

This article was downloaded by: [University of Texas Libraries]

On: 23 April 2015, At: 14:55

Publisher: Taylor & Francis

Informa Ltd Registered in England and Wales Registered Number: 1072954 Registered office: Mortimer House, 37-41 Mortimer Street, London W1T 3JH, UK



Journal of the American Statistical Association

Publication details, including instructions for authors and subscription information:

<http://www.tandfonline.com/loi/uasa20>

Decoupling Shrinkage and Selection in Bayesian Linear Models: A Posterior Summary Perspective

P. Richard Hahn & Carlos M. Carvalho

Accepted author version posted online: 22 Jan 2015. Published online: 22 Apr 2015.



CrossMark

[Click for updates](#)

To cite this article: P. Richard Hahn & Carlos M. Carvalho (2015) Decoupling Shrinkage and Selection in Bayesian Linear Models: A Posterior Summary Perspective, Journal of the American Statistical Association, 110:509, 435-448, DOI: [10.1080/01621459.2014.993077](https://doi.org/10.1080/01621459.2014.993077)

To link to this article: <http://dx.doi.org/10.1080/01621459.2014.993077>

PLEASE SCROLL DOWN FOR ARTICLE

Taylor & Francis makes every effort to ensure the accuracy of all the information (the "Content") contained in the publications on our platform. However, Taylor & Francis, our agents, and our licensors make no representations or warranties whatsoever as to the accuracy, completeness, or suitability for any purpose of the Content. Any opinions and views expressed in this publication are the opinions and views of the authors, and are not the views of or endorsed by Taylor & Francis. The accuracy of the Content should not be relied upon and should be independently verified with primary sources of information. Taylor and Francis shall not be liable for any losses, actions, claims, proceedings, demands, costs, expenses, damages, and other liabilities whatsoever or howsoever caused arising directly or indirectly in connection with, in relation to or arising out of the use of the Content.

This article may be used for research, teaching, and private study purposes. Any substantial or systematic reproduction, redistribution, reselling, loan, sub-licensing, systematic supply, or distribution in any form to anyone is expressly forbidden. Terms & Conditions of access and use can be found at <http://www.tandfonline.com/page/terms-and-conditions>

Decoupling Shrinkage and Selection in Bayesian Linear Models: A Posterior Summary Perspective

P. Richard HAHN and Carlos M. CARVALHO

Selecting a subset of variables for linear models remains an active area of research. This article reviews many of the recent contributions to the Bayesian model selection and shrinkage prior literature. A posterior variable selection summary is proposed, which distills a full posterior distribution over regression coefficients into a sequence of sparse linear predictors.

KEY WORDS: Decision theory; Linear regression; Loss function; Model selection; Parsimony; Shrinkage prior; Sparsity; Variable selection.

1. INTRODUCTION

This article revisits the venerable problem of variable selection in linear models. The vantage point throughout is Bayesian: a normal likelihood is assumed and inferences are based on the posterior distribution, which is arrived at by conditioning on observed data.

In applied regression analysis, a “high-dimensional” linear model can be one which involves tens or hundreds of variables, especially when seeking to compute a full Bayesian posterior distribution. Our review will be from the perspective of a data analyst facing a problem in this “moderate” regime. Likewise, we focus on the situation where the number of predictor variables, p , is fixed.

In contrast to other recent papers surveying the large body of literature on Bayesian variable selection (Liang et al. 2008; Bayarri et al. 2012) and shrinkage priors (O’Hara and Silanpää 2009; Polson and Scott 2012), our review focuses specifically on the relationship between variable selection priors and shrinkage priors. Selection priors and shrinkage priors are related both by the statistical ends they attempt to serve (e.g., strong regularization and efficient estimation) and also in the technical means they use to achieve these goals (hierarchical priors with local scale parameters). We also compare these approaches on computational considerations.

Finally, we turn to variable selection as a problem of posterior summarization. We argue that if variable selection is desired primarily for parsimonious communication of linear trends in the data, that this can be accomplished as a post-inference operation irrespective of the choice of prior distribution. To this end, we introduce a posterior variable selection summary, which distills a full posterior distribution over regression coefficients into a sequence of sparse linear predictors. In this sense “shrinkage” is decoupled from “selection.”

We begin by describing the two most common approaches to this scenario and show how the two approaches can be seen as special cases of an encompassing formalism.

1.1 Bayesian Model Selection Formalism

A now-canonical way to formalize variable selection in Bayesian linear models is as follows. Let \mathcal{M}_ϕ denote a normal linear regression model indexed by a vector of binary indicators $\phi = (\phi_1, \dots, \phi_p) \in \{0, 1\}^p$ signifying which predictors are included in the regression. Model \mathcal{M}_ϕ defines the data distribution as

$$(Y_i | \mathcal{M}_\phi, \beta_\phi, \sigma^2) \sim N(X_i^\phi \beta_\phi, \sigma^2), \quad (1)$$

where X_i^ϕ represents the p_ϕ -vector of predictors in model \mathcal{M}_ϕ . For notational simplicity, (1) does not include an intercept. Standard practice is to include an intercept term and to assign it a uniform prior.

Given a sample $\mathbf{Y} = (Y_1, \dots, Y_n)$ and prior $\pi(\beta_\phi, \sigma^2)$, the inferential target is the set of posterior model probabilities defined by

$$p(\mathcal{M}_\phi | \mathbf{Y}) = \frac{p(\mathbf{Y} | \mathcal{M}_\phi)p(\mathcal{M}_\phi)}{\sum_\phi p(\mathbf{Y} | \mathcal{M}_\phi)p(\mathcal{M}_\phi)}, \quad (2)$$

where $p(\mathbf{Y} | \mathcal{M}_\phi) = \int p(\mathbf{Y} | \mathcal{M}_\phi, \beta_\phi, \sigma^2)\pi(\beta_\phi, \sigma^2)d\beta_\phi d\sigma^2$ is the marginal likelihood of model \mathcal{M}_ϕ and $p(\mathcal{M}_\phi)$ is the prior over models.

Posterior inferences concerning a quantity of interest Δ are obtained via Bayesian model averaging (or BMA), which entails integrating over the model space

$$p(\Delta | \mathbf{Y}) = \sum_\phi p(\Delta | \mathcal{M}_\phi, \mathbf{Y})p(\mathcal{M}_\phi | \mathbf{Y}). \quad (3)$$

As an example, optimal predictions of future values of \tilde{Y} under squared-error loss are defined through

$$E(\tilde{Y} | \mathbf{Y}) \equiv \sum_\phi E(\tilde{Y} | \mathcal{M}_\phi, \mathbf{Y})p(\mathcal{M}_\phi | \mathbf{Y}). \quad (4)$$

An early reference adopting this formulation is Raftery, Madigan, and Hoeting (1997); see also Clyde and George (2004).

Despite its straightforwardness, carrying out variable selection in this framework demands attention to detail: priors over model-specific parameters must be specified, priors over

P. Richard Hahn is Assistant Professor, Booth School of Business, The University of Chicago, Chicago, IL 60611 (E-mail: richard.hahn@chicagobooth.edu). Carlos M. Carvalho is Associate Professor of Statistics, McCombs School of Business, The University of Texas, 2110 Speedway Stop, Austin, TX 78712 (E-mail: carlos.carvalho@mcombs.utexas.edu).

models must be chosen, marginal likelihood calculations must be performed and a 2^p -dimensional discrete space must be explored. These concerns have animated Bayesian research in linear model variable selection for the past two decades (George and McCulloch 1993, 1997; Clyde and George 2004; Hans, Dobra, and West 2007; Liang et al. 2008; Scott and Berger 2010; Clyde, Ghosh, and Littman 2011; Bayarri et al. 2012).

Regarding model parameters, the consensus default prior for model parameters is $\pi(\beta_\phi, \sigma^2) = \pi(\beta | \sigma^2)\pi(\sigma^2) = N(0, g\Omega) \times \sigma^{-1}$. The most widely studied choice of prior covariance is $\Omega = \sigma^2(\mathbf{X}'_\phi \mathbf{X}_\phi)^{-1}$, referred to as “Zellner’s g -prior” (Zellner 1986), a “ g -type” prior or simply g -prior. Notice that this choice of Ω dictates that the prior and likelihood are conjugate normal-inverse-gamma pairs (for a fixed value of g).

For reasons detailed by Liang et al. (2008), it is advised to place a prior on g rather than use a fixed value. Several recent papers describe priors $p(g)$ that still lead to efficient computations of marginal likelihoods; see Cui and George (2008), Liang et al. (2008), Maruyama and George (2011), and Bayarri et al. (2012). Each of these papers (as well as the earlier literature cited therein) study priors of the form

$$p(g) = a[\rho_\phi(b+n)]^a g^d (g+b)^{-(a+c+d+1)} \mathbb{1}\{g > \rho_\phi(b+n) - b\} \quad (5)$$

with $a > 0$, $b > 0$, $c > -1$, and $d > -1$. Specific configurations of these hyperparameters recommended in the literature include: $\{a = 1, b = 1, d = 0, \rho_\phi = 1/(1+n)\}$ (Cui and George 2008), $\{a = 1/2, b = 1 (b = n), c = 0, d = 0, \rho_\phi = 1/(1+n)\}$ (Liang et al. 2008), and $\{a = 1, b = 1, c = -3/4, d = (n-5)/2 - \rho_\phi/2 + 3/4, \rho_\phi = 1/(1+n)\}$ (Maruyama and George 2011).

Bayarri et al. (2012) motivated the use of such priors from a testing perspective, using a variety of formal desiderata based on Jeffreys (1961) and Berger and Pericchi (2001), including consistency criteria, predictive matching criteria and invariance criteria. Their recommended prior uses $\{a = 1/2, b = 1, c = 0, d = 0, \rho_\phi = 1/p_\phi\}$. This prior is termed the *robust prior*, in the tradition following Strawderman (1971) and Berger (1980), who examine the various senses in which such priors are “robust.” This prior will serve as a benchmark in the examples of Section 3.

Regarding prior model probabilities, see Scott and Berger (2010), who recommend a hierarchical prior of the form $\phi_j \stackrel{\text{iid}}{\sim} \text{Ber}(q)$, $q \sim \text{Unif}(0, 1)$.

1.2 Shrinkage Regularization Priors

Although the formulation above provides a valuable theoretical framework, it does not necessarily represent an applied statistician’s first choice. To assess which variables contribute dominantly to trends in the data, the goal may be simply to mitigate—rather than categorize—spurious correlations. Thus, faced with many potentially irrelevant predictor variables, a common first choice would be a powerful *regularization prior*.

Regularization—understood here as the intentional biasing of an estimate to stabilize posterior inference—is inherent to most Bayesian estimators via the use of proper prior distributions and is one of the often-cited advantages of the Bayesian approach. More specifically, regularization priors refer to priors explicitly

designed with a strong bias for the purpose of separating reliable from spurious patterns in the data. In linear models, this strategy takes the form of zero-centered priors with sharp modes and simultaneously fat tails.

A well-studied class of priors fitting this description will serve to connect continuous priors to the model selection priors described above. *Local scale mixture* of normal distributions are of the form (West 1987; Carvalho, Polson, and Scott 2010; Griffin and Brown 2012)

$$\pi(\beta_j | \lambda) = \int N(\beta_j | 0, \lambda^2 \lambda_j^2) \pi(\lambda_j^2) d\lambda_j, \quad (6)$$

where different priors are derived from different choices for $\pi(\lambda_j^2)$.

The last several years have seen tremendous interest in this area, motivated by an analogy with penalized likelihood methods (Tibshirani 1996). Penalized likelihood methods with an additive penalty term lead to estimating equations of the form

$$\sum_i h(Y_i, \mathbf{X}_i, \beta) + \alpha Q(\beta), \quad (7)$$

where h and Q are positive functions and their sum is to be minimized; α is a scalar tuning variable dictating the strength of the penalty. Typically, h is interpreted as a negative log-likelihood, given data \mathbf{Y} , and Q is a penalty term introduced to stabilize maximum likelihood estimation. A common choice is $Q(\beta) = \|\beta\|_1$, which yields sparse optimal solutions β^* and admits fast computation (Tibshirani 1996); this choice underpins the *lasso* estimator, an initialism for “least absolute shrinkage and selection operator.”

Park and Casella (2008) and Hans (2009) “Bayesified” these expressions by interpreting $Q(\beta)$ as the negative log prior density and developing algorithms for sampling from the resulting Bayesian posterior, building upon work of earlier Bayesian authors (Spiegelhalter 1977; West 1987; Pericchi and Walley 1991; Pericchi and Smith 1992). Specifically, an exponential prior $\pi(\lambda_j^2) = \text{Exp}(\alpha^2)$ leads to independent Laplace (double-exponential) priors on the β_j , mirroring expression (7).

This approach has two implications unique to the Bayesian paradigm. First, it presented an opportunity to treat the global scale parameter λ (equivalently the regularization penalty parameter α) as a hyperparameter to be estimated. Averaging over λ in the Bayesian paradigm has been empirically observed to give better prediction performance than cross-validated selection of α (e.g., Hans 2009). Second, a Bayesian approach necessitates forming point estimators from posterior distributions; typically the posterior mean is adopted on the basis that it minimizes mean squared prediction error. Note that posterior mean regression coefficient vectors from these models are nonsparse with probability one. Ironically, the two main appeals of the penalized likelihood methods—efficient computation and sparse solution vectors β^* —were lost in the migration to a Bayesian approach. See, however, Hans (2010) for an application of double-exponential priors in the context of model selection.

Nonetheless, wide interest in “Bayesian lasso” models paved the way for more general local shrinkage regularization priors of the form (6). In particular, Carvalho, Polson, and Scott (2010) developed a prior over location parameters that attempts to shrink irrelevant signals strongly toward zero while avoiding

excessive shrinkage of relevant signals. To contextualize this aim, recall that solutions to ℓ_1 penalized likelihood problems are often interpreted as (convex) approximations to more challenging formulations based on ℓ_0 penalties: $\|\gamma\|_0 = \sum_j \mathbb{1}(\gamma_j \neq 0)$. As such, it was observed that the global ℓ_1 penalty “overshrinks” what ought to be large magnitude coefficients. For one example, the Carvalho, Polson, and Scott (2010) prior may be written as

$$\pi(\beta_j | \lambda) = N(0, \lambda^2 \lambda_j^2),$$

$$\lambda_j \stackrel{\text{iid}}{\sim} C^+(0, 1). \tag{8}$$

with $\lambda \sim C^+(0, 1)$ or $\lambda \sim C^+(0, \sigma^2)$. The choice of half-Cauchy arises from the insight that for scalar observations $y_j \sim N(\theta_j, 1)$ and prior $\theta_j \sim N(0, \lambda_j^2)$, the posterior mean of θ_j may be expressed:

$$E(\theta_j | y_j) = \{1 - E(\kappa_j | y_j)\}y_j, \tag{9}$$

where $\kappa_j = 1/(1 + \lambda_j^2)$. The authors observe that U-shaped Beta(1/2,1/2) distributions (like a horseshoe) on κ_j imply a prior over θ_j with high mass around the origin but with polynomial tails. That is, the “horseshoe” prior encodes the assumption that some coefficients will be very large and many others will be very nearly zero. This U-shaped prior on κ_j implies the half-Cauchy prior density $\pi(\lambda_j)$. The implied marginal prior on β has Cauchy-like tails and a pole at the origin which entails more aggressive shrinkage than a Laplace prior.

Other choices of $\pi(\lambda_j)$ lead to different “shrinkage profiles” on the “ κ scale.” Polson and Scott (2012) provided an excellent taxonomy of the various priors over β that can be obtained as scale-mixtures of normals. The horseshoe and similar priors (e.g., Griffin and Brown 2012) have proven empirically to be fine default choices for regression coefficients: they lack hyperparameters, forcefully separate strong from weak predictors, and exhibit robust predictive performance.

1.3 Model Selection Priors as Shrinkage Priors

It is possible to express model selection priors as shrinkage priors. To motivate this reframing, observe that the posterior mean regression coefficient vector is not well-defined in the model selection framework. Using the model-averaging notion, the posterior average β may be defined as

$$E(\beta | \mathbf{Y}) \equiv \sum_{\phi} E(\beta | \mathcal{M}_{\phi}, \mathbf{Y})p(\mathcal{M}_{\phi} | \mathbf{Y}), \tag{10}$$

where $E(\beta_j | \mathcal{M}_{\phi}, \mathbf{Y}) \equiv 0$ whenever $\phi_j = 0$. Without this definition, the posterior expectation of β_j is undefined in models where the j th predictor does not appear. More specifically, as the likelihood is constant in variable j in such models, the posterior remains whatever the prior was chosen to be.

To fully resolve this indeterminacy, it is common to set β_j identically equal to zero in models where the j th predictor does not appear, consistent with the interpretation that $\beta_j \equiv \partial E(Y)/\partial X_j$. A hierarchical prior reflecting this choice may be expressed as

$$\pi(\beta | g, \Lambda, \Omega) = N(0, g \Lambda \Omega \Lambda^t). \tag{11}$$

In this expression, $\Lambda \equiv \text{diag}((\lambda_1, \lambda_2, \dots, \lambda_p))$ and Ω is a positive semidefinite matrix, both of which may depend on ϕ and/or σ^2 . When Ω is the identity matrix, one recovers (6).

To set $\beta_j = 0$ when $\phi_j = 0$, let $\lambda_j \equiv \phi_j s_j$ for $s_j > 0$, so that when $\phi_j = 0$, the prior variance of β_j is set to zero (with prior mean of zero). George and McCulloch (1997) developed this approach in detail, including the g -prior specification, $\Omega(\phi) = \sigma^2(\mathbf{X}'_{\phi} \mathbf{X}_{\phi})^{-1}$. Priors over the s_j induce a prior on Λ . Under these definitions of the λ_j and Ω , the component-wise marginal distribution for β_j , $j = 1, \dots, p$, may be written as

$$\pi(\beta_j | \phi_j, \sigma^2, g, \lambda_j) = (1 - \phi_j)\delta_0 + \phi_j N(0, g \lambda_j^2 \omega_j), \tag{12}$$

where δ_0 denotes a point mass distribution at zero and ω_j is the j th diagonal element of Ω . Hierarchical priors of this form are sometimes called “spike-and-slab” priors (δ_0 is the spike and the continuous full-support distribution is the slab) or the “two-groups model” for variable selection. References for this specification include Mitchell and Beauchamp (1988) and Geweke (1996), among others.

It is also possible to think of the component-wise prior over each β_j directly in terms of (11) and a prior over λ_j (marginalizing over ϕ):

$$\pi(\lambda_j | q) = (1 - q)\delta_0 + q P_{\lambda_j}, \tag{13}$$

where $\text{Pr}(\phi_j = 1) = q$, and P_{λ_j} is some continuous distribution on \mathbb{R}^+ . Of course, q can be given a prior distribution as well; a uniform distribution is common. This representation transparently embeds model selection priors within the class of local scale mixture of normal distributions. An important article exploring the connections between shrinkage priors and model selection priors is Ishwaran and Rao (2005), who considered a version of (11) via a specification of $\pi(\lambda_j)$ which is bimodal with one peak at zero and one peak away from zero. In many respects, this article anticipated the work of Park and Casella (2008), Hans (2009), Carvalho, Polson, and Scott (2010), Griffin and Brown (2012), Polson and Scott (2012), and the like.

1.4 Computational Issues in Variable Selection

Because posterior sampling is computation-intensive and because variable selection is most desirable in contexts with many predictor variables, computational considerations are important in motivating and evaluating the approaches above. The discrete model selection approach and the continuous shrinkage prior approach are both quite challenging in terms of posterior sampling.

In the model selection setting, for $p > 30$, enumerating all possible models (e.g., to compute marginal likelihoods) is beyond the reach of modern capability. As such, stochastic exploration of the model space is required, with the hope that the unvisited models comprise a vanishingly small fraction of the posterior probability. George and McCulloch (1997) were frank about this limitation; noting that a Markov chain run of length less than 2^p steps cannot have visited each model even once, they wrote hopefully that “it may thus be possible to identify at least some of the high probability values.”

Garcia-Donato and Martinez-Beneito (2013) carefully evaluated methods for dealing with this problem and came to compelling conclusions in favor of some methods over others. Their analysis is beyond the scope of this article, but anyone interested in the variable selection problem in large p settings should consult their advice. In broad strokes, they found that MCMC approaches based on Gibbs samplers

(i.e., George and McCulloch 1997) appear better at estimating posterior quantities—such as the highest probability model, the median probability model (MPM), etc.—compared to methods based on sampling without replacement (i.e., Hans, Dobra, and West 2007 and Clyde, Ghosh, and Littman 2011).

Regarding shrinkage priors, there is no systematic study in the literature suggesting that the above computational problems are alleviated for continuous parameters. In fact, the results of Garcia-Donato and Martinez-Beneito (2013) (see Section 6) suggest that posterior sampling in finite sample spaces is easier than the corresponding problem for continuous parameters, in that convergence to stationarity occurs more rapidly.

Moreover, if one is willing to entertain an extreme prior with $\pi(\phi) = 0$ for $\|\phi\|_0 > M$ for a given constant M , model selection priors offer a tremendous practical benefit: one never has to invert a matrix larger than $M \times M$, rather than the $p \times p$ dimensional inversions required of a shrinkage prior approach. Similarly, only vectors up to size M need to be saved in memory and operated upon. In extremely large problems, with thousands of variables, setting $M = \mathcal{O}(\sqrt{p})$ or $M = \mathcal{O}(\log p)$ saves considerable computational effort (Hans, Dobra, and West 2007). According to personal communications with researchers at Google, this approach is routinely applied to large-scale Internet data. Should M be chosen too small, little can be said; if M truly represents one's computational budget, the best model of size M will have to do.

1.5 Selection: From Posteriors to Sparsity

To go from a posterior distribution to a sparse point estimate requires additional processing, regardless of what prior is used. The specific process used to achieve sparse estimates will depend on the underlying purpose for which the sparsity is desired.

In some cases, identifying sparse models (subsets of nonzero coefficients) might be an end in itself, as in the case of trying to isolate scientifically important variables in the context of a controlled experiment. In this case, a prior with point-mass probabilities at the origin is necessary in terms of defining the implicit (multiple) testing problem. For this purpose, the use of posterior model probabilities is a well-established methodology for evaluating evidence in the data in favor of various hypotheses. Indeed, the highest posterior probability model (HPM) is optimal under model selection loss: $\mathcal{L}(\gamma, \phi) = \mathbb{1}(\gamma = \phi)$, where γ denotes the “action” of selecting a particular model. Under symmetric variable selection loss, $\mathcal{L}(\gamma, \phi) = \sum_j |\gamma_j - \phi_j|$, it is easy to show that the optimal model is the one which includes all and only variables with marginal posterior inclusion probabilities greater than 1/2. This model is commonly referred to as the *median probability model* (MPM).

Note that many authors discuss model selection in terms of Bayes factors with respect to a “base model” \mathcal{M}_{ϕ^*} :

$$\frac{p(\mathcal{M}_\phi | \mathbf{Y})}{p(\mathcal{M}_{\phi^*} | \mathbf{Y})} = \text{BF}(\mathcal{M}_\phi, \mathcal{M}_{\phi^*}) \frac{p(\mathcal{M}_\phi)}{p(\mathcal{M}_{\phi^*})}, \quad (14)$$

where \mathcal{M}_{ϕ^*} is typically chosen to be the full model with no zero coefficients or the null model with all zero coefficients. This notation should not obscure the fact that posterior model probabilities underlie subsequent model selection decisions.

As an alternative to posterior model probabilities, many ad hoc *hard thresholding* methods have been proposed, which can be employed when $\pi(\lambda_j)$ has a non-point-mass density. Such methods derive classification rules for selecting subsets of nonzero coefficients, on the basis of the posterior distribution over β_j and/or λ_j . For example, Carvalho, Polson, and Scott (2010) suggested setting to zero those coefficients for which

$$E(\kappa_j = 1/(1 + \lambda_j^2) | \mathbf{Y}) < 1/2.$$

Ishwaran and Rao (2005) discussed a variety of posterior thresholding rules and relate them to conventional thresholding rules based on ordinary least-square estimates of β . An important limitation of most commonly used thresholding approaches is that they are applied separately to each coefficient, irrespective of any dependencies that arise in the posterior between the elements of $\lambda_1, \dots, \lambda_p$.

In other cases, the goal—rather than isolating all and only relevant variables, no matter their absolute size—is simply to describe the “important” relationships between predictors and response. In this case, the model selection route is simply a means to an end. From this perspective, a natural question is how to fashion a sparse vector of regression coefficients which parsimoniously characterizes the available data. Leamer (1978) is a notable early effort advocating ad hoc model selection for the purpose of human comprehensibility. Fouskakis and Draper (2008), Fouskakis, Ntzoufras, and Draper (2009), and Draper (2013) represent efforts to define variable importance in real-world terms using subject matter considerations. A more generic approach is to gauge predictive relevance (Gelfand, Dey, and Chang 1992).

A widely cited result relating variable selection to predictive accuracy is that of Barbieri and Berger (2004). Consider mean squared prediction error (MSPE), $n^{-1}E\{\sum_i (\tilde{Y}_i - \tilde{X}_i \hat{\beta})^2\}$, and recall that the model-specific optimal regression vector is $\hat{\beta}_\phi \equiv E(\beta | \mathcal{M}_\phi, \mathbf{Y})$. Barbieri and Berger (2004) showed that for $\mathbf{X}'\mathbf{X}$ diagonal, the best predicting model according to MSPE is again the median probability model. Their result holds both for a fixed design $\tilde{\mathbf{X}}$ of prediction points or for stochastic predictors with $E\{\tilde{X}'\tilde{X}\}$ diagonal. However, the main condition of their theorem— $\mathbf{X}'\mathbf{X}$ diagonal—is almost never satisfied in practice. Nonetheless, they argue that the median probability model tends to outperform the highest probability model on out-of-sample prediction tasks. Note that the HPM and MPM can be substantially different models, especially in the case of strong dependence among predictors.

Broadly speaking, the difference in the two situations described above is one between “statistical significance” and “practical significance.” In the former situation, posterior model probabilities are the preferred alternative, with thresholding rules being an imperfect analogue for use with continuous (non-point-mass) priors on β . In the latter case, predictive relevance is a commonly invoked operational definition of “practical,” but theoretical results are not available for the case of correlated predictor variables.

2. POSTERIOR SUMMARY VARIABLE SELECTION

In this section, we describe a posterior summary based on an expected loss minimization problem. The loss function is

designed to balance prediction ability (in the sense of mean square prediction error) and narrative parsimony (in the sense of sparsity). The new summary checks three important boxes:

- it produces sparse vectors of regression coefficients for prediction,
- it can be applied to a posterior distribution arising from any prior distribution, and
- it explicitly accounts for co-linearity in the matrix of prediction points and dependencies in the posterior distribution of β .

2.1 The Cost of Measuring Irrelevant Variables

Suppose that collecting information on individual covariates incurs some cost; thus the goal is to make an accurate enough prediction subject to a penalty for acquiring predictively irrelevant facts.

Consider the problem of predicting an n -vector of future observables $\tilde{Y} \sim N(\tilde{X}\beta, \sigma^2\mathbf{I})$ at a pre-specified set of design points \tilde{X} . Assume that a posterior distribution over the model parameters (β, σ^2) has been obtained via Bayesian conditioning, given past data \mathbf{Y} and design matrix \mathbf{X} ; denote the density of this posterior by $\pi(\beta, \sigma^2 | \mathbf{Y})$.

It is crucial to note that \tilde{X} and \mathbf{X} need not be the same. That is, the locations in predictor space where one wants to predict need not be the same points at which one has already observed past data. For notational simplicity, we will write \mathbf{X} instead of \tilde{X} in what follows. Of course, taking $\tilde{X} = \mathbf{X}$ is a conventional choice, but distinguishing between the two becomes important in certain cases such as when $p > n$.

Define an optimal action as one which minimizes expected loss $E(\mathcal{L}(\tilde{Y}, \gamma))$ over all model selection vectors γ , where the expectation is taken over the predictive distribution of unobserved values:

$$f(\tilde{Y} | \mathbf{Y}) = \int f(\tilde{Y} | \beta, \sigma^2)\pi(\beta, \sigma^2 | \mathbf{Y})d(\beta, \sigma^2). \quad (15)$$

As a widely applicable loss function, consider

$$\mathcal{L}(\tilde{Y}, \gamma) = \lambda\|\gamma\|_0 + n^{-1}\|\mathbf{X}\gamma - \tilde{Y}\|_2^2, \quad (16)$$

where again $\|\gamma\|_0 = \sum_j \mathbb{1}(\gamma_j \neq 0)$. This loss sums two components, one of which is a ‘‘parsimony penalty’’ on the action γ and the other of which is the squared prediction loss of the linear predictor defined by γ . The scalar utility parameter λ dictates how severely we penalize each of these two components, relatively. Integrating over \tilde{Y} conditional on (β, σ^2) (and overloading the notation of \mathcal{L}) gives

$$\mathcal{L}(\beta, \sigma, \gamma) \equiv E(\mathcal{L}(\tilde{Y}, \gamma)) = \lambda\|\gamma\|_0 + n^{-1}\|\mathbf{X}\gamma - \mathbf{X}\beta\|_2^2 + \sigma^2. \quad (17)$$

Because (β, σ^2) are unknown, an additional integration over $\pi(\beta, \sigma^2 | \mathbf{Y})$ yields

$$\mathcal{L}(\gamma) \equiv E(\mathcal{L}(\beta, \sigma, \gamma)) = \lambda\|\gamma\|_0 + \bar{\sigma}^2 + n^{-1}\text{tr}(\mathbf{X}'\mathbf{X}\Sigma_\beta) + n^{-1}\|\mathbf{X}\bar{\beta} - \mathbf{X}\gamma\|_2^2, \quad (18)$$

where $\bar{\sigma}^2 = E(\sigma^2)$, $\bar{\beta} = E(\beta)$ and $\Sigma_\beta = \text{cov}(\beta)$, and all expectations are with respect to the posterior.

Dropping constant terms, one arrives at the ‘‘decoupled shrinkage and selection’’ (DSS) loss function:

$$\mathcal{L}(\gamma) = \lambda\|\gamma\|_0 + n^{-1}\|\mathbf{X}\bar{\beta} - \mathbf{X}\gamma\|_2^2. \quad (19)$$

Optimization of the DSS loss function is a combinatorial programming problem depending on the posterior distribution via the posterior mean of β ; the DSS loss function explicitly trades off the number of variables in the linear predictor with its resulting predictive performance. Denote this optimal solution by

$$\beta_\lambda \equiv \arg \min_\gamma \lambda\|\gamma\|_0 + n^{-1}\|\mathbf{X}\bar{\beta} - \mathbf{X}\gamma\|_2^2. \quad (20)$$

Note that the above derivation applies straightforwardly to the selection prior setting via expression (10) or (equivalently) via the hierarchical formulation in (12), which guarantee that $\bar{\beta}$ is well-defined marginally across different models.

2.2 Analogy With High Posterior Density Regions

Although orthographically (19) resembles expressions used in penalized likelihood methods, the better analogy is a Bayesian highest posterior density (HPD) region, defined as the shortest contiguous interval encompassing some fixed fraction of the posterior mass. The insistence on reporting the *shortest* interval is analogous to the DSS summary being defined in terms of the *sparsest* linear predictor which still has reasonable prediction performance. Like HPD regions, DSS summaries are well defined under any prior giving a proper posterior.

To amplify, the DSS optimization problem is well-defined for any *posterior* as long as $\bar{\beta}$ exists. Different priors may lead to very different posteriors, potentially with very different means. However, regardless of the precise nature of the posterior (e.g., the presence of multimodality), $\bar{\beta}$ is the optimal summary under squared-error prediction loss, which entails that expression (20) represents the sparsified solution to the optimization problem given in (16).

An important implication of this analogy is the realization that a DSS summary can be produced for a prior distribution directly, in the same way that a prior distribution has an HPD (with the posterior trivially equal to the prior). The DSS summary requires the user to specify a matrix of prediction points \tilde{X} , but conditional on this choice one can extract sparse linear predictors directly from a prior distribution.

In Section 3, we discuss strategies for using additional features of the posterior $\pi(\beta, \sigma^2 | \mathbf{Y})$ to guide the choice of picking λ .

2.3 Computing and Approximating β_λ

The counting penalty $\|\gamma\|_0$ yields an intractable optimization problem for even tens of variables ($p \approx 30$). This problem has been addressed in recent years by approximating the counting norm with modifications of the ℓ_1 norm, $\|\gamma\|_1 = \sum_h |\gamma_h|$, leading to a surrogate loss function which is convex and readily minimized by a variety of software packages. Crucially, such approximations still yield a sequence of sparse actions (the solution path as a function of λ), simplifying the 2^p dimensional selection problem to a choice between at most p alternatives. The goodness of these approximations is a natural and relevant concern. Note, however, that the goodness of approximation is

a computational rather than inferential concern. This is what is meant by “decoupled” shrinkage and selection.

More specifically, recall that DSS requires finding the optimal solution defined in (20). The most simplistic and yet widely used approximation is to replace the ℓ_0 norm with the ℓ_1 norm, which leads to a convex optimization problem for which many implementations are available, in particular the `lars` algorithm (Efron et al. 2004). Using this approximation, β_λ can be obtained simply by running the `lars` algorithm using $\bar{\mathbf{Y}} = \mathbf{X}\bar{\beta}$ as the “data.”

It is well-known that the ℓ_1 approximation may unduly “shrink” all elements of β_λ beyond the shrinkage arising naturally from the prior over β . To avoid this potential “double-shrinkage” it is possible to explicitly adjust the ℓ_1 approach toward the desired ℓ_0 target. Specifically, the local linear approximation argument of Zou and Li (2008) and Lv and Fan (2009) advises to solve a surrogate optimization problem (for any w_j near the corresponding ℓ_0 solution)

$$\beta_\lambda \equiv \arg \min_{\gamma} \sum_j \frac{\lambda}{|w_j|} |\gamma_j| + n^{-1} \|\mathbf{X}\bar{\beta} - \mathbf{X}\gamma\|_2^2. \quad (21)$$

This approach yields a procedure analogous to the adaptive lasso of Zou (2006) with $\bar{\mathbf{Y}} = \mathbf{X}\bar{\beta}$ in place of \mathbf{Y} . In what follows, we use $w_j = \bar{\beta}_j$ (whereas the adaptive lasso uses the least-squares estimate $\hat{\beta}_j$). The `lars` package in R can then be used to obtain solutions to this objective function by a straightforward rescaling of the design matrix.

In our experience, this approximation successfully avoids double-shrinkage. In fact, as illustrated in the U.S. crime example below, this approach is able to un-shrink coefficients depending on which variables are selected into the model.

For a fixed value of λ , expression (19) uniquely determines a sparse vector β_λ as its corresponding Bayes estimator. However, choosing λ to define this estimator is a nontrivial decision in its own right. Section 3 considers how to use the posterior distribution $\pi(\beta, \sigma^2 | \mathbf{Y})$ to illuminate the trade-offs implicit in the selection of a given value of λ .

3. SELECTION SUMMARY PLOTS

How should one think about the generalization error across possible values of λ ? Consider first two extreme cases. When $\lambda = 0$, the solution to the DSS optimization problem is simply the posterior mean: $\beta_{\lambda=0} \equiv \bar{\beta}$. Conversely, for very large λ , the optimal solution will be the zero vector, $\beta_{\lambda=\infty} = 0$, which will have expected prediction loss equal to the marginal variance of the response Y (which will depend on the predictor points in question).

A sensible way to judge the goodness of β_λ in terms of prediction is relative to the predictive performance of β —were it known—which is the optimal linear predictor under squared-error loss. The relevant scale for this comparison is dictated by σ^2 , which quantifies the best one can hope to do even if β were known. With these benchmarks in mind, one wants to address the question: how much predictive deterioration is a result of sparsification?

The remainder of this section defines three plots that can be used by a data analyst to visualize the predictive deterioration across various values of λ . The first plot concerns a measure

of “variation-explained,” the second plot considers the excess prediction loss on the scale of the response variable, and the final plot looks at the magnitude of the elements of β_λ .

In the following examples, the outcome variable and covariates are centered and scaled to mean zero and unit variance. This step is not strictly required, but it does facilitate default prior specification. Likewise the solution to (19) is invariant to scaling, but the approximation based on (21) is sensitive to the scale of the predictors. Finally, although the exposition above assumed no intercept, in the examples an intercept is always fit; the intercept is given a flat prior and does not appear in the formulation of the DSS loss function.

3.1 Variation-Explained of a Sparsified Linear Predictor

Define the “variation-explained” at design points \mathbf{X} (perhaps different than those seen in the data sample used to form the posterior distribution) as

$$\rho^2 = \frac{n^{-1} \|\mathbf{X}\beta\|^2}{n^{-1} \|\mathbf{X}\beta\|^2 + \sigma^2}. \quad (22)$$

Denote by

$$\rho_\lambda^2 = \frac{n^{-1} \|\mathbf{X}\beta\|^2}{n^{-1} \|\mathbf{X}\beta\|^2 + \sigma^2 + n^{-1} \|\mathbf{X}\beta - \mathbf{X}\beta_\lambda\|^2} \quad (23)$$

the analogous quantity for the sparsified linear predictor β_λ . The gap between β_λ and β due to sparsification is tallied as a contribution to the noise term, which decreases the variation-explained. This quantity has the benefit of being directly comparable to the ubiquitous R^2 metric of model fit familiar to users of statistical software and least-squares theory.

Posterior samples of ρ_λ^2 can be obtained as follows.

1. First, solve (21) by applying the `lars` algorithm with inputs $w_j = \bar{\beta}_j$ and $\mathbf{Y} = \mathbf{X}\bar{\beta}$. A single run of this algorithm will produce a sequence of solutions β_λ for a range of λ values. (Obtaining draws of ρ_λ^2 using a model selection prior requires posterior samples from (β, σ^2) marginally across models.)
2. Second, for each element in the sequence of β_λ 's, convert posterior samples of (β, σ^2) into samples of ρ_λ^2 via definition (23).
3. Finally, plot the expected value and 90% credible intervals of ρ_λ^2 against the model size, $\|\beta_\lambda\|_\lambda$. The posterior mean of ρ_0^2 may be overlaid as a horizontal line for benchmarking purposes; note that even for $\lambda = 0$ (so that $\beta_{\lambda=0} = \bar{\beta}$), the corresponding variation-explained, ρ_0^2 , will have a (non-degenerate) posterior distribution induced by the posterior distribution over (β, σ^2) .

Variation-explained sparsity summary plots depict the posterior uncertainty of ρ_λ^2 , thus, providing a measure of confidence concerning the predictive goodness of the sparsified vector. In these plots, one often observes that the sparsified variation-explained does not deteriorate “statistically significantly” in the sense that the credible interval for ρ_λ^2 overlaps the posterior mean of the unsparsified variation-explained.

3.2 Excess Error of a Sparsified Linear Predictor

Define the “excess error” of a sparsified linear predictor β_λ as

$$\psi_\lambda = \sqrt{n^{-1} \|\mathbf{X}\beta_\lambda - \mathbf{X}\beta\|^2 + \sigma^2} - \sigma. \quad (24)$$

This metric of model fit, while less widely used than variation-explained, has the virtue of being on the same scale as the response variable. Note that excess error attains a minimum of zero precisely when $\beta_\lambda = \beta$. As with the variation-explained, the excess error is a random variable and so has a posterior distribution. By plotting the mean and 90% credible intervals of the excess error against model size (corresponding to increasing values of λ), one can see at a glance the degree of predictive deterioration incurred by sparsification. Samples of ψ_λ can be obtained analogously to the procedure for producing samples of ρ_λ^2 , but using (24) in place of (23).

3.3 Coefficient Magnitude Plot

In addition to the two previous plots, it is instructive to examine which variables remain in the model at different levels of sparsification, which can be achieved simply by plotting the magnitude of each element of β_λ as λ (hence model size) varies. However, using λ or $\|\beta_\lambda\|_0$ for the horizontal axis can obscure the real impact of the sparsification because the predictive impact of sparsification is nonconstant. That is, the jump from a model of size 7 to one of size 6, for example, may correspond to a negligible predictive impact, while the jump from model of size 3 to a model of size 2 could correspond to considerable predictive deterioration. Plotting the magnitude of the elements of β_λ against the corresponding excess error ψ_λ gives the horizontal axis a more interpretable scale.

3.4 A Heuristic for Reporting a Single Model

The three plots described above achieve a remarkable consolidation of information hidden within the posterior samples of $\pi(\beta, \sigma^2 | \mathbf{Y})$. They relate sparsification of a linear predictor to the associated loss in predictive ability, while keeping the posterior uncertainty in these quantities in clear view. Nonetheless, in many situations one would like a procedure that yields a *single* linear predictor.

For producing a single-model linear summary, we propose the following heuristic: report the sparsified linear predictor corresponding to the smallest model whose 90% ρ_λ^2 credible interval contains $E(\rho_0^2)$. In words, we want the smallest linear predictor whose predictive ability (practical significance) is not statistically different than the full model's.

Certainly, this approach requires choosing a credibility level—there is nothing privileged about the 90% level rather than say the 75% or 95%. However, this is true of alternative methods such as hard thresholding or examination of marginal inclusion probabilities, which both require similar conventional choices to be determined. The DSS model selection heuristic offer a crucial benefit over these approaches, though—it explicitly includes a design matrix of predictors into its very formulation. Standard thresholding rules and methods such as the median probability model approach are instead defined on a one-by-one basis, which does not explicitly account for collinearity in the predictor space. (Recall that both the thresholding rules studied

in Ishwaran and Rao (2005) and the median probability theorems of Barbieri and Berger (2004) restrict their analysis to the orthogonal design situation.)

In the DSS approach to model selection, dependencies in the predictor space appear both in the formation of the posterior and also in the definition of the loss function. In this sense, while the response vector \mathbf{Y} is only “used once” in the formation of the posterior, the design information may be “used twice,” both in defining the posterior and also in defining the loss function. Note that this is reasonable in the sense that the model is conditional on \mathbf{X} in the first place. Note also that the DSS loss function may be based on a predictor matrix *different* than the one used in the formation of the posterior.

Example: U.S. Crime Dataset ($p = 15$, $n = 47$)

The U.S. crime data of Vandaele (1978) appears in Raftery, Madigan, and Hoeting (1997) and Clyde, Ghosh, and Littman (2011) among others. The dataset consists of $n = 47$ observations on $p = 15$ predictors. As in earlier analyses we log transform all continuous variables. We produce DSS selection summary plots for three different priors: (i) the horseshoe prior, (ii) the robust prior of Bayarri et al. (2012) with uniform model probabilities, and (iii) a g -prior with $g = n$ and model probabilities as suggested in Scott and Berger (2006). With $p = 15 < 30$, we are able to evaluate marginal likelihoods for all models under the model selection priors (ii) and (iii).

We use these particular priors not to endorse them, but merely as representative examples of widely used specifications.

Figures 1 and 2 show the resulting DSS plots under each prior. Notice that with this dataset the prior choice has an impact; the resulting posteriors for ρ^2 are quite different. For example, under the horseshoe prior we observe a significantly larger amount of shrinkage, leading to a posterior for ρ^2 that concentrates around smaller values as compared to the results in Figure 2. Despite this difference, a conservative reading of the plots would lead to the same conclusion in either situation: the 7-variable model is essentially equivalent (in both suggested metrics, ρ^2 and ψ) to the full model.

To use these plots to produce a *single* sparse linear predictor for the purpose of data summary, we employ the heuristic described in Section 3.4. Table 1 compares the resulting summary to the model chosen according to the median probability model criterion. Notably, the DSS heuristic yields the same 7-variable model under all three choices of prior. In contrast, the HPM is the full model, while the MPM gives either an 11-variable or a 7-variable model depending on which prior is used. Both the HPM and MPM under the robust prior choice would include variables with low statistical and practical significance.

Notice also that the MPM under the robust prior contains four variables with marginal inclusion probabilities near 1/2. The precise numerics of these quantities is highly prior dependent and sensitive to search methods when enumeration is not possible. Accordingly, the MPM model in this case is highly unstable. By focusing on metrics more closely related to practical significance, the DSS heuristic provides more stable selection, returning the same 7-variable model under all prior specifications in this example. As such, this dataset provides a clear example of statistical significance—as evaluated by standard posterior

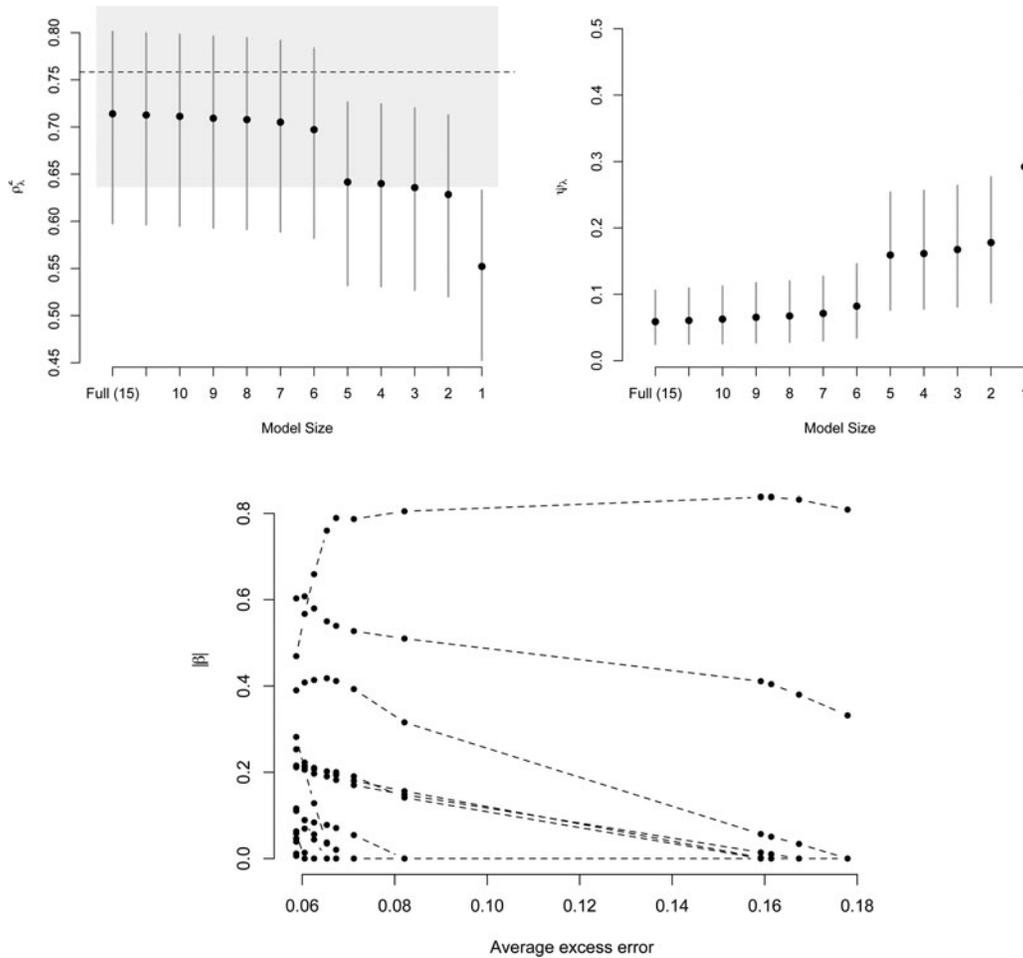


Figure 1. U.S. crime data: DSS plots under the horseshoe prior.

quantities—overwhelming practical relevance. The summary provided by a selection summary plot makes an explicit distinction between the two notions of relevance, providing a clear sense of the predictive cost associated with dropping a predictor.

Finally, notice that there is no evidence of “double-shrinkage.” That is, one might suppose that DSS penalizes coefficients twice, once in the prior and again in the sparsification process, leading to unwanted attenuation of large signals. However, double-shrinkage would not occur if the ℓ_0 penalty were being applied exactly, so any unwanted attenuation is attributable to the imprecision of the surrogate optimization in (21). In practice, we observe that the adaptive lasso-based approximation exhibits minimal evidence of double-shrinkage. Figure 3 displays the resulting values of β_λ in the U.S. crime example plotted against the posterior mean (under the horseshoe prior). Notice that, moving from larger to smaller models, no double-shrinkage is apparent. In fact, we observe reinflation or “unshrinkage” of some coefficients as one progresses to smaller models, as might be expected under the ℓ_0 norm.

Example: Diabetes Dataset ($p = 10, n = 447$)

The diabetes data was used to demonstrate the `lars` algorithm in Efron et al. (2004). The data consist of $p = 10$ baseline measurements on $n = 442$ diabetic patients; the response vari-

able is a numerical measurement of disease progression. As in Efron et al. (2004), we work with centered and scaled predictor and response variables. In this example we only used the robust prior of Bayarri et al. (2012). The goal is to focus on the sequence in which the variables are included and to illustrate how DSS provides an attractive alternative to the median probability model.

Table 2 shows the variables included in each model in the DSS path up to the 5-variable model. The DSS plots in this example (omitted here) suggest that this should be the largest model under consideration. The table also reports the median probability model.

Notice that marginal inclusion probabilities do not necessarily offer a good alternative to rank variable importance, particularly in cases where the predictors are highly colinear. This is evident in the current example in the “dilution” of inclusion probabilities of the variables with the strongest dependencies in this dataset: TC, LDL, HDL, TCH, and LTG. It is possible to see the same effect in the rank of high probability models, as most models on the top of the list represent distinct combinations of correlated predictors. In the sequence of models from DSS, variables LTG and HDL are chosen as the representatives for this group.

Meanwhile, a variable such as Sex appears with marginal inclusion probability of 0.98, and yet its removal from DSS

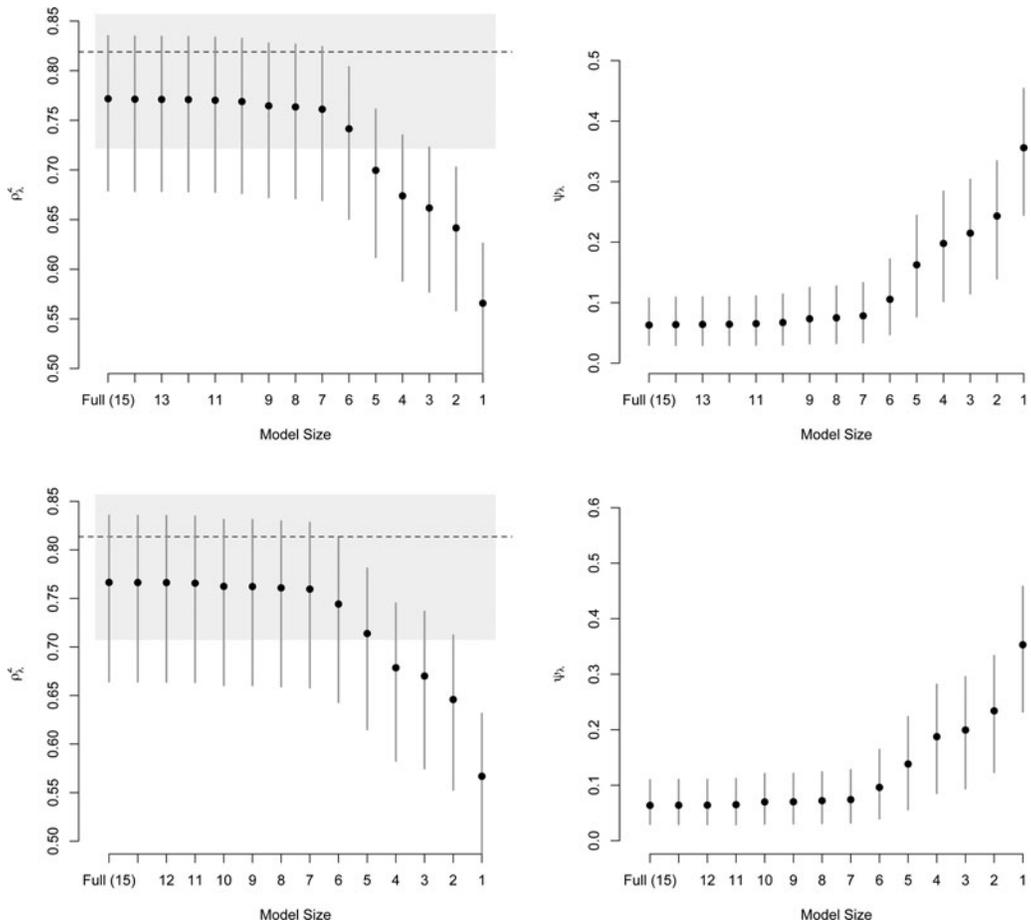


Figure 2. U.S. crime data: DSS plot under the “robust” prior of Bayarri et al. (2012) (top row) and under a g -prior with $g = n$ (bottom row). All 2^{15} models were evaluated in this example.

(five-variable) leads to only a minor decrease in the model’s predictive ability. Thus the diabetes data offer a clear example where statistical significance can overwhelm practical relevance if one looks only at standard Bayesian outputs. The summary provided by DSS makes a distinction between the two notions of relevance, providing a clear sense of the predictive cost associated with dropping a predictor.

Example: Protein Activation Dataset ($p = 88, n = 96$)

The protein activity dataset is from Clyde, Ghosh, and Littman (2011). This example differs from the previous example in that with $p = 88$ predictors, the model space can no longer be exhaustively enumerated. In addition, correlation between the potential predictors is as high as 0.99, with 17 pairs of variables having correlations above 0.95. For this example, the horseshoe prior and the robust prior are considered. To search the model space, we use a conventional Gibbs sampling strategy as in Garcia-Donato and Martinez-Beneito (2013, Appendix A), based on George and McCulloch (1997).

Figure 4 shows the DSS plots under the two priors considered. Once again, the horseshoe prior leads to smaller estimates of ρ^2 . And once again, despite this difference, the DSS heuristic returns the same six predictors under both priors. On this dataset,

the MPM under the Gibbs search (as well as the HPM and MPM given by BAS) coincide with the DSS summary model.

Example: Protein Activation Dataset ($p = 88, n = 80$)

To explore the behavior of DSS in the $p > n$ regime, we modify the previous example by randomly selecting a subset of $n = 80$ observations from the original dataset. These 80 observations are used to form our posterior distribution. To define the DSS summary, we take $\tilde{\mathbf{X}}$ to be the entire set of 96 predictor values. For simplicity we only use the robust model selection prior. Figure 5 shows the results; with fewer observations, smaller models do not give up as much in the ρ^2 and ψ scales as the original example. A conservative read of the DSS plots leads to the same 6-variable model, however, in this limited information situation, the models with 5 or 4 variables are competitive. One important aspect of Figure 5 is that even working in the $p > n$ regime, DSS is able to evaluate the predictive performance and provide a summary of models of any dimension up to the full model. Because the robust prior was used, posterior model probabilities are defined only for models of dimension less than $n - 1$. Despite this, the DSS optimization problem is well defined as long as the number of points in \mathbf{X} is greater than p . If the dataset itself has fewer than p unique points, one

Table 1. Selected models by different methods in the U.S. crime example

	DSS-HS	DSS-Robust	DSS-g-prior	MPM(Robust)	MPM(g-prior)	HS(th)	<i>t-stat</i>
M	•	•	•	0.89	0.85	•	•
So	—	—	—	0.39	0.27	—	—
Ed	•	•	•	0.97	0.96	•	•
Po1	•	•	•	0.71	0.68	•	—
Po2	—	—	—	0.52	0.45	•	—
LF	—	—	—	0.36	0.22	—	—
M.F	—	—	—	0.38	0.24	—	—
Pop	—	—	—	0.51	0.40	—	—
NW	•	•	•	0.77	0.70	•	•
U1	—	—	—	0.39	0.27	—	—
U2	•	•	•	0.71	0.63	—	•
GDP	—	—	—	0.52	0.39	—	—
Ineq	•	•	•	0.99	0.99	•	•
Prob	•	•	•	0.91	0.88	•	•
Time	—	—	—	0.52	0.40	—	—
R^2_{MLE}	82.6%	82.6%	82.6%	85.4%	82.6%	80.0%	69.0%

NOTE: The MPM column displays marginal inclusion probabilities with the numbers in bold associated with the variables included in the median probability model. The HS(th) column refers to the hard thresholding in Section 1.5 under the horseshoe prior. The *t-stat* column is the model defined by OLS *p*-values smaller that 0.05. The R^2_{mle} row reports the traditional in-sample percentage of variation-explained of the least-squares fit based on only the variables in a given column.

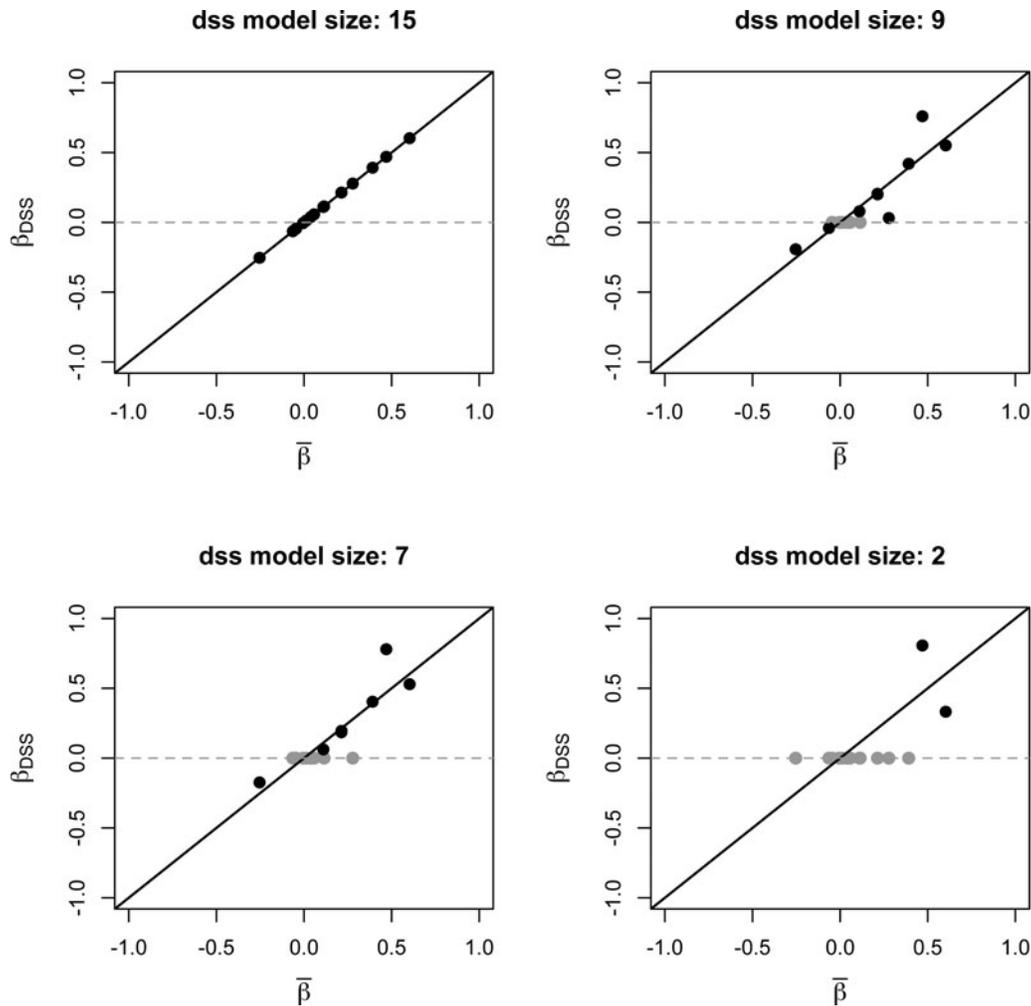


Figure 3. U.S. Crime data under the horseshoe prior: $\bar{\beta}$ refers to the posterior mean while β_{DSS} is the value of β_λ under different values of λ such that different number of variables are selected.

Table 2. Selected models by DSS and model selection prior in the diabetes example

	DSS-Robust(5)	DSS-Robust(4)	DSS-Robust(3)	DSS-Robust(2)	DSS-Robust(1)	MPM(Robust)	<i>t-stat</i>
Age	—	—	—	—	—	0.08	—
Sex	•	—	—	—	—	0.98	•
BMI	•	•	•	•	•	0.99	•
MAP	•	•	•	—	—	0.99	•
TC	—	—	—	—	—	0.66	•
LDL	—	—	—	—	—	0.46	—
HDL	•	•	—	—	—	0.51	—
TCH	—	—	—	—	—	0.26	—
LTG	•	•	•	•	—	0.99	•
GLU	—	—	—	—	—	0.13	—
R^2_{MLE}	50.8%	49.2%	48.0%	45.9%	34.4%	51.3%	50.0%

NOTE: The MPM column displays marginal inclusion probabilities, and the numbers in bold are associated with the variables included in the median probability model. The *t-stat* column is the model defined by OLS *p*-values smaller that 0.05. The R^2_{MLE} row reports the traditional in-sample percentage of variation-explained of the least-squares fit based on only the variables in a given column.

may specify additionally representative points at which to make predictions.

DISCUSSION

A detailed examination of the previous literature reveals that sparsity can play many roles in a statistical analysis—model selection, strong regularization, and improved computation, for example. A central, but often implicit, virtue of sparsity is that human beings find fewer variables easier to think about.

When one desires sparse model summaries for improved comprehensibility, prior distributions are an unnatural vehicle for furnishing this bias. Instead, we describe how to use a decision theoretical approach to induce sparse posterior model summaries. Our new loss function resembles the popular penalized likelihood objective function of the lasso estimator, but its interpretation is very different. Instead of a regularizing tool for estimation, our loss function is a posterior summarizer with an explicit parsimony penalty. To our knowledge this is the first such loss function to be proposed in this capacity. Conceptually,

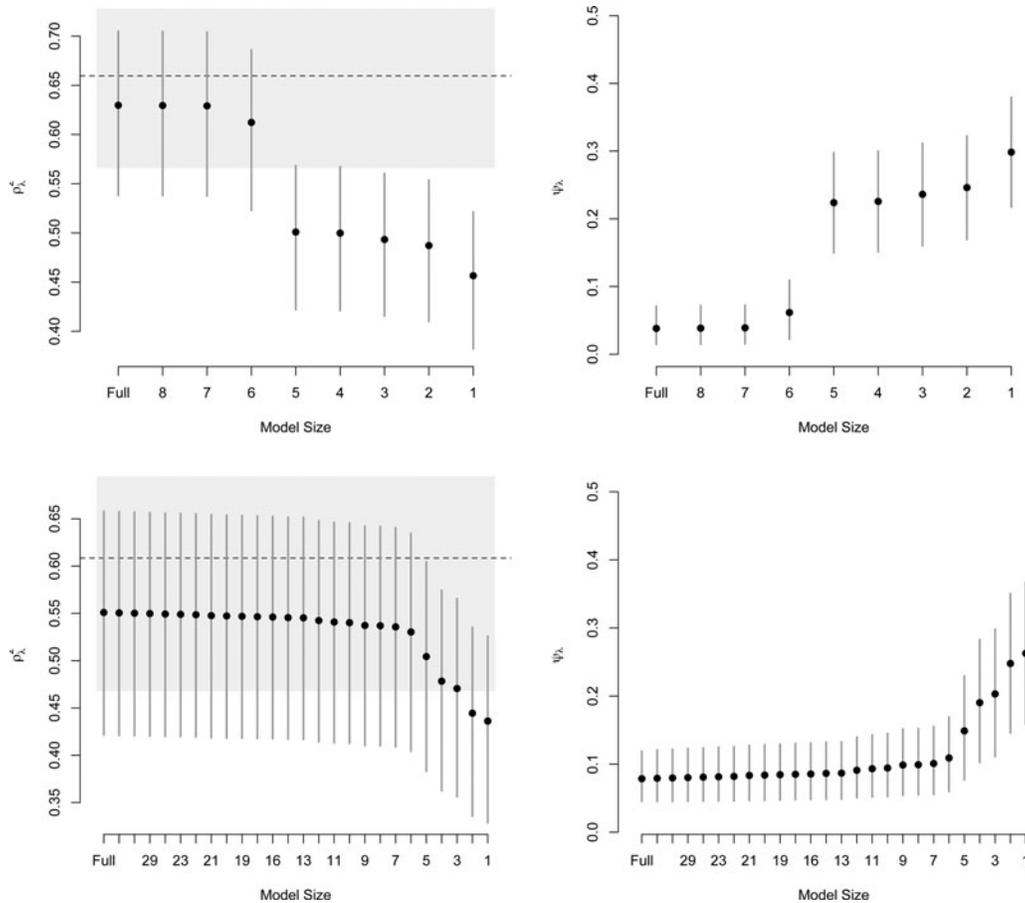


Figure 4. Protein activation data: DSS plots under model selection priors (top row) and under shrinkage priors (bottom row).

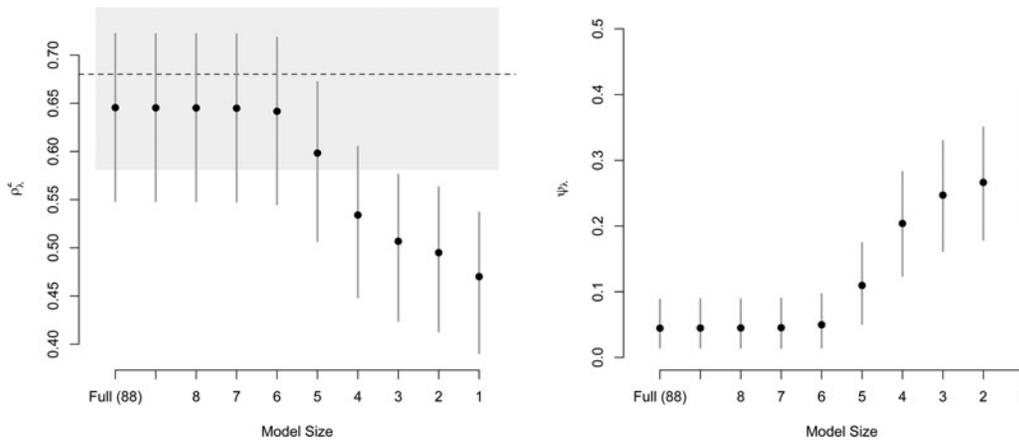


Figure 5. Protein activation data ($p > n$ case): DSS plots under model selection priors.

its nearest forerunner would be HPD regions, which summarize a posterior density while satisfying a compactness constraint.

Unlike hard thresholding rules, our selection summary plots convey posterior uncertainty associated with the provided sparse summaries. In particular, posterior correlation between the elements of β impacts the posterior distribution of the sparsity degradation metrics ρ^2 and ψ . While the DSS summary plots do not “automate” the problem of determining λ (and hence β_λ), they do manage to distill the posterior distribution into a graphical summary that reflects the posterior uncertainty in the predictive degradation due to sparsification. Furthermore, they explicitly integrate information about the possibly nonorthogonal design space in ways that standard thresholding rules and marginal probabilities do not.

As a summary device, these plots can be used in conjunction with whichever prior distribution is most appropriate to the applied problem under consideration. As such, they complement recent advances in Bayesian variable selection and shrinkage estimation and will benefit from future advances in these areas.

We demonstrate how to apply the summary selection concept to logistic regression and Gaussian graphical models in a brief appendix.

APPENDIX A: EXTENSIONS

A.1 Selection Summary in Logistic Regression

Selection summary can be applied outside the realm of normal linear models as well. This section explicitly shows how to extend the approach to logistic regression and provides an illustration on real data.

Although one has many choices for judging predictive accuracy, it is convenient to note that squared prediction loss is precisely the negative log-likelihood in the normal linear model setting, which suggests the following generalization of (16):

$$\mathcal{L}(\tilde{Y}, \gamma) = \lambda \|\gamma\|_0 - n^{-1} \log[f(\tilde{Y}, \mathbf{X}, \gamma)], \quad (\text{A.1})$$

where $f(\tilde{Y}, \gamma)$ denotes the likelihood of \tilde{Y} with “parameters” γ .

In the case of a binary outcome vector using a logistic link function, the generalized DSS loss becomes

$$\mathcal{L}(\tilde{Y}, \gamma) = \lambda \|\gamma\|_0 + n^{-1} \sum_{i=1}^n (\tilde{Y}_i \mathbf{X}_i \gamma - \log(1 + \exp(\mathbf{X}_i \gamma))). \quad (\text{A.2})$$

Taking expectations yields

$$\mathcal{L}(\tilde{Y}, \bar{\pi}) = \lambda \|\gamma\|_0 + n^{-1} \sum_{i=1}^n (\bar{\pi}_i \mathbf{X}_i \gamma - \log(1 + \exp(\mathbf{X}_i \gamma))), \quad (\text{A.3})$$

where $\bar{\pi}_i$ is the posterior mean probability that $\tilde{Y}_i = 1$. To help interpret this formula, note that it can be rewritten as a weighted logistic regression as follows. For each observed X_i , associate a pair of pseudo-responses $Z_i = 1$ and $Z_{i+n} = 0$ with weights $w_i = \bar{\pi}_i$ and $w_{i+n} = 1 - \bar{\pi}_i$ respectively. Then $\bar{\pi}_i \mathbf{X}_i \gamma - \log(1 + \exp(\mathbf{X}_i \gamma))$ may be written as

$$[w_i Z_i \mathbf{X}_i \gamma - w_i \log(1 + \exp(\mathbf{X}_i \gamma))] + [w_{i+n} Z_{i+n} \mathbf{X}_i \gamma - w_{i+n} \log(1 + \exp(\mathbf{X}_i \gamma))]. \quad (\text{A.4})$$

Thus, optimizing the DSS logistic regression loss is equivalent to finding the penalized maximum likelihood of a weighted logistic regression where each point in predictor space has a response $Z_i = 1$, given weight $\bar{\pi}_i$, and a counterpart response $Z_i = 0$, given weight $1 - \bar{\pi}_i$. The observed data determines $\bar{\pi}_i$ via the posterior distribution. As before, if we replace (A.3) by the surrogate ℓ_1 norm

$$\mathcal{L}(\tilde{Y}, \bar{\pi}) = \lambda \|\gamma\|_1 + n^{-1} \sum_{i=1}^n (\bar{\pi}_i \mathbf{X}_i \gamma - \log(1 + \exp(\mathbf{X}_i \gamma))), \quad (\text{A.5})$$

then an optimal solution can be computed via the R package GLMNet (Friedman, Hastie, and Tibshirani 2010).

The DSS summary selection plot may be adapted to logistic regression by defining the excess error as

$$\psi_\lambda = \sqrt{n^{-1} \sum_i \pi_i - 2\pi_{\lambda,i} \pi_i + \pi_{\lambda,i}^2} - \sqrt{n^{-1} \sum_i \pi_i (1 - \pi_i)}, \quad (\text{A.6})$$

where π_i is the probability that $\tilde{y}_i = 1$ given the true model parameters, and $\pi_{\lambda,i}$ is the corresponding quantity under the λ -sparsified model. This expression for the logistic excess error relates to the linear model case in that each expression can be derived from

$$\psi_\lambda = \sqrt{n^{-1} \text{E}(\|\tilde{Y} - \hat{Y}_\lambda\|^2)} - \sqrt{n^{-1} \text{E}(\|\tilde{Y} - \text{E}(\tilde{Y})\|^2)}, \quad (\text{A.7})$$

where the expectation is with respect to the predictive distribution of \tilde{Y} conditional on the model parameters, and \hat{Y}_λ denotes the optimal λ -sparse prediction. In particular, $\hat{Y}_\lambda \equiv \mathbf{X} \beta_\lambda$ for the linear model and $\hat{y}_{\lambda,i} \equiv \pi_{\lambda,i} = (1 + \exp(-X_i \beta_\lambda))^{-1}$ for the logistic regression model. One notable difference between the expressions for excess error under the linear model and the logistic model is that the linear model has constant variance whereas the variance term depends on the predictor point in the logistic model as a result of the Bernoulli likelihood.

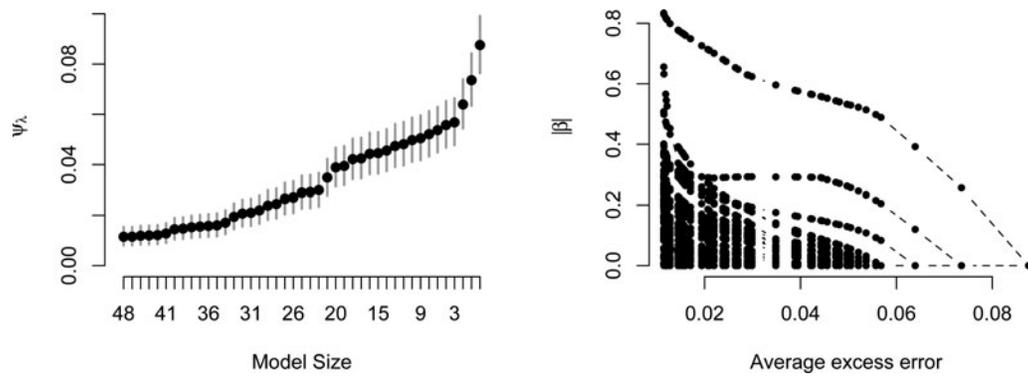


Figure 6. DSS plots for the German credit data. For these data, each included variable seems to add an incremental amount, as the excess error plot builds steadily until reaching the null model with no predictors.

Example: German Credit Data ($n = 1000$, $p = 48$)

To illustrate selection summary in the logistic regression context, we use the German Credit data from the UCI repository, where $n = 1000$ and $p = 48$. In each record we have available covariates associated with a loan applicant, such as credit history, checking account status, car ownership and employment status. The outcome variable is a judgment of whether or not the applicant has “good credit.” A natural objective when analyzing these data would be to develop a good model for assessing creditworthiness of future applicants. A default shrinkage prior over the regression coefficients is used, based on the ideas described in Polson, Scott, and Windle (2013) and the associated R package *BayesLogit*. The DSS selection summary plots (adapted to a logistic regression) are displayed in Figure 6. The plot suggests a high degree of “prevariable selection,” in that all of the predictor variables appear to add an incremental amount of prediction accuracy, with no single predictor appearing to dominate. Nonetheless, several of the larger models (smaller than the full 48 variable model) do not give up much in excess error, suggesting that a moderately reduced model (≈ 35), may suffice in practice. Depending on the true costs associated with measuring those ten least valuable covariates, relative to the cost associated with an increase of 0.01 in excess error, this reduced model may be preferable.

A.2 Selection Summary for Gaussian Graphical Models

Covariance estimation is yet another area where a sparsifying loss function can be used to induce a parsimonious posterior summary.

Consider a $(p \times 1)$ vector $(x_1, x_2, \dots, x_p) = \mathbf{X} \sim N(0, \Sigma)$. Zeros in the precision matrix $\Omega = \Sigma^{-1}$ imply conditional independence among certain dimensions of \mathbf{X} . As sparse precision matrices can be represented through a labelled graph, this modeling approach is often referred to as Gaussian graphical modeling. Specifically, for a graph $G = (V, E)$, where V is the set of vertices and E is the set of edges, let each edge represent a nonzero element of Ω . See Jones et al. (2005) for a thorough overview. This problem is equivalent to finding a sparse representation in p separate linear models for $X_j|X_{-j}$, making the selection summary approach developed above directly applicable.

As with linear models, one has the option of modeling the entries in the precision matrix via shrinkage priors or via selection priors with point masses at zero. Regardless of the specific choice of prior, summarizing patterns of conditional independence favored in the posterior distribution remains a major challenge.

A DSS parsimonious summary can be achieved via a multivariate extension of (16) by once again leveraging the notion of “predictive accuracy” as defined by the negative log-likelihood:

$$\mathcal{L}(\tilde{\mathbf{X}}, \Gamma) = \lambda \|\Gamma\|_0 - \log \det(\Gamma) + \text{tr}(n^{-1} \tilde{\mathbf{X}} \tilde{\mathbf{X}}' \Gamma), \quad (\text{A.8})$$

where Γ represents the decision variable for Ω and $\|\Gamma\|_0$ represents the sum of nonzero entries in off-diagonal elements of Γ . Taking expectations with respect to the posterior predictive of $\tilde{\mathbf{X}}$ yields

$$\mathcal{L}(\Gamma) = E(\mathcal{L}(\tilde{\mathbf{X}}, \Gamma)) = \lambda \|\Gamma\|_0 - \log \det(\Gamma) + \text{tr}(\bar{\Sigma} \Gamma), \quad (\text{A.9})$$

where $\bar{\Sigma}$ represents the posterior mean of Σ .

As before, an approximate solution to the DSS graphical model posterior summary optimization problem can be obtained by employing the surrogate ℓ_1 penalty

$$\mathcal{L}(\Gamma) = E(\mathcal{L}(\tilde{\mathbf{X}}, \Gamma)) = \lambda \|\Gamma\|_1 - \log \det(\Gamma) + \text{tr}(\bar{\Sigma} \Gamma), \quad (\text{A.10})$$

as developed by penalized likelihood methods such as the graphical lasso (Friedman, Hastie, and Tibshirani 2008).

[Received February 2014. Revised November 2014.]

REFERENCES

- Barbieri, M., and Berger, J. (2004), “Optimal Predictive Model Selection,” *The Annals of Statistics*, 32, 870–897. [438,441]
- Bayarri, M., Berger, J., Forte, A., and Garcia-Donato, G. (2012), “Criteria for Bayesian Model Choice With Application to Variable Selection,” *The Annals of Statistics*, 40, 1550–1577. [435,436,441,442]
- Berger, J. (1980), “A Robust Generalized Bayes Estimator and Confidence Region for a Multivariate Normal Mean,” *The Annals of Statistics*, 8, 716–761. [436]
- Berger, J., and Pericchi, L. (2001), *Objective Bayesian Methods for Model Selection: Introduction and Comparison*, volume 38 of *Lecture Notes-Monograph Series*, pp. 135–207, Hayward, CA: Institute of Mathematical Statistics. [436]
- Carvalho, C., Polson, N., and Scott, J. (2010), “The Horseshoe Estimator for Sparse Signals,” *Biometrika*, 97, 465–480. [436,437,438]
- Clyde, M., and George, E. (2004), “Model Uncertainty,” *Statistical Science*, 19, 81–94. [435]
- Clyde, M., Ghosh, J., and Littman, M. (2011), “Bayesian Adaptive Sampling for Variable Selection and Model Averaging,” *Journal of Computational and Graphical Statistics*, 20, 80–101. [436,438,441,443]
- Cui, W., and George, E. I. (2008), “Empirical Bayes vs. Fully Bayes Variable Selection,” *Journal of Statistical Planning and Inference*, 138, 888–900. [436]
- Draper, D. (2013), “Bayesian Model Specification: Heuristics and Examples,” in *Bayesian Theory and Applications*, eds. P. Damien, P. Dellaportas, N. Polson and D. Stephens, New York: Oxford University Press, pp. 409–431. [438]
- Efron, B., Hastie, T., Johnstone, I., and Tibshirani, R. (2004), “Least Angle Regression,” *The Annals of Statistics*, 32, 407–499. [440,442]
- Fouskakis, D., and Draper, D. (2008), “Comparing Stochastic Optimization Methods for Variable Selection in Binary Outcome Prediction, With Application to Health Policy,” *Journal of the American Statistical Association*, 103, 1367–1381. [438]

- Fouskakis, D., Ntzoufras, I., and Draper, D. (2009), "Bayesian Variable Selection Using Cost-Adjusted BIC, With Application to Cost-Effective Measurement of Quality of Health Care," *Annals of Applied Statistics*, 3, 663–690. [438]
- Friedman, J., Hastie, T., and Tibshirani, R. (2008), "Sparse Inverse Covariance Estimation With the Graphical Lasso," *Biostatistics*, 9, 432–441. [447]
- (2010), "Regularization Paths for Generalized Linear Models via Coordinate Descent," *Journal of Statistical Software*, 33, 1–22. [446]
- Garcia-Donato, G., and Martinez-Beneito, M. (2013), "On Sampling Strategies in Bayesian Variable Selection Problems With Large Model Spaces," *Journal of the American Statistical Association*, 108, 340–352. [437,438,443]
- Gelfand, A. E., Dey, D. K., and Chang, H. (1992), "Model Determination Using Predictive Distributions With Implementation via Sampling-based Methods," Technical Report, DTIC Document. [438]
- George, E., and McCulloch, R. (1993), "Variable Selection via Gibbs Sampling," *Journal of the American Statistical Association*, 88, 881–889. [436]
- (1997), "Approaches for Bayesian Variable Selection," *Statistica Sinica*, 7, 339–373. [436,437,443]
- Geweke, J. (1996), "Variable Selection and Model Comparison in Regression," in *Bayesian Statistics 5*, eds. J. M. Bernardo, J. O. Berger, and A. Smith, New York: Oxford University Press, pp. 609–620. [437]
- Griffin, J., and Brown, P. (2012), "Structuring Shrinkage: Some Correlated Priors for Regression," *Biometrika*, 99, 481–487. [436,437]
- Hans, C. (2009), "Bayesian Lasso Regression," *Biometrika*, 96, 835–845. [436,437]
- (2010), "Model Uncertainty and Variable Selection in Bayesian Lasso Regression," *Statistics and Computing*, 20, 221–229. [436]
- Hans, C., Dobra, A., and West, M. (2007), "Shotgun Stochastic Search in Regression With Many Predictors," *Journal of the American Statistical Association*, 102, 507–516. [436,438]
- Ishwaran, H., and Rao, J. S. (2005), "Spike and Slab Variable Selection: Frequentist and Bayesian Strategies," *The Annals of Statistics*, 33, 730–773. [437,438,441]
- Jeffreys, H. (1961), *Theory of Probability*, Oxford, UK: Oxford University Press. [436]
- Jones, B., Carvalho, C., Dobra, A., Hans, C., Carter, C., and West, M. (2005), "Experiments in Stochastic Computation for High-Dimensional Graphical Models," *Statistical Science*, 20, 388–400. [447]
- Leamer, E. (1978), *Specification Searches: ad hoc Inference With Nonexperimental Data*, Wiley Series in Probability and Mathematical Statistics, New York: Wiley. Available at <http://books.google.com/books?id=sYVYAAAAMAAJ>. [438]
- Liang, F., Paulo, R., Molina, G., Clyde, M., and Berger, J. (2008), "Mixtures of g Priors for Bayesian Variable Selection," *Journal of the American Statistical Association*, 103, 410–423. [435,436]
- Lv, J., and Fan, Y. (2009), "A Unified Approach to Model Selection and Sparse Recovery Using Regularized Least Squares," *The Annals of Statistics*, 37, 3498–3528. [440]
- Maruyama, Y., and George, E. I. (2011), "Fully Bayes Factors With a Generalized g-prior," *The Annals of Statistics*, 39, 2740–2765. [436]
- Mitchell, T. J., and Beauchamp, J. J. (1988), "Bayesian Variable Selection in Linear Regression," *Journal of the American Statistical Association*, 83, 1023–1032. [437]
- O'Hara, R. B., and Sillanpää, M. J. (2009), "A Review of Bayesian Variable Selection Methods: What, How and Which," *Bayesian Analysis*, 4, 85–117. [435]
- Park, T., and Casella, G. (2008), "The Bayesian Lasso," *Journal of the American Statistical Association*, 103, 681–686. [436,437]
- Pericchi, L., and Smith, A. (1992), "Exact and Approximate Posterior Moments for a Normal Location Parameter," *Journal of the Royal Statistical Society, Series B*, 793–804. [436]
- Pericchi, L. R., and Walley, P. (1991), "Robust Bayesian Credible Intervals and Prior Ignorance," *International Statistical Review*, 1–23. [436]
- Polson, N. G., and Scott, J. G. (2012), "Local Shrinkage Rules, Lévy Processes and Regularized Regression," *Journal of the Royal Statistical Society, Series B*, 74, 287–311. [435,437]
- Polson, N. G., Scott, J. G., and Windle, J. (2013), "Bayesian Inference for Logistic Models Using Poly-Gamma Latent Variables," *Journal of the American Statistical Association*, 108, 1339–1349. [447]
- Raftery, A., Madigan, D., and Hoeting, J. (1997), "Bayesian Model Averaging for Linear Regression Models," *Journal of the American Statistical Association*, 92, 1197–1208. [435,441]
- Scott, J., and Berger, J. (2006), "An Exploration of Aspects of Bayesian Multiple Testing," *Journal of Statistical Planning and Inference*, 136, 2144–2162. [441]
- (2010), "Bayes and Empirical-Bayes Multiplicity Adjustment in the Variable-Selection Problem," *The Annals of Statistics*, 38, 2587–2619. [436]
- Spiegelhalter, D. (1977), "A Test for Normality Against Symmetric Alternatives," *Biometrika*, 64, 415–418. [436]
- Strawderman, W. E. (1971), "Proper Bayes Minimax Estimators of the Multivariate Normal Mean," *The Annals of Mathematical Statistics*, 385–388. [436]
- Tibshirani, R. (1996), "Regression Shrinkage and Selection via the Lasso," *Journal of the Royal Statistical Society, Series B*, 58, 267–288. [436]
- Vandaele, W. (1978), "Participation in Illegitimate Activities," in *Deterrence and Incapacitation*, eds. A. Blumstein, J. Cohen, and D. Nagin, Washington, DC: National Academy of Sciences Press, pp. 270–335. [441]
- West, M. (1987), "On Scale Mixtures of Normal Distributions," *Biometrika*, 74, 646–648. [436]
- 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*, Amsterdam: North-Holland, pp. 233–243. [436]
- Zou, H. (2006), "The Adaptive Lasso and Its Oracle Properties," *Journal of the American Statistical Association*, 101, 1418–1429. [440]
- Zou, H., and Li, R. (2008), "One-Step Sparse Estimates in Nonconcave Penalized Likelihood Models," *The Annals of Statistics*, 36, 1509–1533. [440]