

χ^2 and Linear Fits

Andrew Gould (Dept. of Astronomy, Ohio State University)

Abstract

The mathematics of linear fits is presented in covariant form. Topics include: correlated data, covariance matrices, joint fits to multiple data sets, constraints, and extension of the formalism to non-linear fits. A brief summary at the end provides a convenient crib sheet. These are somewhat amplified notes from a 90 minute lecture in a first-year graduate course. None of the results are new. They are presented here because they do not appear to be elsewhere available in compact form.

Expectations and Covariances

Let y be random variable, which is drawn from a random distribution $g(y)$. Then, we define the “expected value” or “mean” of y as

$$\langle y \rangle \equiv \int y g(y) dy / \int g(y) dy.$$

Using this definition, it is straightforward to prove the following identities,

$$\langle y_1 + y_2 \rangle = \langle y_1 \rangle + \langle y_2 \rangle, \quad \langle ky \rangle = k \langle y \rangle, \quad \langle k \rangle = k, \quad \langle \langle y \rangle \rangle = \langle y \rangle,$$

where k is a constant. We now motivate the idea of a “covariance” of two random variables by noting

$$\langle y_1 y_2 \rangle = \langle y_1 \rangle \langle y_2 \rangle + \text{cov}(y_1, y_2)$$

where

$$\text{cov}(y_1, y_2) \equiv \langle (y_1 - \langle y_1 \rangle)(y_2 - \langle y_2 \rangle) \rangle = \langle y_1 y_2 \rangle - \langle y_1 \rangle \langle y_2 \rangle,$$

the last step following from the identities above. If, when y_1 is above its mean, then y_2 also tends to be above its mean (and similarly for below), then $\text{cov}(y_1, y_2) > 0$, and then y_1 is said to be “correlated” with y_2 . If, when y_1 is above then y_2 is below, then $\text{cov}(y_1, y_2) < 0$, and the y_1 and y_2 are said to be “anti-correlated”. If $\text{cov}(y_1, y_2) = 0$, they are said to be “uncorrelated”. Only in this case is it true that $\langle y_1 y_2 \rangle = \langle y_1 \rangle \langle y_2 \rangle$. Three other identities that are easily proven,

$$\text{cov}(ky_1, y_2) = k \cdot \text{cov}(y_1, y_2), \quad \text{cov}(y_1, y_2 + y_3) = \text{cov}(y_1, y_2) + \text{cov}(y_1, y_3), \quad \text{cov}(y, k) = 0.$$

The covariance of a random variable with itself is called its variance. The error, σ , is defined to be the square root of the variance,

$$\text{var}(y) \equiv \text{cov}(y, y) = \langle y^2 \rangle - \langle y \rangle^2, \quad \sigma(y) \equiv \sqrt{\text{var}(y)}$$

Definition of χ^2

Suppose I have a set of N data points y_k with associated (uncorrelated for now) errors σ_k , and I have a model that makes predictions for the values of these data points $y_{k,\text{mod}}$. Then χ^2 is defined to be

$$\chi^2 \equiv \sum_{k=1}^N \frac{(y_k - y_{k,\text{mod}})^2}{\sigma_k^2}.$$

If the errors are Gaussian, then the likelihood is given by $\mathcal{L} = \exp(-\chi^2/2)$, so that minimizing χ^2 is equivalent to maximizing \mathcal{L} . However, even if the errors are not Gaussian, χ^2 minimization is a well-defined procedure, and none of the results given below (with one explicitly noted exception) depend in any way on the errors being Gaussian. This is important because sometimes little is known about the error distributions beyond their variances.

More generally, the errors might be correlated. Although in practice this is the exception, it makes the math much easier to consider the more general case of correlated errors. In this case, the covariance matrix, \mathcal{C}_{kl} , of the correlated errors (and its inverse \mathcal{B}_{kl}) are defined by

$$\mathcal{C}_{kl} \equiv \text{cov}(y_k, y_l), \quad \mathcal{B} \equiv \mathcal{C}^{-1},$$

Both \mathcal{C} and \mathcal{B} are symmetric. Then χ^2 is written as

$$\chi^2 = \sum_{k=1}^N \sum_{l=1}^N (y_k - y_{k,\text{mod}}) \mathcal{B}_{kl} (y_l - y_{l,\text{mod}})$$

Note that for the special case of uncorrelated errors, $\mathcal{C}_{kl} = \delta_{kl} \sigma_k^2$, where the Kronecker delta is defined by

$$\delta_{kl} = 1 \quad (k = l), \quad \delta_{kl} = 0 \quad (k \neq l).$$

In this case, $\mathcal{B}_{kl} = \delta_{kl} \sigma_k^{-2}$, so

$$\chi^2 = \sum_{k=1}^N \sum_{l=1}^N (y_k - y_{k,\text{mod}}) \frac{\delta_{kl}}{\sigma_k^2} (y_l - y_{l,\text{mod}}) = \sum_{k=1}^N \frac{(y_k - y_{k,\text{mod}})(y_k - y_{k,\text{mod}})}{\sigma_k^2},$$

which is the original definition.

Linear model

A linear model of n parameters is given by

$$y_{\text{mod}} \equiv \sum_{i=1}^n a_i f_i(x),$$

where the $f_i(x)$ are n arbitrary functions of the independent variable x . The independent variable is something that is known exactly (or very precisely), such as time. If the

independent variable is subject to significant uncertainty, the approach presented here must be substantially modified.

We can now write χ^2 ,

$$\chi^2 = \sum_{k=1}^N \sum_{l=1}^N \left[y_k - \sum_{i=1}^n a_i f_i(x_k) \right] \mathcal{B}_{kl} \left[y_l - \sum_{j=1}^n a_j f_j(x_l) \right].$$

The proliferation of summation signs is getting annoying. We can get rid of all of them using the ‘‘Einstein summation convention’’, whereby we just agree to sum over repeated indices. In the above cases, these are k, l, i, j . We then write,

$$\chi^2 = [y_k - a_i f_i(x_k)] \mathcal{B}_{kl} [y_l - a_j f_j(x_l)],$$

which is a lot simpler. Note that k and l are summed over the N data points, while i and j are summed over the n parameters.

Minimizing χ^2

The general problem of minimizing χ^2 with respect to the parameters a_i can be difficult, but for linear models it is straightforward. We first rewrite,

$$\chi^2 = y_k \mathcal{B}_{kl} y_l - 2a_i d_i + a_i b_{ij} a_j,$$

where the d_i and b_{ij} are defined by

$$d_i \equiv y_k \mathcal{B}_{kl} f_i(x_l), \quad b_{ij} \equiv f_i(x_k) \mathcal{B}_{kl} f_j(x_l).$$

We find the minimum by setting all the derivatives of χ^2 with respect to the parameters equal to zero,

$$0 = \frac{\partial \chi^2}{\partial a_m} = -2\delta_{im} d_i + \delta_{im} b_{ij} a_j + a_i b_{ij} \delta_{jm} = -2d_m + b_{mj} a_j + a_i b_{im} = 2(b_{mj} a_j - d_m),$$

where I have used $\partial a_i / \partial a_m = \delta_{im}$ and where I summed over one dummy index (i or j) in each step. This equation is easily solved:

$$d_m = b_{mj} a_j \Rightarrow a_i = c_{ij} d_j,$$

where c_{ij} is defined as the inverse of b_{ij}

$$c \equiv b^{-1}.$$

Note that the d_i are random variables because they are linear combinations of other random variables (the y_k), but that the b_{ij} (and so the c_{ij}) are combinations of constants, and so are not random variables.

Covariances of the Parameters

We would like to evaluate the covariances of the a_i , i.e., $\text{cov}(a_i, a_j)$, and derive their errors $\sigma(a_i) \equiv \sqrt{\text{cov}(a_i, a_i)}$. The first step to doing so is to evaluate the covariances of the d_i ,

$$\text{cov}(d_i, d_j) = \text{cov}[y_k \mathcal{B}_{kl} f_i(x_l), y_p \mathcal{B}_{pq} f_j(x_q)] = \mathcal{B}_{kl} f_i(x_l) \mathcal{B}_{pq} f_j(x_q) \text{cov}(y_k, y_p).$$

Then, using the definition of \mathcal{C}_{kp} and successively summing over all repeated indices,

$$\text{cov}(d_i, d_j) = \mathcal{B}_{kl} f_i(x_l) \mathcal{B}_{pq} f_j(x_q) \mathcal{C}_{kp} = \mathcal{B}_{kl} f_i(x_l) \delta_{kq} f_j(x_q) = f_i(x_l) \mathcal{B}_{kl} f_j(x_k) = b_{ij}.$$

Next, we evaluate the covariance of a_i with d_j ,

$$\text{cov}(a_i, d_j) = \text{cov}(c_{im} d_m, d_j) = c_{im} \text{cov}(d_m, d_j) = c_{im} b_{mj} = \delta_{ij}.$$

Finally,

$$\text{cov}(a_i, a_j) = \text{cov}(c_{im} d_m, a_j) = c_{im} \text{cov}(d_m, a_j) = c_{im} \delta_{mj} = c_{ij}.$$

That is, the c_{ij} , which were introduced only to solve for the a_i , now actually turn out to be the covariances of the a_i ! In particular,

$$\sigma(a_i) = \sqrt{c_{ii}}.$$

This is a very powerful result. It means that one can figure out the errors in an experiment, without having any data (just so long as one knows what data one is planning to get, and what the measurement errors will be).

A simple example

Let us consider the simple example of a two-parameter straight-line fit to some data. In this case,

$$y_{\text{mod}} = a_1 f_1(x) + a_2 f_2(x) = a_1 + a_2 x,$$

i.e., $f_1(x) = 1$ and $f_2(x) = x$. Let us also assume that the measurement errors are uncorrelated. Hence

$$d_1 = \sum_{k=1}^N \frac{y_k}{\sigma_k^2}, \quad d_2 = \sum_{k=1}^N \frac{y_k x_k}{\sigma_k^2},$$

$$b_{11} = \sum_{k=1}^N \frac{1}{\sigma_k^2}, \quad b_{12} = b_{21} = \sum_{k=1}^N \frac{x_k}{\sigma_k^2}, \quad b_{22} = \sum_{k=1}^N \frac{x_k^2}{\sigma_k^2}.$$

Let us now further specialize to the case for which all the σ_k are equal. Then

$$b = \frac{N}{\sigma^2} \begin{pmatrix} 1 & \langle x \rangle \\ \langle x \rangle & \langle x^2 \rangle \end{pmatrix}$$

where the “expectation” signs now mean simply averaging over the distribution of the independent variable at the times when the observations actually take place. This matrix is easily inverted,

$$c = \frac{\sigma^2}{N(\langle x^2 \rangle - \langle x \rangle^2)} \begin{pmatrix} \langle x^2 \rangle & -\langle x \rangle \\ -\langle x \rangle & 1 \end{pmatrix}.$$

In particular, the error in the slope is given by

$$[\sigma(a_2)]^2 = \text{var}(a_2) = c_{22} = \frac{\sigma^2}{N \text{var}(x)} \rightarrow \frac{12}{N} \frac{\sigma^2}{(\Delta x)^2},$$

where I have used $\text{var}(x)$ as a shorthand for $\langle x^2 \rangle - \langle x \rangle^2$, and where in the last step I have evaluated this quantity for the special case of data uniformly distributed over an interval Δx . That is, for data with equal errors, the variance of the slope is equal to the variance of the individual measurements divided by the variance of the independent-variable distribution, and then divided by N .

Nonlinear fits

Consider a more general model,

$$y_{\text{mod}} = F(x; a_1 \dots a_n),$$

in which the function F is not linear in the a_i , e.g., $F(x; a_1, a_2) = \cos(a_1 x) \exp(-a_2 x)$. The method given above cannot be used directly to solve for the a_i . Nevertheless, once the minimum a_i^0 is found, by whatever method, one can write

$$F(x; a_1 \dots a_n) = F_0 + (a_i - a_i^0) f_i(x) + \dots$$

where

$$f_i(x) \equiv \left. \frac{\partial F(x; a_1 \dots a_n)}{\partial a_i} \right|_{a_j = a_j^0}.$$

Then, one can calculate the b_{ij} , and so the c_{ij} , and thus obtain the errors. In fact, this formulation can also be used to find the minimum (using Newton’s method), but the description of this approach goes beyond the scope of these notes.

Expected Value of χ^2

We evaluate the expected value of χ^2 after it has been minimized

$$\langle \chi^2 \rangle = \langle y_k \mathcal{B}_{kl} y_l \rangle - 2 \langle a_i d_i \rangle + \langle a_i b_{ij} a_j \rangle = \langle y_k \mathcal{B}_{kl} y_l \rangle - \langle a_i d_i \rangle,$$

where I have used the minimizing condition $d_i = b_{ij} a_j$. To put this expression in an equivalent form whose physical meaning is more apparent, it is better to retreat to the original definition,

$$\langle \chi^2 \rangle = \langle [y_k - a_i f_i(x_k)] \mathcal{B}_{kl} [y_l - a_j f_j(x_l)] \rangle$$

or

$$\langle \chi^2 \rangle = \text{cov}([y_k - a_i f_i(x_k)], \mathcal{B}_{kl}[y_l - a_j f_j(x_l)]) + \langle [y_k - a_i f_i(x_k)] \rangle \mathcal{B}_{kl} \langle [y_l - a_j f_j(x_l)] \rangle.$$

The first term can be evaluated,

$$\begin{aligned} \text{cov}([y_k - a_i f_i(x_k)], \mathcal{B}_{kl}[y_l - a_j f_j(x_l)]) &= \mathcal{B}_{kl} \text{cov}(y_k, y_l) - 2\text{cov}(a_i, d_i) + b_{ij} \text{cov}(a_i, a_j) \\ &= \mathcal{B}_{kl} \mathcal{C}_{kl} - 2\delta_{ii} + b_{ij} c_{ij} = \delta_{kk} - \delta_{ii}, \end{aligned}$$

using $\text{cov}(a_i, d_j) = \delta_{ij}$. Since $\delta_{kk} = N$ and $\delta_{ii} = n$ (note that repeated indices still indicate summation), we have

$$\langle \chi^2 \rangle = N - n + \langle y_k - a_i f_i(x_k) \rangle \mathcal{B}_{kl} \langle y_l - a_j f_j(x_l) \rangle.$$

The last term is composed of the product of the expected difference between the model and the data at all pairs of points. For uncorrelated errors,

$$\langle y_k - a_i f_i(x_k) \rangle \mathcal{B}_{kl} \langle y_l - a_j f_j(x_l) \rangle \rightarrow \sum_{k=1}^N \frac{\langle y_k - a_i f_i(x_k) \rangle^2}{\sigma_k^2}.$$

If the model space spans the physical situation that is being modeled, then this expected difference is exactly zero. In this case, $\langle \chi^2 \rangle = N - n$, the number of data points less the number of parameters. If this relation fails, it can only be for one of three reasons: 1) normal statistical fluctuations, 2) misestimation of the errors, 3) problems in the model. For Gaussian statistics, one can show that

$$\text{var}(\chi^2) = 2(N - n).$$

So if there are, say 58 data points, and 8 parameters, then χ^2 should be 50 ± 10 . So if it is 57 or 41, that is ok. If it is 25 or 108, that is not ok. If it is 25, the only effect that could have produced this would be that the errors were overestimated. If it is 108, there are two possible causes. The errors could have been underestimated or the model could be failing to represent the physical situation. For example, if the physical situation causes the data to trace a parabola (which requires 3 parameters), but the model has only two parameters (say a 2-parameter straight line fit), then the model cannot adequately represent the data and the third term will be non-zero and so cause χ^2 to go up.

Combining Covariance Matrices

Suppose one has two sets of measurements, which had been analyzed as separate χ^2 's, χ_1^2 and χ_2^2 . Minimizing each with respect to the a_i yields best-fits a_i^1 and a_i^2 , and associated vectors and matrices $d_i^1, b_{ij}^1, c_{ij}^1, d_i^2, b_{ij}^2$, and c_{ij}^2 . Now imagine minimizing the sum of these two χ^2 's with respect to the a_m ,

$$0 = \frac{1}{2} \frac{\partial(\chi_1^2 + \chi_2^2)}{\partial a_m} = -(d_m^1 + d_m^2) + (b_{mj}^1 + b_{mj}^2) a_j,$$

which yields a combined solution,

$$a_i = c_{ij}d_j, \quad d_i \equiv d_i^1 + d_i^2, \quad b_{ij} \equiv b_{ij}^1 + b_{ij}^2, \quad c \equiv b^{-1}.$$

But it is not actually necessary to go back to the original χ^2 's. If you were given the results of the previous analyses, i.e. the best fit a_i^1 and c_{ij}^1 and a_i^2 and c_{ij}^2 for the two separate fits, you could calculate $b^1 = (c^1)^{-1}$ and $d_i^1 = b_{ij}^1 a_j^1$ and similarly for "2", and then directly calculate the combined best fit.

Linear Constraints

Suppose that you have found a best fit to your data, but now you have obtained some additional information that fixes a linear relation among your parameters. Such a constraint can be written

$$\kappa_i a_i = z.$$

For example, suppose that your information was that a_1 was actually equal to 3. Then $\kappa = (1, 0, 0, \dots)$ and $z = 3$. Or suppose that the information was that a_1 and a_2 were equal. Then $\kappa = (1, -1, 0, 0, \dots)$ and $z = 0$. How is the best fit solution and covariance matrix affected by this constraint? The answer is

$$\tilde{a}_i = a_i^0 - D\alpha_i \quad D \equiv \frac{[\kappa_j a_j^0 - z]}{\alpha_p \kappa_p}, \quad \alpha_i \equiv c_{ij} \kappa_j.$$

I derive this in a more general context below. For now just note that by

$$\text{cov}(D, D) = \frac{\kappa_i \kappa_j}{(\alpha_p \kappa_p)^2} \text{cov}(a_i^0, a_j^0) = \frac{\kappa_i \kappa_j}{(\alpha_p \kappa_p)^2} c_{ij} = \frac{1}{\alpha_p \kappa_p},$$

and

$$\text{cov}(a_i^0, D) = \frac{\kappa_j}{\alpha_p \kappa_p} \text{cov}(a_i^0, a_j^0) = \frac{\alpha_i}{\alpha_p \kappa_p},$$

we obtain,

$$\tilde{c}_{ij} \equiv \text{cov}(\tilde{a}_i, \tilde{a}_j) = \text{cov}(a_i^0, a_j^0) - 2\text{cov}(a_i^0, D)\alpha_j + \text{cov}(D, D)\alpha_i\alpha_j = c_{ij} - \frac{\alpha_i\alpha_j}{\alpha_p\kappa_p}.$$

How does imposing this constraint affect the expected value of χ^2 ? Substituting $a_i \rightarrow \tilde{a}_i$, we get all the original terms plus a $\Delta\langle\chi^2\rangle$,

$$\Delta\langle\chi^2\rangle = 2\alpha_i\langle Dd_i\rangle - 2\alpha_i b_{ij}\langle a_j D\rangle + \alpha_i\alpha_j b_{ij}\langle D^2\rangle = \alpha_i c_{pj} \kappa_p b_{ij} (\text{cov}(D, D) + \langle D\rangle^2),$$

or

$$\Delta\langle\chi^2\rangle = 1 + \langle D\rangle^2 \alpha_p \kappa_p.$$

That is, if the constraint really reflects reality, i.e. $\langle D\rangle = 0$, then imposing the constraint increases the expected value of χ^2 by exactly unity. However, if the constraint is untrue,

then imposing it will cause χ^2 to go up by more. Hence, if the original model space represents reality and the constrained space continues to do so, then

$$\langle \chi^2 \rangle = N - n + m,$$

where N is the number of data points, n is the number of parameters, and m is the number of constraints.

General Constraint and Proof

As a practical matter, if one has $m > 1$ constraints,

$$\kappa_i^k a_i = z^k, \quad (k = 1 \dots m),$$

these can be imposed sequentially in a computer program using the above formalism. However, suppose in a fit of mathematical purity, you decided you wanted to impose them all at once. Then, following Gould & An (2002 ApJ 565 1381), but with slight notation changes, you should first rewrite,

$$\chi^2(a_i) = (a_i - a_i^0)b_{ij}(a_j - a_j^0) + \chi_0^2,$$

where a_i^0 is the unconstrained solution and χ_0^2 is the value of χ^2 for that solution. At the constrained solution, the gradient of χ^2 must lie in the m dimensional subspace defined by the m constraint vectors κ_i^k . Otherwise it would be possible to reduce χ^2 further while still obeying the constraints. Mathematically,

$$b_{ij}(a_j - a_j^0) + D^k \kappa_i^k = 0,$$

where the D^k are m so far undetermined coefficients. While we do not yet know these parameters, we can already solve this equation for the \tilde{a}_i in terms of them, i.e.,

$$\tilde{a}_i = a_i^0 - D^l \alpha_i^l, \quad \alpha_i^l \equiv c_{ij} \kappa_j^l.$$

Multiplying this equation by the κ_i^k yields m equations,

$$\kappa_i^k \tilde{a}_i = \kappa_i^k a_i^0 - C^{kl} D^l, \quad C^{kl} \equiv \kappa_i^k \alpha_i^l = \kappa_i^k c_{ij} \kappa_j^l.$$

Then applying the m constraints gives,

$$C^{kl} D^l = A^k, \quad A^k \equiv \kappa_i^k a_i^0 - z^k,$$

so that the solution for the D^k is

$$D^k = B^{kl} A^l, \quad B \equiv C^{-1},$$

where now the implied summation is over the m constraints $l = 1, \dots, m$. Note that,

$$\begin{aligned}\text{cov}(A^k, A^l) &= \kappa_i^k \kappa_j^l \text{cov}(a_i^0, a_j^0) = \kappa_i^k \kappa_j^l c_{ij} = C^{kl}, \\ \text{cov}(A^k, D^l) &= B^{lq} \text{cov}(A^k, A^q) = B^{lq} C^{kq} = \delta_{kl}, \\ \text{cov}(D^k, D^l) &= B^{kq} \text{cov}(A^q, D^l) = B^{kq} \delta_{ql} = B^{kl},\end{aligned}$$

and that,

$$\text{cov}(a_i^0, D^k) = B^{kl} \kappa_j^l \text{cov}(a_i^0, a_j^0) = B^{kl} \alpha_i^l.$$

Hence,

$$\text{cov}(\tilde{a}_i, \tilde{a}_j) = c_{ij} - 2\alpha_i^k \text{cov}(D^k, a_j^0) + \alpha_i^k \alpha_j^l \text{cov}(D^k, D^l) = c_{ij} - 2\alpha_i^k B^{kl} \alpha_j^l + \alpha_i^k \alpha_j^l B^{kl}.$$

That is,

$$\tilde{c}_{ij} = c_{ij} - \alpha_i^k B^{kl} \alpha_j^l.$$

Finally, we can evaluate the expected value of χ^2 directly taking into account the m constraints,

$$\langle \chi^2 \rangle = \mathcal{B}_{kl} C_{kl} - \text{cov}(a_i, d_i) + \text{cov}(A^q, D^q) + \mathcal{B}_{kl} \langle y_k \rangle \langle y_l \rangle - \langle a_i \rangle \langle d_i \rangle + \langle A^q \rangle \langle D^q \rangle.$$

That is,

$$\langle \chi^2 \rangle = N - n + m + \langle y_k - y_{k,\text{mod}} \rangle \mathcal{B}_{kl} \langle y_l - y_{l,\text{mod}} \rangle + \langle A^p \rangle B^{pq} \langle A^q \rangle.$$

Summary

The d_i (products of the data y_k with the trial functions $f_i(x_l)$) and the b_{ij} (products of the trial functions with each other),

$$d_i \equiv y_k \mathcal{B}_{kl} f_i(x_l), \quad b_{ij} \equiv f_i(x_k) \mathcal{B}_{kl} f_j(x_l), \quad (\mathcal{B} \equiv \mathcal{C}^{-1}, \quad \mathcal{C}_{kl} \equiv \langle y_k y_l \rangle - \langle y_k \rangle \langle y_l \rangle),$$

are conjugate to the fit parameters a_i and their associated covariance matrix c_{ij} ,

$$\text{cov}(d_i, d_j) = b_{ij}, \quad \text{cov}(a_i, a_j) = c_{ij}, \quad d_i = b_{ij} a_j, \quad a_i = c_{ij} d_j, \quad c = b^{-1}.$$

The parameter errors and covariances c_{ij} can be determined just from the trial functions, without knowing the data values y_k .

Similarly the A^p (products of the unconstrained parameters a_i^0 with the constraints κ_i^p) and the C^{pq} (products of the constraints with each other),

$$A^p \equiv \kappa_i^p a_i^0 - z^p, \quad C^{pq} \equiv \kappa_i^p c_{ij} \kappa_j^q,$$

are conjugate to coefficients of the constrained-parameter adjustments D^p and their associated covariance matrix B^{pq} ,

$$\text{cov}(A^p, A^q) = C^{pq}, \quad \text{cov}(D^p, D^q) = B^{pq}, \quad A^p = C^{pq} D^q, \quad D^p = B^{pq} A^q, \quad B = C^{-1}.$$

And, while the constrained parameters \tilde{a}_i of course require knowledge of the data, the constrained covariance matrix \tilde{c}_{ij} does not,

$$\tilde{a}_i = a_i^0 - D^p \alpha_i^p, \quad \tilde{c}_{ij} = c_{ij} - \alpha_i^p B^{pq} \alpha_j^q.$$

Note that the vector adjustments α_i^p have the same relation to the constraints κ_i^p that the a_i have to the d_i ,

$$\alpha_i^p \equiv c_{ij} \kappa_j^p, \quad \kappa_i^p = b_{ij} \alpha_j^p$$

Finally, the expected value of χ^2 is

$$\langle \chi^2 \rangle = N - n + m + \langle y_k - y_{k,\text{mod}} \rangle \mathcal{B}_{kl} \langle y_l - y_{l,\text{mod}} \rangle + \langle A^p \rangle B^{pq} \langle A^q \rangle.$$

That is, the number of data points, less the number of parameters, plus the number of constraints, plus two possible additional terms. The first is zero if the model space spans the system being measured ($\langle y_{k,\text{mod}} \rangle = \langle y_k \rangle$), but otherwise is strictly positive. The second is zero if the constraints are valid ($\langle A^p \rangle = 0$), but otherwise is strictly positive.

acknowledgment:

I thank David Weinberg and Zheng Zheng for making many helpful suggestions including the idea to place these notes on astro-ph.