tion values. The fitted values shown above used the (apparently) global maximum of the approximate restricted likelihood but would have been nearly identical using the maximum of the exact restricted likelihood. This is not always the case with other series; I was outlandishly lucky with my first example.

Perhaps the moral of the story is that this model is too complex for a data series of length 650. This points to a related puzzle: why is the error variance always zero? In all the models I've fit to this series, with many starting values, I have only rarely found even a secondary maximum at which σ_e^2 was not effectively zero. It appears this DLM differs from Chapter 16's two-variance models in that for all j, some non-error component competes with error to explain $\hat{v}_j^2$, i.e., no range of j is distinctively informative about σ_e^2. Based on Figure 17.3, I hypothesize that the heartbeat component competes with error because heartbeat's $a_{jh} > 1$ for large j, unlike the other components' a_{jk}. In a variant of Model 2 omitting the heartbeat component, error absorbed variation previously captured by heartbeat, while the rest of the model's fit was essentially unchanged. If the data series were much longer, the a_{jk} for all components would become small enough for the largest j so that those $\hat{v}_j^2$ would provide information about σ_e^2. Nothing, however, can alleviate the confounding of signal and mystery: signal's fit is necessarily sensitive to the first few $\hat{v}_j^2$.

Exercises

Regular Exercises

1. (Section 17.1.1) For the additive model with two predictors, construct some examples, compute $\mathbf{H}'\mathbf{Z}_1\mathbf{Z}_1'\mathbf{H}\mathbf{H}'\mathbf{Z}_2\mathbf{Z}_2'\mathbf{H}$ for each example, and check its symmetry. Construct at least one balanced example and one unbalanced example, as "balanced" is defined in Section 17.1.1, and show that $\mathbf{H}'\mathbf{Z}_1\mathbf{Z}_1'\mathbf{H}\mathbf{H}'\mathbf{Z}_2\mathbf{Z}_2'\mathbf{H}$ is symmetric for the balanced example but not the unbalanced example. If you can construct an unbalanced example for which $\mathbf{H}'\mathbf{Z}_1\mathbf{Z}_1'\mathbf{H}\mathbf{H}'\mathbf{Z}_2\mathbf{Z}_2'\mathbf{H}$ *is* symmetric, you've proven that balance is sufficient but not necessary; please send me that example!
2. (Section 17.1.1) For the additive model with two predictors, show that the restricted likelihood can be put in the scalar form if the predictors are balanced as defined in Section 17.1.1. When I proved this, I didn't show that $\mathbf{H}'\mathbf{Z}_1\mathbf{Z}_1'\mathbf{H}\mathbf{H}'\mathbf{Z}_2\mathbf{Z}_2'\mathbf{H}$ is symmetric; rather, I followed Section 15.1's derivation by constructing the matrices Γ_X, Γ_Z, $\mathbf{M}$, $\mathbf{A}$, and $\mathbf{P}$, the latter four of which partition into parts corresponding to the two predictors and their respective splines, so in effect this model is two two-variance models joined together.
3. (Section 17.1.2) Prove all the assertions made in the two paragraphs defining general balance.
4. (Section 17.1.2) For the two-crossed-random-effects model in Section 12.1.2, show that the restricted likelihood, equation (12.18), is the likelihood for a gamma-errors GLM with identity link by identifying the "observations," linear predictor, etc., for that GLM.
5. (Section 17.1.2) Derive the restricted likelihood for the tidied-up epidermal nerve