STOCHASTIC SEARCH VARIABLE SELECTION IN VECTOR ERROR CORRECTION MODELS WITH AN APPLICATION TO A MODEL OF THE UK MACROECONOMY

MARKUS JOCHMANN,^{a,b} GARY KOOP,^{*b,c} ROBERTO LEON-GONZALEZ^{b,d}
AND RODNEY W. STRACHAN^{b,e}

^a *Newcastle University, Newcastle upon Tyne, UK*

^b *Rimini Centre for Economic Analysis, Italy*

^c *Department of Economics, University of Strathclyde, Glasgow, UK*

^d *National Graduate Institute for Policy Studies, Tokyo, Japan*

^e *The Australian National University, Canberra, Australia*

SUMMARY

This paper develops methods for stochastic search variable selection (currently popular with regression and vector autoregressive models) for vector error correction models where there are many possible restrictions on the cointegration space. We show how this allows the researcher to begin with a single unrestricted model and either do model selection or model averaging in an automatic and computationally efficient manner. We apply our methods to a large UK macroeconomic model. Copyright © 2011 John Wiley & Sons, Ltd.

Supporting information may be found in the online version of this article.

1. INTRODUCTION

Empirical macroeconomic modeling is typically a complicated affair, involving many time series variables. For instance, the UK macroeconomic models considered in this paper are based on the influential model of Garratt *et al.* (2003), hereafter GLPS, and Garratt *et al.* (2006). This involves nine variables. Even if one remains in a conventional multivariate time series framework involving vector autoregressive (VAR) and vector error correction models (VECMs), there can be hundreds or more modeling choices. Some of these are driven by econometric considerations (e.g. choice of cointegrating rank, lag length), while others can be driven by economic theory considerations (e.g. GLPS develop a theoretical model of the UK macroeconomy which implies certain restrictions on the cointegrating relationships). In macroeconomic modeling exercises involving structural breaks, time variation in parameters or regime switching such modeling choices proliferate (see, among many others, Primiceri, 2005; Sims and Zha, 2006).

The challenge facing the empirical macroeconomist is how to navigate through this morass of choices. One common strategy, used by GLPS, is to select a single model. This can be done through sequential hypothesis testing procedures or by choosing the model with the highest value for an information criterion or marginal likelihood. Alternatively, economic theory can be used to guide model choice by giving special consideration to models consistent with a particular economic theory.

A second strategy is to do model averaging and present empirical results which are a weighted average over all models. Common choices for weights are marginal likelihoods or functions of information criteria. Empirical papers adopting this approach include Sala-i-Martin *et al.*

---

* Correspondence to: Gary Koop, Department of Economics, University of Strathclyde, Sir William Duncan Building, 130 Rottenrow, Glasgow G4 0GE, UK. E-mail: Gary.Koop@strath.ac.uk

(2004), Fernández *et al.* (2001a,b) and Koop *et al.* (2008). The literature is replete with extensive discussion of the comparative advantages and disadvantages of model averaging and model selection (e.g. Leamer, 1978; Draper, 1995; Raftery *et al.*, 1997). A problem with model averaging is that it can be extremely computationally demanding, especially if, as in the current application, Markov chain Monte Carlo (MCMC) methods are required to carry out inference in each model and the number of models is high.

In this paper, using the modeling framework of GLPS, we use a third strategy which shares similarities with a type of approach that has become increasingly popular in macroeconomics. This strategy can informally be described as working with a single very flexible model. The problems with such models are that they risk over-fitting and typically involve a large number of parameters which can be difficult to estimate with any precision. The solution to such problems is prior information. This prior information can be purely subjective or involve some data information (as in the classic Minnesota prior for VAR models; see Doan *et al.*, 1984). Alternatively the prior can be hierarchical (i.e. can be expressed in terms of a set of unknown parameters which are estimated from the data). Indeed the state equations from state space models such as those used by Primiceri (2005) can be interpreted as hierarchical priors. The mixture innovation models of Giordani and Kohn (2008) or Koop *et al.* (2009), for example, can also be interpreted in this manner. Yet another recent hierarchical approach is explored and adapted in this paper. This is based on the stochastic search variable selection (SSVS) approach to VAR models developed in George *et al.* (2008). The basic idea of this approach can be explained quite simply. Let λ be a parameter (e.g. a coefficient in a VAR). A conventional Bayesian approach would specify a prior for λ such as $\lambda \sim N(0, \underline{V}_\lambda)$. By setting $\underline{V}_\lambda$ to a large value, a relatively noninformative prior is obtained. Smaller values of $\underline{V}_\lambda$ will shrink the coefficient towards zero—something which has been found to be useful in many empirical applications. Instead of such an approach, the SSVS prior is a hierarchical one, involving a mixture of two normal distributions:

$$\lambda|\delta \sim (1 - \delta) N(0, \underline{V}_{\lambda 0}) + \delta N(0, \underline{V}_{\lambda 1}) \quad (1)$$

where δ is a dummy variable that equals 0 if λ is drawn from the first normal and 1 if it is drawn from the second. The prior is hierarchical since δ is treated as an unknown parameter and estimated in a data-based fashion. The SSVS aspect of this prior arises by choosing the first prior variance, $\underline{V}_{\lambda 0}$, to be extremely ‘small’ (so that the parameter is virtually zero) and the second prior variance, $\underline{V}_{\lambda 1}$, to be ‘large’ (implying a relatively noninformative prior for the parameter). The way in which ‘small’ and ‘large’ prior variances can be chosen is discussed in George *et al.* (2008) and below.¹

SSVS can be interpreted as working with a single very flexible model and then (in a data-based fashion) choosing a prior which ensures over-parameterization problems do not occur. But SSVS can also be used as a way of implementing the two strategies described above. The researcher can do model selection by selecting a single restricted model based on some metric involving δ (e.g. if $\Pr(\delta = 1|\text{Data}) > 0.5$ then the explanatory variable corresponding to λ is included in the model, else it is excluded). Bayesian model averaging (BMA) can be done by using an unrestricted model with SSVS prior. This will implicitly average over models which include and exclude the explanatory variable associated with λ . The weights in this average will be $\Pr(\delta = 1|\text{Data})$ and $\Pr(\delta = 0|\text{Data})$, respectively.

¹ Formally, the SSVS prior allows for coefficients to be shrunk to ‘virtually zero’ but not ‘precisely zero’ since $\underline{V}_{\lambda 0}$ is not precisely zero. In the remainder of this paper, we will sometimes omit such qualifying phrases like ‘virtually zero’ and simply refer to SSVS as allowing for parameters to be included (if $\delta = 1$) or excluded (if $\delta = 0$) from the model.

The SSVS approach is very attractive for VAR modeling since VARs can have a large number of coefficients and many of them have no explanatory power. For instance, the trivariate VAR(4) used in Jochmann *et al.* (2010) has 39 VAR coefficients and SSVS picks out only 14 of these as having important explanatory power. In such a situation, the number of hypothesis tests required in a sequential testing procedure could lead the researcher to serious worries about pre-test problems. Alternatively, the researcher who wants to select the single restricted VAR with highest marginal likelihood or highest value for an information criteria would face the serious problem of estimating at least 2^{39} models (i.e. if each restricted model is defined according to whether each of the 39 coefficients is included or excluded, the result is 2^{39} models). Regardless of whether SSVS is used for model selection or model averaging, it seems an attractive approach to VAR modeling which will shrink unimportant coefficients to zero and thus yield a much more parsimonious model. In essence, the SSVS prior can pick out restrictions automatically so that the researcher can begin with a flexible unrestricted model and let SSVS pick out the appropriate restrictions.

As a digression, we mention that the SSVS approach of George *et al.* (2008) is only one of several approaches to stochastic variable selection that have been gaining popularity in econometrics (see, among others, Smith and Kohn, 1996; Fernández *et al.*, 2001a,b; Villani *et al.*, 2009; Frühwirth-Schnatter and Wagner, 2010). Unlike the approach adopted in this paper, which uses the so-called normal-normal prior of equation (1), some of the alternative approaches set the first element in the mixture to be a spike at zero (leading to a so-called point-normal prior). Discussion of the pros and cons of normal-normal and point-normal priors are given, for example, in Section 3.4 of Chipman *et al.* (2001). For instance, the point-normal prior has the potential drawback that in a regression context the posterior will asymptotically favor the inclusion of any explanatory variable even if its coefficient is very small (i.e. below the level of practical significance). Furthermore, in the VECM the ideas underlying the normal-normal prior used for the VAR generalize nicely, but under the point-normal prior they do not. In the statistical literature, discussion of the theoretical properties of various approaches are given in papers such as Ishwaran and Rao (2005) and Casella *et al.* (2009).

The major theoretical econometric contribution of this paper is to extend the SSVS approach, which is developed for multivariate normal linear models such as the VAR, to the non-linear VECM. As we shall see, some subtle econometric issues arise in this regard. A related theoretical contribution is to extend the SSVS approach to restrictions other than ones like (1) where each coefficient is either included or excluded. In particular, macroeconomists such as GLPS are often interested in overidentifying restrictions involving several coefficients. We develop an SSVS approach which can pick out such restrictions automatically.

The major empirical contribution is to use our methods in an important macroeconomic application involving a high-dimensional model. The GLPS model fits the bill perfectly. It involves nine dependent variables, uncertainty over many modeling choices (e.g. cointegrating rank, lag length and treatment of deterministic terms) and many theoretical restrictions motivated by macroeconomic theory. We find SSVS to be computationally efficient and to lead to quite parsimonious modeling. We find support for interest rate parity and Fisher inflation parity conditions which are imposed by GLPS. However, we find less support for purchasing power parity and none at all for a restriction motivated by neoclassical growth models. These latter two restrictions are imposed by GLPS. Nevertheless, when we turn to impulse response analysis we find economically sensible results which are not too different from those of GLPS.

2. MODELS

A Bayesian model is defined by a likelihood function and a prior. We deal with each of these in turn in this section.

2.1. Likelihoods

The VECM for an n -dimensional vector, y_t , is written as

$$\Delta y_t = \Pi y_{t-1} + \sum_{j=1}^{p-1} \Gamma_j \Delta y_{t-j} + \mu d_t + \varepsilon_t \quad (2)$$

for $t = 1, \dots, T$, where the $n \times n$ matrix Π is of rank $r \leq n$ and d_t denotes deterministic terms. ε_t is i.i.d. $N(0, \Sigma)$. The framework described in (2) defines a set of models which differ in the number of cointegrating vectors (r), lag length (p) and the specification of deterministic terms. In addition, the researcher may wish to entertain various over-identifying restrictions on the cointegrating vectors.

Written in matrix form, the unrestricted VECM used in this paper is

$$Y = X\beta\alpha + W\Gamma + E \quad (3)$$

where Y is $T \times n$ with t th row given by $\Delta y'_t$, X is $T \times (n+2)$ with t th row given by $(1, t, y'_{t-1})$, β is $(n+2) \times r_{\max}$, α is $r_{\max} \times n$, W is $T \times [n(p_{\max} - 1) + 2]$ with t th row given by $(1, t, \Delta y'_{t-1}, \dots, \Delta y'_{t-p_{\max}+1})$, Γ is $[n(p_{\max} - 1) + 2] \times n$ and E is $T \times n$ with $\text{vec}(E) \sim N(0, \Sigma \otimes I)$. $r_{\max}$ and $p_{\max}$ are the maximum number of cointegrating relationships and lag length, respectively, that the researcher is willing to entertain. Note that a deterministic trend in the cointegrating residuals has very different implications from a deterministic trend in the levels of the series. Accordingly, following Johansen (1995, Section 5.7), we potentially allow for deterministic terms in the cointegrating residual and directly in the conditional mean of Δy_t .

GLPS work with additional restrictions motivated by macroeconomic theory. A detailed explanation of this theory is provided in GLPS. Here we only sketch the details. The nine variables GLPS work with are the log of the oil price (p_t^o), the log of the effective exchange rate (e_t), the domestic and foreign nominal interest rates (i_t and i_t^*), the inflation rate (Δp_t), the logs of domestic and foreign real per capita GDP (q_t and q_t^*), the log of the ratio of domestic to foreign prices ($p_t - p_t^*$) and the log of the ratio of real high-powered money to real GDP per capita ($h_t - q_t$). Thus $y_t = (p_t^o, e_t, i_t^*, i_t, \Delta p_t, q_t, p_t - p_t^*, h_t - q_t, q_t^*)'$. Exact definitions and sources are given as online supporting information in Appendix C.

The model GLPS call their core model assumes five cointegrating relationships of which the following four place restrictions on β . These are

$$\begin{aligned} p_t - p_t^* - e_t &= b_1 + e_{1t}, \\ i_t - i_t^* &= b_2 + e_{2t}, \\ q_t - q_t^* &= b_3 + e_{3t}, \\ i_t - \Delta p_t &= b_4 + e_{4t} \end{aligned} \quad (4)$$

where b_i are intercepts in the cointegrating relationships and e_{it} denote mean zero stationary errors.²

GLPS derive and explain these cointegrating relationships in detail. Briefly, the first restriction is motivated by the purchasing power parity relationship; the second from an interest rate parity condition; the third by neoclassical growth models for the UK and the rest of the world; the fourth is based on Fisher inflation parity arguments. These equations define over-identifying restrictions and we consider variants of our models which impose some or all of them.

In sum, we have so far developed a very general unrestricted model. However, there are many restricted models of interest which differ in choice of $p \leq p_{\max}$, $r \leq r_{\max}$ and the number of over-identifying restrictions imposed (i.e. restrictions on the cointegration space). Furthermore, setting individual coefficients in α and/or Γ to zero (if empirically warranted) would lead to the desirable goal of a more parsimonious model. Note, in particular, that (3) allows for deterministic terms (both intercept and time trend) to enter both the cointegrating residual and Δy_t . Such a specification is typically too flexible in macroeconomic empirical work and, thus, restricted models which zero out some of these deterministic terms are potentially of interest.

In most applications, this will lead to a huge model space. For instance, even if we ignore restrictions on α and/or Γ , if we set $r_{\max} = 5$, $p_{\max} = 2$ and consider 18 restrictions on β ,³ we would have hundreds of models to work with. This is the approximate number of models used in the conventional Bayesian cointegration analysis of Koop *et al.* (2008). Given that the methodology of Koop *et al.* (2008) required the running of an MCMC algorithm in each model, the computational requirements of that paper were quite substantial. In the present application, α is a 5×9 matrix and Γ an 11×9 matrix and, thus, in total they contain 144 parameters. If we are to extend our model space by considering models which allow for each individual element of α and/or Γ to be zero, then the number of models is further multiplied by 2^{144} . Clearly, doing a conventional Bayesian model selection or model-averaging exercise which involves estimating each of more than 100×2^{144} models is computationally impossible. This consideration helps motivate our SSVS approach which, as we shall see, allows the researcher to work with the one model given in (3) and select appropriate restrictions in an automatic manner.

2.2. The SSVS Prior

In this paper, we will extend the SSVS prior for VARs of George *et al.* (2008) to VECM models. The new econometric issues involve β and α . However, before describing our treatment of these parameters we will briefly describe the prior for Γ to acquaint the reader with the basic ideas underlying SSVS. With one exception (noted below), our prior for Γ is identical to that used in George *et al.* (2008). Our treatment of Σ (which, apart from one exception, is also identical to that of George *et al.*, 2008) is discussed in Appendix A.

² GLPS use a long-run solvency requirement to motivate a fifth cointegrating relationship. However, it implies a cointegrating vector dependent on unknown parameters (i.e. it only implies the exclusion restriction that a cointegrating relationship must exist between two specific dependent variables and does not specify the exact cointegrating vector). When multiple cointegrating relationships exist, imposing such exclusion restrictions does not restrict the cointegrating space because by linearly combining cointegrating vectors one can come up with a new cointegrating vector which incorporates the long-run solvency restriction but does not restrict the coefficients on the other dependent variables. Dealing with such issues is possible using an extended version of our approach, but substantially complicates the analysis and, accordingly, we do not consider this fifth restriction in this paper.

³ These restrictions will be explained below. Basically, they arise from imposing the restrictions in (4) singly or in all possible combinations and restricting the deterministic terms in the cointegrating residuals.

2.2.1 Prior for Γ

Note that, conditional on β and α , the VECM can be written as

$$\tilde{Y} = W\Gamma + E \quad (5)$$

where $\tilde{Y} = Y - X\beta\alpha$. This takes the form of a VAR and the SSVS methods for VARs of George *et al.* (2008) can be applied directly in the context of an MCMC algorithm (except using $\tilde{Y}$ as the vector of dependent variables instead of Y).

SSVS can be interpreted as defining a hierarchical prior for Γ . Each element in the hierarchy is a mixture of two normals: one with a small variance (implying the coefficient is not in the model) and one with a large variance (implying the coefficient is included in the model). The SSVS prior for the i th element of $\gamma = \text{vec}(\Gamma)$ can be written as

$$\gamma_i | \delta_i \sim (1 - \delta_i)N(0, \underline{\kappa}_{0i}^2) + \delta_i N(0, \underline{\kappa}_{1i}^2) \quad (6)$$

The dummy variable δ_i is treated as an unknown parameter and is estimated in a data-based fashion. If $\delta_i = 0$, then the prior for γ_i has such a small variance that the posterior for γ_i is concentrated near zero. In this sense, we can say that, if $\delta_i = 0$, then the i th explanatory variable is excluded from the model (formally its coefficient is shrunk to be very near zero). If $\delta_i = 1$, then the i th explanatory variable is included and γ_i is estimated using a relatively noninformative prior. As discussed in the Introduction, priors of this form can be used for model selection by choosing a single model which includes only those explanatory variables for which $\Pr(\delta_i > 1 | \text{Data}) > c$ for some threshold c (e.g. $c = 0.5$). Alternatively, Bayesian model averaging can be done by running the MCMC algorithm for the unrestricted model and simply averaging any feature of interest (e.g. an impulse response function or a predictive moment) over all MCMC draws. Since the MCMC algorithm provides draws of the $\{\delta_i\}$ (as well as all other parameters), this strategy amounts to averaging over models (i.e. in the sense that different drawn values for the $\{\delta_i\}$ define different models). Crucially, though, doing model averaging in this way requires only that the researcher run a single MCMC algorithm for the unrestricted model (as opposed running an MCMC algorithm for each restricted model).

We use what George *et al.* (2008) call the ‘default semi-automatic approach’ to selecting the prior hyperparameters $\underline{\kappa}_{0i}^2$ and $\underline{\kappa}_{1i}^2$ and the reader is referred to their paper for additional justification for this approach. Basically, $\underline{\kappa}_{0i}^2$ should be selected so that γ_i is essentially zero and $\underline{\kappa}_{1i}^2$ should be selected so that γ_i is empirically substantive. The default semi-automatic approach involves choosing $\underline{\kappa}_{0i}^2 = c_0 \hat{V}(\gamma_i)$ and $\underline{\kappa}_{1i}^2 = c_1 \hat{V}(\gamma_i)$, where $\hat{V}(\gamma_i)$ is an estimate of the posterior variance of the coefficient. In our empirical work, this is based on a preliminary MCMC run using a noninformative prior, although other estimates (e.g. based on maximum likelihood estimation) are possible. The pre-selected constants c_0 and c_1 must have $c_0 \ll c_1$ and we set $c_0 = 0.1$ and $c_1 = 10$.

For the $\{\delta_i\}$, the SSVS prior posits that each element has a Bernoulli form (independent of the other elements of $\{\delta_i\}$). George *et al.* (2008), and early work with SSVS priors, simply set the Bernoulli prior hyperparameters to constants. In contrast, motivated by the arguments in Ley and Steel (2009), we treat them as unknown parameters which require their own prior. Thus, we have

$$\delta_i \sim \text{Bernoulli}(q_\delta), \quad q_\delta \sim \text{Beta}(\underline{a}_\delta, \underline{b}_\delta) \quad (7)$$

Ley and Steel (2009) show how such a prior can be used to elicit beliefs about model size. In their case, model size is the number of regressors in a regression model; in our case it is the number of non-zero coefficients in Γ . We set $\underline{a}_\delta = 1$ and $\underline{b}_\delta = 3$, which expresses a prior view that

1/4 of the coefficients in Γ are non-zero (but with a moderately large prior variance, expressing uncertainty about this view).

2.2.2 Prior for α and β

SSVS has never been extended to VECMs. The fact that $\beta\alpha$ enters in product form and the fact that (without further restrictions) the VECM only identifies the cointegrating space precludes the direct use of SSVS ideas for VARs. Before describing how SSVS can be implemented in a manner which overcomes these problems, we must digress and explain some basic issues that arise in Bayesian analyses of VECMs. An identification problem arises in the VECM since $\Pi = \beta\alpha$ and $\Pi = \beta CC^{-1}\alpha$ are identical for any nonsingular C . In early work, a common practice was to impose a linear normalization such as $\beta = \begin{pmatrix} I_r \\ \beta_0 \end{pmatrix}$. However, a report has recently emerged which argues that it is only the cointegrating space which is identified (see Strachan, 2003; Strachan and Inder, 2004; Villani, 2005, 2006) and this should be the focus of interest (rather than a particular identified parameter such as β_0). For instance, Strachan and Inder (2004) show how the use of linear identifying restrictions places a restriction on the estimable region of the cointegrating space. Furthermore, a flat and apparently ‘noninformative’ prior on β_0 in the linear normalization strongly favors regions of the cointegration space near where the linear normalization is invalid. Hence the linear normalization is used under the assumption that it is valid while at the same time the prior puts weight near the region where the normalization is likely to be invalid. These considerations suggest that an alternative identification restriction is required to form the basis for prior elicitation on the cointegration space. As shown in Strachan and Inder (2004), an identification restriction which does not suffer from the drawbacks of the linear and similar normalizations is

$$\beta'\beta = I \quad (8)$$

and we will impose this semi-orthogonality restriction on β throughout the remainder of this paper.

Formally, let $\text{sp}(\beta)$ denote the cointegration space (i.e. the space spanned by the columns of β) which is an element of the Grassmann manifold. The semi-orthogonality restriction (equation (8)), restricts the matrix of cointegrating vectors to the Stiefel manifold. These spaces are compact and hence a uniform distribution over them is appropriate (the integrating constant is given in Strachan and Inder, 2004). Strachan and Inder (2004) show how a sensible noninformative prior for β which does not restrict the cointegrating space is simply equal to this integrating constant with (8) imposed. This can be used as a noninformative prior on the cointegrating space.

Computation is complicated by the fact that β is semi-orthogonal and conventional algorithms such as the Gibbs sampler of Geweke (1996) cannot be used. The parameter-augmented Gibbs sampler developed in Koop *et al.* (2010) overcomes this complication by introducing a non-identified $r \times r$ symmetric positive definite matrix U with the property

$$\Pi = \beta\alpha = \beta U U^{-1}\alpha \equiv \beta^*\alpha^* \quad (9)$$

where $\alpha^* = U^{-1}\alpha$ and $\beta^* = \beta U$. The introduction of the non-identified U facilitates posterior computation because, under sensible priors (including the SSVS prior we adopt), the posterior conditional distributions of α^* and β^* in the MCMC algorithm are normal. The fact that U is not identified causes no problem and, in fact, introducing non-identified parameters in particular ways is a common strategy for improving the efficiency of MCMC algorithms in many contexts (see, for example, Liu and Wu, 1999). For proofs of all these statements, further discussion and exact details of the MCMC algorithm, see Koop *et al.* (2010).

Both for its own sake and since it will form a crucial part of our SSVS prior for VECMs, we also require an informative prior for the cointegrating space. Strachan and Inder (2004) and Koop *et al.* (2010) develop and motivate such a prior. Here we briefly summarize the elements of this strategy and refer the reader to these earlier papers for additional details. Restrictions on the cointegrating space can always be defined through a semi-orthogonal $(n+2) \times s$ matrix⁴ H with $r \leq s \leq n+2$ and $H_{\perp}$ being its orthogonal complement which is also a semi-orthogonal matrix. As an example, suppose we want to impose the four restrictions in (4) but otherwise want to leave the cointegrating space unrestricted. If the variables are ordered so that the t th row of X is $(1, t, p_t^o, e_t, i_t^*, i_t, \Delta p_t, q_t, p_t - p_t^*, h_t - q_t, q_t^*)$, then the following matrix can be seen to do this:⁵

$$H = \begin{pmatrix} 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ -1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & -1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & -1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & -1 & 0 & 0 & 0 & 0 \end{pmatrix} \quad (10)$$

Note that the first four columns express the four restrictions in (4). The fifth column states that any cointegrating relationship can have an intercept in it (but this H does not have a column which picks out t and so does not allow the deterministic trend to enter the cointegrating relationships). Note that the restrictions in (4) do not involve p_t^o and $h_t - q_t$, explaining the last two columns of H which simply say that there can potentially be some (unspecified) cointegrating relationships involving these variables. Koop *et al.* (2010) give some other examples of H and in Appendix A we will show how different combinations of the restrictions in (4) can be written in terms of appropriately defined H matrices.

Consider a prior for the cointegrating space defined through

$$vec(\beta^*) \sim N[0, I_{r_{\max}} \otimes (HH' + \tau H_{\perp} H'_{\perp})] \quad (11)$$

where $r_{\max}$ is the maximum cointegrating rank the researcher is willing to consider (e.g. $r_{\max} = n$). As shown in Strachan and Inder (2004), this is a sensible informative prior for the cointegrating space in that the prior for the cointegrating space will be centered over $sp(H)$. Furthermore, the dispersion of the prior is controlled by the scalar $\tau \in [0, 1]$ with $\tau = 0$ dogmatically imposing the restrictions expressed by H on the cointegrating space and $\tau = 1$ leading to the noninformative prior discussed previously.⁶

To see how this prior can be used to develop an SSVS prior for the cointegrating space, remember that the conventional implementation of SSVS implies a prior which is the mixture of two normals: one with a ‘small’ variance and the other with a ‘large’ variance. Thus a parametric restriction is

⁴ By allowing for an intercept and trend in the cointegrating residual, each column of β and hence H will have $n+2$ elements.

⁵ This H matrix is not semi-orthogonal, but can be transformed into a matrix which is semi-orthogonal and spans the same space via the transformation $H(H'H)^{-1/2}$.

⁶ Note that in our case we should not choose $\tau = 0$, since this would imply a singular prior covariance matrix in the typical case where $s < n+2$.

either imposed (approximately) or not imposed. For the cointegrating space, an SSVS prior could also be a mixture of two distributions. The first distribution would (approximately) restrict the cointegrating space to lie in the space spanned by H and the second would be a noninformative prior which would allow the cointegrating space to be estimated in an unconstrained manner. A consideration of the prior defined by (11) suggests how this can be done. That is, if we introduce a discrete random variable with two points of support—one which sets τ to a small value, the other which sets $\tau = 1$ —then we obtain a prior which can either (approximately) restrict the cointegration space (if τ is small) or not (if $\tau = 1$).

The previous discussion informally motivates how an SSVS prior for the cointegrating space can be developed when there is a single restriction (defined by H) involved. Extensions to multiple restrictions on the cointegrating space (e.g. our empirical work involves four over-identifying restrictions given in equation 4 which are imposed individually and jointly, plus we consider restricting the deterministic terms) can be done by mixing over more distributions. Formally, in our empirical work, we use an SSVS prior for the cointegrating space which takes the form

$$\text{vec}(\beta^*)|\tau, \phi \sim N[0, I_{r_{\max}} \otimes (H_{\phi}H'_{\phi} + \tau^{\varphi}H_{\perp\phi}H'_{\perp\phi})] \quad (12)$$

The discrete random variable $\phi \in \{1, 2, \dots, 18\}$ selects which restriction applies. Precise definitions of H_1 through H_{18} are provided as online supporting information in Appendix B. Briefly, $H_1 = I$ yields the noninformative prior discussed previously. H_2 restricts the deterministic trend in the cointegrating relationship to be zero and H_3 additionally restricts the intercept to be zero. H_4 through H_7 individually impose the four restrictions on the cointegrating space given in (4). H_8 through H_{13} impose all combinations of two of the four restrictions. H_{14} through H_{17} impose all combinations of three of the four restrictions. H_{18} imposes all of the restrictions. The exponent φ ensures that all priors have the same scale: $\varphi = 1$ if $\phi = 1$ and $\varphi = \frac{n+2-\tilde{s}}{n+2-s_{\phi}}$ if $\phi \in \{2, \dots, 18\}$. s_{ϕ} denotes the number of columns of H_{ϕ} and we set $\tilde{s} = \min\{s_{\phi}\} = 7$. Since ϕ is an unknown parameter, it requires a prior:

$$\Pr(\phi = i) = \underline{p}_{\phi i}, \quad \sum_{i=1}^{18} \underline{p}_{\phi i} = 1 \quad (13)$$

In our empirical work, we use the noninformative choice $\underline{p}_{\phi i} = \frac{1}{18}$ for all i .

We treat τ as an unknown parameter and set its prior to

$$\tau \sim \text{iGamma}(\underline{a}_{\tau}, \underline{b}_{\tau}) \quad (14)$$

We set $\underline{a}_{\tau} = 27$ and $\underline{b}_{\tau} = 1.3$ such that $E(\tau) = 0.05$ and $V(\tau) = 0.0001$. This choice is consistent with the ideas underlying SSVS. That is, τ plays a similar role to $\underline{\kappa}_{0j}^2$ and must be set to a small (but non-zero) value to ensure that the restriction is (to all intents and purposes) imposed on the cointegrating space. Our prior restricts τ (with high probability) to be less than 0.10, with most probability in the region of $\tau = 0.05$.

Finally, we require a prior for α^* . For this we use a standard SSVS prior of the sort used for VARs. In particular, we specify

$$a_i^*|\rho_i \sim (1 - \rho_i)N(0, \underline{\nu}_{0i}^2) + \rho_i N(0, \underline{\nu}_{1i}^2), \quad i = 1, \dots, nr_{\max} \quad (15)$$

where a_i^* is the i th element of $\text{vec}(\alpha^*)$ and $\underline{\nu}_{0i}^2$ and $\underline{\nu}_{1i}^2$ are set to small and big values (comparable to $\underline{\kappa}_{0i}^2$ and $\underline{\kappa}_{1i}^2$). For motivation of this prior (in terms of the shrinkage it achieves), see Koop *et al.*

(2010). ρ_i is a dummy variable (comparable to δ_i) which selects whether a_i^* is (approximately) zero or not. For similar reasons as given previously (see equation 7 and surrounding discussion), its prior is given by

$$\rho_i \sim \text{Bernoulli}(q_\rho), \quad q_\rho \sim \text{Beta}(\underline{a}_\rho, \underline{b}_\rho) \quad (16)$$

with $\underline{a}_\rho = 1$ and $\underline{b}_\rho = 3$. In our empirical work, we set $\underline{v}_{0i}^2 = c_0 \hat{V}(a_i^*)$ and $\underline{v}_{1i}^2 = c_1 \hat{V}(a_i^*)$, where $\hat{V}(a_i^*)$ is an estimate of the variance of a_i^* obtained from a preliminary MCMC run using a noninformative prior.⁷ The pre-selected constants c_0 and c_1 are set (as before) to be $c_0 = 0.1$ and $c_1 = 10$.

2.3. Posterior Simulation

The previous subsections outlined the likelihood and prior we use in this paper. Complete details of the MCMC algorithm used for posterior simulation are provided as online supporting information in Appendix B. Here we note that the algorithm combines blocks from the MCMC algorithm for VECMs described in Koop *et al.* (2010) and the MCMC algorithm for SSVS in VARs of George *et al.* (2008). That is, conditional on the SSVS dummy variables $\{\delta_i\}$, $\{\rho_i\}$, $\{\omega_i\}$ and ϕ (where the $\{\omega_i\}$ will be defined in Appendix A) we have a VECM with a particular prior and the algorithm of Koop *et al.* (2010) can be used. The dummy variables and ϕ can be drawn using the methods of George *et al.* (2008), with minor alterations in the case of ϕ . Exact formulae are given as online supporting information in Appendix B.

3. EMPIRICAL WORK

Our dataset is an updated version of the one constructed by GLPS and runs from 1965:Q1 to 2008:Q1. Remember that $y_t = (p_t^o, e_t, i_t^*, i_t, \Delta p_t, q_t, p_t - p_t^*, h_t - q_t, q_t^*)'$ and these variables have been described in Section 2 (with precise definitions and data sources given as online supporting information in Appendix C). Motivated by GLPS and preliminary experimentation with the data we select $r_{\max} = 5$ and $p_{\max} = 2$ as being reasonable values for the maximum possible number of cointegrating relationships and maximum lag length.⁸ All other modeling details, including prior hyperparameter values, are described in the previous section.

We divide our empirical work into two parts. First, using our SSVS approach, we present evidence on which models are supported by the data. Second, we present impulse responses using our SSVS approach (which we label ‘SSVS for Everything’) and two important special cases (which we label ‘Best Model’ and ‘GLPS Model’, respectively). The Best Model approach uses an initial SSVS run to select the single best model and then we estimate that model. We choose this best model as follows. For Γ and the off-diagonal elements of Σ^{-1} we include all coefficients with posterior inclusion probabilities greater than 0.5 (all other coefficients are set to zero). By ‘posterior inclusion probability’ we mean $\Pr(\delta_i | \text{Data})$ (when referring to the i th element of Γ). For the cointegrating space, we choose the value of ϕ which has highest posterior probability (this turns out to be $\phi = 12$). For α^* we retain the SSVS prior.⁹ The second special case involves using

⁷ To be precise, $\hat{V}(a_i^*)$ is the average of the variances from all elements of α^* that belong to the same equation as a_i^* .

⁸ These are the choices made by GLPS using a shorter dataset. With our updated dataset, the Johansen test (including an intercept) also selects $r = 5$ and the BIC selects $p = 1$, whereas AIC selects $p = 2$. We make the conservative choice of $p = 2$ since SSVS can, if it wishes, set coefficients for the longer lag length to zero.

⁹ We do this since our initial run of SSVS averages over different values of ϕ , not just a single best choice for ϕ . Given the fact that α^* and β^* enter in product form and ϕ determines β^* , SSVS treatment of α^* for fixed ϕ and SSVS treatment of α^* for random ϕ are conceptually and empirically very different. Hence it would be questionable to choose restrictions on α^* using an initial SSVS run treating ϕ as random and then use these restrictions in a model where ϕ is fixed.

our SSVS approach, but imposing the economic theory derived in GLPS on the model. We do this by using a dogmatic prior which attaches all prior weight to H_{18} (which imposes all four of the restrictions in equation 4). Formally, this amounts to setting $\Pr(\phi = 18) = 1$. All other modeling choices, including priors, are exactly the same as with SSVS for Everything.

3.1. Which Models are Supported by the Data?

Using SSVS for Everything we can calculate the probability associated with the various restrictions on the cointegration space (i.e. $\phi = 1, \dots, 18$), the probability associated with each cointegrating rank and the posterior inclusion probability of any coefficient.

Table I presents the posterior distribution of ϕ . It can be seen that there is no one restriction that is predominant, but rather the posterior probability is spread over many different values for ϕ . Imposing some restrictions clearly is preferred by the data since $\Pr(\phi = 1|\text{Data}) = 0$, which indicates no support for the noninformative prior. However, there is also little support for all of the restrictions holding since $\Pr(\phi = 18|\text{Data})$ is also near zero. The most support is for $\phi = 12$, which imposes interest rate parity and Fisher inflation parity. However, there is also substantial support for $\phi = 7$, which imposes only Fisher inflation parity. The next most popular choice is $\phi = 15$ which adds to $\phi = 12$ the additional restriction that purchasing power parity holds. There is never support for any value of ϕ which involves imposing the restriction inspired by neoclassical growth models for the UK and the rest of the world.

Much of the support for the various restrictions occurs jointly, e.g. the joint probability of purchasing power parity and uncovered interest rate parity is 0.059. To obtain the marginal probability of an individual restriction, therefore, we need to sum over all models which impose the restriction. Doing this we obtain: a marginal probability of purchasing power parity of 27.4%; a marginal probability of uncovered interest rate parity of 63.5%; a marginal probability of the neoclassical growth restriction of 0.9%; and a marginal probability of Fisher inflation parity of 87.9%.

When using these models for economic policy (e.g. to do impulse response analysis or forecasting), SSVS for Everything will average over these different choices for ϕ with weights given in Table I.

In terms of the cointegrating rank, note that $\Pi = \beta^* \alpha^*$ and our posterior simulation algorithm provides draws of β^* and α^* . The cointegrating rank is the rank of Π . Accordingly, we can shed light on the posterior of the cointegrating rank using these draws of Π . Formally, our posterior for the cointegrating rank is the posterior of the number of singular values of Π which are greater than 0.05. This is presented in Table II.

Table I. Posterior probabilities of ϕ

ϕ	Probability	ϕ	Probability
1	0.000	10	0.094
2	0.000	11	0.001
3	0.000	12	0.397
4	0.001	13	0.001
5	0.058	14	0.002
6	0.000	15	0.114
7	0.268	16	0.001
8	0.059	17	0.001
9	0.000	18	0.003

Table II. Posterior of cointegration rank

Rank	Probability
1	0.000
2	0.000
3	0.001
4	0.347
5	0.652

Table III. Posterior inclusion probabilities for Γ

	p^o	e	i^*	i	Δp_t	q	$p - p^*$	$h - q$	q^*
Const.	0.018	0.012	0.019	0.039	0.056	0.014	0.015	0.015	0.015
Time	0.019	0.016	0.018	0.027	0.042	0.019	0.012	0.013	0.016
p^o	0.591	0.041	0.053	0.055	0.142	0.034	0.045	0.031	0.037
e	0.042	0.624	0.050	0.045	0.074	0.050	0.170	0.044	0.035
i^*	0.053	0.038	0.997	0.583	0.042	0.042	0.047	0.031	0.072
i	0.080	0.100	0.082	0.299	0.041	0.050	0.043	0.041	0.094
Δp_t	0.053	0.043	0.097	0.041	0.085	0.174	0.048	0.083	0.048
q	0.062	0.045	0.083	0.139	0.055	0.053	0.039	0.043	0.046
$p - p^*$	0.075	0.067	0.181	0.040	0.679	0.078	0.793	0.063	0.040
$h - q$	0.036	0.090	0.034	0.076	0.079	0.052	0.076	0.744	0.056
q^*	0.053	0.036	0.572	0.043	0.041	0.031	0.084	0.050	0.333

Table II shows that there is strong support for a cointegrating rank of either four or five. This is consistent with the non-Bayesian analysis of GLPS who (with a shorter dataset) find a cointegrating rank of five.

Table III presents posterior inclusion probabilities for each coefficient in the 11×9 matrix Γ . Of these 99 coefficients, only eight have posterior inclusion probabilities of greater than 0.5.¹⁰ Often it is the case that the own lag of the dependent variable is the only explanatory variable selected, but there are a few exceptions to this pattern. This reinforces our story that SSVS is an effective method of achieving parsimony in multivariate time series models such as the VECM.

Table IV also reveals a similar pattern of parsimony. It presents posterior inclusion probabilities for elements of the upper-triangular matrix Ψ , where $\Sigma^{-1} = \Psi\Psi'$ (see Appendix A for details). For most of these the posterior inclusion probability is much less than 0.5.

Table IV. Posterior inclusion probabilities for Ψ

	p^o	e	i^*	i	Δp_t	q	$p - p^*$	$h - q$	q^*
p^o		0.235	0.603	0.353	0.188	0.370	0.998	0.385	0.215
e			0.187	0.994	0.194	0.206	0.261	0.259	0.260
i^*				0.918	0.298	0.195	0.483	0.180	0.403
i					0.229	0.259	0.187	0.186	0.252
Δp_t						0.444	0.997	0.589	0.194
q							0.659	1.000	0.991
$p - p^*$								0.592	0.270
$h - q$									0.214
q^*									

¹⁰ The reason that the inclusion probabilities for the intercept are so low is that we are also including an intercept in the cointegrating residual.

Figure 1. Responses to a monetary shock. Solid line, SSVS for Everything; dots, Best Model; dashes, GLPS model

The results in these tables are empirically sensible. Where comparable, they are broadly consistent with those of GLPS. To provide further evidence that our approach is yielding sensible results, Appendix D (online supporting information) presents Bayes factors comparing a subset of our models¹¹ using a non-SSVS prior. We refer the reader there for details, and note here only that results are quite similar to those presented here.

¹¹ For reasons outlined above, to calculate Bayes factors comparing all of the models in our model space would be computationally infeasible.

Figure 2. Responses to an oil price shock. Solid line, SSVS for Everything; dots, best model; dashes, GLPS model

3.2. Impulse Response Analysis

We have seen how SSVS does successfully allow us to work with a very flexible unrestricted model and allow for a more parsimonious restricted model to be picked out automatically. We stress that the alternative strategy of evaluating every restricted model with the goal of either selecting a single one or doing model averaging is computationally infeasible when facing the number of models that we are dealing with here. Now we present evidence that SSVS is picking out economically sensible empirical models through a consideration of impulse responses.

GLPS use a particular set of restrictions which allows them to identify a monetary policy shock and an oil shock. We adopt the same identifying scheme and refer the reader to GLPS for

precise details. Point estimates (posterior means) of the impulse responses to a monetary shock and oil price shock for our three cases (SSVS for Everything, Best Model and GLPS model) are plotted in Figures 1 and 2, respectively. It can be seen that the three approaches are giving very similar impulse responses. Furthermore, these impulse responses are all of reasonable shape and magnitude. Qualitatively, they are similar to those given in Figures 2 and 4 of GLPS despite the fact that we are using a substantially longer data span. The unrestricted VECM produced quite different impulse responses. We find this reassuring in the sense that it is providing evidence that SSVS is successfully choosing a parsimonious and economically sensible model, but in an automatic and computationally efficient manner. We are not carrying out a long series of sequential hypothesis tests nor are we doing Bayesian model averaging in the conventional and computationally demanding manner of Koop *et al.* (2008).

A more detailed set of impulse response functions, including 95% credible intervals, is provided as online supporting information in Appendix D. These credible intervals (similar to the 95% bootstrapped confidence intervals used in GLPS) tend to be fairly wide. Nevertheless, the same pattern of similarity across approaches (and to GLPS) is retained.

4. CONCLUSIONS

SSVS methods typically involve, for each parameter, choosing between a tight prior and a loose prior in a data-based fashion. The tight prior is constructed so as to shrink the parameter so that it is effectively zero. In this paper, we have extended these ideas to VECMs and to the case where the researcher is interested in restrictions involving several parameters. These extensions required us to define carefully what a ‘tight prior’ and a ‘loose prior’ was in the case of the cointegration space. An MCMC algorithm was developed for carrying out empirical inference in the resulting model.

In an empirical exercise involving a nine-variate macroeconomic model of the UK economy, we find our SSVS methods for VECMs to work well and stress that conventional BMA methods would simply be computationally infeasible in this case. Despite beginning with an over-parameterized unrestricted model and facing a huge model space, SSVS quickly moved towards much more parsimonious models. Regardless of whether MCMC output was used for model averaging or model selection, the resulting impulse responses were sensible and similar to those found using the model of GLPS.

ACKNOWLEDGEMENT

We acknowledge financial support from the Leverhulme Trust under Grant F/00 273/J.

REFERENCES

- Casella G, Giron F, Martinez M, Moreno E. 2009. Consistency of Bayesian procedures for variable selection. *Annals of Statistics* **37**: 1207–1228.
- Chipman H, George E, McCulloch R. 2001. The practical implementation of Bayesian model selection. In *Institute of Mathematical Statistics Lecture Notes: Monograph Series*, Vol. 38, Lahiri P (ed.); 65–134.
- Doan T, Litterman R, Sims C. 1984. Forecasting and conditional projection using realistic prior distributions. *Econometric Reviews* **3**: 1–100.
- Draper D. 1995. Assessment and propagation of model uncertainty (with discussion). *Journal of the Royal Statistical Society, Series B* **56**: 45–98.
- Fernández C, Ley E, Steel M. 2001a. Model uncertainty in cross-country growth regressions. *Journal of Applied Econometrics* **16**: 563–576.

- Fernández C, Ley E, Steel M. 2001b. Benchmark priors for Bayesian model averaging. *Journal of Econometrics* **100**: 381–427.
- Frühwirth-Schnatter S, Wagner H. 2010. Stochastic model specification search for Gaussian and partially non-Gaussian state space models. *Journal of Econometrics* **154**: 85–100.
- Garratt A, Lee K, Pesaran MH, Shin Y. 2003. A long-run structural macroeconomic model of the UK economy. *Economic Journal* **113**: 412–455.
- Garratt A, Lee K, Pesaran MH, Shin Y. 2006. *Global and National Macroeconometric Modelling: A Long-Run Structural Approach*. Oxford University Press: Oxford.
- George E, Sun D, Ni S. 2008. Bayesian stochastic search for VAR model restrictions. *Journal of Econometrics* **142**: 553–580.
- Geweke J. 1996. Bayesian reduced rank regression in econometrics. *Journal of Econometrics* **75**: 121–146.
- Giordani P, Kohn R. 2008. Efficient Bayesian inference for multiple change-point and mixture innovation models. *Journal of Business and Economic Statistics* **26**: 66–77.
- Ishwaran H, Rao J. 2005. Spike and slab variable selection: frequentist and Bayesian strategies. *Annals of Statistics* **33**: 730–773.
- Jochmann M, Koop G, Strachan R. 2010. Bayesian forecasting using stochastic search variable selection in a VAR subject to breaks. *International Journal of Forecasting* **26**: 326–347.
- Johansen S. 1995. *Likelihood-Based Inference in Cointegrated Vector Autoregressive Models*. Oxford University Press: Oxford.
- Koop G, Potter S, Strachan R. 2008. Re-examining the consumption–wealth relationship: the role of model uncertainty. *Journal of Money, Credit and Banking* **40**: 341–367.
- Koop G, Leon-Gonzalez R, Strachan R. 2009. On the evolution of the monetary policy transmission mechanism. *Journal of Economic Dynamics and Control* **33**: 997–1017.
- Koop G, Leon-Gonzalez R, Strachan R. 2010. Efficient posterior simulation for cointegrated models with priors on the cointegration space. *Econometric Reviews* **29**: 224–242.
- Leamer E. 1978. *Specification Searches*. Wiley: New York.
- Ley E, Steel M. 2009. On the effect of prior assumptions in Bayesian model averaging with applications to growth regression. *Journal of Applied Econometrics* **24**: 651–674.
- Liu J, Wu Y. 1999. Parameter expansion for data augmentation. *Journal of the American Statistical Association* **94**: 1264–1274.
- Primiceri G. 2005. Time varying structural vector autoregressions and monetary policy. *Review of Economic Studies* **72**: 821–852.
- Raftery A, Madigan D, Hoeting J. 1997. Bayesian model averaging for linear regression models. *Journal of the American Statistical Association* **92**: 179–191.
- Sala-i-Martin X, Doppelhofer G, Miller R. 2004. Determinants of long-term growth: a Bayesian averaging of classical estimates (BACE) approach. *American Economic Review* **94**: 813–835.
- Sims C, Zha T. 2006. Were there regime switches in macroeconomic policy? *American Economic Review* **96**: 54–81.
- Smith M, Kohn R. 1996. Nonparametric regression using Bayesian variable selection. *Journal of Econometrics* **75**: 317–343.
- Smith M, Kohn R. 2002. Parsimonious covariance matrix estimation for longitudinal data. *Journal of the American Statistical Association* **97**: 1141–1153.
- Strachan R. 2003. Valid Bayesian estimation of the cointegrating error correction model. *Journal of Business and Economic Statistics* **21**: 185–195.
- Strachan R, Inder B. 2004. Bayesian analysis of the error correction model. *Journal of Econometrics* **123**: 307–325.
- Villani M. 2005. Bayesian reference analysis of cointegration. *Econometric Theory* **21**: 326–357.
- Villani M. 2006. Bayesian point estimation of the cointegration space. *Journal of Econometrics* **134**: 645–664.
- Villani M, Kohn R, Giordani P. 2009. Regression density estimation using smooth adaptive gaussian mixtures. *Journal of Econometrics* **153**: 155–173.

APPENDIX A: FURTHER DETAILS ON PRIOR ELICITATION

Prior for Σ

Our SSVS prior for Σ is the same as in George *et al.* (2008). We begin by decomposing it as

$$\Sigma^{-1} = \Psi\Psi' \quad (\text{A1})$$

where Ψ is upper-triangular. The SSVS prior involves using a standard Gamma prior for the square of each of the diagonal elements of Ψ and the SSVS mixture of normals prior for each element above the diagonal. Note that this implies that the diagonal elements of Ψ are always included in the model, ensuring a positive definite error covariance matrix. This form for the prior also greatly simplifies posterior computation. Precise details are provided in the next paragraph.

Let the non-zero elements of Ψ be labeled as ψ_{ij} . For the diagonal elements, we assume prior independence with

$$\psi_{jj}^2 \sim \text{Gamma}(\underline{a}_j, \underline{b}_j) \quad (\text{A2})$$

where $\text{Gamma}(\underline{a}_j, \underline{b}_j)$ denotes the Gamma distribution with mean $\underline{a}_j/\underline{b}_j$ and variance $\underline{a}_j/\underline{b}_j^2$. We specify $\underline{a}_j = 0.1$ and $\underline{b}_j = 0.001$, which are relatively noninformative choices, but center the prior at 100.0. The latter is a sensible value given the scale of the variables and the fact that ψ_{jj}^2 is the inverse of the error variance in equation j .

The SSVS prior for the off-diagonal elements has

$$\psi_{ij}|\omega_{ij} \sim (1 - \omega_{ij})\text{N}(0, \underline{\xi}_{0ij}^2) + \omega_{ij}\text{N}(0, \underline{\xi}_{1ij}^2) \quad (\text{A3})$$

where ω_{ij} is a dummy variable. As before, we use a semi-automatic default approach to selecting $\underline{\xi}_{0ij}^2$ and $\underline{\xi}_{1ij}^2$. That is, we set $\underline{\xi}_{0ij}^2 = c_0 \hat{V}(\psi_{ij})$ and $\underline{\xi}_{1ij}^2 = c_1 \hat{V}(\psi_{ij})$, where $\hat{V}(\psi_{ij})$ is based on an estimate of the variance of the appropriate element of Ψ . This estimate is obtained from a preliminary MCMC run using a noninformative prior. The pre-selected constants c_0 and c_1 are set (as before) to be $c_0 = 0.1$ and $c_1 = 10$. This treatment of the error covariance matrix is the same as in George *et al.* (2008). However, it is worth noting that alternative treatments are possible (see, for example, Smith and Kohn, 2002).

Finally, we take

$$\omega_{ij} \sim \text{Bernoulli}(q_\omega), \quad q_\omega \sim \text{Beta}(\underline{a}_\omega, \underline{b}_\omega) \quad (\text{A4})$$

with $\underline{a}_\omega = 1$ and $\underline{b}_\omega = 3$.

Prior for the Cointegration Space

Our SSVS prior for VECMs requires the construction of H_i for $i = 1, \dots, 18$ that impose the over-identifying restrictions or restrict the manner in which deterministic trends enter the cointegrating residuals. The manner in which such matrices are constructed is described in Strachan and Inder (2004) and the reader is referred to that paper for details. The matrices below are not semi-orthogonal, but since $sp(H) = sp(H\kappa)$ for any full rank square κ , after constructing H , we may innocuously make it semi-orthogonal by the transformation $H \rightarrow H(H'H)^{-1/2}$.

Noting that the t th row of X is ordered as $(1, t, p_t^o, e_t, i_t^*, i_t, \Delta p_t, q_t, p_t - p_t^*, h_t - q_t, q_t^*)$, the following restriction matrices can be constructed. The matrix $H_1 = I$ imposes no restrictions at all and simply leads to the noninformative prior for the cointegrating space defined in the body of the text. H_2 deletes the deterministic trend term, but otherwise leaves the cointegrating

H_{14} imposes the first, second and third of the restrictions in (4), H_{15} imposes the first, second and fourth.

$$H_{14} = \begin{pmatrix} 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ -1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & -1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & -1 & 0 & 0 & 0 & 0 \end{pmatrix}, \quad H_{15} = \begin{pmatrix} 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ -1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & -1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & -1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{pmatrix}$$

H_{16} imposes the first, third and fourth of the restrictions in (4), H_{17} the second, third and fourth.

$$H_{16} = \begin{pmatrix} 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ -1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & -1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & -1 & 0 & 0 & 0 & 0 & 0 \end{pmatrix}, \quad H_{17} = \begin{pmatrix} 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ -1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & -1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & -1 & 0 & 0 & 0 & 0 & 0 & 0 \end{pmatrix}$$

Finally, H_{18} is the matrix which imposes all of the restrictions in (4) and is given in (10).

Note that, with the exception of H_1 , all of these matrices delete the deterministic term from the cointegrating relationship. Furthermore, it is common (e.g. Johansen, 1995, Section 5.7), to avoid identification concerns by imposing the restriction that the constant and trend do not enter in both X and W . In the few cases this issue is relevant, note that we do not impose this restriction in our algorithm. But in practice SSVS virtually never chooses constants and trends in both X and W at the same time.