Assessing Homeostasis through Circadian Patterns

Rafael A. Irizarry*, Clarke Tankersley, Robert Frank and Susan Flanders†

June 19, 2001

SUMMARY. An organism is thought to be in a dynamic state of homeostasis when each physiological and behavioral system reaches a delicate balance within the framework of other regulatory processes. Many biological systems target specific set-point variables and generate circadian patterns. In this paper, we focus on specific measurements representative of two systems, namely deep body temperature and activity counts. We examine data collected every 30 minutes in mice, assume there are underlying circadian patterns, and extend the approach presented in Brumback and Rice (1998) in order to obtain estimates in the presence of correlated data. We then assess homeostasis using these estimates and their statistical properties.

KEY WORDS: Circadian patterns; Correlated errors; Mixed-effect models; Nested Curves.

*Department of Biostatistics, Johns Hopkins University, Baltimore, MD 21205. E-mail: rafa@jhu.edu, phone: 410-614-5157, fax: 410-955-0958

†Department of Environmental Health Sciences, Johns Hopkins University, Baltimore, MD 21205.

1 Introduction

Most biological systems function under complex regulatory processes. Under optimal conditions, distinct biological systems target specific *set-point variables*, which allow organisms to withstand certain environmental constraints (Moffett, Moffett, and Schuaf 1993). For example, most homeotherms maintain a deep-body temperature (T_{db}) set-point that approximates 37^{o}C. In addition to set-point variables, organisms generate biological oscillations, i.e. periodic changes in set-point through time (Refinetti and Menaker 1992). It is quite common for these oscillations to have a 24 hours period. In this case, we refer to them as *circadian patterns* or *circadian rhythms*.

An organism is considered to be in a dynamic state of homeostasis when each system reaches a delicate balance within the framework of other regulatory processes. Since it is impossible to simultaneously assess all regulatory processes in a time-dependent manner, empirical biologists often assume that a state of homeostasis is maintained in young adulthood. In the current study, we employ an alternative strategy focusing on specific biological systems to represent homeostasis. We illustrate our strategy by examining measurements related to two specific systems that are known to have set-points and circadian oscillations, namely T_{db} and activity counts (ActC). This methodology, however, can also be applied to measurements, such as heart rate and hormone levels, from other biological systems. The objectives of this paper are to describe how statistical methodology can be used to obtain a practical set of biological indicators useful for assessing homeostasis, and to propose potential research questions motivated by our analysis. To achieve these objectives, T_{db} and ActC were recorded from inbred mice during young adulthood through senescence, and including imminent death.

In general, sources of biological variation can be partitioned into genetic and environmental (i.e. non-genetic) components. Genetic determinants, for example, predispose homeotherms to regulate internal temperature independent of environmental conditions. Variation in temperature regulation can also be influenced by changing environmental conditions, such as occurs during fever. Moreover, natural aging and chronic disease states have been shown to alter an organism's ability to regulate T_{db} thereby creating a departure from homeostasis. We elected to study inbred mouse strains to reduce the biological variation attributed to genetic determinants. The AKR/J strain was chosen because of its relatively short lifespan.

In our analysis, we use a statistical model to evaluate the sources of variation and to quantify departures in an automated way. If an individual organism regulates a measure representative of a biological system with a set point M and a circadian oscillation described by a function of time $C(t)$, then a statistical model for observed measurements is

$$y(t) = M + C(t) + \varepsilon(t) \text{ with } C(t) = C(t+k), k = 1, 2, \ldots, \tag{1}$$

with t time in days and $\varepsilon(t)$ a stochastic process, possibly correlated in time, representing environmental variation and measurement error. If we describe the typical variation of M and $C(t)$ within a population of organisms in homeostasis, we should be able to quantify departures from homeostasis empirically using model (1). In this paper we describe how one can use the fits of models like (1) to construct indicators of homeostatic competence for different biological systems. These indicators can then be combined in many ways to create a specific index that operationally defines homeostasis. In Section 5 we describe a specific application in which we test the hypothesis that mortality risk to environmental stress is proportional to the degree of homeostatic loss.

Various techniques have been suggested to estimate circadian patterns from data; for example,

Greenhouse, Kass, and Tsay (1987) propose fitting a harmonic model, while Wang and Brown (1996) propose a method based on amplitude and phase modulated spline functions. Various approaches have been suggested for the analysis of data observed from different individuals in a population and believed to follow models similar to (1). Wang (1998), Zhang et. al. (1998) and Zhang, Lin, and Sowers (2000) present approaches based on mixed models. In this paper we use an approach, presented by Brumback and Rice (1998), that permits us to model a smooth population circadian pattern from which each individual deviates in a "smooth" way.

2 Data Exploration

(T_{db}) and (ActC) were measured in 17 mice housed in individual cages. The measurements were acquired at 30-minute intervals starting two weeks after the implantation of a radiotelemeter and ending with death. For practical reasons related to the life-time of the battery-operated radiotelemeter, measurements were usually taken only for 48 hours during weekends. In the time series plots shown in Figures 1a and 1b for mouse 4 note that as death approaches, daily mean levels of T_{db} and ActC decrease from what is typical of early adulthood. The regularity of circadian oscillations for both variables was progressively lost late in life.

The first step in our analysis was to describe the set-point and circadian oscillations in a population of healthy mice. We restricted this portion of the analysis to data obtained during consecutive weekends when the mice were relatively young adults, had recuperated from surgery, and exhibited normal behavior by visual inspection. We wanted to maximize the number of days and number of mice included, but some mice had few consecutive weekends as defined by these criteria. We

found a good balance by considering 6 consecutive weekends for which 12 out of 17 mice met the criteria. The other 5 were excluded from this stage of the data analysis. Data from the remaining mice, numbered with $i = 1, \ldots, 12$, are hereafter referred to as *control data*. Figures 1c and 1d show the control data for T_{db} and transformed ActC (The transformation $y = (\text{ActC} + 1)^{1/4}$ was selected from the Box-Cox family), plotted against time. These plots show that all mice follow a similar circadian pattern which can be accounted for by a population $C(t)$ in a model like (1). It is not clear from the plots, however, whether or not the data suggest different $C(t)$ and Ms for different mice. We now describe why for our study we consider diffrent $C(t)$ and M for each mice. The analysis of variance (ANOVA) results in Table 1 show that the population circadian pattern effect C and mouse effect M have the largest mean squares (MS). The MS of the day effect and all interactions were relatively smaller than the C and M effects and of similar size. A useful statistical model may only need to account for C and M. However, the mouse/circadian pattern interaction seems to be important from a physiological perspective. Figure 2 plots the circadian pattern/day interactions (DC), and circadian pattern/mouse interactions (MC). This type of plot is helpful in cases showing apparent "systematic" differences that are too small to substantially affect the MS. For example, if the interactions are viewed as functions of time, MC appears less steep and less noisy than DC. This provides some empirical evidence of systematic differences between the circadian patterns among mice in ActC and T_{db}. Another possible interpretation is that C and M are the only important effects and that the errors for measurements taken from the same mouse are autocorrelated. However, as mentioned, from a physiological point-of-view, it makes sense to include the (MC) interaction term. We therefore choose the following statistical model for our

data:

$$y_{ijk} = M_i + C_k + (MC)_{ik} + \varepsilon_{ijk}, i = 1, \ldots, I,\ j = 1, \ldots, J,\ k = 1, \ldots, K, \qquad (2)$$

with $I = 12$ the number of mice, $J = 12$ the number of days, $K = 48$ the number of observations taken each day, and ε_{ijk} considered random errors. To explore the correlation structure of the errors we examine the residuals obtained from fitting model (2) using least squares (LS). In Figures 2e and 2f we see sample correlations corr$(\widehat{\varepsilon}_{ijk}, \widehat{\varepsilon}_{ijl})$ plotted against $|k - l|$. For T_{db} a positive correlation for smaller values of $|k - l|$ is evident. This may be attributed to the time it takes heat to dissipate. For simplicity, we model the errors in model (2) as an AR(1) for T_{db} . With ActC we do not expect a positive correlation, since over-activity is likely to be associated with fatigue and followed soon after by relative inactivity. The correlation plot for ActC shows less than conclusive evidence of a negative correlation at lag 1 (30 minutes later). Thus we assume the errors are independent and identically distributed (IID). Histograms and quantile-quantile plots (not shown) support our approximating the distribution of errors with a Gaussian model.

3 Statistical Methods

Most biologists agree that mice are homeostatically robust in young adulthood. Thus, we used the control data to characterize homeostatic status. The parameter estimates, obtained from fitting model (2) to the control data, and their statistical variation can be considered "typical" for mice in robust homeostasis. We may consider the fit as a "template" or "measuring rod". However, model (2) contains two terms, C_k and $(MC)_{ik}$, that may be considered to be samples from smooth functions of time. Model (2) does not take advantage of this *a priori* belief. The techniques

described in Greenhouse et. al. (1987) and Wang and Brown (1996) take this into account. The latter is used to produce the estimates shown as unbroken lines in Figures 3a and 3b and is further discussed in Section 5. Wang (1998) and Zhang et. al. (1998) and Zhang et. al. (2000) present approaches for models with smooth C_ks and interaction terms $(MC)_{ik}$ based on mixed models. The method presented in Brumback and Rice (1998) permits consideration of C_k and $(MC)_{ik}$ as twice-differentiable smooth functions of time. This is a natural assumption for our example since it can be interpreted as a smooth circadian pattern C_k for the population from which each healthy mouse deviates in a smooth way. Brumback and Rice's approach requires the errors to be independent. For $\mathrm{T_{db}}$ we demonstrate how their results can be easily extended to the correlated data case.

If we consider one day of data for one mouse, then model (2) reduces to the "non-parametric" regression problem $y_k = s(t_k) + \varepsilon_k$, where we have observations at times $t_k = k/K, k = 1, \ldots, K$ with $s(t)$ a smooth curve and $\boldsymbol{\varepsilon} = (\varepsilon_1, \ldots, \varepsilon_K)' \sim$ multivariate normal$[MVN](0, \sigma^2\boldsymbol{\Sigma}_1)$. For ActC, $\boldsymbol{\Sigma}_1 = \mathbf{I}_K$ and for $\mathrm{T_{db}}$, $\{\boldsymbol{\Sigma}_1\}_{kl} = \rho^{|k-l|}$. We assume $s(t)$ is in the Sobolev space of functions with support $[0, 1]$ whose second derivatives are square integrable. Under this assumption, the function $\widehat{s}(t)$ minimizing the weighted residual sum of squares plus a roughness penalty,

$$(\mathbf{y} - \mathbf{s})'\boldsymbol{\Sigma}_1^{-1}(\mathbf{y} - \mathbf{s}) + \lambda \int_0^1 \{s''(t)\}^2 \, dt, \tag{3}$$

with $\mathbf{y} = (y_1, \ldots, y_K)'$ and $\mathbf{s} = \{s(t_1), \ldots, s(t_K)\}'$ is a natural cubic spline. Because we are interested in circadian patterns, it is natural to assume that $s(t)$ is a periodic function. As pointed out by Wahba (1990), for even K, $s(t)$ is well approximated by

$$s(t) = a_0 + \sum_{\nu=1}^{K/2-1} a_\nu \sqrt{2}\cos(2\pi\nu t) + \sum_{\nu=1}^{K/2-1} b_\nu \sqrt{2}\sin(2\pi\nu t) + a_{K/2}\cos(\pi K t). \tag{4}$$

Furthermore, because the data are equally spaced in time, integration and simple trigonometric properties yield the following simple representation of the penalty term:

$$\lambda \int_0^1 \{s''(t)\}^2 \, dt = \lambda \left\{ \sum_{\nu=1}^{K/2-1} (a_\nu + b_\nu)^2 (2\pi\nu)^4 + \frac{1}{2} a_{K/2} (\pi n)^4 \right\}. \tag{5}$$

Notice that the harmonic models used by Greenhouse, et. al. (1987) are equivalent to (4) with the constraint $a_\nu = 0, b_\nu = 0$ for $\nu > 7$.

It can be shown that the fitted smoothing spline evaluated at the design points, $\widehat{\mathbf{s}} = \{\widehat{s}(t_1), \ldots, \widehat{s}(t_K)\}'$, equals the BLUP solution to a mixed effects model

$$\mathbf{y} = a_0 \mathbf{1} + \mathbf{Z}\mathbf{u} + \boldsymbol{\varepsilon} \tag{6}$$

where a_0 is a fixed effect and $\mathbf{u}$ is a vector of random effects corresponding to a known design matrix $\mathbf{Z}$. Results (4) and (5) permit us to write $\mathbf{Z}$ in a form that provides not only simplicity in the presentation of the results but computational advantages. The mixed-effect model is specified by: $\mathbf{Z}$ a matrix with entries $\{\mathbf{Z}\}_{2\nu-1,k} = \sqrt{2}/(2\pi\nu)^2 \cos(2\pi\nu\, k/K)$, $\{\mathbf{Z}\}_{2\nu,k} = \sqrt{2}/(2\pi\nu)^2 \sin(2\pi\nu\, k/K), \nu = 1, \ldots, K/2-1, \mathbf{Z}_{K/2-1,k} = 1/(\sqrt{2}\pi K)^2 \cos(\pi k)$; $\mathbf{u} \sim MVN\{\mathbf{0}, (\sigma^2/\lambda)\mathbf{I}_K\}$; and $\boldsymbol{\varepsilon} \sim MVN(\mathbf{0}, \sigma^2\boldsymbol{\Sigma}_1)$ independent of $\mathbf{u}$. The calculations for obtaining $\mathbf{Z}$ are the same as those appearing in Brumback and Rice (1998)

The BLUP solutions $\widehat{a}_0$ and $\widehat{\mathbf{u}}$ are defined, for example, in Robinson (1991) and under our assumption that the random components are normally distributed, they are the maximum likelihood (ML) estimates of a_0 and $\mathbf{u}$. In the penalized regression setting, cross-validation is the most popular method for choosing λ. However, in the context of mixed-effects models, it is natural to use restricted ML (REML). See Wang (1998) for insightful and more extensive discussions of these issues.

For observations from a nested sample of curves, the basic idea is to consider an additive model in which the observations are outcomes of

$$y_{ijk} = M_{ij} + s_p(t_k) + s_i(t_k) + \varepsilon_{ijk}, i = 1, \ldots, I, j = 1, \ldots, J, k = 1, \ldots, K \tag{7}$$

with $s_p(t)$ and $s_i(t)$ representing smooth population circadian patterns and smooth individual mouse departures from the population circadian pattern, respectively. Notice that we are fitting a different mean for each mouse i and day j. However, we are not considering M_{ij} to be a day/mouse interaction because we expect $M_{ij} = M_i$ for all j. Later in this section we describe why we fit a model where M_{ij} is permitted to vary with j.

Brumback and Rice (1998), note that the appropriate analog of (3) is not *a-priori* clear and, borrow from the "long history of the ANOVA with mixed-effects models" to suggest an appropriate criterion. In this section we use a simple extension of their result to arrive at a useful solution to our problem.

To build a mixed-effect model from (7) we define

$$\mathbf{y} = \mathbf{X}\boldsymbol{\beta} + \mathbf{Z}_1\mathbf{u}_1 + \mathbf{Z}_2\mathbf{u}_2 + \boldsymbol{\varepsilon} \tag{8}$$

with $\mathbf{y} = (y_{111}, \ldots, y_{IJK})'$; $\mathbf{X} = \mathbf{I}_{IJ} \otimes \mathbf{1}_K$, $\mathbf{Z}_1 = \mathbf{I}_I \otimes \mathbf{1}_J \otimes \mathbf{Z}$, and $\mathbf{Z}_2 = \mathbf{I}_{IJ} \otimes \mathbf{Z}$; $\boldsymbol{\beta} = (M_{1,1}, \ldots, M_{I,J})'$; $\mathbf{u}_1 \sim MVN\{\mathbf{0}, (\sigma^2/\lambda_1)\mathbf{I}_{K-1}\}$; $\mathbf{u}_2 \sim MVN\{\mathbf{0}, (\sigma^2/\lambda_2)\mathbf{I}_{I(K-1)}\}$; $\boldsymbol{\varepsilon} \sim MVN(\mathbf{0}, \sigma^2\boldsymbol{\Sigma})$ with $\boldsymbol{\Sigma} = \mathbf{I}_{IJ} \otimes \boldsymbol{\Sigma}_1$; with $\mathbf{u}_1$, independent of $\mathbf{u}_2$ independent of $\boldsymbol{\varepsilon}$.

Using Theorem 2 in Brumback and Rice (1998), we can show that the BLUP solution summarized by $\widehat{\mathbf{y}} = \mathbf{X}\widehat{\boldsymbol{\beta}} + \mathbf{Z}_1\widehat{\mathbf{u}}_1 + \mathbf{Z}_2\widehat{\mathbf{u}}_2$ to the nested mixed-effect model (8) is equivalent to fitted curves given by a corresponding penalized weighted regression. See the Appendix for details.

4 Assessing homeostasis

In this section, we explain how the estimates $\widehat{M}_{ij}$, $\widehat{s}_p(t)$, and $\widehat{s}_i(t)$, $i = 1, \ldots, I$, $j = 1, \ldots, J$ obtained from fitting (7) to the control data are used to assess homeostasis. In Figures 3c and 3d we present the resulting $\widehat{s}_p(t) + \widehat{s}_i(t), i = 1, \ldots, I$ for ActC and T_{db} respectively, with $\widehat{s}_p(t)$ represented as a bolded line. The model fits the data well. Notice in particular that $\widehat{\rho}^{|k-l|}$ is close to the sample correlations in Figure 2f.

Assume that for a given mice i used to obtain the control data we also have data from one other day; we want to determine if this mouse is in homeostasis on that day. Let $y^*_{ik}, k = 1, \ldots, K$ denote these data. In Section 2 we noted how, as death approaches, daily mean levels decrease. The control data can be used to obtain the estimates $\widehat{M}_{ij}$ and can be used as a "measuring rod" to quantify how far $K^{-1}\sum_{k=1}^{K} y^*_{ik}$ is from the homeostatic set-point. We also noted that the regularity of circadian oscillations are progressively lost late in life. We now describe how to quantify a loss of circadian pattern regulation. We model the new data as

$$y^*_{ik} = M^*_i + s^*_i(t_k) + \varepsilon^*_{ik}, k = 1, \ldots, K \tag{9}$$

with $s^*_i(t_k)$, and ε^*_{ik} satisfying the assumptions of model (7). Because we only have one day of data the population and individual departure curves are not identifiable; thus we combine the two components into $s^*_i(t_k)$. For the procedure we propose below, we do not need to separate the two components. We therefore represent model (9) as (6) with $\mathbf{u} \sim MVN\{\mathbf{0}, (\sigma^2/\lambda_1 + \sigma^2/\lambda_2)\mathbf{I}_K\}$; and $\boldsymbol{\varepsilon} \sim MVN(\mathbf{0}, \sigma^2\mathbf{I}_K)$ for ActC and $\boldsymbol{\varepsilon} \sim MVN\{\mathbf{0}, \boldsymbol{\Sigma}_1(\rho)\}$ for T_{db} . Using the variance component estimates obtained for the control data, the estimates $\widehat{M^*_i}$ and $\widehat{s^*_i}(t)$ are then easily obtained in the form of BLUP solutions. If a mouse is still homeostatically competent then we expect $s^*_i(t) =$

$s_p(t) + s_i(t)$. We can think of this as the "null hypothesis" of homeostatic competence as related to circadian pattern regulation. As a test statistic, we suggest using $d_i = \int_0^1 \{\widehat{s}_i^*(t) - \widehat{s}_p(t) - \widehat{s}_i(t)\}^2\, dt$. We construct a null distribution for d_i from the control data using a bootstrap procedure. For the correlated data, we use the procedure described in Efron and Tibshirani (1993, p. 92–99).

Because five mice were not used to establish normal performance, d_i cannot be constructed for them. However, because we believe there is a strong genetic component to circadian regulation it makes sense to suggest that each individual will mimic the population average. In Figures 3e and 3f, we compare the estimated population averages $\widehat{s}_p$ for ActC and T_{db} respectively for AKR and B6 (an alternative inbred strain). We surround the population averages with bootstrap estimates obtained using the sampling mechanism described by Brumback and Rice (1998). These figures suggest that, for these five mice, a useful test statistic compares $s_i^*(t)$ to the population curve $s_p(t)$ -namely $e_i = \int_0^1 \{\widehat{s}_i^*(t) - \widehat{s}_p(t)\}^2\, dt$. Again, we use a bootstrap procedure to obtain a null distribution for e_i. Notice in Figures 4c and 4d that the shown quantiles for the null distribution for the d_is and e_is are quite close. We may simplify the diagnostic tool by only considering e_i. In Figure 4f we see that for mouse 3 the values of d_i and e_i are close and would yield equivalent results in practice. Similarly, we compare $\widehat{M}_h^*$ of mice not used in the control data to the $\widehat{M}_{ij}$s.

We fit (9) to the data from all 17 mice. Figures 4a and 4b show $\widehat{M}^*$ against weeks left until death for ActC and T_{db} , respectively, for three mice: one of the mice (3) used for the control data and two of the mice (14 and 15) not used. We compare $\widehat{M}^*$ to the estimates $\widehat{M}_{ij}$ obtained by dropping the $M_i = M_{ij}$ constraint as opposed to the $\widehat{M}_i$ obtained by keeping it, since $\widehat{M}^*$ and $\widehat{M}_{ij}$ both represent day averages and thus provide a fair comparison. The horizontal lines shown represent the 0.05 and 0.95 quantiles of the estimates, $\widehat{M}_{ij}$. One possible way of establishing

normal homeostasis for the adult is to specify that $\widehat{M}^*$ is within appropriately chosen quantiles of the $\widehat{M}_{ij}$s.

Figures 4c and 4d show the distances d_i for mouse 3 and e_i for 14 and 15 plotted against weeks left until death. The circadian pattern function starts to deviate from the "control" curve as death approaches. This measure may also be considered an indicator of homeostatic competence. Moreover, these deviations, expressed as distances, should also be within appropriately chosen quantiles of the null distribution.

5 Discussion and Extensions

Based on regulatory processes common to many mammalian species, we have described an automated statistical procedure useful for assembling a set of indicators of homeostasis. Specifically, four different indicators related to the set-point and circadian oscillation for T_{db} and ActC have been defined and are shown in Figures 4a-4d. These figures include only 3 mice. The other 14 mice had results similar to mice 3 and 15 but were not included for graphical clarity.

In Figure 4a, mouse 14 appears "over-active" from the outset whereas its T_{db} (Figure 4b) appears "normal" for about the first seven weeks of observations. During this period, one might infer that temperature regulation was similar to other AKR mice in the face of over-activity. Also in Figure 4b, mouse 15 is outside the dotted lines when it still has 10 weeks to live, yet it is mostly within the dotted line limits in Figures 4a,4b, and 4d for those same weeks. Collectively, the data suggest that when death is imminent, all indicators leave the defined normal boundaries, suggesting a loss of all regulatory processes. Thus, it is useful to consider multiple indicators if possible. Depend-

ing on specific goals, these can be combined in a number of ways to assess/evaluate homeostatic capacity.

For example, the goal of a current project is to test the hypothesis that mortality risk to environmental stress increases with age and is proportional to the degree of homeostatic loss. We propose testing this hypothesis by comparing responses of mice at different levels of homeostatic competence when exposed to specific levels of an air pollutant. We may need to decide on a specific day whether a mouse is in homeostasis or not. If we look at days when the mice (all 17) had about 14 weeks to live (probably still healthy), then the indicator presented in Figure 4a suggests loss of homeostatic competence 38% of the time. The indicators presented in Figures 4b, 4c, and 4d suggest loss of homeostatic competence 54%, 46%, and 38% of the time, respectively. In our current projects we are using an index that requires all 4 measures to be outside their limits. This index makes the aformentioned incorrect assessment 8% of the time. This index has given us a useful operational definition of homeostasis.

The analysis presented in this paper motivates many potential research questions. First, There may be other useful ways to combine the indicators to form indices of homeostatic competence. For example, we may gain power by considering the values of indicators for multiple consecutive days. The use of ROC curves for this is the subject of current research. Second, notice that the indicators are qualitative measures. We may be able to compare the health of two populations by comparing some function of these indicators. In Figure 4e we show distance measurements obtained for T_{db} plotted against age. The curve shown with a solid line is a loess smooth of this data. We superimpose the empirical survival curve of the 17 AKR mice. Further research will explore survival regression models with these indicators as possible covariates. We feel that n must be

large for such a regression to yield interesting results and we are currently collecting more data. Third, we may be able to make the comparison of set-points more precise by considering the M term to be a random effect. As more data become available, we may find a useful approximation that would yield an appropriate assumption for the distribution of M modeled as a random effect. Finally, it would be useful to extend our methods to cases where the errors are not normally distributed.

A number of approaches might be taken to assess homeostasis by estimating circadian curves. We applied the procedure of Wang and Brown (1996) and found that $\widehat{\alpha}$, the circadian pattern amplitude estimate, provided a useful diagnostic since it becomes close to 0 as death approaches. The main reason for using the approach of Brumback and Rice (1998) is its greater flexibility, as demonstrated by Figures 3a and 3b. Wang and Brown's model forces the shape of the circadian pattern to remain the same for all mice. The code for all the programs used in this paper can be found at `http://biosun01.biostat.jhsph.edu/~ririzarr/software.html`.

Appendix

Calculations for EM algorithm. The penalized weighted regression fits M_{ij} and the collection of smooth curves $s_p(t)$ and $s_i(t), i = 1, \ldots, I$ by minimizing :

$$\sum_{j=1}^{J}\sum_{i=1}^{I}(\mathbf{y}_{ij} - M_{ij} - \mathbf{s}_p - \mathbf{s}_i)'\mathbf{\Sigma}_1^{-1}(\mathbf{y}_{ij} - M_{ij} - \mathbf{s}_p - \mathbf{s}_i) + \lambda_1 \int_0^1 \{s_p''(t)\}^2\, dt + \lambda_2 \sum_{i=1}^{I}\int_0^1 \{s_i''(t)\}^2\, dt$$

with $\mathbf{y}_{ij} = (y_{ij1}, \ldots, y_{ijK})'$, $\mathbf{s}_p = \{s_p(t_1), \ldots, s_p(t_K)\}'$ and $\mathbf{s}_i = \{s_i(t_1), \ldots, s_i(t_K)\}', i = 1, \ldots, I$. To obtain estimates of our curves, the variance components must be estimated. Writing $\mathbf{\Sigma}$ as a function of ρ, we have $\text{var}(\mathbf{y}) = \mathbf{V}(\boldsymbol{\theta}) = (\sigma^2/\lambda_1)\mathbf{Z}_1\mathbf{Z}_1' + (\sigma^2/\lambda_2)\mathbf{Z}_2\mathbf{Z}_2' + \sigma^2\mathbf{\Sigma}(\rho)$ where $\boldsymbol{\theta} = (\sigma^2, \rho, \sigma^2/\lambda_1, \sigma^2/\lambda_2)$. We use REML and the EM algorithm of Dempster, Laird, and Rubin (1977).

For the ActC we assume IID errors, so ρ is 0 and we can proceed exactly as in Brumback and Rice (1998). To perform REML, we select a matrix $\mathbf{K}'$ with the highest rank possible such that $\mathbf{K}'\mathbf{X} = \mathbf{0}$ and $\mathbf{K}'$ has full row rank, then consider the transformation $\mathbf{K}'\mathbf{y}$ which does not depend on $\boldsymbol{\beta}$. We then define the complete data set $\mathbf{K}'\mathbf{y}, \mathbf{u}_1, \mathbf{u}_2$, and $\boldsymbol{\varepsilon}$ and use the complete data sufficient statistics for the variance components to define the E-step and M-step.

For T_{db} , we assume the errors are a Gaussian AR(1) process and must alter the procedure to permit estimation ρ. For computational convenience we will assume that measurements form different days are independent, although measurements for consecutive days for the same mice should be correlated (using this assumption produces negligible differences in the results but great improvements in computation time). If we define $t_{0m} = \boldsymbol{\varepsilon}'\mathbf{A}_m\boldsymbol{\varepsilon}, m = 0, \ldots, 3$ (the $\mathbf{A}_m$'s defined in the Appendix), $t_1 = \mathbf{u}_1'\mathbf{u}_1$, and $t_2 = \mathbf{u}_2'\mathbf{u}_2$, then letting $\mathbf{V}_K$ denote var($\mathbf{K}'y$), E-step expectations of the ts conditional on the actual data $\mathbf{K}'\mathbf{y}$ are computed as

$$\begin{aligned} \mathrm{E}(t_{0m}|\mathbf{K}'\mathbf{y}) &= \mathrm{tr}(\sigma^2\mathbf{A}_m\boldsymbol{\Sigma} - \sigma^4\mathbf{A}_m\boldsymbol{\Sigma}\mathbf{K}\mathbf{V}_K^{-1}\mathbf{K}'\boldsymbol{\Sigma}) + \sigma^4\mathbf{y}'\mathbf{K}\mathbf{V}_K^{-1}\mathbf{K}'\boldsymbol{\Sigma}\mathbf{A}_m\boldsymbol{\Sigma}\mathbf{K}\mathbf{V}_K^{-1}\mathbf{K}'\mathbf{y}, m = 0, \ldots, 3 \qquad (10) \\ \mathrm{E}(t_m|\mathbf{K}'\mathbf{y}) &= \mathrm{tr}\left\{(\sigma^2/\lambda_m)\mathbf{I}_n - (\sigma^2/\lambda_m)^2\mathbf{Z}_m\mathbf{K}\mathbf{V}_K^{-1}\mathbf{K}'\mathbf{Z}_m'\right\} + (\sigma^2/\lambda_m)^2\mathbf{y}'\mathbf{K}\mathbf{V}_K^{-1}\mathbf{K}'\mathbf{Z}_m\mathbf{Z}_m'\mathbf{K}\mathbf{V}_K^{-1}\mathbf{K}'\mathbf{y}, m = 1, 2. \end{aligned}$$

The M-step, ML solution for ρ is obtained by solving the cubic equation

$$(n - IJ)(t_{02} - t_{03})\rho^3 - (n - 2IJ)t_{01}\rho^2 - n\{t_{02} - t_{03} - IJ/n(t_{00} + t_{03})\}\rho + nt_{01} = 0. \quad (11)$$

For the remaining parameters, the ML solutions are $\widehat{\sigma}^2 = n^{-1}\left\{(1 - \widehat{\rho}^2)^{-1}(t_{00} - 2\widehat{\rho}t_{01} + \widehat{\rho}^2t_{02}) + t_{03}\right\}$, $(\widehat{\sigma}^2/\widehat{\lambda}_1) = t_1/(K - 1)$, and $(\widehat{\sigma}^2/\widehat{\lambda}_2) = t_2/(IK - I)$.

The part of the log-likelihood related to parameters ρ and σ^2 is $-2l(\sigma^2, \rho|\boldsymbol{\varepsilon}) \sim n\log\sigma^2 + IJ(K - 1)\log(1 - \rho^2) + \sigma^{-2}S$ with $S = \sum_{i=1}^{I}\sum_{j=1}^{J}\{(1 - \rho^2)^{-1}\sum_{k=2}^{K}(\varepsilon_{i,j,k} - \rho\varepsilon_{i,j,k-1})^2 + \varepsilon_{i,j,1}^2\}$. Taking derivatives, we find that the above is minimized when $\sigma^2 = n^{-1}S$, which is equivalent to the expression given in Section 3, and by the ρ that minimizes (11). By defining $\mathbf{A}_1$ as a diagonal

matrix with 1s in each row except the first, $\mathbf{A}_2$ as the matrix having $\{\mathbf{A}_2\}_{ij} = 1$ for $j - i = 1$ and 0 otherwise, $\mathbf{A}_3$ the diagonal matrix with 1s in each row except the last and $\mathbf{A}_4$ a 0 matrix except for $\{\mathbf{A}_4\}_{11} = 1$, we can then compute $\mathrm{E}\{-2l(\sigma^2, \rho|\boldsymbol{\varepsilon})|\mathbf{K}'\mathbf{y}\}$ to obtain (10).

Computational issues. Notice that for our example the matrix $\mathbf{V}_K$ is of dimension 6768×6768. This makes the task of inverting $\mathbf{V}_K$ difficult if not impossible for most desktop computers. However, one can take advantage of the model's structure to reduce the problem to the inversion of a few matrices of size 47×47, reducing computation time to fractions of a second on a Sun Ultra 10. Notice that $\boldsymbol{\Sigma}$ is block diagonal and that we can choose $\mathbf{K}$ to be block diagonal. By not imposing the $M_i = M_{ij}$ constraint, constructing $\mathbf{K}$ is simple because $\mathbf{X}'\mathbf{X} = \mathbf{I}_{IJ} \otimes \mathbf{1}_K \mathbf{1}'_K$. We denote $\boldsymbol{\Sigma}_1$ and $\mathbf{K}_1$ as the respective diagonal entries. Using exercises 2.7 and 2.8 from Rao (1973, p. 33) we can show that $V_K^{-1} = \mathbf{I}_{IJ} \otimes (\mathbf{O} - \mathbf{P}) + \mathbf{I}_I \otimes \mathbf{1}_J \mathbf{1}'_J \otimes (\mathbf{P} - \mathbf{Q}) + \mathbf{1}_{IJ} \mathbf{1}'_{IJ} \otimes \mathbf{Q}$ with $\mathbf{P} = -\mathbf{C}(\mathbf{L} + 2\mathbf{G} + \mathbf{F})\mathbf{C}'$, $\mathbf{O} = \mathbf{A} - \mathbf{P}$, $\mathbf{Q} = \mathbf{C}(\mathbf{H} + 2\mathbf{G} + \mathbf{F})\mathbf{C}'$ where $\mathbf{A} = (\mathbf{K}_1'\boldsymbol{\Sigma}_1\mathbf{K}_1)^{-1}$, $\mathbf{B} = \mathbf{K}_1'\mathbf{Z}$, $\mathbf{C} = \mathbf{A}\mathbf{B}$, $\mathbf{D} = J\mathbf{B}'\mathbf{C}$, $\mathbf{E} = (\mathbf{D} + (\lambda_2/\sigma^2)\mathbf{I})^{-1}$, $\mathbf{F} = [I\mathbf{D} + (\lambda_1/\sigma^2)\mathbf{I} - I\mathbf{D}\{\mathbf{D} + (\lambda_2/\sigma^2)\mathbf{I}]^{-1}\mathbf{D}\}^{-1}$, $\mathbf{G} = -\mathbf{E}\mathbf{D}\mathbf{F}$, $\mathbf{H} = -\mathbf{G}\mathbf{D}\mathbf{E}$, $\mathbf{L} = \mathbf{D} + \mathbf{H}$. Notice that here all that is needed are a few inversion of $(K-1) \times (K-1)$ matrices.

References

Brumback, B. and Rice, J. (1998). Smoothing spline models for the analysis of nested and crossed samples of curves. *Journal of the American Statistical Association* **93,** 961–976.

Dempster, A. P., Laird, N. M., and Rubin, D. B. (1977). Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society, Series B, Methodological* **39,** 1–22.

Efron, B. and Tibshirani, R. (1993). *An Introduction to the Bootstrap*. New York: Chapman & Hall.

Greenhouse, J. B., Kass, R. E., and Tsay, R. S. (1987). Fitting nonlinear models with ARMA errors to biological rhythm data. *Statistics in Medicine* **6,** 167–183.

Moffett, D. F., Moffett, S. B., and Schuaf, C.L. (1993). *Human Physiology Foundations and Frontiers*. St. Louis: Moshby.

Rao, C.R. (1973). *Linear Statistical Inference and Its Applications*. New York: John Wiley & Sons.

Robinson, G. K. (1991). That BLUP is a good thing: The estimation of random effects. *Statistical Science* **6,** 15–32.

Refinetti, R. and Menaker, M. (1992). The circadian rhythm of body temperature. *Physiology & Behavior* **51,** 135–140.

Wang, Y. and Brown, M. B. (1996). A flexible model for human circadian rhythms. *Biometrics* **52,** 588–596.

Wang, Y. (1998). Mixed Effects Smoothing Spline Analysis of Variance. *Journal of the Royal Statistical Society, Series B* **60,** 159–174,

Wahba, G. (1990). *Spline Models for Observational Data*. Philadelphia: Society for Industrial and Applied Mathematics.

Zhang, D. and Lin, X. Raz, J., and Sowers, M. (1998). Semiparametric stochastic mixed models for longitudinal data. *Journal of the American Statistical Association* **93,** 710–719.

Zhang, D. and Lin, X. and Sowers, M. (2000). Semiparametric stochastic mixed models for longitudinal data, *Biometrics* **56,** 31–39.

Table 1: ANOVA tables for ActC and T_{db} . Time of the day (denoted the circadian pattern effect), mouse, and day are considered main effects.

Figures:

1. a) T_{db} measurements time series plots for each weekend for mouse 4. b) Activity count time series plots for each weekend for mouse 4. The vertical lines separate the weekends. c) T_{db} control data measurements. Different symbols represent different mice and the line is the average over mice/day for each time. d) Transformed ActC control data measurements. Different symbols represent different mice and the line is the average over mice/day for each time.

2. a) Mouse/Circadian pattern interaction for ActC with different line types for the different mice. The square root of the mean square and a roughness measure $D = \sum(x_i - x_{i-1})^2 / \sum(x_i - \bar{x})^2$ are shown in the titles. The bold line is circadian pattern effect. b) Day/Circadian pattern interaction for activity counts with different line types for different days. The square root of the mean square and a roughness measure $D \sum(x_i - x_{i-1})^2 / \sum(x_i - \bar{x})^2$ are shown in the titles. The bold line is the circadian pattern effect. c) As a) for T_{db} . d) As b) for T_{db} . e) Sample autocorrelation plots for activity counts. f) Sample autocorrelation plots for T_{db} measurements with the estimate $\widehat{\rho}^{\text{lag}}$ shown.

3. a) Dotted lines represent activity circadian pattern estimate for mouse 14 using Rice and Brumback (1998). Solid line is estimate obtained using Wang and Brown (1996). b) As a) for mouse 9. c) ActC circadian pattern estimates obtained for 12 mice in control data. Bold line is the population curve. d) As c) but for T_{db} . e) Circadian pattern estimates for two

strains, ActC for AKR and B6 strains with 25 bootstrap curves. f) As e) for T_{db} .

4. a) ActC mean levels for mice 3, 14, and 15, with horizontal lines representing 0.05 and 0.95 quantiles of control data estimates. Bold line is population average b) As a) for T_{db} . c) Distance plots for ActC for mice 3, 14, and 15. Horizontal lines represent 0.05 and 0.95 quantiles of null distribution for these distances, solid lines are for d_i dotted are for e_i. d) as c) for T_{db} . e) Value of body temperature circadian indicator plotted against mouse age in days. The solid line curve is a loess smooth. The dotted line is the empirical survival distribution of the 17 mice. e) Comparison of d_i, dotted lines, and e_i, solid lines for mouse 3.

Activity Counts			Body Temperature		
Effect	df	MS	Effect	df	MS
Circadian Pattern	47	32	Circadian Pattern	47	74
Mouse	11	16	Mouse	11	16.6
Mouse/Day	121	0.83	Day	11	0.63
Mouse/Circadian Pattern	517	0.75	Mouse/Day	121	0.58
Day	11	0.5	Mouse/Circadian Pattern	517	0.46
Mouse/Day/Circadian pattern	5687	0.41	Day/Circadian Pattern	517	0.15
Day/Circadian Pattern	517	0.38	Mouse/Day/Circadian pattern	5687	0.130
Total	6912	0.69	Total	6912	0.7

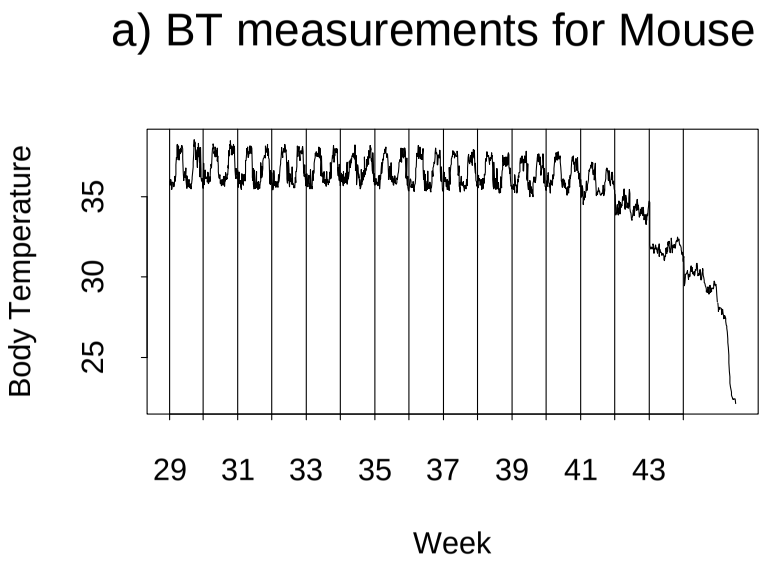
a) BT measurements for Mouse 4
Body Temperature
25
30
35
29
31
33
35
37
39
41
43
Week

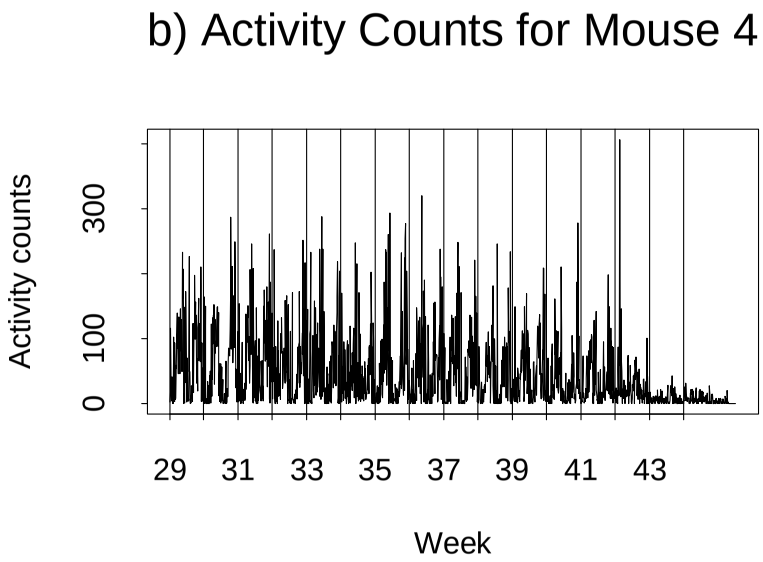
b) Activity Counts for Mouse 4
Activity counts
0
100
300
29
31
33
35
37
39
41
43
Week

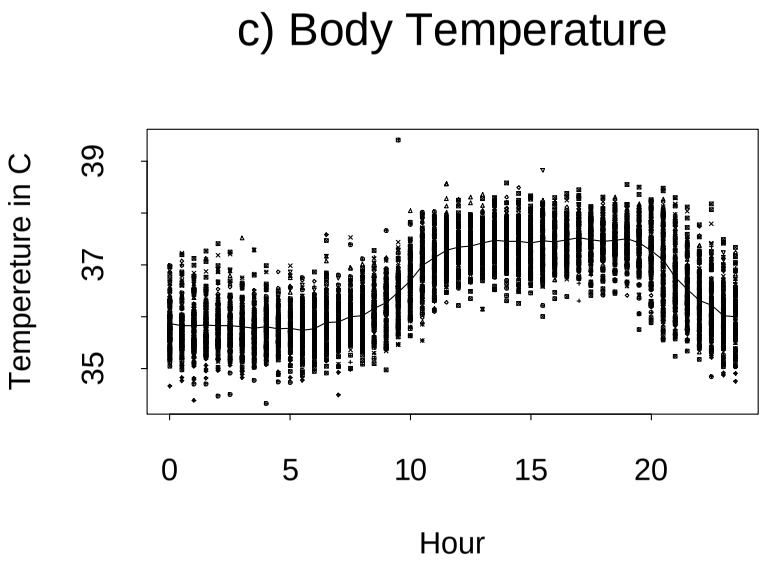
c) Body Temperature
Tempereture in C
35
37
39
0
5
10
15
20
Hour

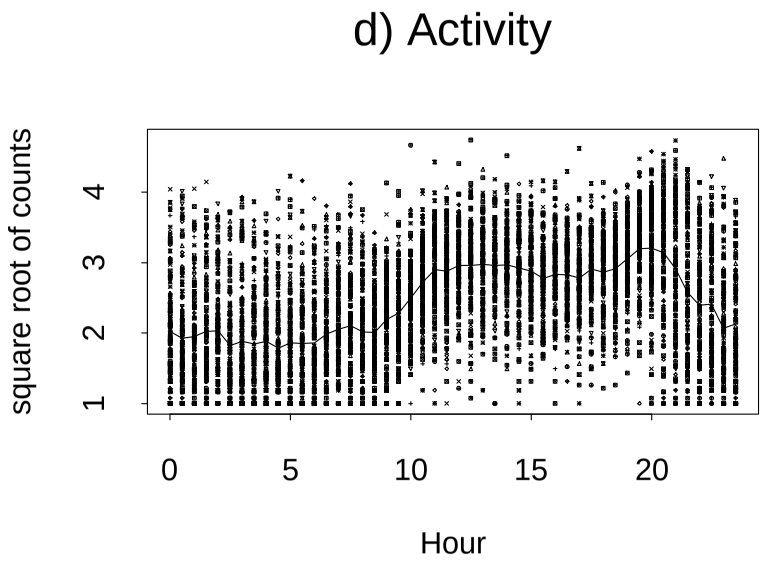
d) Activity
square root of counts
1
2
3
4
0
5
10
15
20
Hour

a) Activity Mf: MS=0.75 D=0.0039

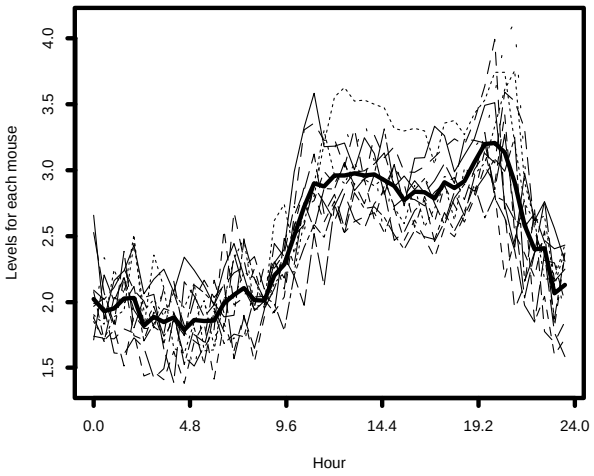

b) Activity Df: MS=0.38 D=0.0044

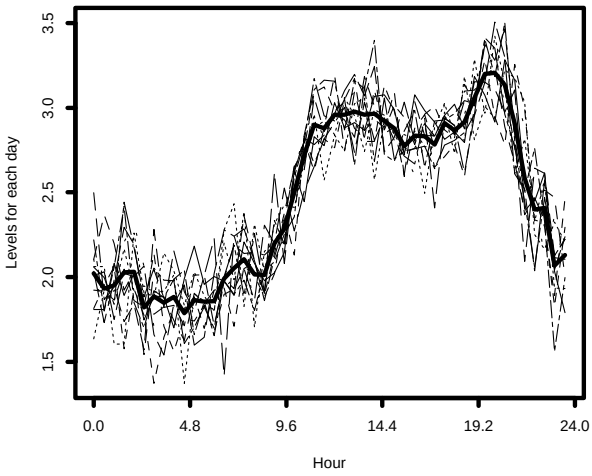

c) BT Mf: MS=0.46 D=0.0016

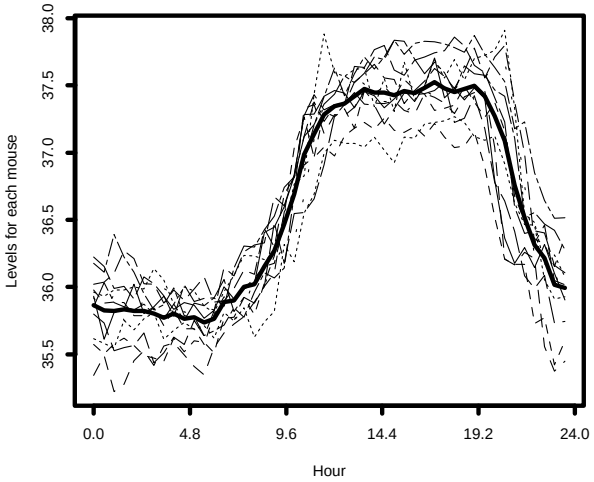

d) BT Df: MS=0.15 D=0.0049

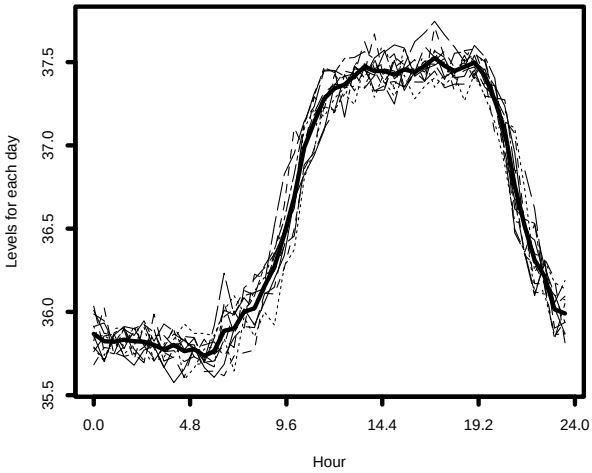

e) Sample correlation for Activity

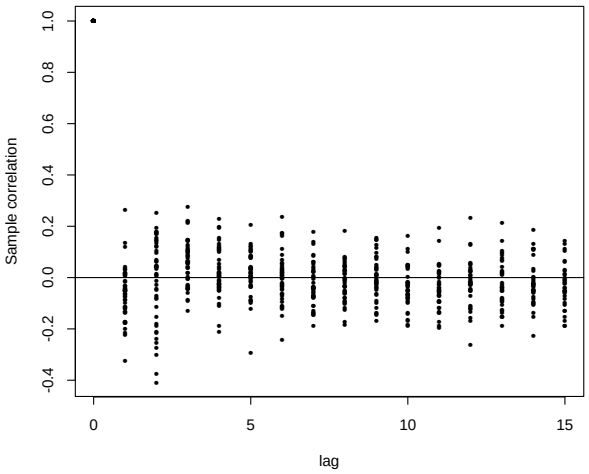

f) Sample correlation for BT

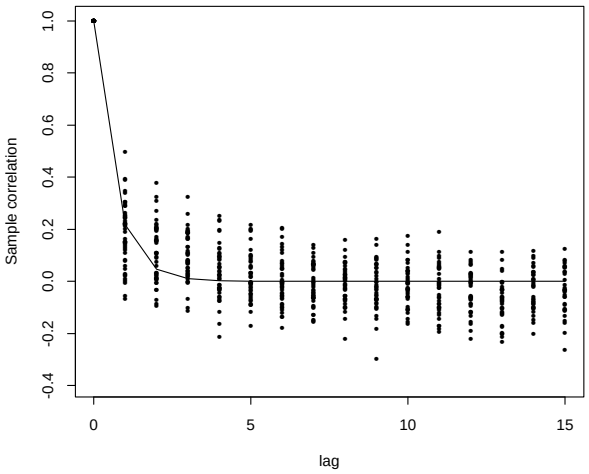

a) Coparison, Mouse 6

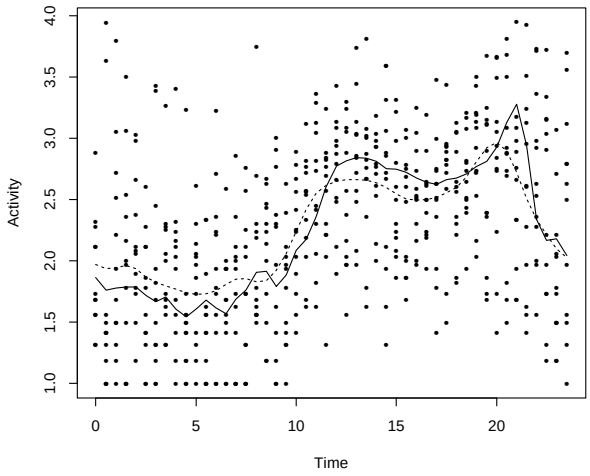

b) Coparison, Mouse 11

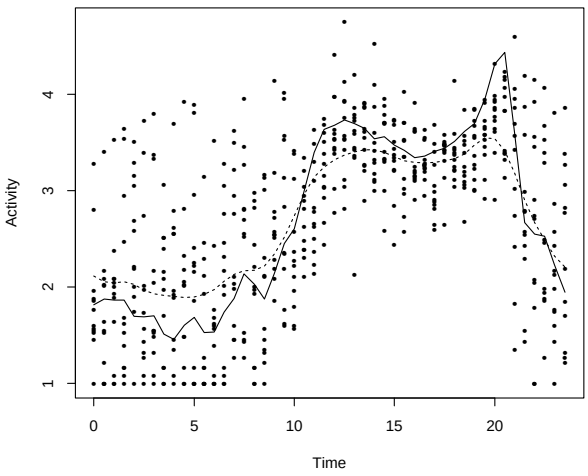

c) Activity for each AKR

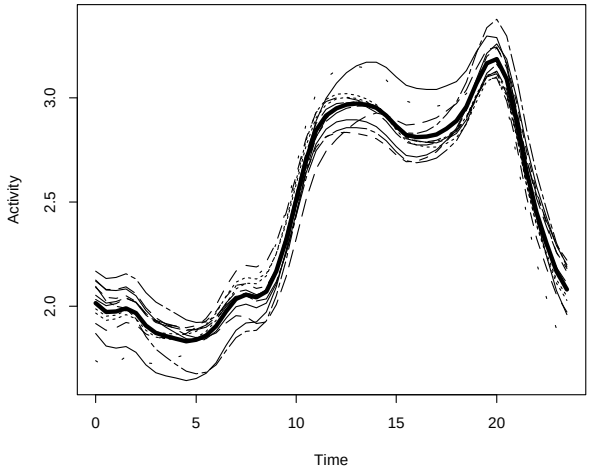

d) BT for each AKR

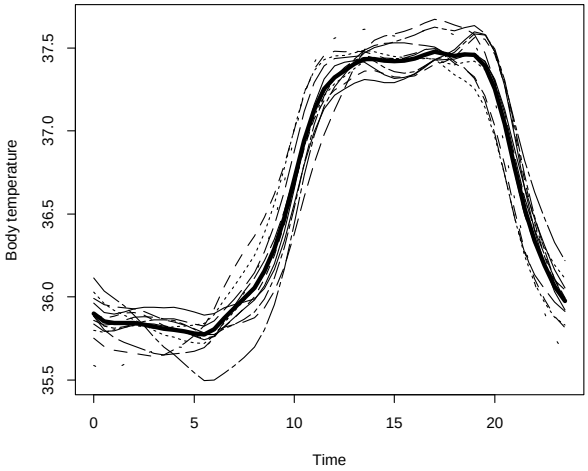

e) Activity: Strain comparison

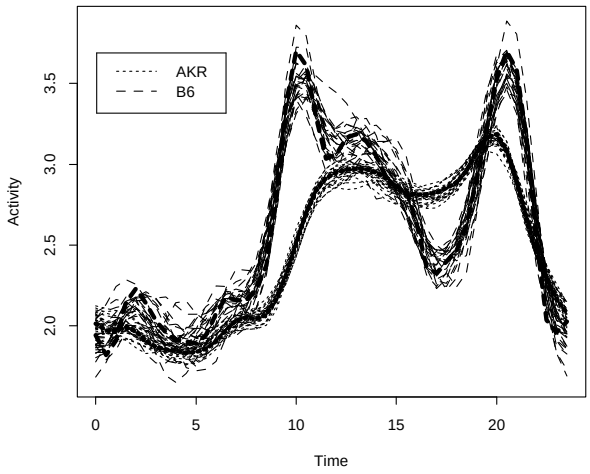

f) BT: Strain comparison

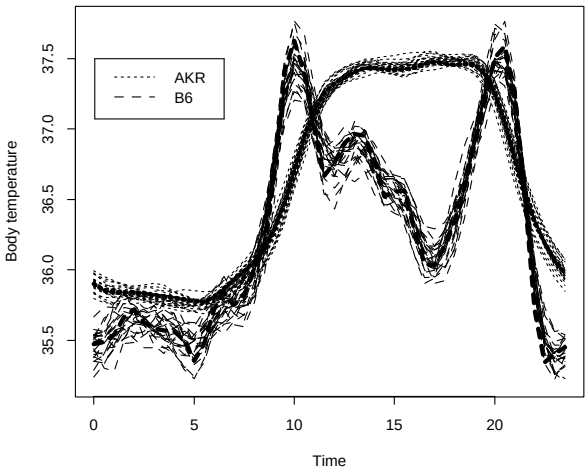

a) Activity

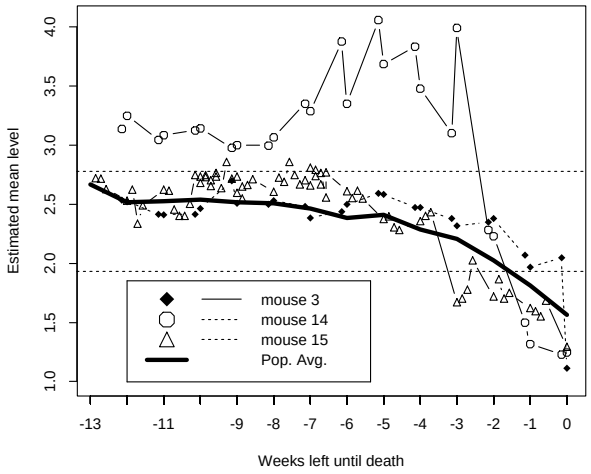

b) Body Temperature

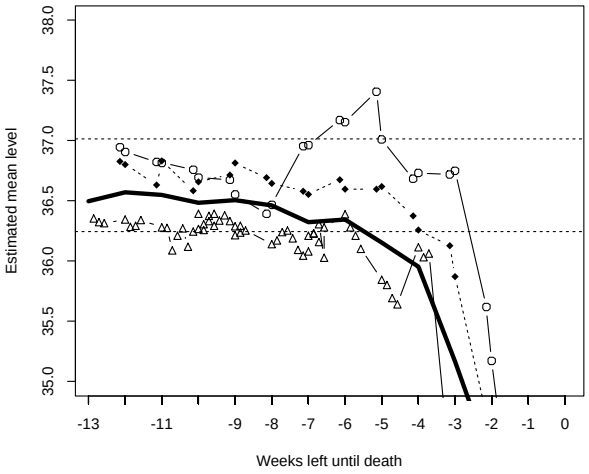

c) Distance

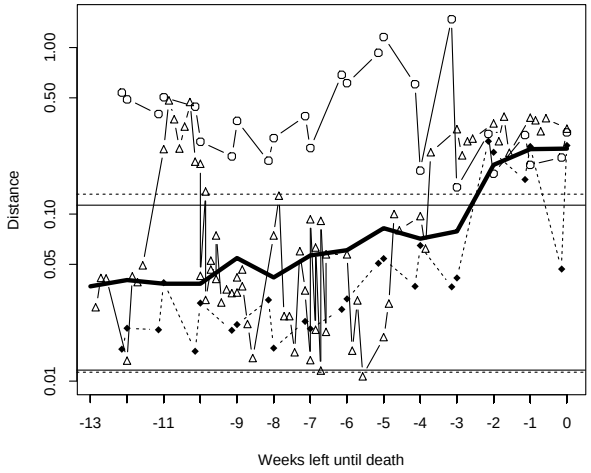

d) Distance

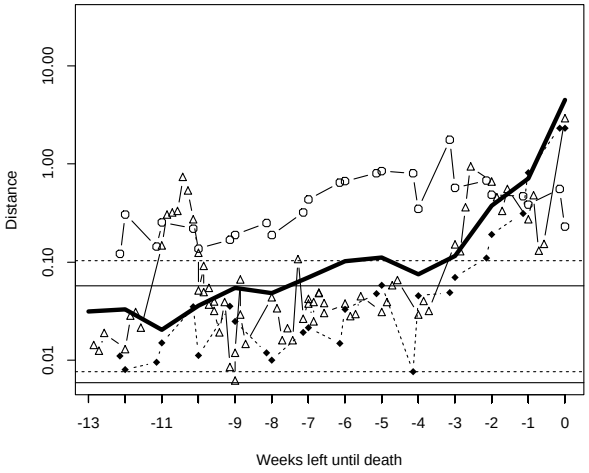

e) Body temp circadian indicador vs. Age

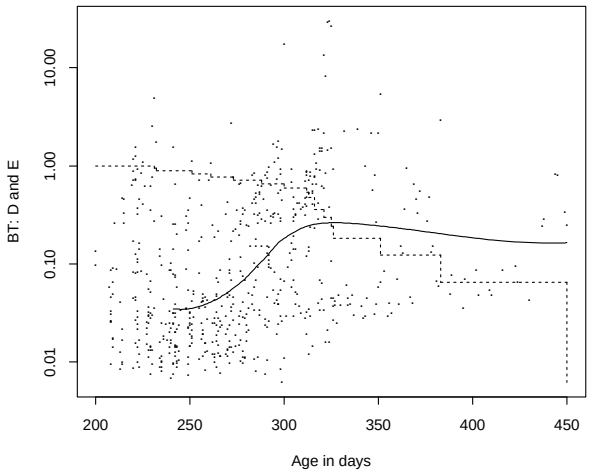

f) Comparison