Supplementary Material for Morris and Carroll (2005)

1 Metropolis-Hastings Method for Covariance Parameters

Here we outline the details of the Metropolis-Hastings procedure we use to update the covariance parameters $\mathbf{\Omega}_P, \mathbf{\Omega}_R, \mathbf{\Omega}_Q$, and $\mathbf{\Omega}_S$.

First, the covariance parameters q^*_{jk} and s^*_{jk} of $\mathbf{\Omega}_Q$ and $\mathbf{\Omega}_S$ are updated as a block for each $j = 1, \ldots, J$ and $k = 1, \ldots, K_j$, as follows. The objective function is

$$f(q^*_{jk}, s^*_{jk} | \mathbf{d}_{jk}, \boldsymbol{\beta}^*_{jk}, \mathbf{\Omega}_P, \mathbf{\Omega}_R) \propto |\Sigma_{jk}|^{-1/2} [\exp\{-1/2(\mathbf{d}_{jk} - X\boldsymbol{\beta}^*_{jk})' \Sigma^{-1}_{jk} (\mathbf{d}_{jk} - X\boldsymbol{\beta}^*_{jk})\}] f(q^*_{jk}, s^*_{jk}), , \quad (1)$$

where recall $\Sigma_{jk} = ZP(\mathbf{\Omega}_P)Z' * q^*_{jk} + R(\mathbf{\Omega}_R) * s^*_{jk}$. If these parameters are only indexed by j (see Section **??**), then this objective function takes the product of these terms over $k = 1, \ldots, K_j$.

Candidate values $(q1)_{jk}$ and $(s1)_{jk}$ are sampled from a independent truncated normal proposal densities centered at the previous values $(q0)_{jk}$ and $(s0)_{jk}$ with proposal variances $(vq)_{jk}$ and $(vs)_{jk}$, respectively. The density is truncated at 0 to ensure positivity of the variance components. Thus, the proposal density is given by

$$p\{(q1)_{jk}, (s1)_{jk} | (q0)_{jk}, (s0)_{jk}\} \propto \Phi\{-(q0)_{jk}/\sqrt{(vq)_{jk}}\}\{(vq)^{-1/2}_{jk}\} \exp[-\{(q1)_{jk} - (q0)_{jk}\}^2/2(vq)_{jk}] * \Phi\{-(s0)_{jk}/\sqrt{(vs)_{jk}}\}\{(vs)^{-1/2}_{jk}\} \exp[-\{(s1)_{jk} - (s0)_{jk}\}^2/2(vs)_{jk}],$$

where $\Phi(x)$ is the normal cdf evaluated at x.

The sampling proceeds as follows for each j and k at each iteration of the MCMC.

1. Generate proposal values $(q1)_{jk}$ and $(s1)_{jk}$ from TrN$\{(q0)_{jk}, (vq)_{jk}\}$ and TrN$\{(s0)_{jk}, (vq)_{jk}\}$, respectively, where TrN refers to a normal truncated at 0 and $(q0)_{jk}$ and $(s0)_{jk}$ are the values of q^*_{jk} and s^*_{jk} from the previous iteration.

2. Compute a_{jk}, given by

$$a_{jk} = \frac{f((q1)_{jk}, (s1)_{jk} | \mathbf{d}_{jk}, \boldsymbol{\beta}^*_{jk}, \mathbf{\Omega}_P, \mathbf{\Omega}_R) p((q0)_{jk}, (s0)_{jk} | (q1)_{jk}, (s1)_{jk})}{f((q0)_{jk}, (s0)_{jk} | \mathbf{d}_{jk}, \boldsymbol{\beta}^*_{jk}, \mathbf{\Omega}_P, \mathbf{\Omega}_R) p((q1)_{jk}, (s1)_{jk} | (q0)_{jk}, (s0)_{jk})}.$$

3. Generate u from a U(0,1), if $u < a_{jk}$, then let $q^*_{jk} = (q1)_{jk}$ and $s^*_{jk} = (s1)_{jk}$ for the current iteration, otherwise $q^*_{jk} = (q0)_{jk}$ and $s^*_{jk} = (s0)_{jk}$.

It is important to carefully choose the proposal variances $(vq)_{jk}$ and $(vs)_{jk}$ since if they are too large, the acceptance probabilities will be low and the sampler inefficient, while if too small, the acceptance probabilities will be too high and the sampler may not adequately explore the space of the objective function. Often trial and error is used to come up with proposal variances with good properties. This is infeasible here because of the very large number of parameters, so we instead determine them automatically from the data. We compute rough estimates of the asymptotic variances for the maximum likelihood estimators of q^*_{jk} and s^*_{jk}, conditional on initial estimates of $\boldsymbol{\beta}^*_{jk}$. Since the amount of variability in estimating the MLE should be a reasonable ballpark estimate for the variance of the objective function given weak prior information, this should be a good starting point. In order to err on one side of taking slightly larger jumps, we multiplied these estimates by a factor of 1.5 to get the proposal variances. We found that this strategy led to reasonable acceptance probabilities for all of our motivating examples.

Similar steps can be followed to update the parameters contained in $\boldsymbol{\Omega}_P$ and $\boldsymbol{\Omega}_R$.

Some trace plots from the MCMC

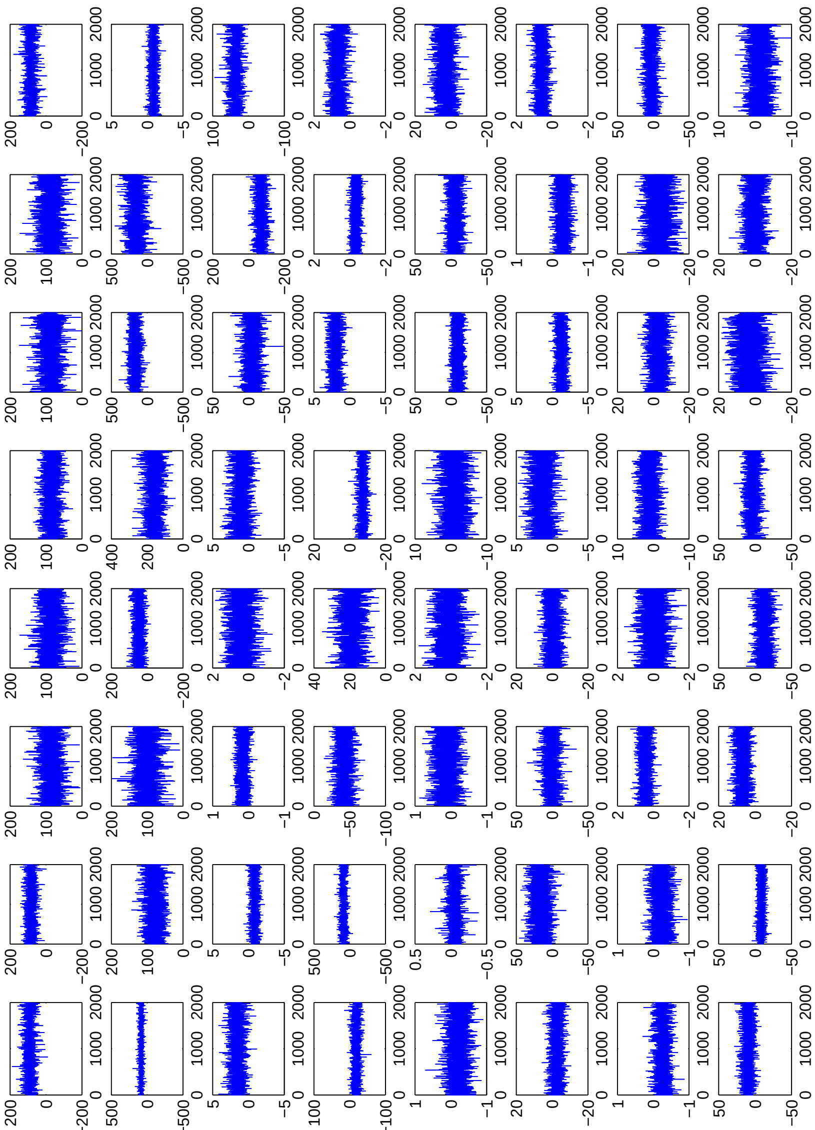

Figure 1: *Trace plot of subset of wavelet coefficients for* 1^{st} *fixed effect function.* Trace plots of 64 wavelet coefficients for the first fixed effect function, the mean MGMT profiles for rats in the fish oil, time 0 group. There are over 3000 wavelet coefficients for the 14 fixed effect functions altogether. These are the 64 lowest frequency coefficients, which account for a vast majority of the variability in the functions (99.79% of the total energy). These trace plots suggest the chain has converged and is mixing nicely.

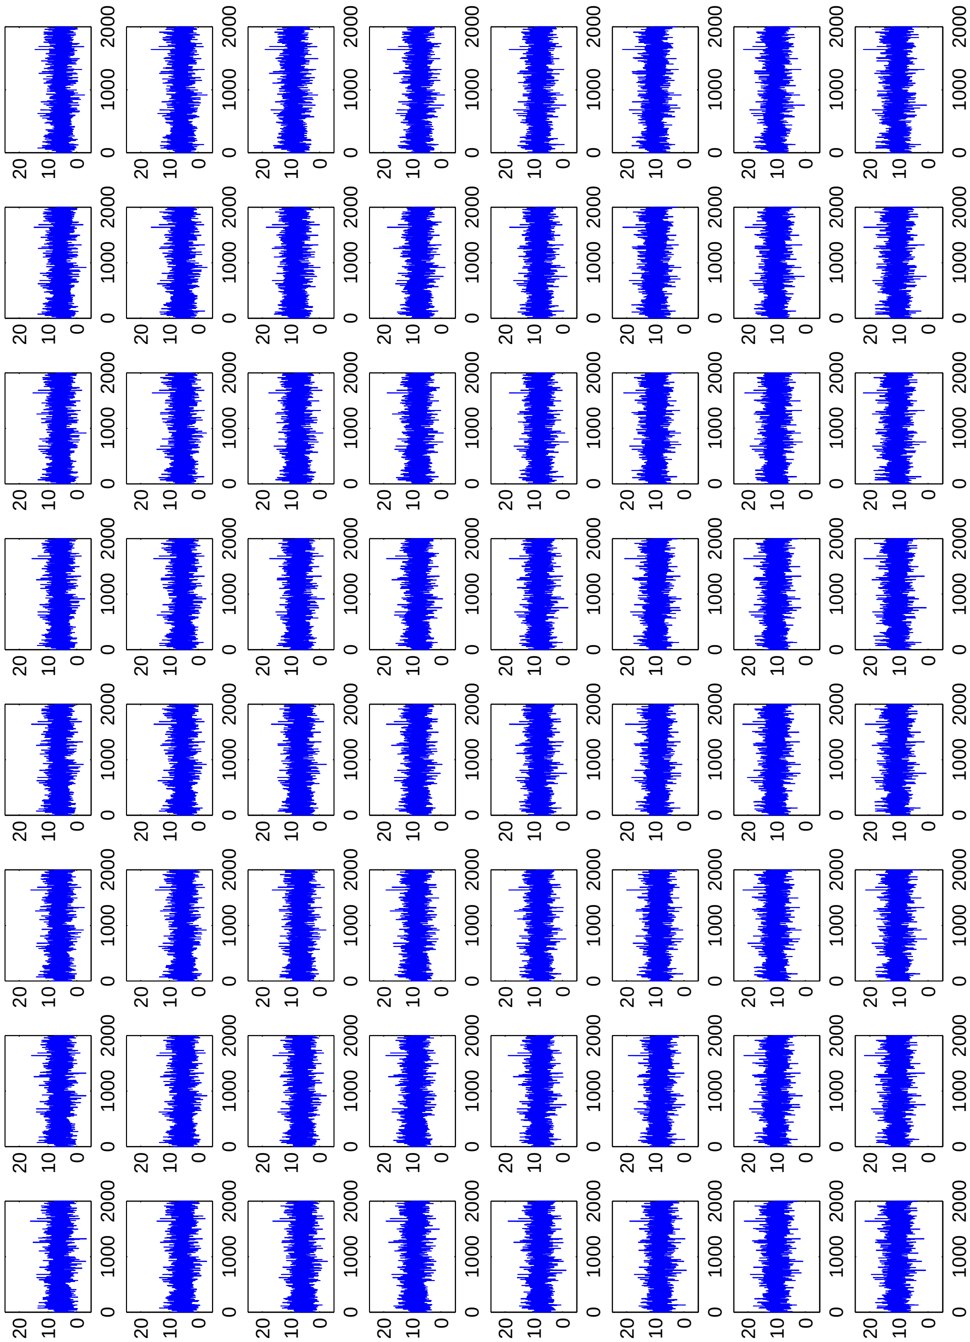

Figure 2: *Trace plot of* 1^{st} *fixed effect function at each of 64 grid points on the interval (0,0.25).*

These trace plots suggest again, that the chain has converged and is mixing nicely.

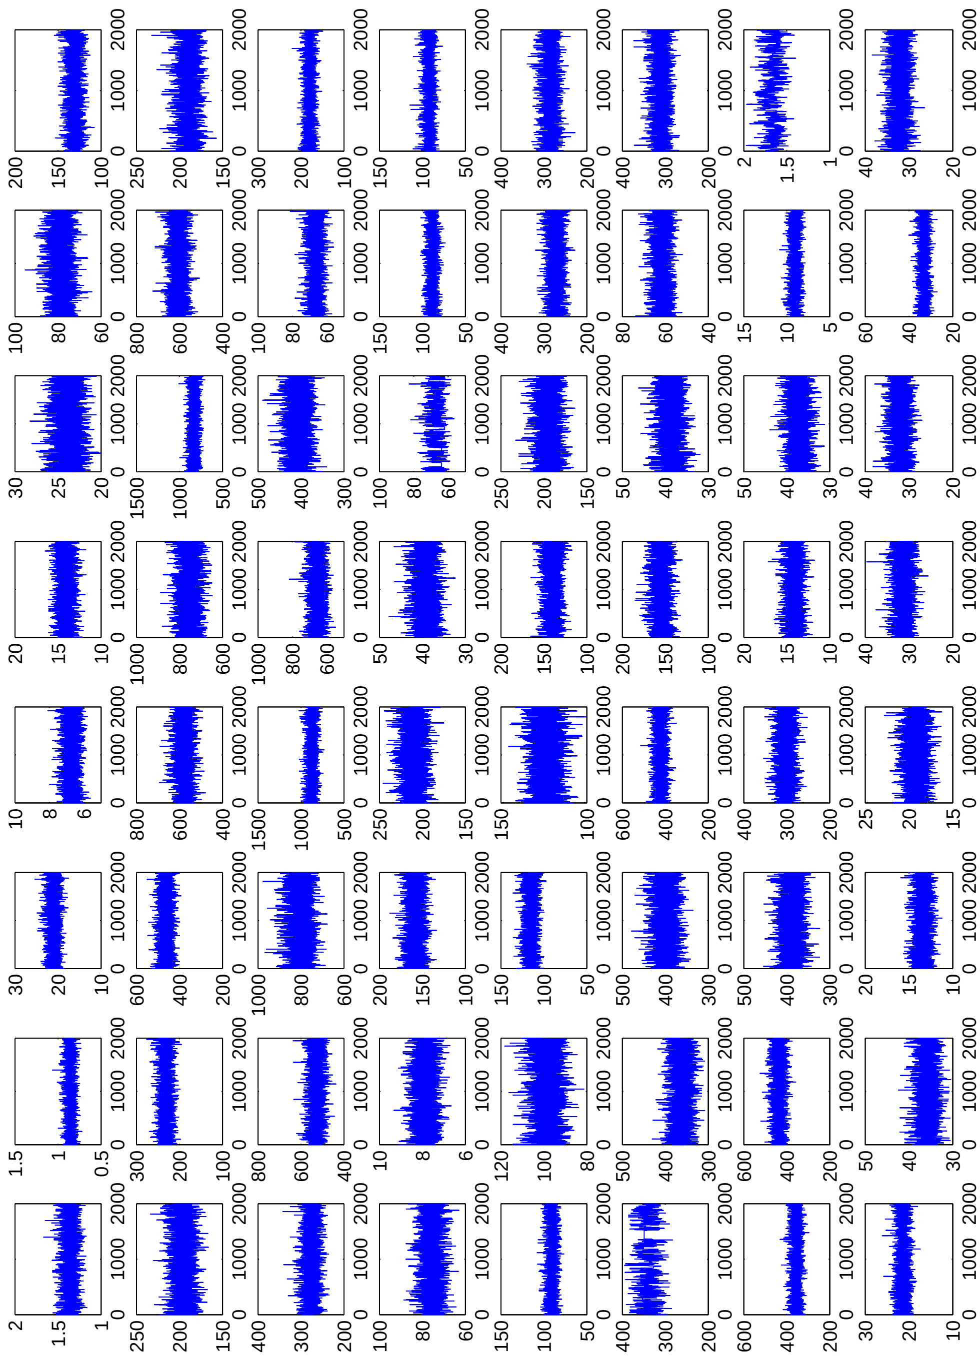

Figure 3: *Trace plot of subset of variance components.* Trace plot of the 64 lowest frequency curve-to-curve wavelet space variance components in the model. These 64 coefficients account for 99.65% of the total energy in terms of the curve-to-curve variability. This plot contains just 64 of the over 500 variance components in the model. These trace plots suggest the chain has converged and is mixing nicely.