

Dynamic Variable Selection with Spike-and-Slab Process Priors

Veronika Rockova^{*,‡} and Kenichiro McAlinn[†]

Abstract. We address the problem of dynamic variable selection in time series regression with unknown residual variances, where the set of active predictors is allowed to evolve over time. To capture time-varying variable selection uncertainty, we introduce new dynamic shrinkage priors for the time series of regression coefficients. These priors are characterized by two main ingredients: smooth parameter evolutions and intermittent zeroes for modeling predictive breaks. More formally, our proposed *Dynamic Spike-and-Slab (DSS)* priors are constructed as mixtures of two processes: a spike process for the irrelevant coefficients and a slab autoregressive process for the active coefficients. The mixing weights are themselves time-varying and depend on lagged values of the series. Our *DSS* priors are probabilistically coherent in the sense that their stationary distribution is fully known and characterized by *spike-and-slab marginals*. For posterior sampling over dynamic regression coefficients, model selection indicators as well as unknown dynamic residual variances, we propose a *Dynamic SSVS* algorithm based on forward-filtering and backward-sampling. To scale our method to large data sets, we develop a *Dynamic EMVS* algorithm for MAP smoothing. We demonstrate, through simulation and a topical macroeconomic dataset, that *DSS* priors are very effective at separating active and noisy coefficients. Our fast implementation significantly extends the reach of spike-and-slab methods to big time series data.

Keywords: Autoregressive mixture processes, Dynamic sparsity, MAP smoothing, Spike and Slab, Stationarity.

1 Dynamic Sparsity

For dynamic linear modeling with many potential predictors, the assumption of a static generative model with a fixed subset of regressors (albeit with time-varying regressor effects) may be misleadingly restrictive. By obscuring variable selection uncertainty over time, confinement to a single inferential model may lead to poorer predictive performance, especially when the actual effective subset at each time is sparse. The potential for dynamic model selection techniques in time series modeling has been recognized (Frühwirth-Schnatter and Wagner 2010; Groen et al. 2013; Nakajima and West 2013a; Kalli and Griffin 2014; Chan et al. 2012; Kowal et al. 2019). In inflation forecasting, for example, large sets of predictors are available and it is expected that the forecasting model changes over time, not only its coefficients (Koop and Korobilis 2012a; Groen

^{*}Booth School of Business, University of Chicago, 5807 S Woodlawn Ave, Chicago, IL 60637

[†]Fox School of Business, Temple University, 1801 Liacouras Walk, Philadelphia, PA 19122, kenichiro.mcalinn@temple.edu

[‡]The author gratefully acknowledge the support from the James S. Kemper Foundation Faculty Research Fund at the University of Chicago Booth School of Business.

et al. 2013; Kalli and Griffin 2014; Wright 2009). In particular, in recessions we might see distress related factors be effective, while having no predictive power in expansions (Koop and Korobilis 2012a). Motivated by such contexts, we develop a new dynamic shrinkage approach for time series models that exploits time-varying predictive subset sparsity.

We present our approach in the context of dynamic linear models (West and Harrison 1997) (or varying coefficient models with a time effect modifier (Hastie and Tibshirani 1993)) that link a scalar response y_t at time t to a set of p known regressors $\mathbf{x}_t = (x_{t1}, \dots, x_{tp})'$ through the relation

$$y_t = \mathbf{x}_t' \boldsymbol{\beta}_t^0 + \varepsilon_t, \quad t = 1, \dots, T, \quad (1.1)$$

where $\boldsymbol{\beta}_t^0 = (\beta_{t1}^0, \dots, \beta_{tp}^0)'$ is a time-varying vector of regression coefficients and where the innovations ε_t come from $\mathcal{N}(0, v_t)$. The observational variances v_t are assumed to be unknown, where the precisions $\nu_t = 1/v_t$ arise from the following Markov evolution model (Chapter 10.8.2 of West and Harrison 1997)

$$\nu_t = c_t \nu_{t-1} / \delta, \quad \text{where } c_t \sim \mathcal{B}(\delta n_{t-1} / 2, (1 - \delta) n_{t-1} / 2) \quad \text{and} \quad n_t = \delta n_{t-1} + 1 \quad (1.2)$$

with a discount parameter $\delta \in (0, 1]$.

The challenge of estimating the $T \times p$ coefficients in (1.1), with merely T observations, is typically made feasible with a smoothness inducing state-space model that treats $\{\boldsymbol{\beta}_t^0\}_{t=1}^T$ as realizations from a (vector autoregressive) stochastic process $\boldsymbol{\beta}_t^0 = f(\boldsymbol{\beta}_{t-1}^0) + \mathbf{e}_t$ with $\mathbf{e}_t \sim \mathcal{N}(0, \boldsymbol{\Lambda}_t)$ for some $\boldsymbol{\Lambda}_t$ and $f(\cdot)$. Nevertheless, any regression model with a large number of potential predictors will still be vulnerable to overfitting. This phenomenon is perhaps even more pronounced here, where the regression coefficients are forced to be dynamically intertwined. The major concern is that overfitted coefficient evolutions disguise true underlying dynamics and provide misleading representations with poor out-of-sample predictive performance. For long term forecasts, this concern is exacerbated by the proliferation of the state space. As the model propagates forward, the non-sparse state innovation accumulates noise, further hindering the out-of-sample forecast ability. With many potentially irrelevant predictors, seeking sparsity is a natural remedy against the loss of statistical efficiency and forecast ability.

We shall assume that p is potentially very large, where possibly only a small portion of predictors is relevant for the outcome at any given time. Besides time-varying regressor effects, we adopt the point of view that the regressors are allowed to enter and leave the model as time progresses, rendering the subset selection problem ultimately dynamic. This anticipation can be reflected by the following sparsity manifestations in the matrix of regression coefficients $\mathbf{B}_{p \times T}^0 = [\boldsymbol{\beta}_1^0, \dots, \boldsymbol{\beta}_T^0]$: (a) *horizontal sparsity*, where each individual time series $\{\beta_{tj}^0\}_{t=1}^T$ (for $j = 1, \dots, p$) allows for intermittent zeroes for when j^{th} predictor is not a persisting predictor at all times, (b) *vertical sparsity*, where only a subset of coefficients $\boldsymbol{\beta}_t^0 = (\beta_{t1}^0, \dots, \beta_{tp}^0)'$ (for $t = 1, \dots, T$) will be active at the t^{th} snapshot in time.

This problem has been addressed in the literature by multiple authors including, for example, Groen et al. (2013); Belmonte et al. (2014); Koop and Korobilis (2012b);

Kalli and Griffin (2014); Nakajima and West (2013a). We should like to draw particular attention to the latent threshold process of Nakajima and West (2013a), a related regime switching scheme for either shrinking coefficients exactly to zero or for leaving them alone on their autoregressive path:

$$\beta_{tj} = b_{tj}\gamma_{tj}, \quad \text{where } \gamma_{tj} = \mathbb{I}(|b_{tj}| > d_j), \quad (1.3)$$

$$b_{tj} = \phi_{0j} + \phi_{1j}(b_{t-1j} - \phi_{0j}) + e_t, \quad |\phi_{1j}| < 1, \quad e_t \stackrel{iid}{\sim} \mathcal{N}(0, \lambda_1). \quad (1.4)$$

The model assumes a latent autoregressive process $\{b_{tj}\}_{t=1}^T$, giving rise to the actual coefficients $\{\beta_{tj}\}_{t=1}^T$ only when it meanders away from a latent basin around zero $[-d_j, d_j]$. This process is reminiscent of a dynamic extension of point-mass mixture priors that exhibit exact zeros (Mitchell and Beauchamp 1988). Other related works include shrinkage approaches towards static coefficients in time-varying models (Frühwirth-Schnatter and Wagner 2010; Bitto and Frühwirth-Schnatter 2019; Lopes et al. 2016). We approach the dynamic sparsity problem through the lens of Bayesian variable selection and develop it further for varying coefficient models. Namely, we assume the traditional spike-and-slab setup by assigning each regression coefficient β_{tj} a mixture prior underpinned by a binary latent indicator γ_{tj} , which flags the coefficient as being either active or inert. While static variable selection with spike-and-slab priors has received a considerable attention (Carlin and Chib 1995; Clyde et al. 1996; George and McCulloch 1993, 1997; Mitchell and Beauchamp 1988; Rockova and George 2014, to name a few), dynamic incarnations are yet to be fully explored (George et al. 2008; Frühwirth-Schnatter and Wagner 2010; Nakajima and West 2013a; Groen et al. 2013). To narrow this gap, this work proposes several new dynamic extensions of popular spike-and-slab priors.

The main thrust of this work is to introduce *Dynamic Spike-and-Slab (DSS) priors*, a new class of time series priors, which induce either smoothness or shrinkage towards zero. These processes are formed as mixtures of two (stationary) time series: one for the active and another for the negligible coefficients. The *DSS* priors pertain closely to the broader framework of mixture autoregressive (*MAR*) processes with a given lag, where the mixing weights are allowed to depend on time. Despite the reported success of *MAR* processes (and variants thereof) for modeling non-linear time series (Wong and Li 2000, 2001; Kalliovirta et al. 2015; Wood et al. 2011), their potential as dynamic sparsity inducing priors has been unexplored. Here, we harness this potential within a dynamic variable selection framework. One feature of stationary variants of our *DSS* priors, that sets it apart from the latent threshold model, is that it yields benchmark continuous spike-and-slab priors (such as the Spike-and-Slab LASSO of Rockova 2018) as its *marginal stationary distribution*. This property guarantees marginal stability in the selection/shrinkage dynamics and probabilistic coherence. Non-stationary variants with a random walk slab process are also possible within our framework.

For efficient posterior sampling under the Gaussian spike-and-slab process, we develop *Dynamic SSVS*, a new extension of SSVS of George and McCulloch (1993) for time series regression with closed-form forward-smoothing and backward-sampling updates (Frühwirth-Schnatter 1994). To scale our method to big data settings, we then develop a MAP smoother called *Dynamic EMVS*, a time series incarnation of EMVS originally conceived for static regression (Rockova and George 2014). Dynamic EMVS

is very fast and uses closed-form updates for both mean and variance parameters. We also consider Laplace spike distributions and turn these mixture processes into dynamic penalty constructs. We formalize the notion of prospective and retrospective shrinkage through doubly adaptive shrinkage terms that pull together past, current, and future information. We introduce asymmetric dynamic thresholding rules –extensions of existing rules for static symmetric regularizers (Fan and Li 2001; Antoniadis and Fan 2001)– to characterize the behavior of joint posterior modes for MAP smoothing. For calculations under the Laplace spike, we implement a one-step-late EM algorithm of Green (1990), that capitalizes on fast closed-form one-site updates. Our dynamic penalties can be regarded as natural extensions of the spike-and-slab penalty functions introduced by Rockova (2018) and further developed by Rockova and George (2018).

We demonstrate the effectiveness of our introduced *DSS* priors with a thorough simulation study and a topical macroeconomic application. Both studies highlight the comparative improvements –in terms of inference, forecasting, and computational time– of *DSS* priors over conventional and recent methods in the literature. In particular, the macroeconomic application, using a large number of economic indicators to forecast inflation and infer on underlying economic structures, serves as a motivating example as to why dynamic sparsity is effective, and even necessary, in these contexts.

The paper is structured as follows: Section 2 and Section 2.1 introduce the *DSS* processes and their variants. Sections 3 and 4 introduce Dynamic SSVS and EMVS, respectively. Section 5 develops the penalized likelihood perspective, introducing the prospective and retrospective shrinkage terms. Section 5.3 develops the one-step-late EM algorithm for Spike-and-Slab Fused LASSO MAP smoothing. Section 6 illustrates the MAP smoothing deployment of *DSS* on simulated examples and Section 7 on a macroeconomic dataset. Section 8 concludes with a discussion.

2 Dynamic Spike-and-Slab Priors

In this section, we introduce the class of Dynamic Spike-and-Slab (*DSS*) priors that constitute a coherent extension of benchmark spike-and-slab priors for dynamic selection/shrinkage. We will assume that the p time series $\{\beta_{tj}\}_{t=1}^T$ (for $j = 1, \dots, p$) in (1.1) follow independent and identical *DSS* priors and thereby we suppress the subscript j (for notational simplicity).

We start with a conditional specification of the *DSS* prior. Given a binary indicator $\gamma_t \in \{0, 1\}$, which encodes the spike/slab membership at time t , and a lagged value β_{t-1} , we assume that β_t arises from a mixture of the form

$$\pi(\beta_t | \gamma_t, \beta_{t-1}) = (1 - \gamma_t)\psi_0(\beta_t | \lambda_0) + \gamma_t\psi_1(\beta_t | \mu_t, \lambda_1), \quad (2.1)$$

where

$$\mu_t = \phi_0 + \phi_1(\beta_{t-1} - \phi_0) \quad \text{with} \quad |\phi_1| < 1 \quad (2.2)$$

and

$$P(\gamma_t = 1 | \beta_{t-1}) = \theta_t. \quad (2.3)$$

For Bayesian variable selection, it has been customary to specify a zero-mean spike density $\psi_0(\beta | \lambda_0)$, such that it concentrates at (or in a narrow vicinity of) zero. Regarding the slab distribution $\psi_1(\beta_t | \mu_t, \lambda_1)$, we require that it be moderately peaked around its mean μ_t , where the amount of spread is regulated by a concentration parameter $\lambda_1 > 0$. The conditional *DSS* prior formulation (2.1) generalizes existing *continuous* spike-and-slab priors (George and McCulloch 1993; Ishwaran and Rao 2005; Rockova 2018) in two important ways. First, rather than centering the slab around zero, the *DSS* prior anchors it around an actual model for the *time-varying* mean μ_t . The non-central mean is defined as an autoregressive lag polynomial of the first order with hyper-parameters (ϕ_0, ϕ_1) . While our framework can be extended to higher-order autoregressive polynomials where μ_t may also depend on values older than β_{t-1} , we outline our method for the first-order autoregression with $\phi_0 = 0$ due to its ubiquity in practice (Tibshirani et al. 2005; West and Harrison 1997; Prado and West 2010). The autoregressive parameter ϕ_1 will be treated as unknown and estimated.

It is illuminating to view the conditional prior (2.1) as a “multiple shrinkage” prior (George 1986b,a) with two shrinkage targets: (1) zero (for the gravitational pull of the spike), and (2) μ_t (for the gravitational pull of the slab). It is also worthwhile to emphasize that the spike distribution $\psi_0(\beta_t | \lambda_0)$ *does not depend on* β_{t-1} , only the slab does. The *DSS* formulation thus induces separation of regression coefficients into two groups, where only the *active* ones are assumed to walk on an autoregressive path.

The second important generalization is implicitly hidden in the hierarchical formulation of the mixing weights θ_t in (2.3), which casts them as a smoothly evolving process (as will be seen in Section 2.2 below). Before turning to this formulation, we discuss several special cases of *DSS* priors.

2.1 Spike and Slab Pairings

One possible choice of the spike distribution is the Laplace density $\psi_0(\beta | \lambda_0) = \frac{\lambda_0}{2} e^{-|\beta|\lambda_0}$ (with a relatively large penalty parameter $\lambda_0 > 0$) due to its ability to threshold via sparse posterior modes, as will be elaborated on in Section 5.3. Under the Laplace spike distribution (i.e. conditionally on $\gamma_t = 0$) the series $\{\beta_t\}_{t=1}^T$ is stationary, iid with a marginal density $\psi_0(\beta | \lambda_0)$. Another natural choice, a Gaussian spike, would impose no new computational challenges due to its conditional conjugacy. However, additional thresholding would be required to obtain a sparse representation.

Regarding the slab distribution, we will focus primarily on the Gaussian slab $\psi_1(\beta_t | \mu_t, \lambda_1)$ (with mean μ_t and variance λ_1) due to its ability to smooth over past/future values. Under the Gaussian slab distribution, $\{\beta_t\}_{t=1}^T$ follow a stationary Gaussian *AR*(1) process

$$\beta_t = \phi_0 + \phi_1(\beta_{t-1} - \phi_0) + e_t, \quad |\phi_1| < 1, \quad e_t \stackrel{iid}{\sim} \mathcal{N}(0, \lambda_1), \quad (2.4)$$

whose stationary distribution is characterized by univariate marginals

$$\psi_1^{ST}(\beta_t | \lambda_1, \phi_0, \phi_1) \equiv \psi_1 \left(\beta_t \mid \phi_0, \frac{\lambda_1}{1 - \phi_1^2} \right); \quad (2.5)$$

a Gaussian density with mean ϕ_0 and variance $\frac{\lambda_1}{1-\phi_1^2}$. The availability of this tractable stationary distribution (2.5) is another appeal of the conditional Gaussian slab distribution.

Rather than shrinking to the *vicinity* of the past value, one might like to entertain the possibility of shrinking *exactly* to the past value (Tibshirani et al. 2005) to obtain piece-wise constant reconstructions. Such a property would be appreciated, for instance, in dynamic sparse portfolio allocation models to mitigate transaction costs associated with negligible shifts in the portfolio weights (Irie and West 2016; Brodie et al. 2009; Jagannathan and Ma 2003; Puelz et al. 2016). This extension has also desirable consequences for h -step ahead forecasting, where $\beta_{j_{t+h}}$ would be prevented from decaying (albeit slowly) over time. One way of attaining the desired effect would be replacing the Gaussian slab $\psi_1(\cdot)$ in (2.1) with a Laplace distribution centered at μ_t , i.e.

$$\psi_1(\beta_t | \mu_t, \lambda_1) = \frac{\lambda_1}{2} e^{-|\beta_t - \mu_t| \lambda_1} \quad (2.6)$$

and by considering $\phi_0 = 0$ and $\phi_1 = 1$. While both the Gaussian and Laplace slab will lead to a conditional posterior mean which shrinks towards the past value, the conditional posterior mode will shrink exactly to the past value for the Laplace (and not the Gaussian). This relates the non-stationary extensions discussed further in Remark 2. A similar effect could be achieved with coefficient specific-autoregressive parameters by allowing for $\phi_{j0} \neq 0$ and $\phi_{j1} = 0$ for $1 \leq j \leq p$ (Lopes et al. 2016).

The stationary Laplace conditional construction (2.6) (with $|\phi_1| < 1$), however, does not imply the Laplace distribution marginally. The univariate marginals are defined through the characteristic function given in (2.7) of Andel (1983). The lack of availability of the marginal density in a simple form thwarts the specification of transition weights in our *DSS* framework. There are, however, avenues for constructing an autoregressive process with Laplace marginals, e.g., through the normal-gamma-autoregressive (*NGAR*) process by Kalli and Griffin (2014). We define the following Laplace autoregressive (*LAR*) process as a special case.

Definition 1. We define the Laplace autoregressive (*LAR*) process by

$$\beta_t = \sqrt{\frac{\psi_t}{\psi_{t-1}}} \phi_1 \beta_{t-1} + \eta_t, \quad \eta_t \sim \mathcal{N}(0, (1 - \phi_1^2) \psi_t),$$

where $\{\psi_t\}_{t=1}^T$ follow an exponential autoregressive process specified through $\psi_t | \kappa_{t-1} \sim \text{Gamma}(1 + \kappa_{t-1}, \lambda_1^2 / [2(1 - \rho)])$ and $\kappa_{t-1} | \psi_{t-1} \sim \text{Poisson}\left(\frac{\rho}{2(1-\rho)} \lambda_1^2 \psi_{t-1}\right)$ with a marginal distribution $\text{Exp}(\lambda_1^2/2)$.

The *LAR* process exploits the scale-normal-mixture representation of the Laplace distribution, yielding Laplace marginals $\beta_t \sim \tilde{\psi}_{ST}(\beta_t | \lambda_1) \equiv \text{Laplace}(\lambda_1)$. This coherence property can be leveraged within our *DSS* framework as follows. If we replace the slab Gaussian *AR*(1) process in (2.1) with the *LAR* process and deploy $\tilde{\psi}_{ST}(\beta_t | \lambda_1)$ instead of $\psi_{ST}(\beta_t | \lambda_1)$ in (2.8), we obtain a Laplace *DSS* variant with the *Spike-and-Slab LASSO* prior of Rockova (2018) as its marginal distribution (according to Theorem 1).

It is worth pointing out an alternative autoregressive construction with Laplace marginals proposed by Andel (1983), where the following $AR(1)$ scheme is considered.

$$\beta_t = \begin{cases} \phi_1 \beta_{t-1} & \text{with probability } \phi_1^2, \\ \phi_1 \beta_{t-1} + \eta_t & \text{with probability } 1 - \phi_1^2, \end{cases} \quad \text{where } \eta_t \sim \text{Laplace}(\lambda_1). \quad (2.7)$$

The innovations in (2.7) come from a mixture of a point mass at zero, providing an opportunity to settle at the previous value, and a Laplace distribution. Again, by deploying this process in the slab, we obtain the *Spike-and-Slab LASSO* marginal distribution (Rockova 2018). While MCMC implementations can be obtained for the dynamic Spike-and-Slab LASSO method (e.g. embedding the sampler of Kalli and Griffin (2014) within our MCMC approach outlined in Section 3), the slab extensions with Laplace marginals are more challenging for optimization. Throughout the rest of the paper, we thereby focus exclusively on the Gaussian $AR(1)$ slab process. We will, however, consider both a Gaussian spike (Section 3 and 4) and a Laplace spike distribution (Section 5).

2.2 Evolving Inclusion Probabilities

A very appealing feature of *DSS* priors that makes them suitable for dynamic subset selection is the opportunity they afford for obtaining “smooth” spike/slab memberships. Recall that the binary indicators in (2.3) determine which of the spike or slab regimes is switched on at time t , where $P(\gamma_t = 1 | \beta_{t-1}) = \theta_t$. It is desirable that the sequence of slab probabilities $\{\theta_t\}_{t=1}^T$ evolves smoothly over time, allowing for changes in variable importance as time progresses and, at the same time, avoiding erratic regime switching. Because the series $\{\theta_t\}_{t=1}^T$ is a key driver of the sparsity pattern, it is important that it be (marginally) stable and that it reflects all relevant information, including not only the previous value θ_{t-1} , but also the previous value β_{t-1} . Many possible constructions of θ_t could be considered. We turn to the implied stationary distribution as a guide for a principled construction of θ_t .

For our formulation, we introduce a *marginal* importance weight $0 < \Theta < 1$, a scalar parameter which controls the overall balance between the spike and the slab distributions. Given $(\Theta, \lambda_0, \lambda_1, \phi_0, \phi_1)$, the conditional inclusion probability θ_t (or a transition function $\theta(\beta_{t-1})$) is defined as

$$\theta_t \equiv \theta(\beta_{t-1}) = \frac{\Theta \psi_1^{ST}(\beta_{t-1} | \lambda_1, \phi_0, \phi_1)}{\Theta \psi_1^{ST}(\beta_{t-1} | \lambda_1, \phi_0, \phi_1) + (1 - \Theta) \psi_0(\beta_{t-1} | \lambda_0)}. \quad (2.8)$$

The conditional mixing weight θ_t can be interpreted as the posterior probability of classifying the past coefficient β_{t-1} as arriving from the *stationary* slab distribution as opposed to the (stationary) spike distribution. This interpretation reveals how the weights $\{\theta_t\}_{t=1}^T$ proliferate parsimony throughout the process $\{\beta_t\}_{t=1}^T$. Suppose that the past value $|\beta_{t-1}|$ was large, then $\theta(\beta_{t-1})$ will be close to one, signaling that the *current* observation β_t is more likely to be in the slab. The contrary occurs when $|\beta_{t-1}|$ is small, where β_t will be discouraged from the slab because the inclusion weight $\theta(\beta_{t-1})$ will be small (close to zero). Let us also note that the weights in (2.8) are *different from* the

conditional probabilities for classifying β_{t-1} as arising from the *conditional slab* in (2.1). These weights will be introduced later in Section 5.

Now that we have elaborated on all the layers of the hierarchical model, we are ready to formally define the Dynamic Spike-and-Slab Process.

Definition 2. Equations (2.1), (2.2), (2.3) and (2.8) define a *Dynamic Spike-and-Slab Process (DSS)* with parameters $(\Theta, \lambda_0, \lambda_1, \phi_0, \phi_1)$. We will write

$$\{\beta_t\}_{t=1}^T \sim DSS(\Theta, \lambda_0, \lambda_1, \phi_0, \phi_1).$$

The *DSS* process relates to the Gaussian mixture of autoregressive (*GMAR*) process of Kalliovirta et al. (2015), which was conceived as a model for time series data with regime switches. Here, we deploy it as a prior on time-varying regression coefficients within the spike-and-slab framework, allowing for distributions other than Gaussian. The *DSS*, being an instance/elaboration of the *GMAR* process, inherits elegant marginal characterizations (as will be seen below)

The *DSS* construction has a strong conceptual appeal in the sense that its marginal probabilistic structure is fully known. This property is rarely available with conditionally defined non-Gaussian time series models, where not much is known about the stationary distribution beyond just the mere fact that it exists. The *DSS* process, on the other hand, guarantees well behaved stable marginals that can be described through benchmark spike-and-slab priors. The marginal distribution can be used as a prior for the initial vector at time $t = 0$, which is typically estimated with the remaining coefficients. The following theorem is an elaboration of Theorem 1 of Kalliovirta et al. (2015).

Theorem 1. Assume $\{\beta_t\}_{t=1}^T \sim DSS(\Theta, \lambda_0, \lambda_1, \phi_0, \phi_1)$ with $|\phi_1| < 1$. Then $\{\beta_t\}_{t=1}^T$ has a stationary distribution characterized by the following univariate marginal distributions:

$$\pi^{ST}(\beta | \Theta, \lambda_0, \lambda_1, \phi_0, \phi_1) = \Theta \psi_1^{ST}(\beta | \lambda_1, \phi_0, \phi_1) + (1 - \Theta) \psi_0(\beta | \lambda_0), \quad (2.9)$$

where $\psi_1^{ST}(\beta | \lambda_1, \phi_0, \phi_1)$ is the stationary slab distribution (2.5).

Proof. We assume an initial condition $\beta_{t=0} \sim \pi^{ST}(\beta_0 | \Theta, \lambda_0, \lambda_1, \phi_0, \phi_1)$. Recall that the conditional density of β_1 given β_0 can be written as

$$\pi(\beta_1 | \beta_0) = (1 - \theta_1) \psi_0(\beta_1 | \lambda_0) + \theta_1 \psi_1(\beta_1 | \mu_1, \lambda_1). \quad (2.10)$$

From the definition of θ_1 in (2.8), we can write the joint distribution as

$$\pi(\beta_1, \beta_0) = \Theta \psi_1^{ST}(\beta_0 | \lambda_1, \phi_0, \phi_1) \psi_1(\beta_1 | \mu_1, \lambda_1) + (1 - \Theta) \psi_0(\beta_0 | \lambda_0) \psi_0(\beta_1 | \lambda_0).$$

Integrating $\pi(\beta_1, \beta_0)$ with respect to β_0 , we obtain

$$\begin{aligned} \pi(\beta_1) &= \int \pi(\beta_1, \beta_0) d\beta_0 \\ &= \Theta \left[\int_{\beta_0} \psi_1(\beta_1 | \mu_1, \lambda_1) \psi_1\left(\beta_0 \mid \phi_0, \frac{\lambda_1}{1 - \phi_1^2}\right) d\beta_0 \right] + (1 - \Theta) \psi_0(\beta_1 | \lambda_0) \\ &= \Theta \psi_1^{ST}(\beta_1 | \lambda_1, \phi_0, \phi_1) + (1 - \Theta) \psi_0(\beta_1 | \lambda_0). \quad \square \end{aligned}$$

Theorem 1 describes the very elegant property of *DSS* that the univariate marginals of this mixture process are Θ -weighted mixtures of marginals. It also suggests a more general recipe for mixing multiple stationary processes through the construction of mixing weights (2.8).

Remark 1. *For autoregressive polynomials of higher order $h > 1$, the transition weights θ_t could be defined in terms of a multivariate stationary distribution evaluated at the last h values of the process, not only the last one. The marginals of such process could be then characterized in terms of a mixture of multivariate Gaussian distributions (Theorem 1 of Kalliovirta et al. 2015).*

It is tempting to regard Θ as the marginal proportion of nonzero coefficients. Such an interpretation is a bit misleading since the sparsity levels are ultimately determined by the θ_t sequence, which is influenced by the component stationary distributions $\psi_0(\cdot)$ and $\psi_1^{ST}(\cdot)$, in particular by the amount of their overlap around zero. With *continuous* spike-and-slab mixtures considered here, more caution is needed for calibration (Rockova 2018). This issue will be revisited in Section 5. One can nevertheless regard Θ as a global sparsity parameter, as we now show.

Unlike with point-mass spike and slab priors (Mitchell and Beauchamp 1988), which assign prior mass directly on sparse vectors, our prior is continuous where exact sparsity can be achieved through posterior modes (using the Laplace spike) or through thresholding. As with other continuous priors (Bhattacharya et al. 2015; Rockova 2018) one can quantify the “effective dimensionality” defined as the number of coefficients which are large enough to be non-negligible. The availability of the stationary distribution is helpful for understanding the *marginal* prior effective dimensionality at each time t . For example, with the Laplace spike and the Gaussian slab (used throughout Section 5) the (smaller) intersection points $\pm\delta$ between the stationary spike and slab densities satisfy (choosing $\phi_0 = 0$ for simplicity)

$$\delta = \frac{\lambda_1 \lambda_0}{1 - \phi_1^2} - \sqrt{\left(\frac{\lambda_1 \lambda_0}{1 - \phi_1^2}\right)^2 - \frac{2\lambda_1}{1 - \phi_1^2} \log\left(\frac{1 - \Theta}{\Theta} \sqrt{\frac{2\pi\lambda_1}{1 - \phi_1^2}} \frac{\lambda_0}{2}\right)}.$$

Defining $\gamma(\beta) = \mathbb{I}(|\beta| > \delta)$ as the indicator for whether or not the coefficient is important, one obtains

$$\mathbb{P}[\gamma(\beta) = 1] = \Theta \left[2 \left(1 - \Phi\left(\delta, 0, \frac{\lambda_1}{1 - \phi_1^2}\right) \right) + \frac{1}{\lambda_0} \phi\left(\delta, 0, \frac{\lambda_1}{1 - \phi_1^2}\right) \right], \tag{2.11}$$

where $\Phi(x, \mu, \sigma^2)$ and $\phi(x, \mu, \sigma^2)$ are the cumulative distribution function and the density of the Gaussian distribution with mean μ and variance σ^2 . Standard Gaussian tail bounds yield $\mathbb{P}[\gamma(\beta) = 1] < \Theta \left[2 + \frac{1}{\lambda_0 \sqrt{2\pi\lambda_1/(1-\phi_1^2)}} \right]$, from which one deduces that the parameter Θ takes the role of a *global sparsity parameter*. Defining the *effective dimensionality* at time t as $|\gamma(\beta_t)| = \sum_{j=1}^p \gamma(\beta_{tj})$, it is desirable that $|\gamma(\beta_t)|$ accumulates roughly around the true dimensionality $p_t = \sum_{j=1}^p \mathbb{I}[\beta_{tj}^0 \neq 0]$. Using the Chernoff bound for binomial random variables, one obtains

$$\mathbb{P}(|\gamma(\beta_t)| > C p_t) \leq \exp(-p_t C \log 2) \quad \text{when} \quad \Theta \leq C_1 p_t/p \quad \text{for} \quad C_1 > 0$$

and for any $C > 2C_1 e \left[2 + \frac{1}{\lambda_0 \sqrt{2\pi\lambda_1/(1-\phi_1^2)}} \right]$. This means that as long as the parameter Θ does not overshoot the true sparsity proportion, the prior will concentrate on small subsets up to a constant multiple of the true model size. This property will be ultimately reflected in the posterior. Expressions similar to (2.11) can also be obtained for a Laplace mixture ($\mathbb{P}[\gamma(\beta) = 1] = 2\Theta \left(\frac{1}{\lambda_1} + \frac{1}{\lambda_0} \right) \psi_1(\delta|\lambda_1)$ where $\delta = \frac{1}{\lambda_0 - \lambda_1} \log \left(\frac{1-\Theta}{\Theta} \frac{\lambda_0}{\lambda_1} \right)$ is the (positive) intersection point between Θ -weighted Laplace densities with penalties $\lambda_1 < \lambda_0$) and a Gaussian mixture ($\mathbb{P}[\gamma(\beta) = 1] \leq 2\Theta (\sqrt{2\pi\lambda_1} + \sqrt{2\pi\lambda_0}) \phi(\delta, 0, \lambda_1)$ where $\delta = \sqrt{2 \log \left(\frac{1-\Theta}{\Theta} \sqrt{\frac{\lambda_1}{\lambda_0}} \right)} / \left(\frac{1}{\lambda_0} - \frac{1}{\lambda_1} \right)$ and where $\lambda_0 < \lambda_1$ are the Gaussian variances).

Remark 2 (Random Walk Extensions). *The definition of DSS transition weights in (2.8) requires stationary distributions under the two spike and slab regimes. It is possible to extend our framework to non-stationary random walk slab process (obtained with $\phi_1 = 1$) by modifying transition weights $\{\theta_t\}_{t=1}^T$. Because the series $\{\theta_t\}_{t=1}^T$ is a key driver of sparsity, it is important that it be stable (not too erratic over time) and that it reflects all relevant information, including not only the previous value θ_{t-1} , but also the previous value β_{t-1} . One viable strategy would be to treat $\{\theta_t\}_{t=1}^T$ as random and relate θ_t to the previous value θ_{t-1} via the conditional beta autoregressive process (Casarin et al. 2012a,b) or a marginal beta autoregressive process (McKenzie 1985). However, the weights may be prone to transitioning too often between the spike/slab states when treated as random. For the random walk extensions, one can set θ_t equal to some deterministic sequence (e.g. as in Nakajima and West 2013b) or to a fixed value $\theta_t = \Theta$ for $1 \leq t \leq T$.*

3 Dynamic SSVS

In this section, we develop an MCMC algorithm (Table 1) for dynamic spike-and-slab priors which can be regarded as the dynamic extension of SSVS of George and McCulloch (1993). The DSS prior specification here departs slightly from our previous setup. The Laplace spike distribution $\psi_0(\beta|\lambda_0) = \lambda_0/2e^{-\lambda_0|\beta|}$ yields sparse posterior modes. Since MCMC ultimately reports the posterior mean (which is non-sparse even under the Laplace prior), we will assume the Gaussian spike to capitalize on its direct conditional conjugacy for posterior updating. In particular, we assume the following spike density for $\lambda_0 \ll \lambda_1$

$$\psi_0(\beta | \lambda_0) = \exp\{-\beta^2/(2\lambda_0)\}/\sqrt{2\pi\lambda_0}. \quad (3.1)$$

This yields the following conditional Gaussian distribution

$$\beta_t | \gamma_t, \beta_{t-1} \sim \mathcal{N}(\gamma_t \mu_t, \gamma_t \lambda_1 + (1 - \gamma_t) \lambda_0)$$

and transition weights θ_t in (2.8) with the Gaussian stationary spike distribution $\psi_0^{ST}(\beta_{t-1}|\lambda_0) = \psi_0(\beta|\lambda_0)$. An extension to the Laplace spike is possible with an additional augmentation step, casting the Laplace distribution as a scale mixture of Gaussians with an exponential mixing distribution (Park and Casella 2008). The MCMC algorithm has a Gibbs structure, sampling iteratively from the conditional posteriors

Algorithm: MCMC algorithm for DSS with a Gaussian spike	
	Initialize γ_{tj} and v_0 for $0 \leq t \leq T$ and $1 \leq j \leq p$ and choose n_0, d_0 .
	Sampling Regression Coefficients
<i>Forward filtering</i>	For $1 \leq t \leq T$ Compute $\mathbf{a}_t = \mathbf{H}_t + \mathbf{\Gamma}_t(\mathbf{m}_{t-1} - \mathbf{H}_t)$. Compute $\mathbf{R}_t = \mathbf{\Gamma}_t \mathbf{C}_{t-1} \mathbf{\Gamma}'_t + \mathbf{W}_t$. Compute $f_t = \mathbf{x}'_t \mathbf{a}_t$. Compute $q_t = \mathbf{x}'_t \mathbf{R}_t \mathbf{x}_t + v_t$ and $e_t = y_t - f_t$. Compute $\mathbf{m}_t = \mathbf{a}_t + \mathbf{A}_t e_t$ and $\mathbf{C}_t = \mathbf{R}_t - \mathbf{A}_t \mathbf{A}'_t q_t$ with $\mathbf{A}_t = \mathbf{R}_t \mathbf{x}_t / q_t$.
<i>Backward sampling</i>	Simulate $\beta_T \sim \mathcal{N}(\mathbf{m}_T, \mathbf{C}_T)$. For $t = T-1, \dots, 0$ Compute $\mathbf{a}_T(t-T) = \mathbf{m}_t + \mathbf{B}_t[\beta_{t+1} - \mathbf{a}_{t+1}]$. Compute $\mathbf{R}_T(t-T) = \mathbf{C}_t - \mathbf{B}_t \mathbf{R}_{t+1} \mathbf{B}'_t$, where $\mathbf{B}_t = \mathbf{C}_t \mathbf{\Gamma}'_{t+1} \mathbf{R}_{t+1}^{-1}$. Simulate $\beta_t \sim \mathcal{N}(\mathbf{a}_T(t-T), \mathbf{R}_T(t-T))$.
	Sampling Indicators
	For $j = 1, \dots, p$ Compute $\theta_{tj} = \theta(\beta_{t-1j})$ for $1 \leq t \leq T$ from (2.8). Compute $p_{tj}^* = p_{tj}^*(\beta_{tj})$ for $1 \leq t \leq T$ from (5.5). Compute $p_{0j}^* = \theta(\beta_{0j})$ from (2.8). Sample $\gamma_{tj} \sim \text{Bernoulli}[p_{tj}^*(\beta_{tj})]$ for $0 \leq t \leq T$.
	Sampling Precisions $\nu_t = 1/v_t$
<i>Forward filtering</i> <i>Backward sampling</i>	For $t = 1, \dots, T$ Compute $n_t = \delta n_{t-1} + 1$ and $d_t = \delta d_{t-1} + r_t^2$, where $r_t = y_t - \mathbf{x}'_t \beta_t$. Sample $\nu_T \sim G(n_T/2, d_T/2)$. For $t = 1, \dots, T$ Sample $\eta_{T-t} \sim G[(1-\delta)n_{T-t}/2, d_{T-t}/2]$. Set $\phi_{T-t} = \eta_{T-t} + \delta \phi_{T-t+1}$.

Table 1: An MCMC algorithm with *DSS* priors and a Gaussian spike. Note that $G(a, b)$ denotes a gamma distribution with a mean a/b .

of the regression coefficients $\beta_{0:T}$, latent indicators $\gamma_{0:T}$ and variances $v_{0:T}$ (Frühwirth-Schnatter 1994; West and Harrison 1997, Sect 15.2; Prado and West 2010, Sect 4.5).

For the stationary *DSS* prior, we assume that the autoregressive parameter $|\phi_1| < 1$ is assigned the following beta prior (as in Kim et al. (1998))

$$\pi(\phi_1) \propto \left(\frac{1 + \phi_1}{2}\right)^{a_0-1} \left(\frac{1 - \phi_1}{2}\right)^{b_0-1} \mathbb{I}(|\phi_1| < 1) \quad \text{with } a_0 = 20 \text{ and } b_0 = 1.5, \quad (3.2)$$

implying a prior mean of $2a_0/(a_0 + b_0) - 1 = 0.86$. As was pointed out by Phillips (1991), a non-informative prior on ϕ_1 might result in instability. Zellner (1971) in Chapter 7 recommends a subjective beta prior peaked around one (see also Kastner et al. 2017; Nakajima and West 2013b). Alternatively, Lopes et al. (2016) considered a grid of possible values for ϕ_1 through a discretized Gaussian prior distribution centered at one with a small variance. We will update ϕ_1 with a Metropolis step, using a uniform proposal density on the interval $[0.8, 1]$. While we assume $\phi_0 = 0$ throughout, one can update ϕ_0 in a similar vein. A detailed derivation of the MCMC algorithm is presented in the Appendix (Rockova and McAlinn 2020). In addition to the MCMC algorithm, we also derive MAP smoothers using a penalized likelihood approach.

4 Dynamic EMVS

Unlike previous developments (Nakajima and West 2013a; Kalli and Griffin 2014), this paper also views Bayesian dynamic shrinkage through the lens of optimization. Rather than distilling posterior samples to learn about $\beta_{1:T} = [\beta_1, \dots, \beta_T]$, we focus on finding the MAP trajectory $\hat{\beta}_{1:T} = \arg \max \pi(\beta_{1:T} | \mathbf{y}_{1:T})$. MAP sequence estimation problems (for non-linear non-Gaussian dynamic models) were addressed previously with, e.g., Viterbi-style algorithms (Godsill et al. 2001). Our optimization strategy is conceptually very different and builds on the EMVS procedure of Rockova and George (2014). First, we focus on the Gaussian spike prior variant (3.1) which allows for very fast block updates in closed form.

A (local) posterior mode $\hat{\beta}_{0:T}$ can be obtained indirectly through an EM algorithm, treating Γ and precision parameters $\nu_t = 1/v_t$ as the missing data. The initial vector $\beta_{t=0} = (\beta_{01}, \dots, \beta_{0p})'$ at time $t = 0$ will be estimated together with all the remaining coefficients $\beta_{1:T}$. We assume that β_0 comes from the stationary distribution described in Theorem 1,

$$\pi(\beta_0 | \gamma_0) = \prod_{j=1}^p [\gamma_{0j} \psi_1^{ST}(\beta_{0j} | \lambda_1, \phi_0, \phi_1) + (1 - \gamma_{0j}) \psi_0(\beta_{0j} | \lambda_0)], \quad (4.1)$$

where $\gamma_0 = (\gamma_{01}, \dots, \gamma_{0p})'$ are independent binary indicators with $\mathbb{P}[\gamma_{0j} = 1 | \Theta] = \Theta$ for $1 \leq j \leq p$. Knowing the stationary distribution is thereby useful for specifying the initial conditions. The goal is obtaining the mode $\hat{\beta}_{0:T}$ of the functional $\pi(\beta_{0:T} | \mathbf{y}_{1:T})$. To this end, we proceed iteratively by augmenting this objective function with the missing data $\gamma_{0:T}$, as prescribed by Rockova and George (2014), and then maximizing w.r.t. $\beta_{0:T}$. An important observation, that facilitates the derivation of the algorithm, is that the prior distribution $\pi(\beta_{0:T}, \gamma_{0:T}, \mathbf{v}_{1:T})$ can be factorized into the following products

$$\pi(\beta_{0:T}, \gamma_{0:T}, \mathbf{v}_{1:T}) = \pi(\beta_0 | \gamma_0) \pi(\gamma_0) \prod_{t=1}^T \left[\pi(v_t | v_{t-1}) \prod_{j=1}^p \pi(\beta_{tj} | \gamma_{tj}, \beta_{t-1j}) \pi(\gamma_{tj} | \beta_{t-1j}) \right],$$

where $\pi(\beta_{tj} | \gamma_{tj}, \beta_{t-1j})$ and $\pi(\gamma_{tj} | \beta_{t-1j})$ are defined in (2.1) and (2.3), respectively. For simplicity, we will outline the procedure assuming $\phi_0 = 0$ and thereby $\mu_{tj} = \phi_1 \beta_{t-1j}$. Then, we can write

$$\begin{aligned} & \log \pi(\beta_{0:T}, \gamma_{0:T}, \mathbf{v}_{1:T} | \mathbf{y}_{1:T}) \\ &= C(\mathbf{v}_{1:T}, \phi_1) + \sum_{t=1}^T \sum_{j=1}^p [\gamma_{tj} \log \theta_{tj} + (1 - \gamma_{tj}) \log(1 - \theta_{tj})] \\ & \quad - \sum_{t=1}^T \left\{ \frac{(y_t - \mathbf{x}'_t \beta_t)^2}{2v_t} + \sum_{j=1}^p \left[\gamma_{tj} \frac{(\beta_{tj} - \phi_1 \beta_{t-1j})^2}{2\lambda_1} + (1 - \gamma_{tj}) \frac{\beta_{tj}^2}{2\lambda_0} \right] + \log \pi(v_t | v_{t-1}) \right\} \\ & \quad - \sum_{j=1}^p \left[\gamma_{0j} \frac{\beta_{0j}^2 (1 - \phi_1^2)}{2\lambda_1} + (1 - \gamma_{0j}) \frac{\beta_{0j}^2}{2\lambda_0} - \gamma_{0j} \log \Theta - (1 - \gamma_{0j}) \log(1 - \Theta) \right]. \quad (4.2) \end{aligned}$$

We will endow the parameters $\beta_{0:T}$ with a superscript m to designate their most recent values at the m^{th} iteration. In the E-step, we compute the conditional expectation of (4.2) with respect to the conditional distribution of $[\gamma_{0:T}, \nu_{1:T}]$, given $\beta_{0:T}^{(m)}$ and $\mathbf{y}_{1:T}$. This boils down to computing conditional inclusion probabilities $p_{tj}^* = \text{P}(\gamma_{tj} = 1 | \beta_{tj}^{(m)}, \beta_{t-1j}^{(m)}, \theta_{tj})$ from (A.8), when $t > 0$, and $p_{0j}^* \equiv \theta_{1j} \equiv \theta(\beta_{0j})$ from (2.8), and replacing all the γ_{tj} 's in (4.2) with p_{tj}^* 's. Additionally, one replaces $1/v_t$ with the conditional expectation $\text{E}[\nu_t | \beta_{0:T}, \mathbf{y}_{1:T}]$ available in closed form from the recurrent relations (West and Harrison 1997 on page 364)

$$\text{E}[\nu_t | \beta_{0:T}^{(m)}, \mathbf{y}_{1:T}] = (1 - \delta)n_t/d_t + \delta \text{E}[\nu_{t+1} | \beta_{0:T}^{(m)}, \mathbf{y}_{1:T}] \quad \text{for } 1 \leq t < T,$$

where n_t and d_t are obtained from (A.10) and where $\text{E}[\nu_T | \beta_{0:T}^{(m)}, \mathbf{y}_{1:T}] = n_T/d_T$. In the M-step, we set out to maximize $\text{E}_{\gamma_{0:T}, \nu_{1:T}} \cdot \log \pi(\beta_{0:T}, \gamma_{0:T}, \nu_{1:T} | \mathbf{y}_{1:T})$ w.r.t. $\beta_{0:T}$. This is achieved in a block-wise fashion, where we update β_t given the most recent updates of β_{t-1} and β_{t+1} . Given the conjugacy of the Gaussian distribution, these updates have closed forms (similarly as in the EMVS procedure of Rockova and George 2014). We summarize the steps in the Table 2. It is worth pointing out that the matrix inversion Σ_t^{-1} in the step M1 in Table 2 can be avoided using the fact that $\mathbf{x}_t \mathbf{x}_t'$ is a rank-one matrix. Denote with $D_t = \text{diag}\{\frac{p_{1j}^*}{\lambda_1} + \frac{1-p_{1j}^*}{\lambda_0} + \mathbb{I}(t < T) \frac{\phi_1^2 p_{t+1j}^*}{\lambda_1}\}_{j=1}^p$. Then the Woodbury-Shermann matrix inversion lemma yields

$$\Sigma_t^{-1} = D_t^{-1} - \nu_t^* D_t^{-1} \frac{\mathbf{x}_t \mathbf{x}_t'}{1 + \nu_t^* \mathbf{x}_t' D_t^{-1} \mathbf{x}_t} D_t^{-1}.$$

Due to this trick, the computation of the M-step is extremely fast. Since each update β_t is conditional on all $\beta_j, j \neq t$, we are performing conditional maximization in the spirit of Expectation-Conditional-Maximization (Meng 1993). In order to speed up convergence, we can afford to loop over these simple updates inside each M-step. We found loops of size 100 to perform well.

Additionally, we can estimate the autoregressive parameter ϕ_1 under (a discretized version) of the prior (3.2) by updating ϕ_1 at each iteration with the value that maximizes the expected log-complete posterior $\text{E}_{\gamma_{0:T}, \nu_{1:T}} \cdot \log \pi(\beta_{0:T}, \gamma_{0:T}, \nu_{1:T} | \mathbf{y}_{1:T})$. One can compute this criterion for a grid of values ϕ_1 and pick the one value that maximizes the expected log-complete posterior. Estimation of ϕ_0 can be incorporated in a similar vein.

In the next section, we develop a penalized likelihood approach to MAP smoothing using a Laplace spike prior.

5 Dynamic Spike-and-Slab Penalty

Spike-and-slab priors give rise to self-adaptive penalty functions for MAP estimation, as detailed in Rockova (2018) and Rockova and George (2018). Here, we introduce elaborations for dynamic shrinkage implied by the *DSS* priors.

Algorithm: Dynamic EMVS algorithm	
Initialize β_{tj} for $t = 0, \dots, T$ and $j = 1, \dots, p$.	
E-Step	
For $j = 1, \dots, p$	
E1: Compute mixing weights	Compute $\theta_{tj} = \theta(\beta_{t-1j})$ for $1 \leq t \leq T$ from (2.8). Compute $p_{tj}^* = p_{tj}^*(\beta_{tj})$ for $1 \leq t \leq T$ from (5.5). Compute $p_{0j}^* = \theta(\beta_{0j})$ from (2.8).
E1: Compute precisions	For $t = 1, \dots, T$ Compute $n_t = \delta n_{t-1} + 1$ and $d_t = \delta d_{t-1} + r_t^2$, where $r_t = y_t - \mathbf{x}'_t \beta_t$. Set $\nu_t^* = n_T / d_T$. For $t = T - 1, \dots, 1$ set $\nu_t^* = (1 - \delta)n_t / d_t + \delta \nu_{t+1}^*$.
M-Step: Gaussian spike version	
M1: Compute regression coefficients	For $t = 1, \dots, T$ Compute $\Sigma_t = \nu_t^* \mathbf{x}_t \mathbf{x}'_t + \text{diag}\left\{\frac{p_{tj}^*}{\lambda_1} + \frac{1-p_{tj}^*}{\lambda_0} + \mathbb{I}(t < T) \frac{\phi_1^2 p_{t+1j}^*}{\lambda_1}\right\}_{j=1}^p$ Compute $\boldsymbol{\mu}_t = \nu_t^* y_t \mathbf{x}_t + \frac{\phi_1}{\lambda_1} \beta_{t-1} \odot \mathbf{p}_t^* + \mathbb{I}(t < T) \frac{\phi_1}{\lambda_1} \beta_{t+1} \odot \mathbf{p}_{t+1}^*$ Update $\beta_t = \Sigma_t^{-1} \boldsymbol{\mu}_t$ Compute $\Sigma_0 = \text{diag}\left\{\frac{(1-\phi_1^2)p_{0j}^*}{\lambda_1} + \frac{1-p_{0j}^*}{\lambda_0} + \frac{\phi_1^2 p_{1j}^*}{\lambda_1}\right\}_{j=1}^p$ Update $\beta_0 = \frac{\phi_1}{\lambda_1} \Sigma_0^{-1} \beta_1 \odot \mathbf{p}_1^*$
M-Step: Laplace spike version	
M2: Update regression coefficients	For $j = 1, \dots, p$ and $t = 1, \dots, T$ Compute β_{0j} using (5.18). Compute β_{tj} using (5.17).

Table 2: Dynamic EMVS algorithm for both the Gaussian spike (3.1) and the Laplace spike. The notation $\mathbf{a} \odot \mathbf{b}$ denotes elementwise vector multiplication.

Definition 3. For a given set of parameters $(\Theta, \lambda_0, \lambda_1, \phi_0, \phi_1)$, we define a prospective penalty function implied by (2.1) and (2.8) as follows:

$$\text{pen}(\beta | \beta_{t-1}) = \log[(1 - \theta_t) \psi_0(\beta | \lambda_0) + \theta_t \psi_1(\beta | \mu_t, \lambda_1)]. \quad (5.1)$$

Similarly, we define a retrospective penalty $\text{pen}(\beta_{t+1} | \beta)$ as a function of the second argument β in (5.1). The Dynamic Spike-and-Slab (DSS) penalty is then defined as

$$\text{Pen}(\beta | \beta_{t-1}, \beta_{t+1}) = \text{pen}(\beta | \beta_{t-1}) + \text{pen}(\beta_{t+1} | \beta) + C, \quad (5.2)$$

where $C \equiv -\text{Pen}(0 | \beta_{t-1}, \beta_{t+1})$ is a norming constant.

Remark 3. Note that the dependence on the previous value β_{t-1} in $\text{pen}(\beta | \beta_{t-1})$ is hidden in θ_t and μ_t . Throughout the paper, we will write $\partial\theta_t/\partial\beta_{t-1}$ and $\partial\mu_t/\partial\beta_{t-1}$ without reminding ourselves of this implicit relationship.

As an example, we consider the Laplace spike prior $\psi_0(\beta | \lambda_0) = \lambda_0/2e^{-\lambda_0|\beta|}$. Figure 1 portrays the prospective penalty for two choices of β_{t-1} and two sets of tuning parameters $\phi_1, \lambda_1, \lambda_0$ and Θ (assuming $\phi_0 = 0$). Because the conditional transfer equation (2.1) is a mixture, $\text{pen}(\beta | \beta_{t-1})$ is apt to be multimodal. Figure 1(a) shows an obvious peak at zero (due to the Laplace spike), but also a peak around $\mu_t = 0.9 \times \beta_{t-1}$, prioritizing values in the close vicinity of the previous value (due to the non-central slab). From an implementation viewpoint, however, it is more desirable that the penalty be uni-modal, reflecting the size of the previous coefficient without ambiguity by suppressing one of the peaks. Such behavior is illustrated in Figure 1(b) and Figure 1(c), where the penalty flexibly adapts to $|\beta_{t-1}|$ by promoting either zero or a value close to β_{t-1} . This effect

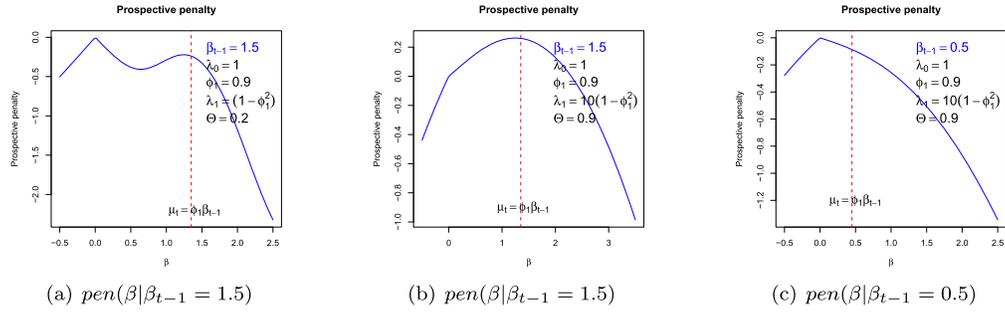


Figure 1: Plots of the prospective penalty function under the Laplace spike.

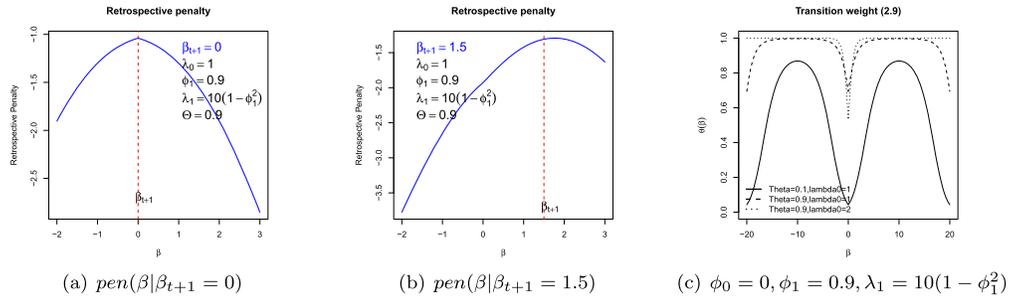


Figure 2: Plots of the retrospective penalty function and the mixing weight (2.8) under the Laplace spike.

is achieved with a relatively large stationary slab variance, such as $\lambda_1/(1 - \phi_1^2) = 10$, a mild Laplace peak $\lambda_0 = 1$ and the marginal importance weight $\Theta = 0.9$. Smaller values Θ would provide an overwhelming support for the zero mode. The parameter Θ , thus should not be regarded as a proportion of active coefficients (as is customary with point-mass mixtures), but rather an interpretation-free tuning parameter.

Figure 1 plots $pen(\beta | \beta_{t-1})$ prospectively as a function of β , given the previous value β_{t-1} . It is also illuminating to plot $pen(\beta_{t+1} | \beta)$ retrospectively as a function of β , given the future value β_{t+1} . Two such retrospective penalty plots are provided in Figure 2(a) and Figure 2(b). When the future value is relatively large ($\beta_{t+1} = 1.5$ in Figure 2(b)), the penalty $pen(\beta_{t+1} | \beta)$ has a peak near β_{t+1} , signaling that the value β_t must be large too. When the future value is small ($\beta_{t+1} = 0$ in Figure 2(a)), the penalty has a peak at zero signaling that the current value β_t must have been small. Again, this balance is achieved with a relatively large stationary slab variance and a large Θ . Note that under the Gaussian spike (3.1), the penalty functions will be differentiable at zero.

The behavior of the prospective and retrospective penalties is ultimately tied to the mixing weight $\theta_t \equiv \theta(\beta)$ in (2.8). It is desirable that $\theta(\beta)$ is increasing with $|\beta|$. However, Laplace tails will begin to dominate for large enough $|\beta|$, where the probability $\theta(\beta)$ will begin to drop (for $|\beta|$ greater than $\delta \equiv (\lambda_0 + \sqrt{2C/A})A$, where $A = \lambda_1/(1 - \phi_1^2)$)

and $C = \log[(1 - \Theta)/\Theta\lambda_0/2\sqrt{2\pi A}]$). However, we can make the turning point δ large enough with larger values Θ and smaller values λ_0 , as indicated in Figure 2(c).

To describe the shrinkage dynamics implied by the penalty (5.2), it is useful to study the partial derivative $\partial Pen(\beta | \beta_{t-1}, \beta_{t+1})/\partial|\beta|$. This term encapsulates how much shrinkage we expect at time t , conditionally on $(\beta_{t-1}, \beta_{t+1})$. We will separate the term into two pieces: a prospective shrinkage effect $\lambda^*(\beta | \beta_{t-1})$, driven by the past value β_{t-1} , and a retrospective shrinkage effect $\tilde{\lambda}^*(\beta | \beta_{t+1})$, driven by the future value β_{t+1} . More formally, we write

$$\frac{\partial Pen(\beta | \beta_{t-1}, \beta_{t+1})}{\partial|\beta|} \equiv -\Lambda^*(\beta | \beta_{t-1}, \beta_{t+1}),$$

where

$$\Lambda^*(\beta | \beta_{t-1}, \beta_{t+1}) = \lambda^*(\beta | \beta_{t-1}) + \tilde{\lambda}^*(\beta | \beta_{t+1}), \quad (5.3)$$

and

$$\lambda^*(\beta | \beta_{t-1}) = -\frac{\partial pen(\beta | \beta_{t-1})}{\partial|\beta|} \quad \text{and} \quad \tilde{\lambda}^*(\beta | \beta_{t+1}) = -\frac{\partial pen(\beta_{t+1} | \beta)}{\partial|\beta|}.$$

5.1 Shrinkage “from the Past”

The prospective shrinkage term $\lambda^*(\beta | \beta_{t-1})$ pertains to Bayesian penalty mixing introduced by Rockova (2018) and Rockova and George (2018) in the sense that it can be characterized as an adaptive linear combination of individual spike and slab shrinkage terms. In particular, we can write

$$\lambda^*(\beta | \beta_{t-1}) = -p_t^*(\beta) \frac{\partial \log \psi_1(\beta | \mu_t, \lambda_1)}{\partial|\beta|} - [1 - p_t^*(\beta)] \frac{\partial \log \psi_0(\beta | \lambda_0)}{\partial|\beta|}, \quad (5.4)$$

where

$$p_t^*(\beta) \equiv \frac{\theta_t \psi_1(\beta | \mu_t, \lambda_1)}{\theta_t \psi_1(\beta | \mu_t, \lambda_1) + (1 - \theta_t) \psi_0(\beta | \lambda_0)}. \quad (5.5)$$

For example, using the Laplace spike, one obtains

$$\lambda^*(\beta | \beta_{t-1}) = p_t^*(\beta) \left(\frac{\beta - \mu_t}{\lambda_1} \right) \text{sign}(\beta) + [1 - p_t^*(\beta)] \lambda_0.$$

Two observations are in order: first, by writing $p_t^*(\beta) = \mathbf{P}(\gamma_t = 1 | \beta_t = \beta, \beta_{t-1}, \theta_t)$, (5.5) can be viewed as a posterior probability for classifying β as arising from the *conditional* slab (versus the spike) at time t , given the previous value β_{t-1} . Second, these weights are *very different from* θ_t in (2.8), which are classifying β as arising from the *marginal* slab (versus the spike). From (5.5), we can see how $p_t^*(\beta)$ hierarchically transmits information about the past value β_{t-1} (via θ_t) to determine the right shrinkage for β_t . This is achieved with a doubly-adaptive chain reaction. Namely, if the previous value β_{t-1} was large, θ_t will be close to one signaling that the next coefficient β_t is prone to be in the slab. Next, if β_t is in fact large, $p_t^*(\beta_t)$ will be close to one, where the first summand

in (5.1) becomes the leading term and shrinks β_t towards μ_t . If β_t is small, however, $p_t^*(\beta_t)$ will be small as well, where the second term in (5.1) takes over to shrink β_t towards zero. This gravitational pull is accelerated when the previous value β_{t-1} was negligible (zero), in which case θ_t will be even smaller, making it even more difficult for the next coefficient β_t to escape the spike. This mechanism explains how the prospective penalty adapts to both (β_{t-1}, β_t) , promoting smooth forward proliferation of spike/slab allocations and coefficients.

5.2 Shrinkage “from the Future”

While the prospective shrinkage term promotes smooth forward proliferation, the retrospective shrinkage term $\tilde{\lambda}^*(\beta | \beta_{t+1})$ operates backwards. For the Laplace spike, we can write

$$\begin{aligned} \tilde{\lambda}^*(\beta | \beta_{t+1}) = & -\frac{\partial \theta_{t+1}}{\partial |\beta|} \left[\frac{p_{t+1}^*(\beta_{t+1})}{\theta_{t+1}} - \frac{1 - p_{t+1}^*(\beta_{t+1})}{1 - \theta_{t+1}} \right] \\ & - p_{t+1}^*(\beta_{t+1}) \phi_1 \text{sign}(\beta) \left[\frac{\beta_{t+1} - \mu_{t+1}}{\lambda_1} \right], \end{aligned} \tag{5.6}$$

where

$$\frac{\partial \theta_{t+1}}{\partial |\beta|} = \theta_{t+1}(1 - \theta_{t+1}) \left[\lambda_0 - \text{sign}(\beta) \left(\frac{\beta - \phi_0}{\lambda_1/(1 - \phi_1^2)} \right) \right]. \tag{5.7}$$

For simplicity, we will write $p_{t+1}^* = p_{t+1}^*(\beta_{t+1})$. Then we have

$$\begin{aligned} \tilde{\lambda}^*(\beta | \beta_{t+1}) = & \left[\lambda_0 - \text{sign}(\beta) \left(\frac{\beta - \phi_0}{\lambda_1/(1 - \phi_1^2)} \right) \right] \left[(1 - p_{t+1}^*)\theta_{t+1} - p_{t+1}^*(1 - \theta_{t+1}) \right] \end{aligned} \tag{5.8}$$

$$- p_{t+1}^* \phi_1 \text{sign}(\beta) \left(\frac{\beta_{t+1} - \mu_{t+1}}{\lambda_1} \right). \tag{5.9}$$

The retrospective term synthesizes information from both (β_{t+1}, β_t) to contribute to shrinkage at time t . When (β_{t+1}, β_t) are both large, we obtain $p_t^*(\beta_{t+1})$ and θ_{t+1} that are both close to one. The shrinkage is then driven by the second summand in (5.9), forcing β_t to be shrunk towards the future value β_{t+1} (through $\mu_{t+1} = \phi_0 + \phi_1(\beta_t - \phi_0)$). When either β_{t+1} or β_t are small, shrinkage is targeted towards the stationary mean through the dominant term (5.8).

5.3 Dynamic Spike-and-Slab Fused LASSO

As we now show, the Laplace spike has the advantage of shrinking coefficient directly to zero, where no additional thresholding is needed for variable selection (Rockova 2018). This has beneficial consequences for computation, where calculations can be narrowed down to active sets of coefficients. In this section, we develop a dynamic coordinate-wise strategy, building on the Spike-and-Slab LASSO method of Rockova and George (2018) for static high-dimensional variable selection.

The key to our approach will be drawing upon the penalized likelihood perspective developed in Section 5. To illustrate the functionality of the dynamic penalty from Section 5, we start by assuming $p = 1$ and $x_t = 1$ in (1.1). This simple case corresponds to a sparse normal-means model, where the means are dynamically intertwined. We begin by characterizing some basic properties of the *conditional* posterior mode

$$\widehat{\boldsymbol{\beta}} = \arg \max_{\boldsymbol{\beta}} \pi(\boldsymbol{\beta} \mid \mathbf{y}, \mathbf{v}),$$

given the variances $\mathbf{v} = (v_1, \dots, v_T)'$, where $\mathbf{y} = (y_1, \dots, y_T)'$ arises from (1.1) and $\boldsymbol{\beta} = (\beta_1, \dots, \beta_T)'$ is assigned the *DSS* prior. One of the attractive features of the Laplace spike in (2.1) is that $\widehat{\boldsymbol{\beta}}$ has a thresholding property. This property is revealed from necessary characterizations for each $\widehat{\beta}_t$ (for $t = 1, \dots, T$), once we condition on the rest of the directions through $(\widehat{\beta}_{t-1}, \widehat{\beta}_{t+1})$. The *conditional* thresholding rule can be characterized using standard arguments, as with similar existing regularizers (Zhang 2010; Fan and Li 2001; Antoniadis and Fan 2001; Zhang and Zhang 2012; Rockova and George 2018). While the typical sparsity-inducing penalty functions are symmetric, the penalty (5.2) is not, due to its dependence on the previous and future values $(\beta_{t-1}, \beta_{t+1})$. Thereby, instead of a single selection threshold, we have two:

$$\Delta^-(x, \beta_{t-1}, \beta_{t+1}) = \sup_{\beta < 0} \left\{ \frac{\beta x^2}{2} - \frac{v_t \text{Pen}(\beta \mid \beta_{t-1}, \beta_{t+1})}{\beta} \right\}, \quad (5.10)$$

$$\Delta^+(x, \beta_{t-1}, \beta_{t+1}) = \inf_{\beta > 0} \left\{ \frac{\beta x^2}{2} - \frac{v_t \text{Pen}(\beta \mid \beta_{t-1}, \beta_{t+1})}{\beta} \right\}. \quad (5.11)$$

The following necessary characterization links the behavior of $\widehat{\boldsymbol{\beta}}$ to the shrinkage terms characterized in Section 5.1 and Section 5.2.

Lemma 1. Denote by $\widehat{\boldsymbol{\beta}} = (\widehat{\beta}_1, \dots, \widehat{\beta}_T)'$ the global mode of $\pi(\boldsymbol{\beta}_{1:T} \mid \mathbf{y}_{1:T}, \mathbf{v}_{1:T})$ and by Δ_t^- and Δ_t^+ the selection thresholds (5.10) and (5.11) with $x = 1$, $\beta_{t-1} = \beta_{t-1}$ and $\beta_{t+1} = \beta_{t+1}$. Then, conditionally on $(\widehat{\beta}_{t-1}, \widehat{\beta}_{t+1})$, we have for $1 < t < T$

$$\widehat{\beta}_t = \begin{cases} 0 & \text{if } \Delta_t^- < y_t < \Delta_t^+ \\ \left[|y_t| - v_t \Lambda^*(\widehat{\beta}_t \mid \widehat{\beta}_{t-1}, \widehat{\beta}_{t+1}) \right]_+ \text{sign}(y_t) & \text{otherwise,} \end{cases} \quad (5.12)$$

where $\Lambda^*(\widehat{\beta}_t \mid \widehat{\beta}_{t-1}, \widehat{\beta}_{t+1})$ was defined in (5.3).

Proof. We begin by noting that $\widehat{\beta}_t$ is a maximizer in t^{th} direction while keeping $(\widehat{\beta}_{t-1}, \widehat{\beta}_{t+1})$ fixed, i.e.

$$\widehat{\beta}_t = \arg \max_{\beta} \left\{ -\frac{1}{2v_t} (y_t - \beta)^2 + \text{Pen}(\beta \mid \widehat{\beta}_{t-1}, \widehat{\beta}_{t+1}) \right\}. \quad (5.13)$$

It turns out that $\widehat{\beta}_t = 0$ iff $\beta \left(y_t - \frac{\beta}{2} + v_t \frac{\text{Pen}(\beta \mid \widehat{\beta}_{t-1}, \widehat{\beta}_{t+1})}{\beta} \right) < 0$, $\forall \beta \in \mathbb{R} \setminus \{0\}$ (Zhang and Zhang 2012). The rest of the proof follows from the definition of Δ_t^+ and Δ_t^- in (5.10) and (5.11). Conditionally on $(\widehat{\beta}_{t-1}, \widehat{\beta}_{t+1})$, the global mode $\widehat{\beta}_t$, once nonzero, has to satisfy (5.12) from the first-order necessary condition. \square

Lemma 1 formally certifies that the posterior mode *under the Laplace spike* exhibits both (a) sparsity and (b) smoothness (through the prospective/retrospective shrinkage terms).

Remark 4. While Lemma 1 assumes $1 < t < T$, the characterization applies also for $t = 1$, once we specify the initial condition $\beta_{t=0}$. The value $\beta_{t=0}$ is not assumed known and will be estimated together with all the remaining parameters. For $t = T$, an analogous characterization exists, where the shrinkage term and the selection threshold only contain the prospective portion of the penalty.

When $p > 1$, there is a delicate interplay between the multiple series, where overfitting in one direction may impair recovery in other directions. As will be seen in Section 6, anchoring on sparsity is a viable remedy to these issues. We obtain analogous characterizations of the global mode. We will denote with Δ_{tj}^- and Δ_{tj}^+ the selection thresholds (5.10) and (5.11) with $x = x_{tj}$, $\beta_{t-1} = \hat{\beta}_{t-1j}$, and $\beta_{t+1} = \hat{\beta}_{t+1j}$.

Lemma 2. Denote by $\hat{\mathbf{B}} = \{\hat{\beta}_{tj}\}_{t,j=1}^{T,p}$ the global mode of $\pi(\beta_{1:T} | \mathbf{y}_{1:T}, \mathbf{v}_{1:T})$ and $\hat{\mathbf{B}}_{\setminus t,j}$ all but the $(t, j)^{\text{th}}$ entry in $\hat{\mathbf{B}}$. Let $z_{tj} = y_t - \sum_{i \neq j} x_{ti} \hat{\beta}_{ti}$ and $Z_{tj} = x_{tj} z_{tj}$. Then $\hat{\beta}_{tj}$ satisfies the following necessary condition

$$\hat{\beta}_{tj} = \begin{cases} \frac{1}{x_{tj}^2} [|Z_{tj}| - v_t \Lambda^*(\hat{\beta}_{tj} | \hat{\beta}_{t-1j}, \hat{\beta}_{t+1j})]_+ \text{sign}(Z_{tj}) & \text{otherwise,} \\ 0 & \text{if } \Delta_{tj}^- < Z_{tj} < \Delta_{tj}^+. \end{cases}$$

Proof. Follows from Lemma 1, noting that $\hat{\beta}_{tj}$ is a maximizer in $(t, j)^{\text{th}}$ direction while keeping $\hat{\mathbf{B}}_{\setminus t,j}$ fixed, i.e.

$$\hat{\beta}_{tj} = \arg \max_{\beta} \left\{ -\frac{1}{2v_t} (z_{tj} - x_{tj}\beta)^2 + \text{Pen}(\beta | \hat{\beta}_{t-1j}, \hat{\beta}_{t+1j}) \right\}. \quad (5.14) \quad \square$$

Lemma 2 evokes coordinate-wise optimization for obtaining the posterior mode. However, the computation of selection thresholds $(\Delta_{tj}^-, \Delta_{tj}^+)$ (as well as the one-site maximizers (5.13)) requires numerical optimization. The lack of availability of closed-form thresholding hampers practicality when T and p are even moderately large. In the next section, we propose an alternative strategy which capitalizes on closed-form thresholding rules.

A (local) posterior mode $\hat{\beta}_{0:T}$ can be obtained either directly, by cycling over one-site updates (5.14), or indirectly through an EMVS algorithm outlined in the previous section. The direct algorithm consists of integrating out $\gamma_{0:T}$ and solving a sequence of non-standard optimization problems (5.14), which necessitate numerical optimization. The EMVS algorithm, on the other hand, obviates the need for numerical optimization by offering closed form one-site updates. The E-step is very similar to the Gaussian case. The expected previsions \mathbf{v}_t^* can be calculated as before. In the calculation of p_{tj}^* and θ_{tj} , we now have to replace the Laplace spike density. For updating $\beta_{0:T}$, we proceed coordinate-wise, iterating over the following single-site updates while keeping all the remaining parameters fixed. For $1 < t < T$, we have

$$\beta_{tj}^{(m+1)} = \arg \max_{\beta} Q_{tj}(\beta),$$

where

$$Q_{tj}(\beta) = -\frac{\nu_t^*}{2}(z_{tj} - x_{tj}\beta)^2 - \frac{p_{tj}^*}{2\lambda_1}(\beta - \phi_1\beta_{t-1j}^{(m)})^2 - \frac{p_{t+1j}^*}{2\lambda_1}(\beta_{t+1j}^{(m)} - \phi_1\beta)^2 \\ - (1 - p_{tj}^*)\lambda_0|\beta| + p_{t+1j}^* \log \theta_{t+1j} + (1 - p_{t+1j}^*) \log(1 - \theta_{t+1j}), \quad (5.15)$$

and where $z_{tj} = y_t - \sum_{i \neq j} x_{ti}\beta_{ti}^{(m)}$. From the first-order condition, the solution $\beta_{tj}^{(m+1)}$, if nonzero, needs to satisfy $\partial Q_{tj}(\beta)/\partial \beta|_{\beta=\beta_{tj}^{(m+1)}} = 0$. To write the derivative slightly more concisely, we introduce the following notation:

$$Z_{tj} = \nu_t^* x_{tj} z_{tj} + \frac{p_{tj}^* \phi_1}{\lambda_1} \beta_{t-1j}^{(m+1)} + \frac{p_{t+1j}^* \phi_1}{\lambda_1} \beta_{t+1j}^{(m+1)} \quad \text{and} \quad W_{tj} = \left(\nu_t^* x_{tj}^2 + \frac{p_{tj}^*}{\lambda_1} + \frac{p_{t+1j}^* \phi_1^2}{\lambda_1} \right).$$

Then we can write for $\beta \neq 0$

$$\frac{\partial Q_{tj}(\beta)}{\partial \beta} = Z_{tj} - W_{tj}\beta - (1 - p_{tj}^*)\lambda_0 \text{sign}(\beta) + \frac{\partial \theta_{t+1j}}{\partial \beta} \left[\frac{p_{t+1j}^*}{\theta_{t+1j}} - \frac{1 - p_{t+1j}^*}{1 - \theta_{t+1j}} \right], \quad (5.16)$$

where

$$\frac{\partial \theta_{t+1j}}{\partial \beta} = \theta_{t+1j}(1 - \theta_{t+1j}) \left[\lambda_0 \text{sign}(\beta) - \frac{\beta(1 - \phi_1^2)}{\lambda_1} \right]$$

is obtained from (5.7). Recall that θ_{t+1j} , defined in (2.8), depends on β_{tj} (denoted by β above). This complicates the tractability of the M-step. If θ_{t+1j} was fixed, we could obtain a simple closed-form solution $\beta_{tj}^{(m+1)}$ through an elastic-net-like update (Zou and Hastie 2005). We can take advantage of this fact with a one-step-late (OSL) adaptation of the EM algorithm (Green 1990). The OSL EM algorithm bypasses intricate M-steps by evaluating the intractable portions of the penalty derivative at the most recent value, rather than the new value. We apply this trick to the last summand in (5.16). Instead of treating θ_{t+1j} as a function of β in (5.16), we fix it at the most recent value $\beta_{tj}^{(m)}$. The solution for β , implied by (5.16), is then (when $\Lambda_{tj} > 0$)

$$\beta_{tj}^{(m+1)} = \frac{1}{W_{tj} + (1 - \phi_1^2)/\lambda_1 M_{tj}} [|Z_{tj}| - \Lambda_{tj}]_+ \text{sign}(Z_{tj}), \quad \text{for } 1 < t < T, \quad (5.17)$$

where $M_{tj} = p_{t+1j}^*(1 - \theta_{t+1j}) - \theta_{t+1j}(1 - p_{t+1j}^*)$ and $\Lambda_{tj} = \lambda_0[(1 - p_{tj}^*) - M_{tj}]$. The update (5.17) is a thresholding rule, with a shrinkage term that reflects the size of $(\beta_{t-1j}^{(m)}, \beta_{tj}^{(m)}, \beta_{t+1j}^{(m)})$. The exact thresholding property is obtained from sub-differential calculus, because $Q_{tj}(\cdot)$ is not differentiable at zero (due to the Laplace spike). A very similar update is obtained also for $t = T$, where all the terms involving p_{t+1j}^* and θ_{t+1j} in Λ_{tj} , W_{tj} and Z_{tj} disappear. For $t = 0$, we have

$$\beta_{0j}^{(m+1)} = \frac{1}{p_{1j}^* \phi_1^2 + p_{0j}^*(1 - \phi_1^2)} [p_{0j}^* |\beta_{1j}| \phi_1 - (1 - p_{0j}^*) \lambda_0 \lambda_1]_+ \text{sign}(\beta_{1j}). \quad (5.18)$$

The updates (5.17) and (5.18) can be either cycled-over at each M-step, or performed just once for each M-step.

It is straightforward to implement a random-walk variant of this procedure with $\theta_{tj} = \Theta$ by setting $M_{tj} = 0$ in (5.17).

Remark 5. For autoregression with a higher order $h > 1$, the retrospective penalty would be similar where μ_t would depend on h lagged values. The prospective penalty would consist of not just one, but h terms. For the EM implementation, one would proceed analogously by evaluating the derivatives of $\theta_{t+1}, \dots, \theta_{t+h}$ w.r.t. β at the most recent update of the process from the previous iteration and keep them fixed for each one-site update.

To illustrate the ability of the *DSS* priors to suppress noise and recover true signal, we consider a high-dimensional synthetic dataset and a topical macroeconomic dataset.

6 Synthetic High-Dimensional Data

We first illustrate our dynamic variable selection procedure on a simulated example with $T = 100$ observations generated from the model (1.1) with $p = 50$ predictors and with $v_t = 0.25$. The predictor values x_{tj} are obtained independently from a standard normal distribution. Out of the 50 predictors, 46 never contribute to the model (predictors x_{t5} through x_{t50}), where $\beta_{t5}^0 = \beta_{t6}^0 = \dots = \beta_{t50}^0 = 0$ at all times. The predictor x_{t1} is a persisting predictor, where $\{\beta_{t1}\}_{t=1}^T$ is generated according to an *AR*(1) process (2.4) with $\phi_0 = 0$ and $\phi_1 = 0.98$ and where $|\beta_{t1}^0| > 0.5$. The remaining three predictors are allowed to enter and leave the model as time progresses. The regression coefficients $\{\beta_{t2}^0\}_{t=1}^T$, $\{\beta_{t3}^0\}_{t=1}^T$ and $\{\beta_{t4}^0\}_{t=1}^T$ are again generated from an *AR*(1) process ($\phi_0 = 0$ and $\phi_1 = 0.98$). However, the values are rescaled and thresholded to zero whenever the absolute value of the process drops below 0.5, creating zero-valued periods. The true sparse series of coefficients are depicted in Figure 3 (black lines).

We begin with the standard DLM approach, which is equivalent to *DSS* when the selection indicators are switched on at all times, i.e., $\gamma_{tj} = 1$ for $t = 0, \dots, T$ and $j = 1, \dots, p$. This is equivalent to setting $\Theta = 1$ in our prior. The autoregressive parameter ϕ_1 is assigned the prior (3.2) and estimated. We also estimate the variances v_t using the discount stochastic volatility model (1.2) with $\delta = 0.9$ and $n_0 = d_0 = 10$. Plots of the estimated posterior mode trajectories of the first 6 series (including the 4 active ones) are in Figure 3 (red broken lines). With the absence of the spike, the estimated series of coefficients cannot achieve sparsity. By failing to discern the coefficients as active or inactive, the state process confuses the source of the signal, distributing it across the redundant covariates. This results in loss of efficiency and poor recovery.

With the hope to improve on this recovery, we deploy the *DSS* process with a sparsity inducing spike. First, we apply Dynamic SSVS with 1000 iterations and 200 burnin time. We set the spike and slab parameters $\Theta = 0.1$, $\lambda_1 = 0.1$ and $\lambda_0 = 0.01$ so that the ratio between spike and slab variances is sufficiently large (George and McCulloch 1993). The autoregressive parameter ϕ_1 is estimated under the prior (3.2) and the stochastic volatilities are also estimated with $\delta = 0.9$ and $n_0 = d_0 = 10$. We plot the posterior mean of the regression coefficients in Figure 3 (black broken line) together with the credible sets (black dotted lines). The recovered series have a strikingly different pattern compared to the non-sparse DLM solution (red dotted lines). First, the

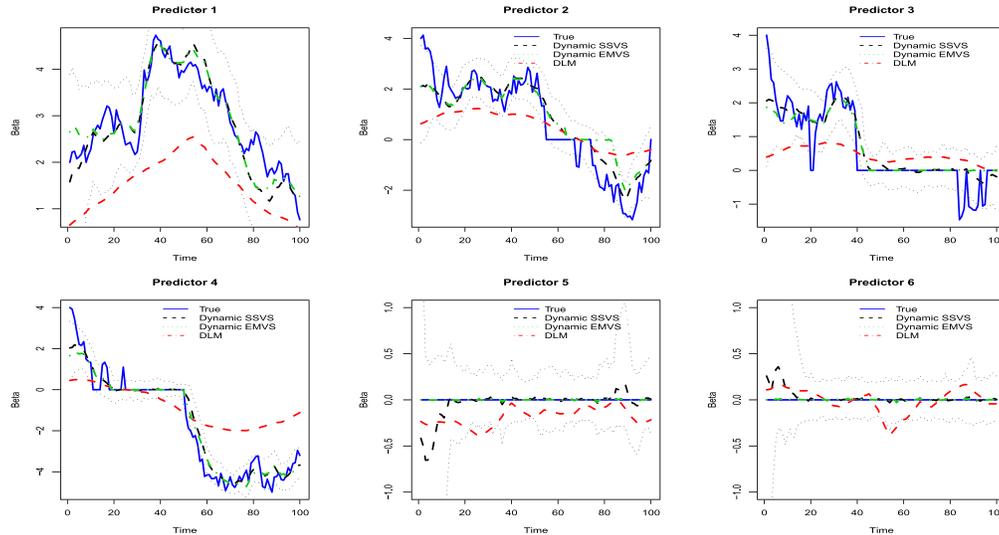


Figure 3: The first six regression coefficients of the true (blue solid lines) and estimated regression coefficients in the simulated example with $p = 50$. The estimates are posterior means from Dynamic SSVS (black broken line) and modes from Dynamic EMVS (green broken line). Comparisons are made with DLM (red dotted line). The black dotted lines denote pointwise credible intervals.

estimated series is seen to track closely the periods of predictor importance/irrelevance, achieving dynamic variable selection. Second, by harnessing sparsity, the *DSS* priors alleviate bias in the nonzero directions, outputting a cleaner representation of the true underlying signal. The posterior mean of the autoregressive parameter ϕ_1 is 0.94. In addition, we plot the posterior inclusion probabilities $P(\gamma_{tj} = 1 \mid \mathbf{y}_{1:T})$ for the first 6 predictors (Figure 4, black lines). These quantities can be used to guide variable selection by focusing on those coefficients whose inclusion probability is at least 0.5 (Barbieri and Berger 2004). Indeed, we can see that these estimated probabilities drop below 0.5 when the true signal is absent, effectively recovering the “pockets of predictability”. The posterior mean of the coefficient ϕ_1 was estimated at 0.981 (very close to the true value 0.98) with posterior a credible interval (0.973, 0.989). The computation took 151.8 seconds in R. We will now turn to Dynamic EMVS to see whether similarly successful recovery can be achieved with less time.

We apply Dynamic EMVS considering the same spike and slab hyper-parameters, i.e. $\lambda_1 = 0.1$ and $\lambda_0 = 0.01$. The global sparsity weight Θ can be regarded as a tempering parameter, where $\Theta = 1$ corresponds to the DLM case. By choosing smaller values Θ , the posterior becomes more multi-modal making it easier for the EM to get trapped. Since the EMVS computation is very fast, we can alleviate local entrapments by applying a deterministic annealing strategy, similar to the one suggested in Rockova and George (2014). We will consider not only one value $\Theta = 0.1$, but a whole sequence of decaying values $\Theta \in \{1, 0.9, 0.5, 0.1\}$ with warm starts. Namely, the output obtained with a larger

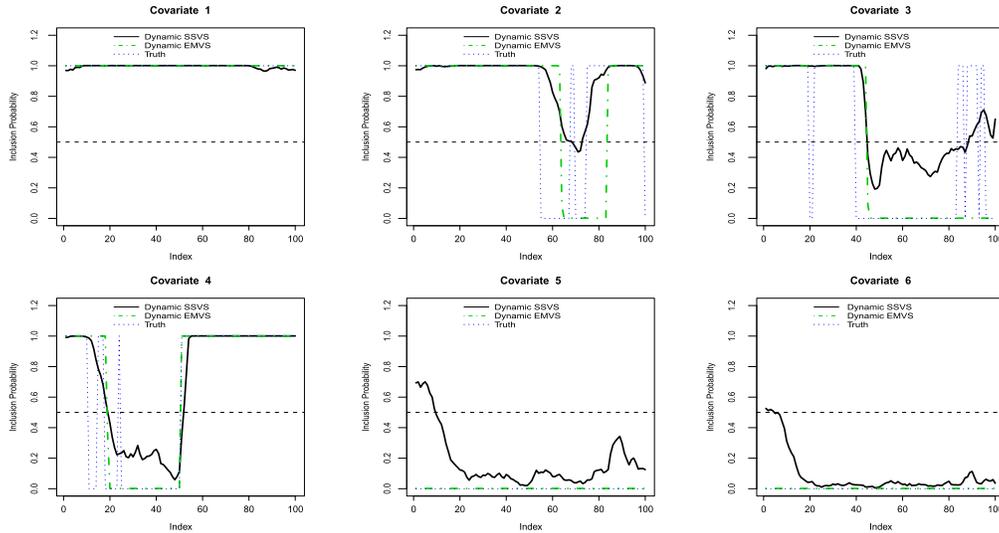


Figure 4: Posterior inclusion probabilities $P(\gamma_{tj} = 1 | \mathbf{y}_{1:T})$ (Dynamic MCMC) and conditional inclusion probabilities $P(\gamma_{tj} = 1 | \hat{\beta}_{tj}, \mathbf{y}_{1:T})$ (Dynamic EMVS) and true pattern of sparsity for the first six series.

value Θ will be used as an initialization for the computation at the next smaller value Θ in the chosen sequence. In this way, we obtain an entire solution path (not only one single solution), we accelerate convergence and increase the chances for the EM to find a promising mode. We successfully apply this strategy for $\Theta \in \{1, 0.9, 0.5, 0.1\}$ and, similarly as before, we estimate ϕ_1 and all the variances v_t under the same priors. The estimated regression coefficients obtained with $\Theta = 0.1$ are depicted in Figure 3 (green broken lines). We can again see dramatic improvements over DLM (obtained with $\Theta = 1$) and, interestingly, a very similar recovery to the posterior mean with Dynamic SSVS. The R computations took 15 seconds for $\Theta = 0.9$, 6 seconds for $\Theta = 0.5$ and 8 seconds for $\Theta = 0.1$, yielding nontrivial computational dividends compared to MCMC (151.8 s for 1000 iterations). In addition, Dynamic EMVS outputs conditional inclusion probabilities $P[\gamma_{tj} = 1 | \hat{\beta}_{tj}, \mathbf{y}_{1:T}]$ which can be regarded as the conditional counterpart to the marginal posterior inclusion probabilities $P[\gamma_{tj} = 1 | \mathbf{y}_{1:T}]$ estimated from MCMC. As can be seen from Figure 4, these conditional probabilities track closely the marginal ones and, again, drop below 0.5 when the true signal is not present. These companion plots are helpful visualizations of the time-varying sparsity profile. In conclusion, the plots for Dynamic SSVS and Dynamic EMVS largely agree.

Next, we deploy the Laplace spike variant of DEMVS with a random walk slab prior and with $\lambda_1 = 0.1, \lambda_0 = 1$ and $\Theta = 0.5$. This hyper-parameter choice corresponds to a very mild separation between the spike and slab distributions. We apply the one-step-late EM algorithm outlined in Section 5.3, initializing the calculation with the output from DLM. We assume that the initial vector $\beta_{t=0}$ is drawn from a Θ -weighted mixture distribution between the Laplace density and a zero-mean Gaussian density with

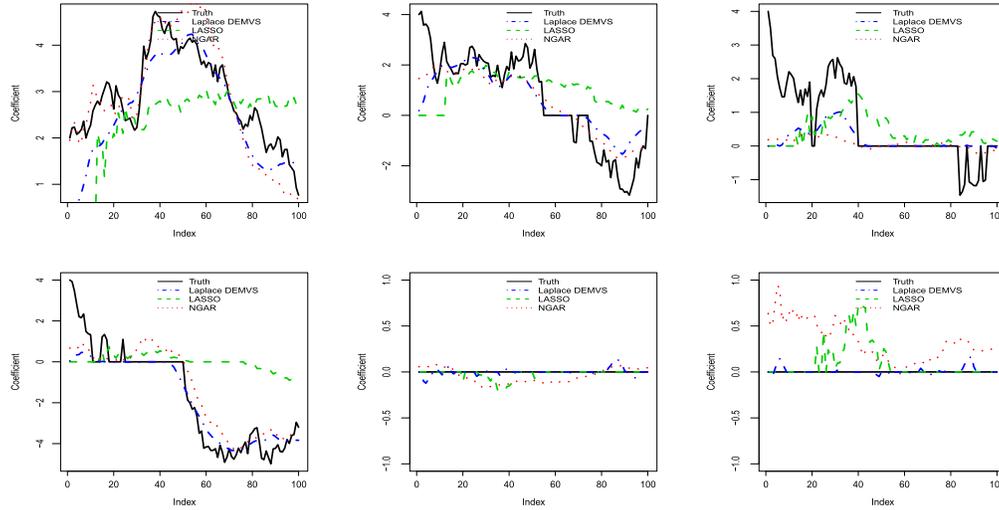


Figure 5: The first six regression coefficients of the true and estimated regression coefficients in the high-dimensional simulated example with $p = 50$. We compare DSS (Laplace version) with NGAR and LASSO.

variance one and we estimate it together with all the other parameters, as prescribed in Section 5.3.

We also compare the performance to the NGAR process of Kalli and Griffin (2014) and the LASSO method. The latter does not take into account the temporal nature of the problem. For NGAR, we use the default settings, $b^* = s^* = 0.1$, with 1,000 burn-in and 2,000 MCMC iterations. For LASSO, we sequentially run a static regression in an extending window fashion, where the LASSO regression is refit using $1:t$ for each $t = 1:T$ to produce a series of quasi-dynamic coefficients; a common practice for using static shrinkage methods for time series data (Bai and Ng 2008; De Mol et al. 2008; Stock and Watson 2012; Li and Chen 2014), choosing λ via 10-fold cross-validation. The estimated trajectories are depicted in Figure 5.

For the first series, the only persistent series, both *DSS* (Laplace) and NGAR succeeds well in tracing the true signal. This is especially true in contrast to DLM and LASSO, which significantly under-valuate the signal. The estimated coefficient evolutions for DLM and LASSO become inconclusive for assessing variable importance, where the coefficient estimates for the relevant variables have been polluted by the elevated estimates for the irrelevant variables. For the second to fourth series with intermittent zeros, we see that *DSS* and NGAR are able to separate the true zero/nonzero signal (noted by the flat coefficient estimates during inactive periods). The LASSO method produces sparse estimates, however the variable selection is not linked over time and thereby erratic. For the two zero series (series five and six), both *DSS* and LASSO truly shrink noise to zero. The *DSS* priors mitigate overfitting by eliminating noisy coefficients and thereby leaving enough room for the true predictors to capture the trend.

$p = 50$			$x_{1:50}$		$x_{1:4}$		$x_{5:50}$		
	Time (s)		SSE	Ham.	SSE	Ham.	SSE	Ham.	
NGAR	564.2		539.2	4708	320.9	108	218.3	4600	
LASSO	9.3		1621.8	281.2	1595.1	186.4	26.7	94.8	
Dynamic SSVS (Gaussian)									
$\lambda_1 = .1$	$\Theta = 1$ (DLM)		164.9	1625.4	4708	1426.8	108	198.5	4600
$\lambda_1 = .1$	$\lambda_0 = .01$	$\Theta = .5$	132.6	1541.6	2447.7	1457.2	108	84.4	2339.7
$\lambda_1 = .1$	$\lambda_0 = .01$	$\Theta = .1$	118.6	108.7	51.4	104.3	40.5	4.5	10.9
$\lambda_1 = .1$	$\lambda_0 = .01$	$\Theta = .1$ (RW)	113.6	683.2	86.5	660.6	63.6	22.6	22.9
$\lambda_1 = .01$	$\lambda_0 = .001$	$\Theta = .1$	117.9	177.6	431	141.2	105.2	36.4	325.8
Dynamic EMVS (Gaussian)									
$\lambda_1 = .1$	$\Theta = 1$ (DLM)		5.1	1241	4708	975.2	108	265.7	4600
$\lambda_1 = .1$	$\lambda_0 = .01$	$\Theta = .9^*$	16.5	286.3	106.3	241.9	54.4	44.3	51.9
$\lambda_1 = .1$	$\lambda_0 = .01$	$\Theta = .5^*$	12.4	294.7	99.6	254.7	56	40	43.6
$\lambda_1 = .1$	$\lambda_0 = .01$	$\Theta = .1^*$	14.3	309.6	93	267.5	59.2	42.1	33.8
$\lambda_1 = .1$	$\lambda_0 = .01$	$\Theta = .1$ (RW)	8	495.8	91.8	494.1	88.2	1.7	3.6

Table 3: Performance evaluation of *DSS*, LASSO, NGAR and DLM on the simulated example with $p = 50$. The results are split for the signal parameters ($x_{1:4}$) and noise parameters ($x_{5:50}$). Dynamic SSVS uses 1000 iterations with 100 burnin. RW stands for a random-walk variant with $\theta_{tj} = \Theta$. EMVS calculations are initialized at DLM solutions besides settings denoted with * where we use warm starts.

We repeat this experiment 10 times, generating different responses and regressors using the same set of coefficients. We compare the sum of squared error (SSE) between the recovered estimates $\hat{\mathbf{B}}$ and the true series \mathbf{B}_0 as well as the Hamming distance between the true and estimated sparsity patterns. For the MCMC version, the sparsity pattern will be estimated from the matrix of posterior inclusions $\Pi = (\pi_{tj})$ where $\pi_{tj} \equiv \mathbb{P}[\gamma_{tj} = 1 \mid \mathbf{y}_{1:T}]$ according to the median probability model rule. We then define the Hamming distance as

$$\text{Ham}(\Pi, \mathbf{B}_0) = \sum_{j=1}^p \sum_{t=1}^T |\mathbb{I}(\pi_{tj} > 0.5) - \mathbb{I}(\beta_{tj}^0 \neq 0)|.$$

For Dynamic EMVS (with Gaussian spike), we can use the conditional inclusion probabilities instead of π_{tj} . Alternatively, one can obtain sparsity patterns by thresholding out coefficients whose magnitude is smaller than the intersection point between the stationary spike and slab densities (as we explained in Section 2.2).

Table 3 reports average performance metrics over the 10 experiments. The performance of *DSS* is compared to the full DLM model (West and Harrison 1997), NGAR (Kalli and Griffin 2014) and LASSO. For DLM and NGAR, we use the same specifications as above. For *DSS*, we now explore a multitude of combinations of hyperparameters using both Dynamic SSVS and Dynamic EMVS. For the Gaussian spike, we choose $\lambda_1 \in \{0.1, 0.01\}$, $\lambda_0 \in \{0.01, 0.001\}$, and $\Theta \in \{0.9, 0.5, 0.1\}$. We also consider a random-walk prior variant with $\theta_{tj} = \Theta$. For Dynamic EMVS, we initialize the calculations at DLM solutions apart from the settings marked with a star, where we use warm starts (as explained above). We focus on the Gaussian EMVS variant, where additional simulations for the Laplace spike are reported in Table 1 in the Appendix.

Looking at Table 3, *DSS* performs better in terms of both SSE and Hamming distance compared to DLM, NGAR, and LASSO for the majority of the hyperparameters considered. To gain more insights, the table is divided into three blocks: overall performance on $\beta_{1:50}$, active coefficients $\beta_{1:4}$ and noise coefficients $\beta_{5:50}$. Because DLM and NGAR only shrink (and do not select), the Hamming distance for the block of noisy coefficients is 100%. For Dynamic SSVS, the sum of squared errors and the Hamming distance is seen to increase with Θ . It is interesting to note the difference in performance between our stationary *DSS* version, where θ_{tj} are dynamically evolving, and the random-walk (RW) version, where $\theta_{tj} = \Theta$. In this stationary situation, there is a clear advantage in linking the weights over time using the deterministic construction (2.8). We found the settings $\lambda_1 = 0.1$, $\lambda_0 = 0.01$ and $\Theta = 0.1$ to work well on this example, where the threshold of practical significance (i.e. the intersection point between the stationary spike and slab densities as discussed in Section 2.2) equals 0.086. Decreasing this threshold to 0.05 with a sharper spike-and-slab prior ($\lambda_1 = 0.01$, $\lambda_0 = 0.001$ and $\Theta = 0.1$), many more false discoveries occur (i.e. increased Hamming distance for the noise coefficients) due to the fact that even very small noisy effects can be assigned to the slab distribution.

Dynamic EMVS reconstructs signal much faster compared to the 1000 iterations of Dynamic SSVS. While the MAP trajectory is not as good in terms of SSE (which is expected from a (local) posterior mode), its performance is still better than LASSO, DLM and NGAR. Again, we found the setting $\lambda_1 = 0.1$, $\lambda_0 = 0.01$ and $\Theta = 0.1$ to work well and we can clearly see dividends of dynamic weighting relative to the random-walk prior. Comparing the results with DLM and LASSO, *DSS* showcases the benefits of combining dynamics and shrinkage, since DLM (only dynamics) and LASSO (only shrinkage) underperform significantly. Regarding timing comparisons with NGAR, we need to point out that NGAR was run with 2,000 iterations and 1,000 burnin, which we found sufficient to obtain nice results.

Now, we explore a far more challenging scenario, repeating the example with $p = 200$ instead of 50. The coefficients and data generating process are the same with $p = 50$, but now instead of 46 noise regressors, we have 196. This high regressor redundancy rate is representative of the “ $p \gg n$ ” paradigm (“ $p \gg T$ ” for time series data) and can test the limits of any sparsity inducing procedure. Note that the number of coefficients to estimate is $p \times T = 20\,000$. This is a very challenging scenario where we will be able to truly evaluate the efficacy of *DSS* when there is a large number of predictors with sparse signals.

The results are collated in Table 4. Across considered hyper-parameter settings, Dynamic EMVS does extremely well also for $p = 200$, dramatically reducing SSE over DLM and LASSO. We have performed the warm start strategy for $\Theta \in \{1, 0.99, 0.9, 0.5, 0.1\}$. Moving from $\Theta = 1$ to $\Theta = 0.99$ already yields considerable improvements in terms of separating the signal from noise. Reducing Θ even further, one obtains reduced Hamming distance for redundant covariates, i.e. noise is being absorbed inside the spike. While LASSO does perform well in terms of the Hamming distance, it does not do so well in terms of SSE. Because LASSO lacks dynamics, the pattern of sparsity is not smooth over time, leading to erratic coefficient evolutions. Because of the smooth nature of its sparsity, *DSS* harnesses the dynamics to discern signal from noise, improving

$p = 200$			Time (s)	$x_{1:50}$		$x_{1:4}$		$x_{5:50}$		FD	FN	DIM
				SSE	Ham.	SSE	Ham.	SSE	Ham.			
LASSO			29.8	1760.2	395	1744.7	197.1	15.5	197.9	34.1	0.1	38
Dynamic EMVS (Gaussian)												
$\lambda_1 = .1$		$\Theta = 1$ (DLM)	16.2	2253.3	19708	2209.9	108	43.4	19600	196	0	200
$\lambda_1 = .1$	$\lambda_0 = .01$	$\Theta = .99^*$	113.5	555.9	580.6	479	91.6	76.8	489	26.2	0.2	30
$\lambda_1 = .1$	$\lambda_0 = .01$	$\Theta = .9^*$	50.9	469.8	153.7	422.4	85.2	47.4	68.5	4.4	0.3	8.1
$\lambda_1 = .1$	$\lambda_0 = .01$	$\Theta = .5^*$	43.8	500.3	154.5	447.3	89.5	53	65	3.9	0.3	7.6
$\lambda_1 = .1$	$\lambda_0 = .01$	$\Theta = .1^*$	51.9	534.2	151.1	473.6	94	60.6	57.1	3.4	0.3	7.1
$\lambda_1 = .1$	$\lambda_0 = .01$	$\Theta = .1^*$ (RW)	34.2	550	122.9	502.9	91.3	47.1	31.6	2.7	0.3	6.4

Table 4: Performance evaluation of *DSS*, LASSO and DLM on the simulated example with $p = 200$. The results are split for the signal parameters ($x_{1:4}$) and noise parameters ($x_{5:50}$). RW stands for a random-walk variant with $\theta_{tj} = \Theta$. EMVS calculations are initialized at previous solutions (warm starts), as designated by the * sign. FD stands for False Discoveries (noise variables which were identified as active at least once), FN stands for False Nondiscoveries (number of true variables which were removed from the model at all time points), DIM is estimated number of covariates identified as active at least once.

in both SSE and Hamming distance. We have also added global variable selection performance metrics: False Discoveries (FD), False Non-discoveries (FN) and Dimension (DIM). FD is defined as the number of noise variables (out of the 196 noise predictors) which were included in the model at least once during the time $t = 1, \dots, 100$. Similarly, FN is the number of true signal variables (out of the 4 true predictors) which were left out of the model at all times points $t = 1, \dots, 100$. Finally, DIM is the estimated number of predictors identified as active at least once. We can see that LASSO includes too many noise variables, while Dynamic EMVS effectively reduces the dimensionality. In this vein, Dynamic EMVS can be regarded as a fast screening rule which can be followed by a more thorough analysis using only a smaller subset of more meaningful predictors.

7 Macroeconomic Data

We further illustrate the effectiveness of *DSS* through a macroeconomic dataset analyzed in Kalli and Griffin (2014). The data consists of quarterly measurements of the US inflation (the personal consumption expenditure (PCE) deflator) and 31 potential explanatory variables including previous lags of inflation, activity variables (such as economic growth rate or output gap), unemployment rate etc. The dataset was obtained from the FRED (Federal Reserve Bank of St. Louis) economic database, the consumer survey database of the University of Michigan, the Federal Reserve Bank of Philadelphia, and the Institute of Supply Management (see Kalli and Griffin 2014 and Figure 2 for more details).

For this example, we will treat the US inflation (Figure 6) as the dependent variable and infer its sources of covariation with the other variables. Inflation forecasting has been of substantial interest within the macroeconomic literature (Stock and Watson 1999; Koop and Korobilis 2012a; Groen et al. 2013; Kalli and Griffin 2014; Wright 2009; Stock and Watson 2007). The primary goal of our analysis is to retrospectively identify underlying economic indicators that are pertinent to inflation. In addition, we evaluate

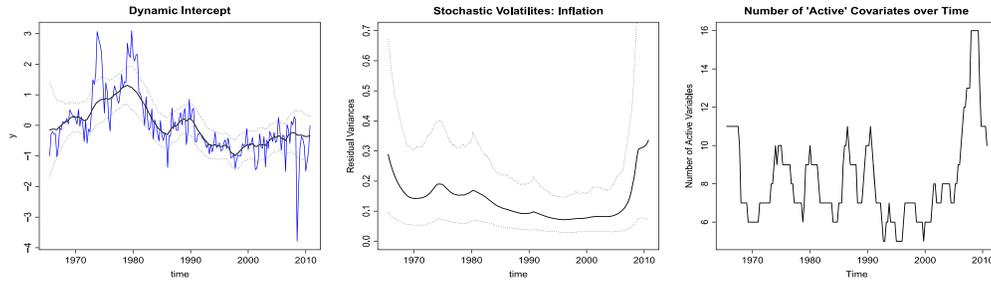


Figure 6: (Left) Observed quarterly US inflation (recentered and rescaled) from 1965/2 to 2011/1 (blue time series). The black lines are the posterior mean of the dynamic intercept together with 95% pointwise credible bands. (Middle) Posterior means of residual variances (together with 95% pointwise credible bands) under the discount stochastic volatility model. (Right) Number of covariates with a posterior inclusion probability above 0.5.

the one-step-ahead forecasting ability of our models. Because the economy is dynamic, it is natural to assume certain indicators to be effective during a certain period but useless during another. For example, one might expect financial indicators to play a significant role in the economy during a financial crisis. The necessity of capturing these dynamic trends have been discussed and explored in Stock and Watson (2007), who point out that forecasting inflation has become harder due to trend cycles and dynamic volatility processes. Unlike Stock and Watson (2007), where they model this trend via an unobserved component trend-cycle, we explore this characteristic through dynamic sparsity in the covariate space. Similar inflation forecasting applications were considered by many other authors including Kalli and Griffin (2014); Koop and Korobilis (2012a). The dataset has a long span (from the second quarter of 1965 to the first quarter of 2011), capturing oil shocks in 1973 and 1979, mild recession in 1990, the dot-com bubble as well as the Great Recession in 2007-2009. We would expect our forecasting model to change during these periods.

To evaluate our method, we first measure its forecasting ability by conducting one-month-ahead point forecasts and computing the mean squared (absolute) cumulative forecast error (MSFE and MAFE). The analysis is done by cutting the data in half, training the methods using the first half of the data from 1965/2 to 1987/3. We then sequentially update the forecasts through the second half from 1987/4 to 2011/1, updating and rerunning estimation to produce 1-month ahead forecasts every time we observe a new data point (using data from $1:t$ to forecast $t + 1$ for $t = 1:T - 1$, where $t = 1$ is 1965/2, and $t = T - 1$ is 2010/4). Namely, we refit the full MCMC analysis (with 500 MCMC iterations and 100 burn-in) for each model to define the posterior based on $\mathbf{y}_{1:t}$ and to obtain forecasts $f_{t+1} = \mathbf{x}'_{t+1} \mathbf{a}_{t+1}$ (using the notation from Section 3). We use conditional forecast densities as explained in the next paragraph. For Dynamic EMVS, we replace posterior means with modes in the forecast calculations. Out-of-sample forecasting is thus conducted in a way that no future information is used to analyze and evaluate the results. At the end of the analysis (2011/1), we estimate the retrospective coefficients using the entire dataset to glean insights into the recovered signals. As

	LPDS	MSFE	MAFE		MSFE	MAFE
Dynamic SSVS				Dynamic EMVS		
$\Theta = 1, \lambda_1 = 0.01$ (DLM) Intercept		38.9	37.03	LASSO	32.97	34.38
$\Theta = 1, \lambda_1 = 0.01$ (DLM) Full	-282.3	51.01	45.56	$\Theta = 1, \lambda_1 = 0.01$ (DLM) Full	46.59	45.25
$\Theta = 0.1, \lambda_1 = 0.01, \lambda_0 = 0.001$	-81.51	36.32	34.28	$\Theta = 0.1, \lambda_1 = 0.01, \lambda_0 = 0.001$	31.61	32.73
$\Theta = 0.5, \lambda_1 = 0.01, \lambda_0 = 0.001$	-93.12	41.31	37.67	$\Theta = 0.5, \lambda_1 = 0.01, \lambda_0 = 0.001$	31.69	32.83
$\Theta = 0.5, \lambda_1 = 0.01, \lambda_0 = 0.001$ (RW)	-117.67	42.69	42.36	$\Theta = 0.5, \lambda_1 = 0.01, \lambda_0 = 0.001$ (RV)	41.37	43.34

Table 5: Mean squared (absolute) one-step-ahead point forecast errors for Dynamic SSVS and Dynamic EMVS. RW stands for the random walk prior variant, “Intercept” stands for a model with only a dynamic intercept and “Full” stands for a full DLM model with no selection shrinkage. LPDS stands for the log-predictive density score measuring the quality of distribution one-step-ahead forecasts.

in Section 7, we compare *DSS* against the full DLM, null DLM (only intercept) and LASSO (expanding window). For *DSS*, we use multiple hyperparameters to discern the combination which produces best forecasts. The initial condition for the SV variance is $1/v_0 \sim G(n_0/2, d_0/2)$ with $n_0 = 1$ and $d_0 = 1$. The discount factor is set to 0.9.

On the comparison of forecast ability (Table 5), it is curious that the null DLM model actually performs better than the full DLM model. While the predictors have some explanatory power, the full DLM model is unable to tease out the signal and badly overfits, clearly hurting forecasts. Dynamic EMVS (SSVS) is able to improve on the null model by capitalizing on the (albeit weak) signal hidden in the predictors. *DSS* thus significantly improves over the full DLM, which is unsurprising since the full DLM model is plagued with overfitting and false discoveries. A surprising result is that Dynamic EMVS outperforms Dynamic SSVS in this example. This can be explained by the fact that the posterior MAP trajectory is sparser (shrunk towards zero more aggressively) and smoother than the posterior mean (which performs model averaging). The added benefit of smoothing (in addition to sparsity) can be seen by comparing Dynamic EMVS to the LASSO, which achieves shrinkage, but does not capture the dynamics of signals. The fact that the forecasting results of LASSO and Dynamic EMVS are similar suggests that the gains from shrinkage are similar. *DSS*, capturing and capitalizing on both dynamics and shrinkage, achieves forecast gains relative to just shrinkage (LASSO) or just dynamics (DLM). The improved performance of Dynamic EMVS is reassuring in the sense that the faster implementation can still yield point forecasts that are very similar, if not better, to the ones obtained from the more time consuming MCMC. We also compare forecasting performance in terms of a metric that involves the entire predictive distribution (not just its mean), namely the sum of log-predictive likelihoods evaluated at observed values y_{t+1} (Koop and Korobilis 2012a). We use a conditional variant of the predictive likelihood $\pi(y_{t+1} | \mathbf{y}_{1:t}, \hat{v}_{t+1}, \hat{\gamma}_{t+1})$, where we condition on the posterior mean of the inclusion indicators and variances, i.e. $\hat{\gamma}_{t+1j} = \mathbf{E}[\gamma_{t+1j} | \mathbf{y}_{1:t}]$ and $\hat{v}_{t+1} = \mathbf{E}[v_{t+1} | \mathbf{y}_{1:t}, \hat{\beta}_{1:t}]$ with $\hat{\beta}_{tj} = \mathbf{E}[\beta_{tj} | \mathbf{y}_{1:t}]$. Again, *DSS* improves on the full model for all the settings considered.

We now deploy *DSS* priors using Dynamic SSVS (2000 posterior samples with a 500 burn-in period) on the entire dataset to recover the series of regression coefficients. In order to capture more subtle signals, we set $\lambda_1 = 0.01, \lambda_0 = 0.001$ and $\Theta = 0.1$ so that the intersection point between the stationary spike-and-slab densities (i.e. our perceived

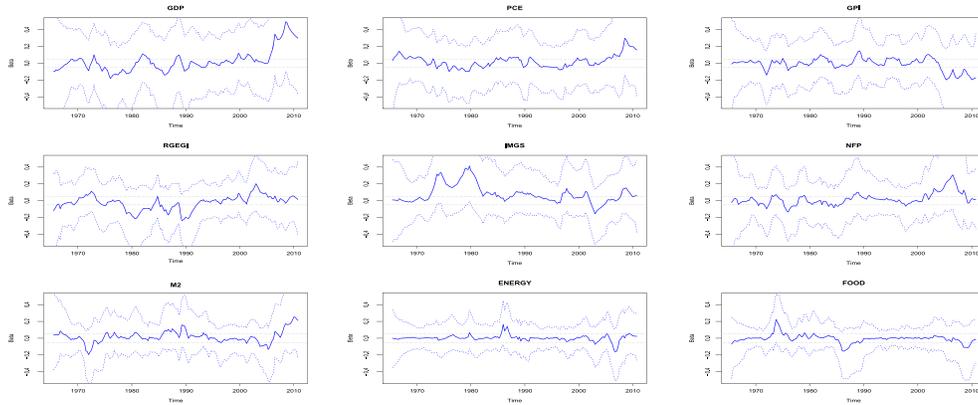


Figure 7: Estimated coefficient evolutions (posterior means) of 9 top predictors together with 95% point-wise credible bands (dotted lines) using Dynamic SSVS. The horizontal lines correspond to the selection threshold 0.05.

selection threshold for practical significance) is 0.05. We assume the discount stochastic volatility model with $\delta = 0.9$ and $n_0 = d_0 = 1$ and include an intercept term which is devoid of shrinkage (i.e. the intercept is in the slab distribution at all times). Out of the 31 indicators (not including the intercept) only 12 (GDP, PCE, GPI, RGEGI, IMGS, NFP, M2, ENERGY, FOOD, MATERIALS, OUTPUT GAP, GS10) had their posterior inclusion probability $P[\gamma_{tj} = 1 | \mathbf{y}_{1:T}]$ above 0.5 at least 10 times throughout the 182 observations. We plot the number of “active” covariates (i.e. with a posterior inclusion probability above 0.5) over time in Figure 6 on the right. Note that the definition of an active coefficient is ultimately tied to our choice of hyper-parameters and our practical significance threshold 0.05, where more strict shrinkage priors would lead to fewer active covariates. This plot presents evidence that the forecasting model is changing over time. More predictors are seen to contribute around the oil shocks, around 1990 and during the financial crisis, mirroring the inflation changes during these periods. A similar conclusion was also found in Koop and Korobilis (2012a).

The coefficient evolutions of the top 9 predictors relevant for inflation are plotted in Figure 7. Many of these predictors were also identified by Kalli and Griffin (2014) with very similar estimated coefficient trajectories. In particular, the explanatory power of IMGS (import of goods and services) growth is seen to peak around the oil shocks in the 1970’s and around late 2000’s. The two largest signals are the production growth indicator GDP and the consumption growth PCE with their coefficients largely stable with a marked increase during crisis in the late 2000’s. Interestingly, conventional indices of the labor market (including unemployment) are not recovered with a very strong signal (see Figure 1 in the Appendix). The characteristics of these coefficients demonstrate how *DSS* is successful in dynamically shrinking coefficients to zero during regime changes (recessions) as we would expect to happen. It is worth noting that the credible intervals absorb zero, indicating inherent sparsity/low signal of the contributing predictors. This is in line with earlier conclusions reached by Koop and Korobilis (2012a) who found

only very few predictors to be relevant for (one-step ahead) inflation forecasting using dynamic model averaging. From the plot of estimated dynamic intercept in Figure 6 (on the left), we can see that the intercept itself is nicely tracking the data, leaving room for other predictors to explain the two shocks in the 1970's and the drop around late 2000's. The plot of the estimated variances (Figure 6 on the right) marks these structural shocks with increased estimated volatility, especially in the late 2000's. This companion plot indicates that structural changes affect both mean and variance.

Through this macroeconomic example, we were able to demonstrate the efficacy of *DSS*, in terms of forecast ability and interpretability, on a real macroeconomic dataset. With *DSS*, we were able to successfully recover interesting signals and reproduce findings discovered by other authors. Dynamic SSVS benefits from available uncertainty statements and full posterior/predictive distributions, which can lead to potentially better informed forecasts/decisions.

8 Discussion

This paper introduces a new class of dynamic shrinkage priors, where the stationary distribution is fully known and characterized by spike-and-slab marginals. A key to obtaining this stabilizing property is the careful hierarchical construction of adaptive mixing weights that allows them to depend on the lagged value of the process, thereby reflecting sparsity of past coefficients. We propose various versions of dynamic spike-and-slab (*DSS*) priors, using Laplace/Gaussian spike/slab distributions. For implementation, we resort to both optimization as well as posterior sampling. For Gaussian *DSS* prior variants, we develop a Dynamic SSVS MCMC algorithm for posterior sampling. For fast MAP smoothing, we develop a complementary procedure called Dynamic EMVS which can quickly glean into the signal structure. For the Laplace spike variant, we implement a one-step-late EM algorithm for MAP estimation which iterates over one-site closed-form thresholding rules. Through simulation and a macroeconomic dataset, we demonstrate that *DSS* are well suited for the dual purpose of dynamic variable selection (through thresholding to exact zero) and smoothing (through an autoregressive slab process) for forecasting and inferential goals.

Many variants and extensions are possible for our *DSS* prototype constructions. While our development has focused on stationary situations, our priors can accommodate random walk evolutions (as we point out in our Remark 2). Schotman and van Dijk (1991) argue that “There is no need to look at the data from the specific viewpoint of stationarity or nonstationarity. Given the data one can determine which of the two is the most likely.” We view stationarity as a modelling assumption which may be suitable for some data sets and less appropriate for others. For instance, Schotman and van Dijk (1991) discovered that for real exchange rate data, stationarity is a posteriori as probable as the random walk hypothesis. Recently, Lopes et al. (2016) propose a mixture prior which switches between stationary and non-stationary specifications and enables the quantification of posterior plausibility of the unit root hypothesis. An extension of our approach along these lines would be very interesting. We provide both stationary and non-stationary variants for the practitioners to choose from. Another interesting

extension will be embedding our *DSS* priors within the TVP-VAR models (Cogley and Sargent 2005; Primiceri 2005; Lopes et al. 2016; Nakajima and West 2013a; Pettenuzzo et al. 2018; Gefang 2014; Giannone et al. 2014; Korobilis 2013; Bańbura et al. 2010).

An R code is available from the first author upon request.

Supplementary Material

Supplementary Material to “Dynamic Variable Selection with Spike-and-Slab Process Priors” (DOI: [10.1214/20-BA1199SUPP](https://doi.org/10.1214/20-BA1199SUPP); .pdf).

References

- Andel, J. (1983). “Marginal distributions of autoregressive processes.” In *Transactions of the Ninth Prague Conference*, 127–135. Springer. [MR0757732](https://doi.org/10.1007/978-1-4613-1111-1_12). 238, 239
- Antoniadis, A. and Fan, J. (2001). “Regularization of wavelet approximations.” *Journal of the American Statistical Association*, 96(455): 939–967. [MR1946364](https://doi.org/10.1198/016214501753208942). doi: <https://doi.org/10.1198/016214501753208942>. 236, 250
- Bai, J. and Ng, S. (2008). “Forecasting economic time series using targeted predictors.” *Journal of Econometrics*, 146(2): 304–317. [MR2465175](https://doi.org/10.1016/j.jeconom.2008.08.010). doi: <https://doi.org/10.1016/j.jeconom.2008.08.010>. 256
- Bańbura, M., Giannone, D., and Reichlin, L. (2010). “Large Bayesian vector auto regressions.” *Journal of Applied Econometrics*, 25(1): 71–92. [MR2751790](https://doi.org/10.1002/jae.1137). doi: <https://doi.org/10.1002/jae.1137>. 264
- Barbieri, M. M. and Berger, J. O. (2004). “Optimal predictive model selection.” *The Annals of Statistics*, 32: 870–897. [MR2065192](https://doi.org/10.1214/009053604000000238). doi: <https://doi.org/10.1214/009053604000000238>. 254
- Belmonte, M. A. G., Koop, G., and Korobilis, D. (2014). “Hierarchical shrinkage in time-varying parameter models.” *Journal of Forecasting*, 33(1): 80–94. [MR3148281](https://doi.org/10.1002/for.2276). doi: <https://doi.org/10.1002/for.2276>. 234
- Bhattacharya, A., Pati, D., Pillai, N., and Dunson, D. (2015). “Dirichlet-Laplace priors for optimal shrinkage.” *Journal of the American Statistical Association*, 110: 1479–1490. [MR3449048](https://doi.org/10.1080/01621459.2014.960967). doi: <https://doi.org/10.1080/01621459.2014.960967>. 241
- Bitto, A. and Frühwirth-Schnatter, S. (2019). “Achieving shrinkage in a time-varying parameter model framework.” *Journal of Econometrics*, 210: 75–97. [MR3944764](https://doi.org/10.1016/j.jeconom.2018.11.006). doi: <https://doi.org/10.1016/j.jeconom.2018.11.006>. 235
- Brodie, J., Daubechies, I., De Mol, C., Giannone, D., and Loris, I. (2009). “Sparse and stable Markowitz portfolios.” *Proceedings of the National Academy of Sciences*, 106(30): 12267–12272. 238
- Carlin, B. P. and Chib, S. (1995). “Bayesian model choice via Markov chain Monte

- Carlo methods.” *Journal of the Royal Statistical Society. Series B (Methodological)*, 473–484. 235
- Casarin, R., Dalla Valle, L., Leisen, F., et al. (2012a). “Bayesian model selection for beta autoregressive processes.” *Bayesian Analysis*, 7(2): 385–410. MR2934956. doi: <https://doi.org/10.1214/12-BA713>. 242
- Casarin, R., Valle, L., and Leisen, F. (2012b). “Bayesian model selection for beta autoregressive processes.” *Bayesian Analysis*, 7: 385–410. MR2934956. doi: <https://doi.org/10.1214/12-BA713>. 242
- Chan, J. C. C., Koop, G., Leon-Gonzalez, R., and Strachan, R. W. (2012). “Time varying dimension models.” *Journal of Business & Economic Statistics*, 30(3): 358–367. MR2969219. doi: <https://doi.org/10.1080/07350015.2012.663258>. 233
- Clyde, M., Desimone, H., and Parmigiani, G. (1996). “Prediction via orthogonalized model mixing.” *Journal of the American Statistical Association*, 91(435): 1197–1208. MR2695016. 235
- Cogley, T. and Sargent, T. J. (2005). “Drifts and volatilities: Monetary policies and outcomes in the post WWII U.S.” *Review of Economic Dynamics*, 8: 262–302. 264
- De Mol, C., Giannone, D., and Reichlin, L. (2008). “Forecasting using a large number of predictors: Is Bayesian shrinkage a valid alternative to principal components?” *Journal of Econometrics*, 146(2): 318–328. MR2465176. doi: <https://doi.org/10.1016/j.jeconom.2008.08.011>. 256
- Fan, J. and Li, R. (2001). “Variable selection via nonconcave penalized likelihood and its oracle properties.” *Journal of the American statistical Association*, 96(456): 1348–1360. MR1946581. doi: <https://doi.org/10.1198/016214501753382273>. 236, 250
- Frühwirth-Schnatter, S. (1994). “Data augmentation and dynamic linear models.” *Journal of Time Series Analysis*, 15: 183–202. MR1263889. doi: <https://doi.org/10.1111/j.1467-9892.1994.tb00184.x>. 235, 243
- Frühwirth-Schnatter, S. and Wagner, H. (2010). “Stochastic model specification search for Gaussian and partial non-Gaussian state space models.” *Journal of Econometrics*, 154(1): 85–100. MR2558953. doi: <https://doi.org/10.1016/j.jeconom.2009.07.003>. 233, 235
- Gefang, D. (2014). “Bayesian doubly adaptive elastic-net Lasso for VAR shrinkage.” *International Journal of Forecasting*, 30(1): 1–11. 264
- George, E. I. (1986a). “Combining minimax shrinkage estimators.” *Journal of the American Statistical Association*, 81(394): 437–445. MR0845881. 237
- George, E. I. (1986b). “Minimax multiple shrinkage estimation.” *The Annals of Statistics*, 14(1): 188–205. 237
- George, E. I. and McCulloch, R. E. (1993). “Variable selection via Gibbs sampling.” *Journal of the American Statistical Association*, 88(423): 881–889. 235, 237, 242, 253

- George, E. I. and McCulloch, R. E. (1997). “Approaches for Bayesian variable selection.” *Statistica Sinica*, 339–373. 235
- George, E. I., Sun, D., and Ni, S. (2008). “Bayesian stochastic search for VAR model restrictions.” *Journal of Econometrics*, 142(1): 553–580. MR2408749. doi: <https://doi.org/10.1016/j.jeconom.2007.08.017>. 235
- Giannone, D., Lenza, M., Momferatou, D., and Onorante, L. (2014). “Short-term inflation projections: a Bayesian vector autoregressive approach.” *International Journal of Forecasting*, 30(3): 635–644. 264
- Godsill, S., Doucet, A., and West, M. (2001). “Maximum a posteriori sequence estimation using Monte Carlo particle filters.” *Annals of the Institute of Statistical Mathematics*, 53(1): 82–96. MR1777255. doi: <https://doi.org/10.1023/A:1017968404964>. 244
- Green, P. J. (1990). “On use of the EM for penalized likelihood estimation.” *Journal of the Royal Statistical Society. Series B (Methodological)*, 443–452. MR1086796. 236, 252
- Groen, J. J. J., Paap, R., and Ravazzolo, F. (2013). “Real-time inflation forecasting in a changing world.” *Journal of Business & Economic Statistics*, 31(1): 29–44. MR3030796. doi: <https://doi.org/10.1080/07350015.2012.727718>. 233, 234, 235, 259
- Hastie, T. and Tibshirani, R. (1993). “Varying-coefficient models.” *Journal of the Royal Statistical Society. Series B (Methodological)*, 55: 757–796. MR1229881. 234
- Irie, K. and West, M. (2016). “Bayesian Emulation for Optimization in Multi-Step Portfolio Decisions.” *Submitted manuscript*. 238
- Ishwaran, H. and Rao, J. S. (2005). “Spike and slab variable selection: frequentist and Bayesian strategies.” *The Annals of Statistics*, 33: 730–773. MR2163158. doi: <https://doi.org/10.1214/009053604000001147>. 237
- Jagannathan, R. and Ma, T. (2003). “Risk reduction in large portfolios: Why imposing the wrong constraints helps.” *The Journal of Finance*, 58: 1651–1684. 238
- Kalli, M. and Griffin, J. E. (2014). “Time-varying sparsity in dynamic regression models.” *Journal of Econometrics*, 178: 779–793. MR3144682. doi: <https://doi.org/10.1016/j.jeconom.2013.10.012>. 233, 234, 238, 239, 244, 256, 257, 259, 260, 262
- Kalliovirta, L., Meitz, M., and Saikkonen, P. (2015). “A Gaussian mixture autoregressive model for univariate time series.” *Journal of Time Series Analysis*, 36(2): 247–266. MR3316469. doi: <https://doi.org/10.1111/jtsa.12108>. 235, 240, 241
- Kastner, G., Frühwirth-Schnatter, S., and Lopes, H. (2017). “Efficient Bayesian inference for multivariate factor stochastic volatility models.” *Journal of Computational and Graphical Statistics*, 26: 905–917. MR3765354. doi: <https://doi.org/10.1080/10618600.2017.1322091>. 243
- Kim, S., Shephard, N., and Chib, S. (1998). “Stochastic volatility: Likelihood inference and comparison with ARCH Models.” *Review of Economic Studies*, 65: 361–393. 243

- Koop, G. and Korobilis, D. (2012a). “Forecasting inflation using dynamic model averaging.” *International Economic Review*, 53(3): 867–886. MR2959266. doi: <https://doi.org/10.1111/j.1468-2354.2012.00704.x>. 233, 234, 259, 260, 261, 262
- Koop, G. and Korobilis, D. (2012b). “Forecasting inflation using dynamic model averaging.” *International Econometrics Reviews*, 30: 867–886. 234
- Korobilis, D. (2013). “VAR forecasting using Bayesian variable selection.” *Journal of Applied Econometrics*, 28(2): 204–230. MR3045864. doi: <https://doi.org/10.1002/jae.1271>. 264
- Kowal, D., Matteson, D., and Ruppert, D. (2019). “Dynamic shrinkage processes.” *Journal of the Royal Statistical Association (to appear)*, 1–30. MR3997101. doi: <https://doi.org/10.1111/rssb.12325>. 233
- Li, J. and Chen, W. (2014). “Forecasting macroeconomic time series: Lasso-based approaches and their forecast combinations with dynamic factor models.” *International Journal of Forecasting*, 30(4): 996–1015. 256
- Lopes, H., McCulloch, R., and Tsay, R. (2016). “Parsimony inducing priors for large scale state-space models.” *Submitted*. 235, 238, 243, 263, 264
- McKenzie, E. (1985). “An Autoregressive process for beta random variables.” *Management Science*, 31: 988–995. 242
- Meng, X. (1993). “Maximum likelihood estimation via the ECM algorithm: A general framework.” *Biometrika*, 80: 267–278. MR1243503. doi: <https://doi.org/10.1093/biomet/80.2.267>. 245
- Mitchell, T. J. and Beauchamp, J. J. (1988). “Bayesian variable selection in linear regression.” *Journal of the American Statistical Association*, 83: 1023–1032. MR0997578. 235, 241
- Nakajima, J. and West, M. (2013a). “Bayesian analysis of latent threshold dynamic models.” *Journal of Business & Economic Statistics*, 31: 151–164. MR3055329. doi: <https://doi.org/10.1080/07350015.2012.747847>. 233, 234, 235, 244, 264
- Nakajima, J. and West, M. (2013b). “Bayesian analysis of latent threshold dynamic models.” *Journal of Business & Economic Statistics*, 31: 151–164. 242, 243
- Park, T. and Casella, G. (2008). “The Bayesian LASSO.” *Journal of the American Statistical Association*, 103: 681–686. MR2524001. doi: <https://doi.org/10.1198/016214508000000337>. 242
- Pettenuzzo, D., Koop, G., Korobilis, D., et al. (2018). “Bayesian compressed vector autoregressions.” *Journal of Econometrics*, Forthcoming. MR3944767. doi: <https://doi.org/10.1016/j.jeconom.2018.11.009>. 264
- Phillips, P. (1991). “To Criticize the Critics: An Objective Bayesian Analysis of Stochastic Trends.” *Journal of Applied Econometrics*, 6: 333–364. 243
- Prado, R. and West, M. (2010). *Time Series: Modelling, Computation & Inference*. Chapman & Hall/CRC Press. 237, 243

- Primiceri, G. E. (2005). “Time varying structural vector autoregressions and monetary policy.” *Review of Economic Studies*, 72: 821–852. MR2148143. doi: <https://doi.org/10.1111/j.1467-937X.2005.00353.x>. 264
- Puelz, D., Hahn, P. R., and Carvalho, C. (2016). “Sparse mean-variance portfolios: A penalized utility approach.” *Submitted manuscript*. 238
- Rockova, V. (2018). “Bayesian Estimation of Sparse Signals with a continuous spike-and-slab prior.” *The Annals of Statistics*, 46: 401–437. MR3766957. doi: <https://doi.org/10.1214/17-AOS1554>. 235, 236, 237, 238, 239, 241, 245, 248, 249
- Rockova, V. and George, E. I. (2014). “EMVS: The EM approach to Bayesian variable selection.” *Journal of the American Statistical Association*, 109: 828–846. MR3223753. doi: <https://doi.org/10.1080/01621459.2013.869223>. 235, 244, 245, 254
- Rockova, V. and George, E. I. (2018). “The spike-and-slab LASSO.” *Journal of the American Statistical Association (to appear)*, 113: 431–444. MR3803476. doi: <https://doi.org/10.1080/01621459.2016.1260469>. 236, 245, 248, 249, 250
- Rockova, V. and McAlinn, K. (2020). “Dynamic Variable Selection with Spike-and-Slab Process Priors – Supplementary Material.” *Bayesian Analysis*. doi: <https://doi.org/10.1214/20-BA1199SUPP>. 243
- Schotman, P. and van Dijk, H. (1991). “A Bayesian analysis of the unit root in real exchange rates.” *Journal of Econometrics*, 49: 195–238. 263
- Stock, J. H. and Watson, M. W. (1999). “Forecasting inflation.” *Journal of Monetary Economics*, 44(2): 293–335. 259
- Stock, J. H. and Watson, M. W. (2007). “Why has US inflation become harder to forecast?” *Journal of Money, Credit and Banking*, 39: 3–33. 259, 260
- Stock, J. H. and Watson, M. W. (2012). “Generalized shrinkage methods for forecasting using many predictors.” *Journal of Business & Economic Statistics*, 30(4): 481–493. MR2990492. doi: <https://doi.org/10.1080/07350015.2012.715956>. 256
- Tibshirani, R., Saunders, M., Rosset, S., Zhu, J., and Knight, K. (2005). “Sparsity and smoothness via the fused LASSO.” *Journal of the Royal Statistical Society. Series B*, 67(1): 91–108. MR2136641. doi: <https://doi.org/10.1111/j.1467-9868.2005.00490.x>. 237, 238
- West, M. and Harrison, J. (1997). *Bayesian Forecasting & Dynamic Models*. Springer Verlag, 2nd edition. MR1482232. 234, 237, 243, 245, 257
- Wong, C. S. and Li, W. K. (2000). “On a mixture autoregressive model.” *Journal of the Royal Statistical Society: Series B*, 62(1): 95–115. MR1747398. doi: <https://doi.org/10.1111/1467-9868.00222>. 235
- Wong, C. S. and Li, W. K. (2001). “On a logistic mixture autoregressive model.” *Biometrika*, 88(1): 833–846. MR1859413. doi: <https://doi.org/10.1093/biomet/88.3.833>. 235

- Wood, S., Rosen, O., and Kohn, R. (2011). “Bayesian mixtures of autoregressive models.” *Journal of Computational and Graphical Statistics*, 20(2): 174–195. MR2816544. doi: <https://doi.org/10.1198/jcgs.2010.09174>. 235
- Wright, J. H. (2009). “Forecasting US inflation by Bayesian model averaging.” *Journal of Forecasting*, 28(2): 131–144. MR2518403. doi: <https://doi.org/10.1002/for.1088>. 233, 259
- Zellner, A. (1971). *An Introduction to Bayesian Inference in Econometrics*. Wiley, New York. MR0433791. 243
- Zhang, C.-H. (2010). “Nearly unbiased variable selection under minimax concave penalty.” *The Annals of Statistics*, 38: 894–942. MR2604701. doi: <https://doi.org/10.1214/09-AOS729>. 250
- Zhang, C.-H. and Zhang, T. (2012). “A general theory of concave regularization for high-dimensional sparse estimation problems.” *Statistical Science*, 27: 576–593. MR3025135. doi: <https://doi.org/10.1214/12-STS399>. 250
- Zou, H. and Hastie, T. (2005). “Regularization and variable selection via the elastic net.” *Journal of the Royal Statistical Society. Series B*, 67: 301–320. MR2137327. doi: <https://doi.org/10.1111/j.1467-9868.2005.00503.x>. 252

Acknowledgments

The authors would like to thank the Reviewers and the Associate Editor for providing thorough and thoughtful feedback which lead to substantial improvements of our paper.