.

Value

Returns a list:

<code>pi0</code>	A numeric that is the estimated proportion of true null p-values.
<code>pi0.lambda</code>	A vector of the proportion of null values at the λ values (see vignette).
<code>lambda</code>	A vector of λ value(s) utilized in calculating <code>pi0.lambda</code> .
<code>pi0.smooth</code>	A vector of fitted values from the smoother fit to the π_0 estimates at each lambda value (<code>pi0.method="bootstrap"</code> returns NULL).

Author(s)

John D. Storey

References

Storey JD. (2002) A direct approach to false discovery rates. *Journal of the Royal Statistical Society, Series B*, 64: 479-498.

<http://onlinelibrary.wiley.com/doi/10.1111/1467-9868.00346/abstract>

Storey JD and Tibshirani R. (2003) Statistical significance for genome-wide experiments. *Proceedings of the National Academy of Sciences*, 100: 9440-9445.

Storey JD. (2003) The positive false discovery rate: A Bayesian interpretation and the q-value. *Annals of Statistics*, 31: 2013-2035.

http://projecteuclid.org/DPubS/Repository/1.0/Disseminate?view=body&id=pdf_1&handle=euclid.aos/1074290335

Storey JD, Taylor JE, and Siegmund D. (2004) Strong control, conservative point estimation, and simultaneous conservative consistency of false discovery rates: A unified approach. *Journal of the Royal Statistical Society, Series B*, 66: 187-205.

<http://onlinelibrary.wiley.com/doi/10.1111/j.1467-9868.2004.00439.x/abstract>

Storey JD. (2011) False discovery rates. In *International Encyclopedia of Statistical Science*.

http://genomine.org/papers/Storey_FDR_2011.pdf

<http://www.springer.com/statistics/book/978-3-642-04897-5>

See Also

[qvalue](#)

Examples

```
# import data
data(hedenfalk)
p <- hedenfalk$p

# proportion of null p-values
nullRatio <- pi0est(p)
nullRatioS <- pi0est(p, lambda=seq(0.40, 0.95, 0.05), smooth.log.pi0="TRUE")
nullRatioM <- pi0est(p, pi0.method="bootstrap")

# check behavior of estimate over lambda
# also, pi0est arguments can be passed to qvalue
qobj = qvalue(p, lambda=seq(0.05, 0.95, 0.1), smooth.log.pi0="TRUE")
hist(qobj)
plot(qobj)
```

plot.qvalue

*Plotting function for q-value object***Description**

Graphical display of the q-value object

Usage

```
## S3 method for class 'qvalue'
plot(x, rng = c(0, 0.1), ...)
```

Arguments

x	A q-value object.
rng	Range of q-values to show. Optional
...	Additional arguments. Currently unused.

Details

The function plot allows one to view several plots:

1. The estimated π_0 versus the tuning parameter λ .
2. The q-values versus the p-values.
3. The number of significant tests versus each q-value cutoff.
4. The number of expected false positives versus the number of significant tests.

This function makes four plots. The first is a plot of the estimate of π_0 versus its tuning parameter λ . In most cases, as λ gets larger, the bias of the estimate decreases, yet the variance increases. Various methods exist for balancing this bias-variance trade-off (Storey 2002, Storey & Tibshirani 2003, Storey, Taylor & Siegmund 2004). Comparing your estimate of π_0 to this plot allows one to gauge its quality. The remaining three plots show how many tests are called significant and how many false positives to expect for each q-value cut-off. A thorough discussion of these plots can be found in Storey & Tibshirani (2003).

Value

Nothing of interest.

Author(s)

John D. Storey, Andrew J. Bass

References

Storey JD. (2002) A direct approach to false discovery rates. *Journal of the Royal Statistical Society, Series B*, 64: 479-498.

<http://onlinelibrary.wiley.com/doi/10.1111/1467-9868.00346/abstract>

Storey JD and Tibshirani R. (2003) Statistical significance for genome-wide experiments. *Proceedings of the National Academy of Sciences*, 100: 9440-9445.

<http://www.pnas.org/content/100/16/9440.full>

Storey JD. (2003) The positive false discovery rate: A Bayesian interpretation and the q-value. *Annals of Statistics*, 31: 2013-2035.

http://projecteuclid.org/DPubS/Repository/1.0/Disseminate?view=body&id=pdf_1&handle=euclid.aos/1074290335

Storey JD, Taylor JE, and Siegmund D. (2004) Strong control, conservative point estimation, and simultaneous conservative consistency of false discovery rates: A unified approach. *Journal of the Royal Statistical Society, Series B*, 66: 187-205.

Storey JD. (2011) False discovery rates. In *International Encyclopedia of Statistical Science*.

http://genomine.org/papers/Storey_FDR_2011.pdf

<http://www.springer.com/statistics/book/978-3-642-04897-5>

See Also

[qvalue](#), [write.qvalue](#), [summary.qvalue](#)

Examples

```
# import data
data(hedenfalk)
p <- hedenfalk$p
qobj <- qvalue(p)

plot(qobj, rng=c(0.0, 0.3))
```

qvalue

Estimate the q-values for a given set of p-values

Description

Estimate the q-values for a given set of p-values. The q-value of a test measures the proportion of false positives incurred (called the false discovery rate) when that particular test is called significant.

Usage

```
qvalue(p, fdr.level = NULL, pfdR = FALSE, ...)
```

Arguments

p	A vector of p-values (only necessary input).
fdr.level	A level at which to control the FDR. Must be in (0,1]. Optional; if this is selected, a vector of TRUE and FALSE is returned that specifies whether each q-value is less than fdr.level or not.
pfdR	An indicator of whether it is desired to make the estimate more robust for small p-values and a direct finite sample estimate of pFDR – optional.
...	Additional arguments passed to <code>pi0est</code> and <code>lfdr</code> .

Details

The function `pi0est` is called internally and calculates the estimate of π_0 , the proportion of true null hypotheses. The function `lfdr` is also called internally and calculates the estimated local FDR values. Arguments for these functions can be included via `...` and will be utilized in the internal calls made in `qvalue`. See http://genomine.org/papers/Storey_FDR_2011.pdf for a brief introduction to FDRs and q-values.

Value

A list of object type "qvalue" containing:

call	Function call.
pi0	An estimate of the proportion of null p-values.
qvalues	A vector of the estimated q-values (the main quantity of interest).
pvalues	A vector of the original p-values.
lfdr	A vector of the estimated local FDR values.
significant	If fdr.level is specified, and indicator of whether the q-value fell below fdr.level (taking all such q-values to be significant controls FDR at level fdr.level).
pi0.lambda	An estimate of the proportion of null p-values at each λ value (see vignette).
lambda	A vector of the λ values utilized to obtain pi0.lambda.

Author(s)

John D. Storey

References

- Storey JD. (2002) A direct approach to false discovery rates. *Journal of the Royal Statistical Society, Series B*, 64: 479-498.
<http://onlinelibrary.wiley.com/doi/10.1111/1467-9868.00346/abstract> Storey JD and Tibshirani R. (2003) Statistical significance for genome-wide experiments. *Proceedings of the National Academy of Sciences*, 100: 9440-9445.
<http://www.pnas.org/content/100/16/9440.full>
- Storey JD. (2003) The positive false discovery rate: A Bayesian interpretation and the q-value. *Annals of Statistics*, 31: 2013-2035.

http://projecteuclid.org/DPubS/Repository/1.0/Disseminate?view=body&id=pdf_1&handle=euclid.aos/1074290335

Storey JD, Taylor JE, and Siegmund D. (2004) Strong control, conservative point estimation, and simultaneous conservative consistency of false discovery rates: A unified approach. *Journal of the Royal Statistical Society, Series B*, 66: 187-205.

<http://onlinelibrary.wiley.com/doi/10.1111/j.1467-9868.2004.00439.x/abstract>

Storey JD. (2011) False discovery rates. In *International Encyclopedia of Statistical Science*.

http://genomine.org/papers/Storey_FDR_2011.pdf

<http://www.springer.com/statistics/book/978-3-642-04897-5>

See Also

[pi0est](#), [lfdr](#), [summary.qvalue](#), [plot.qvalue](#), [hist.qvalue](#), [write.qvalue](#)

Examples

```
# import data
data(hedenfalk)
p <- hedenfalk$p

# get q-value object
qobj <- qvalue(p)
plot(qobj)
hist(qobj)

# options available
qobj <- qvalue(p, lambda=0.5, pfd=TRUE)
qobj <- qvalue(p, fdr.level=0.05, pi0.method="bootstrap", adj=1.2)
```

summary.qvalue	<i>Display q-value object</i>
----------------	-------------------------------

Description

Display summary information for a q-value object.

Usage

```
## S3 method for class 'qvalue'
summary(object, cuts = c(1e-04, 0.001, 0.01, 0.025, 0.05,
  0.1, 1), digits = getOption("digits"), ...)
```

Arguments

object	A q-value object.
cuts	Vector of significance values to use for table (optional).
digits	Significant digits to display (optional).
...	Additional arguments; currently unused.

Details

summary shows the original call, estimated proportion of true null hypotheses, and a table comparing the number of significant calls for the p-values, estimated q-values, and estimated local FDR values using a set of cutoffs given by cuts.

Value

Invisibly returns the original object.

Author(s)

John D. Storey, Andrew J. Bass, Alan Dabney

References

Storey JD. (2002) A direct approach to false discovery rates. *Journal of the Royal Statistical Society, Series B*, 64: 479-498.

<http://onlinelibrary.wiley.com/doi/10.1111/1467-9868.00346/abstract>

Storey JD and Tibshirani R. (2003) Statistical significance for genome-wide experiments. *Proceedings of the National Academy of Sciences*, 100: 9440-9445.

<http://www.pnas.org/content/100/16/9440.full>

Storey JD. (2003) The positive false discovery rate: A Bayesian interpretation and the q-value. *Annals of Statistics*, 31: 2013-2035.

http://projecteuclid.org/DPubS/Repository/1.0/Disseminate?view=body&id=pdf_1&handle=euclid.aos/1074290335

Storey JD, Taylor JE, and Siegmund D. (2004) Strong control, conservative point estimation, and simultaneous conservative consistency of false discovery rates: A unified approach. *Journal of the Royal Statistical Society, Series B*, 66: 187-205.

<http://onlinelibrary.wiley.com/doi/10.1111/j.1467-9868.2004.00439.x/abstract>

Storey JD. (2011) False discovery rates. In *International Encyclopedia of Statistical Science*.

http://genomine.org/papers/Storey_FDR_2011.pdf

<http://www.springer.com/statistics/book/978-3-642-04897-5>

See Also

[qvalue](#), [plot.qvalue](#), [write.qvalue](#)

Examples

```
# import data
data(hedenfalk)
p <- hedenfalk$p

# get summary results from q-value object
qobj <- qvalue(p)
summary(qobj, cuts=c(0.01, 0.05))
```

write.qvalue	<i>Write results to file</i>
--------------	------------------------------

Description

Write the results of the q-value object to a file.

Usage

```
write.qvalue(x, file = NULL, sep = " ", eol = "\n", na = "NA",  
            row.names = FALSE, col.names = TRUE)
```

Arguments

x	A q-value object.
file	Output filename (optional).
sep	Separation between columns.
eol	Character to print at the end of each line.
na	String to use when there are missing values.
row.names	logical. Specify whether row names are to be printed.
col.names	logical. Specify whether column names are to be printed.

Details

The output file includes: (i) p-values, (ii) q-values (iii) local FDR values, and (iv) the estimate of π_0 , one per line. If an FDR significance level was specified in the call to [qvalue](#), the significance level is printed and an indicator of significance is included.

Value

Nothing of interest.

Author(s)

John D. Storey, Andrew J. Bass

See Also

[qvalue](#), [plot.qvalue](#), [summary.qvalue](#)

Examples

```
# import data  
data(hedenfalk)  
p <- hedenfalk$p  
  
# write q-value object  
qobj <- qvalue(p)  
write.qvalue(qobj, file="myresults.txt")
```

Index

*Topic **Discovery**
lfd, 6

*Topic **False**
lfd, 6

*Topic **Rate,**
lfd, 6

*Topic **dataset,**
hedenfalk, 3

*Topic **hedenfalk**
hedenfalk, 3

*Topic **histogram**
hist.qvalue, 4

*Topic **lfd**
lfd, 6

*Topic **local**
lfd, 6

*Topic **nulls**
pi0est, 7

*Topic **pi0est,**
pi0est, 7

*Topic **plot**
plot.qvalue, 9

*Topic **proportion**
pi0est, 7

*Topic **pvalues**
empPvals, 2

*Topic **qvalue**
qvalue, 10

*Topic **summary**
summary.qvalue, 12

*Topic **true**
pi0est, 7

*Topic **write.qvalue**
write.qvalue, 14

empPvals, 2, 4

hedenfalk, 3
hist, (hist.qvalue), 4
hist.qvalue, 4, 7, 12

lfd, 6, 11, 12

pi0est, 6, 7, 7, 11, 12

plot, (plot.qvalue), 9
plot.qvalue, 5, 9, 12–14

qvalue, 3–5, 7, 8, 10, 10, 11, 13, 14

summary, (summary.qvalue), 12
summary.qvalue, 5, 10, 12, 12, 14

write.qvalue, 10, 12, 13, 14