Modeling Multilevel Sleep Transitional Data Via Poisson Log-Linear Multilevel Models

Bruce J. Swihart*

*Johns Hopkins University, bruce.swihart@gmail.com

Modeling Multilevel Sleep Transitional Data Via Poisson Log-Linear Multilevel Models

Bruce J. Swihart

Abstract

This paper proposes Poisson log-linear multilevel models to investigate population variability in sleep state transition rates. We specifically propose a Bayesian Poisson regression model that is more flexible, scalable to larger studies, and easily fit than other attempts in the literature. We further use hierarchical random effects to account for pairings of individuals and repeated measures within those individuals, as comparing diseased to non-diseased subjects while minimizing bias is of epidemiologic importance. We estimate essentially non-parametric piecewise constant hazards and smooth them, and allow for time varying covariates and segment of the night comparisons. The Bayesian Poisson regression is justified through a re-derivation of a classical algebraic likelihood equivalence of Poisson regression with a log(time) offset and survival regression assuming piecewise constant hazards. This relationship allows us to synthesize two methods currently used to analyze sleep transition phenomena: stratified multi-state proportional hazards models and log-linear models with GEE for transition counts. An example data set from the Sleep Heart Health Study is analyzed.

Modeling Multilevel Sleep Transitional Data Via Poisson Log-Linear Multilevel Models

BRUCE J. SWIHART[*]

Department of Biostatistics,
Johns Hopkins School of Public Health, Baltimore, MD 21205, USA
bswihart@jhsph.edu

BRIAN S. CAFFO

Department of Biostatistics,
Johns Hopkins School of Public Health, Baltimore, MD 21205, USA

CIPRIAN CRAINICEANU

Department of Biostatistics,
Johns Hopkins School of Public Health, Baltimore, MD 21205, USA

NARESH M. PUNJABI

Department of Medicine,
Johns Hopkins University, MD 21205, USA

SUMMARY

This paper proposes Poisson log-linear multilevel models to investigate population variability in sleep

state transition rates. We specifically propose a Bayesian Poisson regression model that is more flexible,

scalable to larger studies, and easily fit than other attempts in the literature. We further use hierarchical

random effects to account for pairings of individuals and repeated measures within those individuals, as

comparing diseased to non-diseased subjects while minimizing bias is of epidemiologic importance. We

estimate essentially non-parametric piecewise constant hazards and smooth them, and allow for time vary-

[*]To whom correspondence should be addressed.

ing covariates and segment of the night comparisons. The Bayesian Poisson regression is justified through a re-derivation of a classical algebraic likelihood equivalence of Poisson regression with a log(time) offset and survival regression assuming piecewise constant hazards. This relationship allows us to synthesize two methods currently used to analyze sleep transition phenomena: stratified multi-state proportional hazards models and log-linear models with GEE for transition counts. An example data set from the Sleep Heart Health Study is analyzed.

Keywords: Bayesian, multi-state, recurrent event, competing risk, hierarchical, stratified, survival analysis

1. Introduction

Hypnograms are time series of a subject's *sleep states* from a single night's sleep. In this manuscript we consider methods for the analysis of hypnogram data. We focus on methods that scale to large cohort studies and complex covariance structures. We show how log-linear random effect models can be derived and used to synthesize existing methods for analyzing hypnogram transition data from large cohort studies and extend it to multilevel settings, unearthing data features classical measures bury. We further discuss model-based methods for exploratory data analysis. We begin with a motivating discussion of two subjects from the community based cohort study prompting this work.

1.1 *Motivating example*

Summaries of the measurement of sleep for two subjects with intrinsically different sleep behavior can highlight or mask these differences. To illustrate, Subject A of Figure 1 has severe Sleep Disordered Breathing (SDB, discussed further below), as indicated by a respiratory disturbance index (RDI) of 52.28 apneic or hypopnic events per hour, while Subject B does not (RDI 0.57). Each subject was monitored

overnight during sleep via a polysomnogram for eight hours. The classical summary of their sleep states, sleep architecture, is similar across the two subjects: Subject A spent 69, 16, 15 and Subject B spent 70, 16, 14 percent of the night in the Non-Rapid Eye Movement (NREM), Rapid Eye Movement (REM) and Wake states, respectively. Whereas their sleep architecture is similar, the temporal evolution of their sleep may not be. Sleep for an individual is often visualized with hypnograms, which are time series of sleep states, depicting states of sleep on the x-axis and time from sleep onset on the y-axis. Subject A and B have similar sleep architecture yet dissimilar hypnograms (see Figure 1). For example, in the zoomed-in portion around hour 7, we see a critical difference in the duration of REM sleep for each subject. Subject A's duration in REM sleep is broken into small chunks, whereas Subject B's is uninterrupted. This is a feature that sleep architecture cannot capture.

We have described population variations of this phenomenon more fully elsewhere (Swihart et al., 2008). Despite severe sleep related disease, sleep architecture remains consistent at the population level. Thus any statistical analysis of sleep architecture as a measure of sleep quality may not account for sleep fragmentation, even in extreme comparisons of severely sleep disordered breathing diseased subjects to healthy.

In our motivating example, Subject A's sleep is more fragmented than Subject B's, with 83 overall transitions between states. Subject B had 47 transitions in total. While, summaries of the overall transition rate are useful, closer study of specific transition types can yield important epidemiological information and directionality of associations with health outcomes (Laffan et al., 2009). A summary of the hypnogram by frequency of transition types for the two subjects is in Table 1. There are 15 as many transitions from REM to NREM for Subject A than Subject B.

We make both scientific and methodologic contributions in this paper. From a scientific perspective,

we 1) develop and substantiate transition rates as an informative population measure for sleep comparisons, 2) report population variations in transition rates for different segments of the night, 3) utilize a very large dataset of sleep biosignals from ∼6400 subjects, and 4) reduce bias in our results via matching. From a methods standpoint, we 1) set forth a framework to view the sleep of a population of individuals as a multi-state survival analysis problem with random effects, 2) re-derive and employ a classical algebraic equivalence between survival analysis and Poisson regression within this framework, 3) smooth the piecewise constant hazards, and 4) accomplish all of this with relative computational ease.

This paper proposes Poisson log-linear multilevel models to investigate population variability in transition rates. We specifically propose a Bayesian Poisson regression model that is more flexible, scalable to larger studies, and easily fit than other attempts in the literature (Sinha, 1993; Clayton, 1991; Sinha & Dey, 1997; Sargent, 1998; Fahrmeir & Klinger, 1998; Sargent, 1997; Kneib & Fahrmeir, 2007; Swihart et al., 2008). We further use hierarchical random effects to account for pairings of individuals and repeated measures within those individuals, as comparing diseased to non-diseased subjects while minimizing bias is of epidemiologic importance. We estimate essentially non-parametric piecewise constant hazards and smooth them, and allow for time varying covariates and segment of the night comparisons. The Bayesian Poisson regression is justified through a re-derivation of a classical algebraic likelihood equivalence of Poisson regression with a log(time) offset and survival regression assuming piecewise constant hazards. This relationship allows us to synthesize two methods currently used to analyze sleep transition phenomena, stratified multi-state proportional hazards models (Therneau & Grambsch, 2000) and log-linear models with GEE (Swihart et al., 2008) for transition counts. Moreover, our suggested Poisson multilevel modelling is more flexible than partial-likelihood based multi-state proportional hazards models and GEE models for transition counts by allowing for nested random effect structures and easily handling time-

varying covariates. We demonstrate that the computational burden of these models is manageable and that the methods can be extended to large cohort studies. Finally, the proposed multilevel models yield random effect predictions of transition risk.

In the next two subsections, a brief yet grounding overview of the science of sleep, sleep data extraction, history of statistical approaches to this type of phenomena, and exploratory data analysis on hypnogram data is given. To close out the introduction is a formulation of the set-up and challenges of analyzing these rich data.
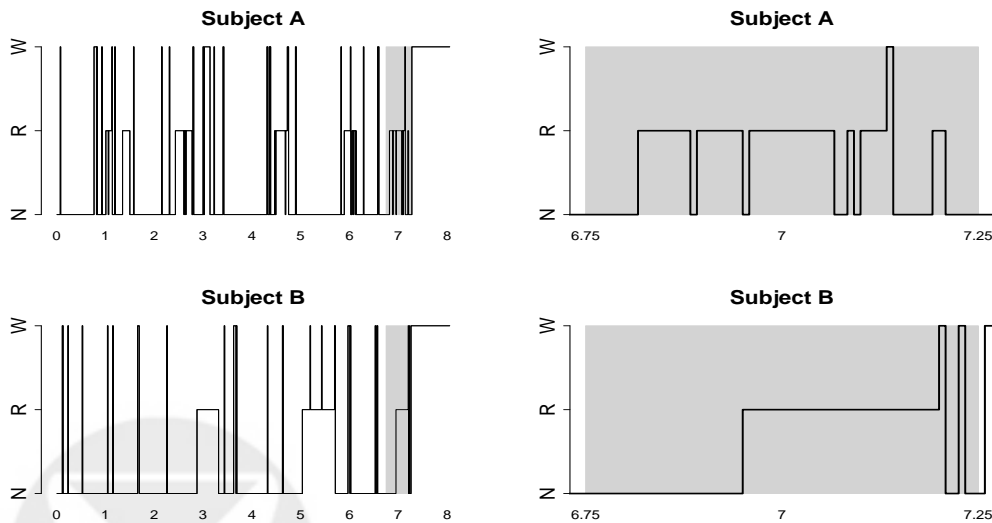

Fig. 1. Left panels, 8 hour sleep hypnograms of Subjects A and B; Right panels, zoomed half-hour portions of the corresponding left panel. On all hypnograms, the x-axis represents the states of sleep (N: Non-REM, R: REM, and W: Wake) a subject can occupy. The y-axis is time of night, with 0 being sleep onset, thus a hypnogram is a state-time graph, showing the trajectory of sleep for an individual.

	Previous state					
	Subject A			Subject B		
Current state	N	R	W	N	R	W
Non-REM (N)	625	15	24	652	1	19
REM (R)	19	138	0	3	155	2
Wake (W)	21	4	119	18	4	111
Total epochs	665	157	143	673	160	132
Total in hours	5.54	1.31	1.91	5.61	1.33	1.10
Sleep Architecture (%)	69	16	15	70	16	14

Table 1. Cross Tabulation of Pairwise Contiguous Epochs for Subject A and B.

1.2 *Biosignals of sleep*

Polysomnography is a multi-faceted sleep monitoring process. A polysomnogram (PSG) is a collection of simultaneous time series that are the measured biological signals related to sleep. The signals that comprise a polysomnogram include the electroencephalogram (EEG), electro-oculogram (EOG), electromyogram (EMG), electrocardiogram (ECG), airflow, chest and abdominal effort, oxyhemoglobin saturation, and body position. A PSG summarizes sleep behavior and is the gold standard for diagnosing somniopathy (sleep disorders), such as SDB and restless leg syndrome. The data from a PSG study is voluminous and complex. Hence feature extraction is usually performed. The previously mentioned RDI or apnea/hypopnea index is an example of clinical feature used in diagnoses and sleep epidemiological research.
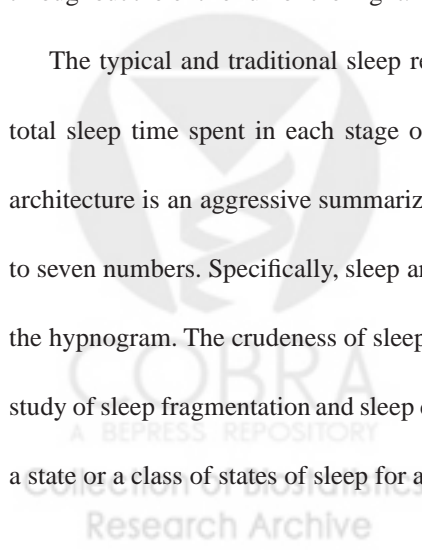
The hypnogram decomposes the PSG into a time series of sleep states via visual classification. The typical six stages of sleep used in this decomposition are known as the R and K system, as put forth by Rechtschaffen and Kales in 1968 and updated by the American Academy of Sleep Medicine (AASM) in 2007 (Rechtschaffen & Kales, 1968). Recently, computer algorithms are supplanting sleep physicians for the task of translating the simultaneous curves of the polysomnogram into the six stages of sleep (Penzel & Conradt, 2000). The summarization that occurs is two-fold: 1) continuous time is discretized

into sequential bins, usually 30 seconds, called "epochs" and 2) within each epoch the information across all time series is combined to declare one of the six stages of sleep. The six stages are: Wake, Stage 1, Stage 2, Stage 3, Stage 4, and Rapid Eye Movement (REM). Often, Stage 1, Stage 2, Stage 3, Stage 4 are grouped into the stage of Non-Rapid Eye Movement (NREM), a convention we adopt in our analysis to limit the number of state-to-state transitions under consideration.

A sleeper passes through these states in a recurrent fashion many times throughout the night. We note that, even though there is a typical progression between states, each of the $\binom{6}{2}$ possible transitions usually occur at least once over the course of a night's sleep. Though possessing limitations, the R and K system facilitates a tremendous data reduction, producing one discrete-time discrete-state process, the hypnogram, from many continuous-time and continuous-state time series.

The hypnogram for an individual is easily visualized and usually is included in a sleep report. It shows that, even when summarized by the R and K system, sleep is a very dynamic discrete-time discrete-state stochastic process, where a transition from any state in an epoch to any other of the six states in the next epoch is possible. Each state has the possibility of being visited for various durations at various times throughout the entire run of the night.

The typical and traditional sleep report also includes sleep architecture, which is the percentage of total sleep time spent in each stage of sleep; this is the canonical summary of the hypnogram. Sleep architecture is an aggressive summarization of the data in the sense that it reduces the entire hypnogram to seven numbers. Specifically, sleep architecture sums over all temporal and transition information from the hypnogram. The crudeness of sleep architecture's description of one's sleep has been limiting for the study of sleep fragmentation and sleep continuity. (Here, continuity of sleep is the concept of remaining in a state or a class of states of sleep for a number of contiguous epochs and fragmentation is the disruption

of this contiguity.)

It has been conjectured that sleep continuity may be an important characteristic of sleep's role in health, especially in the study of sleep disordered breathing and its connections to health outcomes (Norman et al., 2006; Punjabi et al., 1999; Bonnet & Arand, 2003). SDB is a condition where the airway of the throat collapses at least partially, possibly fully, and this collapse stimulates an involuntary response from the sympathetic nervous system: an arousal. An arousal most often does not awaken an individual, but does usually place the sleep process into a state closer to that of wakeful consciousness. These arousals are compiled into indices such as the RDI or apopnea-hyponea index (AHI) as arousal events per hour. Traditionally, the degree of severity of SDB is accorded to a particular range of values for one of the aforementioned indices. An index of less than 5 would often be acknowledged as no-SDB and greater than 30 as severe SDB.

It has been shown that sleep architecture does not necessarily differ between no-SDB and SDB groups (Swihart et al., 2008). This motivates research for a better rubric of sleep fragmentation, for models that can utilize the temporal and transition information of the hypnogram and enable more powerful inferences on the role of sleep and adverse health outcomes.

The structure of the hypnogram lends itself well to multi-state survival models, for which we provide a brief literature review. Sinha (1993) built upon the work of Clayton (1991) in modeling multiple event time data and gave an excellent discourse on the development of survival analysis. Sinha & Dey (1997) give a review of the flexibility and implementation of semiparametric multi-state survival models, giving five ways to model (with assumptions) the non-parametric part of the survival model; however, left to the discussion and future work the methodology of clustered frailties and competing risks. Sargent (1998) outlines a framework for hierarchical Cox proportional hazards regression that leaves the baseline free of

assumptions due to the utilization of the partial likelihood and provides exemplary general notation of the likelihood in preparation for Bayesian modeling. Fahrmeir & Klinger (1998) apply a multi-state survival model with a likelihood derived via counting processes to sleep data which requires integration. They advance the art of modeling time varying covariates in a multi-state framework from that of Sargent (1997). Further work with the counting process likelihood with time varying covariates using mixed models has been put forth by Kneib & Fahrmeir (2007).

One of the first uses of survival analysis involving sleep and SDB was to model hypersomnolence and showed that the more severe the degree of SDB, the greater the daytime sleepiness, as evinced by sleep latency time (Punjabi et al., 1999). Norman et al. (2006) used parametric survival analysis on "sleep runs," where the R and K system was summarized from six states to just two: wake and sleep, and demonstrated that the degree of SDB corresponds to distinct levels of sleep continuity, as represented by a unidimensional estimate. To isolate the effects of SDB on sleep fragmentation, Swihart et al. (2008) fit a log-linear model on the relative frequencies and a multi-state proportional hazards survival model for the hazard ratios describing sleep difference between matched SDB and no-SDB groups.

All aforementioned methods take for granted the R and K system of classification. This summarization of the PSG to the hypnogram perhaps discards useful sleep transitional information on arousals and continuity. While methods for analyzing the PSG, such as the EEG signal (see Crainiceanu et al., to appear in 2009), may fill in important gaps in the R and K summarization, we focus exclusively on the hypnogram data and do not consider the remaining PSG signals.

1.3 *Exploratory Data Analysis and Exploratory Models*

The previous section went through a brief overview of the development of survival analysis and the science

of sleep and SDB. This section demonstrates features of the R and K system data for 102 matched subjects,

51 with SDB and 51 no-SDB (matching details appear later). Sleep architecture shows the percent of night

in REM is statistically different, with the SDB group at 17 percent, no-SDB at 21 (Table 2).

```
Variable                      SDB    no-SDB  p-value
RDI                        40.532     2.114  0.000
BMI                        30.275    30.247  0.972
Age                        61.804    61.804  1.000
Race (% White)             92.160    92.160  1.000
Sex (% Male )              66.667    66.667  1.000
Total Sleep Time          351.397   357.466  0.593
Sleep Efficiency           81.941    83.364  0.743
% Night in Stage 1          5.750     5.577  0.815
% Night in Stage 2         62.693    59.109  0.121
% Night in Stage 3 or 4    13.647    13.908  0.904
% Night in REM             17.909    21.406  0.002
```

```
Table 2. Sleep Architecture
```

To investigate the distribution of transition frequency by type, we can summarize all contiguous pairs

of epochs by a transition type of `previous state` → `current state`. Doing so gives a feel for

which transitions are rare and possibly affected by SDB. For instance, it appears the transition `Wake` →

`REM` (WR) is the least frequent among all pairwise classifications for both groups, yet the SDB group has

over 1.5 times as many such transitions (Table 3).

Collapsing the R and K system into two states of Wake and Sleep, we can plot the probability of each

group being asleep by epoch. Doing so shows similarities between the two groups, but reveals that the

no-SDB group stays asleep longer (Figure 2).

We begin to explore temporal transition models using exploratory two-stage random effect approxima-

Fig. 2. Probability of being asleep over epochs, by disease group.

tions by fitting subject-specific regression models then comparing the fitted parameters across subjects.

Specifically, exploratory transition models can be constructed by using a baseline category logit model

(Agresti, 2003), predicting the stage of the next epoch given the current epoch's stage (Liang & Zeger,

1993). That is, we fit the model

$$\log\{P(Y_{it} = k \mid Y_{i,t-1})/P(Y_{it} = 1 \mid Y_{i,t-1})\} = \beta_t + x_{it2}\delta_2 + x_{it3}\delta_3,$$

separately for each subject, where Y_{it} is the state (taking values $k = 1, 2, 3$ for W, N, R) for subject i at

	Previous state					
	Disease			Controls		
Current state	N	R	W	N	R	W
Non-REM (N)	28,058	124	1,561	27,968	161	1,284
REM (R)	249	6,023	151	364	7,515	96
Wake (W)	1,423	274	7,298	1,070	293	8169
Total epochs	29,730	6,421	9,010	29,402	7,969	9,532
Total in hours	247.75	53.51	75.08	245.02	66.41	79.58

Table 3. Cross Tabulation of Pairwise Contiguous Epochs by Disease Group

epoch t and the design matrix $x_{itj} = I(Y_{i,t-1} = j)$.

Doing this yields fitted values that for the probabilities of being in a stage of sleep one epoch later conditional on the stage experienced in the previous epoch. Using three stages of sleep and applying this model to an individual yields a 3×3 transition matrix with the off-diagonal entries $P(Y_{it} = k \mid Y_{i,t-1} = k')$. We apply this model to each individual of the diseased group and the non-diseased group and plot comparative histograms of the probabilities of 6 different transition types (contiguous epochs of the same state do not consitute a transition). This exploratory model exercise allows for identification of different distributions of transition probabilities between the disease groups (Figure 3).
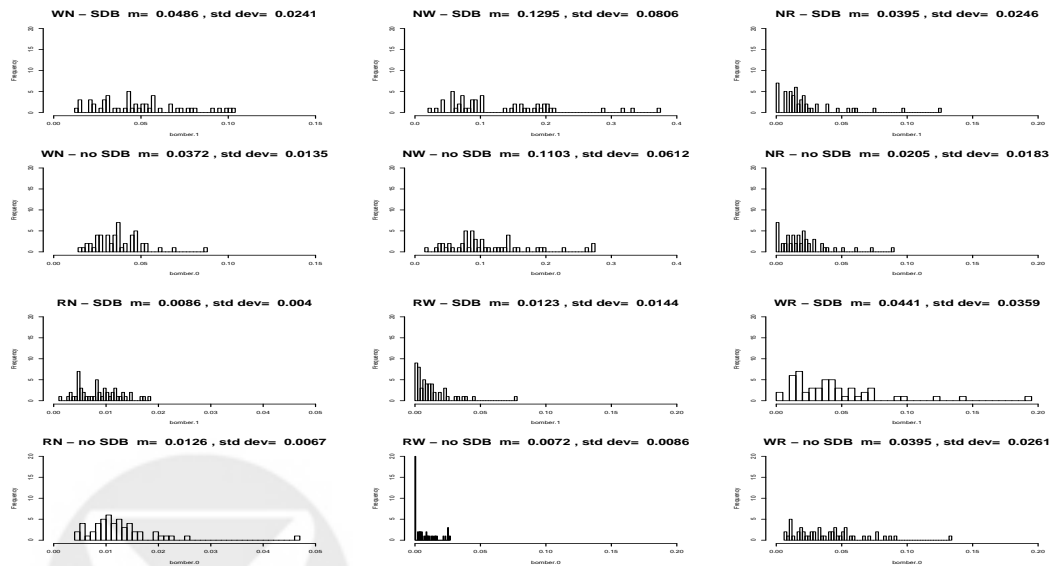

Fig. 3. Multinomial Model Expected Probabilities of Transition by transition-type and disease status

We investigated the inclusion of non-linear non-transitional trends (by epoch) by including natural spline terms. These yield subject-specific time varying probabilities of being in a certain stage in that particular epoch. The multinomial model mandates that its probabilities add to one, and plotting the three probabilities for an individual over epochs shows the trade-offs of the probability of being in a certain

stage. Below is the probability simplex values for a diseased subject, as well as for a non-diseased subject (Figure 4). Note for the diseased individual how fragmented the night becomes with frequent tradeoffs in probability of being in Wake and NREM, and the overall lack of REM probability. Note for the non-diseased individual, the cyclic nature of REM probability.

Fig. 4. The first column is an individual with SDB, the probability of being in the stage by epoch, for one night. The second column is an individual without SDB. Top - NREM, Middle - REM, Bottom - Wake.

Viewing several R and K systems can be done with Lasagna Plots (Figure 5) (our term for advocating heatmaps over traditional spaghetti plots Swihart et al., to be submitted). Each horizontal "layer" across time is a subject's R and K system condensed to Wake, NREM, REM, and Absorbed, with color represent-

ing the subject's sleep state for that epoch. The plots quickly reveal that no-SDB subjects are not without

long stretches of wakefulness in the course of the night and as a whole experience more REM sleep than

the SDB group. These plots are good for visualizing the data of multiple hypnograms.


Fig. 5. Lasagna plot for SDB (top) and no-SDB (bottom).Wake-Black, NREM-Red, REM-Green, Absorbed-Yellow.

1.4 *Set up and challenges*

The sleep transition rate data to be modeled is complex. Our proposed solution is a *multi-state, recurrent*

event, competing risk, hierarchical, stratified survival model fit using Poisson hierarchical models. To

elaborate, it is *Multi-state* as there's more than the traditional 2-states (i.e., alive/dead, wake/sleep, etc.)

in typical survival models. *Recurrent event* because no state is absorbing and all can recur. *Competing*

risk because options exist for the state to which one will transition (from Non-REM to Wake *or* from

Non-REM to REM). *Hierarchical* because of nesting of times-to-event within individuals and individuals

nested within matched pairs. *Stratified* in such a way to render piecewise constant hazards, transition-type

specific inferences between diseased and non-diseased, and inference on segment of night dependence of the transition-specific effects of diseased and non-diseased. Our models are necessarily complex to capture the fine structure of the transition processes that might be of interest. Oversimplification of data, as shown in our first example, can be and is misleading in many applications.

In cohort studies of sleep transitional phenomena, "time" has several meanings which can lead to considerable confusion. We focus on three important distinctions in the discussion of time: duration in state (DIS) time, stopwatch accruing cumulative (SAC) time, and local wall clock (LWC) time. To elucidate, consider an example: a subject falls asleep when the alarm clock on her night stand displays 10:00PM, say. She goes through various states of sleep, and at 11:23pm enters REM sleep. At 11:30pm she exits REM sleep. Consequently, her DIS time for this transition is 7 minutes, her SAC time was 83 minutes when she entered REM, 90 minutes upon exiting. The LWC time of her entering into REM was 11:23PM; of her egress, 11:30PM. Each of these are important, as DIS times are the times-to-event and SAC times help in the segmentation of the night which allows for inference for time-varying transition effects. LWC time is useful to study diurnal effects; for example it has been shown to be important in the studying of sudden death from cardiac causes and sleep disordered breathing (Gami et al., 2005).

We implement the model using MCMC/Gibbs sampling. We show that a segmented SAC time analysis amounts to little adjustment in the model form via minor manipulations of the likelihood. Such segmented SAC time analysis is a vast improvement over the past raw stratification approach of fitting separate models in different portions of the night (Swihart et al., 2008).

The multiple stratifications on transition type, DIS and SAC time interacted with disease status can easily make for high dimension parameterizations as well as binning combinations. Following recent research in smoothing (Di et al., to appear in 2009; Crainiceanu et al., to appear in 2009), we propose a

fine level of binning and allow a smoothing/penalty to prevent over-parameterization, a strategy similar to

the correlated pieces approach (Sinha & Dey, 1997; Sargent, 1998).

Matching is necessary as the data are observational and epidemologic confounding of the disease

effect is of concern. The number of subjects in the Sleep Heart Health Study (SHHS) dataset motivating

this manuscript allow for well populated, well selected sub-groups for the desired comparisons. Matching

is performed via propensity scores (Rosenbaum & Rubin, 1983).

The paper continues with the following sections: Model, Implementation, Application, Results, Dis-

cussion, and Appendix.

2. MODEL

We develop a model in the most general form for the Poisson representation of the hypnogram. As for

notation, $i = 1, ..., I$ indexes individual, $j = 1, ...J_i$ indexes the transitions chronologically, $s = 1, ..., S$

denotes the transition-type, $k = 1, ..., K$ segments SAC time, $l^{(s)} = 1, ..., L_s$ is the transition-type specific

binning of the hazard, w_i is the vector of multiplicative random effects, u_i, as the log of the elements in

w_i, is a vector of additive frailties, z_i is the vector of covariates to linearly combine with the frailties, x_i

is the vector of covariates to linearly combine with the fixed effects. Binary Y_{ijskl} is one if the transition

occurred in the k^{th} segment of SAC time, in the l^{th} bin of the binned hazard for transition-type s; 0

otherwise. Binary δ_{ijskl} is a very useful design variable in the competing risks format. It is one if the j^{th}

transition for individual i is possible as type s in the segment k and bin l, 0 otherwise. To be thorough,

$Y_{ijskl} = \delta_{ijskl} = 1$ for the possible and observed transition, $Y_{ijskl} = 0$, $\delta_{ijskl} = 1$ for possible and

censored, and not possible, $\delta_{ijskl} = 0$ Nonnegative t_{ij} is the duration in state time until the j^{th} transition

occurs for individual i. Nonnegative t_{ijskl} is the amount of time t_{ij} intersected the l^{th} bin of the hazard

for transition-type s for transition j occurring in the k^{th} segment of the night and $\sum_{s,k,l} t_{ijskl} \delta_{ijskl} = t_{ij}$, analogous to Laird & Olivier (1981).

Note that the definition of "segment" of total SAC time is subject-specific and "ragged" in a sense. If a t_{ij} started in one segment and ends in another, the segment k to which it is assigned in its entirety is the greatest (latest) one. With that stated, the binning of SAC supersedes that of the DIS: total SAC time is divided into K segments (i.e. K=2 implies 1st half and 2nd half of night). Then the DIS times are assigned in their entirety to one of the segments. Then the DIS times are partitioned amongst the $l^{(s)} = 1, ..., L_s$ bins. Now, the established relation between survival data and the Poisson likelihood will be reanimated in the outlined framework (Holford, 1976, 1980; Laird & Olivier, 1981). Let the hazard for transition-type s, segment k and bin l be $h_{skl}(t_{ijskl} \mid x_i, z_i, u_i) = h_{0skl}(t_{ijskl}) e^{x_i^\mathsf{T} \beta + z_i^\mathsf{T} u_i}$, where x_i and z_i are covariates that do not depend on transition, SAC time, or DIS time, but easily could. A superscript T denotes a transpose.

The hazard is defined as

$$h_{skl}(t_{ijskl} \mid x_i, z_i, u_i) = \frac{f_{skl}(t_{ijskl}; x_i, z_i, u_i)}{S_{skl}(t_{ijskl}; x_i, z_i, u_i)} = \frac{f_{skl}(t_{ijskl}; x_i, z_i, u_i)}{1 - F_{skl}(t_{ijskl}; x_i, z_i, u_i)},$$

where $f_{skl}(t_{ijskl}; x_i, z_i, u_i)$, $S_{skl}(t_{ijskl}; x_i, z_i, u_i)$, and $F_{skl}(t_{ijskl}; x_i, z_i, u_i)$ are the density, survivor, and distribution functions associated with the survival (DIS) times. The conditional likelihood is therefore:

$$\prod_{i=1}^{I} \prod_{j=1}^{J} \prod_{s=1}^{S} \prod_{k=1}^{K} \prod_{l^{(s)}=1}^{L_s} \left[f(t_{ijskl}; x_i, z_i, u_i)^{y_{ijs}} \{1 - F(t_{ijskl}; x_i, z_i, u_i)\}^{1-y_{ijs}} \right]^{\delta_{ijskl}}$$

$$= \prod_{i=1}^{I} \prod_{j=1}^{J} \prod_{s=1}^{S} \prod_{k=1}^{K} \prod_{l^{(s)}=1}^{L_s} [h(t_{ijskl}; x_i, z_i, u_i)^{y_{ijs}} \{S(t_{ijskl}; x_i, z_i, u_i)\}]^{\delta_{ijskl}} \qquad (2.1)$$

Consider the instance where $\log h_{0skl}(t_{ijskl}) = \mu_{skl}$; hence the strata-specific hazard does not depend on time (t_{ijskl}) and thus f is the exponential density. Utilizing $S(t_{ijskl}; x_i, z_i, u_i) = \exp\{\int_0^{t_{ijskl}} h(r; x_i, z_i, u_i) dr\}$,

the conditional likelihood simplifies to

$$\prod_{i=1}^{I}\prod_{j=1}^{J}\prod_{s=1}^{S}\prod_{k=1}^{K}\prod_{l^{(s)}=1}^{L_s} \{\exp(\mu_{skl} + x_i^\mathsf{T}\beta + z_i^\mathsf{T}u_i)\}^{y_{ijs}\delta_{ijskl}} \exp\{-\delta_{ijskl}t_{ijskl}e^{\mu_{skl}+x_{ijskl}^\mathsf{T}\beta+z_{ijskl}^\mathsf{T}u_i}\}$$

Taking the log and summing over j,

$$= \sum_{i=1}^{I}\sum_{s=1}^{S}\sum_{k=1}^{K}\sum_{l^{(s)}=1}^{L_s} n_{iskl}(\mu_{skl} + x_i^\mathsf{T}\beta + z_i^\mathsf{T}u_i) - e^{\mu_{skl}+x_i^\mathsf{T}\beta+z_i^\mathsf{T}u_i+\log(\Gamma_{iskl})} \qquad (2.2)$$

Noting the general form of the log likelihood for $n \sim \text{Poisson}(\phi)$ is proportional to $nlog(\phi) - \phi$, (2.2) could arise from a Poisson log-linear model with $\phi = \exp\{\mu_{skl} + x_i^\mathsf{T}\beta + z_i^\mathsf{T}u_i + \log(\Gamma_{iskl})\}$. Formally written, the conditional model is:

$$n_{iskl} \mid \mu_{skl}, x_i, \beta, z_i, u_i, \Gamma_{iskl} \sim \text{Poisson}[e^{\mu_{skl}+x_i^\mathsf{T}\beta+z_i^\mathsf{T}u_i+\log(\Gamma_{iskl})}]$$

which is very similar to the Gail, Santner, and Brown rat mammary tumor example (however, in the rat-tumor model considered there, the log offset of aggregated time at risk Γ_{iskl} did not need to be included since it was the same for each rat) (Ibrahim et al., 2001).

Above, n_{iskl} is the count of the number of observed transitions committed during Γ_{iskl}, the total time at risk for person i, of type s, occuring in segment k and bin l. Accounting for Γ_{iskl} is crucial when modeling relative counts, for if a subject makes twice as many transitions as another but had twice as long to do so the rate of transitioning is not truly elevated. If $L_s = 1, \forall s$ and $K = 1$ then (2.2) is equivalent to an exponential survival model. As $L_s \to \infty$ and the model approaches having a completely non-parametric piecewise constant hazard for transition-type s.

The above arguments illustrates how the likelihood equivalence between piecewise exponential survival models synthesizes two methods in practice for analyzing sleep transition data; multi-state proportional hazards models and log-linear models. Sandwich variance estimates were used in Swihart et al.

(2008) to account for within-subject correlation. We instead propose a fully Bayesian approach that utilizes a hierarchical random effect structure.

3. IMPLEMENTATION

For a Bayesian analysis of the model, inference was attained via Markov Chain Monte Carlo. We follow closely the formulation and notation of Sargent (Sargent, 1998). Allowing $\theta_\omega = (x, z, \omega, n, \Gamma)$ and $\theta_{u,\mu} = (x, z, u, \mu, n, \Gamma)$, the posterior distribution is proportional to three components:

$$p(\beta, u, \mu, \xi \mid \theta_\omega) \propto L(\beta \mid \theta_{u,\mu})g(\beta, u, \mu \mid \xi)q(\xi \mid \omega)$$

We choose independent priors and hyper-priors, yielding:

$$p(\beta, u, \mu, \xi \mid \theta_\omega) \propto L(\beta \mid \theta_{u,\mu}) \times$$

$$g_\beta(\beta \mid \xi_\beta)g_{fra,set}(u \mid \xi_{fra,set})g_{haz}(\mu \mid \xi_{haz}) \times$$

$$q_\beta(\xi_\beta \mid \omega_\beta)q_{fra,set}(\xi_{fra,set} \mid \omega_{fra,set})q_{haz}(\xi_{haz} \mid \omega_{haz}).$$

Addressing each piece, the likelihood

$$L(\beta \mid \theta_{u,\mu}) = \prod_{i=1}^{I}\prod_{s=1}^{S}\prod_{k=1}^{K}\prod_{l(s)=1}^{L_s} e^{n_{iskl}(\mu_{skl}+x_i^\mathsf{T}\beta+z_i^\mathsf{T}u_i)-e^{\mu_{skl}+x_i^\mathsf{T}\beta+z_i^\mathsf{T}u_i}+\log\Gamma_{iskl}}$$

and the prior on the regression coefficients $g_\beta(\beta \mid \xi_\beta) = N(\mathbf{0}, \mathbf{\Sigma})$. We partition the vector β into the elements yielding relative transition rates and those adjusting for covariates,

$$\beta = (\beta_{trans}, \beta_{cov})^\mathsf{T} = (\beta_{11}, ..., \beta_{sk}, ..., \beta_{SK}, \beta_{\text{covariate 1}}, ..., \beta_{\text{covariate p}})^\mathsf{T}.$$

We choose Σ to be a diagonal matrix with only two unique non-zero elements: σ_{trans}^2 appearing in the first SK diagonal spots, and the remaining diagonal spots filled with σ_{cov}^2. Hyperprior $q_\beta(\xi_\beta \mid \omega_\beta) =$

$Gamma(\alpha_{trans}, \phi_{trans})Gamma(\alpha_{cov}, \phi_{cov})$, where $\xi_\beta = (\xi_{trans}, \xi_{cov}) = (\frac{1}{\sigma^2_{trans}}, \frac{1}{\sigma^2_{cov}})$. Equivalent

to setting log-normal priors directly on the individual and pair (set) frailties, we appropriate $g_{fra,set}(u \mid \xi_{fra,set}) =$

$N(0, \sigma^2_{fra})N(0, \sigma^2_{set})$ and allot the hyperprior $q_{fra,set}(\xi_{fra,set} \mid \omega_{fra,set}) = Gamma(\alpha_{fra}, \phi_{fra})Gamma(\alpha_{set}, \phi_{set})$,

where ξ_β represents $(\xi_{fra}, \xi_{set}) = (\frac{1}{\sigma^2_{fra}}, \frac{1}{\sigma^2_{set}})$. Similar to the prior for the frailties, $g_{haz}(\mu \mid \xi_{haz}) =$

$N(0, \sigma^2_{haz})$ prior on each μ_{skl}, which is the same as log-normal priors directly on the baseline haz-

ard, h_{0skl}. Lastly $q_{haz}(\xi_{haz} \mid \omega_{haz}) = Gamma(\alpha_{haz}, \phi_{haz})$ serves as the hyper-prior on ξ_{haz} where

$\xi_{haz} = \frac{1}{\sigma^2_{haz}}$.

Which gives the general model:

$$p(\beta, u, \mu, \xi \mid \theta_\omega) \propto \prod_{i=1}^{I}\prod_{s=1}^{S}\prod_{k=1}^{K}\prod_{l(s)=1}^{L_s}\prod_{p=1}^{P} e^{n_{iskl}(\mu_{skl}+x_i^\mathsf{T}\beta+z_i^\mathsf{T}u_i)-e^{\mu_{skl}+x_i^\mathsf{T}\beta+z_i^\mathsf{T}u_i}+\log\Gamma_{iskl}} \times$$

$$\frac{1}{\sqrt{2\pi\sigma^2_{trans}}}e^{-\frac{\beta^2_{sk}}{2\sigma^2_{trans}}}\frac{1}{\sqrt{2\pi\sigma^2_{cov}}}e^{-\frac{\beta^2_{SK+p}}{2\sigma^2_{cov}}}\frac{1}{\sqrt{2\pi\sigma^2_{fra}}}e^{-\frac{u^2_{1i}}{2\sigma^2_{fra}}}\frac{1}{\sqrt{2\pi\sigma^2_{set}}}e^{-\frac{u^2_{2i}}{2\sigma^2_{set}}}\frac{1}{\sqrt{2\pi\sigma^2_{haz}}}e^{-\frac{\mu^2_{skl}}{2\sigma^2_{haz}}} \times$$

$$\frac{\phi^{\alpha_{trans}}_{trans}}{\Gamma(\alpha_{trans})}(\frac{1}{\sigma^2_{trans}})^{\alpha_{trans}-1}e^{\frac{\phi_{trans}}{\sigma^2_{trans}}}\frac{\phi^{\alpha_{cov}}_{cov}}{\Gamma(\alpha_{cov})}(\frac{1}{\sigma^2_{cov}})^{\alpha_{cov}-1}e^{\frac{\phi_{cov}}{\sigma^2_{cov}}} \times$$

$$\frac{\phi^{\alpha_{fra}}_{fra}}{\Gamma(\alpha_{fra})}(\frac{1}{\sigma^2_{fra}})^{\alpha_{fra}-1}e^{\frac{\phi_{fra}}{\sigma^2_{fra}}}\frac{\phi^{\alpha_{set}}_{set}}{\Gamma(\alpha_{set})}(\frac{1}{\sigma^2_{set}})^{\alpha_{set}-1}e^{\frac{\phi_{set}}{\sigma^2_{set}}} \times$$

$$\frac{\phi^{\alpha_{haz}}_{haz}}{\Gamma(\alpha_{haz})}(\frac{1}{\sigma^2_{haz}})^{\alpha_{haz}-1}e^{\frac{\phi_{haz}}{\sigma^2_{haz}}}.$$

We also consider a smoothing of the hazard bins, in which case we supplant the μ_{skl} priors above

with: $g_{haz}(\mu \mid \xi_{haz}) = N(\theta_{skl}, \sigma_{haz})$ prior on each μ_{skl}, where $\theta_{skl} = 0$ if $l = 1$, $\theta_{skl} = \mu_{sk(l-1)}$ if

$l > 1$ This allows bins to be "similar" to each other. This is what we refer to hence forth as the "smoothed"

model.

If no demographic covariates are included in the process ($P = 0$), any density involving the subscript

cov can be eliminated from the posterior joint distribution. Likewise, if it is not desired to keep track of paired/set frailties, then any density involving the subscript *set* should be eliminated from the posterior joint distribution.

4. Application

The application makes use of R and K system data from the Sleep Heart Health Study (SHHS), a multicenter study on SDB and cardiac outcomes (Quan et al., 1997). Subjects for the SHHS were recruited from ongoing cohort studies on respiratory and cardiovascular disease. From the first SHHS cohort of over 6300 subjects, 5614 were identified as having reliable and high quality in home polysomnograms. To assess the independent effects of SDB on sleep structure, a matched subset of the 5614 with and without SDB was selected for the current study. Subjects with severe SDB were identified as those with a RDI > 30 events/hour. Subjects without SDB were identified as those with an RDI < 5 events/hour. Other exclusion criteria included prevalent cardiovascular disease, hypertension, chronic obstructive pulmonary disease, asthma, coronary heart disease, history of stroke, and current smoking.

Propensity score matching was utilized to balance the groups on demographic factors and to minimize confounding. SDB subjects were matched with no-SDB subjects on the factors of age, BMI, race, and sex. Race and sex were exactly matched, while age and BMI were matched using the nearest neighbor Mahalanobis technique with a caliper of 0.10. The resultant match was 51 pairs that met the strict inclusion criteria outlined above and exhibiting very low standardized biases, a vast improvement on the imbalance of BMI between diseased and non-diseased groups of past studies (Swihart et al., 2008). Polar opposites of SDB severity, isolated from comorbities, were used to increase the likelihood of finding 1) differences in sleep architecture (see Table 2) and 2) independent effects of SDB on sleep continuity.

Conceptualizing sleep as a multi-state competing risks process, we analyzed 3-state sleep, collapsing the four stages of non-REM into one state, "NREM", leaving the traditional "Wake" and rapid eye movement "REM" states. From any of the three states one may transition into the others producing six possible transition types: Wake to NREM (WN), NREM to Wake (NW), NREM to REM (NR), REM to Wake (RW), REM to NREM (RN), and Wake to REM (WR).

In the context of the application, $i = 1, ..., 102$ indexes individual, $s = 1, ..., 6$ denotes the transition-type, $k = 1, 2$ segments the night,

$$(L_1, L_2, L_3, L_4, L_5, L_6) = (2, 6, 12, 12, 12, 1)$$

is the transition-type specific binning of the hazard, which was determined by the distinct quantiles of the duration in state times per transition-type s. Finding L_s was done iteratively, first attempting to have 12 bins with approximately the same number of transitions of type s in them for model stability. The number 12 was selected for its versatility: one pass through the data binning hazards into 12ths and one could easily construct 12, 6, 4, 3, 2, or 1 piece models by summing number of transtions and total duration in state time, collapsing 1/12 bins into larger fraction binning. If the type s did not yield distinct quanitles for 12 bins, then bin sizes of 6, 4, 3, 2, and 1 were sequentially tried. The vector $w_i = (w_{1i}, w_{2i})$ of multiplicative random effects, the first for individual and the second for matched pair. The vector $u_i = (u_{1i}, u_{2i})$, as the log of the elements in w_i, is a vector of additive random effects. The vector $z_i = (1, 1)$ in models with individuals nested within matched pair, $(1, 0)$ for models not accounting for pair. The vector x_i is composed of the design variables and (potentially) the demographic covariates. The design variables are the 3-way interaction of disease status, the k^{th} segment of the total SAC time, and transition-type s. The design interaction variables require the data to be at the "cross-binned" $i - s - k - l$ level and this enables the corresponding β vector to have elements β_{sk} which quantify the average transition frequency

of type s in the k^{th} segment of the total SAC time for diseased versus non-diseased. In the case of $K = 2$, this allows sampling from the posterior distribution of the composite quanitity of the rate ratio between the two segments of night $(\frac{\exp(\beta_{s2})}{\exp(\beta_{s1})})$, enabling inference as to whether transition intensities change over the course of sleep.

Models with various combinations of bin smoothing, accounting for pair frailty, and number of included demographic covariates are fit. All models were fit with two segments of total SAC time $(K = 2)$ and the aforementioned $l^{(s)}$. For each model, we ran five chains for 1200 iterations and used the last 200 of each chain, yielding 1000 samples from each relevant full conditional of β, u_i and μ_{skl} (where $u_i = \log(w_i)$, $\mu_{skl} = log(h_{0skl})$). Our hyper-parameter values were selected based on Sargent (1998):

$$\omega = (\alpha_{trans}, \phi_{trans}, \alpha_{cov}, \phi_{cov}, \alpha_{fra}, \phi_{fra}, \alpha_{set}, \phi_{set}, \alpha_{haz}, \phi_{haz})$$

$$= (1.1, 0.1, 0.1, 0.1, 1.1, 0.1, 1.1, 0.1, 0.1, 0.1).$$

5. RESULTS

Upon visual inspection of trace plots, the chains were well mixed and the lag auto-correlation was acceptable (see Appendix). Convergence monitoring was conducted using the Brooks and Gelman diagnostic (Carlin & Louis, 2000; Brooks & Gelman, 1998) (acknowledging the limitations of such convergence diagnostic measures). A vast majority of these univariate diagnostics are greater than but close to 1, suggesting convergence and appropriately overdispersed starting values. From graphical inspection of the diagnostic over iterations, a vast majority not only narrow to 1, but also show the stabilization of the pooled and within interval widths.

| | Model | | | Rate Ratios for SDB vs. no-SDB by Transition Type s | | | | | |
Smoothed	Pair Frailty	No. Covariates	Night Segment	WN	NW	NR	RW	RN	WR
Yes	Yes	4	1	$_{0.99}1.12_{1.28}$	$_{1.1}1.25_{1.42}$	$_{0.56}0.72_{0.92}$	$_{1.02}1.32_{1.73}$	$_{0.67}0.93_{1.24}$	$_{1.57}2.66_{4.95}$
			2	$_{0.87}0.98_{1.11}$	$_{1.11}1.26_{1.42}$	$_{0.54}0.66_{0.81}$	$_{0.87}1.07_{1.31}$	$_{0.74}0.98_{1.3}$	$_{0.78}1.01_{1.32}$
Yes	Yes	2	1	$_{0.99}1.12_{1.29}$	$_{1.1}1.26_{1.43}$	$_{0.55}0.71_{0.92}$	$_{1.01}1.31_{1.72}$	$_{0.67}0.9_{1.23}$	$_{1.59}2.69_{4.55}$
			2	$_{0.87}0.98_{1.10}$	$_{1.11}1.27_{1.43}$	$_{0.53}0.66_{0.81}$	$_{0.87}1.08_{1.33}$	$_{0.76}1.00_{1.31}$	$_{0.78}1.03_{1.38}$
Yes	Yes	0	1	$_{1.00}1.13_{1.27}$	$_{1.11}1.27_{1.43}$	$_{0.56}0.73_{0.93}$	$_{0.98}1.30_{1.7}$	$_{0.70}0.93_{1.29}$	$_{1.57}2.65_{4.36}$
			2	$_{0.88}0.98_{1.11}$	$_{1.12}1.27_{1.43}$	$_{0.53}0.66_{0.81}$	$_{0.89}1.08_{1.31}$	$_{0.75}0.99_{1.32}$	$_{0.77}1.03_{1.36}$
Yes	No	0	1	$_{0.97}1.12_{1.29}$	$_{1.10}1.24_{1.39}$	$_{0.55}0.71_{0.91}$	$_{1.00}1.31_{1.69}$	$_{0.67}0.91_{1.26}$	$_{1.62}2.71_{4.43}$
			2	$_{0.86}0.97_{1.09}$	$_{1.10}1.25_{1.41}$	$_{0.53}0.66_{0.82}$	$_{0.88}1.07_{1.28}$	$_{0.75}0.98_{1.28}$	$_{0.78}1.02_{1.35}$
No	No	0	1	$_{0.98}1.12_{1.26}$	$_{1.07}1.22_{1.39}$	$_{0.53}0.68_{0.86}$	$_{0.96}1.25_{1.61}$	$_{0.64}0.87_{1.14}$	$_{1.57}2.56_{4.42}$
			2	$_{0.85}0.96_{1.09}$	$_{1.10}1.24_{1.42}$	$_{0.5}0.63_{0.78}$	$_{0.87}1.05_{1.29}$	$_{0.72}0.95_{1.25}$	$_{0.77}1.01_{1.33}$
No	Yes	0	1	$_{0.98}1.11_{1.26}$	$_{1.09}1.24_{1.4}$	$_{0.53}0.68_{0.87}$	$_{0.98}1.25_{1.66}$	$_{0.65}0.87_{1.18}$	$_{1.51}2.48_{4.24}$
			2	$_{0.86}0.97_{1.1}$	$_{1.10}1.26_{1.41}$	$_{0.51}0.64_{0.81}$	$_{0.86}1.05_{1.29}$	$_{0.72}0.95_{1.24}$	$_{0.76}1.01_{1.32}$

Table 4. Rate Ratios for SDB vs. no-SDB by Transition Type. Blue indicates diseased transition significantly more than non-diseased. Red indicates diseased transition significantly less than non-diseased. The tables are in a format where the elements are the estimates, credible intervals as the subscripts, the center number the estimate (Louis & Zeger, 2007).

6. DISCUSSION

All models exhibit SDB subjects transitioning significantly more of type NREM → Wake in both halves of the night, Wake → REM in the first half of the night, and significantly less of type NREM → REM for both segments of the night (Table 4). In other words, given a SDB subject is in NREM, he is more likely than a no-SDB subject to transition to Wake and less likely to transition to REM regardless of how long he has been asleep. This is corroboratively linking with findings of SDB subjects having higher all cause mortality (Punjabi et al., 2009) and increases in NREM → Wake and decreases in NREM → REM leading to higher all cause mortality (Laffan et al., 2009).

Model			Relative Rate Ratios for SDB vs. no-SDB of segment 2 vs segment 1 by Transition Type s					
Smoothed	Frailty	No. Covariates	WN	NW	NR	RW	RN	WR
Yes	Yes	4	$_{0.75}0.87_{1.01}$	$_{0.57}0.82_{1.15}$	$_{0.87}1.01_{1.18}$	$_{0.72}1.08_{1.64}$	$_{0.66}0.92_{1.24}$	$\color{red}{_{0.2}0.4_{0.71}}$
Yes	Yes	2	$_{0.75}0.87_{1.02}$	$_{0.59}0.84_{1.14}$	$_{0.86}1.01_{1.18}$	$_{0.74}1.14_{1.72}$	$_{0.69}0.94_{1.27}$	$\color{red}{_{0.21}0.4_{0.72}}$
Yes	Yes	0	$_{0.75}0.87_{1.01}$	$_{0.62}0.84_{1.14}$	$_{0.86}1.01_{1.18}$	$_{0.70}1.08_{1.59}$	$_{0.67}0.92_{1.24}$	$\color{red}{_{0.22}0.41_{0.7}}$
Yes	No	0	$_{0.75}0.87_{1.02}$	$_{0.60}0.82_{1.11}$	$_{0.86}1.01_{1.18}$	$_{0.71}1.10_{1.62}$	$_{0.68}0.94_{1.28}$	$\color{red}{_{0.21}0.39_{0.64}}$
No	No	0	$\color{red}{_{0.74}0.86_{1.00}}$	$_{0.61}0.85_{1.12}$	$_{0.86}1.02_{1.19}$	$_{0.73}1.11_{1.64}$	$_{0.67}0.94_{1.29}$	$\color{red}{_{0.22}0.41_{0.67}}$
No	Yes	0	$_{0.75}0.88_{1.02}$	$_{0.61}0.85_{1.15}$	$_{0.87}1.02_{1.19}$	$_{0.74}1.11_{1.66}$	$_{0.67}0.96_{1.33}$	$\color{red}{_{0.23}0.43_{0.69}}$

Table 5. Comparisons of beta coefficients, 2nd segment of night to 1st segment. Blue indicates the relative rate of 2nd segment of night for diseased transitioning compared to the non-diseased is significantly more than that of the 1st segment. Red indicates the relative rate of 2nd segment of night for diseased transitioning compared to the non-diseased is significantly less than that of the 1st segment. The tables are in a format where the elements are the estimates, credible intervals as the subscripts, the center number the estimate (Louis & Zeger, 2007).

Given a SDB subject is in Wake he is on average ~ 2.6 times as likely as his no-SDB counterpart to transition to REM in the 1st half of the night. However, there is no significant difference between the SDB groups for the WR transition in the second half of the night. The segmented SAC time analysis of the 2nd half of the night to the 1st shows a reduction of 60% of the disparity between average transition frequencies of diseased and non-diseased for type WR (Table 5). This suggests the second half of the night has both groups getting to REM from Wake at more simliar rates than the first half.

As for the accounting for pairing discussion, (Table 4) shows very little difference between models differing only by the accounting of pairs. In those comparisons, the magnitudes and directions mirror well, and the only difference in significant results are due to 95% credible intervals containing 1.00. It appears

that in this analysis, the gain in parsimony would favor the omission of pairing information (Stuart, 2008).

The model described the sleep hypnogram more fully than that of traditional sleep architecture, where only % time in REM differed: SDB 17%, no-SDB 21%. Showing the derivation of the Poisson representation provides motivation for a shift in the conceptualization of modeling sleep. The problem can be thought of as a *multi-state, recurrent event, competing risk, hierarchical, stratified survival* model or a Poisson process with the sufficent statistics of number of transitions arising from time at risk for those transitions. This shift makes concerns about tie handling of DIS times inconsequential. The ability to piecewise model the hazard, segment the night, and account for transition-type allow for a very flexible model that can easily incorporate time varying covariates. The model is very scalable, with analysis on 5,614 individuals taking just under 5 hours on a laptop with a 1.83 GHz processor.

MCMC allowed us to account for the correlation induced by repeated measurements on the same individual nested within matched pairs and would facilitate the examination of the heterogeneity in our population through random intercepts. Heterogeneity of populations is a very crucial topic in epidemiologic studies. Through the assumption of exponential survival times we gain a framework that potentially allows us to eschew/relax parametric assumptions about the hazard. These reasons plus the eloquence of jointly modeling the frequency of transitions and times to transition make the Bayesian Poisson regression framework a powerful and flexible tool in modeling sleep as represented by hypnograms.

REFERENCES

Agresti, A. (2003): *Categorical data analysis*. Wiley-Interscience.

Bonnet, M. & Arand, D. (2003): Clinical effects of sleep fragmentation versus sleep deprivation. *Sleep Medicine Reviews* **7**(4):297–310.

Brooks, S. & Gelman, A. (1998): General Methods for Monitoring Convergence of Iterative Simulations. *Journal of Computational and Graphical Statistics* **7**:434–455.

Carlin, B. & Louis, T. (2000): *Bayes and Empirical Bayes Methods for Data Analysis*. Chapman & Hall/CRC.

Clayton, D. (1991): A Monte Carlo Method for Bayesian Inference in Frailty Models. *Biometrics* **47**(2):467–485.

Crainiceanu, C.; Caffo, B. & Punjabi, N. (to appear in 2009): Nonparametric Signal Extraction and Measurement Error in the Analysis of Electroencephalographic Data. *JASA* .

Di, C.; Crainiceanu, C.; Caffo, B. & Punjabi, N. (to appear in 2009): Multilevel Functional Principal Component Analysis. *The Annals of Applied Statistics* .

Fahrmeir, L. & Klinger, A. (1998): A nonparametric multiplicative hazard model for event history analysis. *Biometrika* **85**(3):581–592.

Gami, A.; Howard, D.; Olson, E. & Somers, V. (2005): Day-Night Pattern of Sudden Death in Obstructive Sleep Apnea.

Holford, T. (1976): Life tables with concomitant information. *Biometrics* **32**(3):587–597.

Holford, T. (1980): The analysis of rates and of survivorship using log-linear models. *Biometrics* **36**(2):299–305.

Ibrahim, J.; Chen, M. & Sinha, D. (2001): *Bayesian Survival Analysis*. Springer.

Kneib, T. & Fahrmeir, L. (2007): A Mixed Model Approach for Geoadditive Hazard Regression. *Scandinavian Journal of Statistics* **34**(1):207–228.

Laffan, A.; Gottlieb, D.; Monahan, K.; Quan, S.; Robbins, J.; Samet, J. & Punjabi, N. (2009): Sleep fragmentation predicts all-cause mortality in a cohort of middle aged and older adults. *Manuscript in progress* **XX**(Y):zz–aa.

Laird, N. & Olivier, D. (1981): Covariance Analysis of Censored Survival Data Using Log-Linear Analysis Tech-

niques. *Journal of the American Statistical Association* **76**(374):231–240.

Liang, K. & Zeger, S. (1993): Regression Analysis for Correlated Data. *Annual Reviews in Public Health* **14**(1):43–68.

Louis, T. & Zeger, S. (2007): Effective communication of standard errors and confidence intervals. *Johns Hopkins University, Dept. of Biostatistics Working Papers* page 151.

Norman, R.; Scott, M.; Ayappa, I.; Walsleben, J. & Rapoport, D. (2006): Sleep continuity measured by survival curve analysis. *Sleep* **29**(12):1625–31.

Penzel, T. & Conradt, R. (2000): Computer based sleep recording and analysis. *Sleep Medicine Reviews* **4**(2):131–148.

Punjabi, N.; Caffo, B.; Goodwin, J.; Gottlieb, D.; Newman, A. et al. (2009): Sleep-Disordered Breathing and Mortality: A Prospective Cohort Study. *PLoS Med* **6**(8):e1000132.

Punjabi, N.; O'Hearn, D.; Neubauer, D.; Nieto, F.; Schwartz, A.; Smith, P. & Bandeen-Roche, K. (1999): Modeling Hypersomnolence in Sleep-disordered Breathing A Novel Approach Using Survival Analysis. *American Journal of Respiratory and Critical Care Medicine* **159**(6):1703–1709.

Quan, S.; Howard, B.; Iber, C.; Kiley, J.; Nieto, F.; O'Connor, G.; Rapoport, D.; Redline, S.; Robbins, J.; Samet, J. et al. (1997): The Sleep Heart Health Study: design, rationale, and methods. *Sleep* **20**(12):1077–85.

Rechtschaffen, A. & Kales, A. (1968): A Manual of Standardized Terminology. *Techniques and Scoring System for Sleep Stages of Human Subjects. Washington, US Government Printing Office* .

Rosenbaum, P. & Rubin, D. (1983): The central role of the propensity score in observational studies for causal effects. *Biometrika* **70**(1):41–55.

Sargent, D. (1997): A Flexible Approach to Time-varying Coefficients in the Cox Regression Setting. *Lifetime Data Analysis* **3**(1):13–25.

Sargent, D. (1998): A general framework for random effects survival analysis in the Cox proportional hazards setting. *Biometrics* **54**(4):1486–97.

Sinha, D. (1993): Semiparametric Bayesian Analysis of Multiple Event Time Data. *Journal of the American Statistical*

Association **88**(423):979–983.

Sinha, D. & Dey, D. K. (1997): Semiparametric Bayesian Analysis of Survival Data. *Journal of the American Statis-*
tical Association **92**(439):1195–1212.

Stuart, E. (2008): Developing practical recommendations for the use of propensity scores: A Discussion. *Stat Med*
27(12):2062–2065.

Swihart, B.; Caffo, B.; Bandeen-Roche, K. & Punjabi, N. (2008): Characterizing sleep structure using the hypnogram.
Journal of Clinical Sleep Medicine **4**(4):349–355.

Swihart, B.; Caffo, B.; James, B.; Schwartz, B.; Strand, M. & Punjabi, N. (to be submitted): Lasagna Plots: A saucy
new alternative to Spaghetti Plots. *Epidemiology* pages XXX–XXX.

Therneau, T. & Grambsch, P. (2000): *Modeling Survival Data: Extending the Cox Model*. Springer.

Appendix

Subset of Chains from MCMC Sampling