
Isqfit Documentation

Release 4.5

G. P. Lepage

July 31, 2013

1	Overview and Tutorial	3
1.1	Introduction	3
1.2	Making Fake Data	6
1.3	Basic Fits	7
1.4	Chained Fits	14
1.5	x has Error Bars	16
1.6	Correlated Parameters; Gaussian Bayes Factor	18
1.7	Tuning Priors and the Empirical Bayes Criterion	19
1.8	Partial Errors and Error Budgets	21
1.9	y has No Error Bars	22
1.10	SVD Cuts and Roundoff Error	26
1.11	Bootstrap Error Analysis	28
1.12	Testing Fits with Simulated Data	29
1.13	Positive Parameters	32
1.14	Troubleshooting	34
1.15	Appendix: A Simple Pedagogical Example	35
2	gvar - Gaussian Random Variables	41
2.1	Introduction	41
2.2	Creating Gaussian Variables	41
2.3	Computing Covariance Matrices	43
2.4	Random Number Generators	44
2.5	Limitations	46
2.6	Implementation Notes; Derivatives; Optimizations	46
2.7	Utilities	47
2.8	Classes	53
2.9	Requirements	58
3	gvar.dataset - Random Data Sets	59
3.1	Introduction	59
3.2	Functions	61
3.3	Classes	63
4	lsqfit - Nonlinear Least Squares Fitting	67
4.1	Introduction	67
4.2	Formal Background	68
4.3	nonlinear_fit Objects	69
4.4	Functions	74
4.5	Utility Classes	76
4.6	Requirements	79

5 Indices and tables	81
Python Module Index	83
Index	85

Contents:

OVERVIEW AND TUTORIAL

1.1 Introduction

The modules defined here are designed to facilitate least-squares fitting of noisy data by multi-dimensional, nonlinear functions of arbitrarily many parameters. The central module is `lsqfit` because it provides the fitting functions. `lsqfit` makes heavy use of auxiliary module `gvar`, which provides tools that facilitate the analysis of error propagation, and also the creation of complicated multi-dimensional Gaussian distributions.

The following (complete) code illustrates basic usage of `lsqfit`:

```
import numpy as np
import gvar as gv
import lsqfit

y = {
    # data for the dependent variable
    'data1' : gv.gvar([1.376, 2.010], [[ 0.0047, 0.01], [ 0.01, 0.056]]),
    'data2' : gv.gvar([1.329, 1.582], [[ 0.0047, 0.0067], [0.0067, 0.0136]]),
    'b/a'    : gv.gvar(2.0, 0.5)
}
x = {
    # independent variable
    'data1' : np.array([0.1, 1.0]),
    'data2' : np.array([0.1, 0.5])
}
prior = dict(a=gv.gvar(0.5, 0.5), b=gv.gvar(0.5, 0.5))

def fcn(x, p):
    # fit function of x and parameters p
    ans = {}
    for k in ['data1', 'data2']:
        ans[k] = gv.exp(p['a'] + x[k] * p['b'])
    ans['b/a'] = p['b'] / p['a']
    return ans

# do the fit
fit = lsqfit.nonlinear_fit(data=(x, y), prior=prior, fcn=fcn)
print(fit.format(100)) # print standard summary of fit

p = fit.p
# best-fit values for parameters
outputs = dict(a=p['a'], b=p['b'])
outputs['b/a'] = p['b']/p['a']
inputs = dict(y=y, prior=prior)
print(gv.fmt_values(outputs)) # tabulate outputs
print(gv.fmt_errorbudget(outputs, inputs)) # print error budget for outputs

# save best-fit values in file 'outputfile.p' for later use
```

```
import pickle
pickle.dump(fit.p, open('outputfile.p', 'wb'))
```

This code fits the function $f(x, a, b) = \exp(a + b \cdot x)$ (see `fcn(x, p)`) to two sets of data, labeled `data1` and `data2`, by varying parameters `a` and `b` until $f(x['data1'], a, b)$ and $f(x['data2'], a, b)$ equal $y['data1']$ and $y['data2']$, respectively, to within the `ys`' errors. The means and covariance matrices for the `ys` are specified in the `gv.gvar(...)`s used to create them: for example,

```
>>> print(y['data1'])
[1.376 +- 0.0685565 2.01 +- 0.236643]
>>> print(y['data1'][0].mean, "+-", y['data1'][0].sdev)
1.376 +- 0.068556546004
>>> print(gv.evalcov(y['data1']))    # covariance matrix
[[ 0.0047  0.01 ]
 [ 0.01   0.056 ]]
```

shows the means, standard deviations and covariance matrix for the data in the first data set (0.0685565 is the square root of the 0.0047 in the covariance matrix). The dictionary `prior` gives *a priori* estimates for the two parameters, `a` and `b`: each is assumed to be 0.5 ± 0.5 before fitting. The parameters `p[k]` in the fit function `fcn(x, p)` are stored in a dictionary having the same keys and layout as `prior`. In addition, there is an extra piece of input data, `y['b/a']`, which indicates that b/a is 2 ± 0.5 . The fit function for this data is simply the ratio b/a (represented by `p['b']/p['a']` in fit function `fcn(x, p)`). The fit function returns a dictionary having the same keys and layout as the input data `y`.

The output from the code sample above is:

```
Least Square Fit:
  chi2/dof [dof] = 0.17 [5]      Q = 0.97      logGBF = 0.65538      itns = 5

Parameters:
      a    0.253 (32)      [ 0.50 (50) ]
      b    0.449 (65)      [ 0.50 (50) ]

Fit:
      key          y[key]      f(p) [key]
-----
      b/a          2.00 (50)      1.78 (30)
  data1 0          1.376 (69)      1.347 (46)
        1          2.01 (24)      2.02 (16)
  data2 0          1.329 (69)      1.347 (46)
        1          1.58 (12)      1.612 (82)

Settings:
  svdcut = (1e-15,1e-15)    svdnum = (None,None)    reltol/abstol = 0.0001/0

Values:
      a: 0.253(32)
  b/a: 1.78(30)
      b: 0.449(65)

Partial % Errors:
              a          b/a          b
-----
          y:      12.75      16.72      14.30
        prior:      0.92       1.58       1.88
-----
        total:      12.78      16.80      14.42
```

The best-fit values for `a` and `b` are 0.253(32) and 0.449(65), respectively; and the best-fit result for b/a is 1.78(30),

which, because of correlations, is slightly more accurate than might be expected from the separate errors for a and b . The error budget for each of these three quantities is tabulated at the end and shows that the bulk of the error in each case comes from uncertainties in the y data, with only small contributions from uncertainties in the priors `prior`. The fit results corresponding to each piece of input data are also tabulated (Fit: ...); the agreement is excellent, as expected given that the χ^2 per degree of freedom is only 0.17.

The last section of the code uses Python's `pickle` module to save the best-fit values of the parameters in a file for later use. They are recovered using `pickle` again:

```
>>> import pickle
>>> p = pickle.load(open('outputfile.p', 'rb'))
>>> print(p['a'])
0.252798 +- 0.0323152
>>> print(p['b'])
0.448762 +- 0.0647224
>>> print(p['b']/p['a'])
1.77518 +- 0.298185
```

The recovered parameters are `gvar.GVars`, with their full covariance matrix intact. (`pickle` works here because the variables in `fit.p` are stored in a special dictionary of type `gvar.BufferDict`; `gvar.GVars` cannot be pickled otherwise.)

Note that the constraint in y on b/a in this example is much tighter than the constraints on a and b separately. This suggests a variation on the previous code, where the tight restriction on b/a is built into the prior rather than y :

```
... as before ...

y = {
    # data for the dependent variable
    'data1' : gv.gvar([1.376, 2.010], [[ 0.0047, 0.01], [ 0.01, 0.056]]),
    'data2' : gv.gvar([1.329, 1.582], [[ 0.0047, 0.0067], [0.0067, 0.0136]]))
}
x = {
    # independent variable
    'data1' : np.array([0.1, 1.0]),
    'data2' : np.array([0.1, 0.5])
}
prior = dict(a=gv.gvar(0.5, 0.5))
prior['b'] = prior['a']*gv.gvar(2.0, 0.5)

def fcn(x, p):
    # fit function of x and parameters p[k]
    ans = {}
    for k in ['data1', 'data2']:
        ans[k] = gv.exp(p['a'] + x[k]*p['b'])
    return ans

... as before ...
```

Here the dependent data y no longer has an entry for b/a , and neither do results from the fit function; but the prior for b is now 2 ± 0.5 times the prior for a , thereby introducing a correlation that limits the ratio b/a to be 2 ± 0.5 in the fit. This code gives almost identical results to the first one — very slightly less accurate, since there is less input data. We can often move information from the y data to the prior or back since both are forms of input information.

There are several things worth noting from this example:

- The input data (y) is expressed in terms of Gaussian random variables — quantities with means and a covariance matrix. These are represented by objects of type `gvar.GVar` in the code; module `gvar` has a variety of tools for creating and manipulating Gaussian random variables.
- The input data is stored in a dictionary (y) whose values can be `gvar.GVars` or arrays of `gvar.GVars`. The use of a dictionary allows for far greater flexibility than, say, an array. The fit function (`fcn(x, p)`) has to return a dictionary with the same layout as that of y (that is, with the same keys and where the value for each

key has the same shape as the corresponding value in `y`). `lsqfit` does allow `y` to be an array instead of a dictionary, which might be preferable for very simple fits (but usually not otherwise).

- The independent data (`x`) can be anything; it is simply passed through the fit code to the fit function `fcn(x, p)`. It can also be omitted altogether, in which case the fit function depends only upon the parameters: `fcn(p)`.
- The fit parameters (`p` in `fcn(x, p)`) are also stored in a dictionary whose values are `gvar.GVars` or arrays of `gvar.GVars`. Again this allows for great flexibility. The layout of the parameter dictionary is copied from that of the prior (`prior`). Again `p` can be a single array instead of a dictionary, if that simplifies the code (which is usually not the case).
- The best-fit values of the fit parameters (`fit.p[k]`) are also `gvar.GVars` and these capture statistical correlations between different parameters that are indicated by the fit. These output parameters can be combined in arithmetic expressions, using standard operators and standard functions, to obtain derived quantities. These operations take account of and track statistical correlations.
- Function `gvar.fmt_errorbudget()` is a useful tool for assessing the origins (inputs) of the statistical errors obtained in various final results (outputs). It is particularly useful for analyzing the impact of the *a priori* uncertainties encoded in the prior (`prior`).

What follows is a brief tutorial that demonstrates in greater detail how to use these modules in some standard variations on the data fitting problem. As above, code for the examples is specified completely and so can be copied into a file, and run as is. It can also be modified, allowing for experimentation. At the very end, in an appendix, there is a very simple pedagogical example that illustrates the nature of priors and demonstrates some of the simpler techniques supported by `lsqfit`.

About Printing: The examples in this tutorial use the `print` function as it is used in Python 3. Drop the outermost parenthesis in each `print` statement if using Python 2; or add

```
from __future__ import print_function
```

at the start of your file.

1.2 Making Fake Data

We need data in order to demonstrate curve fitting. The easiest route is to make fake data. The recipe is simple: 1) choose some well defined function $f(x)$ of the independent variable x ; 2) choose values for the x s, and therefore the “correct” values for $y=f(x)$; and 3) add random noise to the y s, to simulate measurement errors. Here we will work through a simple implementation of this recipe to illustrate how the `gvar` module can be used to build complicated Gaussian distributions (in this case for the correlated noise in the y s). A reader eager to fit real data can skip this section on first reading.

For the function f we choose something familiar (to some people): a sum of exponentials `sum_i=0..99 a[i] exp(-E[i]*x)`. We take as our exact values for the parameters `a[i]=0.4` and `E[i]=0.9*(i+1)`, which are easy to remember. This is simple in Python:

```
import numpy as np

def f_exact(x, nexp=100):
    return sum(0.4 * np.exp(-0.9 * (i + 1) * x) for i in range(nexp))
```

For x s we take 1, 2, 3...10, 12, 14...20, and exact y s are then given by `f_exact(x)`:

```
>>> x = array([1., 2., 3., 4., 5., 6., 7., 8., 9., 10., 12., 14., 16., 18., 20.])
>>> y_exact = f_exact(x)
>>> print(y_exact)                                # correct/exact values for y
[ 2.74047100e-01  7.92134506e-02  2.88190008e-02 ... ]
```

Finally we need to add random noise to the `y_exacts` to obtain our fit data. We do this by forming `y_exact*noise` where

```
noise = 1 + sum_n=0..99 c[n] * (x / x_max) ** n,
```

Here `x_max` is the largest `x` used, and the `c[n]` are Gaussian random variable with means and standard deviations of order 0.01. This is easy to implement in Python using the `gvar` module:

```
import gvar as gv
```

```
def make_data(nexp=100, eps=0.01): # make x, y fit data
    x = np.array([1., 2., 3., 4., 5., 6., 7., 8., 9., 10., 12., 14., 16., 18., 20.])
    cr = gv.gvar(0.0, eps)
    c = [gv.gvar(cr(), eps) for n in range(100)]
    x_max = x/max(x)
    noise = 1 + sum(c[n] * x_max ** n for n in range(100))
    y = f_exact(x, nexp) * noise
    return x, y
```

Gaussian variable `cr` represents a Gaussian distribution with mean 0.0 and width 0.01, which we use here as a random number generator: `cr()` is a number drawn randomly from the distribution represented by `cr`:

```
>>> print(cr)
0 +- 0.01
>>> print(cr())
0.00452180208286
>>> print(cr())
-0.00731564589737
```

We use `cr()` to generate mean values for the Gaussian distributions represented by the `c[n]`s, each of which has width 0.01. The resulting `ys` fluctuate around the corresponding values of `f_exact(x)` and have statistical errors:

```
>>> print(y)
[0.275179 +- 0.0027439 0.0795054 +- 0.000796125 ... ]
>>> print(y-f_exact(x))
[0.00113215 +- 0.0027439 0.000291951 +- 0.000796125 ... ]
```

Different `ys` are also correlated (by construction), which becomes clear if we evaluate the covariance matrix for the `ys`:

```
>>> print(gv.evalcov(y))
[[ 7.52900382e-06  2.18173029e-06  7.95744444e-07 ... ]
 [ 2.18173029e-06  6.33815228e-07  2.31761675e-07 ... ]
 [ 7.95744444e-07  2.31761675e-07  8.49651978e-08 ... ]
 ...
]
```

The diagonal elements of the covariance matrix are the variances of the individual `ys`; the off-diagonal elements are a measure of the correlations:

$$< (y[i] - \langle y[i] \rangle) * (y[j] - \langle y[j] \rangle) >.$$

The Gaussian variables `y[i]` together with the numbers `x[i]` comprise our fake data.

1.3 Basic Fits

Now that we have fit data, `x, y = make_data()`, we pretend ignorance of the exact functional relationship between `x` and `y` (*i.e.*, `y=f_exact(x)`). Typically we *do* know the functional form and have some *a priori* idea about the

parameter values. The point of the fit is to improve our knowledge of the parameter values, beyond our *a priori* impressions, by analyzing the fit data. Here we see how to do this using the `lsqfit` module.

First we need code to represent the fit function. In this case we know that a sum of exponentials is appropriate, so we define the following Python function to represent the relationship between x and y in our fit:

```
import numpy as np

def f(x, p):
    # function used to fit x, y data
    a = p['a']      # array of a[i]s
    E = p['E']      # array of E[i]s
    return sum(ai * np.exp(-Ei * x) for ai, Ei in zip(a, E))
```

The fit parameters, `a[i]` and `E[i]`, are stored in a dictionary, using labels `a` and `b` to access them. These parameters are varied in the fit to find the best-fit values `p=p_fit` for which `f(x, p_fit)` most closely approximates the `ys` in our fit data. The number of exponentials included in the sum is specified implicitly in this function, by the lengths of the `p['a']` and `p['E']` arrays.

Next we need to define priors that encapsulate our *a priori* knowledge about the parameter values. In practice we almost always have *a priori* knowledge about parameters; it is usually impossible to design a fit function without some sense of the parameter sizes. Given such knowledge it is important (usually essential) to include it in the fit. This is done by designing priors for the fit, which are probability distributions for each parameter that describe the *a priori* uncertainty in that parameter. As in the previous section, we use objects of type `gvar.GVar` to describe (Gaussian) probability distributions. Let's assume that before the fit we suspect that each `a[i]` is of order 0.5 ± 0.5 , while `E[i]` is of order $1+i \pm 0.5$. A prior that represents this information is built using the following code:

```
import lsqfit
import gvar as gv

def make_prior(nexp):
    # make priors for fit parameters
    prior = gv.BufferDict()      # prior -- any dictionary works
    prior['a'] = [gv.gvar(0.5, 0.5) for i in range(nexp)]
    prior['E'] = [gv.gvar(i+1, 0.5) for i in range(nexp)]
    return prior
```

where `nexp` is the number of exponential terms that will be used (and therefore the number of `as` and `Es`). With `nexp=3`, for example, one would then have:

```
>>> print(prior['a'])
[0.5 +- 0.5 0.5 +- 0.5 0.5 +- 0.5]
>>> print(prior['E'])
[1 +- 0.5 2 +- 0.5 3 +- 0.5]
```

We use dictionary-like class `gvar.BufferDict` for the prior because it allows us to save the prior if we wish (using Python's `pickle` module). If saving is unnecessary, `gvar.BufferDict` can be replaced by `dict()` or most any other Python dictionary class.

With fit data, a fit function, and a prior for the fit parameters, we are finally ready to do the fit, which is now easy:

```
fit = lsqfit.nonlinear_fit(data=(x, y), fcn=f, prior=prior)
```

So pulling together the entire code, from this section and the previous one, our complete Python program for making fake data and fitting it is:

```
import lsqfit
import numpy as np
import gvar as gv

def f_exact(x, nexp=100):
    # exact f(x)
    return sum(0.4*np.exp(-0.9*(i+1)*x) for i in range(nexp))
```

```

def f(x, p):
    a = p['a']
    E = p['E']
    return sum(ai * np.exp(-Ei * x) for ai, Ei in zip(a, E))

def make_data(nexp=100, eps=0.01): # make x, y fit data
    x = np.array([1., 2., 3., 4., 5., 6., 7., 8., 9., 10., 12., 14., 16., 18., 20.])
    cr = gv.gvar(0.0, eps)
    c = [gv.gvar(cr(), eps) for n in range(100)]
    x_xmax = x/max(x)
    noise = 1+ sum(c[n] * x_xmax ** n for n in range(100))
    y = f_exact(x, nexp) * noise
    return x, y

def make_prior(nexp):
    prior = gv.BufferDict()
    prior['a'] = [gv.gvar(0.5, 0.5) for i in range(nexp)]
    prior['E'] = [gv.gvar(i+1, 0.5) for i in range(nexp)]
    return prior

def main():
    gv.ranseed([2009, 2010, 2011, 2012]) # initialize random numbers (opt.)
    x, y = make_data() # make fit data
    p0 = None # make larger fits go faster (opt.)
    for nexp in range(3, 20):
        print('***** nexp =', nexp)
        prior = make_prior(nexp)
        fit = lsqfit.nonlinear_fit(data=(x, y), fcn=f, prior=prior, p0=p0)
        print(fit) # print the fit results
        E = fit.p['E'] # best-fit parameters
        a = fit.p['a']
        print('E1/E0 =', E[1] / E[0], ' E2/E0 =', E[2] / E[0])
        print('a1/a0 =', a[1] / a[0], ' a2/a0 =', a[2] / a[0])
        print()
        if fit.chi2 / fit.dof < 1.:
            p0 = fit.pmean # starting point for next fit (opt.)

if __name__ == '__main__':
    main()

```

We are not sure *a priori* how many exponentials are needed to fit our data; given that there are only fifteen *y*s, and these are noisy, there may only be information in the data about the first few terms. Consequently we wrote our code to try fitting with each of *nexp*=3, 4, 5...19 terms. (The pieces of the code involving *p0* are optional; they make the more complicated fits go about 30 times faster since the output from one fit is used as the starting point for the next fit — see the discussion of the *p0* parameter for `lsqfit.nonlinear_fit`.) Running this code produces the following output, which is reproduced here in some detail in order to illustrate a variety of features:

```

***** nexp = 3
Least Square Fit:
  chi2/dof [dof] = 6.3e+02 [15]    Q = 0    logGBF = -4465.1    itns = 30

Parameters:
  a 0    0.0288 (11)    [ 0.50 (50) ]
    1    0.0354 (13)    [ 0.50 (50) ]
    2    0.0779 (30)    [ 0.50 (50) ]
  E 0    1.0107 (24)    [ 1.00 (50) ]
    1    2.0200 (27)    [ 2.00 (50) ]
    2    3.6643 (33)    [ 3.00 (50) ] *

```

```
Settings:
  svdcut = (1e-15,1e-15)   svdnum = (None,None)   reltol/abstol = 0.0001/0
```

```
E1/E0 = 1.99859 +- 0.00239496   E2/E0 = 3.62551 +- 0.00618821
a1/a0 = 1.2313 +- 0.000474021   a2/a0 = 2.707 +- 0.00130172
```

```
***** nexexp = 4
```

```
Least Square Fit:
```

```
chi2/dof [dof] = 0.57 [15]   Q = 0.9   logGBF = 220.04   itns = 220
```

```
Parameters:
```

```
  a 0    0.4018 (40)   [ 0.50 (50) ]
    1    0.4055 (42)   [ 0.50 (50) ]
    2    0.4952 (76)   [ 0.50 (50) ]
    3    1.124 (12)    [ 0.50 (50) ] *
  E 0    0.90037 (51)  [ 1.00 (50) ]
    1    1.8023 (13)   [ 2.00 (50) ]
    2    2.7731 (90)   [ 3.00 (50) ]
    3    4.383 (21)    [ 4.00 (50) ]
```

```
Settings:
```

```
svdcut = (1e-15,1e-15)   svdnum = (None,None)   reltol/abstol = 0.0001/0
```

```
E1/E0 = 2.00179 +- 0.00120727   E2/E0 = 3.07996 +- 0.00977469
a1/a0 = 1.00942 +- 0.00299114   a2/a0 = 1.23251 +- 0.0140225
```

```
***** nexexp = 5
```

```
Least Square Fit:
```

```
chi2/dof [dof] = 0.45 [15]   Q = 0.97   logGBF = 220.84   itns = 6
```

```
Parameters:
```

```
  a 0    0.4018 (40)   [ 0.50 (50) ]
    1    0.4049 (44)   [ 0.50 (50) ]
    2    0.478 (26)    [ 0.50 (50) ]
    3    0.63 (28)     [ 0.50 (50) ]
    4    0.62 (35)     [ 0.50 (50) ]
  E 0    0.90036 (51)  [ 1.00 (50) ]
    1    1.8019 (15)   [ 2.00 (50) ]
    2    2.759 (22)    [ 3.00 (50) ]
    3    4.09 (26)     [ 4.00 (50) ]
    4    4.95 (48)     [ 5.00 (50) ]
```

```
Settings:
```

```
svdcut = (1e-15,1e-15)   svdnum = (None,None)   reltol/abstol = 0.0001/0
```

```
E1/E0 = 2.00133 +- 0.00142399   E2/E0 = 3.06486 +- 0.024111
a1/a0 = 1.00754 +- 0.00422523   a2/a0 = 1.18874 +- 0.0632288
```

```
***** nexexp = 6
```

```
Least Square Fit:
```

```
chi2/dof [dof] = 0.45 [15]   Q = 0.97   logGBF = 220.7   itns = 6
```

```
Parameters:
```

```
  a 0    0.4018 (40)   [ 0.50 (50) ]
    1    0.4041 (47)   [ 0.50 (50) ]
    2    0.461 (41)    [ 0.50 (50) ]
    3    0.60 (24)     [ 0.50 (50) ]
    4    0.47 (37)     [ 0.50 (50) ]
```

```

      5      0.45 (46)      [ 0.50 (50) ]
E 0    0.90035 (51)      [ 1.00 (50) ]
      1      1.8015 (17)   [ 2.00 (50) ]
      2      2.746 (34)    [ 3.00 (50) ]
      3      3.98 (32)     [ 4.00 (50) ]
      4      4.96 (49)     [ 5.00 (50) ]
      5      6.01 (50)     [ 6.00 (50) ]

```

Settings:

```
svdcut = (1e-15,1e-15)   svdnum = (None,None)   reltol/abstol = 0.0001/0
```

```
E1/E0 = 2.00084 +- 0.00168855   E2/E0 = 3.04939 +- 0.0374116
```

```
a1/a0 = 1.00551 +- 0.00561391   a2/a0 = 1.14607 +- 0.101175
```

```
***** nexpt = 7
```

Least Square Fit:

```
chi2/dof [dof] = 0.45 [15]   Q = 0.96   logGBF = 220.6   itns = 6
```

Parameters:

```

a 0    0.4018 (40)      [ 0.50 (50) ]
      1    0.4036 (48)   [ 0.50 (50) ]
      2    0.452 (47)    [ 0.50 (50) ]
      3    0.60 (22)     [ 0.50 (50) ]
      4    0.42 (37)     [ 0.50 (50) ]
      5    0.42 (46)     [ 0.50 (50) ]
      6    0.46 (49)     [ 0.50 (50) ]
E 0    0.90035 (51)     [ 1.00 (50) ]
      1    1.8012 (18)   [ 2.00 (50) ]
      2    2.739 (39)    [ 3.00 (50) ]
      3    3.94 (33)     [ 4.00 (50) ]
      4    4.96 (49)     [ 5.00 (50) ]
      5    6.02 (50)     [ 6.00 (50) ]
      6    7.02 (50)     [ 7.00 (50) ]

```

Settings:

```
svdcut = (1e-15,1e-15)   svdnum = (None,None)   reltol/abstol = 0.0001/0
```

```
E1/E0 = 2.00059 +- 0.00181902   E2/E0 = 3.04178 +- 0.0431729
```

```
a1/a0 = 1.0045 +- 0.00626848   a2/a0 = 1.1258 +- 0.116336
```

```

.
.
.

```

```
***** nexpt = 19
```

Least Square Fit:

```
chi2/dof [dof] = 0.46 [15]   Q = 0.96   logGBF = 220.52   itns = 1
```

Parameters:

```

a 0    0.4018 (40)      [ 0.50 (50) ]
      1    0.4033 (49)   [ 0.50 (50) ]
      2    0.447 (51)    [ 0.50 (50) ]
      3    0.60 (21)     [ 0.50 (50) ]
      4    0.38 (37)     [ 0.50 (50) ]
      5    0.40 (46)     [ 0.50 (50) ]
      6    0.45 (49)     [ 0.50 (50) ]
      7    0.48 (50)     [ 0.50 (50) ]
      8    0.49 (50)     [ 0.50 (50) ]
      9    0.50 (50)     [ 0.50 (50) ]

```

```
10      0.50 (50)      [ 0.50 (50) ]
11      0.50 (50)      [ 0.50 (50) ]
12      0.50 (50)      [ 0.50 (50) ]
13      0.50 (50)      [ 0.50 (50) ]
14      0.50 (50)      [ 0.50 (50) ]
15      0.50 (50)      [ 0.50 (50) ]
16      0.50 (50)      [ 0.50 (50) ]
17      0.50 (50)      [ 0.50 (50) ]
18      0.50 (50)      [ 0.50 (50) ]
E 0      0.90035 (51)  [ 1.00 (50) ]
1      1.8011 (19)    [ 2.00 (50) ]
2      2.734 (42)    [ 3.00 (50) ]
3      3.91 (33)     [ 4.00 (50) ]
4      4.97 (49)     [ 5.00 (50) ]
5      6.02 (50)     [ 6.00 (50) ]
6      7.02 (50)     [ 7.00 (50) ]
7      8.01 (50)     [ 8.00 (50) ]
8      9.00 (50)     [ 9.00 (50) ]
9     10.00 (50)     [ 10.00 (50) ]
10     11.00 (50)     [ 11.00 (50) ]
11     12.00 (50)     [ 12.00 (50) ]
12     13.00 (50)     [ 13.00 (50) ]
13     14.00 (50)     [ 14.00 (50) ]
14     15.00 (50)     [ 15.00 (50) ]
15     16.00 (50)     [ 16.00 (50) ]
16     17.00 (50)     [ 17.00 (50) ]
17     18.00 (50)     [ 18.00 (50) ]
18     19.00 (50)     [ 19.00 (50) ]
```

Settings:

```
svdcut = (1e-15,1e-15)   svdnum = (None,None)   reltol/abstol = 0.0001/0
```

```
E1/E0 = 2.00041 +- 0.00190442   E2/E0 = 3.0363 +- 0.046784
a1/a0 = 1.00376 +- 0.00668978   a2/a0 = 1.11149 +- 0.125413
```

----- fit with extra information

There are several things to notice here:

- Clearly three exponentials (`nexp=3`) is not enough. The `chi**2` per degree of freedom (`chi2/dof`) is much larger than one. The `chi**2` improves significantly for `nexp=4` exponentials and by `nexp=6` the fit is as good as it is going to get — there is essentially no change when further exponentials are added.
- The best-fit values for each parameter are listed for each of the fits, together with the prior values (in brackets, on the right). Values for each `a[i]` and `E[i]` are listed in order, starting at the points indicated by the labels `a` and `E`. Asterisks are printed at the end of the line if the mean best-fit value differs from the prior's mean by more than one standard deviation; the number of asterisks, up to a maximum of 5, indicates how many standard deviations the difference is. Differences of one or two standard deviations are not uncommon; larger differences could indicate a problem with the prior or the fit.

Once the fit converges, the best-fit values for the various parameters agree well — that is to within their errors, approximately — with the exact values, which we know since we are using fake data. For example, `a` and `E` for the first exponential are 0.402(4) and 0.9003(5), respectively, from the fit where the exact answers are 0.4 and 0.9; and we get 0.45(5) and 2.73(4) for the third exponential where the exact values are 0.4 and 2.7.

- Note in the `nexp=7` fit how the means and standard deviations for the parameters governing the seventh (and last) exponential are almost identical to the values in the corresponding priors: 0.46(49) from the fit for `a` and 7.0(5) for `E`. This tells us that our fit data has little or no information to add to what we knew *a priori* about these parameters — there isn't enough data and what we have isn't accurate enough.

This situation is truer still of further terms as they are added in the `nexp=8` and later fits. This is why the fit results stop changing once we have `nexp=6` exponentials. There is no point in including further exponentials, beyond the need to verify that the fit has indeed converged.

- The last fit includes `nexp=19` exponentials and therefore has 38 parameters. This is in a fit to 15 `ys`. Old-fashioned fits, without priors, are impossible when the number of parameters exceeds the number of data points. That is clearly not the case here, where the number of terms and parameters can be made arbitrarily large, eventually (after `nexp=6` terms) with no effect at all on the results.

The reason is that the prior that we include for each new parameter is, in effect, a new piece of data (the mean and standard deviation of the *a priori* expectation for that parameter); it leads to a new term in the `chi**2` function. We are fitting both the data and our *a priori* expectations for the parameters. So in the `nexp=19` fit, for example, we actually have 53 pieces of data to fit: the 15 `ys` plus the 38 prior values for the 38 parameters.

The effective number of degrees of freedom (`dof` in the output above) is the number of pieces of data minus the number of fit parameters, or $53-38=15$ in this last case. With priors for every parameter, the number of degrees of freedom is always equal to the number of `ys`, irrespective of how many fit parameters there are.

- The Gaussian Bayes Factor (whose logarithm is `logGBF` in the output) is a measure of the likelihood that the actual data being fit could have come from a theory with the prior and fit function used in the fit. The larger this number, the more likely it is that prior/fit-function and data could be related. Here it grows dramatically from the first fit (`nexp=3`) but then more-or-less stops changing around `nexp=6`. The implication is that this data is much more likely to have come from a theory with `nexp>=6` than with `nexp=3` (which we know to be the actual case).
- In the code, results for each fit are captured in a Python object `fit`, which is of type `lsqfit.nonlinear_fit`. A summary of the fit information is obtained by printing `fit`. Also the best-fit results for each fit parameter can be accessed through `fit.p`, as is done here to calculate various ratios of parameters.

The errors in these last calculations automatically account for any correlations in the statistical errors for different parameters. This is obvious in the ratio `a1/a0`, which would be 1.004(16) if there was no statistical correlation between our estimates for `a1` and `a0`, but in fact is 1.004(7) in this fit. The (positive) correlation is evident in the covariance matrix:


```
>>> print(gv.evalcov([a[0], a[1]]))
[[ 1.61726195e-05  1.65492001e-05]
 [ 1.65492001e-05  2.41547633e-05]]
```

Finally we inspect the fit's quality point by point. The input data are compared with results from the fit function, evaluated with the best-fit parameters, in the following table (obtained in the code by printing the output from `fit.format(100)`):

```
Fit:
      x[k]          y[k]          f(x[k],p)
-----
      1         0.2752 (27)         0.2752 (20)
      2         0.07951 (80)        0.07952 (58)
      3         0.02891 (29)        0.02892 (21)
      4         0.01127 (11)        0.011272 (83)
      5         0.004502 (46)        0.004506 (34)
      6         0.001817 (19)        0.001819 (14)
      7         0.0007362 (79)       0.0007375 (57)
      8         0.0002987 (33)       0.0002994 (24)
      9         0.0001213 (14)       0.00012163 (99)
     10        0.00004926 (57)       0.00004943 (41)
     12         8.13(10)e-06         8.164(72)e-06
     14         1.342(19)e-06        1.348(13)e-06
     16         2.217(37)e-07        2.227(23)e-07
```

18	3.661(85)e-08	3.679(40)e-08
20	6.24(61)e-09	6.078(71)e-09

The fit is excellent over the entire eight orders of magnitude. This information is presented again in the following plot, which shows the ratio $y/f(x, p)$, as a function of x , using the best-fit parameters p . The correct result for this ratio, of course, is one. The smooth variation in the data — smooth compared with the size of the statistical-error bars — is an indication of the statistical correlations between individual y s.



This particular plot was made using the `matplotlib` module, with the following code added to the end of `main()` (outside the loop):

```
import pylab as plt
ratio = y / f(x, fit.pmean)
plt.xlim(0, 21)
plt.xlabel('x')
plt.ylabel('y/f(x,p)')
plt.errorbar(x=x, y=gv.mean(ratio), yerr=gv.sdev(ratio), fmt='ob')
plt.plot([0.0, 21.0], [1.0, 1.0])
plt.show()
```

1.4 Chained Fits

The priors in a fit represent knowledge that we have about the parameters before we do the fit. This knowledge might come from theoretical considerations or experiment. Or it might come from another fit. Imagine that we want to add new information to that extracted from the fit in the previous section. For example, we might learn from some other source that the ratio of amplitudes $a[1]/a[0]$ equals $1 \pm 1e-5$. The challenge is to combine this new information with information extracted from the fit above without rerunning that fit. (We assume it is not possible to rerun the first fit, because, say, the input data for that fit has been lost or is unavailable.)

We can combine the new data with the old fit results by creating a new fit using the best-fit parameters, `fit.p`,

from the old fit as the priors for the new fit. To try this out, we add the following code onto the end of the `main()` subroutine in the previous section:

```
def ratio(p):                                # new fit function
    a = p['a']
    return a[1] / a[0]

prior = fit.p                                # prior = best-fit parameters from 1st fit
data = gv.gvar(1, 1e-5)                     # new data for the ratio

newfit = lsqfit.nonlinear_fit(data=data, fcn=ratio, prior=prior)
print(newfit)
```

The result of the new fit (to one piece of new data) is:

Least Square Fit:

chi2/dof [dof] = 0.32 [1] Q = 0.57 logGBF = 3.9303 itns = 2

Parameters:

a 0	0.4018 (40)	[0.4018 (40)]
1	0.4018 (40)	[0.4033 (49)]
2	0.421 (20)	[0.447 (51)]
3	0.53 (17)	[0.60 (21)]
4	0.46 (34)	[0.38 (37)]
5	0.50 (42)	[0.40 (46)]
6	0.50 (48)	[0.45 (49)]
7	0.50 (50)	[0.48 (50)]
8	0.50 (50)	[0.49 (50)]
9	0.50 (50)	[0.50 (50)]
10	0.50 (50)	[0.50 (50)]
11	0.50 (50)	[0.50 (50)]
12	0.50 (50)	[0.50 (50)]
13	0.50 (50)	[0.50 (50)]
14	0.50 (50)	[0.50 (50)]
15	0.50 (50)	[0.50 (50)]
16	0.50 (50)	[0.50 (50)]
17	0.50 (50)	[0.50 (50)]
18	0.50 (50)	[0.50 (50)]
E 0	0.90030 (51)	[0.90035 (51)]
1	1.80007 (67)	[1.8011 (19)]
2	2.711 (12)	[2.734 (42)]
3	3.76 (18)	[3.91 (33)]
4	5.02 (48)	[4.97 (49)]
5	6.00 (50)	[6.02 (50)]
6	7.00 (50)	[7.02 (50)]
7	8.00 (50)	[8.01 (50)]
8	9.00 (50)	[9.00 (50)]
9	10.00 (50)	[10.00 (50)]
10	11.00 (50)	[11.00 (50)]
11	12.00 (50)	[12.00 (50)]
12	13.00 (50)	[13.00 (50)]
13	14.00 (50)	[14.00 (50)]
14	15.00 (50)	[15.00 (50)]
15	16.00 (50)	[16.00 (50)]
16	17.00 (50)	[17.00 (50)]
17	18.00 (50)	[18.00 (50)]
18	19.00 (50)	[19.00 (50)]

Settings:

```
svdcut = (1e-15,1e-15)    svdnum = (None,None)    reltol/abstol = 0.0001/0
```

Parameters $a[0]$ and $E[0]$ are essentially unchanged by the new information, but $a[i]$ and $E[i]$ are much more precise for $i=2$ and $i=3$. It might seem odd that $E[1]$, for example, is changed at all, since the fit function, `ratio(p)`, makes no mention of it. This is not surprising, however, since `ratio(p)` does depend up $a[1]$, and $a[1]$ is strongly correlated with $E[1]$ through the prior. It is important to include all parameters from the first fit as parameters in the new fit in order to capture the impact of the new information on parameters correlated with $a[1]/a[0]$.

It would have been easy to change the fit code in the previous section to incorporate the new information about $a[1]/a[0]$. The approach presented here is numerically equivalent to that approach insofar as the `chi**2` function for the original fit can be well approximated by a quadratic function in the fit parameters — that is, insofar as $\exp(-\text{chi}^2/2)$ is well approximated by a Gaussian distribution in the parameters, as specified by the best-fit means and covariance matrix (in `fit.p`). This is, of course, a fundamental assumption underlying the use of `lsqfit` in the first place.

Obviously, we can include further fits in order to incorporate more data. The prior for each new fit is the best-fit output (`fit.p`) from the previous fit. The output from the chain's final fit is the cumulative result of all of these fits.

1.5 x has Error Bars

We now consider variations on our basic fit analysis (described above). The first variation concerns what to do when the independent variables, the *xs*, have errors, as well as the *ys*. This is easily handled by turning the *xs* into fit parameters, and otherwise dispensing with independent variables.

To illustrate this, we modify the basic analysis code above. First we need to add errors to the *xs*, which we do by changing `make_data` so that each *x* has a random value within about $\pm 0.001\%$ of its original value and an error:

```
def make_data(nexp=100, eps=0.01): # make x, y fit data
    x = np.array([1.,2.,3.,4.,5.,6.,7.,8.,9.,10.,12.,14.,16.,18.,20.])
    cr = gv.gvar(0.0, eps)
    c = [gv.gvar(cr(), eps) for n in range(100)]
    x_xmax = x/max(x)
    noise = 1+ sum(c[n] * x_xmax ** n for n in range(100))
    y = f_exact(x, nexp) * noise # noisy y[i]s
    xfac = gv.gvar(1.0, 0.00001) # Gaussian distrib'n: 1±0.0001%
    x = np.array([xi * gv.gvar(xfac(), xfac.sdev) for xi in x]) # noisy x[i]s
    return x, y
```

Here `gvar.GVar` object `xfac` is used as a random number generator: each time it is called, `xfac()` is a different random number from the distribution with mean `xfac.mean` and standard deviation `xfac.sdev` (that is, 1 ± 0.00001). The main program is modified so that the (now random) *x* array is treated as a fit parameter. The prior for each *x* is, obviously, specified by the mean and standard deviation of that *x*, which is read directly out of the array of *xs* produced by `make_data()`:

```
def make_prior(nexp, x): # make priors for fit parameters
    prior = gv.BufferDict() # prior -- any dictionary works
    prior['a'] = [gv.gvar(0.5, 0.5) for i in range(nexp)]
    prior['E'] = [gv.gvar(i+1, 0.5) for i in range(nexp)]
    prior['x'] = x # x now an array of parameters
    return prior

def main():
    gv.ranseed([2009, 2010, 2011, 2012]) # initialize random numbers (opt.)
    x, y = make_data() # make fit data
    p0 = None # make larger fits go faster (opt.)
```

```

for nexp in range(3, 20):
    print(' ***** nexp =', nexp)
    prior = make_prior(nexp, x)
    fit = lsqfit.nonlinear_fit(data=y, fcn=f, prior=prior, p0=p0)
    print(fit)                                # print the fit results
    E = fit.p['E']                             # best-fit parameters
    a = fit.p['a']
    print('E1/E0 =', E[1] / E[0], ' E2/E0 =', E[2] / E[0])
    print('a1/a0 =', a[1] / a[0], ' a2/a0 =', a[2] / a[0])
    print()
    if fit.chi2/fit.dof<1.:
        p0 = fit.pmean                        # starting point for next fit (opt.)

```

The fit data now consists of just the y array (data=y), and the fit function loses its x argument and gets its x values from the fit parameters p instead:

```

def f(p):
    a = p['a']
    E = p['E']
    x = p['x']
    return sum(ai*exp(-Ei*x) for ai, Ei in zip(a, E))

```

Running the new code gives, for nexp=6 terms:

```

***** nexp = 6
Least Square Fit:
  chi2/dof [dof] = 0.54 [15]      Q = 0.92      logGBF = 198.93      itns = 6

```

Parameters:

a 0	0.4025 (41)	[0.50 (50)]
1	0.429 (32)	[0.50 (50)]
2	0.58 (23)	[0.50 (50)]
3	0.40 (38)	[0.50 (50)]
4	0.42 (46)	[0.50 (50)]
5	0.46 (49)	[0.50 (50)]
E 0	0.90068 (60)	[1.00 (50)]
1	1.818 (20)	[2.00 (50)]
2	2.95 (28)	[3.00 (50)]
3	3.98 (49)	[4.00 (50)]
4	5.02 (50)	[5.00 (50)]
5	6.01 (50)	[6.00 (50)]
x 0	0.999997 (10)	[0.999997 (10)]
1	1.999958 (20)	[1.999958 (20)]
2	3.000014 (30)	[3.000013 (30)]
3	4.000065 (36)	[4.000064 (40)]
4	5.000047 (34)	[5.000069 (50)]
5	6.000020 (39)	[5.999986 (60)]
6	6.999988 (40)	[6.999942 (70)]
7	7.999956 (42)	[7.999982 (80)]
8	8.999934 (50)	[9.000054 (90)]
9	9.999923 (59)	[9.99991 (10)]
10	11.999929 (79)	[11.99982 (12)]
11	13.99992 (11)	[13.99991 (14)]
12	15.99992 (15)	[15.99998 (16)]
13	18.00022 (18)	[18.00020 (18)]
14	20.00016 (20)	[20.00016 (20)]

Settings:

```

svdcut = (1e-15,1e-15)    svdnum = (None,None)    reltol/abstol = 0.0001/0

```

```
E1/E0 = 2.01801 +- 0.0216404    E2/E0 = 3.27385 +- 0.305544
a1/a0 = 1.06515 +- 0.0765488    a2/a0 = 1.4485 +- 0.573896
```

This looks quite a bit like what we obtained before, except that now there are 15 more parameters, one for each x , and also now all results are a good deal less accurate. Note that one result from this analysis is new values for the x s. In some cases the errors on the x values have been reduced — by information in the fit data.

1.6 Correlated Parameters; Gaussian Bayes Factor

`gvar.GVar` objects are very useful for handling more complicated priors, including situations where we know *a priori* of correlations between parameters. Returning to the *Basic Fits* example above, imagine a situation where we still have a ± 0.5 uncertainty about the value of any individual $E[i]$, but we know *a priori* that the separations between adjacent $E[i]$ s is 0.9 ± 0.01 . We want to build the correlation between adjacent $E[i]$ s into our prior.

We do this by introducing a `gvar.GVar` object `de[i]` for each separate difference $E[i] - E[i-1]$, with `de[0]` being $E[0]$:

```
de = [gvar(0.9, 0.01) for i in range(nexp)]
de[0] = gvar(1, 0.5)      # different distribution for E[0]
```

Then `de[0]` specifies the probability distribution for $E[0]$, `de[0]+de[1]` the distribution for $E[1]$, `de[0]+de[1]+de[2]` the distribution for $E[2]$, and so on. This can be implemented (slightly inefficiently) in a single line of Python:

```
E = [sum(de[:i+1]) for i in range(nexp)]
```

For `nexp=3`, this implies that

```
>>> print(E)
[1 +- 0.5 1.9 +- 0.5001 2.8 +- 0.5002]
>>> print(E[1] - E[0], E[2] - E[1])
0.9 +- 0.01 0.9 +- 0.01
```

which shows that each $E[i]$ separately has an uncertainty of ± 0.5 (approximately) but that differences are specified to within ± 0.01 .

In the code, we need only change the definition of the prior in order to introduce these correlations:

```
def make_prior(nexp):
    # make priors for fit parameters
    prior = gv.BufferDict()
    # prior -- any dictionary works
    prior['a'] = [gv.gvar(0.5, 0.5) for i in range(nexp)]
    de = [gv.gvar(0.9, 0.01) for i in range(nexp)]
    de[0] = gv.gvar(1, 0.5)
    prior['E'] = [sum(de[: i + 1]) for i in range(nexp)]
    return prior
```

Running the code as before, but now with the correlated prior in place, we obtain the following fit with `nexp=7` terms:

```
***** nexp = 7
Least Square Fit:
    chi2/dof [dof] = 0.44 [15]    Q = 0.97    logGBF = 227.47    itns = 3

Parameters:
    a 0    0.4018 (40)    [ 0.50 (50) ]
      1    0.4016 (42)    [ 0.50 (50) ]
      2    0.404 (12)    [ 0.50 (50) ]
      3    0.394 (46)    [ 0.50 (50) ]
```

```

4      0.40 (16)      [ 0.50 (50) ]
5      0.51 (31)      [ 0.50 (50) ]
6      0.52 (42)      [ 0.50 (50) ]
E 0    0.90032 (51)    [ 1.00 (50) ]
1      1.8001 (11)     [ 1.90 (50) ]
2      2.701 (10)     [ 2.80 (50) ]
3      3.601 (14)     [ 3.70 (50) ]
4      4.501 (17)     [ 4.60 (50) ]
5      5.401 (20)     [ 5.50 (50) ]
6      6.301 (22)     [ 6.40 (50) ]

```

Settings:

```
svdcut = (1e-15, None)   svdnum = (None, None)   reltol/abstol = 0.0001/0
```

```
E1/E0 = 1.9994 +- 0.00106276   E2/E0 = 2.99989 +- 0.0110965
a1/a0 = 0.999601 +- 0.00254531   a2/a0 = 1.00498 +- 0.0280721
```

The results are similar to before for the leading parameters, but substantially more accurate for parameters describing the second and later exponential terms, as might be expected given our enhanced knowledge about the differences between $E[i]$ s. The output energy differences are particularly accurate: they range from $E[1]-E[0] = 0.900(1)$, which is ten times more precise than the prior, to $E[6]-E[5] = 0.900(10)$, which is just what was put into the fit through the prior (the fit data adds no new information). The correlated prior allows us to merge our *a priori* information about the energy differences with the new information carried by the fit data x, y .

Note that the Gaussian Bayes Factor (see `logGBF` in the output) is significantly larger with the correlated prior ($\log\text{GBF} = 227$) than it was for the uncorrelated prior ($\log\text{GBF} = 221$). Had we been uncertain as to which prior was more appropriate, this difference says that the data prefers the correlated prior. (More precisely, it says that we would be $\exp(227-221) = 400$ times more likely to get this data from a theory with the correlated prior than from one with the uncorrelated prior.) This difference is significant despite the fact that the χ^2 s in the two cases are almost the same. χ^2 tests goodness of fit, but there are usually more ways than one to get a good fit. Some are more plausible than others, and the Bayes factor helps sort out which.

1.7 Tuning Priors and the Empirical Bayes Criterion

Given two choices of prior for a parameter, the one that results in a larger Gaussian Bayes Factor after fitting (see `logGBF` in fit output or `fit.logGBF`) is the one preferred by the data. We can use this fact to tune a prior or set of priors in situations where we are uncertain about the correct *a priori* value: we vary the widths and/or central values of the priors of interest to maximize `logGBF`. This leads to complete nonsense if it is applied to all the priors, but it is useful for tuning (or testing) limited subsets of the priors when other information is unavailable. In effect we are using the data to get a feel for what is a reasonable prior. This procedure for setting priors is called the *Empirical Bayes* method.

This method is implemented in a driver program

```
fit, z = lsqfit.empbayes_fit(z0, fitargs)
```

which varies numpy array `z`, starting at `z0`, to maximize `fit.logGBF` where

```
fit = lsqfit.nonlinear_fit(**fitargs(z)).
```

Function `fitargs(z)` returns a dictionary containing the arguments for `nonlinear_fit()`. These arguments, and the prior in particular, are varied as some function of `z`. The optimal fit (that is, the one for which `fit.logGBF` is maximum) and `z` are returned.

To illustrate, consider tuning the widths of the priors for the amplitudes, `prior['a']`, in the example from the previous section. This is done by adding the following code to the end of `main()` subroutine:

```
def fitargs(z, nexp=nexp, prior=prior, f=f, data=(x, y), p0=p0):
    z = np.exp(z)
    prior['a'] = [gv.gvar(0.5, 0.5 * z[0]) for i in range(nexp)]
    return dict(prior=prior, data=data, fcn=f, p0=p0)

##
z0 = [0.0]
fit, z = empbayes_fit(z0, fitargs, tol=1e-3)
print(fit)                # print the optimized fit results
E = fit.p['E']             # best-fit parameters
a = fit.p['a']
print('E1/E0 =', E[1] / E[0], ' E2/E0 =', E[2] / E[0])
print('a1/a0 =', a[1] / a[0], ' a2/a0 =', a[2] / a[0])
print("prior['a'] =", fit.prior['a'][0])
print()
```

Function `fitargs` generates a dictionary containing the arguments for `lsqfit.nonlinear_fit`. These are identical to what we have been using except that the width of the priors in `prior['a']` is adjusted according to parameter `z`. Function `lsqfit.empbayes_fit()` does fits for different values of `z` and selects the `z` that maximizes `fit.logGBF`. It returns the corresponding fit and the value of `z`.

This code generates the following output when `nexp=7`:

```
Least Square Fit:
  chi2/dof [dof] = 0.77 [15]      Q = 0.71      logGBF = 233.98      itns = 1

Parameters:
      a 0      0.4026 (40)      [ 0.500 (95) ] *
        1      0.4025 (41)      [ 0.500 (95) ] *
        2      0.4071 (80)      [ 0.500 (95) ]
        3      0.385 (20)       [ 0.500 (95) ] *
        4      0.431 (58)       [ 0.500 (95) ]
        5      0.477 (74)       [ 0.500 (95) ]
        6      0.493 (89)       [ 0.500 (95) ]
      E 0      0.90031 (50)     [ 1.00 (50) ]
        1      1.8000 (10)      [ 1.90 (50) ]
        2      2.7023 (86)      [ 2.80 (50) ]
        3      3.603 (14)       [ 3.70 (50) ]
        4      4.503 (17)       [ 4.60 (50) ]
        5      5.403 (19)       [ 5.50 (50) ]
        6      6.303 (22)       [ 6.40 (50) ]

Settings:
  svdcut = (1e-15,1e-15)      svdnum = (None,None)      reltol/abstol = 0.0001/0

E1/E0 = 1.99934 +- 0.0010096      E2/E0 = 3.0015 +- 0.00944816
a1/a0 = 0.999537 +- 0.00248687      a2/a0 = 1.01093 +- 0.0168333
prior['a'] = 0.5 +- 0.0952405
```

Reducing the width of the `prior['a']`s from 0.5 to 0.1 increased `logGBF` from 227 to 234. The error for `a2/a0` is 40% smaller, but the other results are not much affected — suggesting that the details of `prior['a']` are not all that important, which is confirmed by the error budgets generated in the next section. It is not surprising, of course, that the optimal width is 0.1 since the mean values for the `fit.p['a']`s are clustered around 0.4, which is 0.1 below the mean value of the priors `prior['a']`.

The Bayes factor, `exp(fit.logGBF)`, is useful for deciding about fit functions as well as priors. Consider the following two fits of the sort discussed in the previous section, one using just two terms in the fit function and one using three terms:


```

***** nexp = 2
Least Square Fit:
  chi2/dof [dof] = 0.47 [15]    Q = 0.96    logGBF = 254.15    itns = 6

Parameters:
      a 0      0.4018 (40)      [ 0.50 (50) ]
      1      0.4018 (40)      [ 0.50 (50) ]
      E 0      0.90036 (50)     [ 1.00 (50) ]
      1      1.80036 (50)     [ 1.90 (50) ]

Settings:
  svdcut = (1e-15,1e-15)    svdnum = (None,None)    reltol/abstol = 0.0001/0

***** nexp = 3
Least Square Fit:
  chi2/dof [dof] = 0.5 [15]    Q = 0.94    logGBF = 243.12    itns = 4

Parameters:
      a 0      0.4018 (40)      [ 0.50 (50) ]
      1      0.4018 (40)      [ 0.50 (50) ]
      2      8 (10)e-06         [ 0.50 (50) ]
      E 0      0.90035 (50)     [ 1.00 (50) ]
      1      1.80034 (50)     [ 1.90 (50) ]
      2      2.700 (10)        [ 2.80 (50) ]

Settings:
  svdcut = (1e-15,1e-15)    svdnum = (None,None)    reltol/abstol = 0.0001/0

```

Measured by their χ^2 s, the two fits are almost equally good. The Bayes factor for the first fit, however, is much larger than that for the second fit. It says that the probability that our fit data comes from an underlying theory with exactly two terms is $\exp(254 - 243) = 59,874$ times larger than the probability that it comes from a theory with three terms. In fact, the data comes from a theory with only two terms since it was generated using the same code as in the previous section but with $x, y = \text{make_data}(2)$ instead of $x, y = \text{make_data}()$ in the main program.

1.8 Partial Errors and Error Budgets

We frequently want to know how much of the uncertainty in a fit result is due to a particular input uncertainty or subset of input uncertainties (from the input data and/or from the priors). We refer to such errors as “partial errors” (or partial standard deviations) since each is only part of the total uncertainty in the fit result. The collection of such partial errors, each associated with a different input error, is called an “error budget” for the fit result. The partial errors from all sources of input error reproduce the total fit error when they are added in quadrature.

Given the `fit` object (an `lsqfit.nonlinear_fit` object) from the example in the section on *Correlated Parameters; Gaussian Bayes Factor*, for example, we can extract such information using `gvar.GVar.partialsdev()` — for example:

```

>>> E = fit.p['E']
>>> a = fit.p['a']
>>> print(E[1] / E[0])
1.9994 +- 0.00106276
>>> print((E[1] / E[0]).partialsdev(fit.prior['E']))
0.000419219538523
>>> print((E[1] / E[0]).partialsdev(fit.prior['a']))
0.000158871440271

```

```
>>> print((E[1] / E[0]).partialsdev(y))
0.000952553004005
```

This shows that the total uncertainty in $E[1]/E[0]$, 0.00106, is the sum in quadrature of a contribution 0.00042 due to the priors specified by `prior['E']`, 0.00016 due to `prior['a']`, and 0.00095 from the statistical errors in the input data y .

There are two utility functions for tabulating results and error budgets. They require dictionaries of output results and inputs, and use the keys from the dictionaries to label columns and rows, respectively, in an error-budget table:

```
outputs = {'E1/E0':E[1] / E[0], 'E2/E0':E[2] / E[0],
           'a1/a0':a[1] / a[0], 'a2/a0':a[2] / a[0]}
inputs = {'E':fit.prior['E'], 'a':fit.prior['a'], 'y':y}
print(fit.fmt_values(outputs))
print(fit.fmt_errorbudget(outputs, inputs))
```

This gives the following output:

Values:

```
E2/E0: 3.000 (11)
E1/E0: 1.9994 (11)
a2/a0: 1.005 (28)
a1/a0: 0.9996 (25)
```

Partial % Errors:

	E2/E0	E1/E0	a2/a0	a1/a0
a:	0.09	0.01	1.09	0.02
y:	0.07	0.05	0.77	0.19
E:	0.35	0.02	2.44	0.16
total:	0.37	0.05	2.79	0.25

This table shows, for example, that the 0.37% uncertainty in $E2/E0$ comes from a 0.09% contribution due to `prior['a']`, a 0.07% contribution due to statistical errors in the fit data y , and a 0.35% contribution due to `prior['E']`, where, again, the total error is the sum in quadrature of the partial errors. This suggests that reducing the statistical errors in the input y data would reduce the error in $E2/E0$ only slightly. On the other hand, more accurate y data should significantly reduce the errors in $E1/E0$ and $a1/a0$, where y is the dominant source of uncertainty. In fact a four-fold reduction in the y errors reduces the $E1/E0$ error to 0.02% (from 0.05%) while leaving the $E2/E0$ error at 0.37%.

1.9 y has No Error Bars

Occasionally there are fit problems where values for the dependent variable y are known exactly (to machine precision). This poses a problem for least-squares fitting since the `chi**2` function is infinite when standard deviations are zero. How does one assign errors to exact y s in order to define a `chi**2` function that can be usefully minimized?

It is almost always the case in physical applications of this sort that the fit function has in principle an infinite number of parameters. It is, of course, impossible to extract information about infinitely many parameters from a finite number of y s. In practice, however, we generally care about only a few of the parameters in the fit function. (If this isn't the case, give up.) The goal for a least-squares fit is to figure out what a finite number of exact y s can tell us about the parameters we want to know.

The key idea here is to use priors to model the part of the fit function that we don't care about, and to remove that part of the function from the analysis by subtracting or dividing it out from the input data. To illustrate, consider again the example described in the section on *Correlated Parameters; Gaussian Bayes Factor*. Let us imagine that we know the exact values for y for each of $x=1, 1.2, 1.4 \dots 2.6, 2.8$. We are fitting this data with a sum

of exponentials $a[i] \exp(-E[i] \cdot x)$ where now we will assume that *a priori* we know that: $E[0]=1.0(5)$, $E[i+1]-E[i]=0.9(2)$, and $a[i]=0.5(5)$. Suppose that our goal is to find good estimates for $E[0]$ and $a[0]$.

We know that for some set of parameters

```
y = sum_i=0..inf a[i]*exp(-E[i]*x)
```

for each x-y pair in our fit data. Given that $a[0]$ and $E[0]$ are all we want to know, we might imagine defining a new, modified dependent variable ymod, equal to just $a[0] \exp(-E[0] \cdot x)$:

```
ymod = y - sum_i=1..inf a[i]*exp(-E[i]*x)
```

We know everything on the right-hand side of this equation: we have exact values for y and we have *a priori* estimates for the $a[i]$ and $E[i]$ with $i>0$. So given means and standard deviations for every $i>0$ parameter, and the exact y, we can in principle determine a mean and standard deviation for ymod. The strategy then is to compute the corresponding ymod for every y and x pair, and then fit ymod versus x to the *single* exponential $a[0] \exp(-E[0] \cdot t)$. That fit will give values for $a[0]$ and $E[0]$ that reflect the uncertainties in ymod, which in turn originate in uncertainties in our knowledge about the parameters for the $i>0$ exponentials.

It turns out to be quite simple to implement such a strategy using `gvar.GVars`. We convert our code by first modifying the main program so that it provides prior information to a subroutine that computes ymod. We will vary the number of terms nexp that are kept in the fit, putting the rest into ymod as above (up to a maximum of 20 terms, which is close enough to infinity):

```
def main():
    gv.ranseed([2009, 2010, 2011, 2012]) # initialize random numbers (opt.)
    max_prior = make_prior(20)           # maximum sized prior
    p0 = None                             # make larger fits go faster (opt.)
    for nexp in range(1, 7):
        print('***** nexp =', nexp)
        fit_prior = gv.BufferDict()       # part of max_prior used in fit
        ymod_prior = gv.BufferDict()      # part of max_prior absorbed in ymod
        for k in max_prior:
            fit_prior[k] = max_prior[k][:nexp]
            ymod_prior[k] = max_prior[k][nexp:]
        x, y = make_data(ymod_prior)      # make fit data
        fit = lsqfit.nonlinear_fit(data=(x, y), fcn=f, prior=fit_prior, p0=p0)
        print(fit.format(10))             # print the fit results
        print()
        if fit.chi2/fit.dof<1.:
            p0 = fit.pmean                 # starting point for next fit (opt.)
```

We put all of our *a priori* knowledge about parameters into prior max_prior and then pull out the part we need for the fit — that is, the first nexp terms. The remaining part of max_prior is used to correct the exact data, which comes from a new make_data:

```
def make_data(ymod_prior):               # make x, y fit data
    x = np.arange(1., 10 * 0.2 + 1., 0.2)
    ymod = f_exact(x) - f(x, ymod_prior)
    return x, ymod
```

Running the new code produces the following output, where again nexp is the number of exponentials kept in the fit (and 20-nexp is the number pushed into the modified dependent variable ymod):

```
***** nexp = 1
Least Square Fit (input data correlated with prior):
  chi2/dof [dof] = 0.051 [10]    Q = 1    logGBF = 97.499    itns = 5
```

Parameters:

```

a 0    0.4009 (14)    [ 0.50 (50) ]
E 0    0.90033 (62)   [ 1.00 (50) ]

```

Fit:

	x[k]	y[k]	f(x[k],p)
	1	0.15 (11)	0.16292 (47)
	1.2	0.128 (74)	0.13607 (38)
	1.4	0.110 (52)	0.11365 (30)
	1.6	0.093 (37)	0.09492 (24)
	1.8	0.078 (26)	0.07928 (19)
	2	0.066 (18)	0.06622 (15)
	2.2	0.055 (13)	0.05531 (12)
	2.4	0.0462 (93)	0.046192 (94)
	2.6	0.0387 (66)	0.038581 (74)
	2.8	0.0323 (47)	0.032223 (58)

Settings:

```

svdcut = (1e-15,1e-15)   svdnum = (None,None)   reltol/abstol = 0.0001/0

```

***** nexpt = 2

Least Square Fit (input data correlated with prior):

```

chi2/dof [dof] = 0.053 [10]   Q = 1   logGBF = 99.041   itns = 3

```

Parameters:

```

a 0    0.4002 (13)    [ 0.50 (50) ]
1      0.405 (36)     [ 0.50 (50) ]
E 0    0.90006 (55)   [ 1.00 (50) ]
1      1.803 (30)     [ 1.90 (54) ]

```

Fit:

	x[k]	y[k]	f(x[k],p)
	1	0.223 (45)	0.2293 (44)
	1.2	0.179 (26)	0.1823 (28)
	1.4	0.145 (15)	0.1459 (18)
	1.6	0.1168 (90)	0.1174 (12)
	1.8	0.0947 (53)	0.09492 (74)
	2	0.0770 (32)	0.07711 (47)
	2.2	0.0628 (19)	0.06289 (30)
	2.4	0.0515 (11)	0.05148 (19)
	2.6	0.04226 (67)	0.04226 (12)
	2.8	0.03479 (40)	0.034784 (72)

Settings:

```

svdcut = (1e-15,1e-15)   svdnum = (None,None)   reltol/abstol = 0.0001/0

```

***** nexpt = 3

Least Square Fit (input data correlated with prior):

```

chi2/dof [dof] = 0.057 [10]   Q = 1   logGBF = 99.845   itns = 4

```

Parameters:

```

a 0    0.39998 (93)   [ 0.50 (50) ]
1      0.399 (35)     [ 0.50 (50) ]
2      0.401 (99)     [ 0.50 (50) ]
E 0    0.89999 (36)   [ 1.00 (50) ]

```

```

1      1.799 (26)      [ 1.90 (54) ]
2      2.70 (20)      [ 2.80 (57) ]

```

Fit:

x[k]	y[k]	f(x[k],p)
1	0.253 (19)	0.2557 (54)
1.2	0.1968 (91)	0.1977 (28)
1.4	0.1545 (45)	0.1548 (15)
1.6	0.1224 (22)	0.12256 (76)
1.8	0.0979 (11)	0.09793 (39)
2	0.07885 (54)	0.07886 (20)
2.2	0.06391 (27)	0.06391 (10)
2.4	0.05206 (13)	0.052065 (52)
2.6	0.042602 (67)	0.042601 (26)
2.8	0.034983 (33)	0.034982 (13)

Settings:

```
svdcut = (1e-15,1e-15)   svdnum = (None,None)   reltol/abstol = 0.0001/0
```

```
***** nexp = 4
```

Least Square Fit (input data correlated with prior):

```
chi2/dof [dof] = 0.057 [10]   Q = 1   logGBF = 99.841   itns = 4
```

Parameters:

```

a 0    0.39995 (77)      [ 0.50 (50) ]
  1      0.399 (32)      [ 0.50 (50) ]
  2      0.40 (10)       [ 0.50 (50) ]
  3      0.40 (15)       [ 0.50 (50) ]
E 0    0.89998 (30)      [ 1.00 (50) ]
  1      1.799 (23)      [ 1.90 (54) ]
  2      2.70 (19)       [ 2.80 (57) ]
  3      3.61 (28)       [ 3.70 (61) ]

```

Fit:

x[k]	y[k]	f(x[k],p)
1	0.2656 (78)	0.2666 (22)
1.2	0.2027 (32)	0.20297 (97)
1.4	0.1573 (13)	0.15737 (42)
1.6	0.12378 (54)	0.12381 (18)
1.8	0.09853 (22)	0.098540 (78)
2	0.079153 (93)	0.079155 (34)
2.2	0.064051 (39)	0.064051 (15)
2.4	0.052134 (16)	0.0521344 (64)
2.6	0.0426348 (67)	0.0426347 (29)
2.8	0.0349985 (28)	0.0349985 (13)

Settings:

```
svdcut = (1e-15,1e-15)   svdnum = (None,None)   reltol/abstol = 0.0001/0
```

```
E1/E0 = 1.999(24)   E2/E0 = 3.00(21)
```

```
a1/a0 = 0.997(77)   a2/a0 = 1.01(25)
```

Here we use `fit.format(10)` to print out a table of `x` and `y` (actually `ymod`) values, together with the value of the fit function using the best-fit parameters. There are several things to notice:

- Were we really only interested in $a[0]$ and $E[0]$, a single-exponential fit would have been adequate. This is because we are in effect doing a 20-exponential fit even in that case, by including all but the first term as corrections to y . The answers given by the first fit are correct (we know the exact values since we are using fake data).

The ability to push uninteresting parameters into a `ymod` can be highly useful in practice since it is usually much cheaper to incorporate those fit parameters into `ymod` than it is to include them as fit parameters — fits with smaller numbers of parameters are usually a lot faster.

- The `chi**2` and best-fit parameter means and standard deviations are almost unchanged by shifting terms from `ymod` back into the fit function, as `nexp` increases. The final results for $a[0]$ and $E[0]$, for example, are nearly identical in the `nexp=1` and `nexp=4` fits.

In fact it is straightforward to prove that best-fit parameter means and standard deviations, as well as `chi**2`, should be exactly the same in such situations provided the fit function is linear in all fit parameters. Here the fit function is approximately linear, given our small standard deviations, and so results are only approximately independent of `nexp`.

- The uncertainty in `ymod` for a particular x decreases as `nexp` increases and as x increases. Also the `nexp` independence of the fit results depends upon capturing all of the correlations in the correction to y . This is why `gvar.GVars` are useful since they make the implementation of those correlations trivial.
- Although we motivated this example by the need to deal with y s having no errors, it is straightforward to apply the same ideas to a situation where the y s have errors. Again one might want to do so since fitting uninteresting fit parameters is generally more costly than absorbing them into the y (which then has a modified mean and standard deviation).

1.10 SVD Cuts and Roundoff Error

All of the fits discussed above have (default) *SVD* cuts of $1e-15$. This has little impact in most of the problems, but makes a big difference in the problem discussed in the previous section. Had we run that fit, for example, with an *SVD* cut of $1e-19$, instead of $1e-15$, we would have obtained the following output:

```
Least Square Fit (input data correlated with prior):
  chi2/dof [dof] = 0.057 [10]      Q = 1      logGBF = 100.46      itns = 5
```

```
Parameters:
      a 0    0.39994 (77)      [ 0.50 (50) ]
      1      0.398 (32)      [ 0.50 (50) ]
      2      0.40 (10)      [ 0.50 (50) ]
      3      0.40 (15)      [ 0.50 (50) ]
      E 0    0.89997 (30)      [ 1.00 (50) ]
      1      1.799 (23)      [ 1.90 (54) ]
      2      2.70 (19)      [ 2.80 (57) ]
      3      3.61 (28)      [ 3.70 (61) ]
```

```
Fit:
      x[k]      y[k]      f(x[k],p)
-----
      1      0.2656 (78)      0.267 (16)
     1.2      0.2027 (32)      0.2030 (74)
     1.4      0.1573 (13)      0.1574 (34)
     1.6      0.12378 (54)      0.1238 (15)
     1.8      0.09853 (22)      0.09854 (67)
      2      0.079153 (93)      0.07915 (29)
     2.2      0.064051 (39)      0.06405 (12)
     2.4      0.052134 (16)      0.052134 (41)
```

```

2.6      0.0426348 (67)      0.0426347 (92)
2.8      0.0349985 (28)      0.0349985 (40)

```

Settings:

```
svdcut = (1e-19,None)   svdnum = (None,None)   reltol/abstol = 0.0001/0
```

```

E1/E0 = 2.00(72)   E2/E0 = 3.0(6.1)
a1/a0 = 1.0(2.3)   a2/a0 = 1.0(6.8)

```

The standard deviations quoted for $E1/E0$, *etc.* are much too large compared with the standard deviations shown for the individual parameters, and much larger than what we obtained in the previous section. This is due to roundoff error. The standard deviations quoted for the parameters are computed differently from the standard deviations in `fit.p` (which was used to calculate $E1/E0$). The former come directly from the curvature of the `chi**2` function at its minimum; the latter are related back to the standard deviations of the input data and priors used in the fit. The two should agree, but they will not agree if the covariance matrix for the input y data is too ill-conditioned.

The inverse of the y covariance matrix is used in the `chi**2` function that is minimized by `lsqfit.nonlinear_fit`. Given the finite precision of computer hardware, it is impossible to compute this inverse accurately if the matrix is singular or almost singular, and in such situations the reliability of the fit results is in question. The eigenvalues of the covariance matrix in this example (for `nexp=6`) indicate that this is the case: they range from $7.2e-5$ down to $4.2e-26$, covering 21 orders of magnitude. This is likely too large a range to be handled with the 16–18 digits of precision available in normal double precision computation. The smallest eigenvalues and their eigenvectors are likely to be quite inaccurate, as is any method for computing the inverse matrix.

The standard solution to this common problem in least-squares fitting is to introduce an *SVD* cut, here called `svdcut`:

```
fit = nonlinear_fit(data=(x, ymod), fcn=f, prior=prior, p0=p0, svdcut=1e-15)
```

Then the inverse of the y covariance matrix is computed from its eigenvalues and eigenvectors, but with any eigenvalue smaller than `svdcut` times the largest eigenvalue replaced by the cutoff (that is, by `svdcut` times the largest eigenvalue). This limits the singularity of the covariance matrix, leading to improved numerical stability. The cost is less precision in the final results since we are in effect decreasing the precision of the input y data. This is a conservative move, but numerical stability is worth the tradeoff.

Note that taking `svdcut=-1e-15`, with a minus sign, causes the problematic modes to be dropped. This is a more conventional implementation of *SVD* cuts, but here it results in much less precision than using `svdcut=1e-15` (giving, for example, 1.993(69) for $E1/E0$, which is almost three times less precise). Dropping modes is equivalent to setting the corresponding variances to infinity, which is (obviously) much more conservative and less realistic than setting them equal to the *SVD*-cutoff variance.

The error budget is interesting in this case. There is no contribution from the original y data since it was exact. So all statistical uncertainty comes from the priors in `max_prior`, and from the *SVD* cut, which contributes since it modifies the effective variances of several eigenmodes of the covariance matrix. The *SVD* contribution can be obtained from `fit.svdcorrection` so the full error budget is constructed by the following code,

```

outputs = {'E1/E0':E[1] / E[0], 'E2/E0':E[2] / E[0],
           'a1/a0':a[1] / a[0], 'a2/a0':a[2] / a[0]}
inputs = {'E':max_prior['E'], 'a':max_prior['a'], 'svd':fit.svdcorrection}
print(fit.fmt_values(outputs))
print(fit.fmt_errorbudget(outputs, inputs))

```

which gives:

Values:

```

E2/E0: 3.00(21)
E1/E0: 1.999(24)
a2/a0: 1.01(25)
a1/a0: 0.997(77)

```

Partial % Errors:

	E2/E0	E1/E0	a2/a0	a1/a0
-----	-----	-----	-----	-----
a:	3.76	0.71	11.80	4.39
svd:	0.29	0.10	0.13	0.55
E:	5.87	0.99	22.35	6.30
-----	-----	-----	-----	-----
total:	6.98	1.22	25.27	7.70

Here the contribution from the *SVD* cut is almost negligible.

Note that covariance matrices are rescaled so that all diagonal elements equal one before the *SVD* cut is applied. This means, among other things, that uncorrelated errors — that is, diagonal sub-matrices of the covariance matrix — are unaffected by *SVD* cuts. Applying an *SVD* cut of $1e-4$, for example, to the following singular covariance matrix,

```
[[ 1.0  1.0  0.0 ]
 [ 1.0  1.0  0.0 ]
 [ 0.0  0.0  1e-20]],
```

gives a new, non-singular matrix:

```
[[ 1.0001  0.9999  0.0 ]
 [ 0.9999  1.0001  0.0 ]
 [ 0.0      0.0     1e-20]]
```

`lsqfit.nonlinear_fit` uses a default value for `svdcut` of $1e-15$, and applies *SVD* cuts to the covariance matrices from both the fit data and the prior. This default can be overridden as shown above, but for many problems it is a good choice. Roundoff errors become more accute, however, when there are strong positive correlations between different parts of the fit data or prior. Then much larger `svdcuts` may be needed.

The method `lsqfit.nonlinear_fit.check_roundoff()` can be used to check for roundoff errors by adding the line `fit.check_roundoff()` after the fit. It generates a warning if roundoff looks to be a problem. This check is done automatically if `debug=True` is added to argument list of `lsqfit.nonlinear_fit`.

1.11 Bootstrap Error Analysis

Our analysis above assumes that every probability distribution relevant to the fit is approximately Gaussian. For example, we characterize the input data for y by a mean and a covariance matrix obtained from averaging many random samples of y . For large sample sizes it is almost certainly true that the average values follow a Gaussian distribution, but in practical applications the sample size could be too small. The *statistical bootstrap* is an analysis tool for dealing with such situations.

The strategy is to: 1) make a large number of “bootstrap copies” of the original input data that differ from each other by random amounts characteristic of the underlying randomness in the original data; 2) repeat the entire fit analysis for each bootstrap copy of the data, extracting fit results from each; and 3) use the variation of the fit results from bootstrap copy to bootstrap copy to determine an approximate probability distribution (possibly non-Gaussian) for the each result.

Consider the code from the previous section, where we might reasonably want another check on the error estimates for our results. That code can be modified to include a bootstrap analysis by adding the following to the end of the `main()` subroutine:

```
Nbs = 40                                     # number of bootstrap copies
outputs = {'E1/E0':[], 'E2/E0':[], 'a1/a0':[], 'a2/a0':[]} # results
for bsfit in fit.bootstrap_iter(n=Nbs):
    E = bsfit.pmean['E']                     # best-fit parameter values
```



```

a = bsfit.pmean['a']                                # (ignore errors)
outputs['E1/E0'].append(E[1] / E[0])                # accumulate results
outputs['E2/E0'].append(E[2] / E[0])
outputs['a1/a0'].append(a[1] / a[0])
outputs['a2/a0'].append(a[2] / a[0])
outputs['E1'].append(E[1])
outputs['a1'].append(a[1])
# extract "means" and "standard deviations" from the bootstrap output;
# print using .fmt() to create compact representation of GVars
outputs = gv.dataset.avg_data(outputs, bstrap=True)
print('Bootstrap results:')
print('E1/E0 =', outputs['E1/E0'].fmt(), '    E2/E1 =', outputs['E2/E0'].fmt())
print('a1/a0 =', outputs['a1/a0'].fmt(), '    a2/a0 =', outputs['a2/a0'].fmt())
print('E1 =', outputs['E1'].fmt(), '    a1 =', outputs['a1'].fmt())

```

The results are consistent with the results obtained directly from the fit (when using `svdcut=1e-15`):

```

Bootstrap results:
E1/E0 = 1.999(17)    E2/E1 = 2.98(18)
a1/a0 = 0.995(55)    a2/a0 = 0.97(28)
E1 = 1.799(16)    a1 = 0.398(23)

```

In particular, the bootstrap analysis confirms our previous error estimates (to within 10-30%, since `Nbs=40`). When `Nbs` is small, it is often safer to use the median instead of the mean as the estimator, which is what `gv.dataset.avg_data` does here since flag `bstrap` is set to `True`.

1.12 Testing Fits with Simulated Data

Ideally we would test a fitting protocol by doing fits of data similar to our actual fit but where we know the correct values for the fit parameters ahead of the fit. The `lsqfit.nonlinear_fit` iterator `simulated_fit_iter` creates any number of such simulations of the original fit. Returning again to the fits in the section on *Correlated Parameters; Gaussian Bayes Factor*, we can add three fit simulations to the end of the main program:

```

def main():
    gv.seed([2009, 2010, 2011, 2012]) # initialize random numbers (opt.)
    x, y = make_data()                # make fit data
    p0 = None                          # make larger fits go faster (opt.)
    for nexp in range(3, 20):
        print('***** nexp =', nexp)
        prior = make_prior(nexp)
        fit = lsqfit.nonlinear_fit(data=(x, y), fcn=f, prior=prior, p0=p0)
        print(fit)                    # print the fit results
        E = fit.p['E']                # best-fit parameters
        a = fit.p['a']
        print('E1/E0 =', E[1] / E[0], '    E2/E0 =', E[2] / E[0])
        print('a1/a0 =', a[1] / a[0], '    a2/a0 =', a[2] / a[0])
        print()
        if fit.chi2 / fit.dof < 1.:
            p0 = fit.pmean             # starting point for next fit (opt.)

    # 3 fit simulations based upon last fit
    for sfit in fit.simulated_fit_iter(3):
        print(sfit)
        sE = sfit.p['E']              # best-fit parameters (simulation)
        sa = sfit.p['a']
        E = sfit.pexact['E']          # correct results for parameters

```

```

a = sfit.pexact['a']
print('E1/E0 =', sE[1] / sE[0], ' E2/E0 =', sE[2] / sE[0])
print('a1/a0 =', sa[1] / sa[0], ' a2/a0 =', sa[2] / sa[0])
print('\nSimulated Fit Values - Exact Values:')
print(
    'E1/E0:', (sE[1] / sE[0]) - (E[1] / E[0]),
    ' E2/E0:', (sE[2] / sE[0]) - (E[2] / E[0])
)
print(
    'a1/a0:', (sa[1] / sa[0]) - (a[1] / a[0]),
    ' a2/a0:', (sa[2] / sa[0]) - (a[2] / a[0])
)

# compute chi**2 comparing selected fit results to exact results
sim_results = [sE[0], sE[1], sa[0], sa[1]]
exact_results = [E[0], E[1], a[0], a[1]]
chi2 = gv.chi2(sim_results, exact_results)
print(
    '\nParameter chi2/dof [dof] = %.2f' % (chi2 / gv.chi2.dof),
    '[%d]' % gv.chi2.dof,
    ' Q = %.1f' % gv.chi2.Q
)

```

The fit data for each of the three simulations is the same as the original fit data except that the means have been adjusted (randomly) so the the correct values for the fit parameters are in each case equal to `pexact=fit.pmean`. Simulation fit results will typically differ from the correct values by an amount of order a standard deviation. With sufficiently accurate data, the results from a large number of simulations will be distributed in Gaussians centered on the correct values (`pexact`), with widths that equal the standard deviations given by the fit (`fit.psdev`). (With less accurate data, the distributions may become non-Gaussian, and the interpretation of fit results more complicated.)

In the present example, the output from the three simulations is:

```

***** simulation
Least Square Fit:
    chi2/dof [dof] = 0.27 [15]      Q = 1      logGBF = 228.78      itns = 46

Parameters:
    a 0    0.3968 (40)      [ 0.50 (50) ]
      1    0.3983 (41)      [ 0.50 (50) ]
      2    0.390 (12)       [ 0.50 (50) ]
      3    0.453 (45)       [ 0.50 (50) ]
      4    0.15 (15)        [ 0.50 (50) ]
      5    0.72 (31)        [ 0.50 (50) ]
      6    0.71 (42)        [ 0.50 (50) ]
    E 0    0.89969 (51)     [ 1.00 (50) ]
      1    1.8011 (11)      [ 1.90 (50) ]
      2    2.703 (10)       [ 2.80 (50) ]
      3    3.603 (14)       [ 3.70 (50) ]
      4    4.503 (17)       [ 4.60 (50) ]
      5    5.403 (20)       [ 5.50 (50) ]
      6    6.303 (22)       [ 6.40 (50) ]

Settings:
    svdcut = (None,None)    svdnum = (None,None)    reltol/abstol = 0.0001/0

E1/E0 = 2.00188 +- 0.00106514    E2/E0 = 3.00474 +- 0.0111774
a1/a0 = 1.00397 +- 0.00252259    a2/a0 = 0.984026 +- 0.0282035

```

Simulated Fit Values - Exact Values:

E1/E0: 0.00247635 +- 0.00106514 E2/E0: 0.00485 +- 0.0111774
a1/a0: 0.00436695 +- 0.00252259 a2/a0: -0.0209515 +- 0.0282035

Parameter chi2/dof [dof] = 1.78 [4] Q = 0.1

***** simulation

Least Square Fit:

chi2/dof [dof] = 0.73 [15] Q = 0.76 logGBF = 225.31 itns = 25

Parameters:

a	0	0.4034 (40)	[0.50 (50)]
	1	0.4037 (42)	[0.50 (50)]
	2	0.403 (12)	[0.50 (50)]
	3	0.412 (46)	[0.50 (50)]
	4	0.34 (16)	[0.50 (50)]
	5	0.56 (31)	[0.50 (50)]
	6	0.57 (42)	[0.50 (50)]
E	0	0.90024 (51)	[1.00 (50)]
	1	1.8005 (11)	[1.90 (50)]
	2	2.701 (10)	[2.80 (50)]
	3	3.601 (14)	[3.70 (50)]
	4	4.501 (17)	[4.60 (50)]
	5	5.401 (20)	[5.50 (50)]
	6	6.301 (22)	[6.40 (50)]

Settings:

svdcut = (None,None) svdnum = (None,None) reltol/abstol = 0.0001/0

E1/E0 = 1.99998 +- 0.00105728 E2/E0 = 3.00061 +- 0.0111101
a1/a0 = 1.00066 +- 0.00253149 a2/a0 = 0.999036 +- 0.0280364

Simulated Fit Values - Exact Values:

E1/E0: 0.000580316 +- 0.00105728 E2/E0: 0.000719649 +- 0.0111101
a1/a0: 0.00106294 +- 0.00253149 a2/a0: -0.0059414 +- 0.0280364

Parameter chi2/dof [dof] = 0.13 [4] Q = 1.0

***** simulation

Least Square Fit:

chi2/dof [dof] = 0.92 [15] Q = 0.54 logGBF = 223.86 itns = 2

Parameters:

a	0	0.4003 (40)	[0.50 (50)]
	1	0.4002 (42)	[0.50 (50)]
	2	0.402 (12)	[0.50 (50)]
	3	0.390 (46)	[0.50 (50)]
	4	0.41 (16)	[0.50 (50)]
	5	0.49 (31)	[0.50 (50)]
	6	0.50 (42)	[0.50 (50)]
E	0	0.90040 (51)	[1.00 (50)]
	1	1.8003 (11)	[1.90 (50)]
	2	2.701 (10)	[2.80 (50)]
	3	3.601 (14)	[3.70 (50)]
	4	4.501 (17)	[4.60 (50)]
	5	5.401 (20)	[5.50 (50)]
	6	6.301 (22)	[6.40 (50)]

```
Settings:
    svdcut = (None,None)    svdnum = (None,None)    reltol/abstol = 0.0001/0

E1/E0 = 1.99949 +- 0.0010671    E2/E0 = 2.99932 +- 0.0110989
a1/a0 = 0.999605 +- 0.00255698    a2/a0 = 1.0041 +- 0.0280514

Simulated Fit Values - Exact Values:
E1/E0: 9.04515e-05 +- 0.0010671    E2/E0: -0.000569657 +- 0.0110989
a1/a0: 4.25028e-06 +- 0.00255698    a2/a0: -0.000873644 +- 0.0280514

Parameter chi2/dof [dof] = 0.14 [4]    Q = 1.0
```

The simulations show that the fit values usually agree with the correct values to within a standard deviation or so (the correct results here are the mean values from the last fit discussed in *Correlated Parameters; Gaussian Bayes Factor*). Furthermore the error estimates for each parameter are reproduced by the simulations. We also compute the `chi**2` for the difference between the leading fit parameters and the exact values. This checks parameter values, standard deviations, and correlations. The results are reasonable for four degrees of freedom. Here the first simulation shows results that are off by a third of a standard deviation on average, but this is not so unusual — the `Q=0.1` indicates that it happens 10% of the time.

More thorough testing is possible: for example, one could run many simulations (100?) to verify that the distribution of (simulation) fit results is Gaussian, centered around `pexact`. This is overkill in most situations, however. The three simulations above are enough to reassure us that the original fit estimates, including errors, are reliable.

1.13 Positive Parameters

The priors for `lsqfit.nonlinear_fit` are all Gaussian. There are situations, however, where other distributions would be desirable. One such case is where a parameter is known to be positive, but is close to zero in value (“close” being defined relative to the *a priori* uncertainty). For such cases we would like to use non-Gaussian priors that force positivity — for example, priors that impose log-normal or exponential distributions on the parameter. Ideally the decision to use such a distribution would be made on a parameter- by-parameter basis, when creating the priors, and would have no impact on the definition of the fit function itself.

`lsqfit` provides a decorator, `lsqfit.transform_p`, for fit functions that makes this possible. This decorator only works for fit functions that use dictionaries for their parameters. Given a prior `prior` for a fit, the decorator is used in the following way: for example,

```
@lsqfit.transform_p(prior.keys(), 0)
def fitfcn(p):
    ...
```

when the parameter argument is the first argument of the fit function, or

```
@lsqfit.transform_p(prior.keys(), 1)
def fitfcn(x, p):
    ...
```

when the parameter argument is the second argument of the fit function (see the `lsqfit.transform_p` documentation for more detail). Consider any parameter `p['XX']` used in `fitfcn`. The prior distribution for that parameter can now be turned into a log-normal distribution by replacing `prior['XX']` with `prior['logXX']` (or `prior['log(XX)']`) when defining the prior, thereby assigning a Gaussian distribution to `logXX` rather than to `XX`. Nothing need be changed in the fit function, other than adding the decorator. The decorator automatically detects parameters whose keys begin with `'log'` and adds new parameters to the parameter-dictionary for `fitfcn` that are exponentials of those parameters.

To illustrate consider a simple problem where an experimental quantity y is known to be positive, but experimental errors mean that measured values can often be negative:

```
import gvar as gv
import lsqfit

y = gv.gvar([
    '-0.17(20)', '-0.03(20)', '-0.39(20)', '0.10(20)', '-0.03(20)',
    '0.06(20)', '-0.23(20)', '-0.23(20)', '-0.15(20)', '-0.01(20)',
    '-0.12(20)', '0.05(20)', '-0.09(20)', '-0.36(20)', '0.09(20)',
    '-0.07(20)', '-0.31(20)', '0.12(20)', '0.11(20)', '0.13(20)'
])
```

We want to know the average value a of the y s and so could use the following fitting code:

```
prior = gv.BufferDict(a=gv.gvar(0.02, 0.02))      # a = avg value of y's

def fcn(p, N=len(y)):
    return N * [p['a']]

fit = lsqfit.nonlinear_fit(prior=prior, data=y, fcn=fcn)
print(fit)
print('a =', fit.p['a'].fmt())
```

where we are assuming *a priori* information that suggests the average is around 0.02. The output from this code is:

```
Least Square Fit:
  chi2/dof [dof] = 0.84 [20]      Q = 0.67      logGBF = 5.3431      itns = 2

Parameters:
      a      0.004 (18)      [ 0.020 (20) ]

Settings:
  svdcut = (1e-15,1e-15)      svdnum = (None,None)      reltol/abstol = 0.0001/0

a = 0.004(18)
```

This is not such a useful result since much of the one-sigma range for a is negative, and yet we know that a must be positive.

A better analysis is to use a log-normal distribution for a :

```
prior = gv.BufferDict(loga=gv.log(gv.gvar(0.02, 0.02))) # loga not a

@lsqfit.transform_p(prior.keys(), 0)
def fcn(p, N=len(y)):
    return N * [p['a']]

fit = lsqfit.nonlinear_fit(prior=prior, data=y, fcn=fcn)
print(fit)
print('a =', fit.transformed_p['a'].fmt())      # exp(loga)
```

The fit parameter is now $\log(a)$ rather than a itself, but we are able to use the identical fit function. Here `fit.transformed_p` is the same as `fit.p` but augmented to include the exponentials of any log-normal variables — that is, a as well as $\log a$. Rather than including all keys, the decorator can be written with a list containing just the variables to be transformed: here, `@lsqfit.transform_p(['loga'], 0)`.

The result from this fit is

```
Least Square Fit:
  chi2/dof [dof] = 0.85 [20]      Q = 0.65      logGBF = 5.252      itns = 12

Parameters:
      loga      -4.44 (97)      [ -3.9 (1.0) ]

Settings:
  svdcut = (1e-15,1e-15)      svdnum = (None,None)      reltol/abstol = 0.0001/0

a = 0.012(11)
```

which is more compelling. The “correct” value for a here is 0.015 (from the method used to generate the y s).

`lsqfit.transform_p()` also allows parameters to be replaced by their square roots as fit parameters — for example, define `prior['sqrt(a)']` (or `prior['sqrta']`) rather than `prior['a']` when creating the prior. This again guarantees positive parameters. The prior for `p['a']` is an exponential distribution if the mean of `p['sqrt(a)']` is zero. Using `prior['sqrt(a)']` in place of `prior['a']` in the example above leads to $a = 0.010(13)$, which is almost identical to the result obtained from the log-normal distribution.

1.14 Troubleshooting

`lsqfit.nonlinear_fit` error messages that come from inside the *gsf* routines doing the fits are sometimes less than useful. They are usually due to errors in one of the inputs to the fit (that is, the fit data, the prior, or the fit function). Setting `debug=True` in the argument list of `lsqfit.nonlinear_fit` might result in more intelligible error messages. This option also causes the fitter to check for significant roundoff errors in the matrix inversions of the covariance matrices.

Occasionally `lsqfit.nonlinear_fit` appears to go crazy, with gigantic χ^2 s (e.g., $1e78$). This could be because there is a genuine zero-eigenvalue mode in the covariance matrix of the data or prior. Such a zero mode makes it impossible to invert the covariance matrix when evaluating χ^2 . One fix is to include SVD cuts in the fit by setting, for example, `svdcut=(1e-14, 1e-14)` in the call to `lsqfit.nonlinear_fit`. These cuts will exclude exact or nearly exact zero modes, while leaving important modes mostly unaffected.

Even if the SVD cuts work in such a case, the question remains as to why one of the covariance matrices has a zero mode. A common cause is if the same `gvar.GVar` was used for more than one prior. For example, one might think that

```
>>> import gvar as gv
>>> z = gv.gvar(1, 1)
>>> prior = gv.BufferDict(a=z, b=z)
```

creates a prior 1 ± 1 for each of parameter a and parameter b . Indeed each parameter separately is of order 1 ± 1 , but in a fit the two parameters would be forced equal to each other because their priors are both set equal to the same `gvar.GVar`, z :

```
>>> print(prior['a'], prior['b'])
1 +- 1 1 +- 1
>>> print(prior['a']-prior['b'])
0 +- 0
```

That is, while parameters a and b fluctuate over a range of 1 ± 1 , they fluctuate together, in exact lock-step. The covariance matrix for a and b must therefore be singular, with a zero mode corresponding to the combination $a-b$; it is all 1s in this case:

```
>>> import numpy as np
>>> cov = gv.evalcov(prior.flat)      # prior's covariance matrix
```

```
>>> print(np.linalg.det(cov))      # determinant is zero
0.0
```

This zero mode upsets `nonlinear_fit()`. If `a` and `b` are meant to fluctuate together then an *SVD* cut as above will give correct results (with `a` and `b` being forced equal to several decimal places, depending upon the cut). Of course, simply replacing `b` by `a` in the fit function would be even better. If, on the other hand, `a` and `b` were not meant to fluctuate together, the prior should be redefined:

```
>>> prior = gv.BufferDict(a=gv.gvar(1, 1), b=gv.gvar(1, 1))
```

where now each parameter has its own `gvar.GVar`.

1.15 Appendix: A Simple Pedagogical Example

Consider a problem where we have five pieces of uncorrelated data for a function $y(x)$:


<code>x[i]</code>	<code>y(x[i])</code>
0.1	0.5351 (54)
0.3	0.6762 (67)
0.5	0.9227 (91)
0.7	1.3803 (131)
0.95	4.0145 (399)

We know that $y(x)$ has a Taylor expansion in x :

$$y(x) = \sum_{n=0}^{\infty} p[n] x^n$$

The challenge is to extract a reliable estimate for $y(0)=p[0]$ from the data — that is, the challenge is to extrapolate the data to $x=0$.

One approach that is certainly wrong is to truncate the expansion of $y(x)$ after five terms, because there are only five pieces of data. That gives the following fit:



This fit was generated using the following code:

```
import numpy as np
import gvar as gv
import lsqfit
```

```
# fit data
y = gv.gvar([
    '0.5351(54)', '0.6762(67)', '0.9227(91)', '1.3803(131)', '4.0145(399)'
])
x = np.array([0.1, 0.3, 0.5, 0.7, 0.95])

def f(x, p):
    # fit function
    return sum(pn * x ** n for n, pn in enumerate(p))

p0 = np.ones(5.) # starting value for chi**2 minimization
fit = lsqfit.nonlinear_fit(data=(x, y), p0=p0, fcn=f)
print(fit.format(maxline=5))
```

Note that here the function `gv.gvar` converts the strings `'0.5351(54)'`, etc. into `gvar.GVars`. Running the code gives the following output:

```
Least Square Fit (no prior):
  chi2/dof [dof] = 4.8e-27 [0]      Q = 0      logGBF = None      itns = 2

Parameters:
      0      0.742 (39)      [ 1.0 +- inf ]
      1      -3.86 (59)      [ 1.0 +- inf ]
      2       21.5 (2.4)      [ 1.0 +- inf ]
      3     -39.1 (3.7)      [ 1.0 +- inf ]
      4       25.8 (1.9)      [ 1.0 +- inf ]

Fit:
      x[k]      y[k]      f(x[k],p)
-----
      0.1      0.5351 (54)      0.5351 (54)
      0.3      0.6762 (67)      0.6762 (67)
      0.5      0.9227 (91)      0.9227 (91)
      0.7      1.380 (13)      1.380 (13)
      0.95      4.014 (40)      4.014 (40)
```

This is a perfect fit in that the fit function agrees exactly with the data; the `chi**2` for the fit is zero. The 5-parameter fit gives a fairly precise answer for `p[0]` (`0.74(4)`), but the curve looks oddly stiff. Also some of the best-fit values for the coefficients are quite large (e.g., `p[3] = -39(4)`), perhaps unreasonably large.

The problem with a 5-parameter fit is that there is no reason to neglect terms in the expansion of $y(x)$ with $n > 4$. Whether or not extra terms are important depends entirely on how large we expect the coefficients `p[n]` for $n > 4$ to be. The extrapolation problem is impossible without some idea of the size of these parameters; we need extra information.

In this case that extra information is obviously connected to questions of convergence of the Taylor expansion we are using to model $y(x)$. Let's assume we know, from previous work, that the `p[n]` are of order one. Then we would need to keep at least 91 terms in the Taylor expansion if we wanted the terms we dropped to be small compared with the 1% data errors at $x=0.95$. So a possible fitting function would be:

```
y(x; N) = sum_n=0..N p[n] x**n
```

with `N=90`.

Fitting a 91-parameter formula to five pieces of data is also impossible. Here, however, we have extra (*prior*) information: each coefficient is order one, which we make specific by saying that they equal `0(1)`. So we are actually fitting 91+5 pieces of data with 91 parameters.

The prior information is introduced into the fit as a *prior*:

```
import numpy as np
import gvar as gv
```




```
import lsqfit

# fit data
y = gv.gvar([
    '0.5351(54)', '0.6762(67)', '0.9227(91)', '1.3803(131)', '4.0145(399)'
])
x = np.array([0.1, 0.3, 0.5, 0.7, 0.95])

def f(x, p):
    # fit function
    return sum(pn * x ** n for n, pn in enumerate(p))

prior = gv.gvar(91 * ['0(1)']) # 91-parameter prior for the fit
fit = lsqfit.nonlinear_fit(data=(x, y), prior=prior, fcn=f)
print(fit.format(maxline=5))
```

Note that a starting value `p0` is not needed when a prior is specified. This code also gives an excellent fit, with a χ^2 per degree of freedom of 0.35:



The fit code output is:

```
Least Square Fit:
chi2/dof [dof] = 0.35 [5]    Q = 0.88    logGBF = -0.45508    itns = 2
```

```
Parameters:
0    0.489 (17)    [ 0.0 (1.0) ]
1    0.40 (20)    [ 0.0 (1.0) ]
2    0.60 (64)    [ 0.0 (1.0) ]
3    0.44 (80)    [ 0.0 (1.0) ]
4    0.28 (87)    [ 0.0 (1.0) ]
5    0.19 (87)    [ 0.0 (1.0) ]
6    0.16 (90)    [ 0.0 (1.0) ]
7    0.16 (93)    [ 0.0 (1.0) ]
8    0.17 (95)    [ 0.0 (1.0) ]
9    0.18 (96)    [ 0.0 (1.0) ]
10   0.19 (97)    [ 0.0 (1.0) ]
11   0.19 (97)    [ 0.0 (1.0) ]
12   0.19 (97)    [ 0.0 (1.0) ]
13   0.19 (97)    [ 0.0 (1.0) ]
14   0.18 (97)    [ 0.0 (1.0) ]
15   0.18 (97)    [ 0.0 (1.0) ]
.
```

```

.
.
88    0.0 (1.0)    [ 0.0 (1.0) ]
89    0.0 (1.0)    [ 0.0 (1.0) ]
90    0.0 (1.0)    [ 0.0 (1.0) ]

```

Fit:

x[k]	y[k]	f(x[k],p)
0.1	0.5351 (54)	0.5349 (54)
0.3	0.6762 (67)	0.6768 (65)
0.5	0.9227 (91)	0.9219 (87)
0.7	1.380 (13)	1.381 (13)
0.95	4.014 (40)	4.014 (40)

Settings:

```
svdcut = (1e-15,1e-15)    svdnum = (None,None)    reltol/abstol = 0.0001/0
```

This is a much more plausible fit than the 5-parameter fit, and gives an extrapolated value of $p[0]=0.489(17)$. The original data points were created using a Taylor expansion with random coefficients, but with $p[0]$ set equal to 0.5. So this fit to the five data points (plus 91 *a priori* values for the $p[n]$ with $n<91$) gives the correct result. Increasing the number of terms further would have no effect since the last terms added are having no impact, and so end up equal to the prior value — the fit data are not sufficiently precise to add new information about these parameters.

We can test our priors for this fit by re-doing the fit with broader and narrower priors. Setting `prior = gv.gvar(91 * ['0(3)'])` gives an excellent fit,

Least Square Fit:

```
chi2/dof [dof] = 0.039 [5]    Q = 1    logGBF = -5.0993    itns = 2
```

Parameters:

```

0    0.490 (33)    [ 0.0 (3.0) ]
1    0.38 (48)    [ 0.0 (3.0) ]
2    0.6 (1.8)    [ 0.0 (3.0) ]
...

```

but with a very small `chi2/dof` and somewhat larger errors on the best-fit estimates for the parameters. The logarithm of the (Gaussian) Bayes Factor, `logGBF`, can be used to compare fits with different priors. It is the logarithm of the probability that our data would come from parameters generated at random using the prior. The exponential of `logGBF` is more than 100 times larger with the original priors of 0(1) than with priors of 0(3). This says that our data is more than 100 times more likely to come from a world with parameters of order one than from one with parameters of order three. Put another way it says that the size of the fluctuations in the data are more consistent with coefficients of order one than with coefficients of order three. The `logGBF` values argue for the original prior.

Narrower priors, `prior = gv.gvar(91 * ['0.0(3)'])`, give a poor fit, and also a less optimal `logGBF`:

Least Square Fit:

```
chi2/dof [dof] = 3.7 [5]    Q = 0.0024    logGBF = -3.3058    itns = 2
```

Parameters:

```

0    0.484 (11)    [ 0.00 (30) ] *
1    0.454 (98)    [ 0.00 (30) ] *
2    0.50 (23)    [ 0.00 (30) ] *
...

```

The priors are responsible for about half of the final error in our best estimate of $p[0]$ (with priors of 0(1)); the rest comes from the uncertainty in the data. This can be established by creating an error budget using the code

```
inputs = dict(prior=prior, y=y)
outputs = dict(p0=fit.p[0])
print(gv.fmt_errorbudget(inputs=inputs, outputs=outputs))
```

which prints the following table:

```
Partial % Errors:
                                     p0
-----
                y:      2.67
            prior:      2.23
-----
            total:      3.48
```

There is a second, equivalent way of fitting this data that illustrates the idea of *marginalization*. We really only care about parameter $p[0]$ in our fit. This suggests that we remove $n>0$ terms from the data *before* we do the fit:

```
ymod[i] = y[i] - sum_n=1...inf prior[n] * x[i] ** n
```

Before the fit, our best estimate for the parameters is from the priors. We use these to create an estimate for the correction to each data point coming from $n>0$ terms in $y(x)$. This new data, $y_{\text{mod}}[i]$, should be fit with a new fitting function, $y_{\text{mod}}(x) = p[0]$ — that is, it should be fit to a constant, independent of $x[i]$. The last three lines of the code above are easily modified to implement this idea:

```
import numpy as np
import gvar as gv
import lsqfit

# fit data
y = gv.gvar([
    '0.5351(54)', '0.6762(67)', '0.9227(91)', '1.3803(131)', '4.0145(399)'
])
x = np.array([0.1, 0.3, 0.5, 0.7, 0.95])

def f(x, p):
    # fit function
    return sum(pn * x ** n for n, pn in enumerate(p))

prior = gv.gvar(91 * ['0(1)']) # prior for the fit
priormod = prior[:1]          # restrict fit to p[0]
ymod = y - (f(x, prior) - f(x, priormod)) # correct y
fit = lsqfit.nonlinear_fit(data=(x, ymod), prior=priormod, fcn=f)
print(fit.format(maxline=5))
```

Running this code give:

```
Least Square Fit (input data correlated with prior):
  chi2/dof [dof] = 0.35 [5]    Q = 0.88    logGBF = -0.45508    itns = 2
```


```
Parameters:
          0    0.489 (17)    [ 0.0 (1.0) ]
```

```
Fit:
      x[k]      y[k]      f(x[k],p)
-----
      0.1      0.54 (10)    0.489 (17)
      0.3      0.68 (31)    0.489 (17)
      0.5      0.92 (58)    0.489 (17)
      0.7      1.38 (98)    0.489 (17)
      0.95     4.0 (3.0)    0.489 (17)  *
```

Settings:

```
svdcut = (1e-15,1e-15)    svdnum = (None,None)    reltol/abstol = 0.0001/0
```

Remarkably this one-parameter fit gives results for `p[0]` that are identical (to machine precision) to our 91-parameter fit above. The 90 parameters for $n>0$ are said to have been *marginalized* in this fit. Marginalizing a parameter in this way has no effect if the fit function is linear in that parameter. Marginalization has almost no effect for nonlinear fits as well, provided the fit data have small errors (in which case the parameters are effectively linear). The fit here is:



The constant is consistent with all of the data in `ymod[i]`, even at `x[i]=0.95`, because `ymod[i]` has much larger errors for larger `x[i]` because of the correction terms.

Fitting to a constant is equivalent to doing a weighted average of the data plus the prior, so our fit can be replaced by an average:

```
lsqfit.wavg(list(ymod) + list(priormod))
```

This again gives 0.489(17) for our final result. Note that the central value for this average is below the central values for every data point in `ymod[i]`. This is a consequence of large positive correlations introduced into `ymod` when we remove the $n>0$ terms. These correlations are captured automatically in our code, and are essential — removing the correlations between different `ymods` results in a final answer, 0.564(97), which has much larger errors.

GVAR - GAUSSIAN RANDOM VARIABLES

2.1 Introduction

This module provides tools for representing and manipulating Gaussian random variables numerically. A Gaussian variable is a random variable that represents a *typical* random number drawn from a particular Gaussian (or normal) probability distribution; more precisely, it represents the entire probability distribution, and not, for example, a *particular* random number drawn from that distribution. A given Gaussian variable x is therefore completely characterized by its mean $x.mean$ and standard deviation $x.sdev$.

A mathematical function of a Gaussian variable can be defined as the probability distribution of function values obtained by evaluating the function for random numbers drawn from the original distribution. The distribution of function values is itself approximately Gaussian provided the standard deviation of the Gaussian variable is sufficiently small. Thus we can define a function f of a Gaussian variable x to be a Gaussian variable itself, with

```
f(x).mean = f(x.mean)
f(x).sdev = x.sdev |f'(x.mean)|,
```

which follows from linearizing the x dependence of $f(x)$ about point $x.mean$. (This obviously fails at an extremum of $f(x)$, where $f'(x)=0$.)

The last formula, together with its multidimensional generalization, leads to a full calculus for Gaussian random variables that assigns Gaussian-variable values to arbitrary arithmetic expressions and functions involving Gaussian variables. This calculus is useful for analyzing the propagation of statistical and other random errors (provided the standard deviations are small enough).

A multidimensional collection $x[i]$ of Gaussian variables is characterized by the means $x[i].mean$ for each variable, together with a covariance matrix $cov[i, j]$. Diagonal elements of cov specify the standard deviations of different variables: $x[i].sdev = cov[i, i]**0.5$. Nonzero off-diagonal elements imply correlations between different variables:

```
cov[i, j] = <x[i]*x[j]> - <x[i]> * <x[j]>
```

where $\langle y \rangle$ denotes the expectation value or mean for a random variable y .

2.2 Creating Gaussian Variables

An object of type `gvar.GVar` represents a single Gaussian variable. Such an object can be created for a single variable, with mean $xmean$ and standard deviation $xsdev$ (both scalars), using:

```
x = gvar.gvar(xmean, xsdev).
```

This function can also be used to convert strings like `'-72.374(22)'` or `'511.2 +- 0.3'` into `gvar.GVars`: for example,

```
>>> import gvar
>>> x = gvar.gvar(3.1415, 0.0002)
>>> print(x)
3.1415 +- 0.0002
>>> x = gvar.gvar("3.1415(2)")
>>> print(x)
3.1415 +- 0.0002
```

This `gvar.GVar` can be converted to a more compact string using the `GVar.fmt()` method: for example,

```
>>> print(x.fmt(4))
3.1415(2)
>>> print(x.fmt(5))
3.14150(20)
>>> print(x.fmt())
3.14150(20)
```

where the argument is the number of decimal places retained. (Omit the argument and `fmt()` will adjust the number of decimal places automatically to display the error to two significant figures.)

Function `gvar.asgvar(x)` returns `x` if it is a `gvar.GVar`; otherwise it returns `gvar(x)`.

`gvar.GVars` are far more interesting when used to describe multidimensional distributions, especially if there are correlations between different variables. Such distributions are represented by collections of `gvar.GVars` in one of two standard formats: 1) numpy type arrays of `gvar.GVars` (any shape); or, more flexibly, 2) Python dictionaries whose values are `gvar.GVars` or arrays of `gvar.GVars`. Most functions in `gvar` that handle multiple `gvar.GVars` work with either format, and if they return multidimensional results do so in the same format as the inputs (that is, arrays or dictionaries). Any dictionary is converted internally into a specialized (ordered) dictionary of type `gvar.BufferDict`, and dictionary-valued results are also `gvar.BufferDicts`. `gvar.BufferDicts` are also useful for archiving `gvar.GVars`, since they may be pickled using Python's `pickle` module; `gvar.GVars` cannot be pickled otherwise. A pickled `gvar.BufferDict` preserves any correlations that exist between the different `gvar.GVars` in it.

To create an array of `gvar.GVars` with mean values specified by array `xmean` and covariance matrix `xcov`, use

```
x = gvar.gvar(xmean, xcov)
```

where array `x` has the same shape as `xmean` (and `xcov.shape = xmean.shape+xmean.shape`). Then each element `x[i]` of a one-dimensional array, for example, is a `gvar.GVar` where:

```
x[i].mean = xmean[i]           # mean of x[i]
x[i].val   = xmean[i]           # same as x[i].mean
x[i].sdev  = xcov[i, i]**0.5     # std deviation of x[i]
x[i].var   = xcov[i, i]         # variance of x[i]
```

`gvar.GVars` can be used in arithmetic expressions, just like Python floats. These expressions result in new `gvar.GVars` whose means and standard deviations are determined from the original covariance matrix. The arithmetic expressions can include calls to standard functions including: `exp`, `log`, `sqrt`, `sin`, `cos`, `tan`, `arcsin`, `arccos`, `arctan`, `sinh`, `cosh`, `tanh`, `arcsinh`, `arccosh`, `arctanh`.

As an example,

```
>>> x, y = gvar.gvar([0.1, 10.], [[0.015625, 0.], [0., 4.]])
>>> print('x =', x, 'y =', y)
x = 0.1 +- 0.125    y = 10 +- 2
```

makes `x` and `y` `gvar.GVars` with standard deviations `sigma_x=0.125` and `sigma_y=2`, and, in this case, no correlation between `x` and `y` (since `cov[i, j]=0` when `i!=j`). If now we set, for example,

```
>>> f = x + y
>>> print('f =', f)
f = 10.1 +- 2.0039
```

then `f` is a `gvar.GVar` with

```
f.var = df/dx cov[0, 0] df/dx + df/dx cov[0, 1] df/dy + ...
      = 2.0039**2
```

where `cov` is the original covariance matrix used to define `x` and `y` (in `gvar.gvar`). Note that while `f` and `y` separately have 20% uncertainties in this example, the ratio `f/y` has much smaller errors:

```
>>> print(f / y)
1.01 +- 0.012659
```

This happens, of course, because the errors in `f` and `y` are highly correlated (since the error in `f` comes mostly from `y`).

It is sometimes useful to know how much of the uncertainty in some quantity is due to a particular input uncertainty. Continuing the example above, for example, we might want to know how much of `f`'s standard deviation is due to the standard deviation of `x` and how much comes from `y`. This is easily computed (for the example above):

```
>>> print(f.partialsdev(x))      # uncertainty in f due to x
0.125
>>> print(f.partialsdev(y))      # uncertainty in f due to y
2.0
>>> print(f.partialsdev(x, y))   # uncertainty in f due to x and y
2.00390244274
>>> print(f.sdev)                # should be the same
2.00390244274
```

`gvar.gvar()` can also be used to convert strings or tuples stored in arrays or dictionaries into `gvar.GVars`: for example,

```
>>> garray = gvar.gvar(['2(1)', '10+-5', (99, 3), gvar.gvar(0, 2)])
>>> print(garray)
[2 +- 1 10 +- 5 99 +- 3 0 +- 2]
>>> gdict = gvar.gvar(dict(a='2(1)', b=['10+-5', (99, 3), gvar.gvar(0, 2)]))
>>> print(gdict)
{'a': 2 +- 1, 'b': array([10 +- 5, 99 +- 3, 0 +- 2], dtype=object)}
```

If the covariance matrix in `gvar.gvar` is diagonal, it can be replaced by an array of standard deviations (square roots of diagonal entries in `cov`). The example above, therefore, is equivalent to:

```
>>> x, y = gvar.gvar([0.1, 10.], [0.125, 2.])
>>> print('x =', x, 'y =', y)
x = 0.1 +- 0.125    y = 10 +- 2
```

2.3 Computing Covariance Matrices

The covariance matrix for a set of `gvar.GVars`, `g0 g1 ...`, can be computed using

```
gvar.evalcov([g0, g1...]) -> cov_g
```

where `cov_g[i, j]` gives the covariance between `gi` and `gj`. Instead of a list or array of `gs`, one can also give a dictionary `g` where `g[k]` is a `gvar.GVar`. In this case `gvar.evalcov()` returns a doubly-indexed dictionary `cov_g[k1][k2]` where keys `k1, k2` are in `g`.

Using the example from the previous section, the code

```
>>> x, y = gvar.gvar([0.1, 10.], [[0.015625, 0.], [0., 4.]])
>>> f = x+y
>>> print(gvar.evalcov([x, y, f]))
[[ 0.015625  0.          0.015625]
 [ 0.         4.         4.        ]
 [ 0.015625  4.         4.015625]]
```

confirms that x and y are uncorrelated with each other, but strongly correlated with f .

It is often convenient to group related `gvar.GVars` together in a dictionary rather than an array since dictionaries are far more flexible. `gvar.evalcov` can be used to evaluate the covariance matrix for a dictionary containing `gvar.GVars` and/or arbitrary arrays of `gvar.GVars`:

```
>>> d = dict(x=x, y=y, g=[x+y, x-y])
>>> cov = gvar.evalcov(d)
>>> print(cov['x', 'x'])
0.015625
>>> print(cov['x', 'y'])
0.0
>>> print(cov['x', 'g'])
[ 0.015625  0.015625]
```

2.4 Random Number Generators

`gvar.GVars` represent probability distributions. It is possible to use them to generate random numbers from those distributions. For example, in

```
>>> z = gvar.gvar(2.0, 0.5)
>>> print(z())
2.29895701465
>>> print(z())
3.00633184275
>>> print(z())
1.92649199321
```

calls to `z()` generate random numbers from a Gaussian random number generator with mean `z.mean=2.0` and standard deviation `z.sdev=0.5`.

To obtain random arrays from an array `g` of `gvar.GVars` use `giter=gvar.raniter(g)` (see `gvar.raniter()`) to create a random array generator `giter`. Each call to `next(giter)` generates a new array of random numbers. The random number arrays have the same shape as the array `g` of `gvar.GVars` and have the distribution implied by those random variables (including correlations). For example,

```
>>> a = gvar.gvar(1.0, 1.0)
>>> da = gvar.gvar(0.0, 0.1)
>>> g = [a, a+da]
>>> giter = gvar.raniter(g)
>>> print(next(giter))
[ 1.51874589  1.59987422]
>>> print(next(giter))
[-1.39755111 -1.24780937]
>>> print(next(giter))
[ 0.49840244  0.50643312]
```

Note how the two random numbers separately vary over the region 1 ± 1 (approximately), but the separation between the two is rarely more than 0 ± 0.1 . This is as expected given the strong correlation between `a` and `a+da`.

`gvar.raniter(g)` also works when `g` is a dictionary (or `gvar.BufferDict`) whose entries `g[k]` are `gvar.GVars` or arrays of `gvar.GVars`. In such cases the iterator returns a dictionary with the same layout:

```
>>> g = dict(a=gvar.gvar(0, 1), b=[gvar.gvar(0, 100), gvar.gvar(10, 1e-3)])
>>> print(g)
{'a': 0 +- 1, 'b': [0 +- 100, 10 +- 0.001]}
>>> giter = gvar.raniter(g)
>>> print(next(giter))
{'a': -0.88986130981173306, 'b': array([-67.02994213,  9.99973707])}
>>> print(next(giter))
{'a': 0.21289976681277872, 'b': array([ 29.9351328 , 10.00008606])}
```

One use for such random number generators is dealing with situations where the standard deviations are too large to justify the linearization assumed in defining functions of Gaussian variables. Consider, for example,

```
>>> x = gvar.gvar(1., 3.)
>>> print(cos(x))
0.540302 +- 2.52441
```

The standard deviation for `cos(x)` is obviously wrong since `cos(x)` can never be larger than one. To obtain the real mean and standard deviation, we generate a large number of random numbers `xi` from `x`, compute `cos(xi)` for each, and compute the mean and standard deviation for the resulting distribution (or any other statistical quantity, particularly if the resulting distribution is not Gaussian):

```
# estimate mean,sdev from 1000 random x's
>>> ran_x = numpy.array([x() for in range(1000)])
>>> ran_cos = numpy.cos(ran_x)
>>> print('mean =', ran_cos.mean(), ' std dev =', ran_cos.std())
mean = 0.0350548954142 std dev = 0.718647118869

# check by doing more (and different) random numbers
>>> ran_x = numpy.array([x() for in range(100000)])
>>> ran_cos = numpy.cos(ran_x)
>>> print('mean =', ran_cos.mean(), ' std dev =', ran_cos.std())
mean = 0.00806276057656 std dev = 0.706357174056
```

This procedure generalizes trivially for multidimensional analyses, using arrays or dictionaries with `gvar.raniter()`.

Finally note that *bootstrap* copies of `gvar.GVars` are easily created. A bootstrap copy of `gvar.GVar x ± dx` is another `gvar.GVar` with the same width but where the mean value is replaced by a random number drawn from the original distribution. Bootstrap copies of a data set, described by a collection of `gvar.GVars`, can be used as new (fake) data sets having the same statistical errors and correlations:

```
>>> g = gvar.gvar([1.1, 0.8], [[0.01, 0.005], [0.005, 0.01]])
>>> print(g)
[1.1 +- 0.1 0.8 +- 0.1]
>>> print(gvar.evalcov(g)) # print covariance matrix
[[ 0.01  0.005]
 [ 0.005  0.01 ]]
>>> gbs_iter = gvar.bootstrap_iter(g)
>>> gbs = next(gbs_iter) # bootstrap copy of f
>>> print(gbs)
[1.13881 +- 0.1 0.896066 +- 0.1] # different means
>>> print(gvar.evalcov(gbs)) # same covariance matrix
[[ 0.01  0.005]
 [ 0.005  0.01 ]]
```

Such fake data sets are useful for analyzing non-Gaussian behavior, for example, in nonlinear fits.

2.5 Limitations

The most fundamental limitation of this module is that the calculus of Gaussian variables that it assumes is only valid when standard deviations are small (compared to the distances over which the functions of interest change appreciably). One way of dealing with this limitation is described above in the section on *Random Number Generators*.

Another potential issue is roundoff error, which can become problematic if there is a wide range of standard deviations among correlated modes. For example, the following code works as expected:

```
>>> from gvar import gvar, evalcov
>>> tiny = 1e-4
>>> a = gvar(0., 1.)
>>> da = gvar(tiny, tiny)
>>> a, ada = gvar([a.mean, (a+da).mean], evalcov([a, a+da])) # = a, a+da
>>> print(ada-a)      # should be da again
0.0001 +- 0.0001
```

Reducing `tiny`, however, leads to problems:

```
>>> from gvar import gvar, evalcov
>>> tiny = 1e-8
>>> a = gvar(0., 1.)
>>> da = gvar(tiny, tiny)
>>> a, ada = gvar([a.mean, (a+da).mean], evalcov([a, a+da])) # = a, a+da
>>> print(ada-a)      # should be da again
1e-8 +- 0
```

Here the call to `gvar.evalcov()` creates a new covariance matrix for `a` and `ada = a+da`, but the matrix does not have enough numerical precision to encode the size of `da`'s variance, which gets set, in effect, to zero. The problem arises here for values of `tiny` less than about $2e-8$ (with 64-bit floating point numbers — `tiny**2` is what appears in the covariance matrix).

2.6 Implementation Notes; Derivatives; Optimizations

There are two types of `gvar.GVar`: the underlying independent variables, created with calls to `gvar.gvar()`; and variables which are obtained from functions of the underlying variables. Each `gvar.GVar` must keep track of three pieces of information: 1) its mean value; 2) its derivatives with respect to the underlying variables; and 3) the covariance matrix for the underlying variables. The derivatives and covariance matrix allow one to compute the standard deviation of the `gvar.GVar` as well as correlations between it and any other function of the underlying variables. A `gvar.GVar` can be constructed at a very low level by supplying all three pieces of information — for example,

```
f = gvar.gvar(fmean, fder, cov)
```

where `fmean` is the mean, `fder` is an array where `fder[i]` is the derivative of `f` with respect to the *i*-th underlying variable (numbered in the order in which they were created using `gvar.gvar()`), and `cov` is the covariance matrix for the underlying variables (easily obtained from an existing `gvar.GVar x` using `x.cov`).

The derivatives stored in a `gvar.GVar` are sometimes useful. Consider, for example, an array `x` each of whose elements was created by a call to `gvar.gvar()`: `x[i] = gvar.gvar(xi_mean, xi_sdev)`. Then derivatives of a function `f(x)` with respect to the `x[i]` can be computed from the `gvar.GVar` `fx = f(x)` using `fx.dotder(x[i].der)`, which equals $df(x)/dx[i]$ at the point `x` specified by the means of the `x[i]`s. Note that this trick only works because the `x[i]` are among the underlying (original) `gvar.GVars` (and not combinations of these).

When there are lots of underlying variables, the number of derivatives can become rather large, potentially (though not necessarily) leading to slower calculations. One way to alleviate this problem, should it arise, is to separate the underlying variables into groups that are never mixed in calculations and to use different `gvar.gvar()`s when generating the variables in different groups. New versions of `gvar.gvar()` are obtained using `gvar.switch_gvar()`: for example,

```
import gvar
...
x = gvar.gvar(...)
y = gvar.gvar(...)
z = f(x, y)
... other manipulations involving x and y ...
gvar.switch_gvar()
a = gvar(...)
b = gvar(...)
c = g(a, b)
... other manipulations involving a and b (but not x and y) ...
```

Here the `gvar.gvar()` used to create `a` and `b` is a different function than the one used to create `x` and `y`. A derived quantity, like `c`, knows about its derivatives with respect to `a` and `b`, and about their covariance matrix; but it carries no derivative information about `x` and `y`. Absent the `switch_gvar` line, `c` would have information about its derivatives with respect to `x` and `y` (zero derivative in both cases) and this would make calculations involving `c` slightly slower than with the `switch_gvar` line. Usually the difference is negligible — it used to be more important, in earlier implementations of `gvar.GVar` before sparse matrices were introduced to keep track of covariances. Note that the previous `gvar.gvar()` can be restored using `gvar.restore_gvar()`.

`gvar.GVars` are designed to work well with `numpy` arrays. They can be combined in arithmetic expressions with arrays of numbers or of `gvar.GVars`; the results in both cases are arrays of `gvar.GVars`.

Arithmetic operators `+` `-` `*` `/` `**` `==` `!=` `<>` `+=` `-=` `*=` `/=` are all defined. `gvar.GVars` are not ordered so `>` `>=` `<` `<=` are not defined. Two `gvar.GVars` are equal only if their means and derivatives are equal, and their covariance matrices the same. A `gvar.GVar` `x` is defined to equal a non-`gvar.GVar` `y` only if `x.mean == y` and `x.sdev == 0`.

2.7 Utilities

The function used to create Gaussian variable objects is:

`gvar.gvar(...)`

Create one or more new `gvar.GVars`.

Each of the following creates new `gvar.GVars`:

`gvar.gvar(x, xsdev)`

Returns a `gvar.GVar` with mean `x` and standard deviation `xsdev`. Returns an array of `gvar.GVars` if `x` and `xsdev` are arrays with the same shape; the shape of the result is the same as the shape of `x`.

`gvar.gvar(x, xcov)`

Returns an array of `gvar.GVars` with means given by array `x` and a covariance matrix given by array `xcov`, where `xcov.shape = 2*x.shape`. The result has the same shape as `x`.

`gvar.gvar((x, xsdev))`

Returns a `gvar.GVar` with mean `x` and standard deviation `xsdev`.

`gvar.gvar(xstr)`

Returns a `gvar.GVar` corresponding to string `xstr` which is either of the form `"xmean +- xsdev"` or `"x(xerr)"` (see `GVar.fmt()`).

`gvar.gvar(xgvar)`

Returns `gvar.GVar` `xgvar` unchanged.

`gvar.gvar(xdict)`

Returns a dictionary (`BufferDict`) `b` where `b[k] = gvar(xdict[k])` for every key in dictionary `xdict`. The values in `xdict`, therefore, can be strings, tuples or `gvar.GVars` (see above), or arrays of these.

`gvar.gvar(xarray)`

Returns an array `a` having the same shape as `xarray` where every element `a[i...] = gvar(xarray[i...])`. The values in `xarray`, therefore, can be strings, tuples or `gvar.GVars` (see above).

`gvar.gvar` is actually an object of type `gvar.GVarFactory`.

Means, standard deviations, variances, formatted strings, covariance matrices and correlation/comparison information can be extracted from arrays (or dictionaries) of `gvar.GVars` using:

`gvar.mean(g)`

Extract means from `gvar.GVars` in `g`.

`g` can be a `gvar.GVar`, an array of `gvar.GVars`, or a dictionary containing `gvar.GVars` or arrays of `gvar.GVars`. Result has the same layout as `g`.

`gvar.sdev(g)`

Extract standard deviations from `gvar.GVars` in `g`.

`g` can be a `gvar.GVar`, an array of `gvar.GVars`, or a dictionary containing `gvar.GVars` or arrays of `gvar.GVars`. Result has the same layout as `g`.

`gvar.var(g)`

Extract variances from `gvar.GVars` in `g`.

`g` can be a `gvar.GVar`, an array of `gvar.GVars`, or a dictionary containing `gvar.GVars` or arrays of `gvar.GVars`. Result has the same layout as `g`.

`gvar.fmt(g, ndecimal=None, sep='')`

Format `gvar.GVars` in `g`.

`g` can be a `gvar.GVar`, an array of `gvar.GVars`, or a dictionary containing `gvar.GVars` or arrays of `gvar.GVars`. Each `gvar.GVar` `gi` in `g` is replaced by the string generated by `gi.fmt(ndecimal, sep)`. Result has same structure as `g`.

`gvar.evalcov(g)`

Compute covariance matrix for elements of array/dictionary `g`.

If `g` is an array of `gvar.GVars`, `evalcov` returns the covariance matrix as an array with shape `g.shape+g.shape`. If `g` is a dictionary whose values are `gvar.GVars` or arrays of `gvar.GVars`, the result is a doubly-indexed dictionary where `cov[k1, k2]` is the covariance for `g[k1]` and `g[k2]`.

`gvar.uncorrelated(g1, g2)`

Return True if `gvar.GVars` in `g1` uncorrelated with those in `g2`.

`g1` and `g2` can be `gvar.GVars`, arrays of `gvar.GVars`, or dictionaries containing `gvar.GVars` or arrays of `gvar.GVars`.

`gvar.chi2(g1, g2)`

Compute χ^2 of `g1-g2`.

χ^2 is a measure of whether the multi-dimensional Gaussian distributions `g1` and `g2` (dictionaries or arrays) agree with each other — that is, do their means agree within errors for corresponding elements. The probability is high if $\chi^2(g1, g2) / \chi^2.dof$ is of order 1 or smaller.

Usually `g1` and `g2` are dictionaries with the same keys, where `g1[k]` and `g2[k]` are `gvar.GVars` or arrays of `gvar.GVars` having the same shape. Alternatively `g1` and `g2` can be `gvar.GVars`, or arrays of `gvar.GVars` having the same shape.

One of `g1` or `g2` can contain numbers instead of `gvar.GVars`, in which case `chi**2` is a measure of the likelihood that the numbers came from the distribution specified by the other argument.

One or the other of `g1` or `g2` can be missing keys, or missing elements from arrays. Only the parts of `g1` and `g2` that overlap are used. Also setting `g2=None` is equivalent to replacing its elements by zeros.

`chi**2` is computed from the inverse of the covariance matrix of `g1-g2`. The matrix inversion can be sensitive to roundoff errors. In such cases, SVD cuts can be applied by setting parameters `svdcut` and `svdnum`. See the documentation for `gvar.SVD` for information about these parameters.

The return value is the `chi**2`. Extra data is stored in `chi2` itself:

`chi2.dof`

Number of degrees of freedom (that is, the number of variables compared).

`chi2.Q`

The probability that the `chi**2` could have been larger, by chance, even if `g1` and `g2` agree. Values smaller than 0.1 or so suggest that they do not agree. Also called the *p-value*.

`gvar.GVars` contain information about derivatives with respect to the *independent* `gvar.GVars` from which they were constructed. This information can be extracted using:

`gvar.deriv(g, x)`

Compute first derivatives wrt `x` of `gvar.GVars` in `g`.

`g` can be a `gvar.GVar`, an array of `gvar.GVars`, or a dictionary containing `gvar.GVars` or arrays of `gvar.GVars`. Result has the same layout as `g`.

`x` must be an *independent* `gvar.GVar`, which is a `gvar.GVar` created by a call to `gvar.gvar()` (e.g., `x = gvar.gvar(xmean, xsdev)`) or a function `f(x)` of such a `gvar.GVar`. (More precisely, `x.der` must have only one nonzero entry.)

The following function creates an iterator that generates random arrays from the distribution defined by array (or dictionary) `g` of `gvar.GVars`. The random numbers incorporate any correlations implied by the `gs`.

`gvar.raniter(g, n=None, svdcut=None, svdnum=None, rescale=True)`

Return iterator for random samples from distribution `g`

The gaussian variables (`gvar.GVar` objects) in array (or dictionary) `g` collectively define a multidimensional gaussian distribution. The iterator defined by `raniter()` generates an array (or dictionary) containing random numbers drawn from that distribution, with correlations intact.

The layout for the result is the same as for `g`. So an array of the same shape is returned if `g` is an array. When `g` is a dictionary, individual entries `g[k]` may be `gvar.GVars` or arrays of `gvar.GVars`, with arbitrary shapes.

`raniter()` also works when `g` is a single `gvar.GVar`, in which case the resulting iterator returns random numbers drawn from the distribution specified by `g`.

Parameters

- **g** (array or dictionary or *BufferDict* or *GVar*) – An array (or dictionary) of objects of type `gvar.GVar`; or a `gvar.GVar`.
- **n** – Maximum number of random iterations. Setting `n=None` (the default) implies there is no maximum number.
- **svdcut** (None or number) – If positive, replace eigenvalues of the covariance matrix of `g` with `svdcut*(max eigenvalue)`; if negative, discards eigenmodes with eigenvalues smaller than `svdcut*(max eigenvalue)`; ignore if set to None.

- **svdnum** (*None* or positive *int*) – If positive, keep only the modes with the largest *svdnum* eigenvalues in the covariance matrix for *g*; ignore if set to *None* or negative.
- **rescale** (*bool*) – Covariance matrix is rescaled so that diagonal elements equal 1 before applying *svd* cuts if *rescale*=*True*.

Returns An iterator that returns random arrays or dictionaries with the same shape as *g* drawn from the gaussian distribution defined by *g*.

`gvar.bootstrap_iter(g, n=None, svdcut=None, svdnum=None, rescale=True)`

Return iterator for bootstrap copies of *g*.

The gaussian variables (`gvar.GVar` objects) in array (or dictionary) *g* collectively define a multidimensional gaussian distribution. The iterator created by `bootstrap_iter()` generates an array (or dictionary) of new `gvar.GVars` whose covariance matrix is the same as *g*'s but whose means are drawn at random from the original *g* distribution. This is a *bootstrap copy* of the original distribution. Each iteration of the iterator has different means (but the same covariance matrix).

`bootstrap_iter()` also works when *g* is a single `gvar.GVar`, in which case the resulting iterator returns bootstrap copies of the *g*.

Parameters

- **g** (*array or dictionary or BufferDict*) – An array (or dictionary) of objects of type `gvar.GVar`.
- **n** – Maximum number of random iterations. Setting *n*=*None* (the default) implies there is no maximum number.
- **svdcut** (*None* or number) – If positive, replace eigenvalues of the covariance matrix of *g* with *svdcut**(max eigenvalue); if negative, discards eigenmodes with eigenvalues smaller than *svdcut**(max eigenvalue); ignore if set to *None*.
- **svdnum** (*None* or positive *int*) – If positive, keep only the modes with the largest *svdnum* eigenvalues in the covariance matrix for *g*; ignore if set to *None* or negative.
- **rescale** (*bool*) – Covariance matrix is rescaled so that diagonal elements equal 1 before applying *svd* cuts if *rescale*=*True*.

Returns An iterator that returns bootstrap copies of *g*.

`gvar.ranseed(a)`

Seed random number generators with tuple *seed*.

Argument *seed* is a tuple of integers that is used to seed the random number generators used by *numpy* and *random* (and therefore by `gvar`). Reusing the same *seed* results in the same set of random numbers.

ranseed generates its own seed when called without an argument or with *seed*=*None*. This seed is stored in `ranseed.seed` and also returned by the function. The seed can be used to regenerate the same set of random numbers at a later time.

Parameters *seed* (*tuple or None*) – A tuple of integers. Generates a random tuple if *None*.

Returns The seed.

Two functions that are useful for tabulating results and for analyzing where the errors in a `gvar.GVar` constructed from other `gvar.GVars` come from:

`gvar.fmt_errorbudget(outputs, inputs, ndecimal=2, percent=True, colwidth=10)`

Tabulate error budget for `outputs[ko]` due to `inputs[ki]`.

For each output `outputs[ko]`, `fmt_errorbudget` computes the contributions to `outputs[ko]`'s standard deviation coming from the `gvar.GVars` collected in `inputs[ki]`. This is done for each key combination (*k_o*, *k_i*) and the results are tabulated with columns and rows labeled by *k_o* and *k_i*, respectively. If

a `gvar.GVar` in `inputs[ki]` is correlated with other `gvar.GVars`, the contribution from the others is included in the `ki` contribution as well (since contributions from correlated `gvar.GVars` cannot be resolved). The table is returned as a string.

Parameters

- **outputs** – Dictionary of `gvar.GVars` for which an error budget is computed.
- **inputs** – Dictionary of: `gvar.GVars`, arrays/dictionaries of `gvar.GVars`, or lists of `gvar.GVars` and/or arrays/dictionaries of `gvar.GVars`. `fmt_errorbudget` tabulates the parts of the standard deviations of each `outputs[ko]` due to each `inputs[ki]`.
- **ndecimal** (`int`) – Number of decimal places displayed in table.
- **percent** (`boolean`) – Tabulate % errors if `percent` is `True`; otherwise tabulate the errors themselves.
- **colwidth** (*positive integer*) – Width of each column.

Returns A table (`str`) containing the error budget. Output variables are labeled by the keys in `outputs` (columns); sources of uncertainty are labeled by the keys in `inputs` (rows).

`gvar.fmt_values (outputs, ndecimal=None)`
Tabulate `gvar.GVars` in `outputs`.

Parameters

- **outputs** – A dictionary of `gvar.GVar` objects.
- **ndecimal** (`int` or `None`) – Format values `v` using `v.fmt (ndecimal)`.

Returns A table (`str`) containing values and standard deviations for variables in `outputs`, labeled by the keys in `outputs`.

The following functions creates new functions that generate `gvar.GVars` (to replace `gvar.gvar()`):

`gvar.switch_gvar()`
Switch `gvar.gvar()` to new `gvar.GVarFactory`.

Returns New `gvar.gvar()`.

`gvar.restore_gvar()`
Restore previous `gvar.gvar()`.

Returns Previous `gvar.gvar()`.

`gvar.gvar_factory (cov=None)`
Return new function for creating `gvar.GVars` (to replace `gvar.gvar()`).

If `cov` is specified, it is used as the covariance matrix for new `gvar.GVars` created by the function returned by `gvar_factory (cov)`. Otherwise a new covariance matrix is created internally.

`gvar.GVars` created by different functions cannot be combined in arithmetic expressions (the error message “In-compatible GVars.” results).

The following function can be used to rebuild collections of `gvar.GVars`, ignoring all correlations with other variables. It can also be used to introduce correlations between uncorrelated variables.

`gvar.rebuild (g, gvar=gvar, corr=0.0)`
Rebuild `g` stripping correlations with variables not in `g`.

`g` is either an array of `gvar.GVars` or a dictionary containing `gvar.GVars` and/or arrays of `gvar.GVars`. `rebuild (g)` creates a new collection `gvar.GVars` with the same layout, means and covariance matrix as those in `g`, but discarding all correlations with variables not in `g`.

If `corr` is nonzero, `rebuild` will introduce correlations wherever there aren’t any using


```
cov[i,j] -> corr * sqrt(cov[i,i]*cov[j,j])
```

wherever `cov[i,j]==0.0` initially. Positive values for `corr` introduce positive correlations, negative values anti-correlations.

Parameter `gvar` specifies a function for creating new `gvar.GVars` that replaces `gvar.gvar()` (the default).

Parameters

- **g** (array or dictionary) – `gvar.GVars` to be rebuilt.
- **gvar** (`gvar.GVarFactory` or `None`) – Replacement for `gvar.gvar()` to use in rebuilding. Default is `gvar.gvar()`.
- **corr** (number) – Size of correlations to introduce where none exist initially.

Returns Array or dictionary (`gvar.BufferDict`) of `gvar.GVars` (same layout as `g`) where all correlations with variables other than those in `g` are erased.

Finally there is a utility function and a class for implementing an *svd* analysis of a covariance or other symmetric, positive matrix:

```
gvar.svd(g, svdcut=None, svdnum=None, compute_delta=False, rescale=True)
```

Apply *svd* cuts to collection of `gvar.GVars` in `g`.

`g` is an array of `gvar.GVars` or a dictionary containing `gvar.GVars` and/or arrays of `gvar.GVars`. `svd(g, ...)` returns a copy of `g` whose `gvar.GVars` have been modified so that their covariance matrix is less singular than for the original `g` (the `gvar.GVar` means are unchanged). This is done using an *svd* algorithm which is controlled by three parameters: `svdcut`, `svdnum` and `rescale` (see `gvar.SVD` for more details). *svd* cuts are not applied when the covariance matrix is diagonal (that is, when there are no correlations between different elements of `g`).

The input parameters are :

Parameters

- **g** – An array of `gvar.GVars` or a dictionary whose values are `gvar.GVars` and/or arrays of `gvar.GVars`.
- **svdcut** (`None` or number (`|svdcut|<=1`)) – If positive, replace eigenvalues of the covariance matrix with `svdcut*(max eigenvalue)`; if negative, discard eigenmodes with eigenvalues smaller than `svdcut` times the maximum eigenvalue. Default is `None`.
- **svdnum** (`None` or `int`) – If positive, keep only the modes with the largest `svdnum` eigenvalues; ignore if set to `None`. Default is `None`.
- **rescale** – Rescale the input matrix to make its diagonal elements equal to 1.0 before applying *svd* cuts. (Default is `True`.)
- **compute_inv** – Compute representation of inverse of covariance matrix if `True`; the result is stored in `svd.inv_wgt` (see below). Default value is `False`.

Returns A copy of `g` with the same means but with a covariance matrix modified by *svd* cuts.

Data from the *svd* analysis of `g`'s covariance matrix is stored in `svd` itself:

`svd.val`

Eigenvalues of the covariance matrix after *svd* cuts (and after rescaling if `rescale=True`); the eigenvalues are ordered, with the smallest first.

`svd.vec`

Eigenvectors of the covariance matrix after *svd* cuts (and after rescaling if `rescale=True`), where `svd.vec[i]` is the vector corresponding to `svd.val[i]`.

svd.eigen_range
Ratio of the smallest to largest eigenvalue before *svd* cuts are applied (but after rescaling if *rescale=True*).

svd.D
Diagonal of matrix used to rescale the covariance matrix before applying *svd* cuts (cuts are applied to $D \cdot \text{cov} \cdot D$) if *rescale=True*; *svd.D* is None if *rescale=False*.

svd.logdet
Logarithm of the determinant of the covariance matrix after *svd* cuts are applied.

svd.correction
Vector of the *svd* corrections to *g.flat*;

svd.inv_wgt
The sum of the outer product of vectors *inv_wgt[i]* with themselves equals the inverse of the covariance matrix after *svd* cuts. Only computed if *compute_inv=True*. The order of the vectors is reversed relative to *svd.val* and *svd.vec*

2.8 Classes

The fundamental class for representing Gaussian variables is:

class *gvar.GVar*

The basic attributes are:

mean
Mean value.

sdev
Standard deviation.

var
Variance.

Two methods allow one to isolate the contributions to the variance or standard deviation coming from other *gvar.GVars*:

partialvar (*args)

Compute partial variance due to *gvar.GVars* in args.

This method computes the part of *self.var* due to the *gvar.GVars* in args. If *args[i]* is correlated with other *gvar.GVars*, the variance coming from these is included in the result as well. (This last convention is necessary because variances associated with correlated *gvar.GVars* cannot be disentangled into contributions corresponding to each variable separately.)

Parameters *args[i]* (*gvar.GVar* or array/dictionary of *gvar.GVars*) – Variables contributing to the partial variance.

Returns Partial variance due to all of args.

partialsdev (*args)

Compute partial standard deviation due to *gvar.GVars* in args.

This method computes the part of *self.sdev* due to the *gvar.GVars* in args. If *args[i]* is correlated with other *gvar.GVars*, the standard deviation coming from these is included in the result as well. (This last convention is necessary because variances associated with correlated *gvar.GVars* cannot be disentangled into contributions corresponding to each variable separately.)

Parameters `args[i]` (`gvar.GVar` or array/dictionary of `gvar.GVars`) – Variables contributing to the partial standard deviation.

Returns Partial standard deviation due to `args`.

Partial derivatives of the `gvar.GVar` with respect to the independent `gvar.GVars` from which it was constructed are given by:

deriv (*x*)

Derivative of `self` with respect to *independent* `gvar.GVar` *x*.

x must be an *independent* `gvar.GVar`, which is a `gvar.GVar` created by a call to `gvar.gvar()` (e.g., `x = gvar.gvar(xmean, xsdev)`) or a function `f(x)` of such a `gvar.GVar`. (More precisely, `x.der` must have only one nonzero entry.)

All `gvar.GVars` are constructed from a set of independent `gvar.GVars`. `self.deriv(x)` returns the partial derivative of `self` with respect to independent `gvar.GVar` *x*, holding all of the other independent `gvar.GVars` constant.

Parameters *x* – The independent `gvar.GVar`.

Returns The derivative of `self` with respect to *x*.

There are two methods for converting `self` into a string, for printing:

__str__ ()

Convert to string with format: mean +- std-dev.

fmt (*ndecimal=None, sep=''*)

Convert to string with format: mean (sdev).

Leading zeros in the standard deviation are omitted: for example, 25.67 +- 0.02 becomes 25.67(2). Parameter `ndecimal` specifies how many digits follow the decimal point in the mean. Parameter `sep` is a string that is inserted between the mean and the (sdev). If `ndecimal` is `None` (default), it is set automatically to the larger of `int(2-log10(self.sdev))` or 0; this will display at least two digits of error. Very large or very small numbers are written with exponential notation when `ndecimal` is `None`.

Two attributes and a method make reference to the original variables from which `self` is derived:

cov

Underlying covariance matrix (type `gvar.smat`) shared by all `gvar.GVars`.

der

Array of derivatives with respect to underlying (original) `gvar.GVars`.

dotder (*v*)

Return the dot product of `self.der` and *v*.

The following class is a specialized form of an ordered dictionary for holding `gvar.GVars` (or other scalars) and arrays of `gvar.GVars` (or other scalars) that supports Python pickling:

class `gvar.BufferDict`

Dictionary whose data is packed into a 1-d buffer (`numpy.array`).

A `gvar.BufferDict` object is a dictionary-like object whose values must either be scalars or arrays (like `numpy` arrays, with arbitrary shapes). The scalars and arrays are assembled into different parts of a single one-dimensional buffer. The various scalars and arrays are retrieved using keys, as in a dictionary: e.g.,

```
>>> a = BufferDict()
>>> a['scalar'] = 0.0
>>> a['vector'] = [1., 2.]
>>> a['tensor'] = [[3., 4.], [5., 6.]]
```

```

>>> print(a.flatten())           # print a's buffer
[ 0.  1.  2.  3.  4.  5.  6.]
>>> for k in a:                   # iterate over keys in a
...     print(k,a[k])
scalar 0.0
vector [ 1.  2.]
tensor [[ 3.  4.]
        [ 5.  6.]]
>>> a['vector'] = a['vector']*10   # change the 'vector' part of a
>>> print(a.flatten())
[ 0. 10. 20.  3.  4.  5.  6.]

```

The first four lines here could have been collapsed to one statement:

```

a = BufferDict(scalar=0.0,vector=[1.,2.],tensor=[[3.,4.],[5.,6.]])

or

a = BufferDict([('scalar',0.0),('vector',[1.,2.]),
               ('tensor',[[3.,4.],[5.,6.]])])

```

where in the second case the order of the keys is preserved in a (that is, BufferDict is an ordered dictionary).

The keys and associated shapes in a `gvar.BufferDict` can be transferred to a different buffer, creating a new `gvar.BufferDict`: e.g., using a from above,

```

>>> buf = numpy.array([0.,10.,20.,30.,40.,50.,60.])
>>> b = BufferDict(a,buf=buf)      # clone a but with new buffer
>>> print(b['tensor'])
[[ 30.  40.]
 [ 50.  60.]]
>>> b['scalar'] += 1
>>> print(buf)
[ 1. 10. 20. 30. 40. 50. 60.]

```

Note how b references buf and can modify it. One can also replace the buffer in the original `gvar.BufferDict` using, for example, `a.buf = buf`:

```

>>> a.buf = buf
>>> print(a['tensor'])
[[ 30.  40.]
 [ 50.  60.]]
>>> a['tensor'] *= 10.
>>> print(buf)
[ 1. 10. 20. 300. 400. 500. 600.]

```

`a.buf` is the numpy array used for a's buffer. It can be used to access and change the buffer directly. In `a.buf = buf`, the new buffer `buf` must be a numpy array of the correct shape. The buffer can also be accessed through iterator `a.flat` (in analogy with numpy arrays), and through `a.flatten()` which returns a copy of the buffer.

A `gvar.BufferDict` functions like a dictionary except: a) items cannot be deleted once inserted; b) all values must be either scalars or arrays of scalars, where the scalars can be any noniterable type that works with numpy arrays; and c) any new value assigned to a key must have the same size and shape as the original value.

Note that `gvar.BufferDicts` can be pickled and unpickled even when they store `gvar.GVars` (which themselves cannot be pickled separately).

The main attributes are:

size
Size of buffer array.

flat
Buffer array iterator.

dtype
Data type of buffer array elements.

buf
The (1d) buffer array. Allows direct access to the buffer: for example, `self.buf[i] = new_val` sets the value of the *i*-th element in the buffer to value `new_val`. Setting `self.buf = nbuf` replaces the old buffer by new buffer `nbuf`. This only works if `nbuf` is a one-dimensional `numpy` array having the same length as the old buffer, since `nbuf` itself is used as the new buffer (not a copy).

shape
Always equal to `None`. This attribute is included since `gvar.BufferDicts` share several attributes with `numpy` arrays to simplify coding that might support either type. Being dictionaries they do not have shapes in the sense of `numpy` arrays (hence the shape is `None`).

The main methods are:

flatten()
Copy of buffer array.

slice(*k*)
Return slice/index in `self.flat` corresponding to key *k*.

isscalar(*k*)
Return `True` if `self[k]` is scalar else `False`.

update(*d*)
Add contents of dictionary *d* to `self`.

static load(*fobj*, use_json=False)
Load serialized `gvar.BufferDict` from file object *fobj*. Uses `pickle` unless `use_json` is `True`, in which case it uses `json` (obviously).

static loads(*s*, use_json=False)
Load serialized `gvar.BufferDict` from string object *s*. Uses `pickle` unless `use_json` is `True`, in which case it uses `json` (obviously).

dump(*fobj*, use_json=False)
Serialize `gvar.BufferDict` in file object *fobj*.

Uses `pickle` unless `use_json` is `True`, in which case it uses `json` (obviously). `json` does not handle non-string valued keys very well. This attempts a workaround, but it will only work in simpler cases. Serialization only works when `pickle` (or `json`) knows how to serialize the data type stored in the `gvar.BufferDict`'s buffer (or for `gvar.GVars`).

dumps(*use_json=False*)
Serialize `gvar.BufferDict` into string.

Uses `pickle` unless `use_json` is `True`, in which case it uses `json` (obviously). `json` does not handle non-string valued keys very well. This attempts a workaround, but it will only work in simpler cases (e.g., integers, tuples of integers, etc.). Serialization only works when `pickle` (or `json`) knows how to serialize the data type stored in the `gvar.BufferDict`'s buffer (or for `gvar.GVars`).

SVD analysis is handled by the following class:

```
class gvar.SVD(mat, svdcut=None, svdnum=None, compute_delta=False, rescale=False)
    SVD decomposition of a pos. sym. matrix.
```

`SVD` is a function-class that computes the eigenvalues and eigenvectors of a positive symmetric matrix `mat`. Eigenvalues that are small (or negative, because of roundoff) can be eliminated or modified using `svd` cuts. Typical usage is:

```
>>> mat = [[1.,.25],[.25,2.]]
>>> s = SVD(mat)
>>> print(s.val)           # eigenvalues
[ 0.94098301  2.05901699]
>>> print(s.vec[0])        # 1st eigenvector (for s.val[0])
[ 0.97324899 -0.22975292]
>>> print(s.vec[1])        # 2nd eigenvector (for s.val[1])
[ 0.22975292  0.97324899]

>>> s = SVD(mat,svdcut=0.6) # force s.val[i]>=s.val[-1]*0.6
>>> print(s.val)
[ 1.2354102   2.05901699]
>>> print(s.vec[0])        # eigenvector unchanged
[ 0.97324899 -0.22975292]

>>> s = SVD(mat)
>>> w = s.decomp(-1)        # decomposition of inverse of mat
>>> invmat = sum(numpy.outer(wj,wj) for wj in w)
>>> print(numpy.dot(mat,invmat)) # should be unit matrix
[[ 1.00000000e+00  2.77555756e-17]
 [ 1.66533454e-16  1.00000000e+00]]
```

Input parameters are:

Parameters

- **mat** (2-d sequence (`numpy.array` or `list` or ...)) – Positive, symmetric matrix.
- **svdcut** (None or number (`|svdcut|<=1`)) – If positive, replace eigenvalues of `mat` with `svdcut*(max eigenvalue)`; if negative, discard eigenmodes with eigenvalues smaller than `svdcut` times the maximum eigenvalue.
- **svdnum** (None or int) – If positive, keep only the modes with the largest `svdnum` eigenvalues; ignore if set to None.
- **compute_delta** (*boolean*) – Compute `delta` (see below) if True; set `delta=None` otherwise.
- **rescale** – Rescale the input matrix to make its diagonal elements equal to 1.0 before diagonalizing.

The results are accessed using:

val

An ordered array containing the eigenvalues of `mat`. Note that `val[i]<=val[i+1]`.

vec

Eigenvectors `vec[i]` corresponding to the eigenvalues `val[i]`.

D

The diagonal matrix used to precondition the input matrix if `rescale==True`. The matrix diagonalized is $D M D$ where M is the input matrix. D is stored as a one-dimensional vector of diagonal elements. D is None if `rescale==False`.

nmod

The first `nmod` eigenvalues in `self.val` were modified by the `SVD` cut (equals 0 unless `svdcut > 0`).

kappa

Ratio of the smallest to the largest eigenvector in the unconditioned matrix (after rescaling if `rescale=True`)

delta

A vector of `gvars` whose means are zero and whose covariance matrix is what was added to `mat` to condition its eigenvalues. Is `None` if `svdcut<0` or `compute_delta==False`.

decomp (*n*)

Vector decomposition of input matrix raised to power *n*.

Computes vectors `w[i]` such that

$$\text{mat}^{**n} = \sum_i \text{numpy.outer}(w[i], w[i])$$

where `mat` is the original input matrix to `svd`. This decomposition cannot be computed if the input matrix was rescaled (`rescale=True`) except for `n=1` and `n=-1`.

Parameters *n* (*number*) – Power of input matrix.

Returns Array *w* of vectors.

2.9 Requirements

`gvar` makes heavy use of `numpy` for array manipulations. It also uses the `numpy` code for implementing elementary functions (*e.g.*, `sin`, `exp` ...) in terms of member functions.

GVAR.DATASET - RANDOM DATA SETS

3.1 Introduction

`gvar.dataset` contains a several tools for collecting and analyzing random samples from arbitrary distributions. The random samples are represented by lists of numbers or arrays, where each number/array is a new sample from the underlying distribution. For example, six samples from a one-dimensional gaussian distribution ($1+-1$) might look like

```
>>> random_numbers = [1.739, 2.682, 2.493, -0.460, 0.603, 0.800]
```

while six samples from a two-dimensional distribution ($[1+-1, 2+-1]$) might be

```
>>> random_arrays = [[ 0.494, 2.734], [ 0.172, 1.400], [ 1.571, 1.304],  
...                  [ 1.532, 1.510], [ 0.669, 0.873], [ 1.242, 2.188]]
```

Samples from more complicated multidimensional distributions are represented by dictionaries whose values are lists of numbers or arrays: for example,

```
>>> random_dict = dict(n=random_numbers, a=random_arrays)
```

where list elements `random_dict['n'][i]` and `random_dict['a'][i]` are part of the same multidimensional sample for every `i` — that is, the lists for different keys in the dictionary are synchronized one with the other.

With large samples, we typically want to estimate the mean value of the underlying distribution. This is done using `gvar.dataset.avg_data()`: for example,

```
>>> print(avg_data(random_numbers))  
1.3095 +- 0.452117
```

indicates that 1.31 (45) is our best guess, based only upon the samples in `random_numbers`, for the mean of the distribution from which those samples were drawn. Similarly

```
>>> print(avg_data(random_arrays))  
[0.946667 +- 0.217418 1.66817 +- 0.251002]
```

indicates that the means for the two-dimensional distribution behind `random_arrays` are `[0.95(22), 1.67(25)]`. `avg_data()` can also be applied to a dictionary whose values are lists of numbers/arrays: for example,

```
>>> print(avg_data(random_dict))  
{ 'a': array([0.946667 +- 0.217418, 1.66817 +- 0.251002], dtype=object),  
  'n': 1.3095 +- 0.452117 }
```

Class `gvar.dataset.Dataset` can be used to assemble dictionaries containing random samples. For example, imagine that the random samples above were originally written into a file, as they were generated:

```
# file: datafile
n 1.739
a [ 0.494, 2.734]
n 2.682
a [ 0.172, 1.400]
n 2.493
a [ 1.571, 1.304]
n -0.460
a [ 1.532, 1.510]
n 0.603
a [ 0.669, 0.873]
n 0.800
a [ 1.242, 2.188]
```

Here each line is a different random sample, either from the one-dimensional distribution (labeled `n`) or from the two-dimensional distribution (labeled `a`). Assuming the file is called `datafile`, this data can be read into a dictionary, essentially identical to the data dictionary above, using:

```
>>> data = Dataset("datafile")
>>> print(data['a'])
[array([ 0.494, 2.734]), array([ 0.172, 1.400]), array([ 1.571, 1.304]) ... ]
>>> print(avg_data(data['n']))
1.3095 +- 0.452117
```

The brackets and commas can be omitted in the input file for one-dimensional arrays: for example, `datafile` (above) could equivalently be written

```
# file: datafile
n 1.739
a 0.494 2.734
n 2.682
a 0.172 1.400
...
```

Other data formats may also be easy to use. For example, a data file written using `yaml` would look like

```
# file: datafile
---
n: 1.739
a: [ 0.494, 2.734]
---
n: 2.682
a: [ 0.172, 1.400]
.
.
.
```

and could be read into a `gvar.dataset.Dataset` using:

```
import yaml

data = Dataset()
with open("datafile", "r") as dfile:
    for d in yaml.load_all(dfile.read()): # iterate over yaml records
        data.append(d)                  # d is a dictionary
```

Finally note that data can be binned, into bins of size `binsize`, using `gvar.dataset.bin_data()`. For example, `gvar.dataset.bin_data(data, binsize=3)` replaces every three samples in `data` by the average of those samples. This creates a dataset that is $1/3$ the size of the original but has the same mean. Binning is use-

ful for making large datasets more manageable, and also for removing sample-to-sample correlations. Over-binning, however, erases statistical information.

Class `gvar.dataset.Dataset` can also be used to build a dataset sample by sample in code: for example,

```
>>> a = Dataset()
>>> a.append(n=1.739, a=[ 0.494, 2.734])
>>> a.append(n=2.682, a=[ 0.172, 1.400])
...
```

creates the same dataset as above.

3.2 Functions

The functions defined in the module are:

`gvar.dataset.avg_data(data, median=False, spread=False, bstrap=False)`

Average random data to estimate mean.

`data` is a list of random numbers or random arrays, or a dictionary of lists of random numbers/arrays: for example,

```
>>> random_numbers = [1.60, 0.99, 1.28, 1.30, 0.54, 2.15]
>>> random_arrays = [[12.2, 121.3], [13.4, 149.2], [11.7, 135.3],
...                  [7.2, 64.6], [15.2, 69.0], [8.3, 108.3]]
>>> random_dict = dict(n=random_numbers, a=random_arrays)
```

where in each case there are six random numbers/arrays. `avg_data` estimates the means of the distributions from which the random numbers/arrays are drawn, together with the uncertainties in those estimates. The results are returned as a `gvar.GVar` or an array of `gvar.GVars`, or a dictionary of `gvar.GVars` and/or arrays of `gvar.GVars`:

```
>>> print(avg_data(random_numbers))
1.31 +- 0.203169
>>> print(avg_data(random_arrays))
[11.3333 +- 1.13521 107.95 +- 12.936]
>>> print(avg_data(random_dict))
{'n': 1.31 +- 0.203169, 'a': array([11.3333 +- 1.13521, 107.95 +- 12.936], dtype=object)}
```

The arrays in `random_arrays` are one dimensional; in general, they can have any shape.

`avg_data(data)` also estimates any correlations between different quantities in `data`. When `data` is a dictionary, it does this by assuming that the lists of random numbers/arrays for the different `data[k]`s are synchronized, with the first element in one list corresponding to the first elements in all other lists, and so on. If some lists are shorter than others, the longer lists are truncated to the same length as the shortest list (discarding data samples).

There are four optional arguments. If argument `spread=True` each standard deviation in the results refers to the spread in the data, not the uncertainty in the estimate of the mean. The former is \sqrt{N} larger where N is the number of random numbers (or arrays) being averaged:

```
>>> print(avg_data(random_numbers, spread=True))
1.31 +- 0.497661
>>> print(avg_data(random_numbers))
1.31 +- 0.203169
>>> >>> print((0.497661/0.203169)**2)    # should be 6
6.00001491255
```

This is useful, for example, when averaging bootstrap data. The default value is `spread=False`.

The second option is triggered by setting `median=True`. This replaces the means in the results by medians, while the standard deviations are approximated by the half-width of the interval, centered around the median, that contains 68% of the data. These estimates are more robust than the mean and standard deviation when averaging over small amounts of data; in particular, they are unaffected by extreme outliers in the data. The default is `median=False`.

The third option is triggered by setting `bstrap=True`. This is shorthand for setting `median=True` and `spread=True`, and overrides any explicit setting for these keyword arguments. This is the typical choice for analyzing bootstrap data — hence its name. The default value is `bstrap=False`.

The final option is to omit the error estimates on the averages, which is triggered by setting `noerror=True`. Just the mean values are returned. The default value is `noerror=False`.

`gvar.dataset.autocorr(data, ncorr=None)`

Compute autocorrelation in random data.

`data` is a list of random numbers or random arrays, or a dictionary of lists of random numbers/arrays.

When `data` is a list of random numbers, `autocorr(data)` returns an array where `autocorr(data)[i]` is the correlation between elements in `data` that are separated by distance `i` in the list: for example,

```
>>> print(autocorr([2,-2,2,-2,2,-2]))
[ 1. -1.  1. -1.  1. -1.]
```

shows perfect correlation between elements separated by an even interval in the list, and perfect anticorrelation between elements by an odd interval.

`autocorr(data)` returns a list of arrays of autocorrelation coefficients when `data` is a list of random arrays. Again `autocorr(data)[i]` gives the autocorrelations for `data` elements separated by distance `i` in the list. Similarly `autocorr(data)` returns a dictionary when `data` is a dictionary.

`autocorr(data)` uses FFTs to compute the autocorrelations; the cost of computing the autocorrelations should grow roughly linearly with the number of random samples in `data` (up to logarithms).

`gvar.dataset.bin_data(data, binsize=2)`

Bin random data.

`data` is a list of random numbers or random arrays, or a dictionary of lists of random numbers/arrays. `bin_data(data, binsize)` replaces consecutive groups of `binsize` numbers/arrays by the average of those numbers/arrays. The result is new data list (or dictionary) with $1/\text{binsize}$ times as much random data: for example,

```
>>> print(bin_data([1,2,3,4,5,6,7], binsize=2))
[1.5, 3.5, 5.5]
>>> print(bin_data(dict(s=[1,2,3,4,5], v=[[1,2], [3,4], [5,6], [7,8]]), binsize=2))
{'s': [1.5, 3.5], 'v': [array([ 2.,  3.]), array([ 6.,  7.])]}
```

Data is dropped at the end if there is insufficient data to form complete bins. Binning is used to make calculations faster and to reduce measurement-to-measurement correlations, if they exist. Over-binning erases useful information.

`gvar.dataset.bootstrap_iter(data, n=None)`

Create iterator that returns bootstrap copies of `data`.

`data` is a list of random numbers or random arrays, or a dictionary of lists of random numbers/arrays. `bootstrap_iter(data, n)` is an iterator that returns `n` bootstrap copies of `data`. The random numbers/arrays in a bootstrap copy are drawn at random (with repetition allowed) from among the samples in `data`: for example,

```

>>> data = [1.1, 2.3, 0.5, 1.9]
>>> data_iter = bootstrap_iter(data)
>>> print(next(data_iter))
[ 1.1  1.1  0.5  1.9]
>>> print(next(data_iter))
[ 0.5  2.3  1.9  0.5]

>>> data = dict(a=[1,2,3,4],b=[1,2,3,4])
>>> data_iter = bootstrap_iter(data)
>>> print(next(data_iter))
{'a': array([3, 3, 1, 2]), 'b': array([3, 3, 1, 2])}
>>> print(next(data_iter))
{'a': array([1, 3, 3, 2]), 'b': array([1, 3, 3, 2])}

>>> data = [[1,2],[3,4],[5,6],[7,8]]
>>> data_iter = bootstrap_iter(data)
>>> print(next(data_iter))
[[ 7.  8.]
 [ 1.  2.]
 [ 1.  2.]
 [ 7.  8.]]
>>> print(next(data_iter))
[[ 3.  4.]
 [ 7.  8.]
 [ 3.  4.]
 [ 1.  2.]]

```

The distribution of bootstrap copies is an approximation to the distribution from which `data` was drawn. Consequently means, variances and correlations for bootstrap copies should be similar to those in `data`. Analyzing variations from bootstrap copy to copy is often useful when dealing with non-gaussian behavior or complicated correlations between different quantities.

Parameter `n` specifies the maximum number of copies; there is no maximum if `n` is `None`.

3.3 Classes

`gvar.dataset.Dataset` is used to assemble random samples from multidimensional distributions:

class `gvar.dataset.Dataset`

Dictionary for collecting random data.

This dictionary class simplifies the collection of random data. The random data are stored in a dictionary, with each piece of random data being a number or an array of numbers. For example, consider a situation where there are four random values for a scalar `s` and four random values for vector `v`. These can be collected as follows:

```

>>> data = Dataset()
>>> data.append(s=1.1,v=[12.2,20.6])
>>> data.append(s=0.8,v=[14.1,19.2])
>>> data.append(s=0.95,v=[10.3,19.7])
>>> data.append(s=0.91,v=[8.2,21.0])
>>> print(data['s'])           # 4 random values of s
[ 1.1, 0.8, 0.95, 0.91]
>>> print(data['v'])           # 4 random vector-values of v
[array([ 12.2,  20.6]), array([ 14.1,  19.2]), array([ 10.3,  19.7]), array([  8.2,  21. ])]

```

The argument to `data.append()` could be a dictionary: for example, `dd = dict(s=1.1,v=[12.2,20.6]); data.append(dd)` is equivalent to the first `append` statement

above. This is useful, for example, if the data comes from a function (that returns a dictionary).

One can also append data key-by-key: for example, `data.append('s', 1.1)`; `data.append('v', [12.2, 20.6])` is equivalent to the first append in the example above. One could also achieve this with, for example, `data['s'].append(1.1)`; `data['v'].append([12.2, 20.6])`, since each dictionary value is a list, but `gvar.Dataset`'s `append` checks for consistency between the new data and data already collected and so is preferable.

Use `extend` in place of `append` to add data in batches: for example,

```
>>> data = Dataset()
>>> data.extend(s=[1.1, 0.8], v=[[12.2, 20.6], [14.1, 19.2]])
>>> data.extend(s=[0.95, 0.91], v=[[10.3, 19.7], [8.2, 21.0]])
>>> print(data['s'])      # 4 random values of s
[ 1.1, 0.8, 0.95, 0.91]
```

gives the same dataset as the first example above.

A `Dataset` can also be created from a file where every line is a new random sample. The data in the first example above could have been stored in a file with the following content:

```
# file: datafile
s 1.1
v [12.2, 20.6]
s 0.8
v [14.1, 19.2]
s 0.95
v [10.3, 19.7]
s 0.91
v [8.2, 21.0]
```

Lines that begin with `#` are ignored. Assuming the file is called `datafile`, we create a dataset identical to that above using the code:

```
>>> data = Dataset('datafile')
>>> print(data['s'])
[ 1.1, 0.8, 0.95, 0.91]
```

Data can be binned while reading it in, which might be useful if there the data set is huge. To bin the data contained in file `datafile` in bins of `binsize 2` we use:

```
>>> data = Dataset('datafile', binsize=2)
>>> print(data['s'])
[0.95, 0.93]
```

Finally the keys read from a data file are restricted to those listed in keyword `keys` and those that are matched (or partially matched) by regular expression `grep` if one or the other of these is specified: for example,

```
>>> data = Dataset('datafile')
>>> print([k for k in a])
['s', 'v']
>>> data = Dataset('datafile', keys=['v'])
>>> print([k for k in a])
['v']
>>> data = Dataset('datafile', grep='[^v]')
>>> print([k for k in a])
['s']
>>> data = Dataset('datafile', keys=['v'], grep='[^v]')
>>> print([k for k in a])
[]
```

The main attributes and methods are:

samplesize

Smallest number of samples for any key.

append (*args, **kwargs)

Append data to dataset.

There are three equivalent ways of adding data to a dataset `data`: for example, each of

```
data.append(n=1.739,a=[0.494,2.734])           # method 1
```

```
data.append(n,1.739)                             # method 2
data.append(a,[0.494,2.734])
```

```
dd = dict(n=1.739,a=[0.494,2.734])             # method 3
data.append(dd)
```

adds one new random number (or array) to `data['n']` (or `data['a']`).

extend (*args, **kwargs)

Add batched data to dataset.

There are three equivalent ways of adding batched data, containing multiple samples for each quantity, to a dataset `data`: for example, each of

```
data.extend(n=[1.739,2.682],
            a=[[0.494,2.734],[ 0.172, 1.400]]) # method 1
```

```
data.extend(n,[1.739,2.682])                   # method 2
data.extend(a,[[0.494,2.734],[ 0.172, 1.400]])
```

```
dd = dict(n=[1.739,2.682],
            a=[[0.494,2.734],[ 0.172, 1.400]]) # method 3
data.extend(dd)
```

adds two new random numbers (or arrays) to `data['n']` (or `data['a']`).

This method can be used to merge two datasets, whether or not they share keys: for example,

```
data = Dataset("file1")
data_extra = Dataset("file2")
data.extend(data_extra) # data now contains all of data_extra
```

grep (rexp)

Create new dataset containing items whose keys match `rexp`.

Returns a new `gvar.dataset.Dataset` containing only the items `self[k]` whose keys `k` match regular expression `rexp` (a string) according to Python module `re`:

```
>>> a = Dataset()
>>> a.append(xx=1.,xy=[10.,100.])
>>> a.append(xx=2.,xy=[20.,200.])
>>> print(a.grep('y'))
{'yy': [array([ 10., 100.]), array([ 20., 200.])]}
>>> print(a.grep('x'))
{'xx': [1.0, 2.0], 'xy': [array([ 10., 100.]), array([ 20., 200.])]}
>>> print(a.grep('x|y'))
{'xx': [1.0, 2.0], 'xy': [array([ 10., 100.]), array([ 20., 200.])]}
>>> print a.grep('[^y][^x]')
{'xy': [array([ 10., 100.]), array([ 20., 200.])]}
```

Items are retained even if `rexp` matches only part of the item's key.

slice (*sl*)

Create new dataset with `self[k] -> self[k][sl]`.

Parameter `sl` is a slice object that is applied to every item in the dataset to produce a new `gvar.Dataset`. Setting `sl = slice(0, None, 2)`, for example, discards every other sample for each quantity in the dataset. Setting `sl = slice(100, None)` discards the first 100 samples for each quantity.

arrayzip (*template*)

Merge lists of random data according to `template`.

`template` is an array of keys in the dataset, where the shapes of `self[k]` are the same for all keys `k` in `template`. `self.arrayzip(template)` merges the lists of random numbers/arrays associated with these keys to create a new list of (merged) random arrays whose layout is specified by `template`: for example,

```
>>> d = Dataset()
>>> d.append(a=1,b=10)
>>> d.append(a=2,b=20)
>>> d.append(a=3,b=30)
>>> print(d)           # three random samples each for a and b
{'a': [1.0, 2.0, 3.0], 'b': [10.0, 20.0, 30.0]}
>>> # merge into list of 2-vectors:
>>> print(d.arrayzip(['a','b']))
[[ 1.  10.]
 [ 2.  20.]
 [ 3.  30.]]
>>> # merge into list of (symmetric) 2x2 matrices:
>>> print(d.arrayzip(['b','a'], ['a','b']))
[[[ 10.   1.]
   [ 1.  10.]]

 [[ 20.   2.]
   [ 2.  20.]]

 [[ 30.   3.]
   [ 3.  30.]]]
```

The number of samples in each merged result is the same as the number samples for each key (here 3). The keys used in this example represent scalar quantities; in general, they could be either scalars or arrays (of any shape, so long as all have the same shape).

trim ()

Create new dataset where all entries have same sample size.

toarray ()

Copy `self` but with `self[k]` as numpy arrays.

LSQFIT - NONLINEAR LEAST SQUARES FITTING

4.1 Introduction

This package contains tools for nonlinear least-squares curve fitting of data. In general a fit has four inputs:

1. The dependent data y that is to be fit — typically y is a Python dictionary in an `lsqfit` analysis. Its values $y[k]$ are either `gvar.GVars` or arrays (any shape or dimension) of `gvar.GVars` that specify the values of the dependent variables and their errors.
2. A collection x of independent data — x can have any structure and contain any data (or no data).
3. A fit function $f(x, p)$ whose parameters p are adjusted by the fit until $f(x, p)$ equals y to within y 's errors — parameters p are usually specified by a dictionary whose values $p[k]$ are individual parameters or (numpy) arrays of parameters. The fit function is assumed independent of x (that is, $f(p)$) if $x = \text{False}$ (or if x is omitted from the input data).
4. Initial estimates or *priors* for each parameter in p — priors are usually specified using a dictionary `prior` whose values `prior[k]` are `gvar.GVars` or arrays of `gvar.GVars` that give initial estimates (values and errors) for parameters $p[k]$.

A typical code sequence has the structure:

```
... collect x, y, prior ...

def f(x, p):
    ... compute fit to y[k], for all k in y, using x, p ...
    ... return dictionary containing the fit values for the y[k]s ...

fit = lsqfit.nonlinear_fit(data=(x, y), prior=prior, fcn=f)
print(fit)          # variable fit is of type nonlinear_fit
```

The parameters $p[k]$ are varied until the χ^2 for the fit is minimized.

The best-fit values for the parameters are recovered after fitting using, for example, `p=fit.p`. Then the $p[k]$ are `gvar.GVars` or arrays of `gvar.GVars` that give best-fit estimates and fit uncertainties in those estimates. The `print(fit)` statement prints a summary of the fit results.

The dependent variable y above could be an array instead of a dictionary, which is less flexible in general but possibly more convenient in simpler fits. Then the approximate y returned by fit function $f(x, p)$ must be an array with the same shape as the dependent variable. The prior `prior` could also be represented by an array instead of a dictionary.

The `lsqfit` tutorial contains extended explanations and examples.

4.2 Formal Background

The formal structure of a least-squares problem involves fitting input data y_i with functions $f_i(p)$ by adjusting fit parameters p_a to minimize

$$\begin{aligned}\chi^2 &\equiv \sum_{ij} \Delta y(p)_i (\text{cov}_y^{-1})_{ij} \Delta y(p)_j \\ &\equiv \Delta y(p) \cdot \text{cov}_y^{-1} \cdot \Delta y(p)\end{aligned}$$

where cov_y is the covariance matrix for the input data and

$$\Delta y(p)_i \equiv f_i(p) - y_i.$$

There are generally two types of input data — actual data and prior information for each fit parameter — but we lump these together here since they enter in the same way (that is, the sums over i and j are over all data and priors).

The best-fit values \bar{p}_a for the fit parameters are those that minimize χ^2 :

$$(\partial_a \Delta y(\bar{p}))^T \cdot \text{cov}_y^{-1} \cdot \Delta y(\bar{p}) = 0$$

where the derivatives are $\partial_a = \partial/\partial \bar{p}_a$. The covariance matrix cov_p for these is obtained (approximately) from

$$(\text{cov}_p^{-1})_{ab} \equiv (\partial_a \Delta y(\bar{p}))^T \cdot \text{cov}_y^{-1} \cdot (\partial_b \Delta y(\bar{p})).$$

Consequently the variance for any function $g(\bar{p})$ of the best-fit parameters is given by (approximately)

$$\sigma_g^2 = \partial g(\bar{p}) \cdot \text{cov}_p \cdot \partial g(\bar{p})$$

The definition of the covariance matrix implies that it and any variance σ_g^2 derived from it depend linearly (approximately) on the elements of the input data covariance matrix cov_y , at least when errors are small:

$$\sigma_g^2 \approx \sum_{ij} c(\bar{p})_{ij} (\text{cov}_y)_{ij}$$

This allows us to associate different portions of the output error σ_g^2 with different parts of the input error cov_y , creating an “error budget” for $g(\bar{p})$. Such information helps pinpoint the input errors that most affect the output errors for any particular quantity $g(\bar{p})$, and also indicates how those output errors might change for a given change in input error.

The relationship between the input and output errors is only approximately linear because the coefficients in the expansion depend upon the best-fit values for the parameters, and these depend upon the input errors — but only weakly when errors are small. Neglecting such variation in the parameters, the error budget for any quantity is easily computed using

$$\frac{\partial (\text{cov}_p)_{ab}}{\partial (\text{cov}_y)_{ij}} = D_{ai} D_{bj}$$

where

$$D_{ai} \equiv (\text{cov}_p \cdot \partial \Delta y \cdot \text{cov}_y^{-1})_{ai}$$

and, trivially, $\text{cov}_p = D \cdot \text{cov}_y \cdot D^T$.

This last formula suggests that

$$\frac{\partial \bar{p}_a}{\partial y_i} = D_{ai}.$$

This relationship is true in the limit of small errors, as is easily derived from the minimum condition for the fit, which defines (implicitly) $\bar{p}_a(y)$: Differentiating with respect to y_i we obtain

$$(\partial_a \Delta y(\bar{p}))^T \cdot \text{cov}_y^{-1} \cdot \frac{\partial \Delta y(\bar{p})}{\partial y_i} = 0$$

where we have ignored terms suppressed by a factor of $\Delta y(p)$. This leads immediately to the relationship above.

The data's covariance matrix cov_y is sometimes rather singular, making it difficult to invert. This problem is dealt with using an *svd* cut: the covariance matrix is diagonalized, some number of the smallest (and therefore least-well determined) eigenvalues and their eigenvectors are discarded, and the inverse matrix is reconstituted from the eigenmodes that remain. (Instead of discarding modes one can replace their eigenvalues by the smallest eigenvalue that is retained; this is less conservative and sometimes leads to more accurate results.) Note that the covariance matrix has at most N non-zero eigenvalues when it is estimated from N random samples; zero-modes should always be discarded.

4.3 nonlinear_fit Objects

```
class lsqfit.nonlinear_fit (data=None, fcn=None, prior=None, p0=None, svdcut=(1e-15, 1e-15),
                           svdnum=None, debug=False, **kargs)
```

Nonlinear least-squares fit.

`lsqfit.nonlinear_fit` fits a (nonlinear) function $f(x, p)$ to data y by varying parameters p , and stores the results: for example,

```
fit = nonlinear_fit(data=(x, y), fcn=f, prior=prior)    # do fit
print (fit)                                           # print fit results
```

The best-fit values for the parameters are in `fit.p`, while the `chi**2`, the number of degrees of freedom, the logarithm of Gaussian Bayes Factor, the number of iterations, and the cpu time needed for the fit are in `fit.chi2`, `fit.dof`, `fit.logGBF`, `fit.nit`, and `fit.time`, respectively. Results for individual parameters in `fit.p` are of type `gvar.GVar`, and therefore carry information about errors and correlations with other parameters. The fit data and prior can be recovered using `fit.x` (equals `False` if there is no `x`), `fit.y`, and `fit.prior`; the data and prior are corrected for the *svd* cut, if there is one (that is, their covariance matrices have been modified in accordance with the *svd* cut).

Parameters

- **data** – Data to be fit by `lsqfit.nonlinear_fit`. It can have any of the following formats:

data = x, y `x` is the independent data that is passed to the fit function with the fit parameters: `fcn(x, p)`. `y` is a dictionary (or array) of `gvar.GVars` that encode the means and covariance matrix for the data that is to be fit being fit. The fit function must return a result having the same layout as `y`.

data = y `y` is a dictionary (or array) of `gvar.GVars` that encode the means and covariance matrix for the data being fit. There is no independent data so the fit function depends only upon the fit parameters: `fit(p)`. The fit function must return a result having the same layout as `y`.

data = x, ymean, ycov `x` is the independent data that is passed to the fit function with the fit parameters: `fcn(x, p)`. `ymean` is an array containing the mean values of the fit data. `ycov` is an array containing the covariance matrix of the fit data; `ycov.shape` equals `2*ymean.shape`. The fit function must return an array having the same shape as `ymean`.

data = x, ymean, ysdev *x* is the independent data that is passed to the fit function with the fit parameters: `fcn(x, p)`. *ymean* is an array containing the mean values of the fit data. *ysdev* is an array containing the standard deviations of the fit data; `ysdev.shape` equals `ymean.shape`. The data are assumed to be uncorrelated. The fit function must return an array having the same shape as *ymean*.

Setting `x=False` in the first, third or fourth of these formats implies that the fit function depends only on the fit parameters: that is, `fcn(p)` instead of `fcn(x, p)`. (This is not assumed if `x=None`.)

- **fcn (function)** – The function to be fit to *data*. It is either a function of the independent data *x* and the fit parameters *p* (`fcn(x, p)`), or a function of just the fit parameters (`fcn(p)`) when there is no *x* data or `x=False`. The parameters are tuned in the fit until the function returns values that agree with the *y* data to within the *ys*' errors. The function's return value must have the same layout as the *y* data (a dictionary or an array). The fit parameters *p* are either: 1) a dictionary where each `p[k]` is a single parameter or an array of parameters (any shape); or, 2) a single array of parameters. The layout of the parameters is the same as that of prior `prior` if it is specified; otherwise, it is inferred from of the starting value `p0` for the fit.
- **prior** (dictionary, array, or None) – A dictionary (or array) containing *a priori* estimates for all parameters *p* used by fit function `fcn(x, p)` (or `fcn(p)`). Fit parameters *p* are stored in a dictionary (or array) with the same keys and structure (or shape) as `prior`. The default value is None; `prior` must be defined if `p0` is None.
- **p0** (dictionary, array, string or None) – Starting values for fit parameters in fit. `lsqfit.nonlinear_fit` adjusts `p0` to make it consistent in shape and structure with `prior` when the latter is specified: elements missing from `p0` are filled in using `prior`, and elements in `p0` that are not in `prior` are discarded. If `p0` is a string, it is taken as a file name and `lsqfit.nonlinear_fit` attempts to read starting values from that file; best-fit parameter values are written out to the same file after the fit (for priming future fits). If `p0` is None or the attempt to read the file fails, starting values are extracted from `prior`. The default value is None; `p0` must be defined if `prior` is None.
- **svdcut** (None or float or 2-tuple) – If positive, eigenvalues of the (rescaled) *y* covariance matrix that are smaller than `svdcut` times the maximum eigenvalue are replaced by `svdcut` times the maximum eigenvalue. If negative, eigenmodes with eigenvalues smaller than `|svdcut|` times the largest eigenvalue are discarded. If zero or None, the covariance matrix is left unchanged. If `svdcut` is a 2-tuple, the first entry is `svdcut` for the *y* covariance matrix and the second entry is `svdcut` for the prior's covariance matrix.
- **svdnum** (None or int or 2-tuple) – If positive, at most `svdnum` eigenmodes of the (rescaled) *y* covariance matrix are retained; the modes with the smallest eigenvalues are discarded. `svdnum` is ignored if set to None. If `svdnum` is a 2-tuple, the first entry is `svdnum` for the *y* covariance matrix and the second entry is `svdnum` for the prior's covariance matrix.
- **debug (boolean)** – Set to True for extra debugging of the fit function and a check for roundoff errors. (Default is False.)
- **fitterargs** – Dictionary of arguments passed on to `lsqfit.multifit`, which does the fitting.

The results from the fit are accessed through the following attributes (of `fit` where `fit = nonlinear_fit(...)`):

chi2

The minimum `chi**2` for the fit. `fit.chi2 / fit.dof` is usually of order one in good fits; values

much less than one suggest that the actual standard deviations in the input data and/or priors are smaller than the standard deviations used in the fit.

cov

Covariance matrix of the best-fit parameters from the fit.

dof

Number of degrees of freedom in the fit, which equals the number of pieces of data being fit when priors are specified for the fit parameters. Without priors, it is the number of pieces of data minus the number of fit parameters.

logGBF

The logarithm of the probability (density) of obtaining the fit data by randomly sampling the parameter model (priors plus fit function) used in the fit. This quantity is useful for comparing fits of the same data to different models, with different priors and/or fit functions. The model with the largest value of `fit.logGBF` is the one preferred by the data. The exponential of the difference in `fit.logGBF` between two models is the ratio of probabilities (Bayes factor) for those models. Differences in `fit.logGBF` smaller than 1 are not very significant. Gaussian statistics are assumed when computing `fit.logGBF`.

p

Best-fit parameters from fit. Depending upon what was used for the prior (or `p0`), it is either: a dictionary (`gvar.BufferDict`) of `gvar.GVars` and/or arrays of `gvar.GVars`; or an array (`numpy.ndarray`) of `gvar.GVars`. `fit.p` represents a multi-dimensional Gaussian distribution which, in Bayesian terminology, is the *posterior* probability distribution of the fit parameters.

pmean

Means of the best-fit parameters from fit (dictionary or array).

psdev

Standard deviations of the best-fit parameters from fit (dictionary or array).

palt

Same as `fit.p` except that the errors are computed directly from `fit.cov`. This is faster but means that no information about correlations with the input data is retained (unlike in `fit.p`); and, therefore, `fit.palt` cannot be used to generate error budgets. `fit.p` and `fit.palt` give the same means and normally give the same errors for each parameter. They differ only when the input data's covariance matrix is too singular to invert accurately (because of roundoff error), in which case an *svd* cut is advisable.

transformed_p

Same as `fit.p` but augmented to include the transforms of any log-normal or other parameter implemented using decorator `lsqfit.transform_p`. In the case of a log-normal variable `fit.p['logXX']`, for example, `fit.transformed_p['XX']` is defined equal to `exp(fit.p['logXX'])`.

p0

The parameter values used to start the fit.

Q

The probability that the χ^2 from the fit could have been larger, by chance, assuming the best-fit model is correct. Good fits have *Q* values larger than 0.1 or so. Also called the *p-value* of the fit.

svdcorrection

A dictionary containing the (flattened) *svd* corrections, if any, to the fit data `y` and the prior `prior`. For example, `fit.y` is obtained by adding `fit.svdcorrection['y']` to the (flattened) input `y` data. Similarly `fit.prior` is the input prior plus `fit.svdcorrection['prior']`. When there is no *svd* correction, the entries are set equal to `None`. When the input data and prior are correlated, `fit.svdcorrection['all']` contains the correction for the concatenated data and prior. There are no entries for keys `'y'` and `'prior'` in this case. There is also no entry for key `'prior'` when there is no prior.

time

CPU time (in secs) taken by fit.

The input parameters to the fit can be accessed as attributes. Note in particular attributes:

prior

Prior used in the fit. This may differ from the input prior if an *svd* cut is used (*svdcut*>0). It is either a dictionary (*gvar.BufferDict*) or an array (*numpy.ndarray*), depending upon the input. Equals *None* if no prior was specified.

x

The first field in the input *data*. This is sometimes the independent variable (as in ‘y vs x’ plot), but may be anything. It is set equal to *False* if the *x* field is omitted from the input *data*. (This also means that the fit function has no *x* argument: so *f(p)* rather than *f(x,p)*.)

y

Fit data used in the fit. This may differ from the input data if an *svd* cut is used (*svdcut*>0). It is either a dictionary (*gvar.BufferDict*) or an array (*numpy.ndarray*), depending upon the input.

Additional methods are provided for printing out detailed information about the fit, testing fits with simulated data, doing bootstrap analyses of the fit errors, dumping (for later use) and loading parameter values, and checking for roundoff errors in the final error estimates:

format (*maxline=0, pstyle='v'*)

Formats fit output details into a string for printing.

Parameters

- **maxline** (*integer*) – Maximum number of data points for which fit results and input data are tabulated. *maxline*<0 implies that only *chi2*, *Q*, *logGBF*, and *itns* are tabulated; no parameter values are included. Default is *maxline*=0.
- **pstyle** (*‘vv’, ‘v’, or ‘m’*) – Style used for parameter list. Supported values are ‘vv’ for very verbose, ‘v’ for verbose, and ‘m’ for minimal. When ‘m’ is set, only parameters whose values differ from their prior values are listed.

Returns String containing detailed information about fit.

fmt_errorbudget (*outputs, inputs, ndecimal=2, percent=True*)

Tabulate error budget for *outputs[k_o]* due to *inputs[k_i]*.

For each output *outputs[k_o]*, *fmt_errorbudget* computes the contributions to *outputs[k_o]*’s standard deviation coming from the *gvar.GVars* collected in *inputs[k_i]*. This is done for each key combination (*k_o*, *k_i*) and the results are tabulated with columns and rows labeled by *k_o* and *k_i*, respectively. If a *gvar.GVar* in *inputs[k_i]* is correlated with other *gvar.GVars*, the contribution from the others is included in the *k_i* contribution as well (since contributions from correlated *gvar.GVars* cannot be resolved). The table is returned as a string.

Parameters

- **outputs** – Dictionary of *gvar.GVars* for which an error budget is computed.
- **inputs** – Dictionary of: *gvar.GVars*, arrays/dictionaries of *gvar.GVars*, or lists of *gvar.GVars* and/or arrays/dictionaries of *gvar.GVars*. *fmt_errorbudget* tabulates the parts of the standard deviations of each *outputs[k_o]* due to each *inputs[k_i]*.
- **ndecimal** (*int*) – Number of decimal places displayed in table.
- **percent** (*boolean*) – Tabulate % errors if *percent* is *True*; otherwise tabulate the errors themselves.
- **colwidth** (*positive integer*) – Width of each column.

Returns A table (`str`) containing the error budget. Output variables are labeled by the keys in `outputs` (columns); sources of uncertainty are labeled by the keys in `inputs` (rows).

fmt_values (`outputs`, `ndecimal=None`)
 Tabulate `gvar.GVars` in `outputs`.

Parameters

- **outputs** – A dictionary of `gvar.GVar` objects.
- **ndecimal** (`int` or `None`) – Format values `v` using `v.fmt(ndecimal)`.

Returns A table (`str`) containing values and standard deviations for variables in `outputs`, labeled by the keys in `outputs`.

simulated_fit_iter (`n=None`, `pexact=None`, `**kargs`)

Iterator that returns simulation copies of a fit.

Fit reliability can be tested using simulated data which replaces the mean values in `self.y` with random numbers drawn from a distribution whose mean equals `self.fcn(pexact)` and whose covariance matrix is the same as `self.y`'s. Simulated data is very similar to the original fit data, `self.y`, but corresponds to a world where the correct values for the parameters (*i.e.*, averaged over many simulated data sets) are given by `pexact`. `pexact` is usually taken equal to `fit.pmean`.

Each iteration of the iterator creates new simulated data, with different random numbers and fits it, returning the the `lsqfit.nonlinear_fit` that results. The simulated data has the same covariance matrix as `fit.y`. Typical usage is:

```
...
fit = nonlinear_fit(...)
...
for sfit in fit.simulated_fit_iter(n=3):
    ... verify that sfit.p agrees with fit.pmean within errors ...
```

Only a few iterations are needed to get a sense of the fit's reliability since we know the correct answer in each case (`pexact`, equal to `fit.pmean` in this example). The simulated fit's output results should agree with `pexact` within the simulated fit's errors.

Simulated fits can also be used to correct for biases in the fit's output parameters or functions of them, should non-Gaussian behavior arise. This is possible, again, because we know the correct value for every parameter before we do the fit. Simulated fits provide a faster alternative to a traditional bootstrap analysis.

Parameters

- **n** (`integer` or `None`) – Maximum number of iterations (equals infinity if `None`).
- **pexact** (`None` or array or dictionary of numbers) – Fit-parameter values for the underlying distribution used to generate simulated data; replaced by `self.pmean` if is `None` (default).

Returns An iterator that returns `lsqfit.nonlinear_fits` for different simulated data.

Note that additional keywords can be added to overwrite keyword arguments in `lsqfit.nonlinear_fit`.

bootstrap_iter (`n=None`, `datalist=None`)

Iterator that returns bootstrap copies of a fit.

A bootstrap analysis involves three steps: 1) make a large number of "bootstrap copies" of the original input data that differ from each other by random amounts characteristic of the underlying randomness in the original data; 2) repeat the entire fit analysis for each bootstrap copy of the data, extracting fit results from each; and 3) use the variation of the fit results from bootstrap copy to bootstrap copy to determine an approximate probability distribution (possibly non-gaussian) for the each result.

Bootstrap copies of the data for step 2 are provided in `datalist`. If `datalist` is `None`, they are generated instead from the means and covariance matrix of the fit data (assuming gaussian statistics). The maximum number of bootstrap copies considered is specified by `n` (`None` implies no limit).

Typical usage is:

```
...
fit = lsqfit.nonlinear_fit(...)
...
for bsfit in fit.bootstrapped_fit_iter(n=100, datalist=datalist):
    ... analyze fit parameters in bsfit.pmean ...
```

Parameters

- **n** (*integer*) – Maximum number of iterations if `n` is not `None`; otherwise there is no maximum.
- **datalist** (sequence or iterator or `None`) – Collection of bootstrap data sets for fitter.

Returns Iterator that returns an `lsqfit.nonlinear_fit` object containing results from the fit to the next data set in `datalist`

`dump_p` (*filename*)

Dump parameter values (`fit.p`) into file `filename`.

`fit.dump_p(filename)` saves the best-fit parameter values (`fit.p`) from a `nonlinear_fit` called `fit`. These values are recovered using `p = nonlinear_fit.load_parameters(filename)` where `p`'s layout is the same as that of `fit.p`.

`dump_pmean` (*filename*)

Dump parameter means (`fit.pmean`) into file `filename`.

`fit.dump_pmean(filename)` saves the means of the best-fit parameter values (`fit.pmean`) from a `nonlinear_fit` called `fit`. These values are recovered using `p0 = nonlinear_fit.load_parameters(filename)` where `p0`'s layout is the same as `fit.pmean`. The saved values can be used to initialize a later fit (`nonlinear_fit` parameter `p0`).

`static load_parameters` (*filename*)

Load parameters stored in file `filename`.

`p = nonlinear_fit.load_p(filename)` is used to recover the values of fit parameters dumped using `fit.dump_p(filename)` (or `fit.dump_pmean(filename)`) where `fit` is of type `lsqfit.nonlinear_fit`. The layout of the returned parameters `p` is the same as that of `fit.p` (or `fit.pmean`).

`check_roundoff` (*rtol=0.25, atol=1e-6*)

Check for roundoff errors in `fit.p`.

Compares standard deviations from `fit.p` and `fit.palt` to see if they agree to within relative tolerance `rtol` and absolute tolerance `atol`. Generates a warning if they do not (in which case an *svd* cut might be advisable).

4.4 Functions

`lsqfit.empbayes_fit` (*z0, fitargs, **minargs*)

Call `lsqfit.nonlinear_fit(**fitargs(z))` varying `z`, starting at `z0`, to maximize `logGBF` (empirical Bayes procedure).

The fit is redone for each value of z that is tried, in order to determine $\log\text{GBF}$.

Parameters

- **z0** (*array*) – Starting point for search.
- **fitargs** (*function*) – Function of array z that determines which fit parameters to use. The function returns these as an argument dictionary for `lsqfit.nonlinear_fit()`.
- **minargs** (*dictionary*) – Optional argument dictionary, passed on to `lsqfit.multiminex`, which finds the minimum.

Returns A tuple containing the best fit (object of type `lsqfit.nonlinear_fit`) and the optimal value for parameter z .

`lsqfit.wavg(xa, svdcut=None, svdnum=None, rescale=True, covfac=None)`

Weighted average of `gvar.GVars` or arrays/dicts of `gvar.GVars`.

The weighted average of several `gvar.GVars` is what one obtains from a least-squares fit of the collection of `gvar.GVars` to the one- parameter fit function `def f(p): return N * [p[0]]` where N is the number of `gvar.GVars`. The average is the best-fit value for $p[0]$. `gvar.GVars` with smaller standard deviations carry more weight than those with larger standard deviations. The averages computed by `wavg` take account of correlations between the `gvar.GVars`.

Typical usage is:

```
x1 = gvar.gvar(...)
x2 = gvar.gvar(...)
x3 = gvar.gvar(...)
xavg = wavg([x1, x2, x3])    # weighted average of x1, x2 and x3
```

The individual `gvar.GVars` in this example can be replaced by multidimensional distributions, represented by arrays of `gvar.GVars` or dictionaries of `gvar.GVars` or arrays of `gvar.GVars`. For example,

```
x1 = [gvar.gvar(...), gvar.gvar(...)]
x2 = [gvar.gvar(...), gvar.gvar(...)]
x3 = [gvar.gvar(...), gvar.gvar(...)]
xavg = wavg([x1, x2, x3])
        # xavg[i] is wgted avg of x1[i], x2[i], x3[i]
```

and

```
x1 = dict(a=[gvar.gvar(...), gvar.gvar(...)], b=gvar.gvar(...))
x2 = dict(a=[gvar.gvar(...), gvar.gvar(...)], b=gvar.gvar(...))
x3 = dict(a=[gvar.gvar(...), gvar.gvar(...)], b=gvar.gvar(...))
xavg = wavg([x1, x2, x3])
        # xavg['a'][i] is wgted avg of x1['a'][i], x2['a'][i], x3['a'][i]
        # xavg['b'] is gtd avg of x1['b'], x2['b'], x3['b']
```

Parameters

- **xa** – The `gvar.GVars` to be averaged. xa is a one-dimensional sequence of `gvar.GVars`, or of arrays of `gvar.GVars`, or of dictionaries containing `gvar.GVars` or arrays of `gvar.GVars`. All $xa[i]$ must have the same shape.
- **svdcut** (*None* or *float*) – If positive, eigenvalues of the xa covariance matrix that are smaller than `svdcut` times the maximum eigenvalue are replaced by `svdcut` times the maximum eigenvalue. If negative, eigenmodes with eigenvalues smaller than $|\text{svdcut}|$ times the largest eigenvalue are discarded. If zero or *None*, the covariance matrix is left unchanged.

- **svdnum** (None or int) – If positive, at most `svdnum` eigenmodes of the `xa` covariance matrix are retained; the modes with the smallest eigenvalues are discarded. `svdnum` is ignored if set to None.
- **rescale** (bool) – If True, rescale covariance matrix so diagonal elements all equal 1 before applying *svd* cuts. (Default is True.)
- **covfac** (None or number) – The covariance matrix (or matrices) of `xa` is multiplied by `covfac` if `covfac` is not None.

Returns Weighted average of the `xa` elements. The result has the same type and shape as each element of `xa` (that is, either a `gvar.GVar` or an array of `gvar.GVars`.)

The following function attributes are also set:

`wavg.chi2`

`chi**2` for weighted average.

`wavg.dof`

Effective number of degrees of freedom.

`wavg.Q`

The probability that the `chi**2` could have been larger, by chance, assuming that the data are all Gaussian and consistent with each other. Values smaller than 0.1 suggest that the data are not Gaussian or are inconsistent with each other. Also called the *p-value*.

Quality factor *Q* (or *p-value*) for fit.

`wavg.time`

Time required to do average.

`lsqfit.gammaQ()`

Return the incomplete gamma function $Q(a, x) = 1 - P(a, x)$.

Note that `gammaQ(ndof/2., chi2/2.)` is the probability that one could get a `chi**2` larger than `chi2` with `ndof` degrees of freedom even if the model used to construct `chi2` is correct.

4.5 Utility Classes

class `lsqfit.transform_p` (*priorkeys*, *pindex=None*, *pkey=None*)

Decorate fit function to allow log/sqrt-normal priors.

This decorator can be applied to fit functions whose parameters are stored in a dictionary-like object. It searches the parameter keys for string-valued keys of the form "`log (XX)`", "`logXX`", "`sqrt (XX)`", or "`sqrtXX`" where "`XX`" is an arbitrary string. For each such key it adds a new entry to the parameter dictionary with key "`XX`" where:

```
p["XX"] = exp(p[k])      for k = "log (XX)" or "logXX"
```

or

```
p["XX"] = p[k] ** 2      for k = "sqrt (XX)" or "sqrtXX"
```

This means that the fit function can be expressed entirely in terms of `p["XX"]` even if the actual fit parameter is the logarithm or square root of that quantity. Since fit parameters have gaussian/normal priors, `p["XX"]` has a log-normal or "sqrt-normal" distribution in the first or second cases above, respectively. In either case `p["XX"]` is guaranteed to be positive.

This is a convenience function. It allows for the rapid replacement of a fit parameter by its logarithm or square root without having to rewrite the fit function — only the prior need be changed. The decorator needs the keys from the prior, and it needs to be told which argument (numbered from 0) of the fit function is the parameter dictionary, unless the fit function has only a single argument:

```
@lsqfit.transform_p(prior.keys(), 1)
def fitfcn(x, p):
    ...

or

@lsqfit.transform_p(prior.keys(), 0)
def fitfcn(p, other_arg1, ...):
    ...

or

@lsqfit.transform_p(prior.keys())
def fitfcn(p):
    ...
```

A list of the specific keys that need transforming can be used instead of the list of all keys (`prior.keys()`). The decorator assigns a copy of itself to the function as an attribute: `fitfcn.transform_p`.

Parameters

- **priorkeys** (*sequence*) – The keys in the prior that are to be tranformed. Other keys can be in `priorkeys` provided they do not begin with 'log' or 'sqrt' — they are ignored.
- **pindex** (*integer or None*) – Index of the parameters-dictionary in the argument list of the fit function. Default value is `None`; one of `pkey` or `pindex` must be specified (i.e., not `None`), unless the fit function has only a single argument.
- **pkey** (*string or None*) – Name of the parameters-variable in the argument keyword dictionary of the fit function. Default value is `None`; one of `pkey` or `pindex` must be specified (i.e., not `None`), unless the fit function has only a single argument.

transform(p)

Create transformed copy of dictionary p.

Create a copy of parameter-dictionary p that includes new entries for each "logXX", etc entry corresponding to "XX". The values in p can be any type that supports logarithms, exponentials, and arithmetic.

untransform(p)

Undo `self.transform(p)`.

Reconstruct p0 where `p == self.transform(p0)`; that is remove entries for keys "XX" that were added by `transform_p.transform()` (because "logXX" or "sqrtXX" or ... appeared in p0).

static paramkey(k)

Return parameter key corresponding to prior-key k.

Strip off any "log" or "sqrt" prefix.

static priorkey(prior, k)

Return key in prior corresponding to k.

Add in "log" or "sqrt" as needed to find a key in prior.

class `lsqfit.multifit` (*x0, n, f, reltol=1e-4, abstol=0, maxit=1000, alg='lmsder', analyzer=None*)

Fitter for nonlinear least-squares multidimensional fits.

Parameters

- **x0** (numpy array of floats) – Starting point for minimization.
- **n** (positive integer) – Length of vector returned by the fit function $f(x)$.
- **f** (function) – Fit function: `multifit` minimizes $\sum_i f_i(x)^2$ by varying parameters x . The parameters are a 1-d numpy array of either numbers or `gvar.GVars`.
- **reitol** (float) – The fit stops when $|dx_i| < abstol + reitol * |x_i|$; default value is $1e-4$.
- **abstol** (float) – The fit stops when $|dx_i| < abstol + reitol * |x_i|$; default value is 0.0 .
- **maxit** (integer) – Maximum number of iterations in search for minimum; default is 1000.
- **alg** (string) – *GSL* algorithm to use for minimization. Two options are currently available: "lmsder", the scaled *LMDER* algorithm (default); and "lmdcr", the unscaled *LMDER* algorithm.
- **analyzer** (function) – Optional function of $x, [...f_i(x) ...], [...df_{ij}(x) ...]$ which is called after each iteration. This can be used to inspect intermediate steps in the minimization, if needed.

`multifit` is a function-class whose constructor does a least squares fit by minimizing $\sum_i f_i(x)^2$ as a function of vector x . The following attributes are available:

x
Location of the most recently computed (best) fit point.

cov
Covariance matrix at the minimum point.

f
The fit function $f(x)$ at the minimum in the most recent fit.

J
Gradient $J_{ij} = df_i/dx[j]$ for most recent fit.

nit
Number of iterations used in last fit to find the minimum.

error
None if fit successful; an error message otherwise.

`multifit` is a wrapper for the `multifit` *GSL* routine.

class `lsqfit.multiminex` (*x0*, *f*, *tol*=*1e-4*, *maxit*=*1000*, *step*=*1*, *alg*='nmsimplex2', *analyzer*=None)
Minimizer for multidimensional functions.

Parameters

- **x0** (numpy array of floats) – Starting point for minimization search.
- **f** (function) – Function $f(x)$ to be minimized by varying vector x .
- **tol** (float) – Minimization stops when x has converged to with tolerance *tol*; default is $1e-4$.
- **maxit** (integer) – Maximum number of iterations in search for minimum; default is 1000.
- **step** (number) – Initial step size to use in varying components of x ; default is 1.
- **alg** (string) – *GSL* algorithm to use for minimization. Three options are currently available: "nmsimplex", Nelder Mead Simplex algorithm; "nmsimplex2", an improved version

of "nmsimplex" (default); and "nmsimplex2rand", a version of "nmsimplex2" with random shifts in the start position.

- **analyzer** (*function*) – Optional function of x , $f(x)$, it , where it is the iteration number, which is called after each iteration. This can be used to inspect intermediate steps in the minimization, if needed.

`multiminex` is a function-class whose constructor minimizes a multidimensional function $f(x)$ by varying vector x . This routine does *not* use user-supplied information about the gradient of $f(x)$. The following attributes are available:

- x**
Location of the most recently computed minimum (1-d array).
- f**
Value of function $f(x)$ at the most recently computed minimum.
- nit**
Number of iterations required to find most recent minimum.
- error**
None if fit successful; an error message otherwise.

`multiminex` is a wrapper for the `multimin` *GSL* routine.

4.6 Requirements

`lsqfit` relies heavily on the `gvar`, and `numpy` modules. Several utility functions are in `lsqfit_util`. Also the minimization routines are from the Gnu Scientific Library (*GSL*).

INDICES AND TABLES

- *genindex*
- *modindex*
- *search*

g

`gvar`, [41](#)

`gvar.dataset`, [59](#)

l

`lsqfit`, [67](#)

Symbols

`__str__()` (gvar.GVar method), 54

A

`append()` (gvar.dataset.Dataset method), 65

`arrayzip()` (gvar.dataset.Dataset method), 66

`autocorr()` (in module gvar.dataset), 62

`avg_data()` (in module gvar.dataset), 61

B

`bin_data()` (in module gvar.dataset), 62

`bootstrap_iter()` (in module gvar), 50

`bootstrap_iter()` (in module gvar.dataset), 62

`bootstrap_iter()` (lsqfit.nonlinear_fit method), 73

`buf` (gvar.BufferDict attribute), 56

`BufferDict` (class in gvar), 54

C

`check_roundoff()` (lsqfit.nonlinear_fit method), 74

`chi2` (lsqfit.nonlinear_fit attribute), 70

`chi2` (lsqfit.wavg attribute), 76

`chi2()` (in module gvar), 48

`correction` (gvar.svd attribute), 53

`cov` (gvar.GVar attribute), 54

`cov` (lsqfit.multifit attribute), 78

`cov` (lsqfit.nonlinear_fit attribute), 71

D

`D` (gvar.SVD attribute), 57

`D` (gvar.svd attribute), 53

`Dataset` (class in gvar.dataset), 63

`decomp()` (gvar.SVD method), 58

`delta` (gvar.SVD attribute), 58

`der` (gvar.GVar attribute), 54

`deriv()` (gvar.GVar method), 54

`deriv()` (in module gvar), 49

`dof` (gvar.chi2 attribute), 49

`dof` (lsqfit.nonlinear_fit attribute), 71

`dof` (lsqfit.wavg attribute), 76

`dotder()` (gvar.GVar method), 54

`dtype` (gvar.BufferDict attribute), 56

`dump()` (gvar.BufferDict method), 56

`dump_p()` (lsqfit.nonlinear_fit method), 74

`dump_pmean()` (lsqfit.nonlinear_fit method), 74

`dumps()` (gvar.BufferDict method), 56

E

`eigen_range` (gvar.svd attribute), 52

`empbayes_fit()` (in module lsqfit), 74

`error` (lsqfit.multifit attribute), 78

`error` (lsqfit.multiminex attribute), 79

`evalcov()` (in module gvar), 48

`extend()` (gvar.dataset.Dataset method), 65

F

`f` (lsqfit.multifit attribute), 78

`f` (lsqfit.multiminex attribute), 79

`flat` (gvar.BufferDict attribute), 56

`flatten()` (gvar.BufferDict method), 56

`fmt()` (gvar.GVar method), 54

`fmt()` (in module gvar), 48

`fmt_errorbudget()` (in module gvar), 50

`fmt_errorbudget()` (lsqfit.nonlinear_fit method), 72

`fmt_values()` (in module gvar), 51

`fmt_values()` (lsqfit.nonlinear_fit method), 73

`format()` (lsqfit.nonlinear_fit method), 72

G

`gammaQ()` (in module lsqfit), 76

`grep()` (gvar.dataset.Dataset method), 65

`GVar` (class in gvar), 53

`gvar` (module), 41

`gvar()` (in module gvar), 47, 48

`gvar.dataset` (module), 59

`gvar_factory()` (in module gvar), 51

I

`inv_wgt` (gvar.svd attribute), 53

`isscalar()` (gvar.BufferDict method), 56

J

`J` (lsqfit.multifit attribute), 78

K

kappa (gvar.SVD attribute), 57

L

load() (gvar.BufferDict static method), 56
load_parameters() (lsqfit.nonlinear_fit static method), 74
loads() (gvar.BufferDict static method), 56
logdet (gvar.svd attribute), 53
logGBF (lsqfit.nonlinear_fit attribute), 71
lsqfit (module), 67

M

mean (gvar.GVar attribute), 53
mean() (in module gvar), 48
multifit (class in lsqfit), 77
multiminex (class in lsqfit), 78

N

nit (lsqfit.multifit attribute), 78
nit (lsqfit.multiminex attribute), 79
nmod (gvar.SVD attribute), 57
nonlinear_fit (class in lsqfit), 69

P

p (lsqfit.nonlinear_fit attribute), 71
p0 (lsqfit.nonlinear_fit attribute), 71
palt (lsqfit.nonlinear_fit attribute), 71
paramkey() (lsqfit.transform_p static method), 77
partialsdev() (gvar.GVar method), 53
partialvar() (gvar.GVar method), 53
pmean (lsqfit.nonlinear_fit attribute), 71
prior (lsqfit.nonlinear_fit attribute), 72
priorkey() (lsqfit.transform_p static method), 77
psdev (lsqfit.nonlinear_fit attribute), 71

Q

Q (gvar.chi2 attribute), 49
Q (lsqfit.nonlinear_fit attribute), 71
Q (lsqfit.wavg attribute), 76

R

raniter() (in module gvar), 49
ranseed() (in module gvar), 50
rebuild() (in module gvar), 51
restore_gvar() (in module gvar), 51

S

samplesize (gvar.dataset.Dataset attribute), 65
sdev (gvar.GVar attribute), 53
sdev() (in module gvar), 48
shape (gvar.BufferDict attribute), 56
simulated_fit_iter() (lsqfit.nonlinear_fit method), 73
size (gvar.BufferDict attribute), 55

slice() (gvar.BufferDict method), 56
slice() (gvar.dataset.Dataset method), 66
SVD (class in gvar), 56
svd() (in module gvar), 52
svdcorrection (lsqfit.nonlinear_fit attribute), 71
switch_gvar() (in module gvar), 51

T

time (lsqfit.nonlinear_fit attribute), 71
time (lsqfit.wavg attribute), 76
toarray() (gvar.dataset.Dataset method), 66
transform() (lsqfit.transform_p method), 77
transform_p (class in lsqfit), 76
transformed_p (lsqfit.nonlinear_fit attribute), 71
trim() (gvar.dataset.Dataset method), 66

U

uncorrelated() (in module gvar), 48
untransform() (lsqfit.transform_p method), 77
update() (gvar.BufferDict method), 56

V

val (gvar.SVD attribute), 57
val (gvar.svd attribute), 52
var (gvar.GVar attribute), 53
var() (in module gvar), 48
vec (gvar.SVD attribute), 57
vec (gvar.svd attribute), 52

W

wavg() (in module lsqfit), 75

X

x (lsqfit.multifit attribute), 78
x (lsqfit.multiminex attribute), 79
x (lsqfit.nonlinear_fit attribute), 72

Y

y (lsqfit.nonlinear_fit attribute), 72