

n° 2010-43

**Regularization in Regression :
Comparing Bayesian and
Frequentist Methods in a
Poorly Informative Situation**

**G. CELEUX¹ – M. EL ANBARI²
J.-M. MARIN³ – C. P. ROBERT⁴**

Les documents de travail ne reflètent pas la position de l'INSEE et n'engagent que leurs auteurs.

Working papers do not reflect the position of INSEE but only the views of the authors.

¹ *Project SELECT*, INRIA Saclay, Université Paris Sud, Orsay, France.

² Université Caddi Ayyad, Marrakech, Maroc.

³ Université de Montpellier 2, France et CREST-INSEE, Paris.

⁴ Université Paris-Dauphine, CEREMADE, France et CREST, Paris.

Regularization in regression: comparing Bayesian and frequentist methods in a poorly informative situation *

GILLES CELEUX

Project SELECT, INRIA Saclay, Université Paris Sud, Orsay, France

MOHAMMED EL ANBARI

Université Caddi Ayyad, Marrakech, Maroc

JEAN-MICHEL MARIN

*Institut de Mathématiques et Modélisation de Montpellier,
Université de Montpellier 2, France & CREST, INSEE, Paris*

& CHRISTIAN P. ROBERT

Université Paris-Dauphine, CEREMADE, France & CREST, Paris

October 5, 2010

Abstract

We propose a global noninformative approach for Bayesian variable selection that builds on Zellner's g -priors and is similar to Liang et al. (2008). Our proposal does not require any kind of calibration. In the case of a benchmark, we compare Bayesian and frequentist regularization approaches under a low informative constraint when the number of variables is almost equal to the number of observations. The simulated and real dataset experiments we present here highlight the appeal of Bayesian regularization methods, when compared with alternatives. They dominate frequentist methods in the sense they provide smaller prediction errors while selecting the most relevant variables in a parsimonious way.

Keywords: Model choice, regularization methods, noninformative priors, Zellner's g -prior, calibration, Lasso, elastic net, Dantzig selector.

1 Introduction

Given a response variable, y and a collection of p associated potential predictor variables x_1, \dots, x_p , the classical linear regression model imposes a linear dependence on the conditional expectation (Rao, 1973)

$$\mathbb{E}[y|x_1, \dots, x_p] = \beta_0 + \beta_1 x_1 + \dots + \beta_p x_p.$$

*This paper is part of Mohammed EL Anbari's PhD thesis. This work has been partly supported by the Agence Nationale de la Recherche (ANR, 212, rue de Bercy 75012 Paris) through the 2006-2008 project *Misgepop* and the 2009-2012 project ANR-09-BLAN-01 *EMILE* for the last two authors, and by Institut Universitaire de France for the last author. Jean-Michel Marin and Christian P. Robert are grateful to the participants to the BIRS 07w5079 meeting on "Bioinformatics, Genetics and Stochastic Computation: Bridging the Gap" for their helpful comments. Discussions in 2007 in Banff with Sylvia Richardson and in Roma with Jim Berger and Paul Speckman are also gratefully acknowledged. Given that Arnold Zellner sadly passed away last August, we would like to dedicate this paper to the memory of this leading Bayesian thinker who influenced so much the field and will continue to do so.

A fundamental inferential direction for those models relates to the variable selection problem, namely that only variables of relevance should be kept within the regression while the others should be removed. While we cannot discuss at length the potential applications of this issue, variable selection is particularly relevant when the number p of regressors is larger than the number n of observations (as in microarray and other genetic data analyzes).

To deal with poorly or ill-posed regression problems, many regularization methods have been proposed in the literature, like ridge regression (Hoerl and Kennard, 1970) and Lasso (Tibshirani, 1996). Recently the interest for frequentist regularization methods has increased and this has produced a flurry of methods (see, among others, Candès and Tao, 2007, Zou and Hastie, 2005, Zou, 2006, Yuan and Lin, 2007).

However, a natural approach for regularization is to follow the Bayesian paradigm as demonstrated recently by the Bayesian Lasso of Park and Casella (2008). The amount of literature on Bayesian variable selection is quite enormous (a small subset of which is, for instance, Mitchell and Beauchamp, 1988, George and McCulloch, 1993, Chipman, 1996, Smith and Kohn, 1996, George and McCulloch, 1997, Dupuis and Robert, 2003, Brown and Vannucci, 1998, Philips and Guttman, 1998, George, 2000, Kohn et al., 2001, Nott and Green, 2004, Schneider and Corcoran, 2004, Casella and Moreno, 2006, Cui and George, 2008, Liang et al., 2008, Bottolo and Richardson, 2010). The number of approaches and scenarii that have been advanced to undertake the selection of the most relevant variables given a set of observations is quite large, presumably due to the vague decisional setting induced by the question *Which variables do matter?* Such a variety of resolutions signals a lack of agreement between the actors in the field and we thus feel the need to propose what we deem is a coherent resolution of the Bayesian prior modelling in noninformative settings. We will more closely follow variations with the recent approaches of Cui and George (2008), Liang et al. (2008), and Bottolo and Richardson (2010).

As in Liang et al. (2008) and Bottolo and Richardson (2010), we focus on the use of the g -prior, introduced by Zellner (1986). In fact, while this prior has a long history and while it reduces the prior input to a single integer, g , we stress that the influence of this remaining prior factor is long-lasting and that large values of g are no guarantee of negligible effects, in connection with the Bartlett or Lindley–Jeffreys paradoxes (Bartlett, 1957, Lindley, 1957, Robert, 1993), as illustrated for instance in Celeux et al. (2006) or Marin and Robert (2007). In order to alleviate this influence, we then follow a most natural Bayesian path and simply proceed to the construction of a noninformative alternative that eliminates the impact of the factor g , based on the Jeffreys prior associated with a hierarchical model. We thus claim to have achieved here a calibration-free Bayesian approach to the variable selection problem.

While the hierarchical extension of the g -prior can be found in the earlier literature, either through an implicit modelling as in Zellner and Siow (1980) or through a clear call to an hyperprior as in Liang et al. (2008) and Bottolo and Richardson (2010), two major differences with those recent papers are that we assume that there always is an intercept coefficient β_0 in the regression model and that we further process this parameter β_0 in the same way as we process the other regression parameters of the linear model. As a result and in contrast with those recent papers, we do avoid the need to call for a proper prior on g , which necessarily depends on an hyperparameter. Obviously, the counterpart of this global approach is that the intercept is no longer estimated under standard invariance constraints like location equivariance (Berger et al., 1998). We nonetheless argue that the prior independence between β_0 and the other parameters of β , found in the models of Liang et al. (2008) and Bottolo and Richardson (2010), does not have to hold and that

the lack of location invariance resulting from our approach is a negligible problem (if any) from practical and methodological viewpoints. (Guo and Speckman, 2009, have furthermore established the consistency of the Bayes factors based on such priors.)

Beside this new Bayesian modelling in linear models, the purpose of our paper is to compare the frequentist and the Bayesian point of views in regularization when n remains (slightly) greater than p . As a matter of fact, the use of g -prior implies that the sample size n has to be greater than p . This comparison is considered from both the predictive and the explicative point of views.

The plan of the paper is as follows: we recall the details on Zellner's (1986) original g prior in Section 2, as well as discuss the potential choices of g in Section 2, then proceed to construct our hierarchical noninformative alternative in Section 3 including the re-derivation of the Jeffreys prior in this setting. Sections 4 and 4.3 compare the results of Bayesian and frequentist methods on simulated and real datasets. Section 5 concludes the paper.

2 Zellner's g -priors

Following standard notation, we introduce a variable $\gamma \in \Gamma = \{0, 1\}^{\otimes p}$ that indicates which variables are active in the regression, excluding the constant vector corresponding to the intercept that is assumed to be always present in the linear regression model.

We observe $\mathbf{y}, \mathbf{x}_1, \dots, \mathbf{x}_p \in \mathbb{R}^n$, the model \mathcal{M}_γ is defined as the conditional distribution

$$\mathbf{y}|\mathbf{X}, \gamma, \boldsymbol{\beta}^\gamma, \sigma^2 \sim \mathcal{N}_n(\mathbf{X}^\gamma \boldsymbol{\beta}^\gamma, \sigma^2 I_n), \quad (1)$$

where

- ▶ $p_\gamma = \sum_{i=1}^p \gamma_i$,
- ▶ \mathbf{X}^γ is the $(n, p_\gamma + 1)$ matrix which columns are made of the vector $\mathbf{1}_n$ and of the variables \mathbf{x}_i for which $\gamma_i = 1$,
- ▶ $\boldsymbol{\beta}^\gamma \in \mathbb{R}^{p_\gamma+1}$ and $\sigma^2 \in \mathbb{R}_+^*$ are unknown parameters.

The same symbol for the parameter σ^2 is used in all models. For model \mathcal{M}_γ , Zellner's g -prior is given by

$$\begin{aligned} \boldsymbol{\beta}^\gamma|\mathbf{X}, \gamma, \sigma^2 &\sim \mathcal{N}_{p_\gamma+1}(\tilde{\boldsymbol{\beta}}^\gamma, g_\gamma \sigma^2 ((\mathbf{X}^\gamma)' \mathbf{X}^\gamma)^{-1}), \\ \pi(\sigma^2|\mathbf{X}, \gamma) &\propto \sigma^{-2}. \end{aligned}$$

The experimenter chooses the prior expectation $\tilde{\boldsymbol{\beta}}^\gamma$ and g_γ . For such a prior, we obtain the classical average between prior and observed regressors,

$$\mathbb{E}(\boldsymbol{\beta}^\gamma|\mathbf{X}, \gamma, \mathbf{y}) = \frac{g_\gamma \tilde{\boldsymbol{\beta}}^\gamma + \boldsymbol{\beta}^\gamma}{g_\gamma + 1}.$$

This prior is traditionally called Zellner's g -prior in the Bayesian folklore because of the use of the constant g_γ by Zellner (1986) in front of Fisher's information matrix $((\mathbf{X}^\gamma)' \mathbf{X}^\gamma)^{-1}$. Its appeal is that, by using the information matrix as a global scale,

- ▶ it avoids the specification of a whole prior covariance matrix, which would be a tremendous task;

- it allows for a specification of the constant g_γ in terms of observational units, or virtual prior pseudo-observations in the sense of de Finetti (1972).

A fundamental feature of the g -prior is that this prior is improper, due to the use of an infinite mass on σ^2 . From a theoretical point of view, this should jeopardize the use of posterior model probabilities since these probabilities are not uniquely scaled under improper priors, because there is no way of eliminating the residual constant factor in those priors (DeGroot, 1973, Kass and Raftery, 1995, Robert, 2001). However, under the assumption that σ^2 is a parameter that is common to all models \mathcal{M}_γ , Berger et al. (1998) develop a framework that allows to work with a single improper prior that is common to all models (see also Marin and Robert, 2007). A fundamental appeal of Zellner’s g -prior in model comparison and in particular in variable selection is its simplicity, since it reduces the prior input to the sole specification of a scale parameter g .

At this stage, we need to point out that an alternative g -prior is often used (Berger et al., 1998, Fernandez et al., 2001, Liang et al., 2008, Bottolo and Richardson, 2010), by singling out the intercept parameter in the linear regression. By first assuming a centering of the covariates, i.e. $\mathbf{1}'_n \mathbf{x}_i = 0$ for all i ’s, the intercept α is given a flat prior while the other parameters of β^γ are associated with a corresponding g -prior. Thus, there is an alternative model to model \mathcal{M}_γ , which we denote by model $\mathcal{M}_\gamma^{\text{inv}}$ in this paper to stress the distinctions between both representation and which is such that

$$\mathbf{y}|\mathbf{X}, \gamma, \alpha, \beta_{\text{inv}}^\gamma, \sigma^2 \sim \mathcal{N}_n(\alpha \mathbf{1}_n + \mathbf{X}_{\text{inv}}^\gamma \beta_{\text{inv}}^\gamma, \sigma^2 I_n), \quad (2)$$

where

- $\mathbf{X}_{\text{inv}}^\gamma$ the (n, p_γ) matrix which columns are made of the variables \mathbf{x}_i for which $\gamma_i = 1$,
- $\alpha \in \mathbb{R}$, $\beta_{\text{inv}}^\gamma \in \mathbb{R}^{p_\gamma}$ and $\sigma^2 \in \mathbb{R}_+^*$ are unknown parameters.

The parameters σ^2 and α are denoted the same way across all models and rely on the same prior. Namely, for model $\mathcal{M}_\gamma^{\text{inv}}$, the corresponding Zellner’s g -prior is given by

$$\begin{aligned} \beta_{\text{inv}}^\gamma | \mathbf{X}, \gamma, \sigma^2 &\sim \mathcal{N}_{p_\gamma}(\tilde{\beta}_{\text{inv}}^\gamma, g_\gamma \sigma^2 ((\mathbf{X}_{\text{inv}}^\gamma)' \mathbf{X}_{\text{inv}}^\gamma)^{-1}), \\ \pi(\alpha, \sigma^2 | \mathbf{X}, \gamma) &\propto \sigma^{-2}. \end{aligned}$$

In that case, we obtain

$$\mathbb{E}(\beta_{\text{inv}}^\gamma | \mathbf{X}, \gamma, \mathbf{y}) = \frac{g_\gamma \hat{\beta}_{\text{inv}}^\gamma + \tilde{\beta}_{\text{inv}}^\gamma}{g_\gamma + 1},$$

and

$$\mathbb{E}(\alpha | \mathbf{X}, \gamma, \mathbf{y}) = \bar{y} = \frac{1}{n} \sum_{i=1}^n y_i.$$

For model \mathcal{M}_γ , in a noninformative setting, we can for instance choose $\tilde{\beta}^\gamma = \mathbf{0}_{p_{\gamma+1}}$ and g_γ large. However, as pointed out in Marin and Robert (2007, Chapter 3) among others, there is a lasting influence of g_γ over the resulting inference and it is impossible to “let g_γ go to infinity” to eliminate this influence, because of the Bartlett and Lindley-Jeffreys (Bartlett, 1957, Lindley, 1957, Robert, 1993) paradoxes that an infinite value of g_γ ends up selecting the null model, regardless of the information brought by the data. For this reason, data-dependent versions of g_γ have been proposed with various degrees of justification:

- ▶ Kass and Wasserman (1995) use $g_\gamma = n$ so that the amount of information about the parameters contained in the prior equals the amount of information brought by one observation. As shown by Foster and George (1994), for n large enough this perspective is very close to using the Schwarz (Kass and Wasserman, 1995) or BIC criterion in that the log-posterior corresponding to $g = n$ is equal to the penalized log-likelihood of this criterion.
- ▶ Foster and George (1994) and George and Foster (2000) propose $g_\gamma = p_\gamma^2$, in connection with the Risk Inflation Criterion (RIC) that penalizes the regression sum of squares.
- ▶ Fernandez et al. (2001) gather both perspectives in $g_\gamma = \max(n, p_\gamma^2)$ as a conservative bridge between BIC and RIC, a choice that they christened “benchmark prior”.
- ▶ George and Foster (2000) and Cui and George (2008) resort to empirical Bayes techniques.

These solutions, while commendable since based on asymptotic properties (see in particular Fernandez et al., 2001 for consistency results), are nonetheless unsatisfactory in that they depend on the sample size and involve a degree of arbitrariness, while a complete Bayesian solution is readily available, as demonstrated below.

3 Mixtures of g -priors

The most natural Bayesian approach to solving the uncertainty on the parameter $g_\gamma = g$ is indeed to put a hyperprior on this parameter:

- ▶ This was implicitly proposed by Zellner and Siow (1980) since those authors introduced Cauchy priors on the β^γ 's since this corresponds to a g -prior augmented by a Gamma $\mathcal{G}a(1/2, n/2)$ prior on g^{-1} .
- ▶ For model $\mathcal{M}_\gamma^{\text{inv}}$, Liang et al. (2008), Cui and George (2008) and Bottolo and Richardson (2010) use

$$\beta_{\text{inv}}^\gamma | \mathbf{X}, \gamma, \sigma^2 \sim \mathcal{N}_{p_\gamma}(0_{p_\gamma}, g\sigma^2((\mathbf{X}_{\text{inv}}^\gamma)' \mathbf{X}_{\text{inv}}^\gamma)^{-1})$$

and an hyperprior of the form

$$\pi(\alpha, \sigma^2, g | \mathbf{X}, \gamma) \propto (1 + g)^{-a/2} \sigma^{-2},$$

with $a > 2$. This constraint on a is due to the fact that the hyperprior must be proper, in connection with the separate processing of the intercept α and the use of a Lebesgue measure as a prior on α . We note that a needs to be specified, $a = 3$ and $a = 4$ being the solutions favored by Liang et al. (2008).

- ▶ For model \mathcal{M}_γ , Celeux et al. (2006) and Marin and Robert (2007) used

$$\beta^\gamma | \mathbf{X}, \gamma, \sigma^2 \sim \mathcal{N}_{p_\gamma+1}(0_{p_\gamma+1}, g\sigma^2((\mathbf{X}^\gamma)' \mathbf{X}^\gamma)^{-1})$$

and an hyperprior of the form

$$\pi(\sigma^2, g | \mathbf{X}) \propto \sigma^{-2} g^{-1} \mathbb{I}_{\mathbb{N}^*}(g),$$

as a simple way of circumventing computational difficulties.

In the current paper, we propose a more convincing modelling, based on the remark that the Jeffreys prior is available. Indeed, if

$$\beta^\gamma | \mathbf{X}, \gamma, \sigma^2 \sim \mathcal{N}_{p_\gamma+1}(0_{p_\gamma+1}, g\sigma^2((\mathbf{X}^\gamma)' \mathbf{X}^\gamma)^{-1}),$$

then

$$\mathbf{y} | \mathbf{X}, \gamma, g, \sigma^2 \sim \mathcal{N}_{p_\gamma+1} \left(0_n, \sigma^2 \left[\mathbf{I}_n - \frac{g}{g+1} \mathbf{P}_\gamma \right]^{-1} \right),$$

where \mathbf{P}_γ is the orthogonal projector on the linear subspace spanned by the columns of \mathbf{X}^γ . Since, the Fisher information matrix is

$$\mathfrak{J}(\sigma^2, g) = \left(\frac{1}{2} \right) \begin{bmatrix} n/\sigma^4 & (p_\gamma + 1)/(\sigma^2(g+1)) \\ (p_\gamma + 1)/(\sigma^2(g+1)) & (p_\gamma + 1)/(g+1)^2 \end{bmatrix},$$

the corresponding Jeffreys prior on (σ^2, g) is

$$\pi(\sigma^2, g | \mathbf{X}) \propto \sigma^{-2}(g+1)^{-1}.$$

In contrast with our proposal, the prior of Liang et al. (2008) depends on a tuning parameter a . The choice of this hyperparameter a is sensitive and its influence is unfortunately non-vanishing against an increase of the number of observations n , since g has a significant influence on the Bayesian analysis of the linear model, as discussed earlier. As far as we can judge, it is difficult to propose any guideline for choosing a . However, there also exist arguments to back up this prior modelling, including invariance under location-scale transforms. As obviously deduced from later formulae, our modelling and prior choices ensure scale invariance but not location invariance. In order to ensure location invariance, it would be necessary to center the observation variable y as well as the dependent variables X . Obviously, this centering of the data is completely unjustified from a Bayesian perspective and it further creates artificial correlations between observations. In conclusion, we argue that the lack of location invariance only pertains to quite specific and somehow artificial situations and that it is negligible in most situations. There is therefore no justification in using artefacts to ensure location invariance.

The main consequence of this modelling is that, for the selected prior distributions, there exists a closed-form representation for posterior quantities in that

$$\pi(\gamma, g | \mathbf{X}, \mathbf{y}) \propto (g+1)^{n/2-(p_\gamma+1)/2-1} (1 + g(1 - \mathbf{y}' \mathbf{P}_\gamma \mathbf{y} / \mathbf{y}' \mathbf{y}))^{-n/2}$$

and

$$\pi(\gamma | \mathbf{X}, \mathbf{y}) \propto \frac{{}_2F_1(n/2, 1; (p_\gamma + 3)/2; \mathbf{y}' \mathbf{P}_\gamma \mathbf{y} / \mathbf{y}' \mathbf{y})}{p_\gamma + 1},$$

where ${}_2F_1$ is the Gaussian hypergeometric function (Butler and Wood, 2002). We can thus proceed to undertake Bayesian variable selection without resorting at all to numerical methods like simulation (Marin and Robert, 2007). Moreover, the shrinkage factor due to the Bayesian modelling can also be expressed in closed form as

$$\begin{aligned} \mathbb{E}(g/(g+1) | \mathbf{X}, \gamma, \mathbf{y}) &= \frac{\int_0^\infty g(g+1)^{n/2-(p_\gamma+1)/2-2} (1 + g(1 - \mathbf{y}' \mathbf{P}_\gamma \mathbf{y} / \mathbf{y}' \mathbf{y}))^{-n/2} dg}{\int_0^\infty (g+1)^{n/2-(p_\gamma+1)/2-1} (1 + g(1 - \mathbf{y}' \mathbf{P}_\gamma \mathbf{y} / \mathbf{y}' \mathbf{y}))^{-n/2} dg} \\ &= \frac{{}_2F_1(n/2, 2; (p_\gamma + 3)/2 + 1; \mathbf{y}' \mathbf{P}_\gamma \mathbf{y} / \mathbf{y}' \mathbf{y})}{(p_\gamma + 3) {}_2F_1(n/2, 1; (p_\gamma + 3)/2; \mathbf{y}' \mathbf{P}_\gamma \mathbf{y} / \mathbf{y}' \mathbf{y})}. \end{aligned}$$

This obviously leads to straightforward representations for Bayes estimates. If \mathbf{X}_{new} is a $q \times p$ matrix containing q new values of the explanatory variables for which we would like to predict the corresponding response \mathbf{y}_{new} , the Bayesian predictor of \mathbf{y}_{new} is given by

$$\begin{aligned}\hat{\mathbf{y}}_{\text{new}}^{\gamma} &= \mathbb{E}[\mathbf{y}_{\text{new}}|\mathbf{X}_{\text{new}}, \mathbf{X}, \gamma, \mathbf{y}] \\ &= 2 \frac{{}_2F_1(n/2, 2; (p_{\gamma} + 3)/2 + 1; \mathbf{y}'\mathbf{P}_{\gamma}\mathbf{y}/\mathbf{y}'\mathbf{y})}{(p_{\gamma} + 3) {}_2F_1(n/2, 1; (p_{\gamma} + 3)/2; \mathbf{y}'\mathbf{P}_{\gamma}\mathbf{y}/\mathbf{y}'\mathbf{y})} \mathbf{X}_{\text{new}}\hat{\boldsymbol{\beta}}^{\gamma}.\end{aligned}$$

Similarly, the Bayesian model averaging predictor of \mathbf{y}_{new} is given by

$$\begin{aligned}\hat{\mathbf{y}}_{\text{new}} &= \mathbb{E}[\mathbf{y}_{\text{new}}|\mathbf{X}_{\text{new}}, \mathbf{X}, \mathbf{y}] \\ &= 2 \frac{\sum_{\gamma \in \Gamma} {}_2F_1(n/2, 2; (p_{\gamma} + 3)/2 + 1; \mathbf{y}'\mathbf{P}_{\gamma}\mathbf{y}/\mathbf{y}'\mathbf{y}) / [(p_{\gamma} + 1)(p_{\gamma} + 3)]}{\sum_{\gamma \in \Gamma} {}_2F_1(n/2, 1; (p_{\gamma} + 3)/2; \mathbf{y}'\mathbf{P}_{\gamma}\mathbf{y}/\mathbf{y}'\mathbf{y}) / (p_{\gamma} + 1)} \mathbf{X}_{\text{new}}\hat{\boldsymbol{\beta}}^{\gamma}.\end{aligned}\tag{3}$$

This numerical simplification in the derivation of Bayesian estimates and predictors is also available in Liang et al. (2008) and exploited further in Bottolo and Richardson (2010).

4 Numerical comparisons

In this section, we present the results of several numerical experiments aimed at comparing the behavior of Bayesian variable selection and of some (non-Bayesian) popular recent regularization methods in regression when considered from a variable selection point of view. The regularization methods that we consider are the Lasso, the Dantzig selector, and elastic net, as described in Section 4.1. The Bayesian variable selection procedures we consider in terms of the strategies selecting the hyperparameter g in Zellner's g -priors. They are described in Table 1. We have also included in this comparison the highly standard AIC and BIC penalized likelihood criteria.

4.1 Regularization methods

1) **The Lasso:** The Lasso, introduced by Tibshirani (1996), is a shrinkage method for linear regression. It is defined as the solution to the following ℓ_1 penalized least squares optimization problem

$$\hat{\boldsymbol{\beta}}_{\text{Lasso}} = \arg \min_{\boldsymbol{\beta}} \|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}\|_2^2 + \lambda \sum_{j=1}^p |\beta_j|,$$

where λ is a positive tuning parameter.

2) **The Dantzig Selector:** Candès and Tao (2007) introduced the Dantzig Selector as an alternative to the Lasso. The Dantzig Selector is the solution to the optimization problem

$$\min_{\boldsymbol{\beta} \in \mathbb{R}^p} \|\boldsymbol{\beta}\|_1 \quad \text{subject to} \quad \|\mathbf{X}^t(\mathbf{y} - \mathbf{X}\boldsymbol{\beta})\|_{\infty} \leq \lambda,$$

where λ is a positive tuning parameter. The constraint $\|\mathbf{X}^t(\mathbf{y} - \mathbf{X}\boldsymbol{\beta})\|_{\infty} \leq \lambda$ can be viewed as a relaxation of the normal equation in the classical linear regression.

3) **The Elastic Net (Enet)**: The Lasso has at least two limitations: a) Lasso does not encourage grouped selection in the presence of high correlated covariates and b) for $p > n$ case Lasso can select at most n covariates. To overcome these limitations, Zou and Hastie (2005) proposed an elastic net that combines both ridge ℓ_2 and Lasso ℓ_1 penalties, i. e.

$$\hat{\beta}_{\text{Enet}} = \arg \min_{\beta} \|\mathbf{y} - X\beta\|_2^2 + \lambda \sum_{j=1}^p |\beta_j| + \mu \sum_{j=1}^p \beta_j^2,$$

where λ and μ are two positive tuning parameters.

4.2 Numerical experiments on simulated datasets

We have used five different simulated datasets, chosen as follows:

1. Example 1 (**sparse uncorrelated design**) corresponds to an uncorrelated covariate setting ($\rho = 0$), with $p = 10$ predictors and where the components of \mathbf{x}_i ($i = 1, \dots, 10$) are iid $\mathcal{N}_1(0, 1)$ realizations. The response is simulated as

$$\mathbf{y} \sim \mathcal{N}_n(2 + \mathbf{x}_2 + 2\mathbf{x}_3 - 2\mathbf{x}_6 - 1.5\mathbf{x}_7, I_n).$$

2. Example 2 (**sparse correlated design**) corresponds to a correlated case ($\rho = 0.9$), with $p = 10$ predictors and $\mathbf{x}_i = (\mathbf{z}_i + 3\mathbf{z}_{11})/\sqrt{10}$, for $i = 1, 2$, $\mathbf{x}_i = (\mathbf{z}_i + 3\mathbf{z}_{12})/\sqrt{10}$, for $i = 3, 4, 5$, and $\mathbf{x}_i = (\mathbf{z}_i + 3\mathbf{z}_{13})/\sqrt{10}$ for $i = 6, \dots, 10$, the components of \mathbf{z}_i ($i = 1, \dots, 13$) being iid $\mathcal{N}_1(0, 1)$ realizations. The use of common terms in the \mathbf{x}_i 's obviously induces a correlation among those \mathbf{x}_i 's: the correlation between variables \mathbf{x}_1 and \mathbf{x}_2 is 0.9, as for the variables ($\mathbf{x}_3, \mathbf{x}_4$ and \mathbf{x}_5), and for the variables ($\mathbf{x}_6, \mathbf{x}_7, \mathbf{x}_8, \mathbf{x}_9$ and \mathbf{x}_{10}). There is no correlation between those three groups of variables. The response is simulated as

$$\mathbf{y} \sim \mathcal{N}_n(2 + \mathbf{x}_2 + 2\mathbf{x}_3 - 2\mathbf{x}_6 - 1.5\mathbf{x}_7, I_n).$$

3. Example 3 (**sparse noisy correlated design**) involves $p = 8$ predictors. Those variables are generated using a multivariate Gaussian distribution with correlations

$$\rho(\mathbf{x}_i, \mathbf{x}_j) = 0.5^{|i-j|}.$$

The response is simulated as

$$\mathbf{y} \sim \mathcal{N}_n(3\mathbf{x}_1 + 1.5\mathbf{x}_2 + 2\mathbf{x}_5, 9I_n).$$

4. Example 4 (**saturated noisy correlated design**) is the same as Example 3, except that the response is simulated as

$$\mathbf{y} \sim \mathcal{N}_n \left(0.85 \sum_{i=1}^8 \mathbf{x}_i, 9I_n \right).$$

5. Example 5 (**saturated correlated design**) is the same as Example 4, except that the response is simulated as

$$\mathbf{y} \sim \mathcal{N}_n \left(0.85 \sum_{i=1}^8 \mathbf{x}_i, I_n \right).$$

Each dataset consists of a training set of size $n = 15$, on which the regression model has been fitted and a test set T of size $n_T = 200$ for assessing the performances. Tuning parameters in the Lasso, the Dantzig selector (DZ), and the elastic net (ENET) have been selected by minimizing the cross-validation prediction error through leave-one-out. For each example, 100 independent datasets have been simulated. We use three measures for performances:

1. The root mean squared error (MSE)

$$\text{MSE}_y = \sqrt{\sum_{i=1}^{n_T} (y_i - \hat{y}_i)^2 / n_T},$$

\hat{y}_i being the prediction of y_i in the test set;

2. HITS: the number of correctly identified influential variables;
3. FP (False Positives): the number of non-influential variables declared as influential.

Using those five different datasets, we compare the variable selection methods listed in Table 1. The results of those different methods are summarized in Tables 2–6. In the Bayesian approach, the set of variables is selected according to the maximum posterior probability $\pi(\gamma|\mathbf{X}, \mathbf{y})$ and the predictive is obtained via the Bayesian model averaging predictor of (3). Note that we also ran a test about the modified behavior of our approach (NIMS, which stands for non-informative mixture selection) when the response variable is drifted by 20, i.e. using $\mathbf{y} + 20$ instead of \mathbf{y} . As we expected, a change in the location of the response does not affect at all the statistical performances of the analysis obtained with our approach.

AIC	Akaike Information Criterion
BIC	Bayesian Information Criterion
BRIC	g prior with $g = \max(n, p^2)$
EB-L	Local EB estimate of g in g -prior
EB-G	Global EB estimate of g in g -prior
ZS-N	Base model in Bayes factor taken as the null model (Liang et al., 2008)
ZS-F	Base model in Bayes factor taken as the full model (Liang et al., 2008)
HG-3	Hyper- g prior with $a = 3$ (Liang et al., 2008)
HG-4	Hyper- g prior with $a = 4$ (Liang et al., 2008)
NIMS	Our proposal based on the non-invariant model
LASSO	Lasso (Tibshirani, 1996)
DZ	The Dantzig Selector (Candes and Tao, 2007)
ENET	The elastic-net (Zou and Hastie, 2005)

Table 1: Acronyms used for the variable selection methods.

The Bayesian procedures are clearly much more parsimonious than the regularization procedures in that they almost always avoid overfitting. (Tables corresponding to Examples 1, 3–5 for the performances are not reproduced in this paper since they either exhibit a constant number of hits or a quasi-constant MSE. In all examples, the false positive rate FP is zero for all Bayesian solutions and positive for regularization methods.). Except for the ZS-F scenario which behaves

slightly worse than the others, all the Bayesian procedures tested here produce the same selection of predictors. From a predictive viewpoint, Bayesian approaches also perform better than regularization approaches except for the noisy example (Example 4). We further note that the classical selection procedures based on AIC and BIC do not easily reject variables and are thus slightly worse than Bayesian and regularization procedures (which is not surprising for AIC). In all examples but the noisy one, the NIMS approach leads to optimal performances in that it selects the right covariates and only the right covariates, while achieving close to the minimal root mean squared error compared with all the other Bayesian solutions we considered. It also does almost systematically better than BIC and AIC. The case of Example 4 is rather extreme due to the large noise factor, but the standard regularization procedures manage to reduce the MSE in this case by close to 10% compared with the Bayesian procedures.

A global remark is that all Bayesian procedures have a very similar MSE and thus that they all correspond to the same regularization effect, except for ZS-F which does systematically worse. Obviously, the examples are such that no (Bayesian) procedure performs very poorly compared with the others. Even Steel’s simple BRIC stands its ground against the more elaborate alternatives, including the extreme Example 4. Note that Example 5 whose results are not reported here brings no leeway for comparing those Bayesian procedures among themselves because almost systematically all procedures include all variables.

Variables	1	2	3	4	5	6	7	8	9	10
AIC	0.19	1	1	0.28	0.19	1	1	0.21	0.23	0.18
BIC	0.11	1	1	0.14	0.05	1	1	0.07	0.05	0.04
BRIC	0.03	1	1	0.05	0.02	1	1	0.02	0.01	0.01
EB-L	0.03	1	1	0.05	0.02	1	1	0.03	0.01	0.01
EB-G	0.03	1	1	0.05	0.02	1	1	0.03	0.01	0.01
ZS-N	0.03	1	1	0.05	0.02	1	1	0.03	0.01	0.01
ZS-F	0.12	1	1	0.20	0.10	1	1	0.12	0.10	0.10
HG-3	0.03	1	1	0.04	0.02	1	1	0.03	0.01	0.01
HG-4	0.03	1	1	0.05	0.02	1	1	0.03	0.01	0.01
NIMS	0.03	1	1	0.04	0.01	1	1	0.02	0.01	0.01
LASSO	0.49	1	1	0.64	0.53	1	1	0.52	0.56	0.53
DZ	0.25	0.99	1	0.29	0.21	1	1	0.25	0.25	0.24
ENET	0.61	1	1	0.64	0.49	1	1	0.55	0.54	0.56

Table 2: Example 1: Relative frequencies of selections of the variables

4.3 Real datasets

The two datasets that we consider in this section are associated with a moderate number of variables against the number of observations.

Body fat dataset The body fat dataset has been first used by Penrose et al. (1985). The corresponding study aims at estimating the percentage of body fat from various body circumference measurements observed on 252 men. The thirteen regressor variables are:

	MSE _y	HITS	FP
AIC	1.10	4	1
BIC	1.08	4	0
BRIC	1.04	4	0
EB-L	1.04	4	0
EB-G	1.04	4	0
ZS-N	1.04	4	0
ZS-F	1.06	4	1
HG-3	1.04	4	0
HG-4	1.04	4	0
NIMS	1.07	4	0
LASSO	1.10	4	3
DZ	1.13	4	5
ENET	1.12	4	3

Table 3: Example 2: performances as MSE, HITS and FP

Variables	1	2	3	4	5	6	7	8	9	10
AIC	0.16	1	1	0.17	0.27	1	1	0.28	0.21	0.21
BIC	0.08	1	1	0.06	0.08	1	1	0.12	0.07	0.07
BRIC	0	1	1	0	0	1	1	0	0	0
EB-L	0.01	1	1	0	0.02	1	1	0.02	0.01	0.03
EB-G	0.01	1	1	0	0.02	1	1	0.02	0.01	0.03
ZS-N	0.01	1	1	0	0.02	1	1	0.02	0	0.03
ZS-F	0.10	1	1	0.08	0.13	1	1	0.19	0.13	0.10
HG-3	0.01	1	1	0	0.02	1	1	0.02	0	0.03
HG-4	0.01	1	1	0	0.02	1	1	0.02	0	0.03
NIMS	0.01	1	1	0	0.02	1	1	0.02	0.01	0.03
LASSO	0.55	1	1	0.57	0.53	1	1	0.59	0.53	0.52
DZ	0.81	1	1	0.77	0.79	1	1	0.86	0.75	0.78
ENET	0.59	1	1	0.62	0.56	1	1	0.57	0.58	0.58

Table 4: Example 2: Relative frequencies of selections of the variables

Variables	1	2	3	4	5	6	7	8
AIC	1	0.92	0.11	0.20	0.98	0.22	0.15	0.26
BIC	1	0.84	0.06	0.07	0.96	0.10	0.07	0.11
BRIC	1	0.80	0.05	0.04	0.96	0.06	0.05	0.07
EB-L	1	0.83	0.05	0.06	0.96	0.08	0.06	0.12
EB-G	1	0.84	0.05	0.06	0.97	0.08	0.07	0.13
ZS-N	1	0.82	0.05	0.05	0.96	0.07	0.06	0.10
ZS-F	1	0.88	0.08	0.10	0.97	0.14	0.08	0.17
HG-3	1	0.83	0.05	0.06	0.96	0.08	0.06	0.11
HG-4	1	0.84	0.05	0.06	0.97	0.09	0.07	0.12
NIMS	1	0.83	0.05	0.05	0.96	0.08	0.06	0.12
LASSO	1	0.99	0.34	0.40	1	0.46	0.35	0.44
DZ	1	0.98	0.23	0.32	0.99	0.28	0.18	0.25
ENET	1	0.99	0.47	0.42	0.98	0.45	0.30	0.37

Table 5: Example 3: Relative frequencies of selections of the variables

Variables	1	2	3	4	5	6	7	8
AIC	0.66	0.59	0.62	0.63	0.69	0.58	0.61	0.59
BIC	0.49	0.39	0.50	0.46	0.48	0.44	0.54	0.36
BRIC	0.48	0.36	0.51	0.43	0.48	0.39	0.54	0.33
EB-L	0.49	0.43	0.52	0.50	0.49	0.46	0.54	0.39
EB-G	0.52	0.45	0.52	0.51	0.50	0.47	0.55	0.40
ZS-N	0.49	0.41	0.51	0.47	0.49	0.44	0.55	0.38
ZS-F	0.54	0.47	0.52	0.56	0.50	0.49	0.57	0.42
HG-3	0.49	0.43	0.52	0.49	0.49	0.46	0.54	0.39
HG-4	0.52	0.45	0.52	0.51	0.50	0.47	0.55	0.40
NIMS	0.49	0.43	0.52	0.50	0.49	0.46	0.54	0.39
LASSO	0.91	0.86	0.90	0.93	0.94	0.86	0.87	0.79
DZ	0.78	0.80	0.87	0.89	0.88	0.80	0.88	0.69
ENET	0.86	0.90	0.94	0.96	0.92	0.92	0.92	0.81

Table 6: Example 4: Relative frequencies of selection of the variables

1. age,
2. weight (lbs),
3. height (inches),
4. neck circumference,
5. chest circumference,
6. abdomen 2 circumference,
7. hip circumference,
8. thigh circumference,
9. knee circumference,
10. ankle circumference,
11. biceps (extended) circumference,
12. forearm circumference,
13. wrist circumference.

In order to investigate the performances of the different methods, the Penrose et al. (1985) dataset has been split 25 times into a training set of 151 observations and a test set of 101 observations. Tuning parameter for the frequentist regularization methods have been chosen by minimizing the (ten fold) cross-validated prediction error.

For this dataset, the Bayesian procedures we investigated are much more parsimonious than the standard regularization procedures, as shown in Table 7. There is no variability in the prediction MSE but regularization procedures require about twice to three times as many covariates to achieve this goal, being much less parsimonious than AIC. (We stress that the MSE is computed by model averaging for the Bayesian procedures.) As in the simulation experiment, all Bayesian approaches are highly similar, except for ZS-F which remains more open to incorporating the last two covariates.

Ozone data This second benchmark dataset is used in Marin and Robert (2007) and consists of daily measurements of the maximum ozone concentration and of eight meteorological variables near Los Angeles. Those variables are

1. the daily ozone concentration (maximum one hour average, parts per million) at Upland, CA which is the response variable;
2. the Vandenburg 500 millibar pressure height (m);
3. the wind speed (mph) at Los Angeles International Airport (LAX);
4. the humidity (percent) at LAX;
5. the Sandburg Air Force Base temperature (F°);

	MSE_y	Median number of selected variables
AIC	4.58	6
BIC	4.60	4
BRIC	4.51	3
EB-L	4.52	3
EB-G	4.52	4
ZS-N	4.52	3
ZS-F	4.49	4
HG-3	4.54	3
HG-4	4.56	3
NIMS	4.50	2
LASSO	4.54	8
DZ	4.51	11
ENET	4.54	9

Table 7: Body fat dataset: MSE and median number of selected variables

Variables	1	2	3	4	5	6	7	8	9	10	11	12	13
AIC	0.44	0.84	0.16	0.64	0.04	1.00	0.20	0.16	0.08	0.16	0.44	0.80	0.88
BIC	0.08	0.84	0.08	0.32	0.00	1.00	0.12	0.08	0.04	0.00	0.16	0.28	0.40
BRIC	0.08	0.84	0.08	0.32	0.00	1.00	0.12	0.08	0.04	0.00	0.16	0.24	0.40
EB-L	0.08	0.84	0.08	0.32	0.00	1.00	0.12	0.08	0.04	0.00	0.16	0.28	0.40
EB-G	0.08	0.88	0.08	0.36	0.00	1.00	0.08	0.08	0.04	0.00	0.20	0.36	0.40
ZS-N	0.08	0.84	0.08	0.32	0.00	1.00	0.12	0.08	0.04	0.00	0.16	0.24	0.40
ZS-F	0.20	0.84	0.12	0.40	0.00	1.00	0.12	0.12	0.08	0.04	0.24	0.60	0.68
HG-3	0.08	0.84	0.08	0.32	0.00	1.00	0.12	0.08	0.04	0.00	0.16	0.28	0.40
HG-4	0.08	0.88	0.08	0.32	0.00	1.00	0.08	0.08	0.04	0.00	0.16	0.36	0.40
NIMS	0.04	0.88	0.04	0.08	0.00	1.00	0.04	0.08	0.04	0.00	0.04	0.04	0.12
LASSO	1	0.28	1.00	0.88	0.24	1.00	0.44	0.52	0.28	0.56	0.68	0.84	1.00
DZ	1	0.80	1.00	0.88	0.60	1.00	0.80	0.72	0.40	0.88	0.92	0.88	0.96
ENET	1	0.40	1.00	0.80	0.28	1.00	0.40	0.64	0.44	0.64	0.68	0.84	1.00

Table 8: Body fat dataset: relative frequencies of selections of the variables over the 25 random splits

6. the inversion base height at LAX;
7. the inversion base temperature at LAX;
8. the Daggett Pressure gradient (mm Hg) from LAX to Daggett, CA;
9. the visibility (miles) at LAX.

For this dataset, as shown by Table 9, all Bayesian approaches, as well as AIC and BIC, select four variables, except for our NIMS solution which only selects three variables, while the regularization methods opt for five. The minimal MSE is also obtained when using NIMS, but the difference between all procedures is negligible. (This lack of significant differences in the MSEs is also exhibited through the boxplots of Figure 1.)

Variables	1	2	3	4	5	6	7	8
AIC	5	3	24	25	14	2	11	4
BIC	1	0	24	25	15	0	9	1
BRIC	1	0	24	25	15	0	10	1
EB-L	1	0	24	25	15	0	9	1
EB-G	1	0	24	25	15	0	9	1
ZS-N	1	0	24	25	15	0	9	1
ZS-F	1	2	23	25	15	0	10	2
HG-3	1	0	24	25	15	0	9	1
HG-4	1	0	24	25	15	0	9	1
NIMS	1	0	24	25	15	0	8	1
LASSO	0	0	25	25	25	0	25	25
DZ	0	0	25	25	25	0	25	25
ENET	0	0	25	25	25	0	25	25

Table 9: Ozone dataset: relative frequencies of selections of the variables over the 25 random splits

5 Conclusion

We have shown in this study that an objective Bayesian solution is available for Bayesian variable selection and regularization in linear models. When compared with earlier Bayesian approaches, our approach requires no calibration whatsoever on the prior derivation. While this is not a justification *per se*, we have also shown through a series of examples that NIMS performs as well than the Bayesian alternatives, with a slight tendency to select less variables, and that it mostly overrides the performances of standard regularization tools. Those exhibited here a marked tendency to choose more variables than both Bayesian variable selection methods and standard penalized likelihood criteria such as AIC and BIC without leading to a sensitive decrease in the predictive MSE, due to the model averaging.

A limitation of this study on our objective Bayesian approach is that we do not consider large dimensions as in Bottolo and Richardson (2010), which require different computational tools to

face the enormous number of potential models. This difficulty is obviously faced by all Bayesian solutions considered in this paper and not an issue in terms of the validity of the prior modelling.


References

- Bartlett, M. (1957). A comment on D.V. Lindley's statistical paradox. *Biometrika*, 44:533–534.
- Berger, J., Pericchi, L., and Varshavsky, I. (1998). Bayes factors and marginal distributions in invariant situations. *Sankhya A*, 60:307–321.
- Bottolo, L. and Richardson, S. (2010). Evolutionary stochastic search for Bayesian model exploration. *Bayesian Analysis*, 5(3):583–618.
- Brown, J. and Vannucci, M. (1998). Multivariate Bayesian variable selection and prediction. *J. Royal Statist. Soc. Series B*, 60(3):627–641.
- Butler, R. and Wood, A. (2002). Laplace approximations for hypergeometric functions with matrix arguments. *Ann. Statist.*, 30:1155–1177.
- Candes, E. and Tao, T. (2007). The Dantzig Selector: statistical estimation when p is much larger than n . *Ann. Statist.*, 35(6):2313–2351.
- Casella, G. and Moreno, E. (2006). Objective Bayesian variable selection. *J. American Statist. Assoc.*, 101(473):157–167.
- Celeux, G., Marin, J.-M., and Robert, C. (2006). Sélection bayésienne de variables en régression linéaire. *Journal de la Société Française de Statistique*, 147(1):59–79.
- Chipman, H. (1996). Bayesian variable selection with related predictors. *Canadian Journal of Statistics*, 1:17–36.
- Cui, W. and George, E. (2008). Empirical Bayes vs. fully Bayes variable selection. *Journal of Statistical Planning and Inference*, 138:888–900.
- de Finetti, B. (1972). *Probability, Induction and Statistics*. John Wiley, New York.
- DeGroot, M. (1973). Doing what comes naturally: Interpreting a tail area as a posterior probability or as a likelihood ratio. *J. American Statist. Assoc.*, 68:966–969.
- Dupuis, J. and Robert, C. (2003). Bayesian variable selection in qualitative models by Kullback-Leibler projections. *J. Statist. Plann. Inference*, pages 77–94.
- Fernandez, C., Ley, E., and Steel, M. (2001). Benchmark priors for Bayesian model averaging. *J. Econometrics*, 100:381–427.
- Foster, D. and George, E. (1994). The risk inflation criterion for multiple regression. *Ann. Statist.*, 22:1947–1975.
- George, E. (2000). The variable selection problem. *J. American Statist. Assoc.*, 95:1304–1308.

- George, E. and Foster, D. (2000). Calibration and empirical Bayes variable selection. *Biometrika*, 87(4):731–747.
- George, E. and McCulloch, R. (1993). Variable selection via Gibbs sampling. *J. American Statist. Assoc.*, 88:881–889.
- George, E. and McCulloch, R. (1997). Approaches to Bayesian variable selection. *Statistica Sinica*, 7:339–373.
- Guo, R. and Speckman, P. (2009). Bayes factor consistency in linear models. In *The 2009 International Workshop on Objective Bayes Methodology, Philadelphia, June 5-9, 2009*. <http://www-stat.wharton.upenn.edu/statweb/Conference/OBayes09/AbstractPapers/speckman.pdf>.
- Hoerl, A. and Kennard, R. (1970). Ridge regression: biased estimation for non orthogonal problems. *Technometrics*, 12:55–67.
- Kass, R. and Raftery, A. (1995). Bayes factor and model uncertainty. *J. American Statist. Assoc.*, 90:773–795.
- Kass, R. and Wasserman, L. (1995). A reference Bayesian test for nested hypotheses and its relationship to the Schwarz criterion. *J. American Statist. Assoc.*, 90:928–934.
- Kohn, R., Smith, M., and Chan, D. (2001). Nonparametric regression using linear combinations of basis functions. *Statistics and Computing*, 11:313–322.
- Liang, F., Paulo, R., Molina, G., Clyde, M., and Berger, J. (2008). Mixtures of g-priors for Bayesian variable selection. *J. American Statist. Assoc.*, 103(481):410–423.
- Lindley, D. (1957). A statistical paradox. *Biometrika*, 44:187–192.
- Marin, J. and Robert, C. (2007). *Bayesian Core: A Practical Approach to Computational Bayesian Statistics*. Springer-Verlag, New York.
- Mitchell, T. and Beauchamp, J. (1988). Bayesian variable selection in linear regression. *J. American Statist. Assoc.*, 83:1023–1032.
- Nott, D. J. and Green, P. J. (2004). Bayesian variable selection and the Swendsen-Wang algorithm. *J. Comput. Graph. Statist.*, 13:1–17.
- Park, T. and Casella, G. (2008). The Bayesian lasso. *J. American Statist. Assoc.*, 103(473):681–686.
- Penrose, K., Nelson, A., and Fisher, A. (1985). Generalized body composition prediction equation for men using simple measurement techniques. *Medicine and Science in Sports and Exercise*, 17(2):189.
- Philips, R. and Guttman, I. (1998). A new criterion for variable selection. *Statist. Prob. Letters*, 38:11–19.
- Rao, C. (1973). *Linear Statistical Inference and its Applications*. John Wiley, New York.
- Robert, C. (1993). A note on the Jeffreys-Lindley paradox. *Statistica Sinica*, 3:601–608.

- Robert, C. (2001). *The Bayesian Choice*. Springer-Verlag, 2 edition.
- Schneider, U. and Corcoran, J. (2004). Perfect sampling for Bayesian variable selection in a linear regression model. *J. Statist. Plann. Inference*, 126:153–171.
- Smith, M. and Kohn, R. (1996). Nonparametric regression using Bayesian variable selection. *Journal of Econometrics*, 75:317–343.
- Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. *J. Roy. Statist. Soc. Ser. B*, 58(1):267–288.
- Yuan, M. and Lin, Y. (2007). Model selection and estimation in regression with grouped variables. *J. R. Stat. Soc. Ser. B Stat. Methodol.*, 68(1):49–67.
- Zellner, A. (1986). On assessing prior distributions and Bayesian regression analysis with g -prior distribution regression using Bayesian variable selection. In *Bayesian inference and decision techniques: Essays in Honor of Bruno De Finetti*, pages 233–243. North-Holland / Elsevier.
- Zellner, A. and Siow, A. (1980). Posterior odds ratios for selected regression hypotheses. In *Bayesian Statistics*, pages 585–603. Valencia: University Press. (Proceedings of the first Valencia meeting).
- Zou, H. (2006). the adaptive lasso and its oracle properties. *J. American Statist. Assoc.*, 101:1418–1429.
- Zou, H. and Hastie, T. (2005). Regularization and variable selection via the elastic net. *J. R. Stat. Soc. Ser. B Stat. Methodol.*, 67(2):301–320.

Comparison of methods for 25 random splits of Bodyfat data



Comparison of methods for 25 random splits of Ozone data




Figure 1: Body fat and Ozone datasets: variability of the median test root mean squared errors over 25 random splits for BIC, NIMS, LASSO and ENET methods.