

MAP model selection in Gaussian regression

Felix Abramovich*

Department of Statistics & Operations Research

Tel Aviv University

Tel Aviv 69978, Israel

e-mail: felix@post.tau.ac.il

and

Vadim Grinshtein*

Department of Mathematics

The Open University of Israel

Raanana 43107, Israel

e-mail: vading@openu.ac.il

Abstract: We consider a Bayesian approach to model selection in Gaussian linear regression, where the number of predictors might be much larger than the number of observations. From a frequentist view, the proposed procedure results in the penalized least squares estimation with a complexity penalty associated with a prior on the model size. We investigate the optimality properties of the resulting model selector. We establish the oracle inequality and specify conditions on the prior that imply its asymptotic minimaxity within a wide range of sparse and dense settings for “nearly-orthogonal” and “multicollinear” designs.

AMS 2000 subject classifications: Primary 62C99; secondary 62C10, 62C20, 62G05.

Keywords and phrases: Adaptivity, complexity penalty, Gaussian linear regression, maximum *a posteriori* rule, minimax estimation, model selection, oracle inequality, sparsity.

Received March 2010.

Contents

1	Introduction	933
2	MAP model selection procedure	935
3	Oracle inequality	937
4	Risk bounds for sparse settings	938
5	Asymptotic adaptive minimaxity	939
	5.1 Nearly-orthogonal design	939
	5.2 Multicollinear design	941
6	Computational aspects	943

*Supported by Israel Science Foundation (ISF), grant ISF-248/08.

7 Concluding remarks 943
 8 Appendix 944
 8.1 Proof of Theorem 1 944
 8.2 Proof of Theorem 2 945
 8.3 Proof of Theorem 3 945
 8.4 Proof of Theorem 4 946
 8.5 Proof of Theorem 5 947
 Acknowledgments 948
 References 948

1. Introduction

Consider the standard Gaussian linear regression model

$$\mathbf{y} = X\boldsymbol{\beta} + \boldsymbol{\epsilon}, \tag{1}$$

where $\mathbf{y} \in \mathbb{R}^n$ is a vector of the observed response variable Y , $X_{n \times p}$ is the design matrix of the p explanatory variables (predictors) X_1, \dots, X_p , $\boldsymbol{\beta} \in \mathbb{R}^p$ is a vector of unknown regression coefficients, $\boldsymbol{\epsilon} \sim N(\mathbf{0}, \sigma^2 I_n)$ and the noise variance σ^2 is assumed to be known.

A variety of statistical applications of regression models involves a vast number of potential explanatory variables that might be even large relatively to the amount of available data. It raises a severe “curse of dimensionality” problem. Reducing dimensionality of the model becomes therefore crucial in the analysis of such large data sets. The goal of model (or variable) selection is to select the “best”, parsimonious subset of predictors. The corresponding coefficients are then usually estimated by least squares. The meaning of the “best” subset however depends on the particular aim at hand. One should distinguish, for example, between estimation of regression coefficients $\boldsymbol{\beta}$, estimation of the mean vector $X\boldsymbol{\beta}$, model identification and predicting future observations. Different aims may lead to different optimal model selection procedures especially when the number of potential predictors p might be much larger than the sample size n . In this paper we focus on estimating the mean vector $X\boldsymbol{\beta}$ and the goodness of a model (subset of predictors) M is measured by the quadratic risk $E\|X\hat{\boldsymbol{\beta}}_M - X\boldsymbol{\beta}\|^2 = \|X\boldsymbol{\beta}_M - X\boldsymbol{\beta}\|^2 + \sigma^2|M|$, where $\hat{\boldsymbol{\beta}}_M$ is the least squares estimate of $\boldsymbol{\beta}$ and $X\boldsymbol{\beta}_M$ is the projection of $X\boldsymbol{\beta}$ on the span of M . The first (bias) term of the risk decomposition represents the approximation error of the projection, while the second (variance) term is the price for estimating the projection coefficients $\boldsymbol{\beta}_M$ by $\hat{\boldsymbol{\beta}}_M$ and is proportional to the model size. The “best” model then is the one with the minimal quadratic risk. Note that the true underlying model in (1) is not necessarily the best in this sense since sometimes it is possible to reduce its risk by excluding predictors with small (but still nonzero!) coefficients.

Such a criterion for model selection is obviously impossible to implement since it depends on the unknown $\boldsymbol{\beta}$. Instead, the corresponding ideal minimal

risk can be used as a benchmark for any available model selection procedure. The model selection criteria are typically based on the *empirical* quadratic risk $\|\mathbf{y} - X\hat{\boldsymbol{\beta}}_M\|^2$, which is essentially the least squares. However, direct minimization of the empirical risk evidently leads to a trivial (unsatisfactory!) choice of the saturated model. A typical remedy is then to add a complexity penalty $Pen(|M|)$ that increases with the model size, and to consider *penalized* least squares criterion of the form

$$\|\mathbf{y} - X\hat{\boldsymbol{\beta}}_M\|^2 + Pen(|M|) \rightarrow \min_M \quad (2)$$

The properties of the resulting estimator depends on the proper choice of the complexity penalty function $Pen(\cdot)$ in (2). There exists a plethora of works in literature on this problem. The standard, most commonly used choice is a *linear* type penalty of the form $Pen(k) = 2\sigma^2\lambda k$ for some fixed $\lambda > 0$. The most known examples motivated by different ideas include AIC for $\lambda = 1$ ([4]), BIC for $\lambda = (\ln n)/2$ ([24]) and RIC for $\lambda = \ln p$ ([14]). A series of recent works suggested the so-called $2k \ln(p/k)$ -type *nonlinear* penalties of the form $Pen(k) = 2\sigma^2ck(\ln(p/k) + \zeta_{p,k})$, where $c > 1$ and $\zeta_{p,k}$ is some “negligible” term (see, e.g., [1, 6–8, 19]).

In this paper we present a Bayesian formalism to the model selection problem in Gaussian linear regression (1) that leads to a general penalized model selection rule (2). The proposed Bayesian approach can be used, in fact, as a natural tool for obtaining a variety of penalized least squares estimators with different complexity penalties that accommodate many of the known model selection procedures as particular cases corresponding to specific choices of the prior. Within Bayesian framework, the penalty term in (2) is interpreted as proportional to the logarithm of a prior distribution. Complexity penalties $Pen(|M|)$ imply placing a prior on the model size (the number of nonzero entries of $\boldsymbol{\beta}$). Minimization of (2) corresponds to the maximum *a posteriori* (MAP) rule yielding the resulting MAP model selector to be the posterior mode.

Although there exists a large amount of literature on Bayesian model selection (see [11, 15, 16, 20] for surveys), it mainly focuses on “purely Bayesian” issues (e.g., prior specification, posterior calculations, etc.) and does not investigate the optimality of the resulting Bayesian procedures from a frequentist view. In this paper we study the optimality properties of the proposed MAP model selectors for estimating the mean vector $X\boldsymbol{\beta}$ in (1). First, under mild conditions on the prior we establish the oracle inequality and show that, up to a constant multiplier, they achieve the minimal possible risk among all estimators. We then investigate their asymptotic minimaxity. For “nearly-orthogonal” design they are proved to be simultaneously rate-optimal (in the minimax sense) over a wide range of sparse and dense settings and outperform various existing model selection procedures, e.g. AIC, BIC, RIC, Lasso ([25]) and Dantzig selector ([10]). In a way, these results extend those in [2, 3] for the normal means problem corresponding to the particular case $X = I_n$.

The analysis of “multicollinear” design, which is especially relevant for “ p much larger than n ” setup, is more delicate. We demonstrate that the lower

bounds for the minimax rates for estimating the mean vector in this case are smaller than those for “nearly-orthogonal” design by the factor depending on the design properties. Such “blessing of multicollinearity” can be explained by a possibility of exploiting correlations between predictors to reduce the size of a model (hence, to decrease the variance) without paying much extra price in the bias term. We show that under some additional assumptions on the design and the coefficients vector β in (1), the proposed Bayesian model selectors are still asymptotically rate-optimal.

The paper is organized as follows. The Bayesian model selection procedure that leads to a penalized least squares estimator (2) is introduced in Section 2. In Section 3 we derive an upper bound for the quadratic risk of the resulting MAP model selector, compare it with that of an oracle and find the conditions on the prior where, up to a constant multiplier, it achieves the minimal possible risk among all estimators. In Section 4 we obtain the upper and lower risk bounds of the MAP model selector in a sparse setup that allows us to investigate its asymptotic minimaxity for nearly-orthogonal and multicollinear designs in Section 5. The computational aspects are discussed in Section 6, and the main take-away messages of the paper are summarized in Section 7. All the proofs are given in the Appendix.

2. MAP model selection procedure

Consider the Gaussian linear regression model (1), where the number of possible predictors p might be even larger than the number of observations n . Let $r = \text{rank}(X) (\leq \min(p, n))$ and assume that any r columns of X are linearly independent. For the “standard” linear regression setup, where all p predictors are linearly independent and there are at least p linearly independent design points, $r = p$.

Any model M is uniquely defined by the $p \times p$ diagonal indicator matrix $D_M = \text{diag}(\mathbf{d}_M)$, where $d_{jM} = \mathbb{I}\{X_j \in M\}$ and, therefore, $|M| = \text{tr}(D_M)$. The corresponding least square estimate $\hat{\beta}_M = (D_M X' X D_M)^+ D_M X' \mathbf{y}$, where “+” denotes the generalized inverse matrix.

Assume some prior on the model size $\pi(k) = P(|M| = k)$, where $\pi(k) > 0$, $k = 0, \dots, r$ ($k = 0$ corresponds to a null model with a single intercept) and $\pi(k) = 0$ for $k > r$ since otherwise, there necessarily exists another vector β^* with at most r nonzero entries also satisfying (1), that is, $X\beta = X\beta^*$.

For any $k = 0, \dots, r - 1$ there are $\binom{p}{k}$ different models of a given size k . Assume all of them to be equally likely, that is, conditionally on $|M| = k$,

$$P(M \mid |M| = k) = \binom{p}{k}^{-1}$$

One should be a little bit more careful for $k = r = \text{rank}(X)$. Although there are $\binom{p}{r}$ different sets of predictors of size r , all of them evidently result in the same estimator for the mean vector and, in this sense, are essentially undistinguishable

and associated with a *single* (saturated) model. Hence, in this case, we set

$$P(M \mid |M| = r) = 1 \tag{3}$$

Finally, assume the normal prior on the unknown vector of k coefficients of the model M : $\beta_M \sim N_p(\mathbf{0}, \gamma\sigma^2(D_M X' X D_M)^+)$. This is a well-known conventional g -prior of Zellner [28].

For the proposed hierarchical prior, a straightforward calculus yields the posterior probability of a model M of size $|M| = 0, \dots, r - 1$:

$$P(M|\mathbf{y}) \propto \pi(|M|) \binom{p}{|M|}^{-1} (1 + \gamma)^{-\frac{|M|}{2}} \times \exp \left\{ \frac{\gamma}{\gamma + 1} \frac{\mathbf{y}' X D_M (D_M X' X D_M)^+ D_M X' \mathbf{y}}{2\sigma^2} \right\} \tag{4}$$

Finding the most likely model leads therefore to the following maximum *a posteriori* (MAP) model selection criterion:

$$\mathbf{y}' X D_M (D_M X' X D_M)^+ D_M X' \mathbf{y} + 2\sigma^2(1 + 1/\gamma) \ln \left\{ \binom{p}{|M|}^{-1} \pi(|M|)(1 + \gamma)^{-\frac{|M|}{2}} \right\} \rightarrow \max_M$$

or, equivalently,

$$\|\mathbf{y} - X \hat{\beta}_M\|^2 + 2\sigma^2(1 + 1/\gamma) \ln \left\{ \binom{p}{|M|} \pi(|M|)^{-1} (1 + \gamma)^{\frac{|M|}{2}} \right\} \rightarrow \min_M, \tag{5}$$

which is of the general type (2) with the complexity penalty

$$Pen(k) = 2\sigma^2(1 + 1/\gamma) \ln \left\{ \binom{p}{k} \pi(k)^{-1} (1 + \gamma)^{\frac{k}{2}} \right\}, \quad k = 0, \dots, r - 1 \tag{6}$$

Similarly, for $|M| = r$ from (3) one has

$$Pen(r) = 2\sigma^2(1 + 1/\gamma) \ln \left\{ \pi(r)^{-1} (1 + \gamma)^{\frac{r}{2}} \right\} \tag{7}$$

A specific form of the penalty (6)–(7) depends on the choice of a prior $\pi(\cdot)$. In particular, the (truncated if $p > n$) binomial prior $B(p, \xi)$ corresponds to the prior assumption that the indicators d_{jM} are independent. The binomial prior yields the linear penalty $Pen(k) = 2\sigma^2 \lambda k$, where $\lambda = (1 + 1/\gamma) \ln\{\sqrt{1 + \gamma}(1 - \xi)/\xi\} \sim \ln\{\sqrt{\gamma}(1 - \xi)/\xi\}$ for sufficiently large variance ratio γ . The AIC criterion corresponds then to $\xi \sim \sqrt{\gamma}/(e + \sqrt{\gamma})$, while $\xi \sim \sqrt{\gamma}/(p + \sqrt{\gamma})$ leads to the RIC criterion. These relations indicate that RIC should be appropriate for sparse cases, where the size of the true (unknown) model is believed to be much less than the number of possible predictors, while AIC is suitable for dense cases, where they are of the same order. In fact, any binomial prior or, equivalently, any linear penalty cannot “kill two birds with one stone”. On the other hand,

the (truncated) geometric prior $\pi(k) \propto q^k$, $k = 1, \dots, r$ for some $0 < q < 1$, implies $Pen(k) \sim 2\sigma^2(1 + 1/\gamma)k(\ln(p/k) + c(\gamma, q))$ which is of the $2k \ln(p/k)$ -type introduced above. For large γ it behaves similar to RIC for $k \ll p$ and to AIC for $k \sim p$ and is, therefore, adaptive to both sparse and dense cases. We will discuss these issues more rigorously in Section 5 below.

3. Oracle inequality

In this section we derive an upper bound for the quadratic risk of the proposed MAP model selector and compare it with the ideal minimal quadratic risk often called in literature as an oracle risk.

Assumption (P). Assume that

$$\pi(k) \leq \binom{p}{k} e^{-c(\gamma)k}, \quad k = 0, \dots, r - 1, \quad \text{and} \quad \pi(r) \leq e^{-c(\gamma)r}$$

where $c(\gamma) = 8(\gamma + 3/4)^2 > 9/2$.

Assumption (P) is not restrictive. Indeed, the obvious inequality $\binom{p}{k} \geq (p/k)^k$ implies that for $k < r$ it automatically holds for any prior $\pi(k)$ for all $k \leq pe^{-c(\gamma)}$. Assumption (P) is used to establish an upper bound for the quadratic risk of the MAP model selector.

Theorem 1. Let the model \hat{M} be the solution of (2) with the complexity penalty $Pen(\cdot)$ given in (6)–(7) and $\hat{\beta}_{\hat{M}}$ be the corresponding least squares estimate. Then, under Assumption (P)

$$E\|X\hat{\beta}_{\hat{M}} - X\beta\|^2 \leq c_0(\gamma) \inf_M \{ \|X\beta_M - X\beta\|^2 + Pen(|M|) \} + c_1(\gamma)\sigma^2 \quad (8)$$

for some $c_0(\gamma)$ and $c_1(\gamma)$ depending only on γ .

To assess the quality of the upper bound in (8), we compare it with the oracle risk $\inf_M E\|X\hat{\beta}_M - X\beta\|^2$. Note that the oracle risk is exactly zero when $\beta \equiv \mathbf{0}$ and, evidently, no estimator can achieve it in this case. Hence, an additional, typically negligible term σ^2 , which is, essentially, an error of estimating a single extra parameter, is usually added to the oracle risk for a proper comparison. It is known that no estimator can attain a risk smaller than within $2 \ln p$ factor from that of an oracle (e.g., [9, 13, 14]). The following theorem shows that under certain additional conditions on the prior $\pi(\cdot)$, the resulting MAP model selector achieves this minimal possible risk among all estimators up to a constant multiplier depending on γ :

Theorem 2 (oracle inequality). Let $\pi(k)$ satisfy Assumption (P) and, in addition, $\pi(0) \geq p^{-c}$, $\pi(k) \geq p^{-ck}$, $k = 1, \dots, r$ for some constant $c > 0$. Then, the resulting MAP model selector satisfies

$$E\|X\hat{\beta}_{\hat{M}} - X\beta\|^2 \leq c_2(\gamma) \ln p \left(\inf_M E\|X\hat{\beta}_M - X\beta\|^2 + \sigma^2 \right)$$

for some $c_2(\gamma) \geq 1$.

In particular, it can be easily shown that Theorem 2 holds for the (truncated) binomial $B(p, \xi)$ with $\xi \sim 1/p$ (RIC criterion) and geometric priors (see Section 2). More generally, all priors such that $\ln \pi(k) = O(k \ln(k/p))$ corresponding to the $2k \ln(p/k)$ -type penalties satisfy the conditions of Theorem 2.

4. Risk bounds for sparse settings

In the previous section we considered the global behavior of the MAP estimator without any restrictions on the model size. However, in the analysis of large data sets, it is typically reasonable to assume that the true model in (1) is sparse in the sense that only part of coefficients in β are different from zero. We now show that under such extra sparsity assumption, more can be said on the optimality of the MAP model selection.

For a given $1 \leq p_0 \leq r$, define the sets of models \mathcal{M}_{p_0} that have at most p_0 predictors, that is, $\mathcal{M}_{p_0} = \{M : |M| \leq p_0\}$. Obviously, if a true model in (1) belongs to \mathcal{M}_{p_0} , the l_0 quasi-norm of the corresponding coefficients vector $\|\beta\|_0 \leq p_0$, where $\|\beta\|_0$ is the number of its nonzero entries. In this section we find the upper and lower bounds for the maximal risk of the proposed MAP model selector over \mathcal{M}_{p_0} .

Theorem 3. *Let $1 \leq p_0 \leq r$, $\pi(\cdot)$ satisfy Assumption (P) and, in addition, $\pi(p_0) \geq (p_0/(pe))^{cp_0}$ if $p_0 < r$ or $\pi(r) \geq e^{-cr}$ if $p_0 = r$ for some constant $c > c(\gamma)$. Then, there exists a constant $C_1(\gamma) > 0$ depending only on γ such that*

$$\sup_{\beta: \|\beta\|_0 \leq p_0} E\|X\hat{\beta}_{\hat{M}} - X\beta\|^2 \leq C_1(\gamma)\sigma^2 \min(p_0(\ln(p/p_0) + 1), r) \quad (9)$$

The general upper bound (9) for the maximal risk of the MAP selector over \mathcal{M}_{p_0} in Theorem 3 holds for any design matrix X . To assess its accuracy we establish the lower bound for the minimax risk of estimating the mean vector $X\beta$ in (1).

For any given $k = 1, \dots, r$, let $\phi_{min}[k]$ and $\phi_{max}[k]$ be the k -sparse minimal and maximal eigenvalues of the design defined as

$$\phi_{min}[k] = \min_{\beta: 1 \leq \|\beta\|_0 \leq k} \frac{\|X\beta\|^2}{\|\beta\|^2},$$

$$\phi_{max}[k] = \max_{\beta: 1 \leq \|\beta\|_0 \leq k} \frac{\|X\beta\|^2}{\|\beta\|^2}$$

(see [5, 21]). In fact, $\phi_{min}[k]$ and $\phi_{max}[k]$ are respectively the minimal and maximal eigenvalues of all $k \times k$ submatrices of the matrix $X'X$ generated by any k columns of X . Let $\tau[k] = \phi_{min}[k]/\phi_{max}[k]$, $k = 1, \dots, r$ and set $\tau[k] = \tau[r]$ for all $k > r$. By the definition, $\tau[k]$ is a non-increasing function of k . Obviously, $\tau[k] \leq 1$ and for the orthogonal design the equality holds for all k .

Theorem 4. Consider the model (1) and let $1 \leq p_0 \leq r$. There exists a universal constant $C_2 > 0$ such that

$$\inf_{\hat{\mathbf{y}}} \sup_{\boldsymbol{\beta}: \|\boldsymbol{\beta}\|_0 \leq p_0} E \|\hat{\mathbf{y}} - X\boldsymbol{\beta}\|^2 \geq \begin{cases} C_2 \sigma^2 \tau[2p_0] p_0 (\ln(p/p_0) + 1), & 1 \leq p_0 \leq r/2 \\ C_2 \sigma^2 \tau[p_0] r, & r/2 \leq p_0 \leq r \end{cases} \tag{10}$$

where the infimum is taken over all estimates $\hat{\mathbf{y}}$ of the mean vector $X\boldsymbol{\beta}$.

Theorem 4 shows that the minimax lower bound (10) depends on a specific design matrix X only through the sparse eigenvalues ratios. A computationally simpler but less accurate minimax lower bound can be obtained by replacing $\tau[2p_0]$ and $\tau[p_0]$ in (10) by $\tau[r]$, that for the case $r = p \leq n$ is just the ratio of the minimal and maximal eigenvalues of $X'X$.

For the orthogonal design, where $\tau[\cdot] \equiv 1$, and $p \leq n$ analogous results were obtained in [6]. For a general design and $p_0 \leq r/2$ similar minimax lower bounds were independently obtained in [22] for a design matrix of a full rank and in [23] for a general case within a related aggregation context.

The established upper and lower bounds (9), (10) for the risk of the MAP model selector allow us in the following section to investigate its asymptotic minimaxity as both n and p increase.

5. Asymptotic adaptive minimaxity

5.1. Nearly-orthogonal design

In this section we consider the asymptotic properties of the MAP model selector as the sample size n increases. We allow $p = p_n$ to increase with n as well and look for a projection of the unknown mean vector on an expanding span of predictors. In particular, the most challenging cases intensively studied nowadays in literature are those, where $p > n$ or even $p \gg n$. In such asymptotic settings one should essentially consider a sequence of design matrices X_{n,p_n} , where $r_n \rightarrow \infty$. For simplicity of exposition, in what follows we omit the index n and denote X_{n,p_n} by X_p emphasizing the dependence on the number of predictors p and let r tend to infinity. Similarly, we consider now sequences of coefficients vectors $\boldsymbol{\beta}_p$ and priors $\pi_p(\cdot)$. In these notations, the original model (1) is transformed into a sequence of models

$$\mathbf{y} = X_p \boldsymbol{\beta}_p + \boldsymbol{\epsilon}, \tag{11}$$

where $\text{rank}(X_p) = r$ and any r columns of X_p are linearly independent (hence, $\tau_p[r] > 0$), $\boldsymbol{\epsilon} \sim N(\mathbf{0}, \sigma^2 I_n)$ and the noise variance σ^2 does not depend on n and p . One can also view a sequence of models (11) in a triangular array setup ([17]).

Definition 1. Consider the sequence of design matrices X_p . The design is called nearly-orthogonal if the corresponding sequence of sparse eigenvalues ratios $\tau_p[r]$ is bounded away from zero by some constant $c > 0$. Otherwise, the design is called multicollinear.

Nearly-orthogonality condition essentially means that there is no multicollinearity in the design in the sense that there are no “too strong” linear relationships within any set of r columns of X_p . Intuitively, it is clear that in this case p cannot be “too large” relative to r and, therefore, to n . Indeed, apply the upper and lower bounds (9), (10) for $p_0 = r/2$ to get $(C_2/2)\sigma^2\tau[r]r(\ln(2p/r) + 1) \leq C_1(\gamma)\sigma^2r$ that implies the following remark:

Remark 1. For nearly-orthogonal design, necessarily $p = O(r)$ and, therefore, $p = O(n)$.

The following corollary is an immediate consequence of Theorems 3 and 4:

Corollary 1. Let the design be nearly-orthogonal.

1. As r increases, the asymptotic minimax risk of estimating the mean vector $X_p\beta_p$ over \mathcal{M}_{p_0} is of the order $\min(p_0(\ln(p/p_0) + 1), r)$ or essentially $p_0(\ln(p/p_0) + 1)$ (see Remark 1), that is, there exist two constants $0 < C_1 \leq C_2 < \infty$ such that for all sufficiently large r ,

$$C_1\sigma^2p_0(\ln(p/p_0)+1) \leq \inf_{\hat{y}} \sup_{\beta_p: \|\beta_p\|_0 \leq p_0} E\|\hat{y} - X_p\beta_p\|^2 \leq C_2\sigma^2p_0(\ln(p/p_0)+1)$$

for all $1 \leq p_0 \leq r$.

2. Assume Assumption (P) and, in addition, that $\pi_p(k) \geq (k/(pe))^{c_1k}$, $k = 1, \dots, r - 1$ and $\pi_p(r) \geq e^{-c_2r}$ for some constants $c_1, c_2 > c(\gamma)$. Then, the corresponding MAP model selector attains the minimax convergence rates simultaneously over all \mathcal{M}_{p_0} , $1 \leq p_0 \leq r$.

One can easily verify that the conditions on the prior of Corollary 1 are satisfied, for example, for the truncated geometric prior (see Section 2) for all $k = 1, \dots, r$. The resulting MAP model selector attains, therefore, the minimax rates simultaneously for all \mathcal{M}_{p_0} , $p_0 = 1, \dots, r$. As we have mentioned, the corresponding penalty in (6) is of the $2k \ln(p/k)$ -type. On the other hand, no truncated binomial prior $B(p, \xi_p)$ can satisfy these conditions on the entire range $k = 1, \dots, r$. It is easy to verify that they hold for small ξ_p if $k \ll r$ but for large ξ_p if $k \sim r$. In fact, these arguments go along the lines with the similar results of [6, 7, 14]. Recall that binomial prior corresponds to linear penalties of the type $Pen(k) = 2\sigma^2\lambda k$ in (2) (see Section 2). Foster & George ([14]) and Birgé & Massart ([6], Section 5.2) showed that the best possible risk of such estimators over \mathcal{M}_{p_0} is only of order $\sigma^2p_0 \ln p$ achieved for $\lambda \sim \ln p$ corresponding to the RIC criterion. It is of the same order as the optimal risk $\sigma^2p_0(\ln(p/p_0) + 1)$ for $p_0 \ll p$ (sparse case) but larger for dense case ($p_0 \sim p$). On the other hand, the risk of the AIC estimator ($\lambda = 1$) is of the order σ^2p , which is optimal for dense but much larger for sparse case.

Furthermore, under somewhat stronger nearly-orthogonality conditions, Bickel, Ritov & Tsybakov ([5]) showed that the well-known Lasso ([25]) and Dantzig ([10]) estimators achieve only the same sub-optimal rate $\sigma^2p_0 \ln p$ as RIC. These results are, in fact, not so surprising since both Lasso and Dantzig estimators are essentially based on convex relaxations of the l_0 -norm of regression coefficients

$\|\beta\|_0$ in the linear complexity penalty $2\sigma^2\lambda\|\beta\|_0$ in order to replace the original combinatorial problem (2) by a convex program. Thus, Lasso approximates the l_0 -norm $\|\beta\|_0$ by the corresponding l_1 -norm $\|\beta\|_1$. In particular, for the orthogonal design, linear complexity penalties and Lasso yield respectively hard and soft thresholding of components of β with a *fixed* threshold. RIC estimator and Lasso with the optimally chosen tuning parameter (e.g., [5]) result in this case in the well-known hard and soft universal thresholding of Donoho & Johnstone ([12]) with a fixed threshold $\sigma\sqrt{2\ln p}$ which is rate-optimal for various sparse but not dense settings. On the other hand, the nonlinear MAP penalty corresponds to hard thresholding with a *data-driven* threshold that under conditions on $\pi_p(\cdot)$ in Corollary 1 is simultaneously minimax for both sparse and dense cases ([2, 3]).

Finally, note that for the nearly-orthogonal design, $\|X_p\hat{\beta}_{p\hat{M}} - X_p\beta_p\| \asymp \|\hat{\beta}_{p\hat{M}} - \beta_p\|$, where “ \asymp ” means that their ratio is bounded from below and above. Therefore, all the results of Corollary 1 for estimating the mean vector $X_p\beta_p$ in (11) can be straightforwardly applied for estimating the regression coefficients β_p . This equivalence, however, does not hold for the multicollinear design considered below.

5.2. Multicollinear design

Nearly-orthogonality assumption may be reasonable in the “classical” setup, where p is not too large relatively to n but might be questionable for the analysis of high-dimensional data, where $p \gg n$, due to the multicollinearity phenomenon (see also Remark 1). When this assumption does not hold, the sparse eigenvalues ratios in (10) may tend to zero as p increases and, thus, decrease the minimax lower bound rate relatively to the nearly-orthogonal design. In this case there is a gap between the rates in the lower and upper bounds (10) and (9). Intuitively, one can think of exploiting correlations between predictors to reduce the size of a model (hence, to decrease the variance) without paying much extra price in the bias term, and, therefore, to reduce the risk. We show that under certain additional assumptions on the design and the coefficients vector in (11), the upper risk bound (9) can be indeed reduced to the minimax lower bound rate (10).

For simplicity of exposition we consider the case $p_0 \leq r/2$ although the corresponding conditions for the case $r/2 \leq p_0 \leq r$ can be obtained in a similar way with necessary changes.

We introduce now several definitions that will be used in the sequel (including the proofs in the Appendix). For a given index set \mathcal{J} and $l \geq |\mathcal{J}|$ define a $l \times |\mathcal{J}|$ matrix $G_{l,\mathcal{J}}$ which columns \mathbf{e}_j , $j \in \mathcal{J}$ are the elements of the standard basis in \mathbb{R}^l . Thus, for any matrix A with l columns, $AG_{l,\mathcal{J}}$ selects the columns of A indexed by \mathcal{J} . Similarly, for any $l \times l$ symmetric matrix A , $G'_{l,\mathcal{J}}AG_{l,\mathcal{J}}$ generates a (symmetric) $|\mathcal{J}| \times |\mathcal{J}|$ submatrix of A of the corresponding columns and rows.

For all $k = 1, \dots, r/2$, define $k' = \lceil \tau_p[2k] \cdot k \rceil \geq 1$. Let \mathcal{J}_M be an index set of predictors included in a model M of size $k = |\mathcal{J}_M|$. For any submodel $M' \subset M$

of size $k' < k$ let $\tilde{\phi}_{M,M'}$ be the minimal eigenvalue of the $(k-k') \times (k-k')$ matrix $\Lambda_p(M, M') = G'_{k, \mathcal{J}_{M/M'}} (G'_{p, \mathcal{J}_M} X'_p X_p G_{p, \mathcal{J}_M})^{-1} G_{k, \mathcal{J}_{M/M'}}$. In fact, $\sigma^2 \Lambda_p(M, M')$ is the covariance matrix of the components of the least squares estimate vector $\hat{\beta}_M$ corresponding to a subset of predictors in M/M' .

Finally, define

$$\tilde{\phi}_p[k] = \min_{M:|M|=k} \max_{M' \subset M:|M'|=k'} \tilde{\phi}_{M,M'}$$

As we show later (see the proof of Theorem 5 in the Appendix), $\tilde{\phi}_p^{-1}[k]$ measures an error of approximating mean vectors $X_p \beta_p$, where $\|\beta_p\|_0 = k$, by their projections on lower dimensional subspaces of predictors. The stronger is multicollinearity, the better is the approximation and the larger is $\tilde{\phi}_p[k]$.

Theorem 5. *Let $\tau_p[r] \rightarrow 0$ as $r \rightarrow \infty$ (multicollinear design). Assume the following additional assumptions on the design matrix X_p and the (unknown) vector of coefficients β_p in (11):*

(D) *for all p there exist $1 \leq \kappa_{p1} \leq \kappa_{p2} \leq r/2$ such that*

1. $\tilde{c}_1 \leq \tau_p[2k] \cdot k \leq k - 1, k = \kappa_{p1}, \dots, \kappa_{p2}$
2. $\tau_p[2\kappa_{p2}] \geq (\kappa_{p2}/(pe))^{\tilde{c}_2}$
3. $\phi_{p,min}[2k] \cdot \tilde{\phi}_p[k] \geq \tilde{c}_3, k = \kappa_{p1}, \dots, \kappa_{p2}$

(B) $\|\beta_p\|_\infty^2 \leq \tilde{c}_4 \tau_p[2p_0] \cdot \tilde{\phi}_p[p_0] \cdot (\ln(p/p_0) + 1)$, where $p_0 = \|\beta_p\|_0$

for some positive constants $\tilde{c}_1, \tilde{c}_2, \tilde{c}_3$ and \tilde{c}_4 .

Then, under the above additional restrictions, if the prior $\pi_p(\cdot)$ satisfies Assumption (P) and for all $k = \kappa_{p1}, \dots, \kappa_{p2}$, $\pi_p(k') \geq (k'/(pe))^{ck'}$ for some positive $c > c(\gamma)$, where $k' = \lceil \tau_p[2k] \cdot k \rceil$, the corresponding MAP model selector is asymptotically simultaneously minimax (up to a constant multiplier) over all $\mathcal{M}_{p_0}, \kappa_{p1} \leq p_0 \leq \kappa_{p2}$.

Note that by simple algebra one can verify that $\phi_{p,min}[2k] \cdot \tilde{\phi}_p[k] \leq 1$ and, therefore, the constant \tilde{c}_3 in Assumption (D.3) is not larger than one.

We have argued that multicollinearity typically arises when $p \gg n$. One can easily verify that for $n = O(p^\alpha), 0 \leq \alpha < 1$, Assumption (D.2) always follows from Assumption (D.1) and, therefore, can be omitted in this case.

As we show in the proof, Assumptions (D.1, D.2) and Assumption (B) allow one to reduce the upper bound (9) for the risk of the MAP model selector by the factor $\tau_p[2p_0]$, while Assumption (D.3) is required to guarantee that the additional constraint on β_p in Assumption (B) does not affect the lower bound (10).

To obtain asymptotic minimaxity of the MAP selector within the entire range $1 \leq p_0 \leq r/2$ similar to Corollary 1 for the nearly-orthogonal case, Assumptions (D) on the design matrix are required to be satisfied for all $k = 1, \dots, r/2$ that might be quite restrictive. However, the results of Theorem 5 are more general and show the tradeoff between relaxation of Assumptions (D) to a smaller range of k and the corresponding constriction of the adaptivity range for p_0 .

6. Computational aspects

In practice, minimizing (2) (and (5) in particular) requires generally an NP-hard combinatorial search over all possible models. During the last decade there have been substantial efforts to develop various *approximated* algorithms for solving (2) that are computationally feasible for high-dimensional data (see, e.g. [26] for a survey and references therein). The common remedies involve either greedy algorithms (e.g., forward selection, matching pursuit) approximating the global solution by a stepwise sequence of local ones, or convex relaxation methods replacing the original combinatorial problem by a related convex program (e.g., Lasso and Dantzig selector for linear penalties). The proposed Bayesian formalism allows one instead to use a stochastic search variable selection (SSVS) techniques originated in [15, 16] for solving (5) by generating a sequence of models from the posterior distribution $P(M|\mathbf{y})$ in (4). The key point is that the relevant models with the highest posterior probabilities will appear most frequently and can be identified even for a generated sample of a relatively small size avoiding computations of the entire posterior distribution.

The SSVS algorithm for the problem at hand can be basically described as follows. As we have mentioned in Section 2, every model M is uniquely defined by the corresponding indicator vector \mathbf{d}_M and the joint posterior distribution of \mathbf{d}_M is given by (4) (up to a normalizing constant). SSVS uses the Gibbs sampler to generate a sequence of indicator vectors $\mathbf{d}_1, \dots, \mathbf{d}_m$ *componentwise* by sampling consecutively from the conditional distributions $d_j | (\mathbf{d}_{(-j)}, \mathbf{y})$, $j = 1, \dots, p$, where $\mathbf{d}_{(-j)} = (d_1, \dots, d_{j-1}, d_{j+1}, \dots, d_p)'$. The components d_j can be trivially obtained as simulations of Bernoulli draws, where from (4) the corresponding posterior odds ratio

$$\begin{aligned} \frac{P(d_j = 1 | \mathbf{d}_{(-j)}, \mathbf{y})}{P(d_j = 0 | \mathbf{d}_{(-j)}, \mathbf{y})} &= \frac{P(d_j = 1, \mathbf{d}_{(-j)} | \mathbf{y})}{P(d_j = 0, \mathbf{d}_{(-j)} | \mathbf{y})} \\ &= \frac{\pi(|d_{(-j)}| + 1)}{\pi(|d_{(-j)}|)} \frac{|d_{(-j)}| + 1}{p - |d_{(-j)}|} (1 + \gamma)^{-\frac{1}{2}} \exp \left\{ \frac{\gamma}{\gamma + 1} \frac{\Delta RSS_j}{2\sigma^2} \right\} \end{aligned}$$

and $\Delta RSS_j > 0$ is the increment in the residual sum of squares (RSS) after dropping the j -th predictor from the model $(d_j = 1, \mathbf{d}_{(-j)})'$. The resulting Gibbs sampler is computationally efficient and, as m increases, the empirical distribution of the generated sample converges to the actual posterior distribution of $\mathbf{d} | \mathbf{y}$. After the sequence has reached approximate stationarity, one can identify the most frequently appeared vector(s) \mathbf{d} as potential candidate(s) to solve (5).

7. Concluding remarks

In this paper we considered a Bayesian approach to model selection in Gaussian linear regression. From a frequentist view, the resulting MAP model selector is a penalized least squares estimator with a complexity penalty associated with

a prior $\pi(\cdot)$ on the model size. Although the proposed estimator was originated within Bayesian framework, the latter was used as a natural tool to obtain a wide class of penalized least squares estimators with various complexity penalties. Thus, we believe that the main take-away messages of the paper summarized below are of a more general interest.

The first main take-away message is that neither linear complexity penalties (e.g., AIC, BIC and RIC) corresponding to binomial priors $\pi(\cdot)$, nor closely related Lasso and Dantzig estimators can be simultaneously minimax for both sparse and dense cases. We specify the class of priors and associated nonlinear penalties that do yield such a wide adaptivity range. In particular, it includes $2k \ln(p/k)$ -type penalties.

Another important take-away message is about the effect of multicollinearity of design. Unlike model identification or coefficients estimation, where multicollinearity is a “curse”, it may become a “blessing” for estimating the mean vector allowing one to exploit correlations between predictors to reduce the size of a model (hence, to decrease the variance) without paying much extra price in the bias term. Interestingly, a similar phenomenon also occurs in a testing setup (e.g., [18]).

8. Appendix

Throughout the proofs we use C to denote a generic positive constant, not necessarily the same each time it is used, even within a single equation.

8.1. Proof of Theorem 1

Define

$$L_k = (1/k) \ln \left(\binom{p}{k} \pi^{-1}(k) \right) \geq c(\gamma), \quad k = 1, \dots, r-1 \quad (12)$$

and

$$L_r = (1/r) \ln \pi^{-1}(r) \geq c(\gamma) \quad (13)$$

In terms of L_k the complexity penalty (6)–(7) is $Pen(k) = \sigma^2(1 + 1/\gamma)k(2L_k + \ln(1 + \gamma))$. Following the arguments of the proof of Theorem 1 of [2], under the Assumption (P) one has

$$\sum_{k=1}^{r-1} \binom{p}{k} e^{-kL_k} + e^{-rL_r} = \sum_{k=1}^r \pi(k) = 1 - \pi(0) < 1$$

and

$$(1 + 1/\gamma)(2L_k + \ln(1 + \gamma)) \geq C(1 + \sqrt{2L_k})^2, \quad k = 1, \dots, r$$

The proof of Theorem 1 then follows directly from Theorem 2 of [6]. \square

8.2. Proof of Theorem 2

Let $L^* = \max_{0 \leq k \leq r} L_k$, where L_k , $k = 1, \dots, r$ were defined in (12)–(13) and $L_0 = 2 \ln \pi^{-1}(0)$. Simple calculus shows that the conditions on $\pi(\cdot)$ in Theorem 2 imply $L^* = O(\ln p)$.

Consider first the case $k \geq 1$. From Theorem 1 we have

$$\begin{aligned} E\|X\hat{\beta}_{\hat{M}} - X\beta\|^2 &\leq c_0(\gamma) \inf_M \{ \|X\beta_M - X\beta\|^2 + \sigma^2(1 + 1/\gamma)|M|(2L_{|M|} + \ln(1 + \gamma)) \} + c_1(\gamma)\sigma^2 \\ &\leq c_0(\gamma)(1 + 1/\gamma)(2L^* + \ln(1 + \gamma)) \inf_M \{ \|X\beta_M - X\beta\|^2 + |M|\sigma^2 \} + c_1(\gamma)\sigma^2 \\ &\leq c_2(\gamma)(2L^* + \ln(1 + \gamma)) \left\{ \inf_M E\|X\hat{\beta}_M - X\beta\|^2 + \sigma^2 \right\} \end{aligned} \tag{14}$$

For the degenerative case $k = 0$ ($M = \{0\}$), Theorem 1 implies

$$\begin{aligned} E\|X\hat{\beta}_{\hat{M}} - X\beta\|^2 &\leq c_0(\gamma) (\|X\beta\|^2 + \sigma^2(1 + 1/\gamma)L_0) + c_1(\gamma)\sigma^2 \\ &\leq c_2(\gamma)L^*(\|X\beta\|^2 + \sigma^2) \end{aligned} \tag{15}$$

Combining (14) and (15) completes the proof. □

8.3. Proof of Theorem 3

For all $p_0 \leq r$, Theorem 1 and (7) under the assumption $\pi(r) \geq e^{-cr}$ imply

$$\begin{aligned} \sup_{\beta: \|\beta\|_0 \leq p_0} E\|X\hat{\beta}_{\hat{M}} - X\beta\|^2 &\leq \sup_{\beta: \|\beta\|_0 \leq r} E\|X\hat{\beta}_{\hat{M}} - X\beta\|^2 \leq Pen(r) + c_1(\gamma)\sigma^2 \\ &\leq C_1(\gamma)\sigma^2 r \end{aligned} \tag{16}$$

On the other hand, applying the general upper bound for the risk of MAP model selector established in Theorem 1 for models of size $p_0 < r$ we have

$$\begin{aligned} \sup_{\beta: \|\beta\|_0 \leq p_0} \|X\hat{\beta}_{\hat{M}} - X\beta\|^2 &\leq c_0(\gamma)2\sigma^2(1 + 1/\gamma) \left(\ln \left\{ \binom{p}{p_0} \pi^{-1}(p_0) \right\} + \frac{p_0}{2} \ln(1 + \gamma) \right) + c_1(\gamma)\sigma^2 \end{aligned} \tag{17}$$

Abramovich *et al.* ([3], Lemma 1) showed that $\binom{p}{p_0} \leq (pe/p_0)^{p_0}$. Hence, under the conditions on $\pi(p_0)$ in Theorem 3, for $p_0 = 1, \dots, r - 1$, (17) yields

$$\begin{aligned} \sup_{\beta: \|\beta\|_0 \leq p_0} E\|X\hat{\beta}_{\hat{M}} - X\beta\|^2 &\leq c_0(\gamma)2\sigma^2(1 + 1/\gamma) \left\{ (c + 1)p_0 \ln(pe/p_0) + \frac{p_0}{2} \ln(1 + \gamma) \right\} + c_1(\gamma)\sigma^2 \\ &\leq C_1(\gamma)\sigma^2 p_0 (\ln(p/p_0) + 1) \end{aligned}$$

Finally, note that for $p_0 = r$, as we have already established in (16),

$$\sup_{\beta: \|\beta\|_0 \leq r} E \|X \hat{\beta}_M - X \beta\|^2 \leq C_1(\gamma) \sigma^2 r \leq C_1(\gamma) \sigma^2 r (\ln(p/r) + 1)$$

□

8.4. Proof of Theorem 4

The core of the proof is to find a subset \mathcal{B}_{p_0} of vectors β , where $\|\beta\|_0 \leq p_0$, and the corresponding subset of mean vectors $\mathcal{G}_{p_0} = \{\mathbf{g} \in \mathbb{R}^n : \mathbf{g} = X\beta, \beta \in \mathcal{B}_{p_0}\}$ such that for any $\mathbf{g}_1, \mathbf{g}_2 \in \mathcal{G}_{p_0}$, $\|\mathbf{g}_1 - \mathbf{g}_2\|^2 \geq 4s^2(p_0)$ and the Kullback-Leibler divergence $K(\mathbb{P}_{\mathbf{g}_1}, \mathbb{P}_{\mathbf{g}_2}) = \frac{\|\mathbf{g}_1 - \mathbf{g}_2\|^2}{2\sigma^2} \leq (1/16) \ln \text{card}(\mathcal{G}_{p_0})$. Lemma A.1 of [8] will imply then that $s^2(p_0)$ is the minimax lower bound over \mathcal{M}_{p_0} .

To construct the desired subsets \mathcal{B}_{p_0} and \mathcal{G}_{p_0} we consider three possible cases.

CASE 1. $p_0 \leq r/2$

Define the subset $\tilde{\mathcal{B}}_{p_0}$ of all vectors $\beta \in \mathbb{R}^p$ that have p_0 entries equal to C_{p_0} defined later, while the remaining entries are zeros: $\tilde{\mathcal{B}}_{p_0} = \{\beta : \beta \in \{0, C_{p_0}\}^p, \|\beta\|_0 = p_0\}$. For $p_0 \leq r/2$ from Lemma 8.3 of [23], there exists a subset $\mathcal{B}_{p_0} \subset \tilde{\mathcal{B}}_{p_0}$ such that for some constant $\tilde{c} > 0$, $\ln \text{card}(\mathcal{B}_{p_0}) \geq \tilde{c} p_0 (\ln(p/p_0) + 1)$, and for any pair $\beta_1, \beta_2 \in \mathcal{B}_{p_0}$, the Hamming distance $\rho(\beta_1, \beta_2) = \sum_{j=1}^{p_0} \mathbb{I}\{\beta_{1j} \neq \beta_{2j}\} \geq \tilde{c} p_0$.

Consider the corresponding subset of mean vectors \mathcal{G}_{p_0} , where $\text{card}(\mathcal{G}_{p_0}) = \text{card}(\mathcal{B}_{p_0})$. For any $\mathbf{g}_1, \mathbf{g}_2 \in \mathcal{G}_{p_0}$ and the corresponding $\beta_1, \beta_2 \in \mathcal{B}_{p_0}$ we then have

$$\|\mathbf{g}_1 - \mathbf{g}_2\|^2 = \|X(\beta_1 - \beta_2)\|^2 \geq \phi_{\min}[2p_0] \|\beta_1 - \beta_2\|^2 \geq \tilde{c} \phi_{\min}[2p_0] C_{p_0}^2 p_0 \quad (18)$$

On the other hand, by similar arguments, the Kullback-Leibler divergence satisfies

$$K(\mathbb{P}_{\mathbf{g}_1}, \mathbb{P}_{\mathbf{g}_2}) \leq \frac{\phi_{\max}[2p_0] C_{p_0}^2 \rho(\beta_1, \beta_2)}{2\sigma^2} \leq \frac{\phi_{\max}[2p_0] C_{p_0}^2 p_0}{\sigma^2} \quad (19)$$

Set now $C_{p_0}^2 = (1/16) \sigma^2 \tilde{c} (\ln(p/p_0) + 1) / \phi_{\max}[2p_0]$ and $s^2(p_0) = (1/64) \sigma^2 \times \tilde{c}^2 \tau [2p_0] p_0 (\ln(p/p_0) + 1)$. Then, (18) and (19) yield $\|\mathbf{g}_1 - \mathbf{g}_2\|^2 \geq 4s^2(p_0)$, $K(\mathbb{P}_{\mathbf{g}_1}, \mathbb{P}_{\mathbf{g}_2}) \leq (1/16) \ln \text{card}(\mathcal{G}_{p_0})$, and Lemma A.1 of [8] completes the proof.

CASE 2. $r/2 \leq p_0 \leq r, p_0 \geq 8$

In this case consider the subset $\tilde{\mathcal{B}}_{p_0} = \{\beta \in \mathbb{R}^p : \beta \in \{\{0, C_{p_0}\}^{p_0}, 0, \dots, 0\}\}$ and apply Varshamov-Gilbert bound (see, e.g. Lemma 2.9 in [27]). It guarantees the existence of a subset $\mathcal{B}_{p_0} \subset \tilde{\mathcal{B}}_{p_0}$ such that $\ln \text{card}(\mathcal{B}_{p_0}) \geq (p_0/8) \ln 2$ and the Hamming distance $\rho(\beta_1, \beta_2) \geq p_0/8$ for any pair $\beta_1, \beta_2 \in \mathcal{B}_{p_0}$.

Note also that for any $\beta_1, \beta_2 \in \mathcal{B}_{p_0}$, $\beta_1 - \beta_2$ has at most p_0 non-zero components and repeating the arguments for the Case 1, one achieves the minimax lower bound $s^2(p_0) = C \sigma^2 \tau [p_0] p_0 \geq (C/2) \sigma^2 \tau [p_0] r$.

CASE 3. $r/2 \leq p_0 \leq r$, $2 \leq p_0 < 8$

For this case, obviously, $2 \leq r < 16$. Consider a trivial subset \mathcal{B}_{p_0} containing just two vectors $\beta_1 \equiv 0$ and β_2 that has p_0 nonzero entries equal to $C_{p_0}^2 = (1/64)\sigma^2 \ln 2 / \phi_{max}[p_0]$. For the corresponding mean vectors $\mathbf{g}_1 = X\beta_1 = \mathbf{0}$ and $\mathbf{g}_2 = X\beta_2$, following (18) and (19) one has

$$K(\mathbb{P}_{\mathbf{g}_1}, \mathbb{P}_{\mathbf{g}_2}) \leq \frac{\phi_{max}[p_0]8C_{p_0}^2}{2\sigma^2} = (1/16) \ln \text{card}(\mathcal{G}_{p_0})$$

and

$$\|\mathbf{g}_1 - \mathbf{g}_2\|^2 \geq \phi_{min}[p_0]p_0C_{p_0}^2 = C\sigma^2\tau[p_0]p_0 \geq (C/2)\sigma^2\tau[p_0]r$$

Applying Lemma A.1 of [8] completes the proof. \square

8.5. Proof of Theorem 5

We want to show that under the conditions of Theorem 5 we can reduce the rate in the upper bound (9) for the risk of the MAP model selector established in Theorem 3 over \mathcal{M}_{p_0} by the factor $\tau_p[2p_0]$. Recall that we derived (9) from the general upper bound (8) in Theorem 1 by considering models M of size p_0 . Consider now $1 \leq \kappa_{p_1} \leq p_0 \leq \kappa_{p_2} \leq r/2$ and apply (8) for models M' of less size $p'_0 = \lceil \tau_p[2p_0]p_0 \rceil \geq 1$. Consider an arbitrary β_p with $\|\beta_p\|_0 = p_0$. Under the conditions of Theorem 5 on the prior $\pi_p(\cdot)$, (8) implies

$$E\|X_p\hat{\beta}_{\hat{M}_p} - X_p\beta_p\|^2 \leq C_1(\gamma) \inf_{M'} \|X_p\beta_p - X_p\beta_{pM'}\|^2 + C_2(\gamma)\sigma^2p'_0(\ln(p/p'_0) + 1), \tag{20}$$

where $X_p\beta_{pM'}$ is the projection of the mean vector $X_p\beta_p$ on the span of M' . Comparing (9) and (20) illustrates that reduction of a model size introduces the bias. On the other hand, under Assumptions (D.1) and (D.2), a straightforward calculus shows that the variance term decreases then to the desired order $p_0\tau_p[2p_0](\ln(p/p_0) + 1)$. The idea of the proof will be based on finding a model M'_* such that the resulting bias term will be at most of the same order as the reduced variance.

Consider the model M of size p_0 corresponding to β_p and any of its submodels M' of size p'_0 defined above. Evidently, M/M' is a subset of predictors from M not included in M' and $D_{M/M'} = D_M - D_{M'}$, where diagonal indicator matrices D 's were introduced in Section 2. By a straightforward calculus one then has

$$\begin{aligned} \|X_p\beta_p - X_p\beta_{pM'}\|^2 &= \beta_p' D_{M/M'} (D_{M/M'} (D_M X_p' X_p D_M)^+ D_{M/M'})^+ D_{M/M'} \beta_p \\ &= \beta_p' G_{p, \mathcal{J}_{M/M'}} \Lambda_p^{-1}(M, M') G_{p, \mathcal{J}_{M/M'}}' \beta_p \\ &\leq \tilde{\phi}_{M, M'}^{-1} \|G_{p, \mathcal{J}_{M/M'}}' \beta_p\|^2 \\ &\leq \tilde{\phi}_{M, M'}^{-1} \|\beta_p\|_\infty^2 (p_0 - p'_0), \end{aligned} \tag{21}$$

where the matrices G and Λ and the minimal eigenvalue $\tilde{\phi}_{M, M'}$ were defined in Section 5.2.

Among all submodels $M' \subset M$ of size p'_0 , choose M'_* with the maximal $\tilde{\phi}_{M, M'}$. Then, $\tilde{\phi}_{M, M'_*} \geq \tilde{\phi}_p[p_0]$ and (21) and Assumption (B) yield

$$\|X_p \beta_p - X_p \beta_{pM'_*}\|^2 \leq C \tau_p [2p_0] p_0 \ln((p/p_0) + 1)$$

Hence, we proved that under assumptions on the prior, Assumptions (D.1, D.2) and (B), the upper bound for the risk of the MAP model selector over \mathcal{M}_{p_0} is of the minimax order $\tau_p [2p_0] p_0 \ln((p/p_0) + 1)$. Assumption (D.3) guarantees that the “least-favorable” sets \mathcal{B}_{p_0} constructed in the proof of Theorem 4 satisfy the additional Assumption (B) on β_p and, therefore, the minimax lower bound (10) is not reduced. \square

Acknowledgments

The authors would like to thank Alexander Samarov and Yaácov Ritov for valuable remarks. The authors are especially grateful to an anonymous referee for an excellent constructive review of the first version of the paper.

References

- [1] ABRAMOVICH, F., BENJAMINI, Y., DONOHO, D.L. and JOHNSTONE, I.M. (2006). Adapting to unknown sparsity by controlling the false discovery rate. *Ann. Statist.* **34**, 584–653. [MR2281879](#)
- [2] ABRAMOVICH, F., GRINSHTEIN, V. and PENSKEY, M. (2007). On optimality of Bayesian testimation in the normal means problem. *Ann. Statist.* **35**, 2261–2286. [MR2363971](#)
- [3] ABRAMOVICH, F., GRINSHTEIN, V., PETSÁ, A. and SAPATINAS, T. (2010). On Bayesian testimation and its application to wavelet thresholding. *Biometrika* **97**, 181–198. [MR2594426](#)
- [4] AKAIKE, H. (1973). Information theory and an extension of the maximum likelihood principle. in *Second International Symposium on Information Theory*. (eds. B.N. Petrov and F. Czákí). Akademiai Kiadó, Budapest, 267–281. [MR0483125](#)
- [5] BICKEL, P., RITOV, Y. and TSYBAKOV, A. (2009). Simultaneous analysis of Lasso and Dantzig selector. *Ann. Statist.* **35**, 1705–1732. [MR2533469](#)
- [6] BIRGÉ, L. and MASSART, P. (2001). Gaussian model selection. *J. Eur. Math. Soc.* **3**, 203–268. [MR1848946](#)
- [7] BIRGÉ, L. and MASSART, P. (2007). Minimal penalties for Gaussian model selection. *Probab. Theory Relat. Fields* **138**, 33–73. [MR2288064](#)
- [8] BUNEA, F., TSYBAKOV, A. and WEGKAMP, M.H. (2007). Aggregation for Gaussian regression. *Ann. Statist.* **35**, 1674–1697. [MR2351101](#)
- [9] CANDÉS, E.J. (2006). Modern statistical estimation via oracle inequalities. *Acta Numerica* **15**, 257–325. [MR2269743](#)
- [10] CANDÉS, E.J. and TAO, T. (2007). The Dantzig selector: statistical estimation when p is much larger than n . *Ann. Statist.* **35**, 2313–2351. [MR2382644](#)

- [11] CHIPMAN, H., GEORGE, E.I. and MCCULLOGH, R.E. (2001). *The Practical Implementation of Bayesian Model Selection. IMS Lecture Notes – Monograph Series* **38**. [MR2000752](#)
- [12] DONOHO, D.L. and JOHNSTONE, I.M. (1994). Ideal spatial adaptation via wavelet shrinkage. *Biometrika* **81**, 425–455. [MR1311089](#)
- [13] DONOHO, D.L. and JOHNSTONE, I.M. (1995). Empirical atomic decomposition, *unpublished manuscript*.
- [14] FOSTER, D.P. and GEORGE, E.I. (1994). The risk inflation criterion for multiple regression. *Ann. Statist.* **22**, 1947–1975. [MR1329177](#)
- [15] GEORGE, E.I. and MCCULLOGH, R.E. (1993). Variable selection via Gibbs sampling. *J. Am. Statist. Assoc.* **88**, 881–889.
- [16] GEORGE, E.I. and MCCULLOGH, R.E. (1997). Approaches to Bayesian variable selection. *Statistica Sinica* **7**, 339–373.
- [17] GREENSHTEIN, E. and RITOV, Y. (2004). Persistence in high-dimensional linear predictor selection and the virtue of overparametrization. *Bernoulli* **10**, 971–988. [MR2108039](#)
- [18] HALL, P. and JIN, J. (2010). Innovated higher criticism for detecting sparse signals in correlated noise. *Ann. Statist.* **38**, 1681–1732. [MR2662357](#)
- [19] JOHNSTONE, I.M. (2002). *Function Estimation and Gaussian Sequence Models, unpublished manuscript*.
- [20] LIANG, F., PAULO, R., MOLINA, G., CLYDE, M. and BERGER, J.O. (2008). Mixtures of g priors for Bayesian variable selection. *J. Am. Statist. Assoc.* **103**, 410–423. [MR2420243](#)
- [21] MEINSHAUSEN, N. and YU, B. (2009). Lasso-type recovery of sparse representations for high-dimensional data. *Ann. Statist.* **37**, 246–270. [MR2488351](#)
- [22] RASKUTTI, G., WAINWRIGHT, M.J. and YU, B. (2009). Minimax rates of estimations for high-dimensional regression over l_q balls. *Technical Report, UC Berkeley*, <http://arxiv.org/abs/0910/2042>.
- [23] RIGOLLET, P. and TSYBAKOV, A. (2010). Exponential screening and optimal rates of sparse estimation. <http://arxiv.org/pdf/1003.2654>.
- [24] SCHWARZ, G. (1978). Estimating the dimension of a model. *Ann. Statist.* **6**, 461–464. [MR0468014](#)
- [25] TIBSHIRANI, R. (1996). Regression shrinkage and selection via the Lasso. *J. Roy. Statist. Soc. Ser. B* **58**, 267–288. [MR1379242](#)
- [26] TROPP, J.A. and WRIGHT, S.J. (2010). Computational methods for sparse solution of linear inverse problems. *Proc. IEEE, special issue “Applications of sparse representation and comprehensive sensing”*.
- [27] TSYBAKOV, A. (2009). *Introduction to Nonparametric Estimation*. Springer.
- [28] ZELLNER, A. (1986). On assessing prior distributions and Bayesian regression analysis with g -prior distributions. In *Bayesian Inference and Decision Techniques: Essays in Honor of Bruno de Finetti* (eds. Goel, P.K. and Zellner, A.), North-Holland, Amsterdam, 233–243. [MR0881437](#)