



Is it even rainier in *North Vancouver*? A non-parametric rank-based test for semicontinuous longitudinal data

Harlan Campbell

To cite this article: Harlan Campbell (2018): Is it even rainier in *North Vancouver*? A non-parametric rank-based test for semicontinuous longitudinal data, *Journal of Applied Statistics*

To link to this article: <https://doi.org/10.1080/02664763.2018.1536878>



Published online: 13 Nov 2018.



Submit your article to this journal [↗](#)



View Crossmark data [↗](#)



Is it even rainier in *North Vancouver*? A non-parametric rank-based test for semicontinuous longitudinal data

Harlan Campbell 

Department of Statistics Vancouver, University of British Columbia, Vancouver, BC, Canada

ABSTRACT

When the outcome of interest is semicontinuous and collected longitudinally, efficient testing can be difficult. Daily rainfall data is an excellent example which we use to illustrate the various challenges. Even under the simplest scenario, the popular ‘two-part model’, which uses correlated random-effects to account for both the semicontinuous and longitudinal characteristics of the data, often requires prohibitively intensive numerical integration and difficult interpretation. Reducing data to binary (truncating continuous positive values to equal one), while relatively straightforward, leads to a potentially substantial loss in power. We propose an alternative: using a non-parametric rank test recently proposed for joint longitudinal survival data. We investigate the potential benefits of such a test for the analysis of semicontinuous longitudinal data with regards to power and computational feasibility.

ARTICLE HISTORY

Received 24 November 2017
Accepted 11 October 2018

KEYWORDS

Semicontinuous data; rank test; longitudinal; tobit model; resampling methods

There is booming rain, whispery rain, rain that lulls you to sleep, and rain on the leaves which sings you awake; there is soft rain, hard rain, sideways rain, rain that makes you instantly wet, and rain that leaves soft kisses on your cheek, like the wings of a butterfly.

Vancouver In The Rain,
Regan D’Andrade, 1999

1. Introduction

Semicontinuous data is characterized by a mixture of zeros and continuously distributed positive values. Typically, the proportion of observed values equal to zero is substantial and the positive values observed exhibit right-skewness and heteroscedasticity. What’s more, semicontinuous data is distinct by the fact that zeros represent ‘non-occurrences’, rather than simply being the result of truncated or censored negative values. While this paper only considers data in which the outcome variable is semicontinuous, there are related challenges with data in which predictor variables are semicontinuous (also known as ‘spike at zero’ data) [34].

CONTACT Harlan Campbell  harlan.campbell@stat.ubc.ca  Department of Statistics Vancouver, University of British Columbia, Vancouver, BC, Canada V6T 1Z2

Daily rainfall is an excellent example of semicontinuous data. On dry days, rainfall is zero, while on rainy days, rainfall is a positive continuous number of millimetres (mm). A natural approach is to consider semicontinuous data as the result of a two-part process: the ‘binary part’ determining whether an observation is zero (i.e. ‘Is it raining today? yes/no’), and the ‘continuous part’ determining the positive value of the observation given that it is nonzero (i.e. ‘Given that it is raining today, how much?’). In order to test the null hypothesis of no treatment effect (i.e. ‘Is it rainier over here than over there?’), two-part models were first considered in econometrics [55,63].

By tracking daily rainfall data, one can test whether or not two places (e.g. the city of Vancouver and the neighbouring city of North Vancouver) are equally rainy, but things are complicated by the unavoidable correlation between yesterday’s, today’s, and tomorrow’s weather. To address the inherent correlation within *longitudinal* semicontinuous data, Olsen and Schafer (O&S) [42] proposed a two-part model that accounts for the correlation between a given subject’s repeated observations by including subject-specific random effects. The important idea put forth by O&S [42] is that one can account for the semicontinuous distribution of the outcome variable by allowing the random-effects included in the binary and the continuous parts of the two-part model to be correlated. Despite computational challenges, the two-part model for longitudinal semicontinuous data has inspired numerous extensions and variations.

Liu *et al.* [32] extended the model by allowing positive values to follow different distributions: (a) a generalized gamma distribution, (b) a log-skew-normal distribution, and (c) a normal distribution after the Box-Cox transformation. Tom, Su and Farewell [52] introduced correct formulation for marginal inference. Most recently, Lo [33] fit four versions: (1) a logit-log-normal random effects model, (2) a two-part logit-truncated normal random effects model, (3) a two-part logit-gamma random effects model, and (4) a two-part logit-skew normal random effects model.

Unfortunately, obtaining maximum likelihood parameter estimates for the two-part model remains difficult. In the statistical literature, the process has been described as ‘challenging’ [32], ‘a unique challenge’, [41] and one that can ‘lead to severe computational problems’ [49]. O&S discuss several strategies including the Monte Carlo EM algorithm, which they determine to be ‘far too slow’ [42]. Further details on the EM approach are provided in earlier work with the same conclusion: ‘it is computationally accurate but too slow for practical use’ [41].

Tooze *et al.* [53] recommend using adaptive Gaussian quadrature. This approach is made relatively easy due to implementation available within the SAS PROC NL MIXED procedure. However, the method is computationally demanding, with Su *et al.* [49] noting that: ‘even with properly standardized explanatory variables and the simplest model with 2 correlated random intercepts, it can take several hours to fit using the SAS NL MIXED procedure’ [49]. After much consideration, O&S recommend using Fisher scoring to maximize a sixth-order Laplace-expansion approximation of the likelihood noting that, parameter estimates are obtained relatively quickly under favourable circumstances (when implemented with a Fortran-90 program) [42]. Under less favourable conditions, the Laplace-expansion strategy appears less capable. When fitting the two-part model to artificial data simulated to mimic longitudinal 1,000 subject survey data, O&S find that the algorithm fails to converge in over 60% of runs [42]. Despite this, the authors remain optimistic concluding that: ‘This leads us to believe that when the algorithm does converge,

the estimated coefficients and standard errors are quite reliable'. And in the event that convergence fails, analysts will: 'typically abandon the model and try a simpler one' [42].

What simpler model should analysts turn to in the event that convergence fails? Su *et al.* [49] warn against using two independent models (i.e. the two-part model with non-correlated random-effects). Ignoring the correlation between random effects can introduce bias in the estimation of the regression coefficients and variance components due to an 'informative cluster size' problem. This is due to the fact that parameters in the binary part of the model inevitably impact the number of observations (m_i^*) for the continuous part of the model.

A random-effects Tobit model [9], may also be problematic [25]. According to some, the Tobit model lacks the richness required for semicontinuous data and is therefore an 'unattractive' option [46]. The Tobit model assumes that there is both an unobserved normally distributed latent variable, and an observed outcome variable. The observed outcome variable is equal to the latent variable whenever the latent variable is above a known threshold, and is equal to the threshold otherwise (a result of being censored, or of being below the limit of detection); for more details see [56]. O&S note that, in addition to the problematic, potentially 'dubious' [46] interpretation, the Tobit model is inappropriate if the semicontinuous outcome is indeed the result of two separate processes, rather than a single process and a censoring mechanism [42]. It remains unclear how this 'inappropriateness' and 'unattractiveness' may impact hypothesis testing with regards to validity and efficiency.

A useful alternative could be a non-parametric rank-based test proposed by Lin *et al.* [28] for the analysis of combined survival and quantitative outcomes. Rank tests have the desirable feature that null distributions of test statistics are exact, and do not depend on detailed parametric assumptions about error distributions [4]. This article investigates the merits of such an approach. After a brief overview of the two-part model, in Section 2, and details concerning the proposed rank-based test in Section 3, the results of a simulation study comparing a variety of methods across a range of scenarios will be discussed in Section 4. An application to rainfall data will be presented in Section 5. Section 6 concludes with recommendations for future research.

2. The two-part model

Let us review the most basic implementation of the two-part model. Consider the simplest scenario in which a semicontinuous outcome variable and a binary covariate of interest (e.g. treatment = drug vs. placebo) are observed for each of $i = 1, \dots, n$ subjects at each of $j = 1, \dots, m_i$ time-points. The order of the m_i time-points may or may not be important and the distribution of the outcome may change (linearly) with time.

Let Y_{ij} be the semicontinuous outcome for individual i at time-point j . Let X_i be an m_i by 3 matrix, with a first column of 1s (intercept) and a second column taking values of the covariate of interest (e.g. = 1 for drug, = 0 for placebo) for individual i at time-point j . The third column of X_i can be equal to the m_i time-points. Under the two-part model, Y_{ij} is recoded as two variables defined as follows:

$$U_{ij} = \begin{cases} 1, & \text{if } Y_{ij} \neq 0 \\ 0, & \text{if } Y_{ij} = 0 \end{cases} \quad \text{and} \quad V_{ij} = \begin{cases} g(Y_{ij}), & \text{if } Y_{ij} \neq 0 \\ \text{irrelevant}, & \text{if } Y_{ij} = 0 \end{cases}$$

where $g(\cdot)$ is a monotonic increasing function chosen such that V_{ij} is approximately Gaussian. As such, for i in $1, \dots, n$, U_i is a binary vector of length m_i , and V_i is a positive valued vector of length $m_i^* = \sum_{j=1}^{m_i} I(Y_{ij} > 0)$, where $I(\cdot)$ is the standard indicator function. (If, for a given subject i , only zeros are observed, $m_i^* = 0$.) For U , a mixed logistic regression model [59] is defined by:

$$\eta_i = X_i \boldsymbol{\beta} + Z_i c_i, \quad \text{with } \boldsymbol{\beta} = (\beta_0, \beta_1, \beta_2),$$

where η_i is a m_i -length vector taking values $\eta_{ij} = \log \pi_{ij} / (1 - \pi_{ij})$, $\pi_{ij} = \Pr(U_{ij} = 1)$. The ‘random intercept model’ defines Z_i equal to a m_i -length vector of 1s (*intercept*). Alternatively, the ‘random slope model’ defines Z_i as a m_i by 2 matrix with the first column equal to 1s (*intercept*) and the second column equal to the m_i time-points (*slope*), with $c_i = (c_{i1}, c_{i2})$.

For V , a Gaussian random-effects regression model [27] is defined by:

$$V_i = X_i^* \boldsymbol{\gamma} + Z_i^* d_i + \epsilon_i, \quad \text{where } \boldsymbol{\gamma} = (\gamma_0, \gamma_1, \gamma_2),$$

with $\epsilon_i \sim N(0, \sigma^2 I)$ and X_i^* and Z_i^* are m_i^* -subsets of X_i and Z_i for points at which $Y_{ij} > 0$. Under the ‘random slope model’, $d_i = (d_{i1}, d_{i2})$. In order to account for the relationship between U and V , the two parts of the model are connected by allowing c_i and d_i to be correlated:

$$b_i = \begin{pmatrix} c_i \\ d_i \end{pmatrix} \sim N \left(0, \psi = \begin{pmatrix} \psi_{cc} & \psi_{cd} \\ \psi_{dc} & \psi_{dd} \end{pmatrix} \right) \quad \text{Random intercept model} \quad (1)$$

$$b_i = \begin{pmatrix} c_{i1} \\ c_{i2} \\ d_{i1} \\ d_{i2} \end{pmatrix} \sim N \left(0, \psi = \begin{pmatrix} \psi_{c_1 c_1} & \psi_{c_1 c_2} & \psi_{c_1 d_1} & \psi_{c_1 d_2} \\ \psi_{c_2 c_1} & \psi_{c_2 c_2} & \psi_{c_2 d_1} & \psi_{c_2 d_2} \\ \psi_{d_1 c_1} & \psi_{d_1 c_2} & \psi_{d_1 d_1} & \psi_{d_1 d_2} \\ \psi_{d_2 c_1} & \psi_{d_2 c_2} & \psi_{d_2 d_1} & \psi_{d_2 d_2} \end{pmatrix} \right) \quad \text{Random slope model} \quad (2)$$

It is worth noting that even under the very simple scenario described, the number of parameters required for the two-part model is substantial, see Table 1. Under the two-part model, the null hypothesis of no treatment effect is composite and defined as: $H_0: \beta_1 = 0$ and $\gamma_1 = 0$. One-sided alternatives are not trivial, see [62]. Oftentimes, one is interested in determining whether or not the rate at which the outcome variable changes, is dependant on treatment. In this situation, a treatment by time interaction must also be included in the model. This introduces at least two additional parameters to estimate (β_3 and γ_3), and the null hypothesis is even more complex: $H_0: \beta_1 = 0, \beta_3 = 0, \gamma_1 = 0$, and $\gamma_3 = 0$.

Many more possibilities exist for the two-part model. The complexity can be substantially greater when the number of covariates increases, different random effects setups are required, and changes over time are more elaborate (e.g. time-dependent covariates, lag-effects, serial correlation). In their concluding remarks, O&S note that hierarchical clustering can be quite common (e.g. neighbourhoods nested within cities) [42]. In order to account for this type of correlation structure, a much larger number of variance parameters would need to be included in the two-part model. While such complexities may further complicate inference and increase the computational burden, the authors note that it is important to make these accommodations: ‘failing to account for intra-cluster

Table 1. Summary of the ‘fixed effects’ + ‘random effects’ parameters required of each model.

	Random Intercept	Random Slope
Fixed effects	$\beta_0, \beta_1, \beta_2$	$\beta_0, \beta_1, \beta_2$
Variance	$\gamma_0, \gamma_1, \gamma_2$ σ^2	$\gamma_0, \gamma_1, \gamma_2$ σ^2
	ψ_{cc} ψ_{cd} ψ_{dd}	$\psi_{c_1c_1}, \psi_{c_1c_2}, \psi_{c_1d_1}, \psi_{c_1d_2}$ $\psi_{c_2c_2}, \psi_{c_2d_1}, \psi_{c_2d_2}$ $\psi_{d_1d_1}, \psi_{d_1d_2}, \psi_{d_2d_2}$
Random effects	c_i d_i	c_{i1}, c_{i2} d_{i1}, d_{i2}
Number of Parameters	$10 + 2n$	$17 + 4n$

heterogeneity may lead one to substantially overstate the actual precision of the estimates’.

3. LLT rank based approach

The rank-based test we propose is based on the work of Lin *et al.* (LL&T) [28] for the analysis of joint time-to-event and quantitative data. Joint outcomes are often observed in survival analysis settings, and there are often clear advantages to considering both outcomes together as a composite-endpoint when determining treatment potential [15,51,57]. The issue of how to best analyze joint-outcome data of this type has received considerable attention in the literature within the last two decades and a wide range of strategies have been recommended [26].

A reasonable and rather simple option is that of non-parametric rank analysis. If one can use both time-to-event and quantitative data to compare subjects relative to one another in terms of outcome severity, then the rank of ordered subjects can be used as a single summary measure for the outcome. (It is important to note that the alternative being tested in all rank based tests is a difference in medians between the two populations and not a difference in the means.)

Moyé *et al.* [38] proposed a simple non-parametric rank method based on a *U*-statistic consisting of pairwise comparisons. Building on this idea, Finkelstein and Schoenfeld [13] proposed a widely used ‘joint-rank’ method for composite outcome data. The Finkelstein and Schoenfeld [13] method is attractive as it provides, in a simple manner, a single summary statistic based on the pairwise comparisons of subjects. Simply put, if this statistic is significant and positive, then one can conclude that subjects receiving treatment obtain significantly better (higher ranking) outcomes than their peers. The proposal of Lin, Li and Tan [28] (LL&T) is a similar ‘pair-wise comparison’ method that allows for the inclusion of additional covariates and considers all data collected on each subject.

In the following, consider *p* covariates of interest (with $p \geq 2$) and β , a *p*-length vector of coefficients. Note that no intercept term is included as ranks are invariant to location shift. As in Section 2, for the *i*th observation at the *j*th time-point, let Y_{ij} be the observed outcome and X_{ij} be a 1 by *p* matrix of covariates (i.e. no intercept), for $i = 1, \dots, n$ and $j = 1, \dots, m_i$. In the joint time-to-event and quantitative data setting, the observed outcomes, Y_{ij} , could be either survival times or quantitative measures.

The LL&T approach involves maximizing the following objective function, based on the maximum rank correlation (MRC) originally proposed by Han [17] (see more recently [19]):

$$S(\boldsymbol{\beta}) = \frac{1}{n(n-1)} \sum_{i=1}^n \sum_{j=1}^{m_i} \sum_{k=1}^n \sum_{l=1}^{m_j} I(Y_{ij} > Y_{kl}) I(\mathbf{X}_{ij}^T \boldsymbol{\beta} > \mathbf{X}_{kl}^T \boldsymbol{\beta}), \quad (3)$$

where, for identifiability, $\|\boldsymbol{\beta}\| = 1$. In the joint time-to-event and quantitative data setting, the ranking of different types of outcomes requires additional considerations for defining $I(Y_{ij} > Y_{kl})$; and censored survival times may add additional complications [28]. In our application with semicontinuous outcomes, all outcomes are of the same type (and measured on the same scale) and ranking the outcomes is therefore straightforward.

Maximizing $S(\boldsymbol{\beta})$ is equivalent to maximizing Kendall's tau correlation coefficient between vectors \mathbf{y} and $\mathbf{X}^T \boldsymbol{\beta}$; see [17]. Intuitively, this optimization is looking to obtain an estimate of $\boldsymbol{\beta}$ such that, whenever we have $Y_{ij} \geq Y_{kl}$, it is likely that we also have $\mathbf{X}_{ij}^T \boldsymbol{\beta} \geq \mathbf{X}_{kl}^T \boldsymbol{\beta}$. In other words, we wish to have an estimated $\boldsymbol{\beta}$ such that the ranking of the *observed* outcomes is in agreement with the ranking of the *fitted* outcomes. Ties are irrelevant and are therefore discarded, see Sherman [47].

Unfortunately, due to the discontinuity of the second indicator function which ranks the fitted values, $I(\mathbf{X}_{ij}^T \boldsymbol{\beta} > \mathbf{X}_{kl}^T \boldsymbol{\beta})$, the $S(\boldsymbol{\beta})$ is a discontinuous step function with abrupt changes. This makes optimization difficult, particularly in high dimensions (i.e. when p is large). The most common approach to overcome this sort of challenge is to approximate the discontinuous $S(\boldsymbol{\beta})$ function with a smooth substitute function (e.g. using a sigmoid function [35], or using a standard Gaussian cdf [29]). Other approaches involve clever, yet computationally costly, grid-search type optimization techniques, see Wang [58]. LL&T make use of smoothing with the standard Gaussian cdf, $\Phi(\cdot)$, such that the estimated $\boldsymbol{\beta}$ is defined as:

$$\hat{\boldsymbol{\beta}} = \operatorname{argmax}_{\boldsymbol{\beta}} \frac{1}{n(n-1)} \sum_{i=1}^n \sum_{j=1}^{m_i} \sum_{k=1}^n \sum_{l=1}^{m_j} I(Y_{ij} > Y_{kl}) \Phi((\mathbf{X}_{ij}^T \boldsymbol{\beta} - \mathbf{X}_{kl}^T \boldsymbol{\beta})/h), \quad (4)$$

where h , the bandwidth, is a very small constant that converges to zero as the sample size, n , increases. Based on the work of Lin and Peng [29] (see also [18,47]), LL&T show that, given a sufficiently small value for h (we require $nh^4 \rightarrow 0$) and a finite number of observations per subject (we require $\max_i(m_i) < \infty$), and five other regularity conditions, the estimate $\hat{\boldsymbol{\beta}}$ is \sqrt{n} -consistent and is asymptotically normal [28]. As such, LL&T declare that the selection of h is 'not crucial for the asymptotic performance of the estimate' [28,30]. (See also, Lin and Peng [29]: 'our method is not sensitive to the bandwidth parameter h '.) LL&T [28] suggest taking h equal to $\tilde{\sigma}/n^{1/3}$, where $\tilde{\sigma}$ is the approximated sample standard deviation of $\mathbf{X}^T \hat{\boldsymbol{\beta}}$ (and where $\hat{\boldsymbol{\beta}}$ can be estimated by iterating over Equation (4) with some initial value chosen for h).

In order to determine the statistical significance of $\hat{\boldsymbol{\beta}}$, its sampling distribution is approximated via non-parametric resampling. LL&T make use of the simple and clever resampling technique of Jin *et al.* [24], whereby a statistic is perturbed with a large number (B) of independent draws from a random variable with mean 1 and variance 1. As Zhou *et al.* note, 'it is not clear what is the preferred distribution for the perturbation' [61]. As

such, we will use exponential perturbations, a relatively popular choice (e.g. [11,23,28,44]). Note that, in our implementation, perturbations to the $\hat{\beta}$ statistic are made at the subject-level so as to account for the correlation between repeated measurements. (As an alternative to the perturbation resampling technique, a bootstrap resampling scheme could be implemented [12,50]. In fact, the resampling technique of Jin *et al.* [24] is known to be a special case of Rubin [45]’s ‘Bayesian bootstrap’, see [43]. In a bootstrap-based implementation, in order to account for the correlation between repeated measurements, data must be resampled at the subject level.)

What follows is a step-by-step summary of the LL&T algorithm implemented for longitudinal semicontinuous outcome data in two parts.

The LL&T algorithm for longitudinal semicontinuous outcome data.

Part 1- Establishing the bandwidth and point estimate. Define $\tilde{\sigma}^{[1]} = 1$. Then, for q in $1, \dots, Q$ (or until $\tilde{\sigma}^{[q]} \approx \tilde{\sigma}^{[q-1]}$):

1. With $\Phi()$ as the Normal cdf, and with the restriction that $\|\beta\| = 1$, use numerical methods to maximize:

$$\hat{\beta}^{[q]} = \operatorname{argmax}_{\beta} \frac{1}{n(n-1)} \sum_{i=1}^n \sum_{j=1}^{m_i} \sum_{k=1}^n \sum_{l=1}^{m_j} I(Y_{ij} > Y_{kl}) \times \Phi((\beta^T X_{ij} - \beta^T X_{kl})/h^{[q]})$$

2. Define $\tilde{\sigma}^{[q+1]}$ as the standard deviation of $X^T \hat{\beta}^{[q]}$.
3. Define $h^{[q+1]} = \tilde{\sigma}^{[q+1]}/n^{1/3}$.

Part 2- Calculating p-values. Take the point estimate, $\hat{\beta}$, and the bandwidth, h , from the final iteration of Part 1.

Then, for b in $1, \dots, B$:

1. For i in $1, \dots, n$, sample $\zeta_i^{[b]}$ from the exponential distribution with mean and variance equal 1.
2. With $\Phi()$ as the Normal cdf, and with the restriction that $\|\beta\| = 1$, use numerical methods to maximize:

$$\beta^{[b]} = \operatorname{argmax}_{\beta} \frac{1}{n(n-1)} \sum_{i=1}^n \sum_{j=1}^{m_i} \sum_{k=1}^n \sum_{l=1}^{m_j} \zeta_i^{[b]} \zeta_j^{[b]} I(Y_{ij} > Y_{kl}) \times \Phi((\beta^T X_{ij} - \beta^T X_{kl})/h).$$

Define one-sided p -values, for j in $1, \dots, p$, as follows: $p\text{-value}_j = (1 + \sum_{b=1}^B I(\beta_j^{[b]} > 0))/(1 + B)$.

Define two-sided p -values, for j in $1, \dots, p$, as follows:

$$p\text{-value}_j = 2 \cdot \min \left(\frac{1 + \sum_{b=1}^B I(\beta_j^{[b]} > 0)}{1 + B}, \frac{1 + \sum_{b=1}^B I(\beta_j^{[b]} \leq 0)}{1 + B} \right).$$

Four comments are necessary to clarify the steps outlined above. First, the constrained maximization required can be achieved by first transforming the β vector of length p into a vector θ of length $p-1$ of corresponding polar coordinates. For example, with $p=2$, simply define $\beta = (\sin(\theta), \cos(\theta))$. Transformations for higher dimensional polar coordinates can be obtained similarly [40]. Maximization with respect to the vector θ can then be achieved using standard optimization methods (e.g. Newton-type algorithms). Alternatively, one could maximize with respect to the untransformed β vector using available techniques for optimization subject to nonlinear constraints (e.g. [6]).

Second, in the event of ties (i.e. equal ranks, $Y_{ij} = Y_{kl}$), the indicator function equals zero, (i.e. $I(Y_{ij} > Y_{kl}) = 0$), see [47]. Third, note that the coefficient estimates are scaled ($\|\beta\| = 1$) so that the norm of the estimates is always equal to 1. While this may not be ideal for interpretation, it is necessary to maintain identifiability. We recommend this test only for data with $p \geq 2$. If $p = 1$ (i.e. only one covariate is considered), the only two possibilities for β are values of -1 and 1 and the sampling distribution of the p -value will be poorly behaved.

Finally, we note that, for effective Newton-type optimization, one should consider multiple different initial values so as to avoid local maxima. In our simulation study and application, we select six random initial values as starting points. For each of these six points, we run three Newton-Raphson steps and continue forward with Newton-Raphson optimization from the point for which the objective function is greatest.

4. Simulation study

A small simulation study was conducted comparing the power and type I error of four methods. The simulations were coded in SAS in order to make use of existing procedures and macros. We compared the following four approaches for hypothesis testing:

1. a mixed logistic model (using PROC GLIMMIX);
2. the two-part model (with $g(Y) = \log(Y)$) of O&S [42] (using PROC NLMIXED as in [32,53]);
3. the mixed Tobit model with a log-transformed outcome (i.e. $z = \log(Y)$, if $Y > 0$; $z = 0$, otherwise) (using PROC NLMIXED) [56]; and
4. the proposed LL&T non-parametric rank based test (using PROC IML), with $B = 101$ and $Q = 5$.

Data was simulated from the random intercept model (see Equations (1)–(3)) based in part on the simulation work of O&S [42]. The total number of subjects was equal to n and the number of measurements per subject was varied with $m_i = 5 + p_{i2}$, where $p_{i2} \sim \text{Poisson}(\lambda = 2)$, for i in 1 to n . For each subject, $i = 1, \dots, n$, we generated a non-time-varying covariate $(X_1)_i$ from Bernoulli(0.5). This X_1 was the ‘variable of interest’. We also defined a time-varying covariate, X_2 , as $(X_2)_{ij} = j$, for $j = 1, \dots, m_i$. We then generated the V_i s from a $N(X_i^* \gamma + Z_i^* d_i, \sigma^2 I)$ distribution and U_i s from a Bernoulli($\exp(X_i \beta + Z_i c_i)$) distribution. Finally, we implemented three different scenarios:

- Scenario 1: We define $Y_{ij} = 0$ if $(U_{ij} = 0$ or $V_{ij} \leq 0)$, and $Y_{ij} = \exp(V_{ij})$ otherwise.
- Scenario 2: We define $Y_{ij} = 0$ if $(U_{ij} = 0$ or $V_{ij} \leq 0)$, and $Y_{ij} = \exp(\sqrt{V_{ij}})$ otherwise.

- Scenario 3: We define $Y_{ij} = 0$ if $(U_{ij} = 0 \text{ or } V_{ij} \leq 0)$, and $Y_{ij} = \exp(W_{ij} + \sqrt{V_{ij}})$ otherwise, where $W_{ij} \sim \text{Bernoulli}(0.5)$.

As such, the parametric assumptions required of the two-part model and Tobit model (that the function $g()$ is such that V_{ij} is normally distributed) will be satisfied for Scenario 1, not quite satisfied for Scenario 2, and certainly not satisfied in Scenario 3. Figure 1 plots histograms to illustrate the distribution of the data for each of the scenarios. In practice, upon observing the data, one could conceivably try to use a monotone function to transform the V_{ij} variable as needed. However, for the purposes of this simulation study, we will not allow this strategy as we are interested in investigating the effects of violations to the parametric assumptions (or equivalently, the misspecification of the $g()$ function).

Parameter values for $\beta_0, \gamma_0, \beta_2, \gamma_2, \sigma^2, \psi_{cc}, \psi_{cd}$ and ψ_{dd} were fixed at: $\beta_0 = 0.25, \gamma_0 = 2.5, \beta_2 = 0.15, \gamma_2 = 0.15, \sigma^2 = 0.25, \psi_{cc} = 0.0625, \psi_{cd} = 0.0625, \psi_{dd} = 0.065$. Sample size took one of three values, with $n = 50, 100$ or 150 subjects. Under these settings, approximately 30% of observations were equal to 0.

To study type I error, β_1 and γ_1 were set equal to zero. To study power, β_1 took values 0.1 and 0.25 and γ_1 took values 0.1 and 0.25. For each combination of settings, 500 simulated datasets were generated. The empirical estimate of statistical power is the proportion of simulated datasets in which convergence is achieved and the null hypothesis is rejected (two-sided alternative) at α -level 0.05. Parameter values and sample sizes for the simulations were chosen after a small pilot run of the simulation study suggested they would result in a sufficiently wide range of observed statistical power. Details on the four methods and their implementation are available in the Appendix. Also in the Appendix is the SAS code used to generate the data. SAS Code to implement the four methods is available upon request from the author.

4.1. Simulation study results

Results of the simulation study are presented in Table 2–4. Before reviewing the results with regards to statistical power, consider simulations under the null, when the true $\beta_1 = 0$ and $\gamma_1 = 0$. Several findings merit comment.

- For both scenarios, the rank, Tobit, and logistic approaches show relatively desirable type I error of approximately 0.05 even when sample size is small ($n = 50$). It is difficult to evaluate the type I error for the two-part model, given that convergence is not achieved in a majority of simulations.
- For scenario 1, the two-part model fails to converge in upwards of 60% of simulations. This is in line with the reported findings of O&S [42]. For scenario 2, the two-part model fails to converge in approximately 90% of simulations. This suggests that the two-part model cannot accommodate even modest violations to the stated parametric assumptions.

With regards to statistical power, we obtain a wide range of power across methods which varies considerably with the set values of β_1 and γ_1 . Consider the following. For *Scenario 1* (see Table 2):




Figure 1. Histograms show the distribution of the outcome variable, Y , considered under each of the three scenarios in the simulation study, with $\beta_1 = 0.25$ and $\gamma_1 = 0.25$. The red and blue bars represent observations simulated under $X_1 = 0$ and $X_1 = 1$, respectively. Note that the plots have been truncated on the horizontal axis, such that many larger values in Scenario 1 are not visible.

Table 2. Scenario 1– Simulations study results show power to detect a significant effect at the $\alpha = 0.05$ level. Data conforms to the parametric assumptions of the two-part model. The empirical estimate of statistical power is the proportion of simulated datasets in which convergence is achieved and the null hypothesis is rejected (two-sided alternative) at α -level 0.05. Results are based on 500 simulations runs. Numbers displayed in *bold* indicate best-performing method (within $\pm 1\%$) for a given scenario.

n	β_1	γ_1	rank test	Tobit model	two-part model	logistic model	two-part conv. failures (%)	logistic conv. failures (%)
50	0	0	0.07	0.06	0.02	0.09	62.40	1.20
100	0	0	0.04	0.07	0.02	0.05	60.40	1.80
150	0	0	0.04	0.05	0.02	0.05	59.60	0.60
50	0.25	0.10	0.19	0.20	0.04	0.12	65.60	0.60
100	0.25	0.10	0.36	0.30	0.08	0.28	58.00	0.40
150	0.25	0.10	0.54	0.48	0.09	0.39	65.20	0.60
50	0.25	0.25	0.53	0.71	0.15	0.18	65.40	1.20
100	0.25	0.25	0.83	0.95	0.18	0.28	69.00	0.40
150	0.25	0.25	0.94	0.98	0.23	0.40	71.60	0.40
50	0.10	0.25	0.40	0.73	0.20	0.05	66.60	1.20
100	0.10	0.25	0.66	0.95	0.24	0.10	70.20	1.60
150	0.10	0.25	0.86	0.99	0.28	0.12	70.00	0.20

Table 3. Scenario 2– Simulations study results show power to detect a significant effect at the $\alpha = 0.05$ level. Data does not conform to the parametric assumptions of the two-part model. Results are based on 500 simulation runs. The empirical estimate of statistical power is the proportion of simulated datasets in which convergence is achieved and the null hypothesis is rejected (two-sided alternative) at α -level 0.05. Numbers displayed in *bold* indicate best-performing method (within $\pm 1\%$) for a given scenario.

n	β_1	γ_1	rank test	Tobit model	two-part model	logistic model	two-part conv. failures (%)	logistic conv. failures (%)
50	0	0	0.05	0.05	0.01	0.05	89.80	0.20
100	0	0	0.03	0.06	0.00	0.05	94.20	0.60
150	0	0	0.04	0.06	0.00	0.05	91.80	2.00
50	0.25	0.10	0.21	0.16	0.03	0.16	87.40	0.20
100	0.25	0.10	0.43	0.33	0.02	0.27	96.00	0.20
150	0.25	0.10	0.50	0.38	0.02	0.39	95.60	1.00
50	0.25	0.25	0.52	0.71	0.08	0.14	87.80	1.00
100	0.25	0.25	0.84	0.95	0.03	0.31	95.60	0.80
150	0.25	0.25	0.95	0.99	0.06	0.39	92.80	0.20
50	0.10	0.25	0.42	0.74	0.09	0.06	85.20	1.00
100	0.10	0.25	0.66	0.94	0.06	0.07	92.80	1.60
150	0.10	0.25	0.82	1.00	0.06	0.09	92.80	1.20

- when $\beta_1 = 0.25$ and $\gamma_1 = 0.10$, the rank test appears to have the highest power relative to the other methods.
- when $\beta_1 = 0.25$ and $\gamma_1 = 0.25$, the power of the logistic model remains relatively low, while the other methods all show higher power; the Tobit model shows higher power than the rank test which has higher power than the two-part model.
- when $\beta_1 = 0.10$ and $\gamma_1 = 0.25$, the power of the logistic model is negligible; the Tobit model shows higher power than the rank test and the two-part model.
- for all settings, the rank test and the Tobit model show higher power than the logistic model and the two-part model.

For *Scenario 2* (see Table 3):

- the two-part model fails to converge for more than 85% of simulated datasets.

Table 4. Scenario 3– Simulations study results show power to detect a significant effect at the $\alpha = 0.05$ level. Data is multi-modal and as such does not conform to the parametric assumptions of the two-part model. Results are based on 500 simulation runs. The empirical estimate of statistical power is the proportion of simulated datasets in which convergence is achieved and the null hypothesis is rejected (two-sided alternative) at α -level 0.05. Numbers displayed in *bold* indicate best-performing method (within $\pm 1\%$) for a given scenario.

n	β_1	γ_1	rank test	Tobit model	two-part model	logistic model	two-part conv. failures (%)	logistic conv. failures (%)
50	0	0	0.04	0.05	0.01	0.03	81.40	0.40
100	0	0	0.05	0.06	0.01	0.05	75.40	1.00
150	0	0	0.05	0.05	0.01	0.06	77.00	1.00
50	0.25	0.10	0.18	0.08	0.03	0.14	79.80	0.60
100	0.25	0.10	0.31	0.08	0.01	0.30	83.40	0.40
150	0.25	0.10	0.40	0.11	0.02	0.38	85.60	1.00
50	0.25	0.25	0.32	0.16	0.02	0.19	82.20	0.60
100	0.25	0.25	0.57	0.27	0.02	0.27	77.40	0.40
150	0.25	0.25	0.73	0.40	0.03	0.41	77.60	0.20
50	0.10	0.25	0.19	0.13	0.01	0.09	81.60	0.80
100	0.10	0.25	0.30	0.27	0.01	0.08	80.20	1.20
150	0.10	0.25	0.45	0.41	0.01	0.10	77.00	2.20

- when $\beta_1 = 0.25$ and $\gamma_1 = 0.10$, the rank test appears to have the highest power relative to the other methods.
- when $\beta_1 = 0.25$ and $\gamma_1 = 0.10$, the Tobit model has about the same power as the logistic model; when $\beta_1 = 0.10$ and $\gamma_1 = 0.25$, the Tobit model shows the highest power; and also when $\beta_1 = 0.25$ and $\gamma_1 = 0.25$, the Tobit model shows the highest power.
- for all settings, the rank test shows higher power than the logistic model. The same can not be said of the Tobit model.

For *Scenario 3* (see Table 4):

- the two-part model fails to converge for more than 75% of simulated datasets and when it does converge it very rarely rejects the null.
- when $\beta_1 = 0.25$, the Tobit model shows substantially lower power than the logistic model.
- for all settings, the rank test shows higher power than the three other approaches considered.

Given the relatively successful performance of the random-effects Tobit model, further discussion of this method is appropriate. The likelihood of a Tobit model explicitly incorporates both the probability that an observation is below or equal to zero and the probability distribution of an observation conditional on the fact that it is above zero [48]. This could explain why the Tobit model appears rather efficient for testing semicontinuous data. In our Scenario 2 simulations, when the assumption of normality was invalid, we still observed relatively high power with the Tobit model under most settings. However, under Scenario 3, the Tobit model proved rather inefficient. As such, we remain skeptical as to whether the Tobit model is robust under departures from normality; see also the concerns of [3,20]. Indeed, in many cases for Scenario 3, the Tobit model showed lower power than the logistic model. After a thorough investigation, Arabmazar and Schmidt [1] note of the

Tobit model that: ‘The bias from non-normality can be substantial’ and that the ‘bias due to non-normality depends on the degree of censoring’. For our simulated semicontinuous data, the proportion of observations equal to zero was fairly consistent at approximately 30%.

5. Motivating example: is it even rainier in North Vancouver?

Canadians love to complain about the weather, and those in Vancouver on the We(s)t-coast, are no exception due to the never-ending rainfall. From October to March, it is not uncommon to have three soaking weeks worth of non-stop cold and rainy weather. The daily downpours have precipitated, among umbrella-wielding Vancouverites, a *not-so-heated* debate as to whether neighbouring North Vancouver is even rainier. *Across the water is the weather even wetter?* Environment Canada’s historical weather data is publicly available (www.climate.weather.gc.ca/) and can help us answer this pertinent pluviometric query.

Using four recent years (2013–2016) of daily data from two weather stations (‘Vancouver Harbour CS’ and ‘N. Vancouver Wharves’ at a distance of only 2.23km from one another) we can illustrate how each statistical method compares in terms of efficiency: how many weeks of data are required to correctly conclude that the city North Vancouver is indeed rainier (i.e. reject the null hypothesis)? Table 5 provides a summary splash of the data. (One outlier daily observation (N.Van., 2015.11.30) that was obviously a typo was dropped.)

We randomly sampled 50 (75, and 100) week-city pairs worth of observations and tested for a difference in rainfall between the two cities using the four statistical methods as in the simulation study. The models used in the simulation study were fit nearly identically to the rainfall data (one difference was setting $B = 201$ for added precision). Each week served as an observational unit with days within a week correlated. The X_1 variable specified the city ($X_1 = 0$: ‘Vancouver’; $X_1 = 1$: ‘North Vancouver’). The X_2 variable was set as an indicator of the season such that $X_2 = 0$ represented ‘April to September’ and $X_2 = 1$ represented ‘October to March.’

Due to missing data, not all weeks included seven days. No effort was made to force balance (i.e. in a given dataset, the number of observed days from Vancouver and North Vancouver was not set to be equal, nor was the timing of these observed days). We repeated the exercise 100 times, recording all two-sided p -values. We then evaluated the performance of each model by measuring the proportion of sampled datasets for which the method achieves convergence and correctly rejects the null hypothesis. In this way, we are able to evaluate, to a certain degree, the efficiency of the methods with this real-world data. However, do note that this evaluation of efficiency is imperfect: since we are randomly

Table 5. Rainfall Data Summary– Summary of daily rainfall data (2013-2016) obtained from Environment Canada on two weather stations, ‘Vancouver Harbour CS’ and ‘N. Vancouver Wharves’; ‘rainy season’ is October to March.

	Number of non-rainy days	Total mm recorded	... during rainy season (mm)	Daily mean (mm)	Daily Max. (mm)	NA's (days)
Vancouver	748	4296.00	3242.60	3.20	46.60	114
North Vancouver	657	5851.50	4367.30	4.58	64.00	183

Table 6. Rainfall Data— Results show number of analyses (out of 100) in which convergence is achieved and a significant effect at the $\alpha = 0.05$ level is detected (i.e. in which the null hypothesis is rejected).

n	rank test	Tobit model	two-part model	logistic model	two-part model conv. failures	logistic model conv. failures
50	6	28	7	5	68	0
75	12	39	5	9	61	2
100	44	55	10	11	51	0

sampling observations from a finite dataset (four years of data), each of the 100 evaluations cannot be considered independent.

Table 6 shows the results in terms of the number of times in which each method converged and rejected the null hypothesis. While by no means a deluge, the Tobit model rejected the null the most: in 28 (39, and 55) out of 100 instances. The rank test did not perform quite as well, rejecting the null 6 (12, and 44) times out of 100. In contrast, the logistic model shows negligible efficiency, while, for the majority of samples, convergence failures overcast the two-part model. Clearly, neither the logistic model nor the two-part model can be recommended for distinguishing the drizzly days. The number of times in which the two-part model both converges and successfully rejects the null, is like that of sunny days in the winter months, only 7 (5, and 10) out of 100. While the rank-test does not have any similar convergence issues, it is very slow to compute. On a personal computer (OS X, 8 GB memory, 2.6 GHz processor) running a single analysis of the rain data with $n = 100$ and $B = 201$, takes approximately 25 minutes using SAS Studio and 12 minutes with R (v3.3).

It is worth noting that both covariates in this data analysis are binary. This may be one reason for the relatively limited performance of the rank-based test. If X_1 and X_2 are exactly the same for two sample units, then all corresponding pairwise comparisons will contribute nothing to the estimation of β (see Equation (4)). The rank based test may be best suited for data in which at least one covariate is continuous.

6. Conclusion

How to best address the challenges of longitudinal semicontinuous data continues to be an important and difficult statistical question. Neelon *et al.* [39] provide a summary of the current state of affairs and related issues concerning zero-inflated count data.

In this work, we proposed using an existing testing method for a different application than the one for which it was originally intended. The LL&T rank-test has been proposed and recommended only for applications of joint longitudinal survival data. We believe it has far greater potential. To the best of our knowledge, using a rank-based method for testing longitudinal semicontinuous data has not yet been considered in the literature (one exception might be [54] who very briefly consider the standard Wilcoxon rank sum test). This is no doubt due to the fact that standard rank tests cannot incorporate multiple (both continuous and categorical) covariates; nor can they adjust for the correlation amongst longitudinal (or clustered) observations. Recently, some progress has been made in these regards; see e.g. [7,22].

We believe the LL&T rank-test has much potential for longitudinal semicontinuous data (as well as other data with unorthodox distributions) as it allows for multiple covariates

and correlated observations in a straightforward way. Furthermore, as we established in the simulation studies, the LL&T test shows relatively high efficiency. That being said, this rank based test is a computationally intensive procedure and it remains to be determined how well it will work with a larger number of covariates, p , and a larger number of observations, n . One thing is certain, optimization of the $S(\beta)$ objective function will be much more computationally costly with larger values of p and n ; see the recent work of Fan *et al.* [10]. While the computational cost of the rank-test can be substantial, it may still remain feasible in many applications due to the fact that the algorithm is easily parallelized.

Our simulation study suggested that the advantages, with regards to efficiency, of using the rank-test over alternatives such as the random effects Tobit model, will be dependent on whether the effect is primarily expressed through the binary or continuous components of the data (i.e. will be dependent on the relative magnitude of β_1 and γ_1). Among the four methods tested, the LL&T rank-test is the most powerful in situations when a substantial amount of the treatment effect is expressed through the binary component of the data (i.e. when the magnitude of β_1 is large relative to the magnitude of γ_1).

Our simulation study also confirmed that computational issues can greatly restrict the use of the two-part model in many cases. In fact, we observed that when distributional assumptions of the two-part model do not hold, model convergence is almost always unattainable. Should computational issues be resolved, the two-part model would be, without a doubt, a most desirable approach. Current research on this is promising (e.g. [60] and the Bayesian methods of [8]). For those seeking a simpler approach, the random-effects Tobit model and the LL&T rank-test are good options to consider.

Previous concerns about the suitability of the random-effects Tobit model appear largely unwarranted based on our results (at least with regards to testing). Indeed, the performance of the Tobit model in our simulation study does seem at odds with recommendations in the literature [42,46]. Further research on this is warranted. If the distribution of the continuous outcome appears particularly skewed (even after transformations) such that the assumption of normality may not hold, new research into skew-normal Tobit models may prove useful [48]. With respect to the LL&T rank-test, testing is efficient while interpretation is less straightforward due to the fact that coefficient estimates are scaled. Our simulation study also confirmed that truncating semi-continuous data to binary can be very detrimental in terms of power. In all cases considered, the LL&T rank-test was preferable to the simple mixed logistic model approach. When the data was multi-modal, the LL&T rank-test was clearly the preferred choice for testing among the models we considered.

In many circumstances when testing is the primary objective, the choices required to fit parametric models (e.g. ‘what random-effects to include?’, ‘what function $g()$ is most appropriate?’) can be difficult to justify and the many model assumptions can prove burdensome. In such a setting, a main advantage for the LL&T rank-test is the lack of parametric assumptions and required modelling choices.

The data considered in this paper were chosen to be the simplest possible in order to best illustrate the main challenges involved with longitudinal semicontinuous data. As such, not all aspects of the data were considered in the analysis. For example, we treated each city-week as an independent unit, ignoring any residual correlation between adjacent weeks for the same city. This type of issue and other difficulties may often arise in practice. For example, what to do about informative cluster sizes (i.e. when the number of observations

is related to the outcome) or missing data? How to best accommodate serial correlation (e.g. time-series data)? There are many potential areas for future methodological research. First and foremost, further investigation of the rank based test with regards to its sensitivity to the selection of the bandwidth is required. While the choice of bandwidth is ‘not crucial for the asymptotic performance of the estimate’ [28], it may have the potential to impact the statistical power of the test. In related work, Horowitz [21] investigates, for tests based on smoothed maximum score estimators, how bootstrap critical values are sensitive to the choice of the bandwidth. Less is known about the impact of bandwidth selection when using the perturbation resampling method of Jin *et al.* [24].

Finally, with the daily Vancouver rainfall data, our goal was not to properly model the meteorological process (as others have done, e.g. [14,16]). Rather we would hope that such a simple example, in which established methods show a drought of efficiency and an inundation of computational issues, might precipitate the use alternative methods.

Notes

Note that this work has been previously posted as a pre-print, see [5]. In addition, the rain data and the SAS code for the simulation study has been posted to the publicly available OSF repository (DOI 10.17605/OSF.IO/PZGSA). Finally, to promote the use of the LL&T rank-test we encourage those interested to make use of the R function `LLTranktest` available in the Appendix 3 and within a package on the Github repository ‘harlanhappydog/LLTranktest’. The following four lines of R code will download and install the ‘LLTranktest’ package, and use the LL&T rank-test to test for a treatment effect in the analysis of the longitudinal semicontinuous ‘toenail data’ [2], also analyzed in [37] and [36]:

```
library(devtools)
install_github("harlanhappydog/LLTranktest")
library(LLTranktest)
LLTranktest(UNL_mm~treat_group+month,
            id=toenail$ID, data=toenail, B=1000)
```

Acknowledgements

Thank you to Dr. Kahlil Baker for the friendly and thought-provoking computational assistance. Thank you to Professors Lang Wu, Nancy Heckman, and Paul Gustafson for the thoughtful discussions and suggestions.

Disclosure statement

No potential conflict of interest was reported by the authors.

ORCID

Harlan Campbell  <http://orcid.org/0000-0002-0959-1594>

References

- [1] A. Arabmazar and P. Schmidt, *An investigation of the robustness of the tobit estimator to non-normality*, *Econometrica* 50 (1982), pp. 1055–1063.
- [2] M. Backer, P. Keyser, C. Vroey, and E. Lesaffre, *A 12-week treatment for dermatophyte toe onychomycosis terbinafine 250 mg/day vs. itraconazole 200 mg/day - a double-blind comparative trial*, *Brit. J. Dermatol.* 134 (1996), pp. 16–17.
- [3] M. Barros, M. Galea, V. Leiva, and M. Santos-Neto, *Generalized tobit models: Diagnostics and application in econometrics*, *J. Appl. Stat.* 45 (2018), pp. 145–167.
- [4] B. Brown and Y.G. Wang, *Standard errors and covariance matrices for smoothed rank estimators*, *Biometrika* 92 (2005), pp. 149–158.
- [5] H. Campbell, *Is it even rainier in North Vancouver? A non-parametric rank-based test for semicontinuous longitudinal data*, preprint (2017). Available at arXiv:1711.08876.
- [6] X. Chen and X. Yin, *Nlcoptim: Solve nonlinear optimization with nonlinear constraints* (2017).
- [7] S. Datta and G.A. Satten, *A signed-rank test for clustered data*, *Biometrics* 64 (2008), pp. 501–507.
- [8] E. Dreassi and E. Rocco, *A Bayesian semiparametric model for non negative semicontinuous data*, *Comm. Statist. Theory Methods* 46 (2017), pp. 5133–5146.
- [9] M.P. Epstein, X. Lin, and M. Boehnke, *A tobit variance-component method for linkage analysis of censored trait data*, *Am. J. Hum. Genet.* 72 (2003), pp. 611–620.
- [10] Y. Fan, F. Han, W. Li, and X.H. Zhou, *On rank estimators in increasing dimensions* (2017).
- [11] C. Fan, W. Lu, R. Song, and Y. Zhou, *Concordance-assisted learning for estimating optimal individualized treatment regimes*, *J. R. Stat. Soc. Ser. B* 79 (2017), pp. 1565–1582.
- [12] C.A. Field and A.H. Welsh, *Bootstrapping clustered data*, *J. R. Stat. Soc. Ser. B* 69 (2007), pp. 369–390.
- [13] D.M. Finkelstein and D.A. Schoenfeld, *Combining mortality and longitudinal measures in clinical trials*, *Stat. Med.* 18 (1999), pp. 1341–1354.
- [14] S.W. Fleming, *Climatic influences on Markovian transition matrices for Vancouver daily rainfall occurrence*, *Atmos. Ocean* 45 (2007), pp. 163–171.
- [15] N. Freemantle, M. Calvert, J. Wood, J. Eastaugh, and C. Griffin, *Composite outcomes in randomized trials: Greater precision but with greater uncertainty?*, *JAMA* 289 (2003), pp. 2554–2559.
- [16] J. George, J. Letha, and P. Jairaj, *Daily rainfall prediction using generalized linear bivariate model—a case study*, *Procedia Technol.* 24 (2016), pp. 31–38.
- [17] A.K. Han, *Non-parametric analysis of a generalized regression model: The maximum rank correlation estimator*, *J. Econom.* 35 (1987), pp. 303–316.
- [18] A. Han, *Large sample properties of the maximum rank correlation estimator in generalized regression models*, Preprint, Department of Economics, Harvard University, Cambridge, MA, 1988.
- [19] F. Han, H. Ji, Z. Ji, and H. Wang, *A provable smoothing approach for high dimensional generalized regression with applications in genomics*, *Electron. J. Stat.* 11 (2017), pp. 4347–4403.
- [20] D. Holden, *Testing for heteroskedasticity in the tobit and probit models*, *J. Appl. Stat.* 38 (2011), pp. 735–744.
- [21] J.L. Horowitz, *Bootstrap critical values for tests based on the smoothed maximum score estimator*, *J. Econom.* 111 (2002), pp. 141–167.
- [22] Y. Jiang, X. He, M.L.T. Lee, B. Rosner, and J. Yan, *Wilcoxon rank-based tests for clustered data with r package clusrank*, preprint (2017). Available at arXiv:1706.03409.
- [23] Z. Jin, D. Lin, L. Wei, and Z. Ying, *Rank-based inference for the accelerated failure time model*, *Biometrika* 90 (2003), pp. 341–353.
- [24] Z. Jin, Z. Ying, and L.J. Wei, *A simple resampling method by perturbing the minimand*, *Biometrika* 88 (2001), pp. 381–390.
- [25] Y. Kim and B.O. Muthén, *Two-part factor mixture modeling: Application to an aggressive behavior measurement instrument*, *Struct. Equ. Modeling.* 16 (2009), pp. 602–624.
- [26] B.F. Kurland, L.L. Johnson, B.L. Egleston, and P.H. Diehr, *Longitudinal data with follow-up truncated by death: Match the analysis method to research aims*, *Stat. Sci.* 24 (2009), p. 211.

- [27] N.M. Laird and J.H. Ware, *Random-effects models for longitudinal data*, *Biometrics* 38 (1982), pp. 963–974.
- [28] H. Lin, Y. Li, and M.T. Tan, *Estimating a unitary effect summary based on combined survival and quantitative outcomes*, *Comput. Stat. Data. Anal.* 66 (2013), pp. 129–139.
- [29] H. Lin and H. Peng, *Smoothed rank correlation of the linear transformation regression model*, *Comput. Stat. Data. Anal.* 57 (2013), pp. 615–630.
- [30] H. Lin, L. Zhou, H. Peng, and X.H. Zhou, *Selection and combination of biomarkers using roc method for disease classification and prediction*, *Canad. J. Statist.* 39 (2011), pp. 324–343.
- [31] L. Liu, R.L. Strawderman, M.E. Cowen, and Y.C.T. Shih, *A flexible two-part random effects model for correlated medical costs*, *J. Health. Econ.* 29 (2010), pp. 110–123.
- [32] L. Liu, R.L. Strawderman, B.A. Johnson, and J.M. O’Quigley, *Analyzing repeated measures semi-continuous data, with application to an alcohol dependence study*, *Stat. Methods. Med. Res.* 25 (2012), pp. 133–152.
- [33] Y. Lo, *Assessing effects of an intervention on bottle-weaning and reducing daily milk intake from bottles in toddlers using two-part random effects models*, *J. Data. Sci.* 13 (2015), pp. 1–20.
- [34] E. Lorenz, C. Jenkner, W. Sauerbrei, and H. Becher, *Modeling variables with a spike at zero: Examples and practical recommendations*, *Am. J. Epidemiol.* 185 (2017), pp. 650–660.
- [35] S. Ma and J. Huang, *Regularized ROC method for disease classification and biomarker selection with microarray data*, *Bioinformatics* 21 (2005), pp. 4356–4362.
- [36] S.E. Mahabadi, *A bayesian shared parameter model for incomplete semicontinuous longitudinal data: An application to toenail dermatophyte onychomycosis study*, *J. Stat. Theory Appl.* 13 (2014), pp. 317–332.
- [37] R.I. Mian and M.T. Hasan, *Two-part pattern-mixture model for longitudinal incomplete semi-continuous toenail data*, *Int. J. Stat. Med. Res.* 1 (2012), pp. 120–127.
- [38] L.A. Moyé, B.R. Davis, and C.M. Hawkins, *Analysis of a clinical trial involving a combined mortality and adherence dependent interval censored endpoint*, *Stat. Med.* 11 (1992), pp. 1705–1717.
- [39] B. Neelon, A.J. O’Malley, and V.A. Smith, *Modeling zero-modified count and semicontinuous data in health services research part 1: Background and overview*, *Stat. Med.* 35 (2016), pp. 5070–5093.
- [40] J. Nolan, *Sphericalcubature: Numerical integration over spheres and balls in n-dimensions; multivariate polar coordinates*, R package version 1 (2015).
- [41] M.K. Olsen and J.L. Schafer, *A class of models for semicontinuous longitudinal data*, *American statistical association proceedings on survey research methods*, 1998, pp. 721–726.
- [42] M.K. Olsen and J.L. Schafer, *A two-part random-effects model for semicontinuous longitudinal data*, *J. Am. Stat. Assoc.* 96 (2001), pp. 730–745.
- [43] M. Parzen and S.R. Lipsitz, *Perturbing the minimand resampling with gamma(1,1) random variables as an extension of the Bayesian bootstrap*, *Stat. Probab. Lett.* 77 (2007), pp. 654–657.
- [44] L. Peng and Y. Huang, *Survival analysis with quantile regression models*, *J. Am. Stat. Assoc.* 103 (2008), pp. 637–649.
- [45] D.B. Rubin, *The Bayesian bootstrap*, *Ann. Statist.* 9 (1981), pp. 130–134.
- [46] J.L. Schafer and M.K. Olsen, *Modeling and imputation of semicontinuous survey variables*, *Proceedings of the federal committee on statistical methodology research conference*, Citeseer, 1999, pp. 565–74.
- [47] R.P. Sherman, *The limiting distribution of the maximum rank correlation estimator*, *Econometrica* 61 (1993), pp. 123–137.
- [48] X. Su and S. Luo, *Analysis of censored longitudinal data with skewness and a terminal event*, *Comm. Statist. Simul. Comput.* 46 (2017), pp. 5378–5391.
- [49] L. Su, B.D. Tom, and V.T. Farewell, *Bias in 2-part mixed models for longitudinal semicontinuous data*, *Biostatistics* 10 (2009), pp. 374–389.
- [50] V. Subbotin, *Asymptotic and bootstrap properties of rank regressions*, Working paper, Northwestern University, 2007.
- [51] D.J. Thompson, *Survival models for data arising from multiphase hazards, latent subgroups or subject to time-dependent treatment effects*, Ph.D. diss., Simon Fraser University, 2011.

- [52] B.D. Tom, L. Su, and V.T. Farewell, *A corrected formulation for marginal inference derived from two-part mixed models for longitudinal semi-continuous data*, *Stat. Methods. Med. Res.* 25 (2016), pp. 2014–2020.
- [53] J.A. Tooze, G.K. Grunwald, and R.H. Jones, *Analysis of repeated measures data with clumping at zero*, *Stat. Methods. Med. Res.* 11 (2002), pp. 341–355.
- [54] V. Tran, D. Liu, A.K. Pradhan, K. Li, C.R. Bingham, B.G. Simons-Morton, and P.S. Albert, *Assessing risk-taking in a driving simulator study: Modeling longitudinal semi-continuous driving data using a two-part regression model with correlated random effects*, *Anal. Methods Accident Res.* 5 (2015), pp. 17–27.
- [55] W. Tu and X.H. Zhou, *A Wald test comparing medical costs based on log-normal distributions with zero valued costs*, *Stat. Med.* 18 (1999), pp. 2749–2761.
- [56] J. Twisk and F. Rijmen, *Longitudinal tobit regression: A new approach to analyze outcome variables with floor or ceiling effects*, *J. Clin. Epidemiol.* 62 (2009), pp. 953–958.
- [57] F. van Leth and J.M. Lange, *Use of composite end points to measure clinical events' reply*, *J. Am. Med. Assoc.* 290 (2003), pp. 1456–1457.
- [58] H. Wang, *A note on iterative marginal optimization: A simple algorithm for maximum rank correlation estimation*, *Comput. Stat. Data. Anal.* 51 (2007), pp. 2803–2812.
- [59] P. Wang and M.L. Puterman, *Mixed logistic regression models*, *J. Agric. Biol. Environ. Stat.* 3 (1998), pp. 175–200.
- [60] B. Zhang, W. Liu, H. Zhang, Q. Chen, and Z. Zhang, *Composite likelihood and maximum likelihood methods for joint latent class modeling of disease prevalence and high-dimensional semicontinuous biomarker data*, *Comput. Stat.* 31 (2016), pp. 425–449.
- [61] M. Zhou, *Empirical likelihood analysis of the rank estimator for the censored accelerated failure time model*, *Biometrika* 92 (2005), pp. 492–498.
- [62] G. Zhou, *Multivariate one-sided tests for multivariate normal and mixed effects regression models with missing data, semi-continuous data and censored data*, Ph.D. diss., University of British Columbia, 2017.
- [63] X.H. Zhou and W. Tu, *Comparison of several independent population means when their samples contain log-normal and possibly zero observations*, *Biometrics* 55 (1999), pp. 645–651.

Appendices

Appendix 1. Simulation study details

A.1 The Random effects Tobit model

First, a logarithmic transformation is made, such that $z_{ij} = \log(Y_{ij})$ if $Y_{ij} > 0$; $z_{ij} = 0$ otherwise, for $i = 1, \dots, n$ and $j = 1, \dots, m_i$. As such, the normality assumption is satisfied for Scenario 1 but not so for Scenario 2. Then, using the SAS NLMIXED procedure (with 5 adaptive quadrature points) the likelihood function is maximized:

$$L = (\sigma^2\sqrt{2\pi})^{-1/2} \exp(-(z_{ij} - \mu_{ij})^2/2\sigma^2), \quad \text{if } z_{ij} > 0,$$

$$= \int_{-\infty}^0 (\sigma^2\sqrt{2\pi})^{-1/2} \exp(-(v - \mu_{ij})^2/2\sigma^2) dv, \quad \text{if } z_{ij} = 0;$$

$$\text{where } : \mu_{ij} = \beta_0 + \beta_1 X_{1ij} + \beta_2 X_{2ij} + u_i \quad \text{and} \quad u_i \sim N(0, \sigma_s^2).$$

Initial values for the random effects Tobit model (PROC NLMIXED) were obtained by first fitting a standard Tobit model (using PROC LIFEREG), and by setting $\sigma_s = \sigma = 0.5$.

A.2 The Binary logistic model

The simplest way to approach longitudinal semicontinuous data, is to reduce the data to longitudinal binary data and fit a mixed effects logistic model. This approach is equivalent to only considering U and ignoring V . The downside to this method is that ignoring V will most certainly result in

a reduction of power. Initial values for the logistic mixed effects model (PROC GLIMMIX) were obtained by first fitting a fixed-effects generalized linear model (GLM).

A.3 The two-part model

SAS macros based on the work of Tooze *et al.* [53] and Liu *et al.* [32], were coded for all simulations. Adaptive Gaussian quadrature within the SAS NLMIXED procedure was used (with 5 adaptive quadrature points as in Liu *et al.* [31] who note that: ‘[i]ncreasing the number of quadrature points to 10 substantially increased computation time with negligible changes to the results’). Initial values for the γ_0 , γ_1 , β_0 , and β_1 parameters were obtained from fitting Generalized Estimating Equation (GEE) models (PROC GENMOD) and initial values for σ , ψ_{cc} , ψ_{cd} and ψ_{dd} were all set to equal 0.5.

A.4 The non-parametric rank-based test

A two-sided p -value was obtained as described in steps in Section 3 with B was set at 101 and Q was equal to 5. This B was much smaller than ideal, due to restricted computational capacity. The restricted ($\|\beta\| = 1$) maximization was made simpler by use of the trigonometric identity $\sin(x)^2 + \cos(x)^2 = 1$. We set $\beta_1 = \sin(\theta)$ and $\beta_2 = \cos(\theta)$ and then maximized by Newton-Raphson iteration with respect to θ . In order to avoid simply selecting local optima, we selected six random initial values as starting points. For each of these six points, we ran three Newton-Raphson steps and continued forward with Newton-Raphson optimization from the point for which the objective function was greatest.

Appendix 2. SAS code used to simulate data

```
m = &m; alpha0=0.25;    alpha1=&alpha1; beta0=2.5;
beta1=&beta1;    sigma = 0.5;    psiaa =0.25;    psibb=0.05;
do j=1 to m;
  ID = j;
  ni = 5 + rand("Poisson", 2);
  ai = rand("Normal", 0, psiaa);  bi = rand("Normal", ai, psibb);
  xtest = rand("Bernoulli", 0.5);
  do i = 1 to ni;
    X2 = i;
    linpred = alpha0 + xtest * alpha1 + 0.15*X2 + ai;
    prob = exp(linpred) / (1 + exp(linpred));
    ytest = uniform(0) lt prob;
    logz = exp(rand("Normal", beta0 + xtest*beta1
      + 0.15*X2+ bi, sigma));
    Y = ytest * logz;
  ZZ = log(Y);
  if zz=. then zz=0;
  X1 =xtest;
  output;
end;
end;
```

Appendix 3. R code to use LLTranktest

```
## null example:
## mydat<-data.frame(Y=rnorm(60), X1=rnorm(60), X2=rnorm(60),
##                  X3=rnorm(60), ID=sort(c(rep(1:20,3))))
## LLTrankTtest(Y~X1+X2+X3, id=mydat$ID, data=mydat, B=200)

#####
LLTranktest<-function(formula, id=NULL, data, B=200, Q=10){
```

```

require(SphericalCubature)
##### The setup
cl <- match.call()

mf <- match.call(expand.dots = FALSE)
m <- match(c("formula", "data"), names(mf), 0)
mf <- mf[c(1,m)]
mf$drop.unused.levels <- TRUE
mf[[1]] <- quote(stats::model.frame)
mf <- eval.parent(mf)

mt <- attr(mf, "terms")
X <- model.matrix(mt, mf, contrasts)[-1]
Y <- model.response(mf, 'numeric')

if(is.null(id)){id<-paste(1:length(Y))}

X<-as.matrix(cbind(X))
p<-dim(X)[2]
N<-length(unique(id))
NM<-length(Y)
s<-rank(Y)

if(p==1){warning("Test only available for more than one covariate")}
break}

#####

## Using (p-1) dimensional polar coords
sn_pdim<-function(mythetavec, psi=rep(1,N), h_mult=1){

  h= h_mult*(N^(-1/3))
  reppsi<-c(unlist(apply(cbind(1:N),1,function(x) rep(psi[x],
    sum(id==unique(id)[x])))))
  betavec<-cbind(rev(c(polar2rect(1, mythetavec))))

  okaym<-function(m){ sum(reppsi*reppsi[m]*(s>s[m])*pnorm
    ((X%*%betavec-c(X[m,]%*%betavec))/h))}
  bb<-sum(-apply(cbind(1:NM),1,okaym))/(N)*(N-1)
  return(bb)
}
#####

generateoriginal<-function(my_H){

  sn_pdim_H<-function(x){sn_pdim(x, psi=rep(1,N), h_mult=my_H)}

  init_points<-matrix(runif(6*(p-1), -3.142, 3.142),6,)
  trial_outs<-init_points*0
  trials<-rep(0,6)
  for(jj in 1:dim(init_points)[1]){

    tryouts<-nlm(sn_pdim_H, init_points[jj,], iterlim=3)
    trial_outs[jj,]<-tryouts$estimate
    trials[jj]<-tryouts$minimum

  }
  theta_hat = nlm(sn_pdim_H, trial_outs[which.min(trials),])$estimate
  mybetavec<-rev(polar2rect(1,theta_hat))
  yfit = c(X%*%mybetavec)
  return(list(sd=sd(yfit), point_est=mybetavec))}

#####

myH<-generateoriginal(1)$sd

```

```

print("Part 1: establishing the bandwidth and point estimate.")
for(q in 1:Q){
part1_out<-generateoriginal(myH)
myH_new<-part1_out$sd
myH<-myH_new
print(paste("sigma[",q,"]=" , round(myH,3), sep=""))}
my_betavec<-part1_out$point_est
print(paste("point estimate ="))
print(round(my_betavec,4))

print("Part 2: Calculating p-values.")
mybetavec_b<-matrix(0,B,p)
pb = txtProgressBar(min = 0, max = B, initial = 0)
for(bb in 1:B){
  setTxtProgressBar(pb,bb)

  tryCatch({
    psib=rexp(N)
    snB_pdim<-function(x){sn_pdim(mythetavec=x, psi=psib, h_mult= myH)}

    init_points<-matrix(runif(6*(p-1), -3.142, 3.142),6,)
    trial_outs<-init_points*0
    trials<-rep(0,6)
    for(jj in 1:dim(init_points)[1]){

      tryouts<-nlm(snB_pdim, init_points[jj,], iterlim=3)
      trial_outs[jj,]<-tryouts$estimate
      trials[jj]<-tryouts$minimum

    }

    theta_hat_b = nlm(snB_pdim, trial_outs[which.min(trials),])$estimate
    mybetavec_b[bb,]<-rev(polar2rect(1, theta_hat_b))

  }, error=function(e){cat("ERROR :",conditionMessage(e), "\n")})
}
close(pb)

pval1 <-rep(0,p)
for(pp in 1:dim(X)[2]){
  pval1[pp]<- 2*min(c((1+sum(mybetavec_b[,pp]>0))/(1+B),
                    (1+ sum(mybetavec_b[,pp]<=0))/(1+B) ))
}

CI<-apply(mybetavec_b,2,function(x) quantile(x,c(0.025,0.975)))
result<-data.frame(scaled_estimate=c(my_betavec),
                  lower95CI=CI[1,],upper95CI=CI[2,], pval= pval1)
rownames(result)<-colnames(X)

print(paste("Number of Observations:",NM))
print(paste("Number of Groups:",N))

return(result)}
#####

```