Assessing uncertainty in the estimate of trend

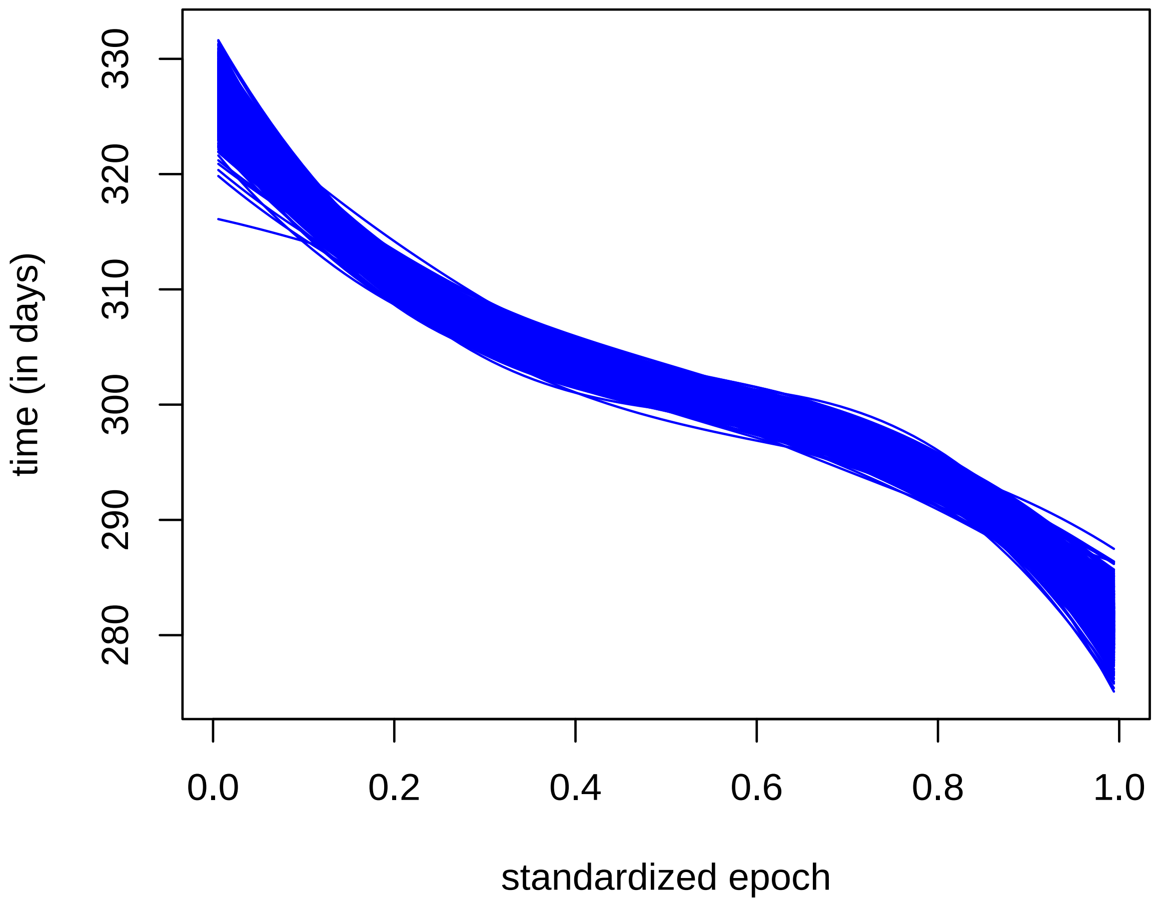

The blue lines are 1000 cubics obtained by selecting 1000 draws at random and with replacement from the 10,000 points generated by MCMC.

Posterior maximizers under parameter transformations

We know that MLEs are invariant under transformations of the parameter. This means that if $\widehat{\boldsymbol{\theta}}$ is the MLE of $\boldsymbol{\theta}$, then $g(\widehat{\boldsymbol{\theta}})$ is the MLE of $g(\boldsymbol{\theta})$.

This invariance property *does not* hold for the maximizers of posterior densities. If $\tilde{\boldsymbol{\theta}}$ maximizes the posterior of $\boldsymbol{\theta}$, then $g(\tilde{\boldsymbol{\theta}})$ does not necessarily maximize the posterior of $g(\boldsymbol{\theta})$.

Consider the case where θ is a scalar, and suppose that the maximizer $\tilde{\theta}$ of π satisfies $\pi'(\tilde{\theta}) = 0$.

If $\tau = e^{\theta}$, then the density of τ is

$$h(\tau) = \pi(\log \tau)\frac{1}{\tau}I_{(0,\infty)}(\tau).$$

Now consider (for $\tau > 0$)

$$\begin{aligned} h'(\tau) &= \frac{\pi'(\log\tau)}{\tau^2} - \frac{\pi(\log\tau)}{\tau^2} \\ &= 0 \Rightarrow \end{aligned}$$

$$\pi'(\log\tau) = \pi(\log\tau).$$

Let $\tilde{\tau} = \exp(\tilde{\theta})$. Then

$$\pi'(\log\tilde{\tau}) = \pi'(\tilde{\theta}) = 0,$$

and so in order for $\tilde{\tau}$ to be a local maximizer of h it would need to be true that $\pi(\tilde{\theta}) = 0$. This, of course, fails to be true for any π that is everywhere positive.

We know that, *under appropriate regularity conditions*, the posterior distribution of a parameter is approximately normal with mean (and hence mode) equal to the MLE.

This implies that, *under appropriate regularity conditions*, the invariance property *does hold* for posterior modes, at least approximately.

> Do the appropriate regularity conditions hold for our *R Aquilae* example?

Probably not!!

An important regularity condition is that the true parameter value be in the interior of the parameter space.

In our example, the parameter r appears to be very close to 0, which is on the boundary of the parameter space.

It may be reasonable to say that the posterior distribution for ρ is approximately normal, but it is definitely **not** reasonable to say that the posterior of r is normal. (Compare the graphs on pp. 292N and 304N.)

Two estimates of the marginal posterior of parameter r

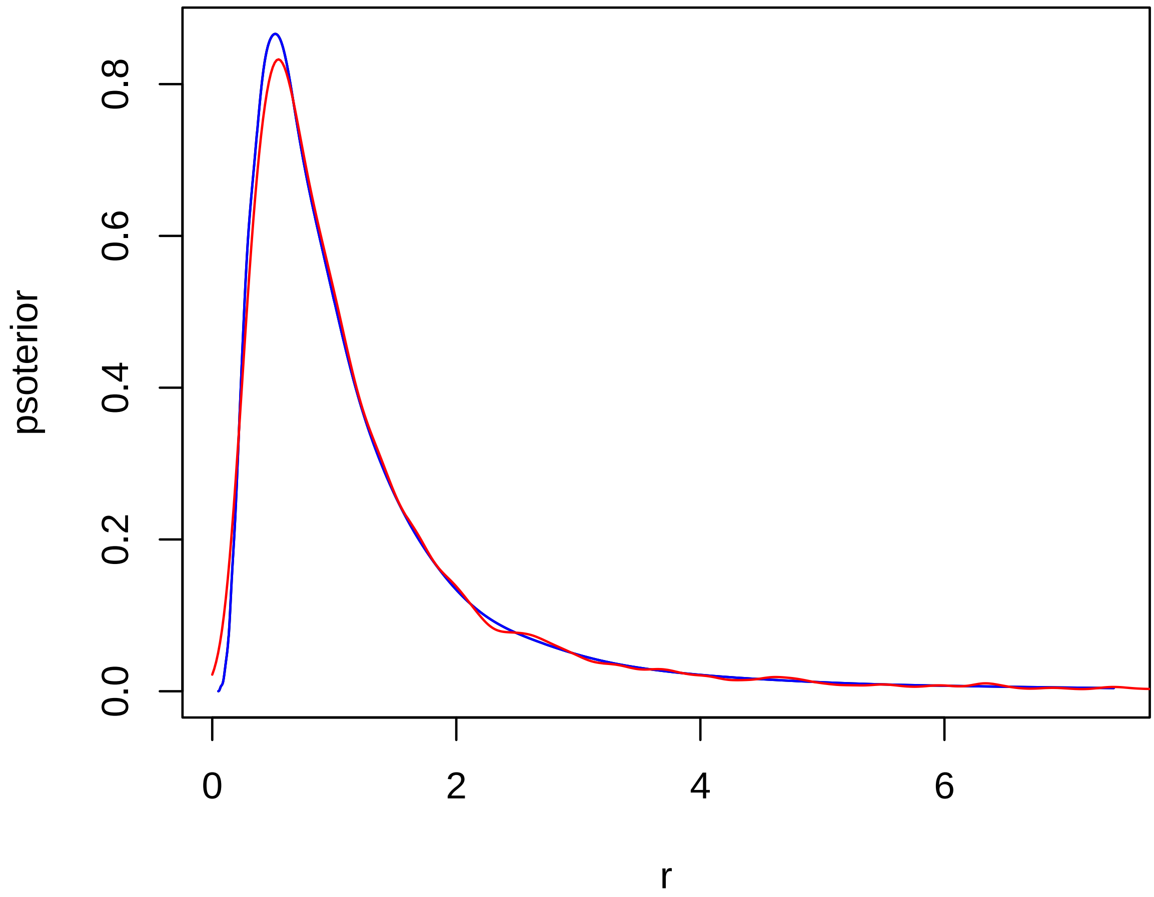

*Transformed: —— Direct: ——

*See the next page for an explanation.

Let $\widehat{\pi}$ be the density estimate of the marginal of ρ, as shown on p. 292N. The "transformed estimate" on p. 304N is

$$\widehat{\pi}(\log r)\frac{1}{r}I_{(0,\infty)}(r),$$

which is appropriate since $\rho = \log r$.

The "direct estimate" on p. 304N is a kernel estimate of the form

$$\widehat{g}(r) = \frac{1}{nh}\sum_{i=1}^{n} K\left(\frac{r - \exp(\rho_i)}{h}\right),$$

where $\rho_1, \ldots, \rho_n$ are the values of ρ generated using MCMC.

The maximizers of the blue and red estimates are 0.519 and 0.545, respectively.

The maximizer of the marginal posterior of ρ on p. 292N is -0.1543, and

$$\exp(-0.1543) = 0.857.$$

The last estimate is substantially different from 0.519 or 0.545.

Also, recall that the overall maximizer of the posterior has $r \approx 0.552$.

Let $\widehat{\theta}_i$ be the ith component of the maximizer $\widehat{\boldsymbol{\theta}}$ of the posterior distribution, and let $\tilde{\theta}_i$ be the maximizer of the marginal posterior of θ_i.

Another important point to keep in mind:

> The quantities $\widehat{\theta}_i$ and $\tilde{\theta}_i$ are not necessarily the same.

More residual analysis of R Aquilae data

Using the fitted cubic on p. 298N, one obtains the following residuals.

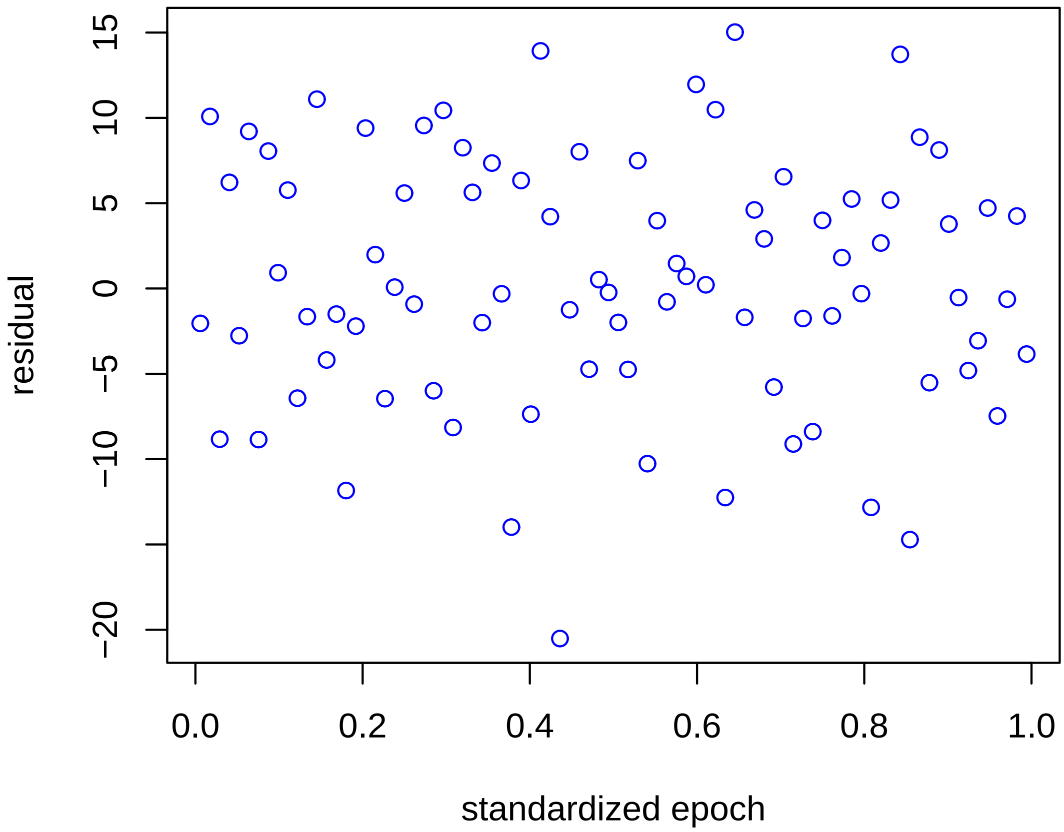

The model used for the error series suggests that the residuals one lag apart will be negatively correlated, while those *more* than one lag apart will be nearly independent.

Consider the following scatterplot of (e_{i-1}, e_i), $i = 2, \ldots, n$, where $e_1, \ldots, e_n$ are the residuals.

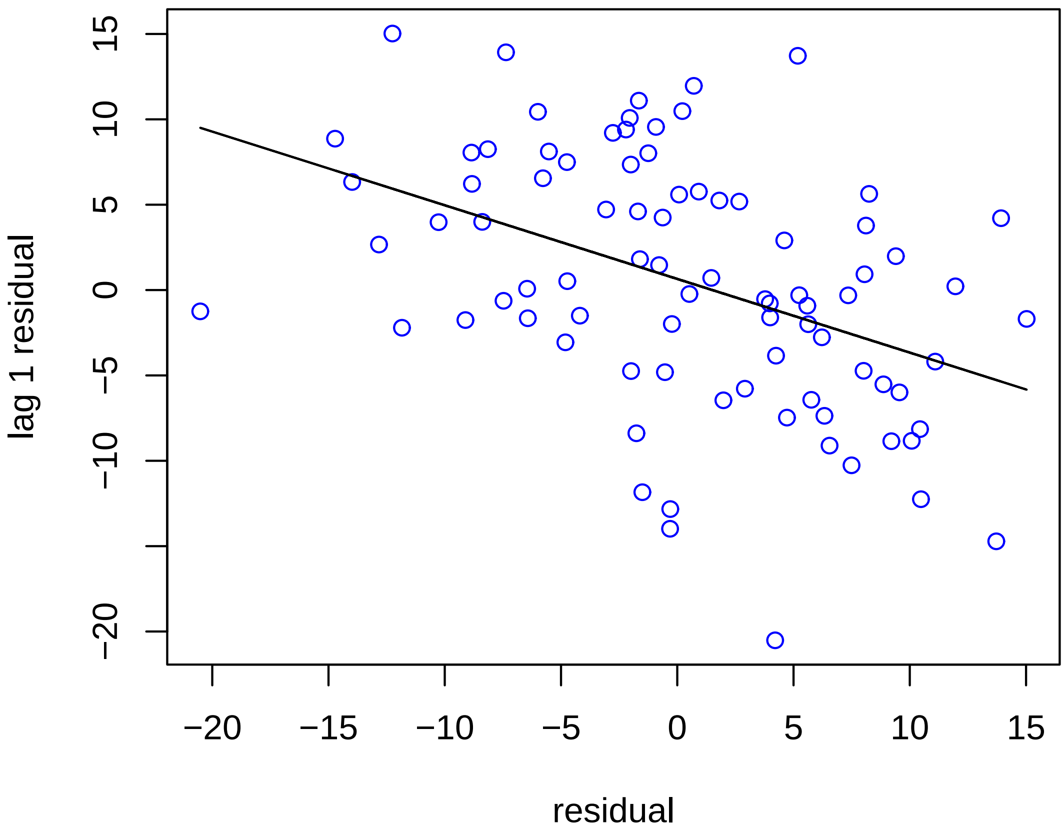

The lag 1 sample autocorrelation is -0.43, and the fitted line has slope significantly different than 0 (P-value < 0.0001).

Scatterplot of (e_{i-2}, e_i), $i = 3, \ldots, n$

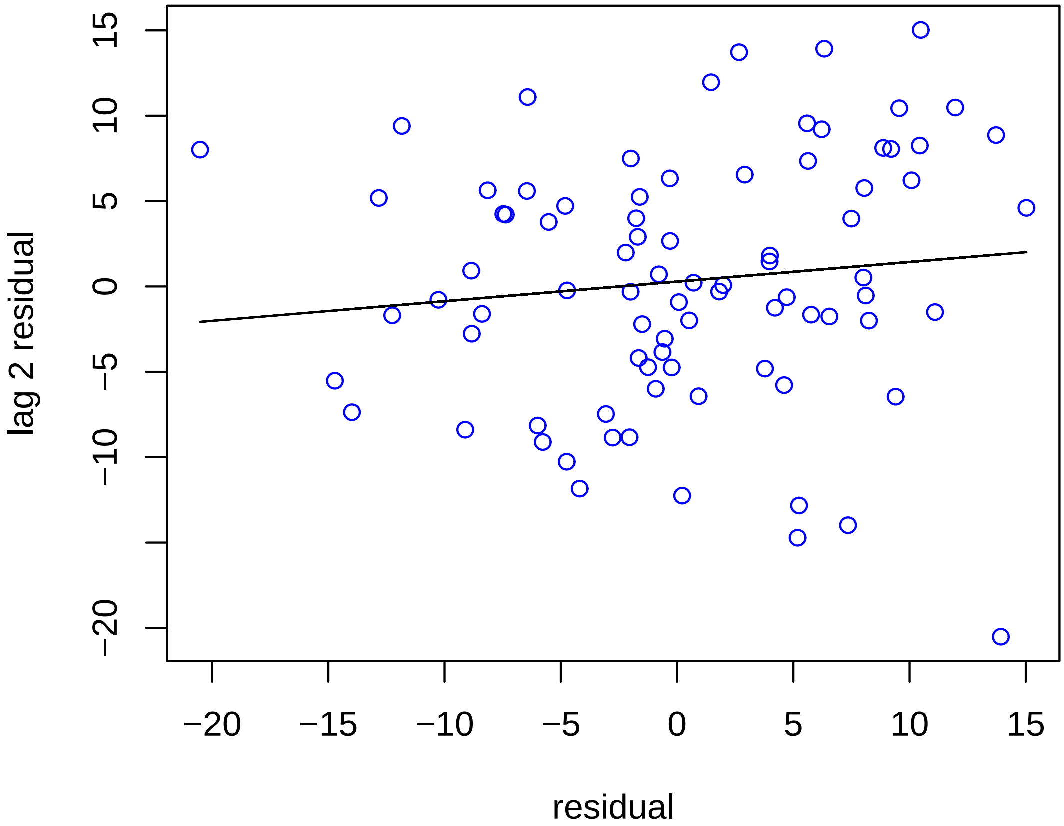

The lag 2 sample autocorrelation is 0.12, and the slope of the line is not significantly different than 0 (P-value = 0.29).