

Lecture 24: Weighted and Generalized Least Squares

1 Weighted Least Squares

When we use ordinary least squares to estimate linear regression, we minimize the mean squared error:

$$MSE(\mathbf{b}) = \frac{1}{n} \sum_{i=1}^n (Y_i - \mathbf{X}_i \cdot \mathbf{b})^2 \quad (1)$$

where \mathbf{X}_i is the i^{th} row of \mathbf{X} . The solution is

$$\hat{\beta}_{OLS} = (\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T \mathbf{Y}. \quad (2)$$

Suppose we minimize the weighted MSE

$$WMSE(\mathbf{b}, w_1, \dots, w_n) = \frac{1}{n} \sum_{i=1}^n w_i (Y_i - \mathbf{X}_i \cdot \mathbf{b})^2. \quad (3)$$

This includes ordinary least squares as the special case where all the weights $w_i = 1$. We can solve it by the same kind of linear algebra we used to solve the ordinary linear least squares problem. If we write \mathbf{W} for the matrix with the w_i on the diagonal and zeroes everywhere else, then

$$WMSE = n^{-1} (\mathbf{Y} - \mathbf{X}\mathbf{b})^T \mathbf{W} (\mathbf{Y} - \mathbf{X}\mathbf{b}) \quad (4)$$

$$= \frac{1}{n} (\mathbf{Y}^T \mathbf{W} \mathbf{Y} - \mathbf{Y}^T \mathbf{W} \mathbf{X} \mathbf{b} - \mathbf{b}^T \mathbf{X}^T \mathbf{W} \mathbf{Y} + \mathbf{b}^T \mathbf{X}^T \mathbf{W} \mathbf{X} \mathbf{b}). \quad (5)$$

Differentiating with respect to \mathbf{b} , we get as the gradient

$$\nabla_{\mathbf{b}} WMSE = \frac{2}{n} (-\mathbf{X}^T \mathbf{W} \mathbf{Y} + \mathbf{X}^T \mathbf{W} \mathbf{X} \mathbf{b}).$$

Setting this to zero at the optimum and solving,

$$\hat{\beta}_{WLS} = (\mathbf{X}^T \mathbf{W} \mathbf{X})^{-1} \mathbf{X}^T \mathbf{W} \mathbf{Y}. \quad (6)$$

But why would we want to minimize Eq. 3?

1. *Focusing accuracy.* We may care very strongly about predicting the response for certain values of the input — ones we expect to see often again, ones where mistakes are especially costly or embarrassing or painful, etc. — than others. If we give the points near that region big weights, and points elsewhere smaller weights, the regression will be pulled towards matching the data in that region.
2. *Discounting imprecision.* Ordinary least squares minimizes the squared error when the variance of the noise terms ϵ is constant over all observations, so we're measuring the regression function with the same precision elsewhere. This situation, of constant noise variance, is called **homoskedasticity**. Often however the magnitude of the noise is not constant, and the data are **heteroskedastic**. Let $\sigma_i^2 = \mathbb{E}[\epsilon_i | X_i] = \sigma_i^2$. Homoskedasticity means that $\sigma_i^2 = \sigma^2$ for all i . Heteroskedasticity means that the σ_i^2 can be different.

When we have heteroskedasticity, ordinary least squares is no longer the optimal estimate — we'll see soon that other estimators can be unbiased and have smaller variance. If however we know the noise variance σ_i^2 at each measurement i , and set $w_i = 1/\sigma_i^2$, we get to minimize the variance of estimation.

3. *Doing something else.* There are a number of other optimization problems which can be transformed into, or approximated by, weighted least squares. The most important of these arises from **generalized linear models**, where the mean response is some nonlinear function of a linear predictor; we will look at them in 402.

2 Heteroskedasticity

Suppose that

$$Y_i = \beta_0 + \beta_1 X_{1i} + \cdots + \beta_p X_{pi} + \epsilon_i$$

where $\mathbb{E}[\epsilon_i] = 0$ and $\text{Var}[\epsilon_i] = \sigma_i^2$. (As usual, we are treating the X_i 's as fixed.) This is called the *Heteroskedastic linear regression model*. **For now, assume we know $\sigma_1, \dots, \sigma_p$.**

The model is

$$\mathbf{Y} = \mathbf{X}\beta + \epsilon$$

where now (treating \mathbf{X} as fixed), $\mathbf{E}[\epsilon] = \mathbf{0}$ and $\text{Var}(\epsilon) = \Sigma$ and Σ is no longer of the form $\sigma^2 I$. The weighted least squares estimator is

$$\hat{\beta} = KY$$

where $K = (\mathbf{X}^T \mathbf{W} \mathbf{X})^{-1} \mathbf{X}^T \mathbf{W}$. So,

$$\hat{\beta} = KY = K(\mathbf{X}\beta + \epsilon) = \beta + K\epsilon.$$

Hence,

$$\mathbb{E}[\hat{\beta}] = \beta, \quad \text{Var}(\hat{\beta}) = K\Sigma K^T. \quad (7)$$

Note that the estimator is unbiased.

3 The Gauss-Markov Theorem

We've seen that when we do weighted least squares, our estimates of β are linear in \mathbf{Y} , and unbiased:

$$\hat{\beta} = (\mathbf{X}^T \mathbf{W} \mathbf{X})^{-1} \mathbf{X}^T \mathbf{W} \mathbf{Y}$$

and $\mathbf{E}[\hat{\beta}] = \beta$.

Let us consider a special case: suppose we take $\mathbf{W} = \Sigma^{-1}$. Then, from (7)

$$\text{Var}(\hat{\beta}) = K\Sigma K^T = \left[(\mathbf{X}^T \Sigma^{-1} \mathbf{X})^{-1} \mathbf{X}^T \Sigma^{-1} \right] \Sigma \left[\Sigma^{-1} \mathbf{X} (\mathbf{X}^T \Sigma^{-1} \mathbf{X})^{-1} \right] = (\mathbf{X}^T \Sigma^{-1} \mathbf{X})^{-1}. \quad (8)$$

We will now show that $\hat{\beta}$ is, in a certain sense, optimal.

Like any optimality result, it is crucial to lay out carefully the range of possible alternatives, and the criterion by which those alternatives will be compared. The classical optimality result for

estimating linear models is the **Gauss-Markov theorem**, which takes the range of possibilities to be *linear, unbiased estimators of β* , and the criterion to be *variance of the estimator*.

Any linear estimator, say $\tilde{\beta}$, could be written as

$$\tilde{\beta} = \mathbf{Q}\mathbf{Y}$$

where \mathbf{Q} would be a $(p+1) \times n$ matrix. We will show that if $\tilde{\beta}$ is unbiased, then it has larger variance than $\hat{\beta}_{\text{WLS}}$.

For $\tilde{\beta}$ to be an unbiased estimator, we must have

$$\mathbb{E}[\mathbf{Q}\mathbf{Y}|\mathbf{X}] = \mathbf{Q}\mathbf{X}\beta = \beta.$$

Since this must hold for all β and all \mathbf{X} , we have to have $\mathbf{Q}\mathbf{X} = \mathbf{I}$. The variance is then

$$\text{Var}[\mathbf{Q}\mathbf{Y}|\mathbf{X}] = \mathbf{Q}\text{Var}[\epsilon|\mathbf{X}]\mathbf{Q}^T = \mathbf{Q}\Sigma\mathbf{Q} \quad (9)$$

where $\Sigma = \text{Var}[\epsilon|\mathbf{X}]$. Let

$$\mathbf{K} \equiv (\mathbf{X}^T\Sigma^{-1}\mathbf{X})^{-1}\mathbf{X}^T\Sigma^{-1}.$$

Now, whatever \mathbf{Q} might be, we can always write

$$\mathbf{Q} = \mathbf{K} + \mathbf{R} \quad (10)$$

for some matrix \mathbf{R} . The unbiasedness constraint on \mathbf{Q} implies that

$$\mathbf{R}\mathbf{X} = \mathbf{0}$$

because $\mathbf{K}\mathbf{X} = \mathbf{I}$. Now we substitute Eq. 10 into Eq. 9:

$$\text{Var}[\tilde{\beta}] = (\mathbf{K} + \mathbf{R})\Sigma(\mathbf{K} + \mathbf{R})^T \quad (11)$$

$$= \mathbf{K}\Sigma\mathbf{K}^T + \mathbf{R}\Sigma\mathbf{K}^T + \mathbf{K}\Sigma\mathbf{R}^T + \mathbf{R}\Sigma\mathbf{R}^T \quad (12)$$

$$= (\mathbf{X}^T\Sigma^{-1}\mathbf{X})^{-1}\mathbf{X}^T\Sigma^{-1}\Sigma\Sigma^{-1}\mathbf{X}(\mathbf{X}^T\Sigma^{-1}\mathbf{X})^{-1} \quad (13)$$

$$+ \mathbf{R}\Sigma\Sigma^{-1}\mathbf{X}(\mathbf{X}^T\Sigma^{-1}\mathbf{X})^{-1} \\ + (\mathbf{X}^T\Sigma^{-1}\mathbf{X})^{-1}\mathbf{X}^T\Sigma^{-1}\Sigma\mathbf{R}^T \\ + \mathbf{R}\Sigma\mathbf{R}^T$$

$$= (\mathbf{X}^T\Sigma^{-1}\mathbf{X})^{-1}\mathbf{X}^T\Sigma^{-1}\mathbf{X}(\mathbf{X}^T\Sigma^{-1}\mathbf{X})^{-1} \quad (14)$$

$$+ \mathbf{R}\mathbf{X}(\mathbf{X}^T\Sigma^{-1}\mathbf{X})^{-1} + (\mathbf{X}^T\Sigma^{-1}\mathbf{X})^{-1}\mathbf{X}^T\mathbf{R}^T \\ + \mathbf{R}\Sigma\mathbf{R}^T$$

$$= (\mathbf{X}^T\Sigma^{-1}\mathbf{X})^{-1} + \mathbf{R}\Sigma\mathbf{R}^T \quad (15)$$

where the last step uses the fact that $\mathbf{R}\mathbf{X} = \mathbf{0}$ (and so $\mathbf{X}^T\mathbf{R}^T = \mathbf{0}^T$).

Hence,

$$\text{Var}[\tilde{\beta}_i] = (\mathbf{X}^T\Sigma^{-1}\mathbf{X})_{ii}^{-1} + R_i^T\Sigma R_i = \text{Var}[\hat{\beta}_i] + R_i^T\Sigma R_i$$

where R_i denotes the i^{th} row of \mathbf{R} . Since Σ is a covariance matrix, it's positive definite, meaning that $a\Sigma a^T \geq 0$ for any vector a . Hence,

$$\text{Var}[\tilde{\beta}_i] \geq \text{Var}[\hat{\beta}_i].$$

We conclude that WLS, with $\mathbf{W} = \Sigma^{-1}$, has the least variance among all possible linear, unbiased estimators of the regression coefficients.

Notes:

1. If all the noise variances are equal, then we've proved the optimality of OLS.
2. The theorem doesn't rule out linear, biased estimators with smaller variance. As an example, albeit a trivial one, $\mathbf{0Y}$ is linear and has variance $\mathbf{0}$, but is (generally) very biased.
3. The theorem also doesn't rule out non-linear unbiased estimators of smaller variance. Or indeed non-linear biased estimators of even smaller variance.
4. The proof actually doesn't require the variance matrix to be diagonal.

4 Finding the Variance and Weights

So far we have assumed that we know $\sigma_1, \dots, \sigma_p$. Here are some cases where this might be true.

Multiple measurements. The easiest case is when our measurements of the response are actually averages over individual measurements, each with some variance σ^2 . If some Y_i are based on averaging more individual measurements than others, there will be heteroskedasticity. The variance of the average of n_i uncorrelated measurements will be σ^2/n_i , so in this situation we could take $w_i \propto n_i$.

Binomial counts Suppose our response variable is a count, derived from a binomial distribution, i.e., $Y_i \sim \text{Binom}(n_i, p_i)$. We would usually model p_i as a function of the predictor variables — at this level of statistical knowledge, a linear function. This would imply that Y_i had expectation $n_i p_i$, and variance $n_i p_i (1 - p_i)$. We would be well-advised to use this formula for the variance, rather than pretending that all observations had equal variance.

Proportions based on binomials If our response variable is a proportion based on a binomial, we'd see an expectation value of p_i and a variance of $\frac{p_i(1-p_i)}{n_i}$. Again, this is not equal across different values of n_i , or for that matter different values of p_i .

Poisson counts Binomial counts have a hard upper limit, n_i ; if the upper limit is immense or even (theoretically) infinite, we may be better off using a Poisson distribution. In such situations, the mean of the Poisson λ_i will be a (possibly-linear) function of the predictors, and the variance will *also* be equal to λ_i .

Other counts The binomial and Poisson distributions rest on independence across “trials” (whatever those might be). There are a range of discrete probability models which allow for correlation across trials (leadings to more or less variance). These may, in particular situations, be more appropriate.

5 Conditional Variance Function Estimation

If we don't know the variances, we might be able to estimate them. There are two common ways to estimate conditional variances, which differ slightly in how they use non-parametric smoothing. (We will discuss non-parametric smoothing in more detail later and in 402).

Method 1:

1. Estimate $m(x)$ with your favorite regression method, getting $\hat{m}(x)$.
2. Construct the **squared residuals**, $u_i = (Y_i - \hat{m}(x_i))^2$.
3. Use your favorite *non-parametric* method to estimate the conditional mean of the u_i , call it $\hat{q}(x)$.
4. Estimate the variance using $\hat{\sigma}_x^2 = \hat{q}(x)$.

Here is method 2:

1. Estimate $m(x)$ with your favorite regression method, getting $\hat{m}(x)$.
2. Construct the **log squared residuals**, $z_i = \log(Y_i - \hat{m}(x_i))^2$.
3. Use your favorite *non-parametric* method to estimate the conditional mean of the z_i , call it $\hat{s}(x)$.
4. Predict the variance using $\hat{\sigma}_x^2 = \exp(\hat{s}(x))$.

The second method ensures that the estimates variances are positive.

We are estimating the variance function to do weighted least squares, but these methods can be used more generally. It's often important to understand variance in its own right, and this is a general method for estimating it. Our estimate of the variance function depends on first having a good estimate of the regression function

5.1 Example

```
> plot(x,y)
> out = lm(y ~ x)
> abline(out)
> summary(out)
Coefficients:
              Estimate Std. Error t value Pr(>|t|)
(Intercept)  0.02934    0.03935   0.745    0.458
x            1.89866    0.07092  26.772 <2e-16 ***

> plot(x,rstudent(out))
> abline(h=0)
>
> u = log((resid(out))^2)
> tmp = loess(u ~ x)
> s2 = exp(tmp$fitted)
```

```

> plot(x,s2)
> w = 1/s2
> out2 = lm(y ~ x,weights=w)
> summary(out2)
Coefficients:
              Estimate Std. Error t value Pr(>|t|)
(Intercept) 0.002967   0.008252   0.36   0.72
x            2.046929   0.043872  46.66 <2e-16 ***

```

6 Correlated Noise and Generalized Least Squares

Sometimes, we might believe the right model is (in matrix form)

$$\mathbf{Y} = \mathbf{X}\beta + \epsilon \quad (16)$$

$$\mathbb{E}[\epsilon|\mathbf{X}] = \mathbf{0} \quad (17)$$

$$\text{Var}[\epsilon|\mathbf{X}] = \mathbf{\Sigma} \quad (18)$$

where the matrix $\mathbf{\Sigma}$ is *not* diagonal. The off-diagonal entries represent covariance in the noise terms, $\text{Cov}[\epsilon_i, \epsilon_j] = \Sigma_{ij}$. How should we estimate β ? Here are two approaches. We assume that $\mathbf{\Sigma}$ is known.

Approach 1. Because $\mathbf{\Sigma}$ is a variance matrix, we know it is square, symmetric, and positive-definite. This implies that there exists a square root matrix \mathbf{S} such that $\mathbf{S}\mathbf{S} = \mathbf{\Sigma}$. Now we multiply our model by \mathbf{S}^{-1} :

$$\mathbf{S}^{-1}\mathbf{Y} = \mathbf{S}^{-1}\mathbf{X}\beta + \mathbf{S}^{-1}\epsilon$$

This is a linear regression of $\mathbf{S}^{-1}\mathbf{Y}$ on $\mathbf{S}^{-1}\mathbf{X}$, with the same coefficients β as our original regression. However, we have improved the properties of the noise. The noise is still zero in expectation,

$$\mathbb{E}[\mathbf{S}^{-1}\epsilon|\mathbf{X}] = \mathbf{S}^{-1}\mathbf{0} = \mathbf{0}$$

but

$$\text{Var}[\mathbf{S}^{-1}\epsilon|\mathbf{X}] = \mathbf{S}^{-1}\text{Var}[\epsilon|\mathbf{X}]\mathbf{S}^{-T} \quad (19)$$

$$= \mathbf{S}^{-1}\mathbf{\Sigma}\mathbf{S}^{-T} \quad (20)$$

$$= \mathbf{S}^{-1}\mathbf{S}\mathbf{S}^T\mathbf{S}^{-T} \quad (21)$$

$$= \mathbf{I}. \quad (22)$$

To sum up, if we know $\mathbf{\Sigma}$, we can estimate β by doing an ordinary least squares regression of $\mathbf{S}^{-1}\mathbf{Y}$ on $\mathbf{S}^{-1}\mathbf{X}$. The estimate is

$$\hat{\beta} = ((\mathbf{S}^{-1}\mathbf{X})^T\mathbf{S}^{-1}\mathbf{X})^{-1}(\mathbf{S}^{-1}\mathbf{X})^T\mathbf{S}^{-1}\mathbf{Y} \quad (23)$$

$$= (\mathbf{X}^T\mathbf{S}^{-T}\mathbf{S}^{-1}\mathbf{X})^{-1}\mathbf{X}^T\mathbf{S}^{-T}\mathbf{S}^{-1}\mathbf{Y} \quad (24)$$

$$= (\mathbf{X}^T\mathbf{\Sigma}^{-1}\mathbf{X})^{-1}\mathbf{X}^T\mathbf{\Sigma}^{-1}\mathbf{Y}. \quad (25)$$

This looks just like our weighted least squares estimate, only with $\mathbf{\Sigma}^{-1}$ in place of \mathbf{W} .

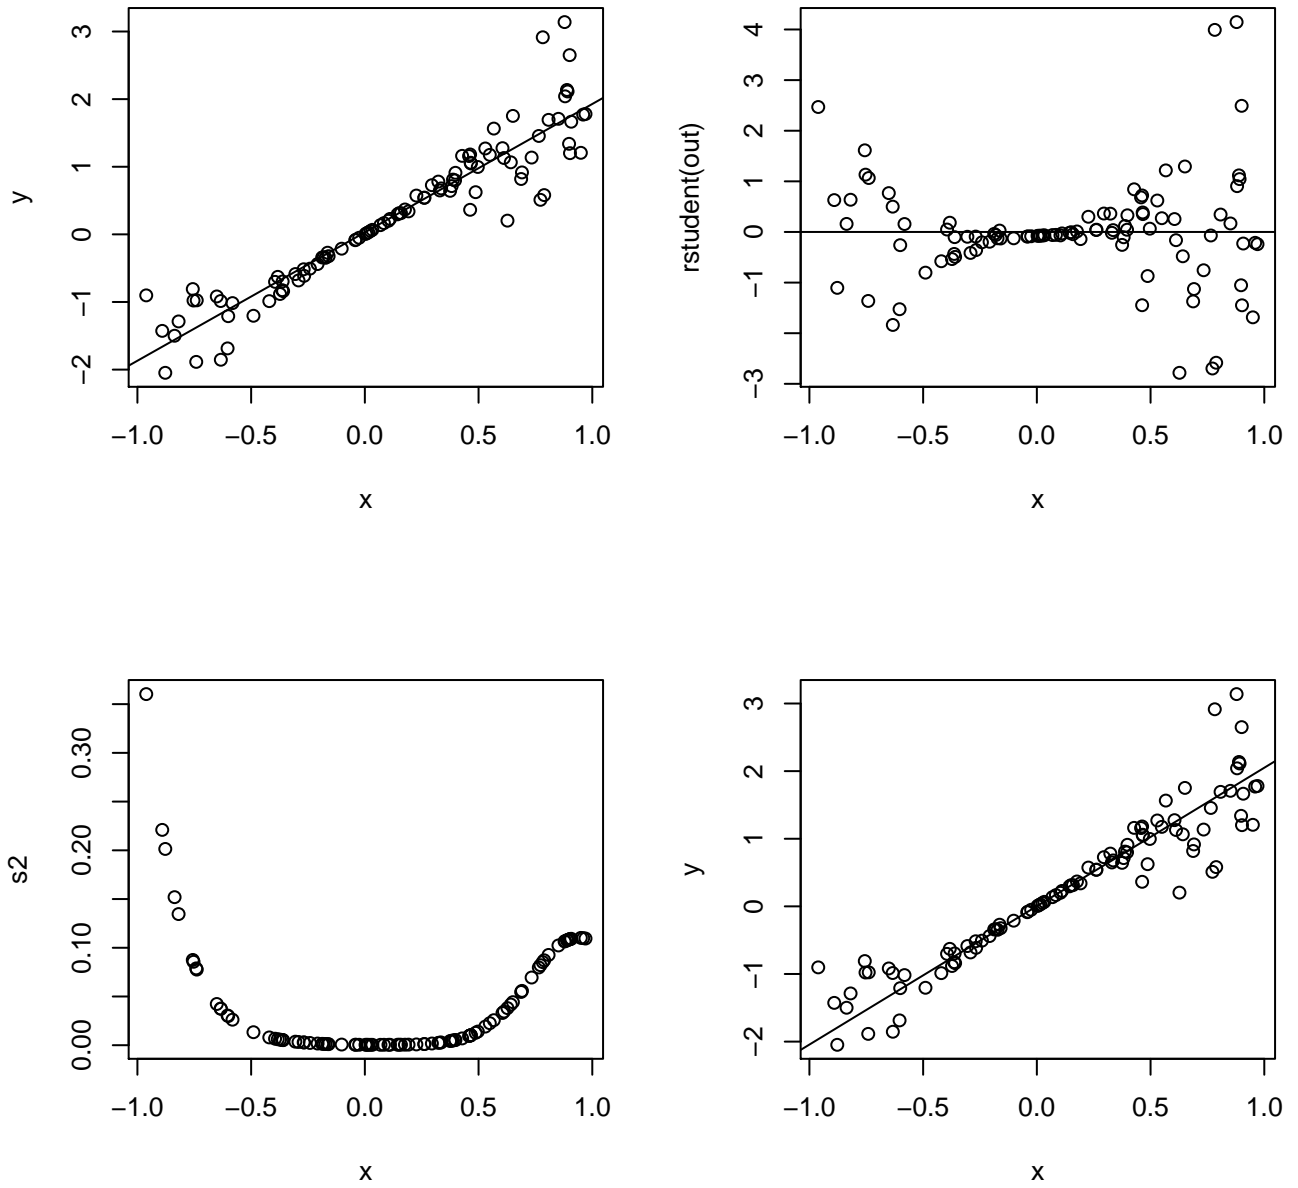


FIGURE 1: *Top left: data with fitted line. Top right: residuals. Bottom left: Plot of estimates variances. Bottom right: Data with weighed least squares line.*

Approach 2. This resemblance is no mere coincidence. We can write the WLS problem as that of minimizing $(\mathbf{Y} - \mathbf{X}\beta)^T \mathbf{W}(\mathbf{Y} - \mathbf{X}\beta)$, for a diagonal matrix \mathbf{W} . Suppose we try instead to minimize

$$(\mathbf{Y} - \mathbf{X}\beta)^T \mathbf{W}(\mathbf{Y} - \mathbf{X}\beta)$$

for a *non-diagonal*, but still symmetric and positive-definite, matrix \mathbf{W} . This is called a **generalized least squares** (GLS) problem. Every single step we went through before is still valid, because none of it rested on \mathbf{W} being diagonal, so

$$\hat{\beta}_{GLS} = (\mathbf{X}^T \mathbf{W} \mathbf{X})^{-1} \mathbf{X}^T \mathbf{W} \mathbf{Y}. \quad (26)$$

We have just seen is that if we set $\mathbf{W} = \Sigma^{-1}$, we also get this solution when we transform the variables so as to de-correlate the noise, and then do ordinary least squares. This should at least make it *plausible* that this is a good way to estimate β in the face of correlated noise.

To go beyond plausibility, refer back to §3. At no point in our reasoning did we actually rely on $\text{Var}[\epsilon|\mathbf{X}]$ being diagonal. It follows that if we set $\mathbf{W} = \text{Var}[\epsilon|\mathbf{X}]^{-1}$, we get the linear, unbiased estimator of minimum variance. If we believe that the noise is Gaussian, then this is also the maximum likelihood estimator.

Where Do the Covariances Come From? In general we cannot estimate Σ . So we can only use this approach when Σ comes from external information. For example, the data may have some time series structure or spatial structure which suggests the form of Σ .

7 WLS versus Model Errors

When you find that your residuals from an initial model have non-constant variance or are correlated with each other, there are (at least) two possible explanations. One is that the fluctuations around the regression line really are heteroskedastic and/or correlated. In that case, you should try to model that variance and those correlations, and use WLS or GLS. The other explanation is that something is wrong with your model. If there's an important predictor variable which is just missing from your model, for example, then its contribution to the response will be part of your residuals. If that omitted variable is larger in some parts of the data than in others, or if the omitted variable has correlations, then that will make your residuals change in magnitude and be correlated. More subtly, having the wrong functional form for a variable you do include can produce those effects as well.