Published in final edited form as:

Psychometrika. 2017 March ; 82(1): 233–244. doi:10.1007/s11336-016-9495-z.

Application of Correlated Time-to-Event Models to Ecological Momentary Assessment Data

Emily A. Scherer,

Division of Biostatistics, Department of Biomedical Data Science, Geisel School of Medicine at Dartmouth, 1 Medical Center Drive, Lebanon, NH 03766

Lin Huang, and

Clinical Research Center, Boston Children's Hospital and Harvard Medical School, Boston, MA

Lydia A. Shrier

Division of Adolescent/Young Adult Medicine, Boston Children's Hospital and Harvard Medical School, Boston, MA

Abstract

Ecological Momentary Assessment (EMA) data consists of in-the-moment sampling several times per day aimed at capturing phenomena that are highly variable. When research questions are focused on the association between a construct measured repeatedly and an event that occurs sporadically over time interspersed between repeated measures, the data consist of correlated observed or censored times to an event. In such a case, specialized time-to-event models that account for correlated observations are required to properly assess the relationships under study. In the current study, we apply two time-to-event analysis techniques, proportional hazards and accelerated failure time modeling, to data from a study of affective states and sexual behavior in depressed adolescents and illustrate differing interpretations from the models.

Introduction

Ecological Momentary Assessment (EMA) is a technique particularly useful when measuring behavioral outcomes (Stone, Shiffman, Atienza, & Nebeling, 2007). EMA consists of in-the-moment sampling several times per day aimed at capturing phenomena that are highly variable, such as affective states or social context. Study participants usually carry an electronic device (e.g., handheld computer or smartphone) that signals at random times throughout the day, prompting completion of a questionnaire. Because questionnaires are answered in the moment, recall bias with respect to behavioral events (e.g., sexual intercourse or substance use) is minimized (Stone & Shiffman, 2002). Analysis of EMA data to model both average levels of behavior (Schwartz & Stone, 1998) and variability within behaviors (Hedeker & Mermelstein, 2007, 2008) via multilevel models is an established area of methodologic study. However, time-to-event modeling with EMA data has received considerably less attention.

A fully observed event time consists of the time between a specified starting point and the occurrence of an event. If all event times are fully observed, these times make up data that can be analyzed with standard statistical models. Time-to-event modeling is necessary when

the outcome under study is the time to a certain event that may or may not be fully observed. Times are censored (i.e., not observed) if the event has not happened for a certain amount of time, but the exact time to event is not known. Such data require analysis techniques that take into account these censored data. For non-EMA data, there are many ways to analyze time-to-event data, otherwise known as “survival” data. Two of the most common ways are to use a proportional hazards (PH) model (Cox, 1972, 1975) or an accelerated failure time (AFT) model (Collett, 1994; Wei, 1992). The interpretation of parameters from these models differs; the PH model assumes that the effect of predictors is on the hazard or likelihood of an event, while the AFT model assumes that the effect of predictors is on the time to the event itself. PH and AFT models can be fit using both parametric and semi-parametric estimation. In parametric models, a probability distribution of the hazard function and survival time, respectively, is assumed while the relationship between predictors and event is estimated. The most common PH model used, the Cox PH model, is semi-parametric and does not assume a distribution of the survival times (Cox, 1975). The AFT model can also be estimated via parametric or semi-parametric means (Kalbfleisch & Prentice, 2002). To allow for the most general case, we choose to focus on the semi-parametric forms of both models.

With EMA data, within an individual, there are many observations of a predictor (from momentary reports) and many observations of times to an event (time from a momentary report to the event of interest), and these observations should be assumed correlated. Both PH and AFT models can incorporate correlated data and thus can be used to analyze repeated observations collected from the same individual. Therefore, PH and AFT models can be adapted for the analysis of EMA data.

Depressed Adolescents and Sexual Risk Behaviors

Depressed adolescents have an increased risk of sexually transmitted infections relative to their healthy peers (Shrier, Harris, & Beardslee, 2002). It is hypothesized that depressed adolescents, who are likely to have decreased positive and increased negative affective states, may use sex to regulate their affective states. There is evidence that both non-depressed (Shrier, Koren, Aneja, & de Moor, 2010) and depressed (Shrier et al., 2012) adolescents experience higher positive affect (PA) and lower negative affect (NA) in the hours immediately before and after a sex event. This pericoital period includes time shortly before sex, when affect may be influenced by anticipation of a planned sex event (Sewell, 2005). Prior to the pericoital period, depressed adolescents experience lower PA in the period of time 6 to 30 hours before a sex event (Shrier et al., 2012). This pattern of lower PA prior to the pericoital period followed by the improved affect that accompanies a sex event is consistent with a hypothesis of affect regulation.

A study of sexual behavior in depressed adolescents was carried out. The goal of the study was to inform interventions to prevent acquisition of sexually transmitted infections in this population. The study implemented EMA methodology to collect in-the-moment affective states and record information about sex events. The resultant data contain many observations from each individual at random intervals over time. These data have been analyzed to determine the temporal association of PA and NA with the occurrence of sex events using an

analysis that aggregated momentary reports of affect into blocks of time preceding and more distant from a sex event (Shrier et al., 2012). The analysis found lower mean PA values preceding sexual intercourse compared to PA values at times more distant from sexual intercourse indicating an association between low PA and sex. The current study re-analyzes these data using time-to-event analysis techniques to compare the results and interpretation from each model, thus demonstrating the application of correlated time-to-event models. The goal of the current study was to propose time-to-event analysis techniques that do not require aggregating momentary reports and may correspond more closely to research questions concerning temporality of association. The proposed techniques may be of use to researchers analyzing EMA data.

Methods

In the current study, we fit two PH models and one AFT model to the data from the depressed adolescent study and compare the results among these models and to those in the original time-block summary analysis. The first PH model included a random individual effect in the model to account for correlation in the EMA data. The second PH model estimated robust standard errors to account for the correlated nature of the data. The AFT model assumes a correlated error structure. For each model, the goal was to estimate the association of PA and NA with time to a subsequent sexual intercourse event. PA and NA were modeled separately; and, for each model type, an unadjusted model and a model adjusted for age, gender, time of day of momentary report, and companionship at the time of the momentary report were fit. Covariates were chosen to match the original analysis (Shrier et al., 2012).

Proportional Hazard (PH) Model

A PH model assumes that the effect of predictors is on the hazard function and is proportional across all times. The hazard function is the risk of an event at a given time, t , and is denoted by $h(t)$. The general equation for a PH model is:

$$h(t) = e^{\beta x} h_0(t), \quad [\text{Equation 1}]$$

where x is a predictor variable, and β is the associated regression coefficient. Note that $h_0(t)$ is referred to as the baseline hazard function.

In the current study, the event of interest was sexual intercourse. When exponentiated, the interpretation of the regression coefficient is that of a multiplicative effect of the predictor on the hazard of sex. That is, the hazard of sex when $x=0$ is $h_0(t)$, while the hazard of sex when $x=1$ is $e^{\beta} h_0(t)$. At any given time, the hazard of sex is e^{β} times higher when $x=1$ than when $x=0$. So, if e^{β} is less than 1 ($\beta < 0$), a lower hazard of sex is evident when $x=1$ compared to when $x=0$. This quantity e^{β} is called a hazard ratio (HR) as it gives the ratio of hazards for differing values of the predictor. For a continuous predictor variable, such as PA or NA, the comparison is between hazard of sex for a unit difference in affect rather than between groups ($x=0$ vs. $x=1$). Note that the “baseline” in baseline hazard refers to the baseline value

of the predictor ($x=0$ for a binary variable), not the time from study enrollment or start time of event.

Given the proportional hazards assumption of the model, it is important to check the assumption in the data under study. This can be done by including an interaction between the predictor of interest and time, to assess whether the effect of the predictor varies by time (Collett, 1994). If the interaction term is not significant, this indicates that the proportional hazards assumption is not violated and use of the PH model may be appropriate. If the interaction term is significant, the proportional hazards assumption is violated and the PH model is not appropriate. Thus, inferences obtained from the model are not valid.

Accelerated Failure Time (AFT) Models

In contrast to a PH model, an AFT model assumes the effect of predictors is on the time to the event rather than on the hazard of an event. The general model equation is: $T = e^{\alpha x} T_0$, where x is a predictor variable, and α is the associated regression coefficient. T_0 is referred to as the baseline time to event. Here again, “baseline” refers to baseline value of the predictor.

The interpretation of regression coefficients is the effect of a predictor on the time to event, usually given as the mean or median time to event. That is, the median time to sex when $x=0$ is T_0 , while the median time to sex when $x=1$ is $e^{\alpha} T_0$. Most often, the model is written out as a linear equation by using a natural log transformation. The equation then becomes:

$$\log T = \alpha x + \log T_0.$$

The $\log T_0$ is an error term. That is, the baseline time to event is assumed to vary randomly according to a statistical distribution. In accordance with statistical notation conventions, this term is usually denoted with an ε , and the equation can be written more familiarly as:

$$\log T = \alpha x + \varepsilon.$$

To compare these two models, it is useful to write both in the same format. The AFT model can also be written equivalently as a function of the hazard function rather than the times to event. The AFT model expressed as a hazard function is:

$$h(t) = e^{-\alpha x} h_0(e^{-\alpha x} t). \quad [\text{Equation 2}]$$

By comparing equation [1] to equation [2], the difference between the PH and AFT models is more evident. Whereas the PH model assumes the effect of a predictor is multiplicative on the baseline ($x=0$) hazard function, the AFT model assumes the effect of a predictor is multiplicative on the baseline ($x=0$) hazard function *evaluated at a time that is also influenced by the predictor*. In this way, the effect of a predictor differs across time, and is

no longer proportional between groups across all times or between levels of a continuous predictor variable across all times.

Distributional Assumptions

Both the PH model and the AFT model can be fit assuming a probability distribution for the time to an event. In both models, though, there are methods available to estimate the parameters without assuming a distribution (Cox, 1975; Jin, Lin, & Ying, 2006) and those are the focus of the current work as they are applicable to all time to event data whether or not the underlying distribution is known. The PH model is estimated via partial maximum likelihood (Cox, 1975). Details of this estimation are beyond the scope of this paper, but a brief description is given here (for more details see Collett, 1994). Parameters are estimated to maximize the likelihood function of the observed data. Specifically, the observed event times are used to construct risk sets of individuals at risk of the event at each of the observed event times. The predictor values of the individuals in the risk set are included in the model-based hazard functions contained in the likelihood to estimate the effect of predictors on the hazard of an event at any given time. In this process, the observed event times are used only to create the risk sets of individuals, and the values of the event times do not directly contribute to the likelihood function. Therefore, the risk of an event rather than the time to an event is the outcome in the estimation.

The AFT model is similar to a linear regression model that can be estimated via least squares. Jin, Lin, and Ying (2006) have developed a method for using least squares estimation in the presence of censoring. Briefly, the value of the observed event times and Kaplan-Meier estimates of survival probabilities are incorporated into the estimation. As with the PH model, a distribution of the event times is not assumed. For details on this estimation procedure see Jin et al. (2006).

Estimation of Effects in Correlated Time-to-event Data

Accounting for the correlation of repeated observations from the same individual can be incorporated into time-to-event models in different ways. We examine two PH models that account for correlated observations in two different ways. The first way is to include a random individual intercept into the model called a “shared frailty” term. This is akin to including a random intercept term in clustered or longitudinal data in a linear model. This model can be fit using the PHREG procedure in SAS and implementing the RANDOM statement. The PH model with the shared frailty term equation is:

$$h(t) = e^{(\beta x + \gamma_i)} h_0(t),$$

where γ_i is the random individual effect or shared frailty term and the i indexes the individual. Similar to linear mixed effects models, the inclusion of the random intercept term results in observations from the same individual being assumed correlated and the variability of this individual-level effect is estimated.

The second PH model we examine accounts for the correlated nature of the data by using a robust standard error estimate to accurately estimate the precision and statistical significance of regression parameters (Lee, Wei, & Amato, 1992; Wei, Lin, & Weissfeld, 1989). This model can be fit using the PHREG procedure in SAS and implementing the COVS(AGGRETATE) option. In contrast to the shared frailty model where an individual-level effect is specified in the model, here a robust standard error is used when making inference about predictors to ensure that the variances are appropriately estimated. This method was originally devised for modeling times-to-event data when the number of correlated observations within an individual is small relative to the number of individuals, but has been successfully applied when the number of correlated observations is larger (Schizophrenia Study Example in Lin, 1994).

The AFT model accounts for correlation among repeated observations from the same individual by assuming a compound symmetric correlation structure between these observations (Jin et al, 2006). This procedure was also devised for relatively few correlated observations within an individual. This model can be fit using an adapted version of the R function, `lss` (Huang & Jin, 2007) which was developed to fit the semi-parametric AFT models using least squares estimation. In the current work, the `lss` function has been adapted to allow for correlated errors among repeated observations from the same individual. The correlation between a set of times to event within individual is assumed to be:

$$\begin{bmatrix} 1 & \rho & \cdots & \rho \\ \rho & 1 & \rho & \rho \\ \vdots & & \ddots & \vdots \\ \rho & \cdots & \rho & 1 \end{bmatrix}$$

Where ρ is the correlation between any two times to event within an individual and the dimension of the correlation matrix is determined by the number of times to event recorded for a given individual. We note this is the same correlation structure obtained when an individual-level random effect is included in a linear model. Sample syntax for fitting each type of model is provided in the Appendix.

Study of Depressed Adolescents Using EMA

Fifty-one adolescents and young adults aged 15 through 22 (mean age 18) were enrolled in the study. Inclusion criteria included clinically significant depressive symptoms defined as a score of ≥ 16 on the Beck Depression Inventory-II (Beck, Steer, & Brown, 1996). Handheld computers were programmed to signal randomly 4-6 times per day for 2 weeks. At each signal, participants were instructed to fill out a questionnaire including questions about momentary affective states, momentary social context (location and companionship), and recent sexual behavior (whether or not sex had occurred, date and time of sex, partner type, and condom use). Affective state was measured separately as PA and NA via an adapted version of the Positive Affect-Negative Affect Schedule (Watson, Clark, & Tellegen, 1998). PA score was computed as the sum of 6 items identifying PA states (strong, proud, alert,

interested, determined, and inspired) and participants were asked the degree to which they felt each state, scored 1-5 for a total score of 5-30. NA score was computed in a similar way using 6 NA states (distressed, upset, guilty, scared, hostile, and irritable).

Prior Analysis of Affect and Sexual Intercourse

The temporal relationship between affective state and sexual intercourse was previously described for this study using a time-block analysis (Shrier et al, 2012). A detailed account of the analytical methods and results of this study can be found in the Shrier et al. paper, but we briefly describe them here. There is a known improvement in affect (increased PA, decreased NA) from approximately 6 hours prior to a sex event, thought to be due to anticipatory effects of knowing a sex event is about to take place, and continuing until approximately 6 hours after sex (Shrier et al., 2012; Shrier et al., 2010). Because the phenomenon currently under study is that of using sex to regulate affective states, we excluded reports in this pericoital period from analysis and research questions examined affect in the hours prior to this pericoital period compared to hours more distant from sex. Therefore, momentary reports within a 6 to 30 hour window (24 hours in duration) prior to a sex event were summarized to obtain a mean PA and a mean NA among all reports in the window. All reports in windows of time more distant (at least 6 hours after and more than 36 hours before) from a sex event were also summarized to obtain a mean PA and mean NA among these reports. After excluding all reports in the pericoital period, the summary and block creation was accomplished by classifying remaining reports as either being prior to a sex event (defined as being within the 6 to 30 hour window prior to a sex event) or distant from a sex event (defined as being more than 6 hours after a sex event and more than 30 hours before a sex event). The mean affect score on all reports coded as prior to a particular sex event is calculated and this value is considered the mean affect in that prior block. Likewise, the mean affect score on all reports coded as distant from sex that occur between two particular sex events is calculated and this value is considered the mean affect in that distant from sex block. Each individual's data consists of multiple prior and multiple distant blocks each with a mean PA and a mean NA value. Figure 1 illustrates how blocks were created from an individual's EMA data. The mean affect in the prior blocks was then compared to the mean affect in the distant blocks of time to evaluate the temporal relationship between affective states and sexual intercourse. In the Shrier et al. study, a logistic generalized estimating equation (GEE) model was fit with block-type as the outcome and mean affect as the predictor. An odds ratio (OR) for the likelihood of a block to be prior to a sex event was therefore obtained. The results were interpreted as a predictor's influence on the likelihood of a sex event initiation. A significantly lower mean PA was observed prior to sexual intercourse while no significant association between mean NA and sexual intercourse was observed. Blocks of time with a unit higher mean PA were less likely to be in the 24 hours prior to a sex event initiation, with an adjusted OR of 0.87 (95% CI: 0.80-0.95). Models were adjusted for age, gender, time of day, and companionship (boy/girlfriend, family, other, or alone) at the time of the momentary signal.

Current Analysis of Affect and Sexual Intercourse

As previously mentioned, the outcome of interest in the current report is time from a momentary report to the start of the next sex event. The predictor of interest is the

momentary PA and NA collected on each report. If a momentary affective state report is followed by a sex event, the outcome of interest is observed and is the time from the last report to the time of sex. If, on the other hand, a momentary affective state report is followed by another momentary affective state report without an intervening sex event, the outcome of interest is censored at the time of the next momentary report. That is, we observe a certain amount of time during which a sex event did not occur (time from last report to current report), but we do not observe the time that the next sex event occurs following the last report. Each individual's data consists of several censored and observed times to an event, each with a momentary PA and a momentary NA value. Figure 2 illustrates the times to sex for an individual. To model the affect regulation process of interest and mirror analyses done in the original report of this research question, we modeled the time from a report to 6 hours prior to a sex event (the start of the pericoital period) thereby excluding the pericoital period known to have altered affective states. We were therefore modeling the time to sex event initiation rather than sex event with the goal of eliminating the anticipatory affective process. For ease of presentation, we refer to this as a 'sex event' hereafter, and note this modeling decision is specific to the current analysis and is not required for performing any of the models described. Unadjusted as well as adjusted models using the covariates used in the time-block analysis were fit.

Results

Among 51 participants, there were 2270 momentary affective state reports, 279 of which were followed directly by a sex event (time to sex observed), and 1991 that were followed by another affective state report (time to sex censored). There were an average of 5.9 sex events per individual (range 1-13), and 59.8 momentary reports per individual (range 15-168).

Tests of the proportional hazards assumption showed that the effect of NA on time to sex violated this assumption, with a significant time-by-NA interaction ($p=0.02$). We therefore do not present results for PH models examining NA and sex. The proportional hazards assumption did not appear to be violated for the relationship between PA and time to sex ($p=0.24$), so the results of these PH models are described below.

PH with Frailty term

Results from all models are contained in Table 1. The PH models including a shared frailty term with risk of sex as the outcome and PA as the predictor indicated that increased PA is significantly associated with decreased risk of sex. Specifically, a unit increase in PA is associated with 0.98 ($=e^{-0.021}$) times the risk of sex at any given time ($p=0.043$). The direction and magnitude of the effects stays the same when adjusting covariates are included in the model, but the PA result becomes non-significant ($HR=0.98$, $p=0.092$).

PH with Robust Standard Errors

The results from the PH models using a robust standard error estimate are similar to those from the shared frailty model. The unadjusted and adjusted PA HRs are both 0.98 with only the unadjusted model being significant ($p=0.035$ vs. $p=0.072$, respectively).

Least Squares AFT

The unadjusted AFT model with time to sex as the outcome and PA as the predictor indicates that time to sex is longer following a higher value of PA. Specifically, in the unadjusted model, a unit increase in PA is associated with a time to sex that is 1.03 ($=e^{0.0321}$) times longer. This association is statistically significant with $p=0.003$. When adjusting covariates are included in the model, the association remains of the same magnitude (1.03) and remains significant ($p=0.025$). Results from the NA model are in the opposite direction. Increased NA is associated with shorter time to sex; a unit increase in NA is associated with a time to sex that is 0.98 times as long in the unadjusted model, although this association does not reach statistical significance in either unadjusted or adjusted models ($p=0.79$ and $p=0.37$, respectively).

Table 2 contains the adjusted PA model results from the original time-block and each additional analysis along with an explanation of how the OR, HR, or time ratio is interpreted. A description in terms of risk of sex has been provided for non-significant HRs from both of the proportional hazard models in order to describe how a HR from these models would be interpreted. We note, however, that due to their lack of significance, in practice these adjusted model results should be interpreted as having no significant association rather than quantifying the size of the risk of sex.

Discussion

There is no standard way of examining the temporal relationship of momentary measurements made on every report and events that are reported multiple times per individual, but interspersed between momentary reports. These data are obtained when examining the relationship between momentary and variable psychological states, such as affect, and behavioral events, such as sexual intercourse. Thus, elaborating on methodological approaches to analyzing these relationships has implications for many areas of research, including affect regulation and health risk behavior. Several possible methods of analysis can be conceived; and, although it was not seen here, it may be the case that different relationships are described between affective state and sex when different models are used to fit the data. In the current report, we explored different ways of examining this unique correlated time-to-event data generated as part of an EMA study.

The results of the current report are consistent with the prior report of the association of PA and sexual intercourse and the lack of significant association between NA and sexual intercourse. The original time-block summary model indicated a decreased likelihood of higher mean PA occurring in the 24 hours prior to a sex event. A temporal interpretation would be that 24-hour periods of higher average PA are less likely to be followed by a sex event. With the time-block approach, we are limited to commenting on the 24-hour period of time chosen for block creation and limited in interpreting the average PA from a set of reports rather than the PA reported on an individual momentary report. Thus, the time-block model may be the preferred method when the effects of momentary experience on a behavioral effect are cumulative or need to be present for a period of time before they affect a behavior. The PH models indicated a decreased hazard of sex with a higher value of PA on a report at any given time. This result can also be interpreted to comment on temporality as a

longer time to sex follows a higher PA measurement. The AFT model indicated increased time from a report of a higher PA and subsequent sexual intercourse. Each of these models gave similar information about the nature of the relationship between affect and sexual intercourse, but the presentation of the results differed. Determining the time to an event following a single momentary observation may be particularly relevant for research questions examining threshold effects. Specifically, regardless of the average of several reports, as long as a single report records a high enough value for the predictor, the event is more likely.

The block method of analysis and the time-to-event methods each have additional benefits and drawbacks. The original time-block model is attractive because it is easily fit and uses readily available and familiar analysis techniques (logistic GEE models) to examine the research questions. However, some of the benefit of momentary assessment is lost when momentary reports are averaged together in blocks of time rather than all contributing directly to the analysis. Also, the results of such an analysis, an odds ratio for block-type, are not as directly interpretable to the research question as other model results.

Rather than averaging across time blocks, the PH models both use all available momentary reports and are easily fit with familiar software analysis techniques (Cox models, the most frequently used time-to-event models, can be fit in all standard statistical software packages). However, as with the logistic GEE, the model result is a hazard ratio and is a less direct measure of the temporal relationship between affect and time to sex. That is, the relative risk of sex rather than the time to sex is summarized in the proportional hazard model. Thus, without further assumptions, an estimate of predicted times to event cannot be obtained. Also, the proportional hazard assumption must be plausible and met in the data under study for this to be a viable analysis technique (i.e., to make valid inferences from the model results). That is, the effect of the predictor on the risk of the event should not vary by time. For example, in examining the association of withdrawal symptoms on time to relapse, it is likely that at earlier times to relapse withdrawal symptoms might be a stronger predictor of time to relapse than at later times. In this case, the effect of withdrawal symptoms on time to relapse would vary over time and violate the proportional hazard assumption making the PH models inappropriate.

The AFT model has the benefit of the result being interpreted as the effect of PA on the time to sex; this is the most direct response to research questions about the temporal association between affect and sex. The model can also easily generate predicted event times where the PH does not. Also, the AFT model does not require a proportional hazards assumption. One disadvantage of the AFT model is that it is less commonly used in time-to-event analyses than Cox models so may be unfamiliar to readers. Also, fitting the semi-parametric version of these models with correlated data requires a special R function (a modified R package (`lss`), available from the authors) and is not yet incorporated into standard statistical software packages. In addition, the error structure in this implementation is limited to compound symmetry and the model does not accommodate random effects. Finally, there were some logistical challenges to fitting the AFT models with these data. Specifically, the time and computer memory required to compute the AFT model estimates was substantial. With an Intel® W3580 computer @ 3.33GHz with 24GB of memory using the 64-bit R

program, the unadjusted models took about 25 minutes to run, whereas the models including the covariates took over 3 hours to run. With a less powerful computer (HP Compaq Elite 8300 with Intel® Core™ i5-3570 CPU @ 3.40GHz and 4GB memory using the 32-bit R program), the models could not even be fit as the memory was exceeded. Desire for the direct interpretation of the effect of predictors on the time to an event (and subsequent event time prediction) and violation of the proportional hazard assumption are the most compelling reasons to consider this model, but must be considered in light of the logistical challenges of fitting this model.

Despite the utility of time to event models for the analysis of EMA data, examples in the literature are scarce. One exception is an analysis of interpersonal problems and negative mood as predictors of time of drinking (Todd, Armeli, & Tennen, 2009). This analysis, which employed a parametric time-to-event frailty model, presents an excellent example of when time-to-event methods are particularly useful over the usual multilevel models used for EMA data. The time of day of drinking initiation is an important outcome as early-day drinking is more problematic than drinking later in the day. However, time of initiation is censored on days on which drinking does not occur. Multilevel models, which do not accommodate censored data, are not appropriate for analysis of this data and time-to event models become necessary. Another notable example of a time-to-event model for the analysis of EMA data is the development of a parametric model for examining the effect of time-varying covariates on smoking relapse (Rathbun, Song, Neustifter, & Shiffman, 2013). In contrast to the models presented in the current application, the Rathbun et al. model is parametric and the EMA data are used for updating time-varying covariates in an analysis of time to smoking relapse rather than to create multiple event times per individual. Other instances where time-to-event models may be useful are cases in which a value of a predictor collected on a momentary report may be used as an intervention point should it indicate significant future risk of an event. For example, suppose certain values of self-efficacy indicate significantly shortened time to lapse from abstinent behavior. In this case, momentary reports could be monitored for concerning self-efficacy values. Essentially, any research question relating to the risk of an event of interest that is interspersed between regular reports of potential predictors of this risk is particularly suited for a time-to-event model.

The limitations of using multilevel models for block-type analysis discussed earlier include the need for analyzing aggregated values of reports in blocks of time. An alternative use of multilevel models would be to use parametric curves or spline-based models to estimate the shape of momentary values surrounding the event (like that performed in Shrier et al., 2012). The curve or spline-based models would serve a different purpose than time-to-event models in this case and thus are limited for making certain inferences. Estimates from these models consist of average predictor values on EMA reports at various periods of time before and after an event and no specific quantification of risk or likelihood of an event is obtained.

Since properties of EMA data may differ from the typical time-to-event data analyzed with these techniques (e.g., clinical trials with few observations per individual), more research on the performance of each of these models with EMA data is warranted. In particular, performance of models in the case of large amounts of censoring (if events are rarely

observed between momentary reports) or in cases of large numbers of observations per individual should be evaluated. This may be done using simulated data with various known relationships and correlations among observations to evaluate whether each technique adequately models the relationships under study. Overall, each of these statistical techniques may be considered for the analysis of time-to-event data collected via EMA, with the choice of method being determined by the framing of the research questions and the balance of specific challenges and advantages of each method.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

References

- Beck, A., Steer, R., Brown, G. BDI-II manual. San Antonio, TX: The Psychological Corporation, Harcourt Brace & Company; 1996.
- Collett, D. Modelling Survival Data in Medical Research. Boca Raton, FL: Chapman & Hall; 1994.
- Cox D. Regression Models and Life-Tables. *Journal of the Royal Statistical Society B*. 1972; 34:187–220.
- Cox D. Partial Likelihood. *Biometrika*. 1975; 62:269–276.
- Hedeker, D., Mermelstein, R. Modeling Contextual Effects in Longitudinal Studies. Mahwah, NJ: Erlbaum; 2007. Mixed-Effects Regression Models with Heterogeneous Variance: Analyzing Ecological Momentary Assessment (EMA) Data of Smoking.
- Hedeker D, Mermelstein R. An Application of a Mixed-Effects Location Scale Model for Analysis of Ecological Momentary Assessment (EMA) Data. *Biometrics*. 2008; 64:627–634. [PubMed: 17970819]
- Huang L, Jin Z. LSS: An S-Plus/R Program for the Accelerated Failure Time Model to Right-Censored Data Based on Least-Squares Principle. *Computer Methods and Programs in Biomedicine*. 2007; 86(45-50)
- Jin Z, Lin DY, Ying Z. On Least-Squares Regression with Censored Data. *Biometrika*. 2006; 91(1): 147–161.
- Kalbfleisch, JD., Prentice, RL. The Statistical Analysis of Failure Time Data. Hoboken, NJ: Joh Wiley & Sons, Inc; 2002.
- Lee, E., Wei, L., Amato, D. Cox-type Regression Analysis for Large Numbers of Small Groups of Correlated Failure Time Observations. In: J. K., Goel, P., editors. *Survival Analysis: State of the Art*. Dordrecht: Kluwer Academic Publishers; 1992. p. 237-247.
- Lin DY. Cox Regression Analysis of Multivariate Failure Time Data: The Marginal Approach. *Statistics in Medicine*. 1994; 13:2233–2247. [PubMed: 7846422]
- Rathbun SL, Song X, Neustifter B, Shiffman S. Survival analysis with time-varying covariates measured at random times by design. *Applied Statistics*. 2013; 62:419–434. [PubMed: 23729866]
- Schwartz J, Stone AA. Strategies for analyzing ecological momentary assessment data. *Health Psychology*. 1998; 17(1):6–16. [PubMed: 9459065]
- Sewell KW. The experience cycle and the sexual response cycle: conceptualization and application of sexual dysfunction. *Journal of Constructivist Psychology*. 2005; 18(1):3–13. DOI: 10.1080/10720530590522973
- Shrier LA, Feldman H, Black S, Walls C, Kendall A, Lops C, Beardslee W. Momentary Affective States Surrounding Sexual Intercourse in Depressed Adolescents and Young Adults. *Archives of Sexual Behavior*. 2012; 41:1161–1171. DOI: 10.1007/s10508-011-9787-4 [PubMed: 21755382]
- Shrier LA, Harris SK, Beardslee WR. Temporal associations between depressive symptoms and self-reported sexually transmitted disease among adolescents. *Archives of Pediatrics and Adolescent Medicine*. 2002; 156:599–606. [PubMed: 12038894]

- Shrier LA, Koren S, Aneja P, de Moor C. Affect regulation, social context, and sexual intercourse in adolescents. *Archives of Sexual Behavior*. 2010; 39(3):695–705. [PubMed: 18818996]
- Stone A, Shiffman S. Capturing momentary, self-report data: a proposal for reporting guidelines. *Annals of Behavioral Medicine*. 2002; 24(3):236–243. [PubMed: 12173681]
- Stone, A., Shiffman, S., Atienza, A., Nebeling, L. *The Science of Real-Time Data Capture*. New York, New York: Oxford University Press; 2007.
- Todd M, Armeli S, Tennen H. Interpersonal Problems and Negative Mood as Predictors of Within-day Time to Drinking. *Psychology of Addictive Behaviors*. 2009; 23(2):205–215. [PubMed: 19586137]
- Watson D, Clark L, Tellegen A. Development and Validation of Brief Measures of Positive and Negative Affect: The PANAS Scales. *Journal of Personality and Social Psychology*. 1998; 54:1063–1070. [PubMed: 3397865]
- Wei L. The Accelerated Failure Time Model: A Useful Alternative to the Cox Regression Model in Survival Analysis. *Statistics in Medicine*. 1992; 11(14-15):1871–1879. [PubMed: 1480879]
- Wei L, Lin D, Weissfeld L. Regression Analysis of Multivariate Incomplete Failure Time Data by Modeling Marginal Distributions. *Journal of the American Statistical Association*. 1989; 84(408): 1065–1073.

Figure 1.

Affect block data collected from an individual. In this example, six momentary reports and 2 sex events are used to create two prior-to-sex blocks and one distant-from-sex block. Mean positive affect and mean negative affect is calculated for each block as the mean of the affect score on all reports contained in the block.

Figure 2.

Time-to-event data collected from an individual. In this example, six momentary reports and 2 sex events are used to create five time-to-event observations; three are censored and two are not.

Table 1
Comparison of results from unadjusted and adjusted* proportional hazards and accelerated failure time model of affective state and time to sexual intercourse

	Proportional Hazard Model with Frailty Term		Proportional Hazard Model with Robust SEs		Accelerated Failure Time Model	
	β (SE)	p-value	β (SE)	p-value	α (SE)	p-value
PA	-0.021 (0.010)	0.043	-0.021 (0.010)	0.035	0.032 (0.011)	0.003
PA Adjusted	-0.019 (0.011)	0.092	-0.019 (0.011)	0.072	0.028 (0.013)	0.025
NA					-0.002 (0.009)	0.79
NA Adjusted					-0.009 (0.013)	0.37

* Models adjusted for gender, age, time of day, and companionship.

Table 2
Association between positive affective state and sexual intercourse from four models and the associated model interpretation

Model	OR (CI), HR (CI) or time ratio (CI)	p-value	Interpretation
Time-block Analysis (OR)	0.87 (0.80-0.95)	0.001	A 24-hour period with a unit higher average PA is 0.87 times as likely to be followed by sex
Proportional Hazard Model with Frailty Term (HR)	0.98 (0.96-1.003)	0.092	A unit higher PA is associated with 0.98 times the risk of sex at any given time (non-significant)
Proportional Hazard Model with Robust SEs (HR)	0.98 (0.96-1.003)	0.072	A unit higher PA is associated with 0.98 times the risk of sex at any given time (non-significant)
Accelerated Failure Time Model (time ratio)	1.03 (1.003-1.05)	0.025	A unit increase in PA is associated with a time to sex that is 1.03 times longer

OR: Odds ratio

HR: Hazard ratio

CI: 95% confidence interval