

Using fitPS to fit number of groups data

In this example we will learn how to use `fitPS` to fit a zeta distribution to data from a survey where the number of groups of glass found is recorded. The data in this example comes from Roux et al. (2001), who surveyed the footwear of 776 individuals in south-eastern Australia, and is summarised in the table below.

% latex table generated in R 4.5.3 by xtable 1.8-8 package % Wed Jun 10 17:21:43 2026

n	r_n
0	754
1	9
2	8
3	4
4	1

This data set is built into the package and can be accessed from the `Psurveys` object. That is, we can type:

```
> data("Psurveys")
> roux = Psurveys$roux
```

The data is stored as an object of class `psData`. This probably will not be important to most users. Details can be found in the **Value** section of the help page for `readData`. There is an S3 `print` method for objects of this type, meaning that if we print the object—either by typing its name at the command prompt, or by explicitly calling `print`—we will get formatted printing of the information contained within the object. Specifically:

- the values of n and r_n are printed in tabular format, where n is the number of groups or the size of the groups, and r_n is the number of times n has been observed in the survey;
- the type of survey is printed, either “Number of Groups” or “Group Size”;
- if the object has a reference or notes attached to it, these are printed as well.

For example:

```
> roux
```

Number of Groups

n	rn
0	754
1	9
2	8
3	4
4	1

Roux C, Kirk R, Benson S, Van Haren T, Petterd C (2001).
"Glass particles in footwear of members of the public in
south-eastern Australia—a survey." *_Forensic Science
International_*, *116*(2), 149–156.
doi:10.1016/S0379-0738(00)00355-8
<<https://doi.org/10.1016/S0379-0738%2800%2900355-8>>.

It is very simple to fit a zeta distribution to this data set. We do this using the `fitDist` function.

```
> fit = fitDist(roux)
```

The function returns an object of class `psFit`, the details of which can be found in the help page for `fitDist`. There are both S3 `print` and S3 `plot` methods for objects of this class. The `print` method displays an estimate of the shape parameter α and the standard error of that estimate, $\widehat{\text{sd}}(\hat{\alpha}) = \text{se}(\hat{\alpha})$. The reported value is the same shape parameter used throughout `fitPS` and stored in the fitted object. The `print` method also displays the first 10 fitted probabilities from the model by default.

```
> fit
```

```
The estimated shape parameter is 4.9544
```

```
The standard error of shape parameter is 0.2366
```

```
The first 10 fitted values are:
```

P0	P1	P2	P3	P4
9.631547e-01	3.106447e-02	4.167082e-03	1.001917e-03	3.316637e-04
P5	P6	P7	P8	P9
1.344002e-04	6.262053e-05	3.231467e-05	1.802885e-05	1.069709e-05

The package provides a `confint` method for fitted values. The method returns both a Wald confidence interval and a profile likelihood interval. The Wald interval has lower and upper bounds given by $\hat{\alpha} \pm z^* \times \text{se}(\hat{\alpha})$. The profile likelihood interval finds the endpoints of the interval that satisfies

$$-2[\ell(\hat{\alpha}; \mathbf{x}) - \ell(\alpha; \mathbf{x})] \leq \chi_1^2(\alpha),$$

where $\ell(\alpha; \mathbf{x})$ is the value of the log-likelihood for shape parameter α . The two intervals are returned as elements of a list named `wald` and `prof`, respectively.

```
> ci = confint(fit)
```

```
> ci$wald
```

2.5%	97.5%
4.490761	5.418099

```
> ci$prof
```

2.5%	97.5%
4.520495	5.451277

References

Roux, C., Kirk, R., Benson, S., Van Haren, T., and Petterd, C. I. (2001). Glass particles in footwear of members of the public in south-eastern Australia: a survey. *Forensic Science International*, 116(2), 149-156.