AMSI HIGHER ED

A comparison of linear mixed models that include time-varying covariates

Alex Ligthart-Smith
Supervised by:
Dr David Farchione; A/Prof Luke Prendergast
La Trobe University

Australian Government
Department of Education and Training

27th February, 2016

Abstract

Longitudinal studies are becoming more prevalent, with new technologies providing us with abilities to collect, store, and analyse more data than ever. With the wealth of new types of information available to researchers and analysts, new approached to analysing and modelling these kinds of data must be developed. One area where researchers often lack the tools to make the most of their data is where time-varying covariates are present in linear mixed models. This paper compares two approaches to this task, one that involves person-mean centring and the other detrending of time-varying covariates, and shows that the person-mean centring approach provides more consistent and precise estimates of the within-effects of these covariates.

1 Introduction

Longitudinal studies are a useful approach to research in a variety of fields, in particular in epidemiology and psychopathology. Linear mixed effects (LME) models are a useful way to make use of the wealth of information available in such data, and are used to great effect by many researchers. However, when it is not just the response variable, but also one or more of the covariates, which vary over time, basic LME models can fall short in disaggregating the within-effects and between-effects (Curran and Bauer, 2011, pp. 584-586).

For the purpose of illustration, imagine you are collecting data on how high kittens can jump, and their weight, over time. If we took a naive approach to modelling this data, we might use the weight of the kittens as a covariate in a standard LME model. However, this approach would not be able to separate the effect of the individually changing weight of each kitten as it ages (the within-effect), from the effect of the comparative weights of kittens in relation to each other (the between-effect), on how high the kittens can jump.

In order to disaggregate these effects, we need to introduce extra terms into the LME to account for the effect of time on the time-varying covariate (TVC). One approach to this problem, what we will call the traditional approach, was introduced by Rabe-Hesketh (Rabe-Hesketh and Skrondall, 2015, pp. 114-122), and involves subject-mean centring on the TVC. A second approach, introduced by Curran and Bauer (2011), in an effort to correct errors in estimation that they had found using the traditional approach to model a simulated data set, detrends the TVC. The purpose of this paper is to compare the performance of these two approaches, particularly in how well they estimate the within-effects, by analysing variability of estimates, and coverage probabilities and expected lengths of 95% confidence intervals.

2 Time-Varying Covariates

The effects of time on time-varying covariates can be modelled using random effects models, where x_{ti} is the value of the covariate, x, for individual i at time t:

$$x_{ti} = \gamma_0 + \gamma_1 t + \mu_{0i} + \mu_{1i} t + r_{ti} \tag{1}$$

where γ_0 and γ_1 are the mean intercept and slope for the population respectively, μ_0 and μ_1 are the random adjustments to the intercept and slope for individual i, and r_{ti} is the random error term.

If we were to use the naive approach, we would take the values of x_{ti} and use these as a simple covariate in an LME model, similar to that used to model x_{ti} in (1), but for our purposes without the random slope component:

$$y_{ti} = \beta_0 + \beta_1(x_{ti}) + \zeta_i + \epsilon_{ti} \tag{2}$$

where y_{ti} is the value of the response variable at time t for individual i, β_0 is the population intercept, β_1 is the effect of the covariate, ζ_i is the random adjustment to the intercept, and ϵ_{ti} is the random error term.

By including only one effect on the covariate, this approach is unable to disaggregate the within- and between-effects of the TVC. The two methods for disaggregating these effects are described in the next section.

3 Two Approaches to Modelling with TVCs

3.1 The Traditional Approach

The traditional method for disaggregating within- and between- effects, introduced by Rabe-Hesketh (Rabe-Hesketh and Skrondall, 2015, pp. 114-122), and similar to the correlated random effects approach described by Woodlridge (2013, pp. 497-499), involves subject-mean centring. The within-effect, β_1, is modelled on the difference of the time-specific TVC measurement, x_{ti}, from the subject mean, $\bar{x}_i$, and the between-effect, β_2, on the subject-mean:

$$y_{ti} = \beta_0 + \beta_1(x_{ti} - \bar{x}_i) + \beta_2\bar{x}_i + \zeta_i + \epsilon_{ti} \tag{3}$$

In a simulation, Curran and Bauer (2011) found that this method produced a poor estimate of the within-effect on a TVC, concluding that that approach was not ideal for modelling longitudinal data with time-varying covariates. Hence, they developed a similar approach, with a correction on the subject-mean centring of the traditional approach.

3.2 The Curran and Bauer Approach

The approach developed by Curran and Bauer involves *detrending* the TVC. This is achieved by first running a linear regression on the TVC itself, to provide a regression line:

$$\hat{x}_{ti} = \alpha_{0i} + \alpha_{1i}t \tag{4}$$

where α_{0i} is the intercept of the TVC for individual i and α_{1i} is its linear slope with respect to time, t.

Curran and Bauer detrend the TVC by taking the difference of measurements from the linear regression line of the TVC, rather than from the subject-mean:

$$y_{ti} = \beta_0 + \beta_1(x_{ti} - \hat{x}_{ti}) + \beta_2\alpha_{0i} + \zeta_i + \epsilon_{ti} \tag{5}$$

where $\hat{x}_{ti}$ and α_{0i} are those values found in (4).

4 Comparing the Approaches

In order to compare the two approaches, we simulated longitudinal data including time-varying covariates using the same method and parameters as were used by Curran and Bauer (2011, p. 605). The number of data points were based on initial values of $n = 500$ and $t = 9$, with time taking integer values with $\bar{t} = 0$. The TVC was generated using formula (1), with baseline values $\gamma_0 = 25$, $\gamma_1 = 1$, μ_0 and r_{ti} normally distributed with mean $= 0$ and $\sigma^2 = 1$, and μ_1 normally distributed with mean $= 0$ and $\sigma^2 = 4$. All simulations were repeated over 1000 trials. y_{ti} values were then generated using these x_{ti} values substituted into model (3), with $\beta_0 = 25$, $\beta_1 = -1$, $\beta_2 = 1.5$, and error terms normally distributed with $\sigma^2 = 1$.

We carried out initial exploration by varying all γ and β parameters one by one to see how this would affect the within- and between-effect estimations of both approaches. This revealed that both produce similar results in between-effect estimation, however there appears to be a marked difference in the variance of within-effect estimates, which becomes larger with an increase in γ_1. Changes to other parameters do not appear to affect the models, apart from an improvement in both with increased number of data points (increases in both n and t improved estimates in both approaches). Based on these findings, we decided that the best way to examine the differences would be to compare the variances of the β_1 estimates, as well as the coverage probabilities and expected confidence interval lengths, of the two approaches.

4.1 Variability of within-effect estimates

The first step in comparing the within-effect estimates of the two approaches was to look at the differences in accuracy and variability of the estimates in simulation. This was carried out over several values of γ_1, the effect of time on the TVC, and over various sample sizes. For a snapshot of these differences, see Figure 1. These boxplots show that the traditional approach provides more accurate mean estimates of the within effect ($\beta_1 = -1$) for small sample sizes, as well as consistently lower variability than the Curran and Bauer approach. It is also apparent that, while the Curran and Bauer approach provides consistent variability over varying values of γ_1, the traditional approach actually provides lower variability as the effect of time on the TVC increases. Both approaches show improvement in estimation with increasing sample size. These results tell us that while both methods provide good mean estimates over a large number of trials, the estimates in each individual trial can vary much more from the true value when using the Curran and Bauer approach as compared to the traditional approach.

4.2 Coverage probabilities and expected lengths

The next step in comparing the two methods was to compare estimates of coverage probabilities and expected lengths of 95% confidence intervals for the two approaches.

We found estimated coverage probabilities by running simulations over 1000 trials, and comparing the number 95% confidence intervals estimates in each method containing the true value of β_1. A good model should include the true value of the parameter in the 95% confidence interval approximately 95% of the time. Figure 2 shows the estimated coverage probabilities for both approaches over a variety of values of both γ_1 and β_1. We can see that the results are similar over changes in γ_1, and for most values of β_1, with the exception of when $\beta_1 = 0$. The traditional approach provides consistent coverage of around 0.95 over all values of the parameters, as would be expected from a good model. However, the Curran and Bauer approach gives coverage of 1.0, except where $\beta_1 = 0$, where it provides coverages closer to 0.95. Of course, when the within-effect is equal to zero there is little difference between the two approaches so it is unsurprising that they give similar results. What these findings tell us is that the Curran and Bauer approach is consistently providing confidence intervals that include the true value of the parameter, which indicates that perhaps the standard error of these estimates is very large, leading to very large, catch-all, interval estimates.

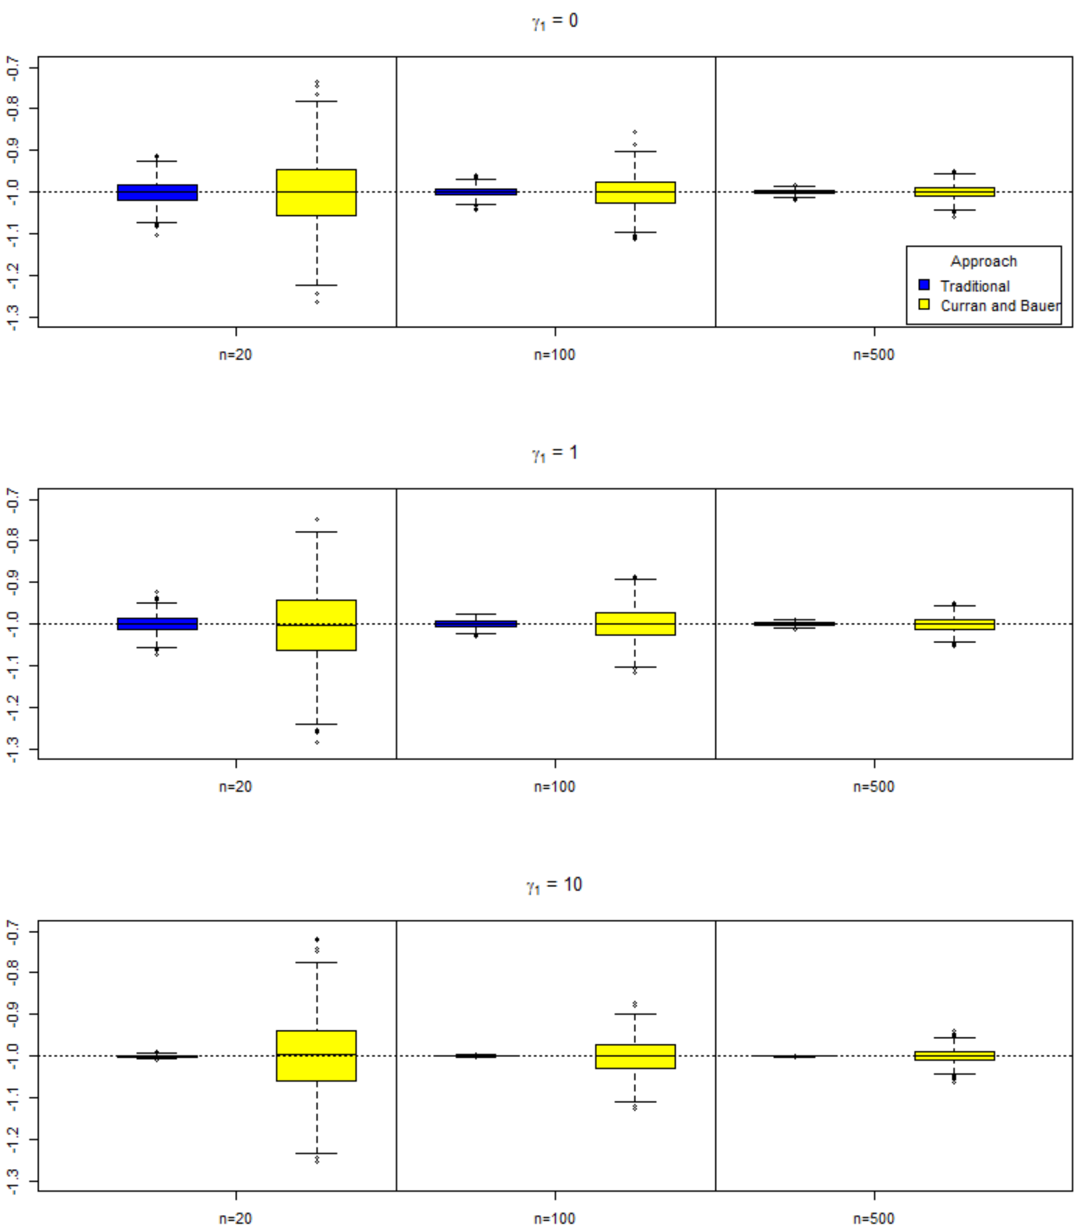

Figure 1: Boxplots of within-effect estimates over varying TVC slopes and sample sizes

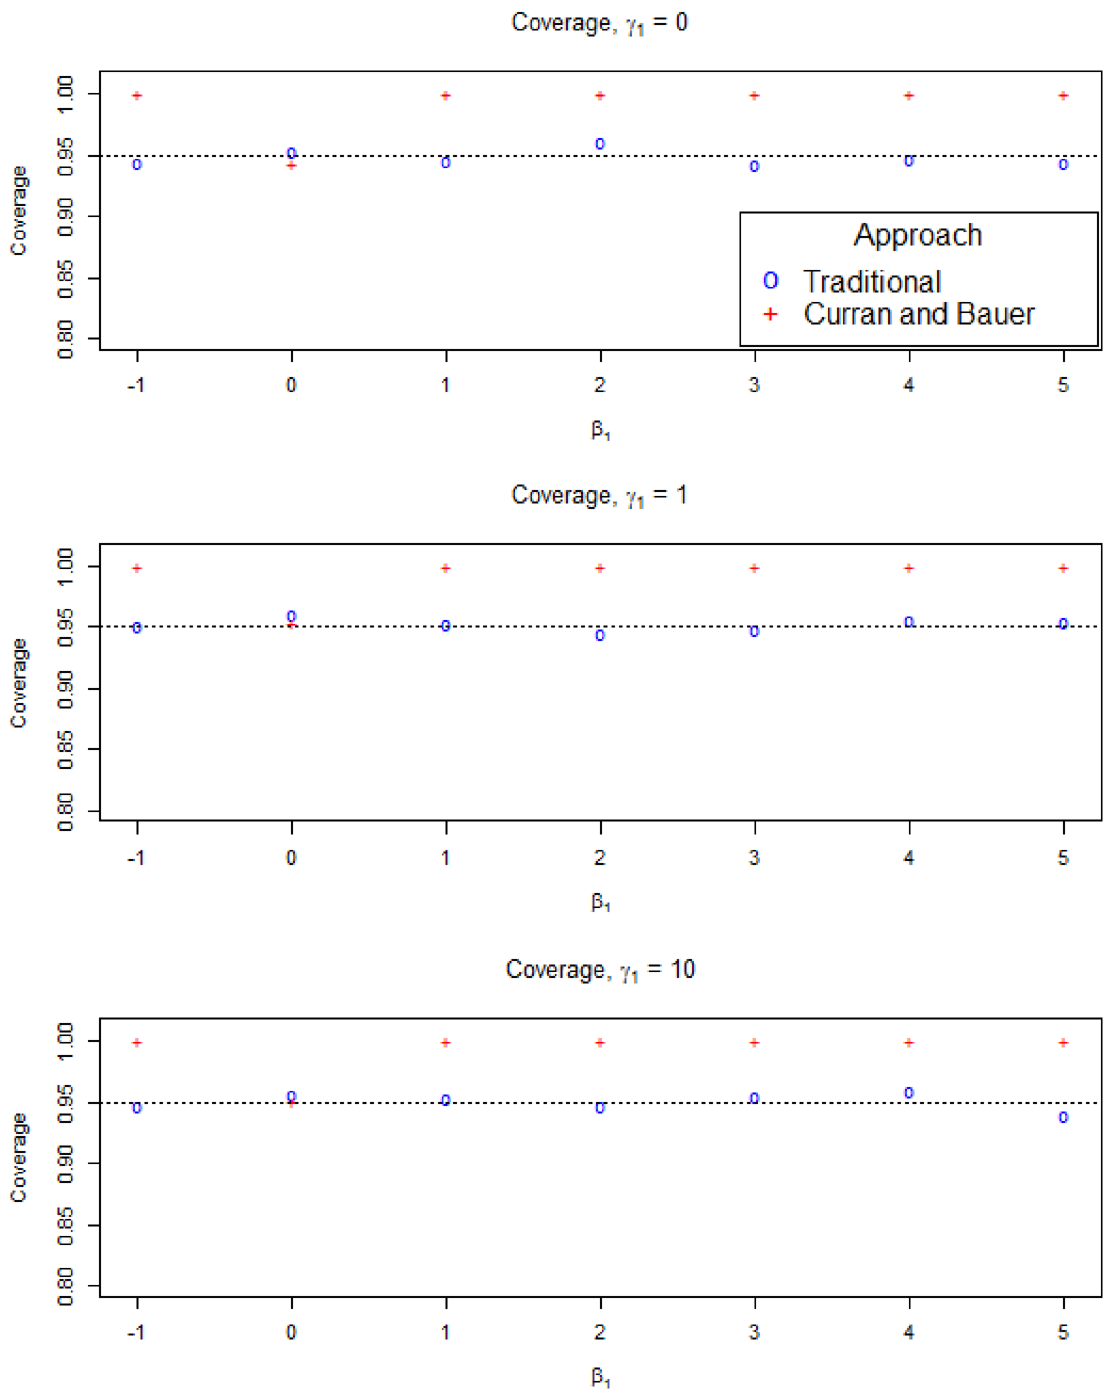

Figure 2: Expected coverage of 95% confidence intervals

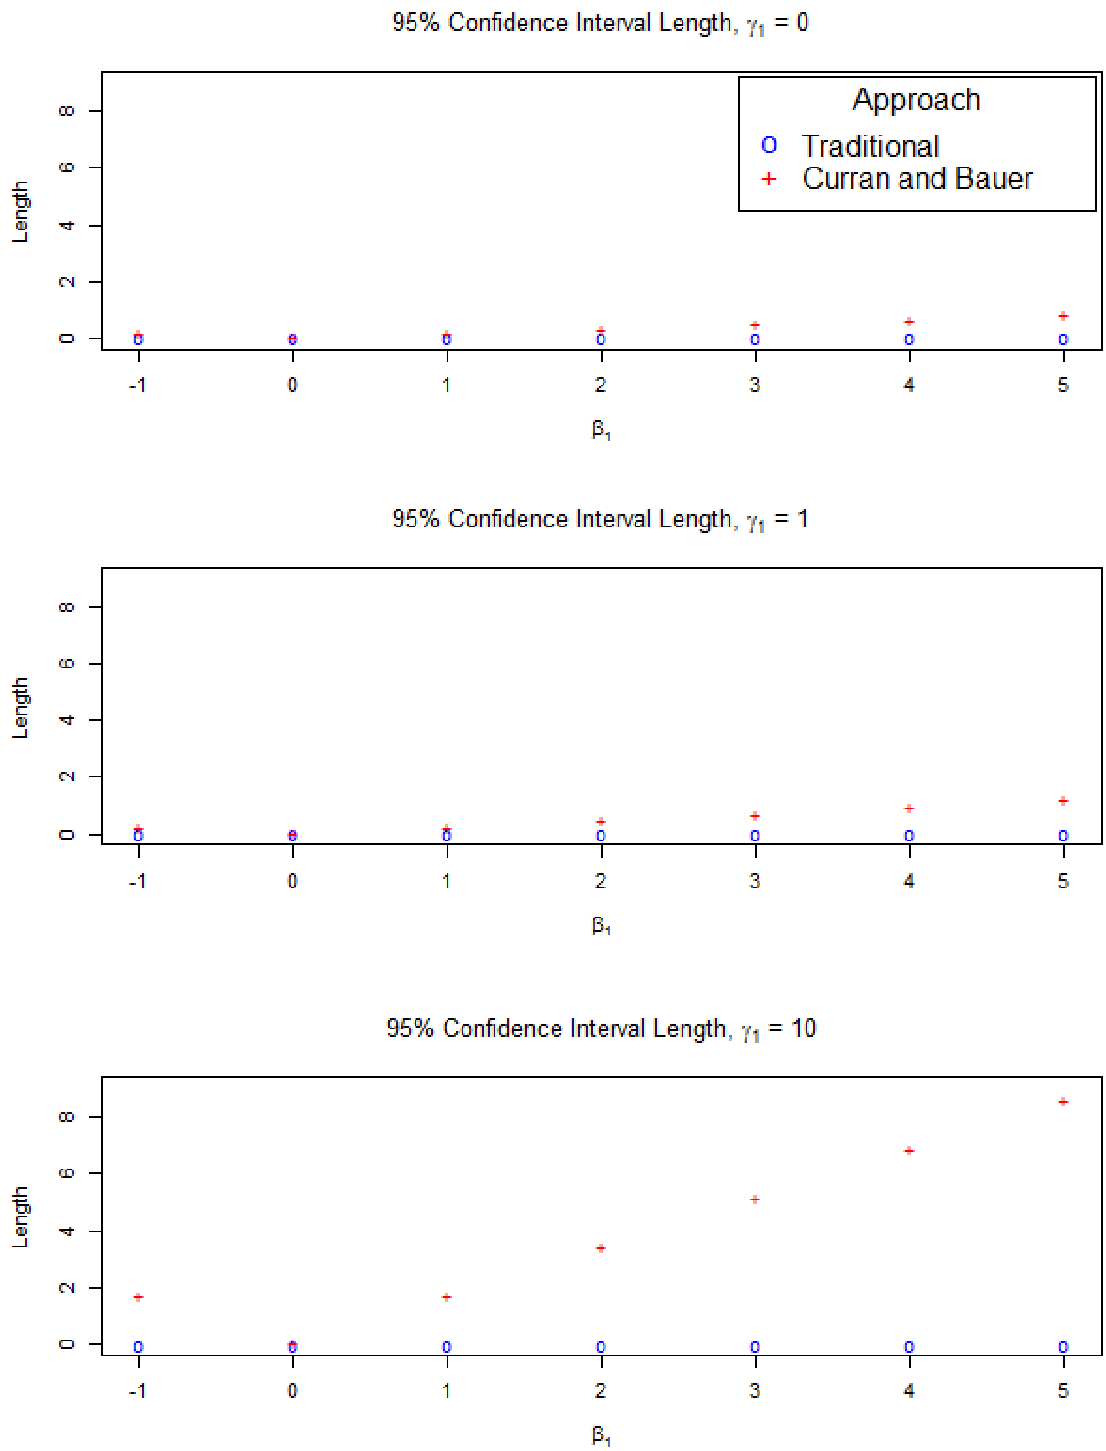

Figure 3: Expected confidence interval lengths of β_1 estimates

In order to further probe this difference, we next compared the expected lengths of the 95% confidence intervals, by taking means over the 1000 trials of the difference between the lower and upper bounds of the estimated intervals. We can see from Figure 3 that while the traditional approach provides consistently small confidence intervals for the estimates of β_1, the Curran and Bauer approach provides larger intervals both as γ_1 and β_1 move away from zero, with an interval as large as 8.6 for estimate of $\beta_1 = 5$ when $\gamma_1 = 10$. These results confirm our instinct from the coverage results that the standard errors of the estimates provided by the Curran and Bauer approach are becoming very large as the within-effects, and the effects of time on the TVC, become larger.

5 Concluding remarks and further questions

While both the traditional, and Curran and Bauer approaches to disagregating within- and between-effects on time-varying covariates in linear mixed models seem to provide consistently good mean estimates of the within-effects in simulations with a large number of trials, comparison of the two approaches indicates that the Curran and Bauer model includes much greater variability in estimation, and larger standard errors of estimates. These findings indicate that the traditional approach is more reliable for estimating the within-effect of time-varying covariates in this kind of data. Further studies should be done to compare the performance of these approaches on more complicated data sets, and to see how the estimates compare when applied to real-life data sets.

Acknowledgments

I would like to thank AMSI for providing me with this fantastic opportunity to experience life as a researcher in statistics, as well as my supervisors Dr. David Farchione and A/Prof Luke Prendergast for offering me this interesting and important project, and for providing support, advice and encouragement over the course of my research. I would also like to thank my best friend and companion Caspian for providing inspiration for my presentation, and for the love and affection he gives me every night after long days at my desk.

References

Curran, P. and Bauer, J. (2011). The disaggregation of within-person and between-person effects in longitudinal models. *Annual Review of Psychology*, 62, pp. 583-619.

Rabe-Hesketh, S. and Skrondal, A. (2015). *Multilevel and Longitudinal Modeling Using Stata*. 2nd ed. College Station: Stata Press.

Wooldridge, J. M. (2013). *Introductory Econometrics: A Modern Approach*. 5th ed. Mason: South-Western Cengage Learning.