

Choosing the Link Function and Accounting for Link Uncertainty in Generalized Linear Models using Bayes Factors

Claudia Czado¹, Adrian E. Raftery²

¹ Zentrum Mathematik, Technische Universität München, Boltzmannstr.3 , D-85747 Garching bei München, Germany ,e-mail: cczado@ma.tum.de

² Department of Statistics, University of Washington, Box 354322, Seattle, WA 98195-4322, USA, e-mail: raftery@stat.washington.edu

Received: date / Revised version: date

Abstract One important component of model selection using generalized linear models (GLM) is the choice of a link function. We propose using approximate Bayes factors to assess the improvement in fit over a GLM with canonical link when a parametric link family is used. The approximate Bayes factors are calculated using the Laplace approximations given in [32], together with a reference set of prior distributions. This methodology can be used to differentiate between different parametric link families, as well as allowing one to jointly select the link family and the independent variables. This involves comparing nonnested models and so standard significance tests cannot be used. The approach also accounts explicitly for uncertainty about the link function. The methods are illustrated using parametric link families studied in [12] for two data sets involving binomial responses.

* The first author was supported by Sonderforschungsbereich 386 *Statistische Analyse Diskreter Strukturen*, and the second author by NIH Grant 1R01CA094212-01 and ONR Grant N00014-01-10745.

Key words Bayes factors; link function; GLM, model selection; reference prior.

1 Introduction

To find an appropriate generalized linear model (GLM) for regression data involves choosing the independent variables, the link function and the variance function ([24]). Typically, many different models are compared using individual significance tests based on the asymptotic distribution of the deviance. As pointed out by [17] and [32], this strategy cannot be used for comparing nonnested models. In addition, model uncertainty is usually ignored, as are power considerations. A Bayesian approach can avoid these difficulties and to implement this, [32] developed approximate Bayes factors for GLM's based on the Laplace method for integrals. These approximations require only the maximum likelihood estimate (MLE), the deviance and the observed or expected Fisher information. [22] and [20] review Bayes factors and discuss different ways to calculate Bayes factors.

In this paper, we extend the approach taken by [32] to calculate approximate Bayes factors for GLM's with a parametric link function. Even though GLM's with canonical links (for definition see [24]), such as the logit link in binomial regression, guarantee maximum information and a simple interpretation of the regression parameters, they do not always provide the best fit available to a given data set. Link misspecification can lead to substantial bias in the regression parameters and the mean response estimates (see [14] for binomial responses). One common approach to guard against link misspecification in generalized linear models is to embed the canonical

link in a wide parametric class of links $\mathfrak{S} = \{F(\cdot, \psi), \psi \in \Psi\}$, which includes the canonical link as a special case when $\psi = \psi_0$. Many such parametric link classes for binary regression data have been proposed in the literature. [26], [8], [2], [19], [27] and [38] proposed one-parameter families, while [31], [30], [34] and [10] considered two-parameter families. Link functions for the non-binary case were studied in [30],[10] and [12].

With the multitude of link families to choose from, the Bayes factor approach is able to compare different link families, regardless of whether they are nested or nonnested. We will illustrate this ability by using the two-parameter link family suggested by [12] in several data sets. In addition, we are able to choose the link family and the set of independent variables jointly.

In Section 2 we define and discuss GLM's with parametric links, while in Section 3 the calculation of approximate Bayes factors including the choice of priors will be discussed. Applications will be given in Section 4 and Section 5 will provide a summary and discussion of the method presented.

2 Generalized Linear Models with Parametric Links

The following model for regression data with response Y_i and independent variables $\mathbf{X}_i = (x_{i1}, \dots, x_{ip})$ for $i = 1, \dots, n$ will be used:

1. **Random Component:** $\{Y_i, 1 \leq i \leq n\}$ are independent and have a density of the form

$$f_{y_i}(y_i, \theta_i, \phi) = \exp\left[\frac{y_i \theta_i - b(\theta_i)}{a(\phi)} + c(y_i, \phi)\right], \quad (2.1)$$

for some specified functions $a(\cdot)$, $b(\cdot)$ and $c(\cdot)$. The scale parameter ϕ is allowed to be known or unknown.

Table 1 Link Families for GLM's

Error Distribution	Parameter Restriction	Canonical Link	Link Family $\mathfrak{S} = \{F(\cdot, \psi) : \psi \in \Psi\}$
Normal	μ real	$F(\eta) = \eta$	$F(\eta, \psi) = h(\eta, \psi)$
Binomial	$\mu \in (0, 1)$	$F(\eta) = \frac{\exp(\eta)}{1 + \exp(\eta)}$	$F(\eta, \psi) = \frac{\exp(h(\eta, \psi))}{1 + \exp(h(\eta, \psi))}$
Poisson	$\mu > 0$	$F(\eta) = \exp(\eta)$	$F(\eta, \psi) = \exp(h(\eta, \psi))$
Gamma	$\mu > 0$	$F(\eta) = \eta^{-1}$	$F(\eta, \psi) = [\exp(h(\eta, \psi))]^{-1}$
Inv. Gaussian	$\mu > 0$	$F(\eta) = \eta^{-.5}$	$F(\eta, \psi) = [\exp(h(\eta, \psi))]^{-.5}$

2. **Systematic Component:** The linear predictors $\eta_i(\boldsymbol{\beta}) = \beta_0 + \beta_1 x_{i1} + \dots + \beta_p x_{ip}$ for $1 \leq i \leq n$ influence the response Y_i . Here $\boldsymbol{\beta} = (\beta_0, \dots, \beta_p)$ are unknown regression parameters.
3. **Parametric Link Component:** The linear predictors $\eta_i(\boldsymbol{\beta})$ are related to the mean μ_i of Y_i by $\mu_i = F(\eta_i(\boldsymbol{\beta}), \psi)$ for some $F(\cdot, \psi)$ in $\mathfrak{S} = \{F(\cdot, \psi) : \psi \in \Psi\}$.

We will restrict attention to link families \mathfrak{S} that contain only strictly monotone continuous functions $F(\cdot, \psi)$. Note that in conventional GLM notation the link g is equal to the inverse of F . An unknown scale parameter ϕ in (2.1) is typically estimated by an appropriate moment estimator involving the Pearson χ^2 Statistic ([24]). For a fixed link parameter ψ we remain in the class of GLM's, while this is no longer true if the link parameter ψ and the regression parameter $\boldsymbol{\beta}$ are jointly estimated by the data. It is shown in [13] that the joint MLE $\hat{\boldsymbol{\delta}} = (\hat{\boldsymbol{\beta}}, \hat{\psi})$ of $\boldsymbol{\delta} = (\boldsymbol{\beta}, \psi)$ is strongly consistent and efficient under regularity conditions.

We will illustrate our approach by using the link families suggested by [12]. They allow separate modifications of the left and/or right tail of the link function and exhibit low variance inflation ([35], [36]) for the regression parameters when the link is estimated from the data. This is due to the fact

that the parametrization is locally orthogonal (see [9]). In addition, they are location and scale invariant (see [12]). For GLM's with parametric links they are defined as in Table 1. In Table 1, $h(\eta, \boldsymbol{\psi})$ is one of the following functions:

$$\text{Both tails: } h_b(\eta, \boldsymbol{\psi} = (\psi_1, \psi_2)) = \begin{cases} +\frac{(\eta+1)^{\psi_1-1}}{\psi_1} & \text{if } \eta \geq 0 \\ -\frac{(-\eta+1)^{\psi_2-1}}{\psi_2} & \text{otherwise} \end{cases} \quad (2.2)$$

$$\text{Right tail: } h_r(\eta, \psi_1) = \begin{cases} +\frac{(\eta+1)^{\psi_1-1}}{\psi_1} & \text{if } \eta \geq 0 \\ \eta & \text{otherwise} \end{cases} \quad (2.3)$$

$$\text{Left tail: } h_l(\eta, \psi_2) = \begin{cases} \eta & \text{if } \eta \geq 0 \\ -\frac{(-\eta+1)^{\psi_2-1}}{\psi_2} & \text{otherwise} \end{cases} \cdot \quad (2.4)$$

The parameter restriction for the mean response makes a right tail modification for the Poisson and a left tail modification for the Gamma and inverse Gaussian cases the only sensible modifications to be considered. In all other cases all modifications of the link function are allowed. In particular, (2.4) is a special case of (2.2) with $\psi_1 = 1$, and (2.3) is a special case of (2.2) with $\psi_2 = 1$. As ψ_1 increases, the right tail of $G(\cdot, \boldsymbol{\psi})$ becomes lighter, while an increasing ψ_2 makes the left tail of $G(\cdot, \boldsymbol{\psi})$ lighter. The specification (2.3) is asymmetric if $\psi_1 \neq 1$, while the specification (2.4) is asymmetric if ($\psi_2 \neq 1$). The both tails specification (2.2) is asymmetric if $\psi_1 \neq \psi_2$. Further, for $\psi_1 < 0$ and $\psi_2 < 0$

$$\lim_{\eta \rightarrow \infty} h_r(\eta, \psi_1) = \frac{1}{|\psi_1|} \quad \text{and} \quad \lim_{\eta \rightarrow -\infty} h_l(\eta, \psi_2) = -\frac{1}{|\psi_2|},$$

$$\lim_{\eta \rightarrow \infty} h_b(\eta, \boldsymbol{\psi} = (\psi_1, \psi_2)) = \frac{1}{|\psi_1|} \quad \text{and} \quad \lim_{\eta \rightarrow -\infty} h_b(\eta, \boldsymbol{\psi} = (\psi_1, \psi_2)) = -\frac{1}{|\psi_2|}.$$

This, together with monotonicity of h_r , h_l and h_b , imply restrictions on the range of allowable means $\mu_i = E(Y_i)$ if $\psi_1 < 0$ or $\psi_2 < 0$. In particular, for binomial links, we have restrictions on the allowable success probabilities

p_i given by

$$p_i \leq \frac{\exp\{\frac{1}{|\psi_1|}\}}{1 + \exp\{\frac{1}{|\psi_1|}\}} \text{ if } \psi_1 < 0 \quad (2.5)$$

$$p_i \leq \frac{\exp\{-\frac{1}{|\psi_2|}\}}{1 + \exp\{-\frac{1}{|\psi_2|}\}} \text{ if } \psi_2 < 0. \quad (2.6)$$

For example if we want an allowable range of success probabilities between .1 and .9, this implies that $\psi_1 \geq -.46$ and $\psi_2 \geq -.46$ if negative link values are allowed, by inversion of (2.5) and (2.6). There are no restrictions on the success probabilities when $\psi_1 \geq 0$ and $\psi_2 \geq 0$.

3 Approximate Bayes Factors for GLM's with Parametric Link

We are interested in assessing the evidence for a GLM with a noncanonical link against the same GLM with a canonical link using Bayes factors. For this, we denote by M_ψ a GLM with a fixed link parameter ψ for a given set of independent variables, while M_c denotes the same GLM using the canonical link. We denote the regression parameter corresponding to model M_ψ by β_ψ to indicate that the regression parameters are on different scales for different ψ 's. We are interested in the Bayes factor for model M_ψ against model M_c given the data $\mathbf{Y} = (Y_1, \dots, Y_n)$, which is defined as the ratio of posterior to prior odds, namely

$$B_\psi := \frac{pr(\mathbf{Y}|M_\psi)}{pr(\mathbf{Y}|M_c)}, \quad (3.1)$$

the ratio of the integrated likelihoods. In equation (3.1),

$$pr(\mathbf{Y}|M_\psi) = \int pr(\mathbf{Y}|M_\psi, \beta_\psi) p(\beta_\psi|M_\psi) d\beta_\psi, \quad (3.2)$$

where β_ψ is the corresponding regression parameter in Model M_ψ and $p(\beta_\psi|M_\psi)$ is its prior density in model M_ψ . Note that M_c corresponds to M_ψ with $\psi = 1$.

The Bayes factor is a summary of the evidence for M_ψ against M_c provided by the data. Sometimes it is useful to consider $2 \log B_\psi$, which is on the same scale as the familiar deviance and likelihood ratio test statistics. We use the rounded scale given in Table 1 of [32] for interpreting B_ψ or $2 \log B_\psi$.

This approach allows us to compare different parametric link families as follows. Let M_θ denote a GLM using a link family indexed by the link parameter θ and construct B_θ in a similar fashion as B_ψ . The quantity $\frac{B_\psi}{B_\theta}$ then provides a summary of the evidence for model M_ψ against model M_θ given the data and the same set of independent variables. In a similar way we can construct comparisons of models with different sets of independent variables and link parameters.

For the link families given in Table 1 it is also of interest to assess whether a right tail, left tail or a both tail modification is needed. For this we can compare $B_{\psi_1}(B_{\psi_2})$ and $B_{\boldsymbol{\psi}=(\psi_1, \psi_2)}$ for individual link parameter values or construct overall Bayes factors for each tail modification, given by

$$\text{Both Tails : } B_b = \int B_{\boldsymbol{\psi}=(\psi_1, \psi_2)} pr(\boldsymbol{\psi} | M_{\boldsymbol{\psi}=(\psi_1, \psi_2)}) d\boldsymbol{\psi} \quad (3.3)$$

$$\text{Right Tail : } B_r = \int B_{\psi_1} pr(\psi_1 | M_{\psi_1}) d\psi_1 \quad (3.4)$$

$$\text{Left Tail : } B_l = \int B_{\psi_2} pr(\psi_2 | M_{\psi_2}) d\psi_2, \quad (3.5)$$

where $pr(\boldsymbol{\psi} | M_{\boldsymbol{\psi}=(\psi_1, \psi_2)})$, $pr(\psi_1 | M_{\psi_1})$ and $pr(\psi_2 | M_{\psi_2})$ denote the corresponding prior densities for $\boldsymbol{\psi}$, ψ_1 and ψ_2 , respectively. If the link parameter values are not chosen in advance, but instead are estimated, B_{ψ_1} , B_{ψ_2} and $B_{\boldsymbol{\psi}}$ will tend to overstate the evidence for a modification. The overall Bayes factors B_r , B_l and B_b are preferable in this case, because they take

into account the fact that the link parameters are unknown and thus take link uncertainty into account. For example, the ratio $\frac{B_b}{B_r}$ will compare a both tails modification to a right tail one. In a similar fashion we can assess the evidence for one link family against another one given the same or different set of independent variables.

To complete the specification of these overall Bayes factors, we have to select prior distributions for the regression parameters given a model with a specified link parameter, as well as the prior distribution to be used for the link parameter.

For the prior distribution of the regression parameters β_ψ in the model M_ψ we use the reference proper prior distributions suggested by [32] for GLM's, since for fixed values of the link parameter ψ we remain in the class of ordinary GLM's. These prior distributions assume little prior information. They are based on adjusted dependent variables to mimic the behavior for ordinary linear regression models. For a $(p + 1)$ -dimensional β_ψ including an intercept, we use the prior

$$\beta_\psi | M_\psi \sim N_{p+1}(\mathbf{v}_\psi, Q_\psi U Q'_\psi), \quad (3.6)$$

where $N_p(\boldsymbol{\mu}, \Sigma)$ denotes a p -dimensional normal distribution with mean vector $\boldsymbol{\mu}$ and covariance matrix Σ . To specify the quantities in (3.6), the adjusted dependent variable $z_i^\psi = g_\psi(\hat{\mu}_i^\psi) + (y_i - \hat{\mu}_i^\psi)g'_\psi(\hat{\mu}_i^\psi)$ with weights w_i^ψ ([24], p.40) has to be considered. Here $\hat{\mu}_i^\psi$ denotes the MLE of the i th mean response in the GLM with link parameter ψ , and $g_\psi(\cdot)$ is the inverse

of $F(\cdot, \psi)$. Define the weighted summary statistics:

$$\bar{z}_\psi = \frac{\sum_{i=1}^n w_i^\psi z_i^\psi}{\sum_{i=1}^n w_i^\psi} \text{ and } s_0^\psi = \sqrt{\frac{\sum_{i=1}^n w_i^\psi (z_i^\psi - \bar{z}_\psi)^2}{\sum_{i=1}^n w_i^\psi}}, \quad (3.7)$$

$$\bar{x}_j^\psi = \frac{\sum_{i=1}^n w_i^\psi x_{ij}}{\sum_{i=1}^n w_i^\psi} \text{ and } s_j^\psi = \sqrt{\frac{\sum_{i=1}^n w_i^\psi (x_{ij} - \bar{x}_j^\psi)^2}{\sum_{i=1}^n w_i^\psi}}, \quad j = 1 \cdots, p. \quad (3.8)$$

Then the prior mean is specified as $\mathbf{v}'_\psi = (\bar{z}_\psi, 0, \dots, 0)'$, U denotes a diagonal matrix with diagonal entries given by $(1, \sigma_p^2, \dots, \sigma_p^2)$ and

$$Q_\psi = s_0^\psi \begin{bmatrix} 1 - \frac{\bar{x}_1^\psi}{s_1^\psi} - \frac{\bar{x}_2^\psi}{s_2^\psi} \cdots - \frac{\bar{x}_p^\psi}{s_p^\psi} \\ 0 \quad \frac{1}{s_1^\psi} \quad 0 \quad \cdots \quad 0 \\ 0 \quad 0 \quad \frac{1}{s_2^\psi} \quad \cdots \quad 0 \\ \vdots \quad \vdots \quad \vdots \quad \vdots \quad \vdots \\ 0 \quad 0 \quad 0 \quad \cdots \quad \frac{1}{s_p^\psi} \end{bmatrix}.$$

It remains to specify σ_p^2 . The arguments of [32] and subsequent experience using Bayes factors for GLM's (e.g. [37]) suggests using the value $\sigma_p^2 = 1$.

We now consider the choice of the prior distribution for the link parameter ψ . For the link families specified in (2.3) and (2.4), we require that $\psi_1 \geq \psi_l$ and $\psi_2 \geq \psi_l$, where ψ_l is chosen in such a way that the restriction on the range of the allowable mean values is reasonable. As noted before for binomial links, $\psi_l = -.46$ restricts the success probabilities to be between .1 and .9 for $\psi_1 < 0$ and $\psi_2 < 0$. If $\psi_l = 0$, there are no restrictions on the success probabilities. Therefore it makes sense to consider prior distributions that are truncated to $[\psi_l, \infty)$ as prior distributions for ψ_1 and ψ_2 , respectively. As a first choice we consider a truncated normal distribution with mean 1, corresponding to the canonical link, and standard deviation σ_ψ . The left column of Figure 1 shows the corresponding prior densities for $\psi_l = -.5$ and $\psi_l = 0$.

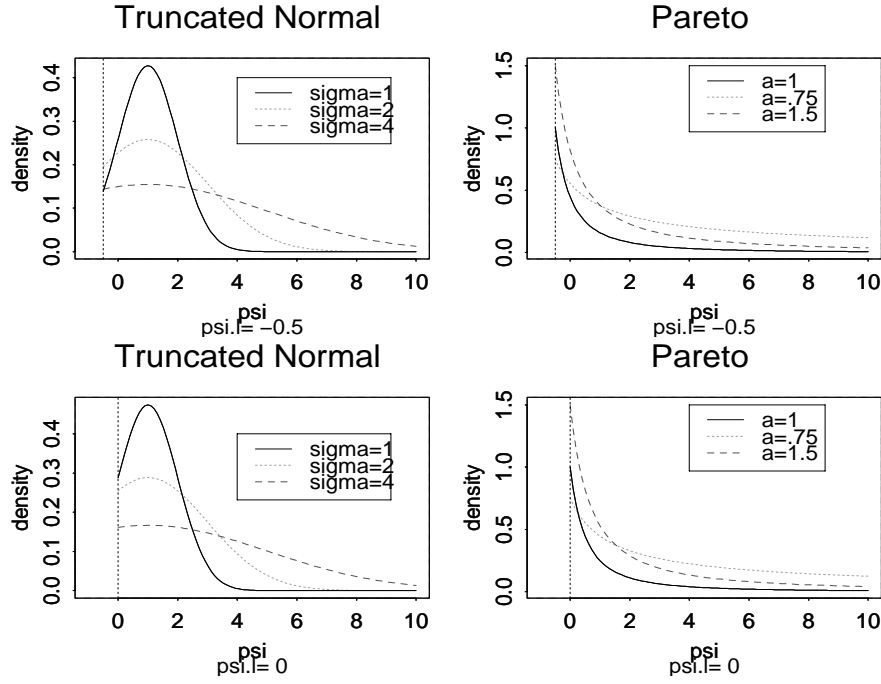


Fig. 1 $N(1, \sigma_\psi)$ priors truncated to $[\psi_l, \infty)$ and Pareto priors defined in (3.9) for the link parameter ψ .

As a second choice we consider Pareto densities given by

$$f(\psi; a, \psi_l) = \frac{a}{(1 + \psi - \psi_l)^{a+1}}, \quad \forall \psi \geq \psi_l, a > 0. \quad (3.9)$$

This prior choice can be motivated by the same arguments as those of [29] on page 11. In particular, the link function $g(\mu; \psi) := F^{-1}(\mu; \psi)$ for (2.3) and (2.4) changes most rapidly for the smallest allowable ψ value, i.e. $\psi = \psi_l$. Therefore [29] chose priors which have highest densities at ψ_l and tails that are monotonically decreasing. The Pareto prior family given in (3.9) satisfies these conditions and is illustrated in the right column of Figure 1.

It seems natural to want the prior distribution of link functions to be centered at, or symmetric about the canonical link function, in some sense, since the canonical link plays a role similar to that of a null hypothesis. As

Table 2 $P(\psi \leq 1)$ for Pareto and truncated normal prior family for ψ .

$P(\psi \leq 1)$ for Pareto prior			$P(\psi \leq 1)$ for $N(1, \sigma_\psi^2)$ trunc. to $[\psi_l, \infty)$		
\mathbf{a}	$\psi_l = -.5$	$\psi_l = 0$	σ_ψ	$\psi_l = -.5$	$\psi_l = 0$
.75	.50	.41	1.00	.46	.41
1.00	.60	.50	2.00	.35	.28
1.50	.75	.65	4.00	.23	.16

a measure of the symmetry of these prior choices about the canonical link parameter we use $P(\psi \leq 1)$, which is tabulated in Table 2 for the different prior choices considered. For the Pareto priors this shows that $a = .75$ can be considered symmetric about the canonical link $\psi = 1$ if $\psi_l = -.5$, while $a = 1$ is symmetric if $\psi_l = 0$. For the truncated normal priors, link values larger than 1 are favored when these parameter values are used.

So far we have considered only single tail modifications. For the both tails case (2.2) with $\boldsymbol{\psi} = (\psi_1, \psi_2)$, we assume independence of the components and use the same priors for ψ_1 and ψ_2 as for the single tail modifications.

To approximate the Bayes factors B_ψ of (3.1) we use the Laplace approximation for Bayes factors for GLM's given in [32], namely

$$2 \log B_\psi \approx \chi_\psi^2 + (E_\psi - E_0), \quad (3.10)$$

where $\chi_\psi^2 = dev(M_c) - dev(M_\psi)$. Here $dev(M)$ denotes the deviance of model M . Let F_ψ denote the observed or expected Fisher information matrix at the MLE $\hat{\boldsymbol{\beta}}_\psi$ in the model M_ψ . Then E_ψ in equation (3.10) is given by

$$E_\psi = \log |G_\psi| - (\hat{\boldsymbol{\beta}}_\psi - \mathbf{v}_\psi)' C_\psi (\hat{\boldsymbol{\beta}}_\psi - \mathbf{v}_\psi) - \log |F_\psi + G_\psi|, \quad (3.11)$$

where $G_\psi = (Q_\psi U Q_\psi')^{-1}$ is the inverse of the prior variance in (3.6) and C_ψ is defined as

$$C_\psi = G_\psi \{I - H_\psi (2I - F_\psi H_\psi) G_\psi\}, \text{ where } H_\psi = (F_\psi + G_\psi)^{-1}.$$

Table 3 Minimal Deviances and Approximate Maximal Bayes Factors for the Generalized Link Functions against the Canonical Link, for the Menarche Data (with prior variance parameter $\sigma_p = 1$)

Model	Minimal Deviance	(ψ_1, ψ_2)	df	Maximal Bayes Factor	(ψ_1, ψ_2)
Right	25.10	(.88,-)	22	2.23	(.75,-)
Left	17.62	(-,1.40)	22	95.05	(-,1.40)
Both	15.38	(1.25,1.67)	21	287.14	(1.27,1.72)

Finally, E_0 is equal to E_ψ , where ψ is taken to be the value corresponding to the canonical link. Equation (3.11) corresponds to equation (9) in [32].

To calculate approximations to the overall Bayes factors specified in (3.3)-(3.5) we use the above approximation and numerically integrate out ψ using the prior specifications for ψ .

4 Applications

4.1 Age of Menarche in Polish girls

In [25] a sample of 3918 Warsaw girls are presented. It is recorded whether or not they had reached menarche together with their age. This is a well known data set and has often been used to demonstrate the need for a link function other than the logistic one. The residual deviance for the logistic regression model with a linear age covariate is 26.70 with 23 degrees of freedom suggesting the possibility of some improving of the fit. Figure 2 gives the deviance profiles and contours when the link families (2.2)-(2.4) are used for binomial regression. The corresponding approximate Bayes factors B_ψ as a function of ψ using $\sigma_p = 1$ are given in Figure 3. The minimal deviances and the maximal approximate Bayes factors are presented in Table 3.

From Table 3 we see that a left tail modification improves the fit using either deviances or Bayes factors. While the deviance indicates that a

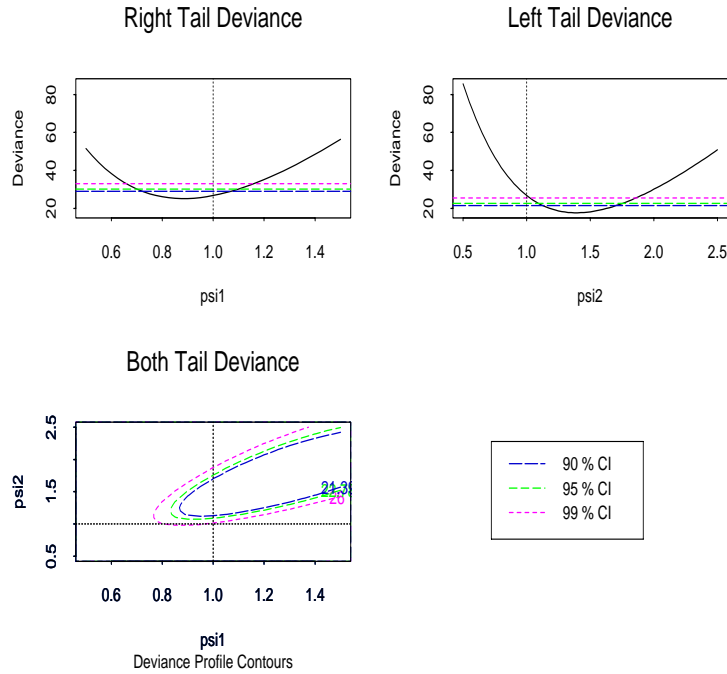


Fig. 2 Minimal Deviance Profiles and Minimal Deviance Contours for the Menarche Data

single tail modification is sufficient, the maximal approximate Bayes factor for the both tail modification is quite high. When looking at the maximal approximate Bayes factor we ignore the error made by estimating the link parameter and therefore it is more appropriate to consider the overall Bayes factor, which accounts for link uncertainty. Since the observed success probabilities vary between 0 and 1 and since a large range of ages was investigated, it seems reasonable not to impose any restriction on the allowable success probabilities. Therefore we assume $\psi_l = 0$. Table 4 gives the overall Bayes factors for the different prior choices for the link parameter.

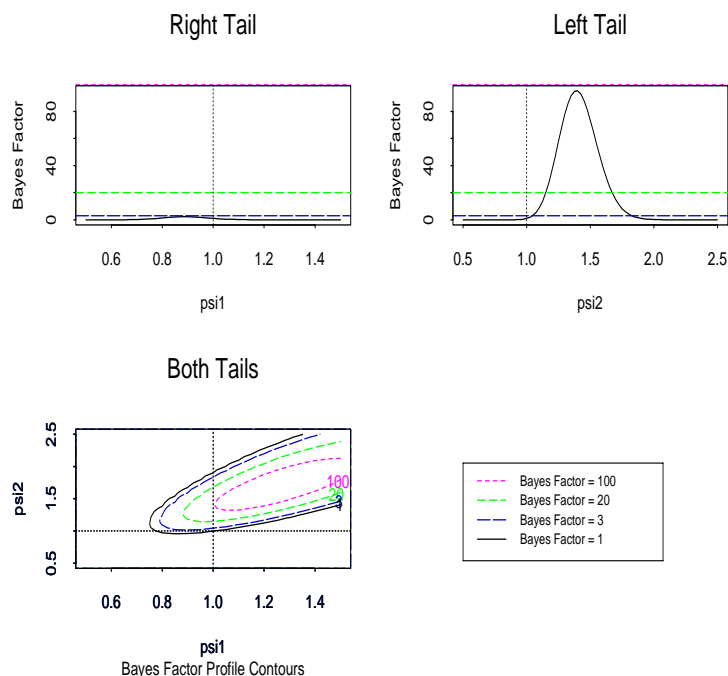


Fig. 3 Approximate Bayes Factor Profiles and Contours for the Menarche Data, with prior variance parameter $\sigma_p = 1$.

Table 4 Approximate Overall Bayes Factors for the Generalized Link Functions Against the Canonical Link, for the Menarche Data with $\psi_l = 0$

a	Pareto prior			$N(1, \sigma_\psi^2)$ truncated to $[\psi_l, \infty)$	σ_ψ	truncated to $[\psi_l, \infty)$		
	Right Tail	Left Tail	Both Tails			Right Tail	Left Tail	Both Tails
.75	.12	5.73	.14	1.00	.23	15.22	9.64	
1.00	.13	5.83	.15	2.00	.14	9.93	4.68	
1.50	.15	5.97	.15	4.00	.08	5.84	1.69	

From Table 4 we see that a left tail modification improves the fit regardless of which prior specification is used for the link parameter and if we account for link uncertainty. The approximate overall Bayes factors also show that a right tail modification gives no improvement for all priors used, while there is slight evidence for a both tail modification when truncated

normal priors are used. Note that these priors give more probability to link values larger than the canonical link value. From Figure 3 we see that B_ψ is largest for link values greater than 1. In contrast, the Pareto priors favor small link values. This explains that the approximate overall Bayes factors are lower for the Pareto priors compared to the truncated normal priors.

Before we can conclude that a link function with a left tail modification is better than the logistic link for these data, we need to consider and exclude possible alternative explanations of what we have observed. It is possible that a transformation of the age variable may be preferable to a link modification, and that the apparently poorer performance of the logistic link is just an artifact due to nonlinearity of the effect of age. We now investigate this possibility.

We consider polynomial models for age of the form

$$\eta_i = \beta_0 + \beta_1 age_i + \dots + \beta_p age_i^p.$$

Calculations of the deviances show that the quadratic model is little better than the linear one, that there is a substantial reduction in deviance when one goes from quadratic to cubic, and little further gain for additional polynomial terms. We therefore restrict ourselves to considering a cubic model for age.

Our question is thus whether a left tail modification with linear age is better than a canonical link with cubic age. The most common approach to this question is to compute deviances and degrees of freedom for the two competing models. The left tail modification with linear age has deviance 17.62 on 22 d.f., while the canonical link with cubic age has deviance 15.04 on 21 d.f. These models are not nested, so a standard likelihood ratio test

cannot be carried out. Nevertheless, the deviance difference is 2.58 with a difference in degrees of freedom of 1, and if this was compared with the standard chi-squared distribution (which cannot validly be done), it would not be significant, and one would typically choose the more parsimonious left tail modification with linear age model.

Bayes factors do allow us to make a formal comparison between these two nonnested models. The Bayes factor for the left tail modification with linear age model against the canonical link model with cubic age is 233 for the Pareto prior with $a = 1$, and 396 for the truncated normal prior with $\sigma = 2$, so that the left tail modification with linear age model is favored. Thus with Bayes factors we reach the same conclusion as with the informal comparison of deviances, with the difference that Bayes factors provide a formal justification for the conclusion.

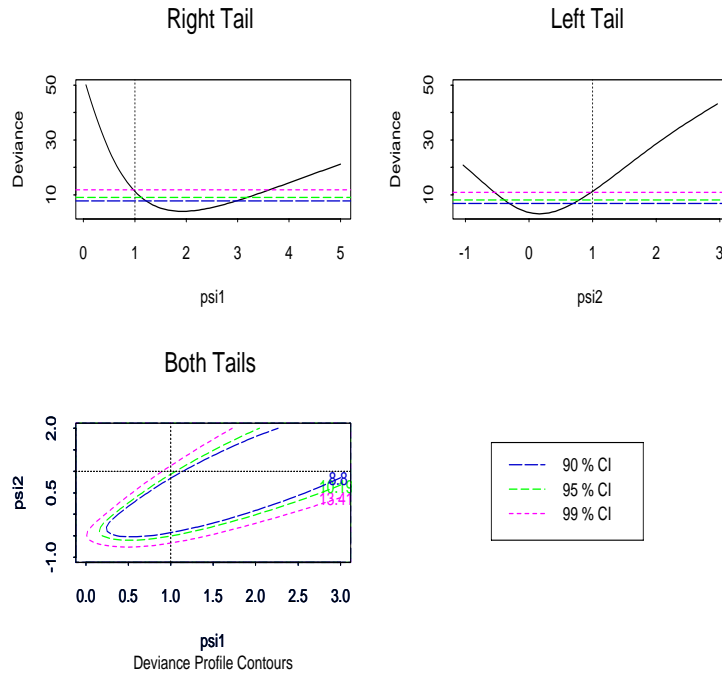
4.2 Beetle Mortality

In [5] the number of insects dead after five hours' exposure to gaseous carbon disulphide at various concentrations are recorded and the data are presented in Table 5. This is also a well known data set for investigating a different link function other than the logistic one. Here, the residual deviance for a logistic model with a centered log dose covariate is 11.23 with 6 degrees of freedom, suggesting some lack of fit.

Figure 4 gives the deviance profiles and contours, when the link families (2.2)-(2.4) are used for binomial regression. They suggest that a link tail modification in this data set is useful and improves the fit. We will now use Bayes factors to decide which specific tail modification is needed. We

Table 5 Beetle Mortality Data

Y_i Number killed	n_i Number of Insects	Dose $\log_{10} CS_2 mg l^{-1}$
6	59	1.6907
13	60	1.7242
18	62	1.7552
28	56	1.7842
52	63	1.8113
53	59	1.8369
61	62	1.8610
60	60	1.8839

**Fig. 4** Deviance Profiles and Deviance Contours for the Beetle Mortality Data

use the prior specification (3.6) with $\sigma_p = 1$ for the regression coefficients. Figure 5 shows the Bayes factors B_ψ as a function of ψ , and in Table 6 we give the minimal deviances and maximal individual Bayes factors B_ψ for each tail modification.

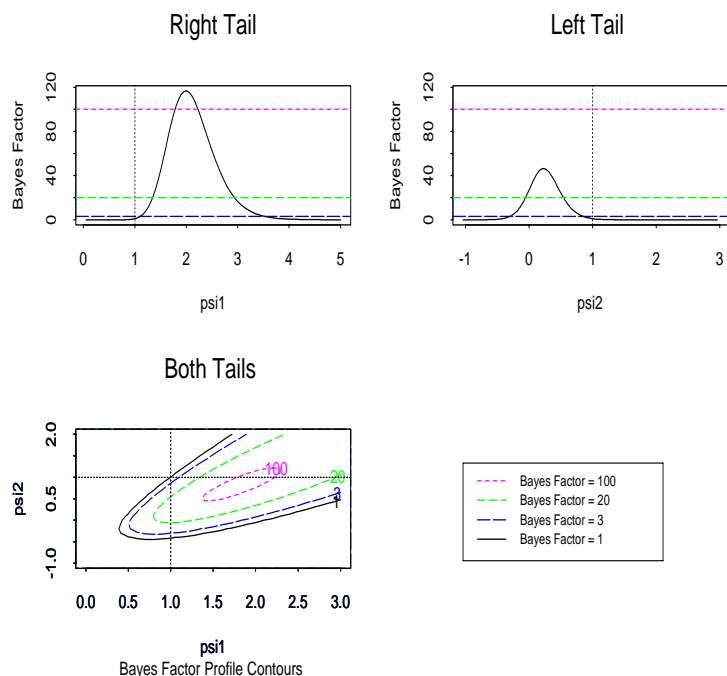


Fig. 5 Approximate Bayes Factor Profiles and Contours for the Beetle Mortality Data with $\sigma_p = 1$

Table 6 Minimal Deviances and Approximate Maximal Individual Bayes Factors for the Beetle Mortality Data (prior variance parameter $\sigma_p = 1$)

Model	Minimal Deviance	(ψ_1, ψ_2)	df	Maximal Bayes Factor	(ψ_1, ψ_2)
Right	3.96	(1.92,-)	5	116.66	(1.99,-)
Left	3.04	(-, .16)	5	46.41	(-, .21)
Both	2.81	(1.2, .3)	4	123.89	(1.8, .8)

From this we conclude that the Bayes factors clearly favor a right tail or both tail modification over a left tail modification. While the likelihood ratio test can be used to show that the reduction in deviance achieved by using a both tail modification over a right/or left tail modification is insignificant, we cannot compare right and left tail modifications, since they are not nested models. Graphically, we see that in Figure 4, the lines determining the

Table 7 Approximate Overall Bayes Factors for the Beetle Data with $\psi_l = -.5$

a	Pareto prior			$N(1, \sigma_\psi^2)$ truncated to $[\psi_l, \infty)$			
	Right Tail	Left Tail	Both Tails	σ_ψ	Right Tail	Left Tail	Both Tails
.75	10.10	8.08	.59	1.00	27.76	8.75	15.37
1.00	9.80	9.48	.60	2.00	26.65	6.55	10.83
1.50	7.77	11.00	.90	4.00	18.14	4.15	5.28

point (1,1) (corresponding to logistic link) intersect the confidence regions, suggesting that single tail modifications are sufficient.

We now take into account the link uncertainty by considering overall Bayes factors, which are given in Table 7 for $\psi_l = -.5$. We use $\psi_l = -.5$ since the observed success probabilities vary between .1 and .9. Therefore we would like to allow for links which take this restriction into account. For the Pareto priors a single tail modification is sufficient, but the difference between a left tail or right tail is minimal. The truncated normal priors favor a right tail modification over a left tail or both tail modification. The difference in the results for the two prior specifications can be explained as follows. Pareto priors favor small link values while truncated normal priors favor large link values. In this data set, this corresponds to left tail modifications for the Pareto priors and right tail modifications for the truncated normal priors.

This data set has also been considered in [6] p. 108-112, where they allowed for the inclusion of a quadratic term on the original CS_2 scale in a logistic model. This yields a residual deviance of 3.08 with 5 degrees of freedom. We can now use Bayes factors to decide if the right tail link fit is preferable over the inclusion of a quadratic term on the original CS_2 scale. Note that these models are again nonnested. The corresponding Bayes factor

is given by

$$B_{\psi_1=1.99} \times \frac{Pr(\mathbf{Y}|M_{\psi_1=1,x=\log(CS_2)})}{Pr(\mathbf{Y}|M_{\psi_1=1,x=(CS_2,CS_2^2)})} = 116.66 \times .0011 = .1280 = \frac{1}{7.80},$$

and the overall Bayes factor using a right tail modification

$$B_r \times \frac{Pr(\mathbf{Y}|M_{\psi_1=1,x=\log(CS_2)})}{Pr(\mathbf{Y}|M_{\psi_1=1,x=(CS_2,CS_2^2)})}$$

varies between $\frac{1}{32}$ and $\frac{1}{117}$ for the different link prior specifications. This shows that a logistic model using a quadratic term on the original scale is favored over a right tail link family. In [6] on p. 140 it was noted that a complementary log-log model for the link parameter fits the data as well as the logistic model using a quadratic term. It is argued that the complementary log-log model would be preferable since it has fewer parameters, but this ignores the uncertainty in the choice of link function.

5 Discussion

We have presented a Bayesian approach to model selection in GLM's with parametric link using Bayes factors to account for structural model uncertainty (see [16]) such as the choice of link in a GLM. This involves a continuous model expansion over ordinary GLM's when a particular link family was considered as well as a discrete model expansion when different link families were compared. In addition we were able to jointly assess the choice of link together with the choice of the set of independent parameters to include in the model. This involves the comparison of nonnested models, which cannot be carried out using classical model selection strategies based on significance tests.

We used reference proper priors for the regression parameters of a GLM with a fixed link function as suggested by [32]. These priors vary with the link parameter, reflecting the fact that the regression parameters are on different scales for different link functions. This reference proper prior avoids the problem of Bartlett's ([4]) or Lindley's ([23]) paradox and thus in this case Bayes factors have the advantage over posterior Bayes factors ([1]), p -values or the AIC criterion that they correctly identify the correct model in large samples, while the other criteria do not ([33]). Different prior distributions of the link parameter were investigated. Finally, the Bayes factors were approximated using the Laplace approximations given in [32]. With regard to prior sensitivity we observed the following. For the large menarche data set the qualitative conclusions are unchanged by the link prior specification, while for the smaller beetle data set we observed a moderate dependency on the prior specification. This kind of behavior is common in Bayesian analysis.

While Markov Chain Monte Carlo (MCMC) methods have been widely applied to GLM's and their hierarchical extensions both in time and space (see for example the books by [18],[15],[7] and [3] and the references within), a fully Bayesian analysis with MCMC of GLM's with data selected link functions have been considered in [11] and [28]. These estimation methods are computer intensive and might require user written software. The methods presented in this paper can be used for a final analysis, or could be used to screen for plausible models, which could then be used as starting points for a complete Bayesian analysis. Note that our methods for calculating these Bayes factors only require software that is able to fit a GLM with

an arbitrary link. In particular, joint maximization over regression parameters and link parameters to determine the maximum likelihood estimator is not needed. Here, calculations were conducted in S-Plus using the `glm()` function together with integration functions in one or two dimensions.

It should be noted that Bayes factors address the issue of model choice, and not parameter estimation. For inference about model-independent quantities such as the log odds ratio of a treatment effect or the mean response at a particular value of the independent variables, methods for taking account of model uncertainty such as Bayesian model averaging (see for example [21]) are needed. This also allows a Bayesian alternative to the quantifications of change to quantities of interest when changing from a GLM with canonical link to one with noncanonical link. This was the goal of a paper by [13].

References

1. M. Aitkin. Posterior Bayes factors (disc: P128-142). *Journal of the Royal Statistical Society, Series B, Methodological*, 53:111–128, 1991.
2. F. J. Aranda-Ordaz. On two families of transformations to additivity for binary response data (corr: V70 p303). *Biometrika*, 68:357–363, 1981.
3. S. Banerjee, B. P. Carlin, and A.E. Gelfand. *Hierarchical Modeling and Analysis for spatial Data*. Chapman & Hall/CRC, Boca Raton, London, New York, Washington, D.C., 2004.
4. M. S. Bartlett. A comment on D. V. Lindley's statistical paradox. *Biometrika*, 44:533, 1957.
5. C.I. Bliss. The calculation of the dose-mortality curve. *Annals of Applied Biology*, 22:134–167, 1935.
6. D. Collett. *Modelling Binary Data*. Chapman & Hall, London, 2nd edition, 2002.
7. P. Congdon. *Applied Bayesian Modelling*. John Wiley & Sons, Chichester, U.K., 2003.
8. T. W. Copenhaver and P. W. Mielke. Quantit analysis: A quantal assay refinement. *Biometrics*, 33:175–186, 1977.
9. D. R. Cox and N. Reid. Parameter orthogonality and approximate conditional inference (c/r: P18-39). *Journal of the Royal Statistical Society, Series B, Methodological*, 49:1–18, 1987.
10. C. Czado. On link selection in generalized linear models. In *Advances in GLIM and Statistical Modelling. Proceedings of the GLIM92 Conference*, pages 60–65. Springer-Verlag (Berlin; New York), 1992.

11. C. Czado. Bayesian inference of binary regression models with parametric link. *Journal of Statistical Planning and Inference*, 41:121–140, 1994.
12. C. Czado. On selecting parametric link transformation families in generalized linear models. *Journal of Statistical Planning and Inference*, 61:125–139, 1997.
13. C. Czado and A. Munk. Noncanonical links in generalized linear models - when is the effort justified? *Journal of Statistical Planning and Inference*, 87:317–345, 2000.
14. C. Czado and T. J. Santner. The effect of link misspecification on binary regression inference. *Journal of Statistical Planning and Inference*, 33:213–231, 1992.
15. D. Dey, S. Ghosh, and B. Mallik. *Generalized Linear Models: A Bayesian Perspective*. Marcel Dekker, New York, 2000.
16. D. Draper. Assessment and propagation of model uncertainty (disc: P71-97). *Journal of the Royal Statistical Society, Series B, Methodological*, 57:45–70, 1995.
17. A. E. Gelfand and D. K. Dey. Bayesian model choice: Asymptotics and exact calculations. *Journal of the Royal Statistical Society, Series B, Methodological*, 56:501–514, 1994.
18. W.R. Gilks, S. Richardson, and D.J. Spiegelhalter. *Markov Chain Monte Carlo in Practice*. Chapman & Hall, London, 1996.
19. V. M. Guerrero and R. A. Johnson. Use of the Box-Cox transformation with binary response models. *Biometrika*, 69:309–314, 1982.
20. C. Han and B.P. Carlin. MCMC methods for computing Bayes factors: a comparative review. *Journal of the American Statistical Association*, 96:1122–1132, 2001.
21. J. A. Hoeting, D. Madigan, A. E. Raftery, and C. T. Volinsky. Bayesian model averaging: A tutorial (disc: P401-417). *Statistical Science*, 14:382–401, 1999. Corrected version available at www.stat.washington.edu/www/research/online/hoeting1999.pdf.
22. R. E. Kass and A. E. Raftery. Bayes factors. *Journal of the American Statistical Association*, 90:773–795, June 1995.
23. D. V. Lindley. On the presentation of evidence. *Biometrika*, 44:187–192, 1957.
24. P. McCullagh and J.A. Nelder. *Generalized Linear Models*. Chapman and Hall, London, 2nd edition, 1989.
25. H. Milicer and F. Szcotka. Age at menarche in Warsaw girls in 1965. *Human Biology*, 40:199–203, 1966.
26. M. A. J. Montfort and A. Otten. Quantal response analysis: Enlargement of the logistic model with a kurtosis parameter. *Biometrical Journal*, 18:371–380, 1976.
27. B. J. T. Morgan. Observations on quantit analysis. *Biometrics*, 39:879–886, 1983.
28. M.A. Newton, C. Czado, and R. Chappell. Bayesian inference for semiparametric binary regression. *J. Amer. Stat. Assoc.*, 91:142–153, 1996.
29. I. Ntzoufras, P. Dellaportas, and J. J. Forster. Bayesian variable and link determination for generalised linear models. *preprint*, 2001.
30. D. Pregibon. Goodness of link tests for generalized linear models. *Applied Statistics*, 29:15–24, 1980.
31. R. L. Prentice. Generalization of the probit and logit methods for dose response curves. *Biometrics*, 32:761–768, 1976.
32. A. E. Raftery. Approximate Bayes factors and accounting for model uncertainty in generalized linear models. *Biometrika*, 83:251–266, 1996.
33. G. Schwartz. Estimating the dimension of a model. *Annals of Statistics*, 6:461–64, 1978.

34. T. Stukel. Generalized logistic models. *Journal of the American Statistical Association*, 83:426–431, 1988.
35. J. M. G. Taylor. The cost of generalizing logistic regression. *Journal of the American Statistical Association*, 83:1078–1083, 1988.
36. J. M. G. Taylor, A.L. Siqueira, and R.E. Weiss. The cost of adding parameters to a model. *Journal of the Royal Statistical Society, Series B, Methodological*, 58:593–607, 1996.
37. V. Viallefont, A. E. Raftery, and S. Richardson. Variable selection and Bayesian model averaging in epidemiological case-control studies. *Statistics in Medicine*, 20:3215–3230, 2001.
38. A.S Whittmore. Transformations to linearity in binary regression. *SIAM Journal on Applied Mathematics*, 43:703–710, 1983.