Diplomarbeit im Institut für Biostatistik
Universität Hannover
Fakultät Naturwissenschaften

Betreuer und Erstprüfer: Prof. Dr. L. Hothorn
Zweitprüfer: Prof. Dr. B. Hau

Poisson confidence intervals for testing the equality of abundance data in generalized linear models

Daniel Gerhard
Matrikelnr.: 2130305
Hornemannweg 4
30167 Hannover

List of abbreviations

α	significance level
β_i	model parameter
Bt	*Bacillus thuringiensis*
δ	equivalence threshold
c	contrasts
χ^2	quantile of the χ^2 distribution
cl	confidence limit
CI	confidence interval
e.g.	abbr. of Latin "'exempli gratia"'
ϵ	residual error
GLM	generalized linear model
GLMM	generalized linear mixed model
H_0	null-Hypothesis
H_A	alternative-Hypothesis
INS	insecticide treated
int	intercept
ISO	isogene
IUT	intersection union test
L	likelihood function
l	log-likelihood function
λ_i	mean of a Poisson distributed sample
μ	sample mean
ML	maximum likelihood
MLE	maximum likelihood estimate
MVN	multivariate normal distribution
n	sample size
ν	link function
p	proportion
ϕ	dispersion factor
PQL	penalized quasi-likelihood
ρ	ratio of two samples
R	correlation matrix
s	standard deviation
t	quantile of the t distribution
Trt	treatment
u_i	random effects parameter
x_i	value in one sample
X	model matrix
z	quantile of the standard normal distribution
ZIP	zero-inflated Poisson model
ZNBM	zero-inflated negative binomial model

Contents

List of Figures

List of Tables

1 Proof of safety for abundance data

For proving safety in a trial, where the influence of genetic modified maize on insect non-target abundance is tested, the use of confidence intervals for estimates of a generalized linear model is examined. To determine the effect of Bt-maize on data of various insect counts not genetic modified maize is used as a control and the use of insecticides on genetic unmodified maize is set as a standard. To compare multiple groups in consideration of time covariates and blocking factors the use of a linear model is recommended. For an easy interpretation confidence intervals should be computed for the ratio of the sample means, to test for equivalence of the treatments. Inference should be made for the ratio of counts to enable the comparability for example of the frequency of different insect species. A bound for the ratio of counts, where two samples are not equal, can be defined for each species as desired.

2 Poisson distribution

When measuring the abundance of insects with catches in traps, the observed values are count data. Count data are non-negative integers with an infinite upper limit, which can follow the Poisson distribution, which probability function is given by

$$pr(Y = y) = \frac{e^{-\mu}\mu^y}{y!}. \tag{1}$$

In the Poisson distribution the variance and all other components depend on the mean. The whole distribution is described by the mean, the only parameter to be modified.

3 Confidence intervals

When for example two sample means should be compared, under the assumptions of normal distributed, independent, continuous and variance homogene data, a Wald t-test can be performed. In the two sided case of this test the hypothesis $H_0 : \mu_i = \mu_j$ can be rejected, if there is a difference between the two groups ($H_A : \mu_i \neq \mu_j$). The test statistic of a Wald t-test can be simply computed by

$$t = \frac{\bar{x}_i - \bar{x}_j}{\sqrt{s^2\left(\frac{1}{n_i} + \frac{1}{n_j}\right)}} \tag{2}$$

where $\bar{x}_i$, with $(i = 1, 2)$, are the two sample means, n_i the sample sizes and s^2 is the common variance estimator of the two samples. This test statistic can be compared with

a critical value of the student t-distribution to an ascertained error probability α, or the quantile of the t-distribution at this value can be used to obtain a p-value, which also can be compared with α. As a result a decision can be made, if the two groups are differing significantly, but it is rather difficult to use a p-value for measurements of the accuracy of discrimination. The equation for the t-test can be adapted for computations of confidence limits, like

$$\mu_1 - \mu_2 \in \left[\bar{x}_i - \bar{x}_j \pm t_{1-\frac{\alpha}{2}} \sqrt{s^2 \left(\frac{1}{n_i} + \frac{1}{n_j} \right)} \right]. \tag{3}$$

Here $t_{1-\alpha/2}$ is the $1 - \alpha/2$ quantile of the student-t-distribution. This confidence interval has the advantage, that its limits are in the scale of the measurements of X_{ij}. If the intervals for the difference of two means not contain zero, H_0 can be rejected, a significant difference between the groups can be detected. In addition effect size can be observed, with the lower limit (with $\bar{x}_i - \bar{x}_j > 0$) showing the minimal difference between the groups and the upper limit showing the maximal difference between the groups at a probability of $1 - \alpha$. If two sided intervals are calculated, further a decision can be made about the the direction of the effect.

4 Test for equivalence

In a typical test for detecting difference between two groups, e.g. a t-test, evidence of this difference can not be proofed directly but it can be shown, that the probability, that the samples are equal is very small. By rejecting the null hypothesis (both samples are equal), the alternative hypothesis (one sample is different from the other) can be assumed.To test for equivalence the hypotheses can be inverted but an equivalence threshold δ has to be inserted, to define at which difference the two samples are expected to be different or equal.

One sided Hypotheses:

Hypotheses	Difference	Equivalence
H_0:	$\mu_1 = \mu_2$	$\mu_1 \neq \mu_2$
H_A:	$\mu_1 - \mu_2 < 0$	$\mu_1 - \mu_2 < \delta$

Considering a two sided test for equivalence, the hypotheses are

$$\begin{aligned} H_A&: \quad \delta_1 < (\mu_1 - \mu_2) < \delta_2 \\ H_0&: \quad (\mu_1 - \mu_2) < \delta_1 \text{ OR } (\mu_1 - \mu_2) > \delta_2 \end{aligned}$$

This null hypothesis can also be written as

$$H_0^{lower}: \quad (\mu_1 - \mu_2) > \delta_1 \text{ AND } \quad H_0^{upper}: \quad (\mu_1 - \mu_2) < \delta_2$$

Both hypotheses H_0^{lower} AND H_0^{upper} have to be rejected to accept H_A, what is called an intersection union test (IUT). This enables us to perform two one sided tests where both of them should show significance to reject H_0 as a level α test.

If δ is known a priori, by transforming the normal test for the difference, a test for equivalence can be performed. But in most cases this threshold is not known. One solution of this problem is to test for equivalence with a confidence interval method. If the confidence limits of a two sided interval are located between δ_1 AND δ_2, H_0 can be rejected. These intervals can be computed at a level of $(1 - 2\alpha)$, as the entire space between two one sided interval limits have to be covered by the interval $[-\delta; \delta]$ and both level α interval limits have to be within this range.

(Wellek (2003), [26])

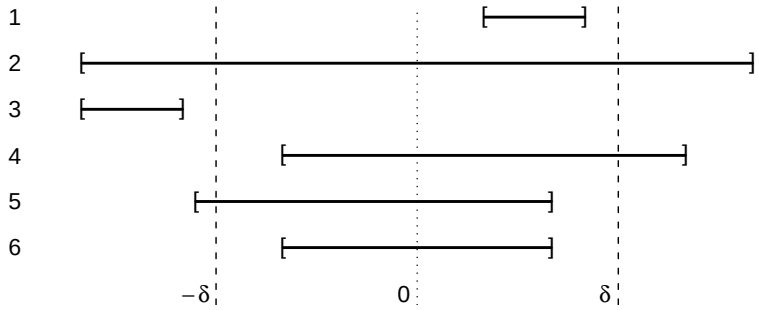

Figure 1: Testing for equivalence with confidence intervals

In Figure 1 with confidence intervals 1 and 6 equivalence can be detected between two samples, because their limits are located between the threshold limits. All the other intervals show no significant equivalence between the groups.

5 Confidence Intervals for two Poisson parameters

Several different confidence intervals can be computed to compare two single Poisson parameters. If these counts were measured in different environments or time intervals,

the rates of the count per environment are of interest, as the observations are made comparable.

5.1 Parameter

In the following x_i, with $(i = 1, 2)$, will be denoted as the response, e.g. one count or the mean of one sample. Further n_i, with $(i = 1, 2)$, is a range, in which the response was measured, for example a time interval or an area. z is the quantile of the standard normal distribution with an error probability of α.

5.2 CIs for the difference of two Poisson parameters

Schwertman and Martinez 1994 [23] constructed eight confidence intervals for the difference of two Poisson parameters in relation to binomial confidence intervals.

5.2.1 Wald interval

Like two proportions also two counts observed over a specific range, e.g. a measurement over a time interval or in a delimited area, can be compared. Simplest way of constructing a confidence interval is Wald's method.

For two binomials:

$$cl_{l,u} = p_1 - p_2 \pm z_{\frac{\alpha}{2}} \sqrt{\frac{p_1(1-p_1)}{n_1} + \frac{p_2(1-p_2)}{n_2}} \tag{4}$$

with the proportion $p_i = \frac{X_i}{n_i}$, X_i as the number of successes in all n_i. Similar to this a Poisson Wald interval for the difference can be derived.

$$cl_{l,u} = \frac{x_1 - x_2 \pm z_{\frac{\alpha}{2}} \sqrt{x_1 + x_2}}{n} \tag{5}$$

where x_i are the counts in a population of size n.

5.2.2 Add4-method

Because the Wald interval is too liberal, it can be improved by adding four pseudo-observations with two successes. The interval construction for proportions is similar to the Wald interval construction, but with

$$p_{add4} = \frac{X_i + 2}{n_i + 4} \tag{6}$$

This continuity correction can be made also for Poisson intervals

$$cl_{l,u} = \frac{(x_1 - x_2) \pm 0.5 \pm z_{\frac{\alpha}{2}} \sqrt{x_1 + x_2 \pm 0.5}}{n} \tag{7}$$

A different approach to construct a CI for Poisson parameters is based on a single parameter interval, which is obtained by solving a quadratic equation in p. (Schwertman, Martinez 1994)

$$cl_{l,u} = \frac{(x_1 - x_2) + \frac{z^2_{\frac{\alpha}{2}}}{2} \pm z_{\frac{\alpha}{2}} \sqrt{x_1 + x_2 + \frac{z^2_{\frac{\alpha}{2}}}{4}}}{n} \tag{8}$$

5.3 CIs for the ratio of two Poisson rates

Poisson parameters are often compared as a ratio of risks, the counts of one population divided by a covariable, which defines the range over which was measured, divided by the rate of counts of a second population.

5.3.1 Wald interval

For the ratio of rates also a simple Wald interval can be constructed.

$$cl_{l,u} = exp\left[log\left(\frac{\frac{x_1}{n_1}}{\frac{x_2}{n_2}}\right) \pm z_{\frac{\alpha}{2}} \sqrt{\frac{1}{x_1} + \frac{1}{x_2}}\right] \tag{9}$$

, where x_1 and x_2 are the positive counts, n_1 and n_2 are the sizes in which range of time or environment x_i were observed. (Graham (2003) [12], Liu (2006) [16])

5.3.2 Jeach's interval

Another method is proposed by Jeach (1970) [13], in which the Poisson statistic is transformed to a standard normal distribution. The lower and upper bounds k can be obtained by:

$$\sqrt{cl_{l,u}} = \sqrt{\left(\frac{\frac{x_1}{n_1}}{\frac{x_2}{n_2}}\right)\left(\frac{\sqrt{(x_1 + 0.5)(x_2 + 0.5)} \pm 0.5 z_{\frac{\alpha}{2}} \sqrt{x_1 + x_2 + 1 - \left(\frac{z^2_{\frac{\alpha}{2}}}{4}\right)}}{x_0 + 0.5 - \left(\frac{z^2_{\frac{\alpha}{2}}}{4}\right)}\right)} \tag{10}$$

(Graham (2003) [12])

5.3.3 Score interval

A third interval proposed by Graham et al. (2003), which behavior is similar to Jeach's CI, is based on the score function. It's limits are computed by

$$cl_{l,u} = \frac{n_2}{n_1}\left(\frac{2x_1x_2 + z^2_{\frac{\alpha}{2}}x. \pm \sqrt{z^2_{\frac{\alpha}{2}}x.\left(4x_1x_2 + z^2_{\frac{\alpha}{2}}x.\right)}}{2x_2^2}\right) \tag{11}$$

, with $x. = x_1 + x_2$.

(Graham (2003) [12])

5.3.4 Example

In the dataset from Eckert (p.49) abundance of four insect species was counted in the years 2003 and 2004. Four plants were observed in overall 24 blocks over 14 dates in 2002 and 8 dates in 2003. Insect counts on isogene maize (control), Bt maize and isogene maize plus insecticide were observed. To deal with the time covariate simply the counts were summed over all dates for each block in each year. The data were split by the factor year (Figure 2).

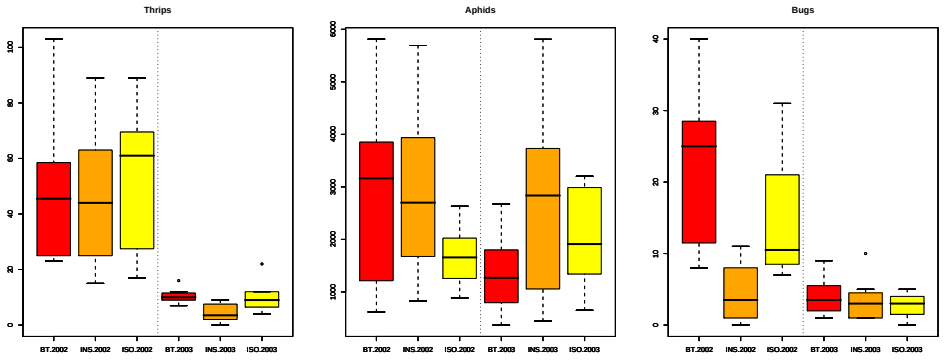

Figure 2: Boxplots for the Eckert dataset for factors type (coloured) and year (left: 2002, right: 2003)

Confidence limits can be computed for the ratio of the rate, which is defined by the mean over all blocks divided by the time interval (Table 1), in which these counts were measured. The time period, where insects were counted, in 2002 is 99 days and in 2003 it is 89 days.

If the count means are high ($\bar{x} > 10$), the intervals are relative narrow, but if not much insects were counted, like for the bugs, especially in 2003, the intervals get non-adequate wide (Table 2). In these situations the interval limit estimates for the three

		$\bar{x}$	
Insect	Type	2002	2003
Thrips	Bt	48.000	10.500
	INS	46.000	4.375
	ISO	52.750	10.125
Aphids	Bt	2861.250	1346.250
	INS	2893.125	2688.250
	ISO	1673.750	2042.125
Bugs	Bt	22.250	4.00
	INS	4.500	3.50
	ISO	14.750	2.75

Table 1: Sample means for each year (dataset Eckert)

confidence intervals differ. The score method, proposed by Graham et al. (2003) [12], has the shortest interval span, the Wald interval range stands between the score and Jeach's interval.This reflects the simulation results by Graham et al. (2003) [12], where the score method is found near nominal level and the Wald based interval is far more conservative. But in these simulation study the interval by Jeach shows nearly the same behaviour as the score interval.

6 Generalized linear model (GLM)

When we are dealing with multiple observations per group, a general linear model is the method of choice. Multiple counts per group can be pooled by adding them up, so that confidence intervals for one parameter can be used, but these are not as accurate as a model, as they are not including any different computed standard errors than the sample means.

The assumptions for a linear least square variance analysis are variance homogeneity, independent, continuous and normal distributed data. To fit a linear model for Poisson data without violating these assumptions, a log transformation can be done. Another recommended method is the use of a generalized linear model. Components of these models are a link function, an error distribution and a variance function.

				Wald CI		Jeach CI		Score CI	
Year	Insect	Comparison	Estimate	lower	upper	lower	upper	lower	upper
2002	Thrips	Bt/Iso	0.910	0.615	1.345	0.614	1.346	0.617	1.342
		Ins/Iso	0.872	0.587	1.295	0.586	1.295	0.589	1.292
	Aphids	Bt/Iso	1.709	1.609	1.816	1.610	1.816	1.609	1.816
		Ins/Iso	1.729	1.628	1.836	1.628	1.836	1.628	1.836
	Bugs	Bt/Iso	1.508	0.781	2.913	0.785	2.961	0.790	2.880
		Ins/Iso	0.305	0.106	0.877	0.096	0.832	0.111	0.839
2003	Thrips	Bt/Iso	1.037	0.437	2.459	0.433	2.495	0.448	2.399
		Ins/Iso	0.432	0.141	1.326	0.128	1.280	0.148	1.259
	Aphids	Bt/Iso	0.659	0.615	0.706	0.615	0.706	0.615	0.706
		Ins/Iso	1.316	1.243	1.394	1.243	1.394	1.243	1.394
	Bugs	Bt/Iso	1.455	0.313	6.753	0.308	7.759	0.354	5.984
		Ins/Iso	1.273	0.262	6.175	0.251	7.013	0.299	5.421

Table 2: Confidence intervals for the ratio of rates (example Eckert)

6.1 Linear predictor

In ordinary linear models for a random variable Y the vector μ gets specified in terms of a number of unknown parameters $\beta_1, \ldots, \beta_p$, with

$$E(Y_i) = \mu_i = \sum_{1}^{p} x_{ij}\beta_j \tag{12}$$

and in matrix notation:

$$\mu = X\beta \tag{13}$$

, where μ_i is the mean of a normal distribution with constant variance σ^2 and X is the model matrix. As an example for a one factorial design with three factor levels this model formulation is represented as

$$\begin{pmatrix} y_1 \\ y_2 \\ y_3 \end{pmatrix} = \begin{pmatrix} 1 & 0 & 0 \\ 1 & 1 & 0 \\ 1 & 0 & 1 \end{pmatrix} \begin{pmatrix} int \\ \beta_1 \\ \beta_2 \end{pmatrix}$$

This is the model like it is fit in R ([22]), where the intercept of this model is the estimate for the mean of the first factor level (in one factorial designs) and the following β are estimates for the difference between the first and the next factor levels. In other statistical

packages, like SAS, the intercept of the model is denoted as the overall mean and a β exists for every factor level, representing the difference from the overall mean.
In generalized linear models the symbol η is introduced for describing the linear structure.

$$\eta = X\beta \tag{14}$$

(Dobson (1990), [8])

6.2 Link function

In classical linear models μ and η are identical. Dealing with count data a link function can be used

$$\eta = log(\mu). \tag{15}$$

In this case the log-link turns the additive effects contributing to η into multiplicative effects contributing to μ (McCullagh and Nelder (1989), [18]). For each different distribution a different link function can be used (table 3,(Everitt and Hothorn (2006), [9])).

Distribution	Canonical link
Normal	$\eta_i = \mu_i$
Poisson	$\eta_i = log(\mu_i)$
Binomial	$\eta_i = log\left(\frac{\pi_i}{1-\pi_i}\right) = log\left(\frac{\mu_i}{n_i-\mu_i}\right)$
Gamma	$\eta_i = \frac{1}{\mu_i}$
Inverse Gaussian	$\eta_i = \frac{1}{\mu_i^2}$

Table 3: Link functions for various distributions

This technique of applying a link function to the mean of the response and fitting a model by maximum likelihood is the main idea of a GLM, instead of transforming the response directly.

6.3 Variance function

The variance function $V(\mu)$ implies, how the variance of a response variable depends on its mean.

$$a(\phi) = \frac{\phi}{w} \tag{16}$$

, where ϕ is called the dispersion factor, which is constant over the observations, and w is a known prior weight, which varies over the observations (Table 4, (Everitt and Hothorn (2006), [9])) (McCullagh and Nelder (1989), [18]).

Distribution	Dispersion parameter ϕ	Variance function $V(\mu)$
Normal	σ^2	1
Poisson	1	μ
Binomial	$1/m$	$\mu(1-\mu)$
Gamma	ν^{-1}	μ^2
Inverse Gaussian	σ^2	μ^3

Table 4: Variance functions for various distributions

6.4 Maximum-likelihood estimation

Generalized linear models have the likelihood function as basis. In likelihood inference, the model that makes the observed data more probable, is said to be more likely (Lindsey (1997), [15]). A vector of parameter values in a model function, that makes the observed data most probable, is called the maximum likelihood estimate. For the assumptions $y_i \sim Poi(\mu_i)$, where $log(\mu_i) = \beta_1 x_{i1} + \ldots + \beta_p x_{ip}$, the likelihood function is

$$L^*(\beta) = L(\mu) = \prod_{i=1}^{n} \left(\frac{\mu^{y_i}}{y_i!} e^{-\mu_i} \right) \tag{17}$$

By differentiating $L(\beta)$ with respect to β and equating the derivate to zero, the values of β that maximize L are found. Maximizing L is the same as maximizing $log(L)$, which is often easier to use. So the transformed log-likelihood function is

$$l(\beta) = log(L(\mu)) = -log\left(\prod y_i!\right) + \left(\sum_{i=1}^{n} y_i\right) log(\mu) - n\mu. \tag{18}$$

The maximization has to be carried out for a permissible range of values of β. For example if β is a variance, the range of permissible values are only non-negative values. If they are not appropriate, adjustments have to be made. (McCulloch, Searle 2001, [19])

7 Confidence intervals for GLM estimates

7.1 Likelihood-ratio test

The likelihood ratio test statistic is a standard hypothesis test. It compares the maximized likelihood values of the full model with a restricted model, where some parameters are treated as nuisance parameters.

$$LR = \frac{L(\hat{\mu}_0)}{L(\hat{\mu})} = \prod_{i=1}^{n} \left(\left[\frac{\hat{\mu}_0}{\hat{\mu}} \right]^{y_i} e^{-(\hat{\mu}_{0i} - \hat{\mu}_i)} \right) \tag{19}$$

and under H_0, $-2log(LR) \sim \chi^2_p$ is approximately true. (McCulloch, Searle 2001, [19])

7.2 Profile-likelihood interval

Based on the asymptotic χ^2 distribution of the likelihood ratio test a confidence interval can be constructed. Let $l(\cdot)$ be the log likelihood function and $\hat{\theta}$ be the maximum likelihood estimate of a parameter vector $\theta \in \Theta \subseteq \Re^p$, then

$$\tilde{l}_j(\hat{\beta}) = \max_{\theta \in \Theta_j(\beta)} l(\theta). \tag{20}$$

This is the profile-likelihood function for a single parameter β, where $\Theta_j(\beta) = \{\theta \in \Theta | \theta_j = \beta\}$. In the profile-likelihood approach $l\theta$ is a function for one single parameter $\beta = \theta_j$ of interest, treating the others as nuisance parameters and maximizing over them. A $(1-\alpha)$ profile likelihood based confidence interval for θ_j is given by

$$\{\beta | 2[l(\hat{\theta}) - \tilde{l}_j(\beta)] \leq \chi^2_{1,1-\alpha}\}. \tag{21}$$

(Venzon and Moolgavkar, 1988, [24])

7.3 Profile-likelihood interval in R [22]

In R the function $confint()$ in the $library(MASS)$ computes confidence intervals for a model parameter β_p from a generalized linear model. A marginal $1-\alpha$ interval for β_p can be written as

$$-t\left(N-P; \frac{\alpha}{2}\right) \leq \delta(\beta_p) \leq t\left(N-P; \frac{\alpha}{2}\right) \tag{22}$$

with

$$\frac{\beta_p - \hat{\beta}_p}{se(\hat{\beta}_p)} = \delta(\beta_p). \tag{23}$$

For a nonlinear model the profile t function $\tau(\theta_p)$ is

$$\tau(\theta_p) = sign(\theta_p - \hat{\theta}_p)\sqrt{\frac{\tilde{S}(\theta_p) - S(\hat{\theta})}{s}} \tag{24}$$

, where

$$\tilde{S}(\beta_p) = min_{\beta_{-p}} S((\beta_p, \beta_{-p}^T)^T) = S((\beta_p, \tilde{\beta}_{-p}^T)^T). \tag{25}$$

The profile t-statistic is defined as the square root of change in sum of squares divided by residual standard error with an appropriate sign ($profile()$ function R).

Now $\tau(\theta_p)$ can be plotted against $\delta(\theta_p)$, where a likelihood interval is read off at the values of τ corresponding to the quantile of the profile t-function with degrees of freedom $N-P$ and a probability $\frac{\alpha}{2}$. These points on the profile-likelihood curve are interpolated by

a spline function fitted for the values of the profile of an GLM fit. (Bates and Watts (1988), [1])This confidence interval derivation can be illustrated by the graphical demonstration of the profile t-function for one parameter in a model. Such a graphic can be easily presented in R.

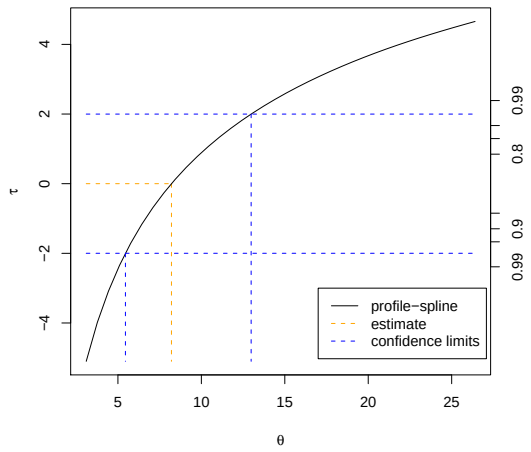

Figure 3: Plot of the profile-t-function

In Figure 3 the point estimate for the ratio is found at $\tau = 0$ and the lower and upper confidence limits are found at the corresponding θ for $\tau = \pm t\left(N - P; \frac{\alpha}{2}\right)$. The resulting confidence limits can be compared with the ones of the function $conf.int()$ from $library(MASS)$ (Table 5).

function	lower cl	upper cl
from graphic	2.639799	9.046682
confint() from library(MASS)	2.667978	8.926303

Table 5: Profile-likelihood intervals for a generated dataset

The differences in the results should be depending on different spline fits.

7.4 Wald interval

For large sample sizes Wald statistic can be found with

$$(\hat{\theta}_1 - \theta_1)'[var_\infty(\hat{\theta}_1)]^{-1}(\hat{\theta}_1 - \theta_1) \leq \chi^2_{\nu,1-\alpha} \tag{26}$$

with the corresponding confidence interval

$$cl_{l,u} = \hat{\beta} \pm z_{\frac{\alpha}{2}}\sqrt{var(\hat{\beta})}. \tag{27}$$

These intervals are easy to calculate and are provided for a wide range of methods. (McCullagh and Nelder (1989), [18])

7.5 Computing Wald intervals for the ratio of factor level marginal means derived from GLM estimates

In a one way design a linear model in R is fit with the marginal mean of the first factor level as the intercept of the model. The next model estimates β_i are the differences between the further means of each factor level with the intercept. The marginal means for each factor level are represented by a linear combination of the parameter estimates from the GLM. The model matrix X, which does not include any replications, can be used to construct linear combinations for computing the marginal means. In a design with two or more levels this can be done similarly, as the only difference to a one way model is the higher complexity of the model matrix. These linear combinations can be used to transform the variance-covariance matrix to derive standard errors for the means and their correlation. After computing the marginal means contrast tests can be used for comparing means of factor level combinations, or for example the mean of two factor levels pooled over all levels of a second factor. The integration of covariates for weighting the marginal means can also be provided. (Yandell (1997), [28])

7.6 Example

The Eckert dataset (Table 21, see page 11) is used to compute Waldintervals and profile-likelihood intervals for GLM estimates.

To simplify the model and by accounting interactions between the maize type and year the dataset was split by the factor year into two separate sets. Now the only repetitions are made over the blocks the variance of the blocks influence the residual error.

Model specification:

$$y_{ij} = type_i + \epsilon_{ij}$$

, where y is the response, type are the three maize treatments (Iso, Bt, Ins) and ϵ denotes the residual error.

The objective of the trial is to show equivalence of insect abundance for Bt- and insecticide treated maize compared to isogene maize. To obtain estimates for the ratio of the means of Iso and Bt and of Iso and Ins the factor level Iso is set as the intercept of

the model, whereby the model parameters β_1 and β_2 are the desired ratios. For verifying equivalence between the groups $(1-2\alpha)$ confidence intervals are computed for the model parameters.

				Wald CI		Profile CI	
Year	Insect	Comparison	Estimate	lower CI	upper CI	lower CI	upper CI
2002	Thrips	Bt/Iso	0.910	0.8103	1.0219	0.8102	1.0218
		Ins/Iso	0.872	0.7755	0.9806	0.7754	0.9805
	Aphids	Bt/Iso	1.7095	1.6792	1.7404	1.6792	1.7404
		Ins/Iso	1.7285	1.6979	1.7597	1.6980	1.7597
	Bugs	Bt/Iso	1.5085	1.2409	1.8337	1.2423	1.8364
		Ins/Iso	0.3051	0.2231	0.4173	0.2210	0.4139
2003	Thrips	Bt/Iso	1.0370	0.8027	1.3398	0.8026	1.3405
		Ins/Iso	0.4321	0.3098	0.6027	0.3073	0.5987
	Aphids	Bt/Iso	0.6592	0.6459	0.6728	0.6459	0.6728
		Ins/Iso	1.3164	1.2941	1.3391	1.2941	1.3391
	Bugs	Bt/Iso	1.4545	0.9223	2.2939	0.9264	2.3133
		Ins/Iso	1.2727	0.9223	2.2939	0.7984	2.0470

Table 6: Confidence intervals for GLM estimates (example Eckert)

For most of the ratios no difference between the Wald interval and the profile-likelihood interval could be observed (Table 6), this corresponds to an analysis in Giminez et al. (2005) [11]. A major difference between the intervals occurs for the bugs data in the year 2003. In this time period only rare events were counted so that the sample means are near the zero truncation. In this example the confidence bounds are lower for the profile interval than for the wald interval. As not many bugs were counted, especially in 2003, the spans of the intervals are the widest. Altogether the intervals are quite narrow, because the variation in abundance over the year is not accounted.

If comparing these intervals from the generalized linear model with these for the ratio of two Poisson parameters, the spans of the intervals derived from a model are noticeable shorter than these ones, assuming a Poisson distributed response by only defining the sample mean. At high abundance all intervals are mostly similar.

8 Simulation study for comparing two means in an one factorial design

8.1 Parameter selection

For examination of coverage probability for the profile likelihood interval and Wald interval in a one-factorial GLM, following variation in parameterization was chosen:

$$n = 5, 10, 20, 50, 100$$

$$\lambda_2 = 0.1, 1, 10$$

$$\rho = 0.1, 0.2, 0.5, 1, 2, 5, 10$$

where

$$\lambda_1 = \rho\lambda_2$$

.

Two samples of random Poisson data with mean λ_i were generated 10000 times for each parameter combination. A seed was set to ascertain comparability of the results.

8.2 Problems with extreme parameter settings

For small numbers of observation or small population means, the probability increases, that in one sample no count was observed. For these situations no profile likelihood interval can be computed, as for example while profiling the GLM model object, better fits could be obtained. Limits for the Wald interval, in cases were one sample contains only zeros, are ranging from zero to infinity. In the simulation study coverage for the profile likelihood interval and the Wald interval was only computed for samples with at least containing one count higher than zero. Additional Wald interval coverage was computed for all samples. Not properly working cases were counted.

8.3 Results

At a low number of observations and a small λ near 100% of the samples are completely filled with zeros the coverage of both intervals is near zero (Figure 4, Table 17). With increasing N the coverage is also increasing, with the likelihood-ratio interval being slightly less liberal than the Wald interval. At higher λ or ρ both intervals reach the 95% at a lower observation number, because in these situations at least one sample is not highly

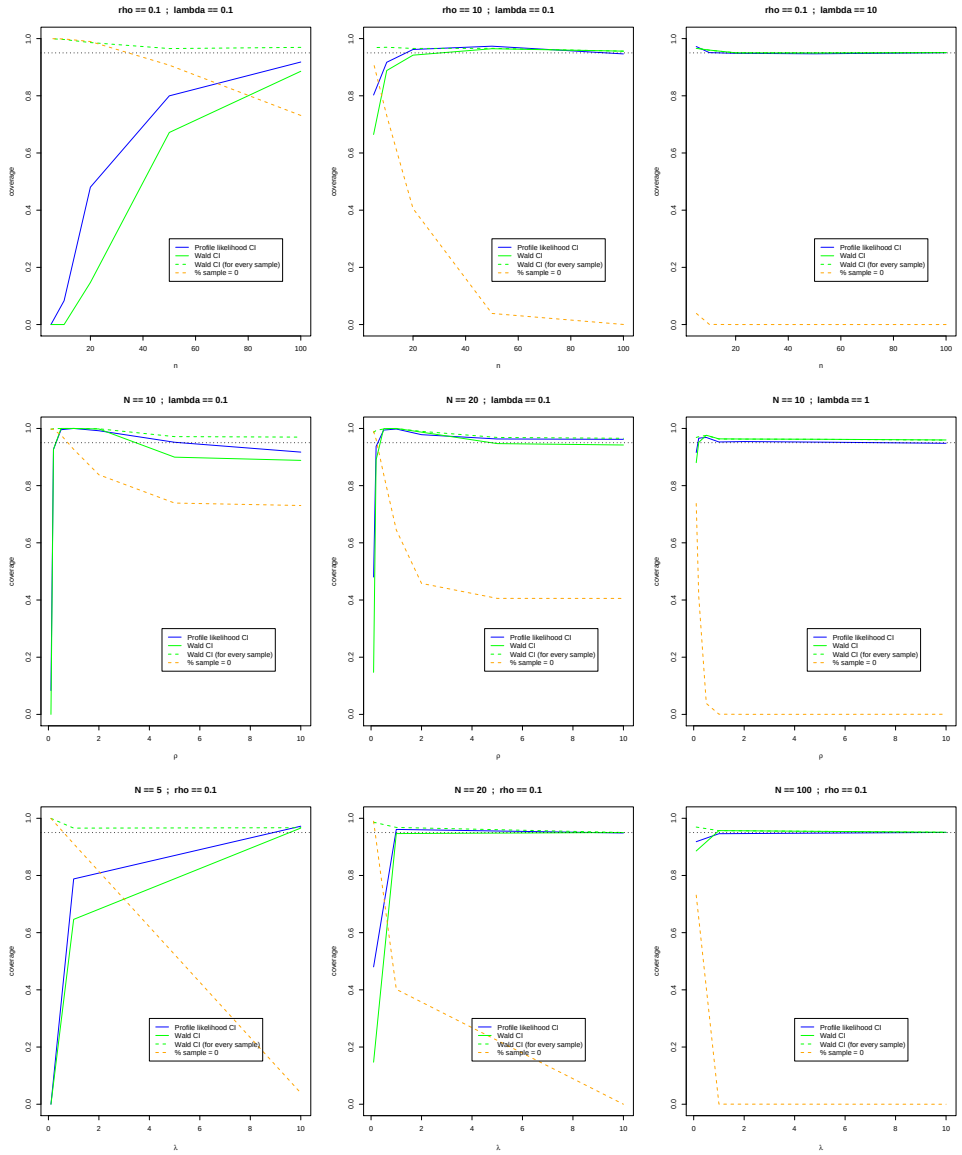

Figure 4: Coverage probability of the two-sided Wald- and the profile-likelihood confidence interval for the ratio of two sample means ($1 - \alpha = 0.95$)

zero inflated. The only observable influence of ρ on the coverage of the intervals is the influence on the height of the sample means. When increasing the difference between the observation means, coverage is located at 0.95 for both $\lambda > 0.5$. For λ near the zero truncation both intervals will perform not well. But if the sample size is large enough, the intervals will show also good results for small sample means ($N = 100$ for $\lambda_{1,2} = 0.1$). Overall the both intervals perform both well for samples following a Poisson distribution without zero inflation. For more extreme parameter settings the profile-likelihood interval has a slightly better behavior.

8.4 Bootstrap confidence interval

The Wald interval and the profile-likelihood interval were compared with a non-parametric bootstrap confidence interval for only a few situations ($\rho = 0.1, 0.2, 0.5$), as it takes some time to compute these. For each parameter setting the bootstrap interval was computed 10000 times with 100 repetitions for each bootstrap. Bootstrap intervals were computed in R [22] with the package *boot* by Angelo Canty and imported to R by Brian Ripley, where the function *boot*() generates the bootstrap replicates of a statistic (GLM) applied to the data and with the function *boot.ci*() a percentile interval is calculated from the bootstrap calculation using the full bootstrap distribution.

Giminez et al. (2005) [11] compared a bootstrap interval with a Wald and a profile-likelihood confidence interval for an example dataset in a capture-recapture model. Their suggestion was that the profile likelihood interval is a compromise between a wider Wald interval and a more narrow bootstrap interval.

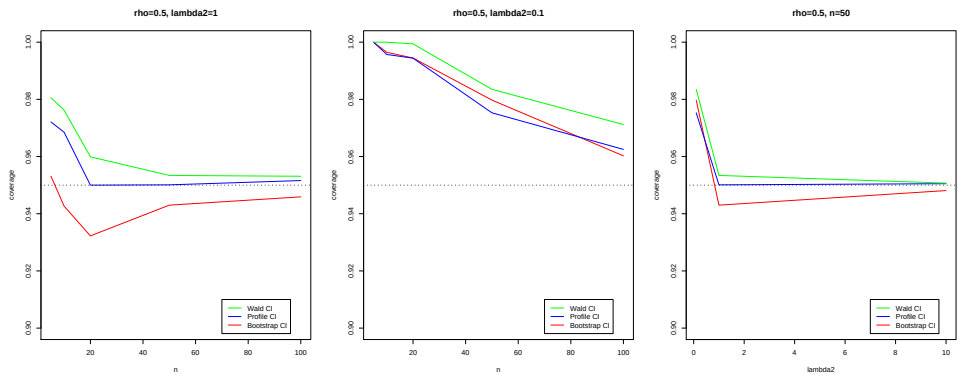

Figure 5: Coverage probability of the Bootstrap-, Wald- and the profile-likelihood interval for the ratio of two sample means

For small sample sizes and sample means, there are also problems to compute the bootstrap intervals for zero inflated data. The situations where the model did not converge, had been omitted. For very small sample sizes ($N < 10$, $\lambda_2 > 1$) the bootstrap interval gets most liberal (coverage< 0.9), the possibility to obtain adequate, exact confidence limits out of the samples is delimited (Figure 5). For parameter settings with large sample sizes and sample means ($N > 100$, $\lambda > 1$, $\rho > 0.5$) all three intervals have coverage probabilities near 95 %. If we are looking at moderate situations between these ($N \approx 50$, $\rho \approx 0.5$ and $\lambda_2 \approx 1$), the Wald interval is slightly conservative, the bootstrap interval is in most cases more liberal and the profile-likelihood interval lies between them at coverage 0.95 or near the bootstrap interval.

9 Generalized linear mixed model (GLMM)

A basic linear model can be written as

$$E\left(Y_i\right) = X\beta \tag{28}$$

with X being the model matrix and β a vector for the fixed effects. This model can be enlarged by a term of random effects:

$$E\left(Y_i|u\right) = X\beta + Zu \tag{29}$$

$$u \sim MVN(0, D) \tag{30}$$

, where Z is a model matrix, like X, for the random effects and u is the vector for the random effects. For a model with one fixed factor and a random one with three levels this can be written as

$$\begin{pmatrix} y_1 \\ y_2 \\ y_3 \end{pmatrix} = \begin{pmatrix} 1 & 0 & 0 \\ 1 & 1 & 0 \\ 1 & 0 & 1 \end{pmatrix} \begin{pmatrix} int \\ \beta_1 \\ \beta_2 \end{pmatrix} + \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} u_1 \\ u_2 \\ u_3 \end{pmatrix}$$

. u follows a multivariate normal distribution with zero mean and variance-covariance matrix D.

For a GLMM the linear predictor η can be written as

$$\eta_i = g(\mu_i) = x_i'\beta + z_i'u \tag{31}$$

with $g(\cdot)$ is a known link function, in the Poisson case the log link, which links together η_i and the conditional mean of y_i and the linear form of predictors. x_i' and z_i' are the ith rows of the model matrices of the fixed and random effects. Because the conditional distribution of y given u is just a notational extension of the generalized linear model, many of the relationships derived for the GLM will hold for a GLMM, too. (McCulloch and Searle 2001, [19])

The link function for a Poisson GLMM is, like in a GLM, the logarithmic transformation, so

$$exp(\eta_i) = \mu_i. \tag{32}$$

The same is true for the variance function: For deriving the marginal variance of y following formula can be used:

$$var(y_i) = var(E[y_i|u]) + E[var(y_i|u)]. \tag{33}$$

Assuming we have a log link and y, given u, following a Poisson distribution gives

$$var(y_i) = var(\mu_i) + E[\mu_i]. \tag{34}$$

This conditional variance for y_i given u is Poisson but the marginal distribution is not Poisson, as for assuming $u_i \sim N(0, \sigma^2)$ the variance is larger than the mean. This occurrence gives the opportunity to deal with overdispersion. (McCulloch and Searle 2001, [19])

9.1 Estimation

The estimation can be done by maximum likelihood for simple situations like a single random effect, two or three nested random effects or longitudinal data with random slopes. With more complicated structures the estimation by maximum likelihood fails.

Alternative methods of estimation are available, which can compute and maximize the likelihood in an effective way. In R for the function $lmer()$ in the package $lme4$ the penalized quasi-likelihood (PQL) method is used as standard for estimation. Quasi-likelihood has the ability to generate highly efficient estimators without making distributional assumptions, specifying only the mean to variance relationship. For the estimation of the variance-covariance structure in a generalized mixed model a penalty function is added to the quasi-likelihood. McCulloch and Searle (2001) [19] recommend only modified PQL algorithms for generalized linear mixed models, as the unmodified method has not been found to work well in practice, especially for binary data in small clusters.

9.2 Confidence intervals

For large samples Waldintervals can be constructed according to McCulloch and Searle (2001, [19]), who proposed following Wald test:

$$\frac{\hat{\beta}_i - \beta_{i,0}}{\sqrt{\hat{var}_\infty\left(\hat{\beta}_i\right)}} \sim N(0,1) \tag{35}$$

Instead of the normal distribution the t-distribution, default setting for interval construction in R ([22]), is used to calculate the intervals, which should be the same for large sample sizes.

9.3 Example

The same dataset from Eckert is used to compare insect counts in isogene- versus Bt- and insecticide treated maize. This time the simplified model gets extended by the year as a fixed or random factor. Model syntax:

$$y_{ijk} = type_i + year_j + (type : year)_{ij} + \epsilon_{ijk}$$

In a generalized linear mixed model, with the factor year is treated as random influencing the variance-covariance matrix, a cell means model can be fitted to obtain the parameter estimates. If the years are taken as fixed, a two factorial model is fitted, where the marginal means for each combination of factor levels can be used to compute the desired ratios for one factor (type) pooled over the second (year). The alculated confidence intervals are Waldintervals for the difference of the logarithmic estimates (marginal means) of the generalized models.

		GLM			GLMM		
Insect	Comparison	lower CI	Estimate	upper CI	lower CI	Estimate	upper CI
Thrips	Bt/Iso	0.8413	0.9714	1.1216	0.8351	0.9304	1.0366
	Ins/Iso	0.5125	0.6138	0.7352	0.7160	0.8012	0.8965
Aphids	Bt/Iso	1.0470	1.0616	1.0764	1.1172	1.1323	1.1476
	Ins/Iso	1.4895	1.5085	1.5276	1.4832	1.5020	1.5211
Bugs	Bt/Iso	1.1497	1.4813	1.9085	1.2485	1.5000	1.8022
	Ins/Iso	0.4671	0.6231	0.8312	0.3547	0.4571	0.5892

Table 7: Confidence intervals for GLMM estimates (example Eckert)

There are large differences between the intervals depending on different models, even between the model estimates (Table 7). The only case where nearly equal results are produced, is for the aphids data between isogene and insecticide treatment. In this situation there are no interactions between *type* and *year* and the marginal means for each year are nearly the same. If there are interactions between the blocks, the block effect can be misinterpreted by pooling over their marginal means to compare the treatment effects in a GLM. With the random effect a GLMM can have a better fit to overdispersed data, which can be caused by extra variance generated by the block effects.

10 Simulation study for comparing two factor levels in a block design

Coverage probabilities for confidence intervals for the ratio of two Poisson samples in the context of a two-factorial general linear model are computed. A blocking factor was added to the generated data, with the blocks having no effect, the means of the two blocks for each factor are equal and with the blocks having an effect, with their sample means were at $0.5\lambda_i$ and $1.5\lambda_i$. A general linear model and a general linear mixed model, with blocks as random factor, were fitted. With the extracted estimates and the variance-covariance matrix, Waldintervals were computed for the ratio of the two sample means. In case of a GLM fit the two blocks for each factor were pooled and in a GLMM the blocks are random, influencing the variance-covariance matrix. For a GLMM the confidence intervals were constructed with the point estimates for the fixed effects, a quantile of the student t distribution and a the estimated variance for the fixed effects.

10.1 Parameter setting

Following parameters for generating Poisson data were chosen:

$$n = 10, 20, 50, 100$$

$$\lambda_2 = 0.1, 1, 10$$

$$\rho = 0.1, 0.2, 0.5, 1, 2, 5, 10$$

where

$$\lambda_1 = \rho\lambda_2$$

.

The simulations had 10000 repetitions and a seed was set.

10.2 Error treatment

In cases where the whole response vector contains only zeros, the models can not converge and no confidence interval can be obtained. These events were counted and ignored at coverage computation. For small sample sizes some iteration errors can occur with the mixed model, as it is possible that the response for some factor combinations is at zero level.

10.3 Results

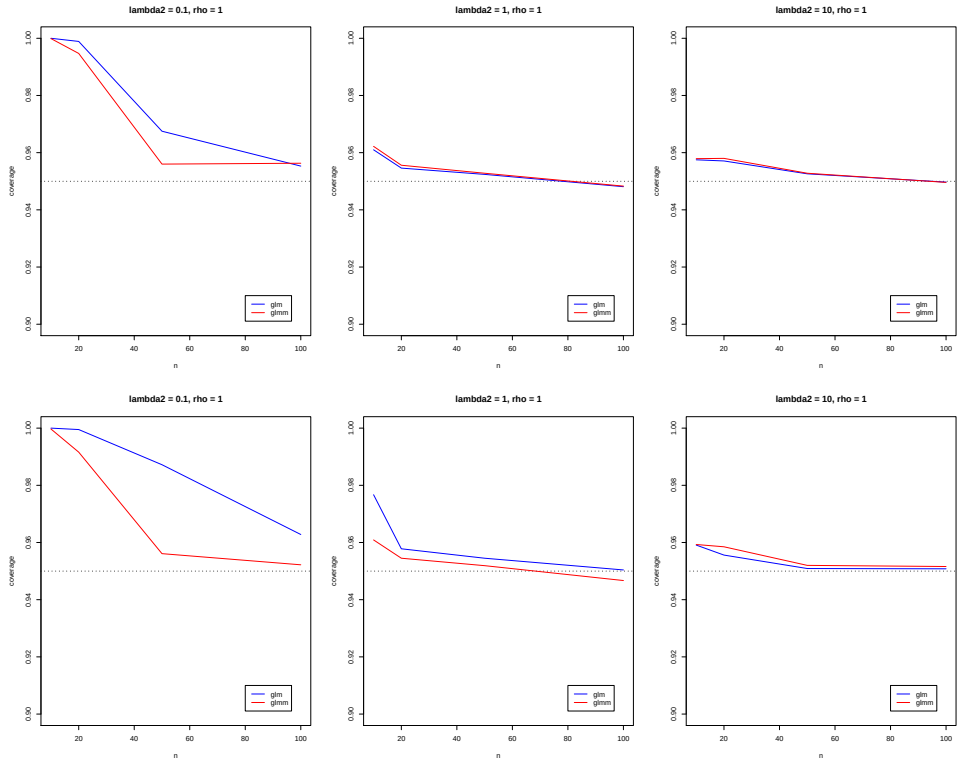

Figure 6: Coverage probability of two-sided Wald intervals ($1 - \alpha = 0.95$) from a generalized linear model and a generalized linear mixed model for the ratio of two sample means in a block design (without/with block effect, row 1/2)

If there is no block effect and λ_i are high, the intervals from the GLM achieve similar results as the GLMM intervals (Figure 7). Only for small λ_i and small sample sizes the GLMM Wald intervals produce a better coverage (Figure 6, Table 18). At these parameter settings the data could be zero inflated resulting in different sample means for each block. If a block is present or absent, there seems to be no difference for the intervals from the GLMM. For the GLM intervals get more conservative with blocks having an effect. The differences to the mixed model increase with decreasing λ and decreasing sample size.

Another setting was observed for a design comparing two groups with ten random distributed block effects each (Figure 8).

For small lambdas the intervals for GLM estimates are far too conservative. On the opposite the mixed model intervals show good coverage (≤ 0.96) probability also for small sample means. If sample means increase, coverage reaches 95% for both intervals.

It can be shown, that for $\lambda \geq 1$ and $n \geq 10$ good coverage probability can be reached for the confidence intervals of both models. If a block effect is present, a mixed model

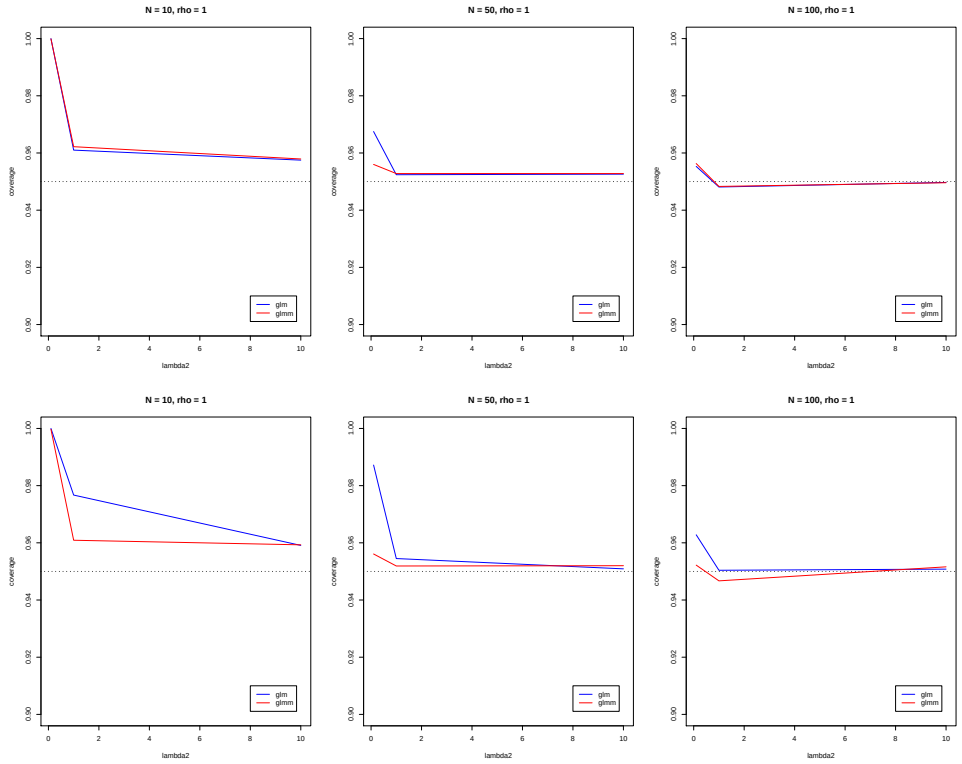

Figure 7: Coverage probability of two-sided Wald intervals ($1-\alpha = 0.95$) from a generalized linear model and a generalized linear mixed model for the ratio of two sample means in a block design (without/with block effect, row 1/2)

with the blocking factor taken as random is recommended.

11 α-Adjustment for multiplicity

If more than one comparison should be made, the probability of seeing a difference between two groups increases with the number of comparisons. A simple but conservative method of correcting this error is to divide the error α by the number of comparisons (Bonferroni method). It is evident that only these levels of a factor should be compared, which are necessary. An enhancement for α-adjustment is the inclusion of the correlation between contrasts using the multivariate normal distribution. For a many-to-one comparison this is called the Dunnett-procedure. In R this is provided by the package *multcomp* for multiple comparisons and *mvtnorm* for the multivariate t distribution. There should be no great difference between the multivariate t distribution and the multivariate normal distribution for large sample sizes. The computed confidence intervals are therefore approximate intervals.

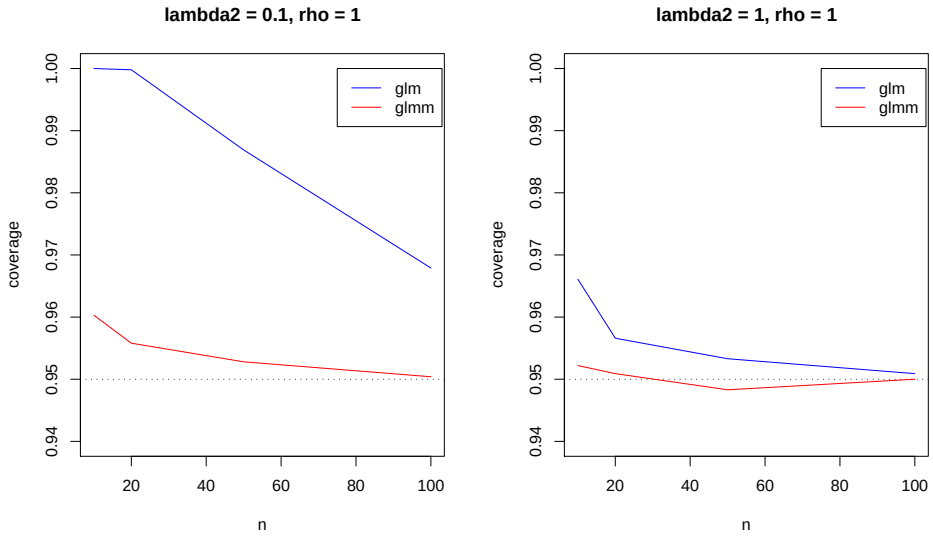

Figure 8: Coverage probability of Waldintervals from a generalized linear model and a generalized linear mixed model for the ratio of two sample means in a block design with random distributed block effects of ten block factors.

11.1 Multiple contrasts for data with heterogeneous variances

Because for Poisson distributed responses the variance depends on the mean, the variance for samples with different means should be differing, too. Therefore the assumption for the application of contrast tests to have equal variances is violated. Simultaneous confidence intervals can be constructed using the unpooled variance estimator to deal with this problem.

$$cl_{l,u} = \sum_{i=1}^{I} c_{ij}\hat{\lambda}_i \pm c_{J,1-\alpha,R}\sqrt{\sum_{i=1}^{I} c_{ij}^2 v\hat{a}r\left(\lambda_i\right)} \tag{36}$$

, where $i = 1, \ldots, I$ are the number of treatment groups and $j = 1, \ldots, J$ represent the number of simultaneous estimated contrasts. $c_{J,1-\alpha,R}$ is the J-variate normal distribution with correlation matrix R. The elements $\rho_{jj'}$ of R are defined as:

$$\rho_{jj'} = \frac{c_{ij}c_{ij'}\lambda_i}{\sqrt{\left(\sum_{i=1}^{I} c_{ij}^2\lambda_i\right)\left(\sum_{i=1}^{I} c_{ij'}^2\lambda_i\right)}}. \tag{37}$$

As the variances are different, the assumption of multivariate normal distribution is violated. Therefore in the following section it is tested by simulation, if nevertheless the approximation can be used to obtain quite good results.

In *library*(*multcomp*) the function *csimint*() compute simultaneous confidence intervals with the multivariate t distribution for provided estimates, contrast matrix and variance-covariance matrix (Bretz, Genz, Hothorn (2001), [3]).

11.2 Example

In the dataset from Eckert (Table 21) a user specified contrast can be used to compare every level of the factor *type* with the average of the two other factors. A generalized linear model with formula:

$$y_{ijk} = type_i + year_j + (type : year)_{ij} + \epsilon_{ijk}$$

, log link and family Poisson was fit. The average contrast for comparisons of the estimates of the marginal means is presented in Table 8 and the computed results are shown in Table 9.

	Bt		Ins		Iso	
Comparison	2002	2003	2002	2003	2002	2003
Bt vs. $\overline{Ins, Iso}$	-0.5	-0.5	0.25	0.25	0.25	0.25
Ins vs. $\overline{Bt, Iso}$	0.25	0.25	-0.5	-0.5	0.25	0.25
Iso vs. $\overline{Ins, Bt}$	0.25	0.25	0.25	0.25	-0.5	-0.5

Table 8: Contrast for comparing one factor level (*type*) versus the average of two others, pooled over two levels of a second factor (*year*) (Example Eckert)

			Confidence interval	
Insect	Comparison	Estimate	lower	upper
Thrips	Bt vs. $\overline{Ins, Iso}$	1.2399	1.0465	1.4689
	Ins vs. $\overline{Bt, Iso}$	0.6228	0.5066	0.7657
	Iso vs. $\overline{Ins, Bt}$	1.2950	1.0925	1.5350
Aphids	Bt vs. $\overline{Ins, Iso}$	0.8643	0.8518	0.8771
	Ins vs. $\overline{Bt, Iso}$	1.4640	1.4448	1.4836
	Iso vs. $\overline{Ins, Bt}$	0.7902	0.7788	0.8019
Bugs	Bt vs. $\overline{Ins, Iso}$	1.8765	1.4307	2.4611
	Ins vs. $\overline{Bt, Iso}$	0.5120	0.3757	0.6976
	Iso vs. $\overline{Ins, Bt}$	1.0409	0.7740	1.3998

Table 9: Simultaneous confidence intervals for Poisson model estimates and average contrast setting (Example Eckert)

For large counts (Aphids) the interval span is very narrow, but if less insects were counted and some zeros occur (Bugs) the confidence limits cover a wider range.

12 Simulation study for simultaneous confidence intervals for user defined contrasts

For some special situations coverage probability was observed for confidence intervals using the multivariate-t-distribution, Bonferroni adjusted intervals and unadjusted confidence intervals.

12.1 Parameter setting

Poisson data was generated for four or five independent samples, with their mean defined by λ and ρ_i, with $i = 1, \ldots, 5$.

$$\rho_i = 1, 1, 1.5, 2, 2.5$$

where

$$\lambda_i = \rho_i \lambda_1$$

Data were generated for:

$$n = 5, 10, 20, 50, 100$$

$$\lambda = 3, 10$$

The simulations had 10000 repetitions and a seed was set.

12.2 Single contrasts

First the behaviour of the Wald interval was observed for testing hypotheses for single linear combinations of MLEs. Every sample mean was compared with the average of the three other sample means and the average of two sample means was compared with the average of the other two.

The coverage probabilities for every contrast nearly equal each other (Figure 9, Table 19). The same results are obtained with generated samples for different means and unbalanced designs. The intervals seem to be almost unaffected by emulating some variance heterogeneity, but the main influence on coverage of these intervals has zero inflation and departure from Poisson distribution.

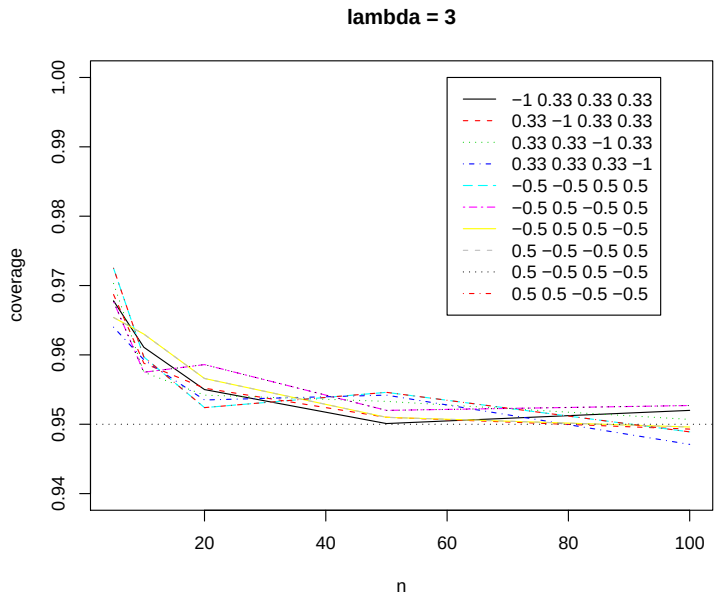

Figure 9: Coverage probability of two-sided Wald intervals ($1-\alpha = 0.95$) for single linear combinations

12.3 Coverage for different contrast settings

A Dunnett contrast (Table 10) is used to perform many-to-one comparisons, for example comparing every treatment versus a control. With the control being the first sample a Dunnett contrast for five groups is

Comparison	Control	Trt1	Trt2	Trt3	Trt4
1 vs. 2	-1	1	0	0	0
1 vs. 3	-1	0	1	0	0
1 vs. 4	-1	0	0	1	0
1 vs. 5	-1	0	0	0	1

Table 10: Dunnett contrast for four groups and a control

The intervals which are not adjusted for multiplicity are far too liberal, whereas the Bonferroni adjusted ones are conservative (Figure 10). The intervals using the multivariate t distribution for implying the correlation between the contrasts lie between the others at a coverage near above 0.95 for larger sample sizes. If decreasing the sample size, all intervals get more conservative, especially for the unadjusted intervals this effect is large. This effect will be delayed by comparing higher sample means to avoid zero inflation. If there are less numbers of groups compared, the Bonferroni-adjusted and the multivariate

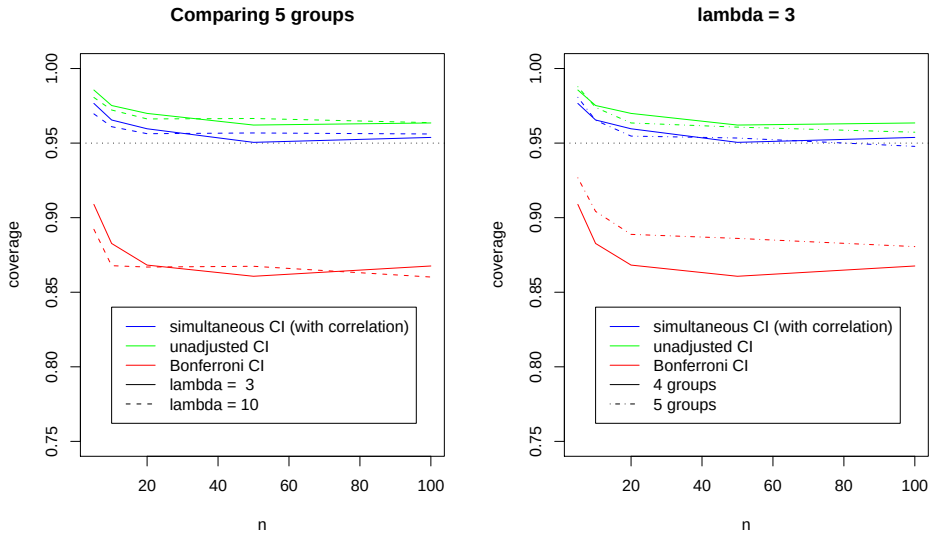

Figure 10: Coverage probability of many-to-one confidence intervals: unadjusted CIs, CIs with multivariate-t-distribution and adjusted intervals by Bonferroni method

intervals get less conservative and the unadjusted ones are less liberal. For fewer compared samples the intervals using the correlation between contrasts gain still the best results.

The Tukey contrast (Table 11) is set for doing all-pair comparisons. In difference to the Dunnett contrast there are some more linear combinations to compute and the group, which first was the control in the many to one setting, has less influence. Tukey contrast for four groups:

Comparison	Trt1	Trt2	Trt3	Trt4
1 vs. 2	-1	1	0	0
1 vs. 3	-1	0	1	0
1 vs. 4	-1	0	0	1
2 vs. 3	0	-1	1	0
2 vs. 4	0	-1	0	1
3 vs. 4	0	0	-1	1

Table 11: Tukey contrast for four groups

In an "'average"'-contrast (Table 12) every single group is compared with the average of the other groups.

Using a Williams contrast (Table 13) results in the comparison of one control with the last group, the average of the last two groups and so on.

Comparison	Trt1	Trt2	Trt3	Trt4
1 vs. $\overline{234}$	1.00	-0.33	-0.33	-0.33
2 vs. $\overline{134}$	-0.33	1.00	-0.33	-0.33
3 vs. $\overline{124}$	-0.33	-0.33	1.00	-0.33
4 vs. $\overline{123}$	-0.33	-0.33	-0.33	1.00

Table 12: Average contrast for four groups

Comparison	Control	Trt1	Trt2	Trt3
1 vs. 4	-1	0.00	0.00	1.00
1 vs. $\overline{34}$	-1	0.00	0.50	0.50
1 vs. $\overline{234}$	-1	0.33	0.33	0.33

Table 13: Williams contrast for four groups

A McDermott contrast (Table 14) is like a reversed Williams contrast comparing the averages of the first, second, third, ... group with the following one.

Comparison	Trt1	Trt2	Trt3	Trt4
1 vs. 2	-1.00	1.00	0.00	0
$\overline{12}$ vs. 3	-0.50	-0.50	1.00	0
$\overline{123}$ vs. 4	-0.33	-0.33	-0.33	1

Table 14: McDermott contrast for four groups

The unadjusted intervals are, as expected, far too liberal with the average contrast at the lowest coverage probability, due to making four comparisons instead of three like the other contrasts (Figure 11). With the chosen parameters the Williams contrast shows the best performance, since the unadjusted intervals perform not such liberal at all, the performance of the Bonferroni adjustment works relatively poor. The intervals involving the correlation between contrasts perform overall well, with the Williams intervals being also better at small sample sizes.

13 Diagnostics for generalized linear models

To explore the adequacy of a model fit a look at the residuals can be useful. For Gaussian models the residuals are $\hat{\varepsilon} = y - \hat{\mu}$. In GLMs the variance of the response is not constant.

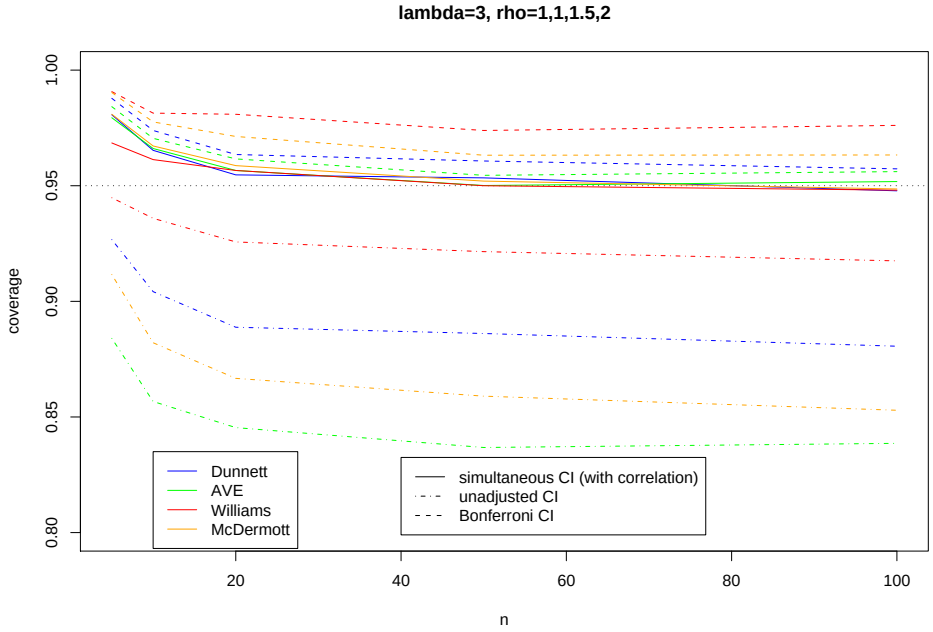

Figure 11: Coverage probability of confidence intervals for different contrasts: unadjusted CIs, CIs with multivariate-t-distribution and adjusted intervals by Bonferroni method

For a similar use, like in a Gaussian linear model, the Pearson residual $y-\hat{\mu}$ gets rescaled:

$$r_P = \frac{y-\hat{\mu}}{\sqrt{V(\hat{\mu})}}. \tag{38}$$

The deviance residuals are defined in analogy:

$$r_D = sign(y-\hat{\mu})\sqrt{d_i}. \tag{39}$$

For the Poisson case, this is

$$r_D = sign(y-\hat{\mu})\sqrt{2\left(y\frac{logy}{\hat{\mu}} - y + \hat{\mu}\right)}. \tag{40}$$

The residuals can be plotted against the predicted values. For detecting outliers, half-normal plots can be created for the jacknife residuals (Figure 12).

For the Gaussian linear model an increasing variance with increasing $\hat{\mu}$ can be observed. This would violate the requirements of variance homogeneity for this model. Because for Poisson data the variance completely depends on the mean, the plot of the deviance residuals of the Poisson GLM shows constant variances over the whole range of $\hat{\mu}$, as the variance function has been already scaled out. By inspecting the response residuals we see the same picture as for the Gaussian model.

(Faraway (2006), [10])

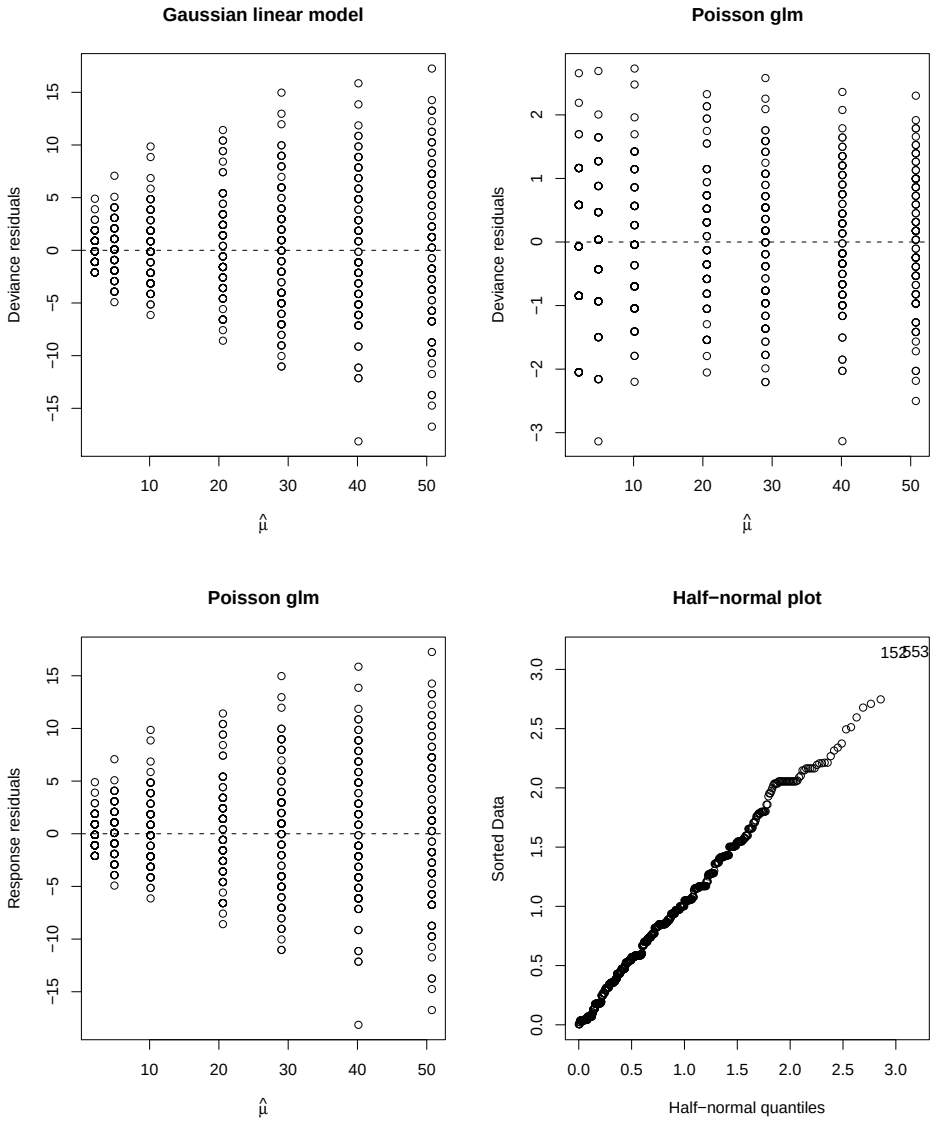

Figure 12: Diagnostic plots for generalized linear models

14 Models for zero inflated data

Often abundance data have many zeros (sometimes until 80 percent of all values), so the Poisson distribution is violated. If using a logarithmic transformation, small values like 0.1 are added up on the response vector to make the transformation possible. Because in generalized linear models the linear predictor is transformed by a link function, these problems do not arise. But if many zeros appear, a normal Poisson model provides no adequate fit. A possible way to deal with this problem is a zero-inflated Poisson model (ZIP). It assumes, that with probability p the only possible observation is zero, and with probability $1 - p$ a $Poisson(\lambda)$ variable is observed, so that

$$
\begin{aligned}
Y_i = 0 \quad & \text{with probability } p_i + (1 - p_i)e^{-\lambda_i} \text{ and} \\
Y_i = k \quad & \text{with probability } (1 - p_i)\frac{e^{-\lambda_i}\lambda_i^k}{k!} \\
& k = 1, 2, \ldots
\end{aligned}
$$

(Lambert (1992) [14]).

Problems can occur with these models, if zero-inflation appears together with overdispersion, what often could happen. For these situation a zero-inflated negative binomial model (ZINB) turned out to be more reliable (Yau et al. (2003), [29]).

Zero-inflated Poisson models and zero-inflated negative binomial models can be fit in R with functions of the package *zicounts*. These results in two sets of estimates: one set for the Poisson or negative binomial part and a second set for the zero-inflated part.

15 Overdispersion

In the Poisson distribution the variance is related to the mean ($V(\mu) = \mu$). If there where e.g. several events recorded on the same unit, the observations are not independent, so the variance could be greater than μ. Overdispersion could also arise, if some variables cause differences in mean for different observations, which are not integrated into the model. Another reason for overdispersion is the dependency of one observation on the former observation in a time series (Lindsey (1997), [15].

A solution to solve the problem of overdispersion is to assume a random effects model, where the parameter which varies in an unknown way has a random distribution. A second way to deal with overdispersion is assuming different distributions.

15.1 Mixed model

One can deal with overdispersion by using a generalized linear mixed model, treating a factor as random, to integrate the evoked overdispersion by this factor into the model (see section 9, p.23).

15.2 Negative binomial distribution

In many cases a gamma mixing distribution, which leads to a negative binomial distribution, is considered for the observed data (Breslow (1990) [2], Warton (2005) [25]). These models can produce an adequate fit to overdispersed count data, with a variance function

$V(\mu) = \mu + \phi\mu^2$. The probability mass function of the two parameter negative binomial family is written as

$$p(x|\kappa, \lambda) = \frac{\Gamma\left(x + \frac{1}{\kappa}\right) \lambda^x}{k^{\frac{1}{\kappa}} \Gamma\left(\frac{1}{\kappa}\right) \Gamma(x+1) \left(\lambda + \frac{1}{\kappa}\right)^{x+\frac{1}{\kappa}}} \tag{41}$$

(McCullagh and Nelder (1989), [18]).

A negative binomial generalized linear model can be fit in R [22] with the function *glm.nb*() or with the family *negative.binomial*, where the specification of the additional parameter θ is needed, while in *glm.nb*() it is estimated from the data.

15.3 Quasilikelihood

The quasi-Poisson model allows the estimation of the dispersion factor ϕ, instead of assuming it to be one, like in the Poisson distribution.

$$V(\mu; \phi) = \phi V(\mu)$$

In this model one parameter more has to be estimated, but it can provide a better fit than the ordinary Poisson GLM. In the quasi-likelihood models the parameter ϕ still is fixed at a certain position. In more complex models ϕ has not to be a constant, as it can vary in a systematic way with other covariates, for example as a quadratic term (McCullagh and Nelder (1989),[18]).

15.3.1 Example

The data from Eckert could be used to estimate the dispersion parameter ϕ by a quasipoisson model to determine if overdispersion is present in this data or if the Poisson distribution can be assumed.

Three models are fitted:

- Model 1: $y_{ijk} = type_i + year_j + (type : year)_{ij} + \epsilon_{ijk}$, with Poisson family and log-link
- Model 2: $y_{ijk} = type_i + year_j + (type : year)_{ij} + \epsilon_{ijk}$, with quasi family, log-link and the dispersion factor of the variance function estimated from the observed data.
- Model 3: $y_{ijk} = type_i + year_j + (type : year)_{ij} + \epsilon_{ijk}$, with negative binomial family, log link.

Type	ϕ
Thrips	7.941
Aphids	716.025
Bugs	3.322

Table 15: Dispersion parameters for different quasipoisson model fits (Example Eckert)

For the normal Poisson model the dispersion parameter is taken to be 1 and for the quasipoisson model the estimated ϕ are shown in Table 15.

In the QQ-Plots in Figure 13 it can be ssen, that the Poisson model does not fit well to the thrips data, but with a negative binomial model an adequate fit can be obtained.

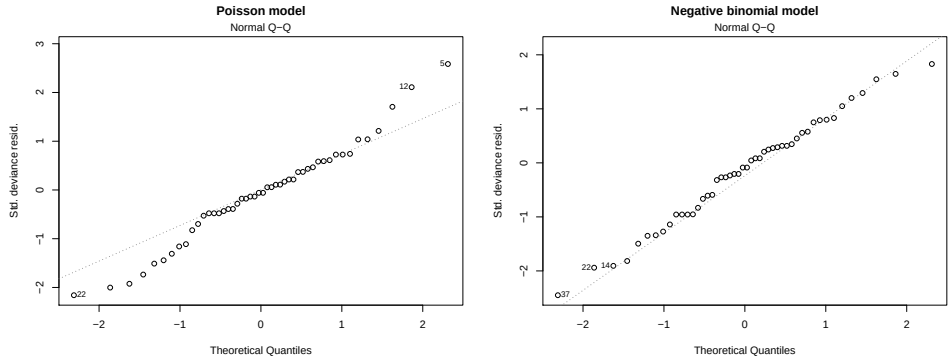

Figure 13: QQ-Plots for the Poisson and the negative binomial model, $Type$ = Thrips (Example Eckert)

These differences in variance assumptions can also be shown, if comparing the related Wald intervals (Table 16).

As the dispersion parameter and therewith the standard errors for the parameter estimates change, the intervals for the quasipoisson model and for the negative binomial model are much wider than the normal Poisson intervals. The differences are largest, if the deviation from the Poisson distribution is greatest. Further, except for one case, the quasi-model contains the 1 (quotient, where both samples are equal) in the interval, while in the Poisson intervals it is only a single time present. Because the negative binomial model shows a better fit to the data and due to the presence of overdispersion, variance functions differing from the Poisson distribution can be assumed, estimates of the quasi or the negative binomial model should be used to calculate confidence intervals.

			Poisson		Quasipoisson		Negative binomial	
Insect	Comparison	Estimate	lower	upper	lower	upper	lower	upper
Thrips	Bt/ISO	0.9714	0.8413	1.1216	0.6478	1.4566	0.7103	1.3286
	INS/ISO	0.6138	0.5125	0.7352	0.3692	1.0205	0.4401	0.8561
Aphid	Bt/ISO	1.0616	1.0470	1.0764	0.7322	1.5391	0.7664	1.4704
	INS/ISO	1.5085	1.4895	1.5276	1.0758	2.1151	1.0891	2.0893
Bugs	Bt/ISO	1.4813	1.1497	1.9085	0.9334	2.3508	0.9814	2.2357
	INS/ISO	0.6231	0.4671	0.8312	0.3685	1.0536	0.4027	0.9643

Table 16: Confidence intervals for Poisson model, Quasipoisson model and negative binomial model estimates (Example Eckert)

16 General discussion

To compare multiple treatments in experiment designs with multiple factors a generalized linear model framework should be chosen. Considering the analysis of count data, a log link is recommended and, neglecting overdispersion and zero-inflation, the Poisson distribution is suitable for this data. The performance of confidence intervals for the ratio (difference in the logarithmic transformation) of model estimates was observed by simulation to predict their benefit for testing of equivalence in practical application. For small sample sizes and very small sample means near the zero truncation no useful confidence limits are obtained. In some extreme situations the profile likelihood confidence intervals perform slightly better than the Wald interval, but the overall behaviour of both intervals is adequate for normal parameter settings ($\lambda > 1$, $n > 10$). The advantage of the Wald interval is its easy calculation and its transferability to a wide range of methods. In more complex designs, here a quite simple block design, a generalized linear mixed model is superior to a normal GLM, so its usage is recommended. If comparing more than one treatment, alpha-adjustment for multiplicity has to be made. For some settings it is shown, that existing methods, which are incorporating the correlation between multiple contrasts by the multivariate-t-distribution, are transferable to confidence intervals for linear combinations of MLEs of generalized linear Poisson models. If the assumption of Poisson distributed data is violated by the occurrence of zero-inflation or overdispersion, the calculated confidence intervals will lose in precision, as the model fits not appropriate to the data. A solution of this problem is the choice of a different model, for example a

zero-inflated Poisson model, a quasi-likelihood model or a negative binomial model.

17 Tables

17.1 Confidence intervals for a parameter of a GLM

			coverage				
$\lambda 2$	n	ρ	Profile likelihood interval	Wald interval (every sample)	Wald interval	Bootstrap interval	% = 0
0.1	100	0.1	0.9181	0.9693	0.8857	0.985	0.7313
0.1	50	0.1	0.7998	0.9652	0.6717	0.981	0.9071
0.1	20	0.1	0.4804	0.9862	0.1471	0.991	0.9898
0.1	10	0.1	0.0833	0.9968	0	1.000	0.9988
0.1	5	0.1	0	0.9997	0	1.000	0.9997
1	100	0.1	0.9461	0.9567	0.9569	0.936	3.00E-04
1	50	0.1	0.9734	0.967	0.9657	0.963	0.0373
1	20	0.1	0.961	0.968	0.9465	0.986	0.4018
1	10	0.1	0.9158	0.9685	0.8804	0.967	0.7375
1	5	0.1	0.7879	0.9656	0.6465	0.952	0.9109
10	100	0.1	0.9506	0.9511	0.9511	0.949	0
10	50	0.1	0.9469	0.9499	0.9499	0.946	0
10	20	0.1	0.9486	0.9503	0.9503	0.927	0
10	10	0.1	0.951	0.9594	0.9599	0.904	5.00E-04
10	5	0.1	0.9723	0.967	0.9659	0.875	0.0393
0.1	100	0.2	0.9646	0.9703	0.95	0.983	0.4061
0.1	50	0.2	0.9432	0.9695	0.8953	0.977	0.7411
0.1	20	0.2	0.9372	0.9931	0.8901	0.977	0.9618
0.1	10	0.2	0.9268	0.9993	0.9268	1.000	0.9959
0.1	5	0.2	0.7778	0.9998	0.7778	0.857	0.9991
1	100	0.2	0.9474	0.9549	0.9549	0.944	0
1	50	0.2	0.9466	0.9572	0.9575	0.939	3.00E-04
1	20	0.2	0.971	0.9664	0.963	0.973	0.0908
1	10	0.2	0.9661	0.9712	0.9513	0.972	0.4132
1	5	0.2	0.945	0.9706	0.8943	0.938	0.7436
10	100	0.2	0.9492	0.9496	0.9496	0.950	0
10	50	0.2	0.9523	0.9528	0.9528	0.945	0
10	20	0.2	0.9494	0.951	0.951	0.931	0
10	10	0.2	0.9469	0.9535	0.9535	0.914	0
10	5	0.2	0.9471	0.9577	0.9581	0.867	5.00E-04
0.1	100	0.5	0.9625	0.9702	0.9712	0.960	0.0392
0.1	50	0.5	0.9753	0.9849	0.9835	0.980	0.3194
0.1	20	0.5	0.9944	0.9982	0.9994	0.994	0.8394
0.1	10	0.5	0.9957	0.9998	1	0.996	0.9769
0.1	5	0.5	1	1	1	1.000	0.9975
1	100	0.5	0.9516	0.9531	0.9531	0.946	0
1	50	0.5	0.9501	0.9534	0.9534	0.943	0
1	20	0.5	0.95	0.9597	0.9599	0.932	3.00E-04
1	10	0.5	0.9685	0.9757	0.9762	0.943	0.0383
1	5	0.5	0.9721	0.983	0.9806	0.953	0.3053
10	100	0.5	0.948	0.9487	0.9487	0.949	0
10	50	0.5	0.9505	0.9507	0.9507	0.948	0

$\lambda 2$	n	ρ	Profile likelihood interval	Wald interval (every sample)	Wald interval	Bootstrap interval	% = 0
10	20	0.5	0.9468	0.9478	0.9478	0.934	0
10	10	0.5	0.9529	0.9541	0.9541	0.925	0
10	5	0.5	0.947	0.9493	0.9493	0.880	0
0.1	100	1	0.9456	0.956	0.9562	0.949	6.00E-04
0.1	50	1	0.9634	0.9809	0.9895	0.971	0.075
0.1	20	1	0.998	0.9996	1	0.998	0.645
0.1	10	1	1	1	1	0.999	0.927
0.1	5	1	1	1	1	1.000	0.9917
1	100	1	0.9507	0.9519	0.9519	NA	0
1	50	1	0.9518	0.9541	0.9541	NA	0
1	20	1	0.9477	0.9544	0.9544	NA	0
1	10	1	0.9527	0.963	0.9636	NA	7.00E-04
1	5	1	0.961	0.9813	0.9881	NA	0.0772
10	100	1	0.9518	0.9518	0.9518	NA	0
10	50	1	0.9501	0.9501	0.9501	NA	0
10	20	1	0.9506	0.9509	0.9509	NA	0
10	10	1	0.9496	0.9501	0.9501	NA	0
10	5	1	0.9475	0.9499	0.9499	NA	0
0.1	100	2	0.9473	0.9557	0.9559	NA	3.00E-04
0.1	50	2	0.9644	0.9714	0.9722	NA	0.0391
0.1	20	2	0.9782	0.9898	0.9871	NA	0.4581
0.1	10	2	0.992	0.9988	0.9981	NA	0.838
0.1	5	2	1	1	1	NA	0.9772
1	100	2	0.9506	0.9513	0.9513	NA	0
1	50	2	0.9513	0.9534	0.9534	NA	0
1	20	2	0.9486	0.9528	0.9528	NA	0
1	10	2	0.9542	0.9628	0.963	NA	2.00E-04
1	5	2	0.9656	0.9738	0.9746	NA	0.0406
10	100	2	0.9504	0.9507	0.9507	NA	0
10	50	2	0.9509	0.951	0.951	NA	0
10	20	2	0.9514	0.9514	0.9514	NA	0
10	10	2	0.9451	0.9454	0.9454	NA	0
10	5	2	0.9485	0.9493	0.9493	NA	0
0.1	100	5	0.9455	0.9561	0.9563	NA	3.00E-04
0.1	50	5	0.9689	0.9666	0.9659	NA	0.0388
0.1	20	5	0.9631	0.9685	0.947	NA	0.4057
0.1	10	5	0.9517	0.9714	0.8996	NA	0.7391
0.1	5	5	0.9368	0.9906	0.8983	NA	0.9351
1	100	5	0.9499	0.951	0.951	NA	0
1	50	5	0.9515	0.9531	0.9531	NA	0
1	20	5	0.9469	0.9506	0.9506	NA	0
1	10	5	0.9521	0.962	0.9622	NA	2.00E-04
1	5	5	0.9701	0.9685	0.968	NA	0.0402
10	100	5	0.951	0.9508	0.9508	NA	0
10	50	5	0.95	0.9501	0.9501	NA	0
10	20	5	0.9492	0.9491	0.9491	NA	0
10	10	5	0.9515	0.952	0.952	NA	0
10	5	5	0.946	0.9493	0.9493	NA	0

$\lambda 2$	n	ρ	Profile likelihood interval	Wald interval (every sample)	Wald interval	Bootstrap interval	% = 0
0.1	100	10	0.9465	0.956	0.9562	NA	3.00E-04
0.1	50	10	0.9735	0.966	0.9646	NA	0.0388
0.1	20	10	0.962	0.9656	0.9421	NA	0.4056
0.1	10	10	0.9173	0.9697	0.8883	NA	0.7305
0.1	5	10	0.8028	0.9686	0.6647	NA	0.9138
1	100	10	0.95	0.9516	0.9516	NA	0
1	50	10	0.9514	0.9521	0.9521	NA	0
1	20	10	0.9499	0.9504	0.9504	NA	0
1	10	10	0.9481	0.9592	0.9597	NA	5.00E-04
1	5	10	0.9725	0.9658	0.9644	NA	0.0393
10	100	10	0.9506	0.9503	0.9503	NA	0
10	50	10	0.9518	0.9513	0.9513	NA	0
10	20	10	0.9526	0.953	0.953	NA	0
10	10	10	0.9492	0.9513	0.9513	NA	0
10	5	10	0.9489	0.9499	0.9499	NA	0

Table 17: Coverage probability of Waldintervals, profile-likelihood-intervals and bootstrap intervals for estimates of a generalized linear model, comparison of two samples

17.2 Block design

			coverage					
			- block effects		+ block effects		random blocks	
$\lambda 2$	n	ρ	GLM	GLMM	GLM	GLMM	GLM	GLMM
0.1	100	0.1	0.9768	0.9683	0.9785	0.9670	0.9987	0.9559
0.1	50	0.1	0.9762	0.9641	0.9797	0.9650	1.0000	0.9700
0.1	20	0.1	0.9884	0.9671	0.9933	0.9679	1.0000	0.9654
0.1	10	0.1	0.9979	0.9841	0.9970	0.9812	1.0000	0.9686
1	100	0.1	0.9574	0.9547	0.9671	0.9559	0.9576	0.9530
1	50	0.1	0.9680	0.9577	0.9728	0.9569	0.9797	0.9518
1	20	0.1	0.9776	0.9705	0.9736	0.9654	0.9958	0.9519
1	10	0.1	0.9790	0.9711	0.9790	0.9710	0.9988	0.9593
10	100	0.1	0.9517	0.9510	0.9497	0.9470	0.9507	0.9511
10	50	0.1	0.9517	0.9527	0.9528	0.9561	0.9529	0.9497
10	20	0.1	0.9574	0.9578	0.9588	0.9542	0.9517	0.9501
10	10	0.1	0.9604	0.9589	0.9724	0.9613	0.9587	0.9531
0.1	100	0.2	0.9766	0.9680	0.9778	0.9659	0.9961	0.9541
0.1	50	0.2	0.9809	0.9677	0.9838	0.9652	0.9991	0.9578
0.1	20	0.2	0.9936	0.9836	0.9969	0.9844	1.0000	0.9668
0.1	10	0.2	0.9980	0.9946	0.9996	0.9938	1.0000	0.9712
1	100	0.2	0.9511	0.9510	0.9570	0.9510	0.9508	0.9520
1	50	0.2	0.9536	0.9526	0.9679	0.9547	0.9593	0.9504
1	20	0.2	0.9740	0.9652	0.9760	0.9667	0.9840	0.9522
1	10	0.2	0.9791	0.9761	0.9821	0.9761	0.9963	0.9555
10	100	0.2	0.9521	0.9523	0.9476	0.9518	0.9455	0.9493
10	50	0.2	0.9545	0.9536	0.9549	0.9559	0.9518	0.9504
10	20	0.2	0.9574	0.9586	0.9566	0.9534	0.9566	0.9552

			- block effects		+ block effects		random blocks	
$\lambda 2$	n	ρ	GLM	GLMM	GLM	GLMM	GLM	GLMM
10	10	0.2	0.9574	0.9572	0.9633	0.9597	0.9567	0.9538
0.1	100	0.5	0.9614	0.9571	0.9780	0.9580	0.9784	0.9536
0.1	50	0.5	0.9808	0.9716	0.9892	0.9712	0.9952	0.9524
0.1	20	0.5	0.9977	0.9908	0.9985	0.9911	1.0000	0.9585
0.1	10	0.5	0.9997	0.9986	1.0000	0.9992	1.0000	0.9761
1	100	0.5	0.9464	0.9476	0.9484	0.9512	0.9500	0.9479
1	50	0.5	0.9503	0.9522	0.9541	0.9557	0.9556	0.9515
1	20	0.5	0.9566	0.9576	0.9684	0.9584	0.9595	0.9481
1	10	0.5	0.9695	0.9627	0.9820	0.9647	0.9803	0.9527
10	100	0.5	0.9520	0.9513	0.9486	0.9509	0.9499	0.9507
10	50	0.5	0.9548	0.9549	0.9502	0.9557	0.9516	0.9488
10	20	0.5	0.9551	0.9561	0.9527	0.9527	0.9499	0.9523
10	10	0.5	0.9556	0.9557	0.9610	0.9571	0.9567	0.9510
0.1	100	1	0.9553	0.9563	0.9628	0.9522	0.9679	0.9504
0.1	50	1	0.9675	0.9560	0.9872	0.9561	0.9869	0.9528
0.1	20	1	0.9989	0.9947	0.9995	0.9916	0.9998	0.9558
0.1	10	1	1.0000	0.9999	1.0000	0.9997	1.0000	0.9603
1	100	1	0.9481	0.9483	0.9504	0.9467	0.9509	0.9500
1	50	1	0.9524	0.9528	0.9545	0.9519	0.9533	0.9483
1	20	1	0.9546	0.9556	0.9578	0.9545	0.9566	0.9509
1	10	1	0.9610	0.9622	0.9767	0.9609	0.9661	0.9522
10	100	1	0.9497	0.9496	0.9508	0.9516	0.9491	0.9500
10	50	1	0.9526	0.9528	0.9509	0.9520	0.9507	0.9475
10	20	1	0.9571	0.9580	0.9556	0.9585	0.9521	0.9509
10	10	1	0.9575	0.9579	0.9591	0.9593	0.9546	0.9527
0.1	100	2	0.9520	0.9533	0.9586	0.9472	0.9637	0.9526
0.1	50	2	0.9648	0.9599	0.9767	0.9622	0.9806	0.9522
0.1	20	2	0.9879	0.9746	0.9932	0.9763	0.9979	0.9564
0.1	10	2	0.9981	0.9934	0.9993	0.9923	1.0000	0.9602
1	100	2	0.9477	0.9485	0.9501	0.9508	0.9505	0.9485
1	50	2	0.9523	0.9536	0.9543	0.9528	0.9510	0.9493
1	20	2	0.9550	0.9554	0.9582	0.9551	0.9548	0.9494
1	10	2	0.9631	0.9626	0.9722	0.9614	0.9628	0.9513
10	100	2	0.9496	0.9494	0.9476	0.9505	0.9494	0.9494
10	50	2	0.9507	0.9510	0.9487	0.9510	0.9516	0.9490
10	20	2	0.9530	0.9536	0.9553	0.9563	0.9548	0.9499
10	10	2	0.9580	0.9581	0.9585	0.9614	0.9538	0.9495
0.1	100	5	0.9532	0.9522	0.9619	0.9501	0.9592	0.9495
0.1	50	5	0.9655	0.9598	0.9707	0.9577	0.9784	0.9536
0.1	20	5	0.9755	0.9680	0.9758	0.9673	0.9954	0.9544
0.1	10	5	0.9837	0.9759	0.9883	0.9729	0.9989	0.9562
1	100	5	0.9479	0.9484	0.9478	0.9504	0.9509	0.9486
1	50	5	0.9541	0.9546	0.9540	0.9513	0.9553	0.9475
1	20	5	0.9548	0.9536	0.9587	0.9543	0.9560	0.9516
1	10	5	0.9642	0.9627	0.9727	0.9613	0.9584	0.9497
10	100	5	0.9496	0.9493	0.9555	0.9526	0.9499	0.9452
10	50	5	0.9550	0.9554	0.9508	0.9497	0.9526	0.9500
10	20	5	0.9565	0.9567	0.9511	0.9507	0.9517	0.9493

			- block effects		+ block effects		random blocks	
$\lambda 2$	n	ρ	GLM	GLMM	GLM	GLMM	GLM	GLMM
10	10	5	0.9593	0.9583	0.9591	0.9578	0.9504	0.9517
0.1	100	10	0.9527	0.9503	0.9650	0.9504	0.9605	0.9510
0.1	50	10	0.9670	0.9589	0.9725	0.9590	0.9790	0.9535
0.1	20	10	0.9741	0.9659	0.9724	0.9653	0.9962	0.9539
0.1	10	10	0.9808	0.9735	0.9821	0.9699	0.9987	0.9572
1	100	10	0.9480	0.9484	0.9513	0.9509	0.9541	0.9521
1	50	10	0.9521	0.9531	0.9558	0.9538	0.9489	0.9453
1	20	10	0.9587	0.9597	0.9588	0.9527	0.9479	0.9464
1	10	10	0.9606	0.9588	0.9706	0.9609	0.9582	0.9512
10	100	10	0.9501	0.9499	0.9495	0.9528	0.9516	0.9514
10	50	10	0.9494	0.9494	0.9519	0.9517	0.9532	0.9520
10	20	10	0.9525	0.9519	0.9545	0.9568	0.9556	0.9545
10	10	10	0.9590	0.9590	0.9565	0.9603	0.9494	0.9511

Table 18: Coverage probability of Waldintervals for GLM and GLMM estimates, comparing two factor levels in a block design

17.3 Simultaneous confidence intervals

17.3.1 Single linear combinations

	contrast				coverage (for different sample sizes)				
λ	grp 1	grp 2	grp 3	grp 4	5	10	20	50	100
3,3,3,3	-1	0.33	0.33	0.33	0.9678	0.9611	0.955	0.9501	0.952
	0.33	-1	0.33	0.33	0.9687	0.9588	0.9552	0.951	0.9493
	0.33	0.33	-1	0.33	0.9703	0.9575	0.9542	0.9533	0.9507
	0.33	0.33	0.33	-1	0.964	0.9593	0.9535	0.9542	0.9471
	-0.5	-0.5	0.5	0.5	0.9725	0.9597	0.9524	0.9546	0.9489
	-0.5	0.5	-0.5	0.5	0.9677	0.9575	0.9586	0.952	0.9527
	-0.5	0.5	0.5	-0.5	0.9654	0.963	0.9566	0.951	0.9496
	0.5	-0.5	-0.5	0.5	0.9654	0.963	0.9566	0.951	0.9496
	0.5	-0.5	0.5	-0.5	0.9677	0.9575	0.9586	0.952	0.9527
	0.5	0.5	-0.5	-0.5	0.9725	0.9597	0.9524	0.9546	0.9489
3,15,15,15	-1	0.33	0.33	0.33	0.9684	0.9618	0.9582	0.9498	0.948
	0.33	-1	0.33	0.33	0.969	0.9601	0.9608	0.952	0.9504
	0.33	0.33	-1	0.33	0.9689	0.9586	0.9549	0.9496	0.9494
	0.33	0.33	0.33	-1	0.9662	0.9613	0.9578	0.9526	0.953
	-0.5	-0.5	0.5	0.5	0.9693	0.9636	0.9577	0.9502	0.9499
	-0.5	0.5	-0.5	0.5	0.9675	0.9599	0.9605	0.952	0.9505
	-0.5	0.5	0.5	-0.5	0.969	0.9629	0.9556	0.9524	0.9478
	0.5	-0.5	-0.5	0.5	0.969	0.9629	0.9556	0.9524	0.9478
	0.5	-0.5	0.5	-0.5	0.9675	0.9599	0.9605	0.952	0.9505
	0.5	0.5	-0.5	-0.5	0.9693	0.9636	0.9577	0.9502	0.9499
15,15,15,3	-1	0.33	0.33	0.33	0.9659	0.9565	0.9544	0.9546	0.9481
	0.33	-1	0.33	0.33	0.9652	0.9613	0.9554	0.9461	0.9518
	0.33	0.33	-1	0.33	0.9673	0.9588	0.955	0.9497	0.9476
	0.33	0.33	0.33	-1	0.9693	0.9587	0.9562	0.9522	0.9499
	-0.5	-0.5	0.5	0.5	0.9659	0.9577	0.953	0.9492	0.9518
	-0.5	0.5	-0.5	0.5	0.9648	0.9599	0.9572	0.9486	0.9506
	-0.5	0.5	0.5	-0.5	0.9681	0.9594	0.9564	0.9509	0.9485
	0.5	-0.5	-0.5	0.5	0.9681	0.9594	0.9564	0.9509	0.9485
	0.5	-0.5	0.5	-0.5	0.9648	0.9599	0.9572	0.9486	0.9506
	0.5	0.5	-0.5	-0.5	0.9659	0.9577	0.953	0.9492	0.9518
3, 15, 15, 15	-1	0.33	0.33	0.33	0.9612	0.9535	0.9554	0.9522	0.9505
unbalanced	0.33	-1	0.33	0.33	0.9619	0.9569	0.9511	0.9544	0.95
n=n,2(n,n,n)	0.33	0.33	-1	0.33	0.961	0.9549	0.9549	0.9497	0.9491
	0.33	0.33	0.33	-1	0.9581	0.9565	0.9513	0.9518	0.9542
	-0.5	-0.5	0.5	0.5	0.9603	0.9539	0.9545	0.9516	0.9515
	-0.5	0.5	-0.5	0.5	0.9634	0.9601	0.9548	0.9537	0.9497
	-0.5	0.5	0.5	-0.5	0.9626	0.9543	0.9539	0.9503	0.9513
	0.5	-0.5	-0.5	0.5	0.9626	0.9543	0.9539	0.9503	0.9513
	0.5	-0.5	0.5	-0.5	0.9634	0.9601	0.9548	0.9537	0.9497
	0.5	0.5	-0.5	-0.5	0.9603	0.9539	0.9545	0.9516	0.9515

Table 19: Coverage probability of confidence intervals for single linear combinations of model estimates

17.3.2 Various contrast settings

Contrast type	n	Interval coverage with correlation	unadjusted	Bonferroni adjusted
Dunnett	5	0.9807	0.9267	0.9878
	10	0.9653	0.9042	0.9739
	20	0.9547	0.8888	0.9635
	50	0.9534	0.8861	0.9607
	100	0.9478	0.8806	0.9573
Tukey	5	0.9805	0.7886	0.9882
	10	0.9670	0.7406	0.9773
	20	0.9600	0.7399	0.9701
	50	0.9539	0.7198	0.9646
Average	5	0.9794	0.8840	0.9842
	10	0.9661	0.8567	0.9705
	20	0.9566	0.8454	0.9616
	50	0.9501	0.8368	0.9545
	100	0.9518	0.8386	0.9561
Williams	5	0.9682	0.9454	0.9883
	10	0.9589	0.9323	0.9811
	20	0.9556	0.9323	0.9787
	50	0.9508	0.9252	0.9733
	100	0.9517	0.9262	0.9730
McDermott	5	0.9809	0.9116	0.9903
	10	0.9672	0.8821	0.9776
	20	0.9587	0.8667	0.9713
	50	0.9520	0.8590	0.9632
	100	0.9486	0.8529	0.9633

Table 20: Coverage probability of simultaneous confidence intervals for different contrasts (λ_1=3, rho_i=1,1.5,2)

18 Example Data: Eckert

		2002			2003		
Type	Block	Aphids	Thrips	Bugs	Aphids	Thrips	Bugs
BT	1	700	41	8	1404	7	9
BT	2	3270	55	30	1456	16	3
BT	3	5814	50	25	2674	9	7
BT	4	3054	62	40	2141	12	3
BT	5	4024	103	13	1132	11	1
BT	6	619	26	27	370	9	1
BT	7	3679	23	10	858	10	4
BT	8	1730	24	25	735	10	4
INS	1	2339	60	8	4284	8	10
INS	2	2524	47	1	2619	2	3
INS	3	4441	66	11	3053	7	5
INS	4	2876	89	1	5809	5	1
INS	5	5692	41	8	3179	0	3
INS	6	829	15	4	444	2	1
INS	7	3432	16	0	1317	2	1
INS	8	1012	34	3	801	9	4
ISO	1	2098	65	18	1778	4	0
ISO	2	2631	62	31	2045	7	5
ISO	3	1948	60	10	3203	12	4
ISO	4	1728	74	11	2855	9	2
ISO	5	1587	89	24	1383	12	4
ISO	6	1097	17	9	651	9	4
ISO	7	884	32	8	1303	6	2
ISO	8	1417	23	7	3119	22	1

Table 21: Example dataset from Eckert (Sum of counts over the time covariate)

19 R Code

19.1 GLM fit

A generalized linear model can be fitted in R [22] by the function glm().

```
glm(formula, family = gaussian, data, weights, subset,
    na.action, start = NULL, etastart, mustart,
    offset, control = glm.control(...), model = TRUE,
    method = "glm.fit", x = FALSE, y = TRUE, contrasts = NULL, ...)
```

A generalized linear poisson model with one factor including three parameters can be fitted by

```
# data generation (3 groups A,B,C with mean 5,10,20)
> n=100
> mu <- rep(c(5, 10, 20), each=n)
> dat <- data.frame(count=rpois(length(mu),mu),
+ trt=rep(c("A", "B", "C"), each=n))

> # Poisson glm
> poi.glm <- glm(count ~ trt, data=dat, family=poisson(link="log"))
> # Cell-means model
> poi.glm.cm <- glm(count ~ trt -1, data=dat, family=poisson(link="log"))
> # Maximum likelihood estimates
> mle <- exp(poi.glm.cm$coefficients)
> mle
 trtA  trtB  trtC
 5.31  9.85 20.11
> # Variances
> diag(vcov(poi.glm.cm))
       trtA        trtB        trtC
0.001883221 0.001015228 0.000497265
```

19.2 Extracting marginal means

To extract the marginal means and the matching variance-covariance matrix following function is used:

```
lsmeans <- function(fit){

  fmd <- fit$model

  fn <- length(names(fmd))
  dl <- length(fmd[,1])
  ord.fmd <- fmd[do.call(order,fmd),]

  sf <- character(length=dl)
```

```
  for (i in 2:fn){sf <- paste(sf, ord.fmd[,i], sep="")}
  ind <- data.frame(sf, z=1:dl)
  ind <- tapply(ind$z, ind$sf, min)
  dat.neu <- ord.fmd[ind,]

  mm <- model.matrix(as.formula(fit$call[2]), dat.neu)
  rownames(mm) <- names(ind)

  para <- fit$coefficients
  est <- mm %*% para
  vcmat <- vcov(fit)
  stdr <- sqrt(diag(mm %*% vcmat %*% t(mm)))
  est.sd <- data.frame(est, stdr)
  vcovmat <- mm %*% vcmat %*% t(mm)
  return(list(estimate = est.sd, vcovm = vcovmat))

}
```

19.3 GLMM fit

A generalized linear mixed model can be fitted in R [22] by the function lmer() in the package lme4 (Bates, D. and Sarkar, D. (2006)).

Generating data for a block design with 10 random distributed blocks and two groups with means λ=5,10 and sample size 20 for each block-treatment combination.

```
> n <- 20
> lambdas <- c(5, 10)
> nb <- 10
> N <- n*nb*length(lambdas)
> A <- rep(as.character(1:length(lambdas)), each=n*nb)
> Bl <- rep(rep(as.character(1:nb), times=n), 2)
> zz <- rnorm(nb, 0, 0.5*lambdas[1])
> lamb2bl <- (lambdas[1] + zz)/lambdas[1]
> rbl <- rep(lamb2bl, times=n*nb)
```

```
> rlam <- rep(rep(lambdas, each=n), each=nb)
> mus <- rlam * rbl
> resp <- rpois(N, mus)
> dat <- data.frame(resp, A, Bl)
```

Fit of a generalized linear mixed Poisson model with lmer() and extracting estimates and the variance covariance matrix:

```
> fit <- lmer(resp ~ A -1 + (1|Bl), data=dat, family=poisson())
> est <- fixef(fit)
> est
      A1       A2
1.511424 2.188384
> vcmat <- as.matrix(vcov(fit))
> vcmat
           A1         A2
A1 0.07779600 0.07690394
A2 0.07690394 0.07735725
```

These parameters can be used to compute intervals, based on normal theory.

19.4 Diagnostics

First some count data are generated for seven groups with different means.

```
> muT <- c(2,5,10,20,30,40,50)
> n <- c(100,100,100,100,100,100,100)
> muT1 <- rep(muT, n)
> dat <- rpois(lambda=muT1, n=sum(n))
> trt <- as.factor(rep(1:length(n),n))
> data <- data.frame(count=dat,trt=trt)
```

Then two generalized linear models with Gaussian- and with Poisson-family are fitted

```
> gmod <- glm(count ~ trt, data=data, family=gaussian(link="identity"))
> pmod <- glm(count ~ trt, data=data, family=poisson(link="log"))
```

Now the residuals can be plotted against the predicted values.

```
> plot(residuals(gmod) ~ predict(gmod, type="response"),
+ xlab=expression(hat(mu)),
+ ylab="Deviance residuals", main="Gaussian linear model")
> abline(h=0, lty=2)

> plot(residuals(pmod) ~ predict(pmod, type="response"),
+ xlab=expression(hat(mu)),
+ ylab="Deviance residuals", main="Poisson glm")
> abline(h=0, lty=2)

> plot(residuals(pmod, type="response") ~ predict(pmod, type="response"),
+ xlab=expression(hat(mu)), ylab="Response residuals", main="Poisson glm")
> abline(h=0, lty=2)
```

For detecting outliers, half-normal plots can be created for the jacknife residuals

```
> library(faraway)
> halfnorm(rstudent(pmod))
```

19.5 Confidence intervals

19.5.1 Profile likelihood interval

Fit and profile of a generalized linear model:

```
> library(MASS)
> fit.glm <- glm(y ~ factor, data=data, family=poisson())
> pr.glm <- profile(fit.glm)
```

Extraction of the profile t statistic (τ) and the relevant range of θ_2 for the parameter β_2 (difference of the parameter "'trt2"' and the intercept ("'trt1"')) from an glm-profile

```
> nam.coef <- names(d.glm$coefficients)
> trt.nr <- 2
> trt <- nam.coef[trt.nr]
> z <- pr.glm[[trt]][,1]
> beh <- exp(pr.glm[[trt]][,2][,trt.nr])
> tdf <- sum(n)-length(nam.coef)
> alpha <- 0.025
```

Plot of the profile t values for the matching θ_2s

```
> plot(z ~ beh, type="n", xlab=expression(theta),
+ ylab=expression(tau))
# profile function
> lines(spline(beh, z))
# estimate
> int.pro.glm <- data.frame(tau=z, beh)
> est <-  as.numeric(int.pro.glm[z==0,])
> lines(x=c(est[2], est[2]), y=c(est[1], min(z)), col="orange", lty=2)
> lines(x=c(min(beh), est[2]), y=c(est[1], est[1]), col="orange", lty=2)
# CIs
> sf <- splinefun(z, beh)
> cil <- sf(qt(alpha, tdf))
> lines(x=c(cil, cil), y=c(qt(alpha,tdf), min(z)), col="blue", lty=2)
> lines(x=c(min(beh), max(beh)), y=c(qt(alpha,tdf),
+ qt(alpha,tdf)), col="blue", lty=2)
> ciu <- sf(qt(1-alpha, tdf))
> lines(x=cdciu, ciu), y=c(qt(1-alpha,tdf), min(z)), col="blue", lty=2)
> lines(x=c(min(beh), max(beh)), y=c(qt(1-alpha,tdf),
+ qt(1-alpha,tdf)), col="blue", lty=2)
> MASS.cis <- exp(confint(pr.glm))
> CI <- data.frame(c(cil, MASS.cis[trt.nr,1]),
+ c(ciu, MASS.cis[trt.nr,2]),
+ row.names=c("aus grafik", "confint()"))
> names(CI) <- c("2.5 %", "97.5 %")
```

19.5.2 Wald interval

```
# data generation + model fit
> muT <- c(2,5,10,20,30,40,50)
> n <- c(100,100,100,100,100,100,100)
> muT1 <- rep(muT, n)
> dat <- rpois(lambda=muT1, n=sum(n))
> trt <- as.factor(rep(1:length(n),n))
```

```
> data <- data.frame(count=dat,trt=trt)
> pmod <- glm(count ~ trt, data=data, family=poisson(link="log"))
# CIs
> est <- data.frame(para=pmod$coefficients,stdr=sqrt(diag(vcov(pmod))))
> df.res <- pmod$df.res
> CI.l <- est[,1]-qt(0.975, df=df.res)*est[,2]
> CI.u <- est[,1]+qt(0.975, df=df.res)*est[,2]
> CI <- data.frame(para=est[,1], lower.limit=CI.l, upper.limit=CI.u)
> exp(CI)
```

19.5.3 Nonparametric bootstrap interval

The bootstrab confidnce interval for the comparison of two groups on page 22 was computed in R [22] with the function boot.ci() of the package boot:

```
library(boot)
> boot.glm2<-function(Y, GROUP, conf.level=0.95, R=100)
+ {
+   ylist <- split(Y, GROUP)
+   data<-data.frame(Y=Y, GROUP=GROUP)
+   k <- length(ylist)
+   ni <- lapply(ylist, length)
+   indlist<-split(1:length(Y), GROUP)
+   bootglm <- function(d, i)
+                 {
+                         datnew<-d[i,1]
+
+                         c.Interest <- glm(datnew~GROUP,
+                         family=poisson())$coefficients[2]
+                         return(c.Interest)
+                 }
+   boot.out<-boot(sim="ordinary", stype="i",
+   R=R, statistic=bootglm, data=data, strata=data[,2])
+   conf.int <- boot.ci(boot.out=boot.out, conf=conf.level,
```

```
+   type =c("perc"))$perc[4:5]
+   estimate=glm(Y~GROUP, data=data, family=poisson())$coefficients[2]
+   return(list(conf.int=conf.int,estimate=estimate
+   ))
+ }

> bstr.ci <- exp(boot.glm2(Y=count, GROUP=as.factor(grp),
+ conf.level=0.95, R=100)$conf.int)
```

19.6 Simultaneous confidence intervals

Dunnett confidence intervals can be computed in R for MLEs of a generalized linear model with the package multcomp (Bretz, F., Hothorn, T., Westfall, P. (2004)) by

```
# data generation
> muT <- c(2,5,10,20,30,40,50)
> n <- c(100,100,100,100,100,100,100)
> muT1 <- rep(muT, n)
> dat <- rpois(lambda=muT1, n=sum(n))
> trt <- as.factor(rep(1:length(n),n))
> data <- data.frame(count=dat,trt=trt)

# Cell means Poisson-model
> pmod <- glm(count ~ trt -1, data=data, family=poisson(link="log"))

# Extracting coefficients of the glm-object
> estpar <- pmod$coefficients
> p <- length(estpar)
> covm <- vcov(pmod)
> df <- pmod$df.res

# Building the Dunnett contrast matrix
> library(multcomp)
> names(n) <- names(estpar)
```

```
> cm <- contrMat(n, type="Dunnett")
> cm
          trt1 trt2 trt3 trt4 trt5 trt6 trt7
trt2-trt1   -1    1    0    0    0    0    0
trt3-trt1   -1    0    1    0    0    0    0
trt4-trt1   -1    0    0    1    0    0    0
trt5-trt1   -1    0    0    0    1    0    0
trt6-trt1   -1    0    0    0    0    1    0
trt7-trt1   -1    0    0    0    0    0    1

# Getting estimates and standard errors for the difference to the control
# from the parameter verctor, the variance-covariance matrix and
# the contrast matrix
> covma <- cm %*% covm %*% t(cm)
> d <- 1/sqrt(diag(covma))
> dd <- diag(d)
> cr <- dd %*% covma %*% dd
> ests <- cm %*% estpar
> ses <- sqrt(diag(covma))

# Computing the quantile of the multivariate standard normal distribution
> conf.level <- 0.95
> dim <- ncol(cr)
> eps = 0.001
> pfct <- function(q, conf = FALSE) {
+   low <- rep(-abs(q), dim)
+   upp <- rep(abs(q), dim)
+   pmvt(lower = low, upper = upp, df = df, corr = cr,
+   abseps = eps/10, maxpts = 1e+06) - conf.level
+ }
> calpha <- uniroot(pfct, lower = 0, upper = 5, tol = eps, conf = TRUE)$root

# Computing confidence intervals for the ratio
> LowerCL <- ests - calpha * ses
```

```
> UpperCL <- ests + calpha * ses

> cint <- cbind(ests, LowerCL, UpperCL)
> colnames(cint) <- c("estimate", "lower", "upper")
> exp(cint)

           estimate     lower     upper
trt2-trt1  2.456311  2.026082  2.977897
trt3-trt1  4.587379  3.834961  5.487420
trt4-trt1  9.475728  7.988913 11.239254
trt5-trt1 14.378641 12.156522 17.006946
trt6-trt1 19.441748 16.460641 22.962748
trt7-trt1 24.004854 20.339835 28.330271
```

This code is provided by the function csimint of the library(multcomp):

```
> cis <- csimint(estpar, df, covm, cmatrix=cm)
> exp(cis$conf.int)

              lower     upper
trt2-trt1  2.026162  2.977778
trt3-trt1  3.835103  5.487217
trt4-trt1  7.989195 11.238858
trt5-trt1 12.156943 17.006357
trt6-trt1 16.461207 22.961959
trt7-trt1 20.340530 28.329303
attr(,"conf.level")
[1] 0.95
```

19.7 Quasipoisson

A quasi-poisson model can be fit in R by

```
> fit.glm <- glm(count ~ factor, data=data,
+ family=quasipoisson(link="log"))
```

For using different variance functions, like a quadratic dependency on the mean a quasi family can be used:

```
> quasi.glm <- glm(count ~ factor, data=data, family=quasi(var="mu^2", link="log"))
```

19.8 Negative binomial GLM

A negative binomial generalized linear model can be fit with the function glm.nb() in the package MASS.

```
> negbin.glm <- glm.nb(count ~ factor, data=data)
```

or with:

```
> glm(count ~ factor, family = negative.binomial(theta), data = data)
```

References

[1] Bates, D.M. and Watts, D.G. (1988): Nonlinear regression analysis and its applications, John Wiley, New York.

[2] Breslow, N. (1990): Tests of Hypotheses in overdispersed Poisson regression and other quasi-likelihood models. JASA 85(410): 565-571.

[3] Bretz, F., Genz, A., Hothorn, L.A. (2001): On the numerical availability of multiple comparison procedures. Biometrical Journal 43(5): 645-656.

[4] Brown, L.D., Cai, T.T., DasGupta, A. (2003): Interval estimation in exponential families. Statistica Sinica 13: 19-49.

[5] Candy, S.G. (2000): The application of generalized linear mixed models to multi-level sampling for insect population. Environmental and Ecological Statistics 7: 217-238.

[6] Chen, J.S. and Jennrich, R.I. (1996): The signed root deviance profile and confidence intervals in maximum likelihood analysis. JASA 91(435): 993-998.

[7] Chen, Y.H. and Zhou, X.H. (2004): Interval estimates for the ratio and difference of two lognormal means. Statistics in medicine (in press)

[8] Dobson, A.J. (1990): An introduction to generalized linear models. Chapman & Hall, London.

[9] Everitt, B. and Hothorn, T.(2006): A Handbook of Statistical Analyses Using R. Chapman & Hall/CRC, Boca Raton, FL.

[10] Faraway, J.J. (2006): Extending the linear model with R. Generalized linear, mixed effects and nonparametric regression models. Chapman & Hall / CRC.

[11] Giminez, Choquet, Lamor, Scofield, Fletcher, Lebreton and Pradel (2005): Profile-likelihood based confidence intervals for capture-recapture models. JABES 10(2): 184-196.

[12] Graham, P.L., Mengersen, K., Morton, A.P. (2003): Confidence limits for the ratio of two rates based on likelihood scores: non-iterative method. Statistics in Medicine 22: 2071-2083.

[13] Jaech J.L. (1970): Comparing two methods of obtaining a confdence interval for the ratio of Poisson parameters. Technometrics 12(2): 383–387.

[14] Lambert, D (1992): Zero-inflated Poisson regression, with an application to defects in manufacturing. Technometrics 34(1): 1-14.

[15] Lindsey, J.K. (1997): Modelling frequency and count data. Oxford Statistical Science Series 15, 2ed., Oxford Science Publications.

[16] Liu, G.F., Wang, J., Liu, K., Snavely, D.B. (2006): Confidence intervals for an exposure adjusted incidence rate difference with applications to clinical trials. Statistics in Medicine 25(8): 1275-1286.

[17] Lui, K.J. (2004): Eight interval estimators of a common rate ratio under stratifed Poisson sampling. Statistics in Medicine 23: 1283-1296.

[18] McCullagh, P. and Nelder, J.A. (1989): Generalized linear models, 2nd edn. Chapman & Hall, London.

[19] McCulloch, C.E. and Searle, S.R. (2001): Generalized, linear and mixed models. John Wiley & Sons, Inc.

[20] Ng, H.K.T. and Tang, M.L. (2005): Testing the equality of two Poisson means using the rate ratio. Statistics in Medicine 24: 955-965.

[21] Price, R.M., Bonnett, G.B. (2000): Estimating the ratio of two Poisson rates. Computational Statistics & Data Analysis 34: 345-356.

[22] R Development Core Team (2005). R: A language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria., URL http://www.R-project.org.

[23] Schwertman, N.C. and Martinez, R. (1994): Approximate confidence intervals for the difference in two poisson parameters. Journal of Statistics in Computation and Simulation 50: 235-247.

[24] Venzon, D.J. and Moolgavkar, S.H. (1988): A method for computing profile-likelihood-based confidence intervals. Applied Statistics 37(1), pp. 87-94.

[25] Warton, D.I. (2005): Many zeros does not mean zero inflation: comparing the goodness-of-fit of parametric models to multivariate abundance data. Environmetrics 16: 275-289.

[26] Wellek, S. (2003): Testing statistical hypotheses of equivalence, Chapman & Hall/CRC, Boca Raton, FL.

[27] Wu, J., Wong, A.C.M., Wei, W. (2006): Interval estimation of the mean response in a log-regression model. Statistics in Medicine 25: 2125-2135.

[28] Yandell, B.S. (1997): Practical data analysis for designed experiments. Chapman & Hall, London.

[29] Yau, K.W., Wang, K., Lee, A.H. (2003): Zero-inflated negative binomial mixed regression modeling of over-dispersed count data with extra zeros. Biometrical Journal 45: 437-452.

Hiermit erkläre ich an Eides Statt, daß ich diese Arbeit nur mit den aufgeführten Hilfsmitteln erstellt habe. Die Arbeit wurde bisher keiner anderen Prüfungskommission vorgelegt.

(Hannover, den 12.Juni.06, Daniel Gerhard)