

Ridge Regularized Estimation of VAR Models for Inference

Giovanni Ballarin*
University of Mannheim

July 7, 2023

Abstract: Ridge regression is a popular method for dense least squares regularization. In this work, ridge regression is studied in the context of VAR model estimation and inference. The implications of anisotropic penalization are discussed and a comparison is made with Bayesian ridge-type estimators. The asymptotic distribution and the properties of cross-validation techniques are analyzed. Finally, the estimation of impulse response functions is evaluated with Monte Carlo simulations and ridge regression is compared with a number of similar and competing methods.

Keywords: ridge regularization, vector autoregression, inference, impulse responses

JEL codes: C32, C51, C52

*Department of Economics, University of Mannheim, L7, 3-5, Mannheim, 68131, Germany. Giovanni.Ballarin@gess.uni-mannheim.de. I am grateful for the comments and suggestions from Lyudmila Grigoryeva, So Jin Lee, Thomasz Olma, Oliver Pfäuti and Mikkel Plagborg-Møller, and the seminar participants at the University of Mannheim, the HKMetrics Workshop and the Young Researchers Workshop on Big and Smart Data Analysis in Finance. I am especially thankful to Claudia Noack for pointing out an important error in a previous version of this paper, as well as Jonas Krampe and Carsten Trenkler for their insightful discussions which helped develop this work significantly. Lastly, I wish to thank Peter C. B. Phillips, Atsushi Inoue and many other colleagues for the suggestion to consider adding a formal analysis of cross-validation in the paper. The author acknowledges support by the state of Baden-Württemberg through bwHPC.

1 Introduction

While the idea of using ridge regression for vector autoregressive model estimation dates back to [Hamilton \(1994\)](#), there seems to be no complete analysis of its properties and asymptotic theory in the literature. This paper fills this gap by analyzing the geometric and distributional properties of ridge in a VAR estimation framework, discussing its comparison to well-known Bayesian approaches and deriving the validity of cross-validation as a selection procedure for the ridge penalty.

First, I show that the shrinkage induced by the ridge estimator, while intuitive in the setting of an isotropic penalty, produces complex effects when estimating a VAR model with a more flexible penalization scheme. This implies that the benefits of the bias-variance trade-off ([Hastie, 2020](#)) may be hard to gauge a priori. I provide a tractable example where ridge can yield estimates that have higher autoregressive dependence than the least squares solution. To better understand how different ridge penalization strategies can be designed, I also make a comparison with Bayesian VAR estimators commonly used in macroeconomic practice.

Second, I generalize the analysis of [Fu and Knight \(2000\)](#) and prove the consistency and asymptotic normality of the ridge estimator, a result that seems to be missing in the literature. For standard inference, the ridge penalty should either be negligible in the limit or its centering converge in probability to the true parameter vector. In both these cases, there is no asymptotic bias and no reduction in variance. Alternatively, in settings where a researcher is willing to assume that a subset of the VAR parameters features small coefficients, one can achieve an asymptotic reduction of variance by correctly tuning the ridge penalty matrix. I further derive the properties of cross-validation, which is a popular approach in practical applications to tune penalized estimators ([Hastie et al., 2009](#), [Bergmeir et al., 2018](#)). More specifically, I show that cross-validation is able to select penalties that are asymptotically valid for inference. In passing, I also prove that in an autoregressive setup the time dependence of regressors has an exponentially small effect on in-sample prediction error evaluation.

Lastly, I use Monte Carlo simulations to study the performance of the different ridge approaches discussed, focusing on impulse response inference. I consider two exercises: one is based on a three-variable VARMA(1,1) data generating process from [Kilian and Kim \(2011\)](#); the other is a VAR(5) model estimated in levels from a set of seven macroeconomic series, following [Giannone et al. \(2015\)](#). The finding is that ridge can lead to improvements over unregularized methods in impulse response confidence interval length, while Bayesian estimators show the best overall performance due to the underlying flexibility of their priors.

Related Literature. This paper does not discuss the high-dimensional setting, where the number of regressors grows together with the sample size. Some important work has been done in this direction already. [Dobriban and Wager \(2018\)](#) derive an explicit expression for the predictive risk of ridge regression assuming a high-dimensional random effects model. Other works in this vein are [Liu and Dobriban \(2020\)](#), [Patil et al. \(2021\)](#) and [Hastie et al. \(2022\)](#), which are mostly focused on penalty selection by cross-validation as well as structural features of ridge. Generally speaking, the complexity of analyzing ridge regression in high dimensions is a challenge to precisely understanding its practical implications. As I show below, in the context of finite-dimensional VARs, asymptotic inference demands that the ridge penalty becomes asymptotically negligible at appropriate rates. Thus, a challenge is understanding in what way high-dimensional time series problems would benefit from the use of ridge. This question is beyond the scope of this paper.

In the time series forecasting literature, ridge regression is commonly used for prediction. I provide a partial list of contributions in this direction. [Inoue and Kilian \(2008\)](#) use ridge regularization for forecasting U.S. consumer price inflation and argue that it compares favorably with bagging techniques; [De Mol et al. \(2008\)](#) use a Bayesian VAR with posterior mean equivalent to a ridge regression in forecasting; [Ghosh et al. \(2019\)](#) again study the Bayesian ridge, this time however in the high-dimensional context; [Goulet Coulombe et al. \(2022\)](#), [Fuleky \(2020\)](#), [Babii et al. \(2021\)](#), and [Medeiros et al. \(2021\)](#) compare LASSO, ridge and other machine learning techniques for forecasting with large economic datasets. [Fuleky \(2020\)](#) gives a textbook treatment of penalized time series estimation, including ridge, but does not discuss inference. The ridge penalty is considered within a more general mixed ℓ_1 - ℓ_2 penalization setting in [Smeekes and Wijler \(2018\)](#), who study the performance and robustness of penalized estimates for constructing forecasts.

Regarding inference, [Li et al. \(2022\)](#) provided a general exploration of shrinkage procedures in the context of structural impulse response estimation. Very recently, [Cavaliere et al. \(2022\)](#) suggested a methodology for inference on ridge-type estimators that relies on bootstrapping. Finally, shrinkage of autoregressive models to constrained sub-models was discussed by [Hansen \(2016b\)](#) in a more general setting.

Finally, various estimation problems can either be cast as or augmented with ridge-type regressions. [Goulet Coulombe \(2023\)](#) shows that the estimation of VARs with time-varying parameters can be written as ridge regression. [Plagborg-Møller \(2016\)](#) and [Barnichon and Brownlees \(2019\)](#) both use ridge to derive smoothed local projection impulse response functions.

Outline. Section 2 provides a discussion of the ridge penalty and the ridge VAR estimator. In Section 3 I deal with the properties of ridge-induced shrinkage in the autoregressive coefficients. I discuss the connections between frequentist and Bayesian ridge for VAR models within Section 4. Section 5 develops the asymptotic theory and inference result in the case where there is no asymptotic shrinkage. This includes studying the property of cross-validation under dependence. Section 6 provides inference and CV results in a setting where some shrinkage of a subset of parameters is possible. Section 7 presents Monte Carlo simulations focused on impulse response estimation. Section 8 concludes. Finally, the Appendix contains more detailed derivations, as well as all proofs and supplementary information.

Notation. Define \mathbb{R}_+ to be the set of strictly positive real numbers. Vectors $v \in \mathbb{R}^N$ and matrices $A \in \mathbb{R}^{N \times M}$ are always denoted with lower and upper-case letters, respectively. Throughout, I will use I_M to represent the identity matrix of dimension M . For any vector $v \in \mathbb{R}^N$, $\|v\|$ is the Euclidean norm. For any matrix $A \in \mathbb{R}^{N \times M}$, $\|A\|$ is the spectral norm unless stated otherwise; $\|A\|_{\max} = \max_{i,j} |a_{ij}|$ is the maximal entry norm; $\|A\|_F = (\text{tr}\{A'A\})^{-1/2}$ is the Frobenius norm; $\text{vec}(\cdot)$ is the vectorization operator and \otimes is the Kronecker product (Lütkepohl, 2005). If a vector represents a vectorized matrix, then it will be written in bold, that is, for $A \in \mathbb{R}^{N \times M}$ I write $\text{vec}(A) = \mathbf{a} \in \mathbb{R}^{NM}$. Let $\Lambda = \text{diag}\{\lambda_1, \dots, \lambda_{K^2p}\}$, $\lambda_i > 0$ for all $i = 1, \dots, K^2p$. To give the partial ordering of diagonal positive semi-definite penalization matrices, let $\Lambda_1 = \text{diag}\{\lambda_{1,j}\}_{j=1}^{K^2p}$ and $\Lambda_2 = \text{diag}\{\lambda_{2,j}\}_{j=1}^{K^2p}$. I write $\Lambda_1 \prec \Lambda_2$ if $\lambda_{1,i} < \lambda_{2,i}$ for all $i = 1, \dots, K^2p$; $\Lambda_1 \preceq \Lambda_2$ if $\lambda_{1,i} \leq \lambda_{2,i}$ for all i and $\exists j \in 1, \dots, K^2p$ such that $\lambda_{1,j} < \lambda_{2,j}$. Symbols \xrightarrow{P} and \xrightarrow{d} are used to indicate convergence in probability and convergence in distribution, respectively.

2 Ridge Regularized VAR Estimation

Let $y_t = (y_{1t}, \dots, y_{Kt})'$ be a K -dimensional vector autoregressive process with lag length $p \geq 1$ and parametrization

$$y_t = \nu_t + A_1 y_{t-1} + A_2 y_{t-2} + \dots + A_p y_{t-p} + u_t, \quad (1)$$

where $u_t = (u_{1t}, \dots, u_{Kt})'$ is additive noise such that u_t are identically, independently distributed with $\mathbb{E}[u_{it}] = 0$ and $\text{Var}[u_t] = \Sigma_u$, and ν_t is a deterministic trend. For simplicity, in the remainder I shall assume that $\nu_t = 0$ so that y_t has no trend component – equivalently, y_t is a de-trended series.

For a given sample size T define $Y = (y_1, \dots, y_T) \in \mathbb{R}^{K \times T}$, $z_t = (y'_t, y'_{t-1}, \dots, y'_{t-p+1})' \in \mathbb{R}^{Kp}$, $Z = (z_0, \dots, z_{T-1}) \in \mathbb{R}^{Kp \times T}$, $B = (A_1, \dots, A_p) \in \mathbb{R}^{K \times Kp}$, $U = (u_1, \dots, u_T) \in \mathbb{R}^{K \times T}$,

and vectorized counterparts $\mathbf{y} = \text{vec}(Y)$, $\boldsymbol{\beta} = \text{vec}(B)$ and $\mathbf{u} = \text{vec}(U)$. Accordingly, $Y = BZ + U$ and $\mathbf{y} = (Z' \otimes I_K)\boldsymbol{\beta} + \mathbf{u}$, where $\Sigma_{\mathbf{u}} = I_K \otimes \Sigma_u$. Importantly, throughout this work, I will assume that the cross-sectional dimension K remains fixed.

Ridge regularization is a modification of the least squares objective by the addition of a term dependent on the Euclidean norm of the coefficient vector. The *isotropic* Ridge-regularized Least Squares (RLS) estimator is therefore defined as

$$\hat{\boldsymbol{\beta}}^R(\lambda) := \arg \min_{\boldsymbol{\beta}} \frac{1}{T} \|\mathbf{y} - (Z' \otimes I_K)\boldsymbol{\beta}\|^2 + \lambda \|\boldsymbol{\beta}\|^2,$$

where $\lambda > 0$ is the scalar regularization parameter or regularizer. When $\lambda \|\boldsymbol{\beta}\|^2$ is replaced with quadratic form $\boldsymbol{\beta}'\Lambda\boldsymbol{\beta}$ for a positive definite matrix Λ the above is often termed Tikhonov regularization. To avoid confusion, I shall also refer to it as “ridge”, since in what follows Λ will always be assumed diagonal. Since Λ does not, in general, penalize coefficients equally, it will yield an *anisotropic* ridge estimator.

By solving the normal equations (see Appendix A.1), the RLS estimator with positive semi-definite regularization matrix $\Lambda \in \mathbb{R}^{K^2 p \times K^2 p}$ is shown to be

$$\hat{\boldsymbol{\beta}}^R(\Lambda) = \left(\frac{ZZ'}{T} \otimes I_K + \Lambda \right)^{-1} \frac{(Z \otimes I_K)\mathbf{y}}{T}.$$

When a centering vector $\boldsymbol{\beta}_0 \neq 0$ is included in penalty $(\boldsymbol{\beta} - \boldsymbol{\beta}_0)'\Lambda(\boldsymbol{\beta} - \boldsymbol{\beta}_0)$, the RLS estimator becomes

$$\hat{\boldsymbol{\beta}}^R(\Lambda, \boldsymbol{\beta}_0) = \left(\frac{ZZ'}{T} \otimes I_K + \Lambda \right)^{-1} \left(\frac{(Z \otimes I_K)\mathbf{y}}{T} + \Lambda\boldsymbol{\beta}_0 \right). \quad (2)$$

In the context of multivariate estimation, one has to make a further distinction between two related types of ridge estimators. I let $\hat{B}^R(\Lambda, \boldsymbol{\beta}_0)$ be the de-vectorized coefficient estimator obtained from reshaping $\hat{\boldsymbol{\beta}}^R(\Lambda, \boldsymbol{\beta}_0)$ to a $K \times Kp$ matrix. But one can also consider the *matrix RLS estimator* $\hat{B}_{\text{mat}}^R(\Lambda_{Kp}, B_0)$ given by

$$\hat{B}_{\text{mat}}^R(\Lambda_{Kp}, B_0) = T^{-1}(Y + B_0 \Lambda_{Kp})Z' \left(T^{-1}ZZ' + \Lambda_{Kp} \right)^{-1},$$

where $\Lambda_{Kp} = \text{diag}\{\lambda_1, \dots, \lambda_{Kp}\}$ and B_0 is a centering matrix. The distinction is important because the vectorized and matrix RLS estimators in general need not coincide. As discussed in Appendix A.2, $\hat{B}^R(\Lambda, \boldsymbol{\beta}_0)$ allows for more general penalty structures compared to $\hat{B}_{\text{mat}}^R(\Lambda_{Kp}, B_0)$. I, therefore, focus on the former rather than the latter.

3 Shrinkage

In this section, I discuss both the isotropic ridge penalty, i.e. the “standard” ridge approach, as well as an anisotropic penalty that is better adapted to the VAR setting. An important result is that, even in simple setups with only two variables, the shrinkage induced by ridge can either increase or reduce bias, as well as the stability of autoregressive estimates.

Throughout this section, I consider *fixed* design matrices and the focus will be on the geometric properties of ridge.

3.1 Isotropic Penalty

The most common way to perform a ridge regression is through isotropic regularization, that is, $\Lambda = \lambda I$ for some scalar $\lambda \geq 0$. Isotropic ridge has been extensively studied, see for example the comprehensive review of [Hastie \(2020\)](#). In regards to shrinkage, an isotropic ridge penalty can be readily studied.

Proposition 1. *Let $Z \in \mathbb{R}^{M \times T}$, $Y \in \mathbb{R}^T$ for $T > M$ be regression matrices. For $\lambda_\bullet > \lambda > 0$ and isotropic RLS estimator $\hat{\beta}^R(\lambda) := (T^{-1}ZZ' + \lambda I_M)^{-1}(T^{-1}ZY)$ it holds*

$$\|\hat{\beta}^R(\lambda_\bullet)\| < \|\hat{\beta}^R(\lambda)\|.$$

Proof. Using the full singular-value decomposition (SVD), decompose $T^{-1/2}Z = UDV' \in \mathbb{R}^{M \times T}$ where U is $M \times M$ orthogonal, D is $M \times T$ diagonal and V is $T \times T$ orthogonal. Write

$$\begin{aligned} \hat{\beta}^R(\lambda_\bullet) &= (T^{-1}ZZ' + \lambda_\bullet I_M)^{-1}(T^{-1}ZY) \\ &= (UDV'VDU' + \lambda_\bullet I_M)^{-1}UDV'(T^{-1/2}Y) \\ &= U(D^2 + \lambda_\bullet I_M)^{-1}DV'(T^{-1/2}Y) \\ &= U(D^2 + \lambda_\bullet I_M)^{-1}(D^2 + \lambda I_M)(D^2 + \lambda I_M)^{-1}DV'(T^{-1/2}Y) \\ &= [U(D^2 + \lambda_\bullet I_M)^{-1}(D^2 + \lambda I_M)U'] \hat{\beta}^R(\lambda). \end{aligned}$$

Since $D^2 = \text{diag}\{\sigma_j^2\}_{j=1}^M$, the term within brackets is $U \text{diag}\{(\sigma_j^2 + \lambda)/(\sigma_j^2 + \lambda_\bullet)\}_{j=1}^M U'$. Moreover, because the spectral norm is unitary invariant and $\lambda_1 > \lambda_2$, it follows that

$$\|U(D^2 + \lambda_\bullet I_M)^{-1}(D^2 + \lambda I_M)U'\| = \|\text{diag}\{(\sigma_j^2 + \lambda)/(\sigma_j^2 + \lambda_\bullet)\}_{j=1}^M\| < 1.$$

Finally, by the sub-multiplicative property it holds

$$\|\hat{\beta}^R(\lambda_\bullet)\| \leq \|U(D^2 + \lambda_1 I_M)^{-1}(D^2 + \lambda I_M)U'\| \cdot \|\hat{\beta}^R(\lambda)\| < \|\hat{\beta}^R(\lambda)\|$$

as claimed. \square

Proposition 1 and its proof expose the main ingredients of ridge regression. From the SVD of Z used above, it is clear that ridge regularization acts uniformly along the orthogonal directions that are the columns of V . The improvement in conditioning of the inverse comes from all diagonal factors $[(D^2 + \lambda_{\bullet} I_M)^{-1} D]_j = \sigma_j / (\sigma_j^2 + \lambda_{\bullet})$ being well-defined even when $\sigma_j = 0$ (as is the case in collinear systems).

However, directly applying isotropic ridge to vector autoregressive models is not necessarily the most effective estimation approach. Stable VAR models show decay in the absolute size of coefficients over lags. So it is reasonable to chose a more general ridge penalty that can accommodate lag decay.

3.2 Lag-Adapted Penalty

I now consider a different form for Λ that is of interest when applying ridge specifically to a VAR model. Define family $\mathcal{F}^{(p)}$ of *lag-adapted* ridge penalty matrices as

$$\mathcal{F}^{(p)} = \{\text{diag}\{\lambda_1, \dots, \lambda_p\} \otimes I_{K^2} \mid \lambda_i \in \mathbb{R}_+, i = 1, \dots, p\},$$

where each λ_i intuitively implies a different penalty for the elements of each coefficient matrix A_i , $i = 1, \dots, p$.¹ Family $\mathcal{F}^{(p)}$ allows imposing a ridge penalty that is coherent with the lag dimension of an autoregressive model. It is parametrized by p distinct penalty factors, meaning that the penalization is *anisotropic*.

Proposition 2. *Let $Z \in \mathbb{R}^{Kp \times T}$, $\mathbf{y} \in \mathbb{R}^{KT}$ for $T > Kp$ be multivariate VAR regression matrices. Given subset $\mathcal{S} \subseteq \{1, \dots, p\}$ of cardinality $s = |\mathcal{S}|$, for $\Lambda^{(p)} \in \mathcal{F}^{(p)}$ define $\hat{\boldsymbol{\beta}}^R(\Lambda^{(p)})_{[\mathcal{S}]}$ as the vector of sK^2 coefficient estimates located at indexes $1 + K^2(j-1), \dots, K^2j$ for $j \in \mathcal{S}$. Let $\mathcal{S}^c = \{1, \dots, p\} \setminus \mathcal{S}$ be the complement of \mathcal{S} .*

- (a) *If $\lambda_1 \geq \lambda_2$, then $\|\hat{\boldsymbol{\beta}}^R(\lambda_1 I_{K^2p})_{[\mathcal{U}]}\| \leq \|\hat{\boldsymbol{\beta}}^R(\lambda_2 I_{K^2p})_{[\mathcal{U}]}\|$ for any $\mathcal{U} \subset \{1, \dots, K^2p\}$. The inequality is strict when $\lambda_1 > \lambda_2$.*
- (b) *Let $\hat{\boldsymbol{\beta}}_{[\mathcal{S}]}^{LS}$ be the least squares estimator of the autoregressive model with only the lags indexed by \mathcal{S} included and zeros as coefficients for the lags indexed by \mathcal{S}^c . Similarly, let $\Lambda_{[\mathcal{S}]}^{(p)}$ be the subset of diagonal elements in $\Lambda^{(p)}$ penalizing the lags in \mathcal{S} . Then*

$$\lim_{\substack{\Lambda_{[\mathcal{S}]}^{(p)} \rightarrow 0 \\ \Lambda_{[\mathcal{S}^c]}^{(p)} \rightarrow \infty}} \hat{\boldsymbol{\beta}}^R(\Lambda^{(p)}) = \hat{\boldsymbol{\beta}}_{[\mathcal{S}]}^{LS},$$

¹Note that with a lag-adapted penalty it is also possible to directly use the matrix ridge estimator since the penalty for $\hat{\boldsymbol{\beta}}^R$ is given by $\text{diag}\{\lambda_1, \dots, \lambda_p\} \otimes I_{K^2} = (\text{diag}\{\lambda_1, \dots, \lambda_p\} \otimes I_K) \otimes I_K$, see Appendix A.2.

where $\Lambda_{[\mathcal{S}]}^{(p)} \rightarrow 0$ and $\Lambda_{[\mathcal{S}^c]}^{(p)} \rightarrow \infty$ are to be intended as the element-wise convergence.

Proposition 2 shows that the limiting geometry of a lag-adapted ridge estimator is thus identical to that of a least squares regression run on the subset specified by \mathcal{S} . By controlling the size of coefficients $\{\lambda_1, \dots, \lambda_p\}$ it is therefore possible to obtain pseudo-model-selection. However, in the next section I show that anisotropic penalization produces complex effects on the models' coefficient estimates.

3.3 Effects of Anisotropic Penalization

In this section, I explore the effect of a lag-adapted ridge penalty on the estimate VAR coefficients, and, more generally, the properties of ridge estimators with anisotropic penalization.

Since ridge operates along the principal components, there is no immediate relationship between a specific subset of estimated coefficients and a given diagonal block of $\Lambda^{(p)}$. For autoregressive modeling, three effects are of interest: the shrinkage of coefficient matrices A_i relative to the choice of λ_i ; the entity of the bias introduced by shrinkage, and the impact of shrinkage on the persistence of the estimated model.

To evaluate these effects, I consider a simple VAR(2) model

$$y_t = A_1 y_{t-1} + A_2 y_{t-2} + u_t, \quad u_t \sim \text{i.i.d. } \mathcal{N}(0, \Sigma_u),$$

where

$$A_1 = \begin{bmatrix} 0.8 & 0.1 \\ -0.1 & 0.7 \end{bmatrix}, \quad A_2 = \begin{bmatrix} 0.1 & -0.2 \\ -0.1 & 0.1 \end{bmatrix}, \quad \Sigma_u = \begin{bmatrix} 0.3 & 0 \\ 0 & 5 \end{bmatrix}.$$

A sample of length $T = 200$ is drawn, demeaned and used to estimate coefficients A_1 and A_2 . The VAR(2) model is fitted using the lag-adapted ridge estimator $\hat{B}^R(\Lambda^{(2)})$, where $\Lambda^{(2)} = \text{diag}\{\lambda_1, \lambda_2\} \otimes I_2$, which can be partitioned into estimates $\hat{A}_1^R(\Lambda^{(2)})$ and $\hat{A}_2^R(\Lambda^{(2)})$.

Shrinkage. To study shrinkage, I consider the restricted case of $\lambda_1 \in [10^{-2}, 10^6]$ and $\lambda_2 = 0$. The ridge estimator is computed for varying λ_1 over a logarithmically spaced grid. Figure 1a shows that $\|\hat{B}^R(\Lambda^{(2)})\|_F \approx \|\hat{B}^{LS}\|_F$ for $\lambda_1 \approx 0$, but as the penalty increases $\|\hat{A}_1^R(\Lambda^{(2)})\|_F$ decreases while $\|\hat{A}_2^R(\Lambda^{(2)})\|_F$ grows. The resulting behavior of $\|\hat{B}^R(\Lambda^{(2)})\|_F$ is non-monotonic in λ_1 , although indeed $\|\hat{B}^R(\Lambda^{(2)})\|_F < \|\hat{B}^{LS}\|_F$ in the limit $\lambda_1 \rightarrow \infty$. This effect is due to the model selection properties of lag-adapted ridge, and the resulting omitted variable bias. Therefore, in practice it is not generally true that anisotropic ridge induces monotonic shrinkage of estimates.




Figure 1: Shrinkage of coefficients estimate in Frobenius norm (a); bias induced by shrinkage (b); change in stability of estimated VAR model at different levels of penalization, measured by the absolute value of the largest companion form eigenvalue (c).

Bias. Since ridge bias is hard to study theoretically, I use a simulation with the same setup of Figure 1a, this time with $\lambda_1, \lambda_2 \in [10^{-2}, 10^4]$. The grid is logarithmic with 150 points. Figure 1b presents a level plot of the sup-norm ridge bias $\|\hat{B}^R(\Lambda^{(2)}) - B\|_\infty$ given multiple combinations of λ_1 and λ_2 . While there can be gains compared to the least squares estimator \hat{B}^{LS} , they are modest. Moreover, level curves of the bias surface show that gains concentrate in a very thin region of the parameter space. Consequently, in practice any (data-driven) ridge penalty selection criterion is unlikely to yield bias improvement over least squares. Yet, in large VAR models with many lags the reduction in variance of the ridge estimator often yields improvements over unregularized procedures (Li et al., 2022). However, the

bias-variance trade-off in ridge is not a free-lunch when performing inference. [Pratt \(1961\)](#) showed that it is not possible to produce a test (equivalently, a CI procedure) which is valid uniformly over the parameter space and yields meaningfully smaller confidence intervals than any other valid method.

Stability. To study the stability of ridge VAR estimates, I reuse the results of the bias simulation above. Let \mathbb{A} be the companion matrix of $[A_1, A_2]$, and $\hat{\mathbb{A}}^R$ the companion matrix of estimates $[\hat{A}_1^R(\Lambda^{(2)}), \hat{A}_2^R(\Lambda^{(2)})]$. For all combinations (λ_1, λ_2) , I compute the largest eigenvalue $\omega_1(\hat{\mathbb{A}}^R)$ of $\hat{\mathbb{A}}^R$. Note that if $|\omega_1(\hat{\mathbb{A}})| < 1$, then the estimated VAR(2) is stable ([Lütkepohl, 2005](#)). [Figure 1c](#) presents the level sets for the surface of maximal eigenvalue moduli, and for comparison $|\omega_1(\hat{B}^{LS})|$ is shown at the origin.² While along the main diagonal there is a clear decrease in $|\omega_1(\hat{\mathbb{A}}^R)|$ as isotropic penalization increases, when λ_1 is large and $\lambda_2 \ll 1$ (or vice versa) the maximal eigenvalue increases instead. Therefore, an estimate of a VAR model obtained with anisotropic ridge may be *closer* to unit root than the least squares estimate.

4 Bayesian and Frequentist Ridge

So far, I have discussed standard ridge penalization schemes. In this section, I study the posterior mean of Bayesian VAR (BVAR) priors commonly applied in the macroeconometrics literature. I show that such posteriors are in fact specific GLS formulations of the ridge estimator. This comparison highlights that ridge can be seen as a way to embed prior knowledge into the least squares estimation procedure by means of centering and rescaling coefficient estimates.

4.1 Litterman-Minnesota Priors

In Bayesian time series modeling, the so-called Minnesota or Litterman prior has found great success ([Litterman, 1986](#)). For stationary processes which one believes to have reasonably small dependence, a zero-mean normal prior can be put on the VAR parameters, with non-zero prior variance. Assuming that the covariance matrix of errors Σ_u is known, the Litterman-Minnesota has posterior mean

$$\bar{\beta} \mid \Sigma_u = \left[V_{\beta}^{-1} + (ZZ' \otimes \Sigma_u^{-1}) \right]^{-1} (Z \otimes \Sigma_u^{-1}) \mathbf{y}, \quad (3)$$

²If $\Lambda^{(2)} \rightarrow 0$, then by continuity of eigenvalues it follows that $|\omega_1(\hat{\mathbb{A}}^R)| \rightarrow |\omega_1(\hat{\mathbb{A}}^{LS})|$, see [Appendix A.3](#).

where $\underline{V}_\beta \succ 0$ is the prior covariance matrix of β (Lütkepohl, 2005). It is common to let \underline{V}_β be diagonal, and often the entries follow a simple pattern which depends on lag, individual components variances, and prior hyperparameters. For example, Bańbura et al. (2010) suggest the following structure for the diagonal

$$v_{i,jk} = \begin{cases} \frac{\lambda^2}{i^2} & \text{if } j = k, \\ \theta \frac{\lambda^2}{i^2} \frac{\sigma_j^2}{\sigma_k^2} & \text{if } j \neq k, \end{cases} \quad (4)$$

where $v_{i,jk}$ is the prior variance for coefficients $(A_i)_{jk}$ for $i = 1, \dots, p$ and $j, k = 1, \dots, K$. Here, σ_j is the j -th diagonal element of Σ_u , $\theta \in (0, 1)$ specifies beliefs on the explanatory importance of own lags relative to other variables' lags, while $\lambda \in [0, \infty]$ controls the overall tightness of the prior. The extreme $\lambda = 0$ yields a degenerate prior centered at $\bar{\beta} = 0$, while $\lambda = \infty$ reduces the posterior mean to OLS estimate $\hat{\beta}^{LS}$. Factor $1/i^2$, which explicitly shrinks variance at higher lags, was originally introduced by De Mol et al. (2008), who formally developed the idea that coefficients at deeper lags should be coupled with more penalizing priors. Note that, in (4), assuming $\sigma_j^2 = \sigma_k^2$ for all $j, k = 1, \dots, K$ and setting $\theta = 1$, produces a \underline{V}_β that has a lag-adapted structure with quadratic lag decay.

Equation (3) more generally demonstrates that the Minnesota posterior mean is equivalent to a ridge procedure. It is important to notice that, while with least squares the OLS and GLS estimators of VAR coefficients coincide, this is not the case with ridge regression. Regularizing a GLS regression will yield

$$\hat{\beta}^{RGLS}(\Lambda) := [\Lambda + (ZZ' \otimes \Sigma_u^{-1})]^{-1} (Z \otimes \Sigma_u^{-1})\mathbf{y}. \quad (5)$$

instead of $\hat{\beta}^R$, which is equivalent to (3) under an appropriate choice of Λ . While I develop the asymptotic results for $\hat{\beta}^R$ assuming a centering parameter $\underline{\beta}_0 \neq 0$ in general, I do not directly study the properties $\hat{\beta}^{RGLS}$. The generalization to GLS ridge employing the least squares error covariance estimator $\hat{\Sigma}_T^{LS}$ should follow from straightforward arguments. In Section 7, I focus on providing evidence on the application $\hat{\beta}^{RGLS}$ in terms of its pointwise impulse response estimation mean-squared error.

4.2 Hierarchical Priors

Recent research on Bayesian vector autoregressions exploit more sophisticated priors compared to the Litterman-Minnesota design. Giannone et al. (2015) develop an advanced BVAR model by setting up hierarchical priors which entail not only model parameters, but

also hyperparameters. They impose

$$\begin{aligned}\Sigma_u &\sim \text{IW}(\underline{\Psi}, \underline{d}), \\ \beta \mid \Sigma_u &\sim \mathcal{N}(\underline{\beta}, \lambda(\Sigma_u \otimes \underline{\Omega})),\end{aligned}$$

for hyperparameters $\underline{\beta}$, $\underline{\Omega}$, $\underline{\Psi}$ and \underline{d} , where IW is the Inverse-Wishart distribution. Here, too, scalar $\lambda \in [0, \infty]$ controls prior tightness. Let \underline{B} be the matrix form of the VAR coefficient prior mean, so that $\text{vec}(\underline{B}) = \underline{\beta}$. The resulting (conditional) posterior mean \overline{B} is given by

$$\overline{B} \mid \Sigma_u = [(\lambda \underline{\Omega})^{-1} + ZZ']^{-1} [ZY + (\lambda \underline{\Omega})^{-1} \underline{B}]. \quad (6)$$

Observe that equation (6) is effectively equivalent to a centered ridge estimator, c.f. (2).

The introduction of a hierarchical prior leaves space to add informative hyperpriors on the model hyperparameters, allowing for a more flexible fit. Indeed, removing the zero centering constraint from the prior on β can improve estimation. It is often the case that economic time series show a high degree of correlation and temporal dependence, therefore imposing $\beta = 0$ as in the Minnesota prior is inadequate. In fact, [Giannone et al. \(2015\)](#) show that their approach yields substantial improvements in forecasting exercises, even when hyperparameter priors are relatively flat and uninformative.

5 Standard Inference

In this section, I state the main asymptotic results for the RLS estimator $\hat{\beta}^R(\Lambda, \beta_0)$ with general regularization matrix Λ . I shall allow Λ and non-zero centering coefficient β_0 to be, under appropriate assumptions, random variables dependent on sample size T . In particular, β_0 may be a consistent estimator of β .

I will impose the following assumptions.

Assumptions

- A. $\{u_t\}_{t=1}^T$ is a sequence of i.i.d. random variables with $\mathbb{E}[u_{it}] = 0$, covariance $\mathbb{E}[u_t u_t'] = \Sigma_u$ non-singular positive definite and $\mathbb{E}|u_{it} u_{jt} u_{mt} u_{nt}| < \infty$, $i, j, m, n = 1, \dots, K$.
- B. There exists $\rho > 1$ such that $\det(I_K - \sum_{i=1}^p A_i z^i) \neq 0$ for all complex z , $|z| \leq \rho$.
- C. There exist $0 < \underline{m} \leq \overline{m} < \infty$ such that $\underline{m} \leq \omega_K(\Gamma) \leq \omega_1(\Gamma) \leq \overline{m}$, where $\Gamma = \mathbb{E}[z_t z_t']$ is the autocovariance matrix of z_t and $\omega_1(\Gamma)$, $\omega_K(\Gamma)$ are its largest and smallest eigenvalues, respectively.

Assumption A is standard and allows to prove the main asymptotic results with well-known theoretical devices. Assuming u_t is white noise or assuming y_t respects strong mixing

conditions (Davidson, 1994) would require more careful consideration in asymptotic arguments but is otherwise a simple generalization, although more involved in terms of notation, see e.g. Boubacar Mainassara and Francq (2011).

Assumption B guarantees that y_t has no unit roots and is stable. Of course, many setups of interest do not satisfy this assumption, the most significant ones being unit roots, cointegrated VARs, and local-to-unity settings. Incorrect identification of unit roots does not invalidate the use of LS or ML estimators (Phillips, 1988, Park and Phillips, 1988, 1989, Sims et al., 1990), however inference is significantly impacted as a result (Pesavento and Rossi, 2006, Mikusheva, 2007, 2012).

Assumption C is standard in the literature regarding penalized estimation and does not imply significant additional constraints on the process y_t , cf. Assumption A. It is sufficient to ensure that for large enough T the plug-in sample autocovariance estimator is invertible even under vanishing Λ .

Before stating the main theorems, let

$$\begin{aligned}\hat{\Gamma} &= T^{-1}ZZ', \\ \hat{U} &= Y - \hat{B}^R Z, \\ \hat{\Sigma}_u^R &= T^{-1}\hat{U}\hat{U}',\end{aligned}$$

be the regression covariance matrix, regression residuals and sample innovation covariance estimator, respectively.

Theorem 1. *Let Assumptions A-C hold and define $\hat{\beta}^R(\Lambda, \beta_0)$ be the centered RLS estimator as in (2). If $\sqrt{T}\Lambda \xrightarrow{P} \Lambda_0$ and $\beta_0 \xrightarrow{P} \underline{\beta}_0$, where Λ_0 is positive semi-definite diagonal matrix and $\underline{\beta}_0$ is a constant vector, then*

- (a) $\hat{\Gamma} \xrightarrow{P} \Gamma$,
- (b) $\hat{\beta}^R(\Lambda, \beta_0) \xrightarrow{P} \beta$,
- (c) $\hat{\Sigma}_u^R \xrightarrow{P} \Sigma_u$,
- (d) $\sqrt{T} \left(\hat{\beta}^R(\Lambda, \beta_0) - \beta \right) \xrightarrow{d} \mathcal{N} \left(\Gamma^{-1} \Lambda_0 (\underline{\beta}_0 - \beta), \Gamma^{-1} \otimes \Sigma_u \right)$.

Theorem 1 considers the most general case, and, as previously mentioned, gives the asymptotic distribution of $\hat{\beta}^R$ under rather weak conditions for the regularizer Λ . The resulting normal limit distribution is clearly dependent on the unknown model parameters β , complicating inference.

However, it is possible – under strengthened assumptions for Λ or β_0 – for $\hat{\beta}^R$ to have a zero-mean Gaussian limit distribution.

Theorem 2. *In the setting of Theorem 1, results (a)-(c) hold and (d) simplifies to*

$$(d') \quad \sqrt{T} \left(\hat{\beta}^R(\Lambda, \beta_0) - \beta \right) \xrightarrow{d} \mathcal{N} \left(0, \Gamma^{-1} \otimes \Sigma_u \right)$$

if either

$$(i) \quad \Lambda = o_P \left(T^{-1/2} \right),$$

$$(ii) \quad \Lambda = O_P \left(T^{-1/2} \right) \text{ and } \beta_0 - \beta = o_p(1).$$

The following corollary is immediate.

Corollary 1. *Let $\hat{\beta}_0$ be a consistent and asymptotically normal estimator of β . Then, under condition (i) or (ii) of Theorem 2 results (a)-(d') hold.*

5.1 Joint Inference

To handle smooth transformations of VAR coefficients, such as impulse responses (Lütkepohl, 1990), I also derive a standard joint limit result for both $\hat{\beta}^R$ and the variance estimator $\hat{\Sigma}_u^R$.

Theorem 3. *Let $\hat{\sigma}^R = \text{vec}(\hat{\Sigma}_u^R)$ and $\sigma = \text{vec}(\Sigma_u)$. Under the assumptions of Theorem 1,*

$$\sqrt{T} \begin{bmatrix} \hat{\beta}^R - \beta \\ \hat{\sigma}^R - \sigma \end{bmatrix} \xrightarrow{d} \mathcal{N} \left(\begin{bmatrix} \Gamma^{-1} \Lambda_0(\beta_0 - \beta) \\ 0 \end{bmatrix}, \begin{bmatrix} \Gamma^{-1} \otimes \Sigma_u & 0 \\ 0 & \Omega \end{bmatrix} \right).$$

Under assumption (1) or (2) of Theorem 2,

$$\sqrt{T} \begin{bmatrix} \hat{\beta}^R - \beta \\ \hat{\sigma}^R - \sigma \end{bmatrix} \xrightarrow{d} \mathcal{N} \left(0, \begin{bmatrix} \Gamma^{-1} \otimes \Sigma_u & 0 \\ 0 & \Omega \end{bmatrix} \right),$$

where $\Omega = \mathbb{E}[\text{vec}(u_t u_t') \text{vec}(u_t u_t')'] - \sigma \sigma'$.

This result is key as it allows, under the stated assumptions on the penalizer, to construct valid asymptotic confidence intervals and, specifically, perform impulse response inference, as done in the simulations of Section 7 using the Delta Method, c.f. Lütkepohl (2005).

5.2 Cross-validation

In practice, the choice of ridge penalty is often data-driven, and cross-validation is an very popular approach to select Λ . I now turn to the properties of CV as applied to $\hat{\beta}^R(\Lambda)$.

For simplicity, assume that y_t is an AR(p) process, that is, $K = 1$. In this setting,

$$\hat{\beta}^R(\Lambda) = \left(\frac{ZZ'}{T} + \Lambda \right)^{-1} \frac{Zy}{T},$$

where $\Lambda = \text{diag}\{\lambda_1, \dots, \lambda_p\}$. Following [Patil et al. \(2021\)](#), the *prediction error* of ridge estimator $\hat{\beta}^R(\Lambda)$ given penalty Λ is

$$\text{Err}(\hat{\beta}^R(\Lambda)) := \mathbb{E}_{\tilde{y}, \tilde{z}} \left[\left(\tilde{y} - \tilde{z}' \hat{\beta}^R(\Lambda) \right)^2 \mid Z, \mathbf{y} \right],$$

where \tilde{y} and \tilde{z} are random variables from an independent copy of y_t . In particular, \tilde{z} is the vector of p lags of \tilde{y} . Moreover, the error curve for Λ is given by

$$\text{err}(\Lambda) := \text{Err}(\hat{\beta}^R(\Lambda)).$$

The prediction error is crucial because it allows to determine the oracle optimal penalization,

$$\Lambda^* := \arg \min_{\Lambda \succeq 0} \text{err}(\Lambda).$$

Clearly, $\text{err}(\Lambda)$ is unavailable in practice and Λ^* must be substituted with a feasible alternative. Cross-validation proposes to construct a collection of paired, non-overlapping subsets of the sample data such that the first subset of the pair (estimation set) is used to estimate the model, while the second (validation set) is used to provide an empirical estimate of the prediction error. The CV penalty is then selected to minimize the total error over validation sets. A very popular approach to build cross-validation subsets is k -fold CV, wherein the sample is split into k blocks, so-called *folds*, of sequential observations (possibly after shuffling the data). Each fold determines a validation set, and is paired with its complement, which gives the estimation set. For more details, see e.g. [Hastie et al. \(2009\)](#).

Again with the intent of keeping complexity low – as this work is not focused on cross-validation – I will make the additional simplifying assumption that CV is implemented with two folds and one pair. Specifically, the first fold is the estimation set, where Z and \mathbf{y} are constructed and $\hat{\beta}^R(\Lambda)$ is estimated. The second fold is the validation set and yields \tilde{Z} , $\tilde{\mathbf{y}}$, where $\tilde{Z} \in \mathbb{R}^{p \times \tilde{T}}$ and $\tilde{\mathbf{y}} \in \mathbb{R}^{\tilde{T}}$. To account for dependence, a buffer of m observations between validation and estimation folds is introduced. The last observation of y_t in the estimation set is y_T , while the first observation in validation set is $\tilde{y}_1 := y_{T+m+1}$, that is, the total number of available observations is $T + m + \tilde{T} + 2p + 1$. This is a stylized version of the CV setup of [Burman et al. \(1994\)](#) – also called *m-block* or *non-dependent cross-validation* in [Bergmeir et al. \(2018\)](#) – and is effectively equivalent to an out-of-sample (OOS) validation scheme. Thus, the 2-fold m -buffered CV error curve is

$$\text{cv}2_m(\Lambda) := \frac{1}{\tilde{T}} \sum_{s=1}^{\tilde{T}} \left(\tilde{y}_s - \tilde{z}'_s \hat{\beta}^R(\Lambda) \right)^2. \quad (7)$$

Theorem 4. *Under Assumptions A-C, for every Λ in the cone of diagonal positive definite*

penalty matrices with diagonal entries in (λ_{\min}, ∞) , $\lambda_{\min} \geq 0$, it holds that

$$\text{cv}2_m(\Lambda) - \text{err}(\Lambda) \xrightarrow{a.s.} 0$$

as $T, \tilde{T} \rightarrow \infty$. Furthermore, the convergence is uniform in Λ over compact subsets of penalty matrices.

In the current setup, the joint limit $T, \tilde{T} \rightarrow \infty$ should be thought as $\tilde{T}/T \rightarrow \gamma \in (0, 1)$, where aspect ratio γ determines the balance of the cross-validation split.

Remark 1. Under Assumption C, $\omega_K(\hat{\Gamma}) > 0$ for T large. Therefore, the bounds derived in the proof of Theorem 4 are finite even if $\Lambda = 0$. In fact, it is easily seen that the behavior of $\text{err}(\Lambda)$ and $\text{cv}2_m(\Lambda)$ is consistent at the endpoints $\Lambda = 0$ and $\Lambda \rightarrow \infty$, c.f. Patil et al. (2021). Observe that

$$\text{cv}2_m(\Lambda) \rightarrow \Sigma_u \quad \text{and} \quad \text{err}(\Lambda) \rightarrow \Sigma_u$$

as $\Lambda \rightarrow 0$, while

$$\text{cv}2_m(\Lambda) \rightarrow \Gamma \quad \text{and} \quad \text{err}(\Lambda) \rightarrow \Gamma$$

as $\Lambda \rightarrow \infty$, as needed.

Theorem 4 thus shows that $\text{cv}2_m(\Lambda)$ gives an asymptotically valid way to evaluate the prediction error curve, and thus tune Λ , over any compact set of diagonal positive semi-definite penalization matrices. Moreover, in Theorem 7, Appendix C.2, I show that the impact of dependence due to the VAR data generating process is exponentially small for m sufficiently large. This property of $\text{cv}2_m(\Lambda)$ is desirable because it lets one choose m small also in applications with moderate sample sizes, and theoretical justifies the prescription of Burman et al. (1994).

5.3 Asymptotically Valid CV

So far, I have shown that a simple 2-fold CV – or, equivalently, out-of-sample validation – correctly estimates the predictive error of the ridge estimator, even under dependence. I turn now to the question of selecting an *asymptotically valid* penalty, that is, a Λ such that condition (1) of Theorem 2 is fulfilled. This enables inference, since one is in a setting where the bias is asymptotically negligible.

The idea is to scale the ridge penalty used at the estimation step of CV by a factor $\sqrt{\tilde{T}}$, so that the validated penalty converges to zero at an appropriate rate as both T and \tilde{T} grow. In other words, an over-smoothed ridge regression turns out to be key when studying

cross-validation. To derive this result, first let

$$\hat{\beta}_{\blacklozenge}^R(\Lambda) := \left(\frac{ZZ'}{T} + \sqrt{T}\Lambda \right)^{-1} \frac{Z\mathbf{y}}{T}$$

be the *over-smoothed ridge estimator*.

Theorem 5. *Under Assumptions A-C, let \mathcal{I}_λ be the compact set of diagonal positive semidefinite penalization matrices Λ such that $\|\Lambda\|_{\max} \leq \lambda < \infty$. It holds*

$$\Lambda_{\blacklozenge}^* := \arg \min_{\Lambda \in \mathcal{I}_\lambda} \text{Err} \left(\hat{\beta}_{\blacklozenge}^R(\Lambda) \right) = o_p(T^{-1/2}).$$

Remark 2. The previous theorem is stated in terms of the oracle predictive error $\text{Err} \left(\hat{\beta}_{\blacklozenge}^R(\tilde{\Lambda}) \right)$, which equals the 2-fold CV error curve up to a factor of order $O_P(\tilde{T}^{-1/2})$. Therefore, assuming that the CV aspect ratio γ is strictly between zero and one, the result of Theorem 5 also directly generalizes to an empirically cross-validated penalty.

6 Inference with Shrinkage

Fu and Knight (2000) have argued that results such as Theorems 1 and 2 portray penalized estimators in a somewhat unfair light, because they result in asymptotic distributions showing no bias-variance trade-off. Indeed, they show that ridge shrinkage yields estimates with asymptotic variance no different than that of least squares. Of course, in finite samples shrinkage has an effect on $\Gamma^{-1} \otimes \Sigma_u$ since $\hat{\Sigma}_T^R$ is used in place of $\hat{\Sigma}_T^{LS}$ to estimate the error term variance matrix. To better understand the value of ridge penalization in practice, one should therefore consider the situation where a number of VAR coefficients are small, but not necessarily zero.

Formally, assume that for some $0 < n \leq p$ one can partition the VAR coefficients as $\beta = (\beta_1', \beta_2')'$, where $\beta_1 \in \mathbb{R}^{K^2(p-n)}$ and $\beta_2 \in \mathbb{R}^{K^2n}$, and assume that $\beta_2 = T^{-(1/2+\delta)} \mathbf{b}_2$ for $\delta > 0$ and $\mathbf{b}_2 \in \mathbb{R}^{K^2n}$ is fixed. Such ordered partitioning of β is without loss of generality.³ In this context, it is desirable to penalize β_1 and β_2 differently when constructing the ridge penalty. From a practical perspective, one can consider, for example, the case of a VAR(∞) model derived by inverting a stable VARMA(p, q) process: for i sufficiently large, coefficient matrices A_i decay exponentially to zero.⁴ In a finite sample, an asymptotic framework with non-negligible penalization of higher-order lag coefficients can be more appropriate than

³The dimensions of β_1 and β_2 are chosen to be multiples of K^2 to better conform to the lag-adapted setting. This choice is also without loss of generality and simplifies exposition.

⁴This result follows from a straightforward generalization of Lemma 1 in Appendix C. The choice of norm to measure such decay is not fundamental as they are equivalent given that dimension K is fixed.

that of Theorem 1. This approach to inference is also in the vein of De Mol et al. (2008), who argue for explicit lag penalization into BVAR priors on similar grounds. Thus, the idea of partitioned penalization exploits prior information on the structure of autoregressive coefficients to asymptotically improve on the bias-variance trade-off. In the context of maximum-likelihood estimation, the use of appropriate and plausible model restrictions to improve efficiency by shrinkage, rather than perform hypothesis testing, has also been discussed by Hansen (2016a).

Let $\Lambda = \text{diag}\{(L'_1, L'_2)'\} \otimes I_K$ where $L_1 \in \mathbb{R}_+^{K^2(p-n)}$ and $L_2 \in \mathbb{R}_+^{K^2n}$. Assume that

$$L_1 = o_P(T^{-1/2}) \quad \text{and} \quad L_2 \xrightarrow{P} \bar{L}_2$$

for a fixed vector $\bar{L}_2 \in \mathbb{R}_+^{K^2n}$. In particular, letting $\Lambda_1 = \text{diag}\{L_1\}$ and $\Lambda_2 = \text{diag}\{L_2\}$,

$$\Lambda = \begin{bmatrix} \Lambda_1 & 0 \\ 0 & \Lambda_2 \end{bmatrix} \otimes I_K \xrightarrow{P} \bar{\Lambda} \otimes I_K \quad \text{where} \quad \bar{\Lambda} = \begin{bmatrix} 0 & 0 \\ 0 & \bar{\Lambda}_2 \end{bmatrix}, \quad \bar{\Lambda}_2 \succ 0. \quad (8)$$

One can now develop an asymptotic result which shows non-negligible shrinkage in the limit distribution of the ridge estimator. For simplicity of exposition, here I will assume that ridge centering β_0 is chosen to be zero.

Theorem 6. *In the setting of Theorem 1, assume that, for $0 < n \leq p$,*

- (i) $\beta = (\beta'_1, \beta'_2)'$ where $\beta_1 \in \mathbb{R}^{K^2(p-n)}$ and $\beta_2 = T^{-(1/2+\delta)} \mathbf{b}_2$ for $\delta > 0$, $\mathbf{b}_2 \in \mathbb{R}^{K^2n}$ fixed.
- (ii) $\Lambda = \text{diag}\{(L'_1, L'_2)'\}$ where $L_1 \in \mathbb{R}_+^{K^2(p-n)}$ and $L_2 \in \mathbb{R}_+^{K^2n}$.
- (iii) $L_1 = o_P(T^{-1/2})$ and $L_2 \xrightarrow{P} \bar{L}_2$ as $T \rightarrow \infty$.
- (iv) $\beta_0 = 0$.

Let $\Gamma_{\bar{\Lambda}} = \Gamma + \bar{\Lambda}$ where $\bar{\Lambda} \succeq 0$ is given by (8). Then, results (a)-(c) hold and

$$(d'') \quad \sqrt{T} \left(\hat{\beta}^R(\Lambda, \beta_0) - \beta \right) \xrightarrow{d} \mathcal{N} \left(0, \Gamma_{\bar{\Lambda}}^{-1} \Gamma \Gamma_{\bar{\Lambda}}^{-1} \otimes \Sigma_u \right)$$

It is easy to see that indeed the term $\Gamma_{\bar{\Lambda}}^{-1} \Gamma \Gamma_{\bar{\Lambda}}^{-1}$ in Theorem 6 is weakly smaller than Γ^{-1} in the positive-definite sense. Note that

$$\begin{aligned} \Gamma_{\bar{\Lambda}}^{-1} \Gamma \Gamma_{\bar{\Lambda}}^{-1} \preceq \Gamma^{-1} &\iff (\Gamma + \bar{\Lambda})^{-1} \Gamma \preceq \Gamma^{-1} (\Gamma + \bar{\Lambda}) \\ &\iff I_{K^2p} - (\Gamma + \bar{\Lambda})^{-1} \bar{\Lambda} \preceq I_{K^2p} + \Gamma^{-1} \bar{\Lambda} \\ &\iff 0 \preceq ((\Gamma + \bar{\Lambda})^{-1} + \Gamma^{-1}) \bar{\Lambda} \end{aligned}$$

The last inequality is true by definition of $\bar{\Lambda}$. Shrinkage gains are concentrated at the components that have non-zero asymptotic shrinkage, i.e. those penalized by \bar{L}_2 .

Remark 3. A key point in the application of Theorem 6 is identification of β_1 and β_2 . In practice, one may then proceed in two ways. As discussed in Section 4, one can see the ridge approach as a frequentist “counterpart” to implementing a Bayesian prior. Therefore, the researcher may split β into subsets of small and large parameters based on economic intuition, domain knowledge or preliminary information. Alternatively, in the following section I show that cross-validation is automatically able to tune Λ appropriately.

6.1 Cross-validation with Partitioned Coefficients

One can use the same approach applied to derive Theorem 5 in order to show that cross-validating the RLS estimator with $\text{Err}(\hat{\beta}_\diamond^R(\Lambda))$ is also asymptotically valid.

Corollary 2. *Consider the setup of Theorem 6 and assume that the assumptions of Theorem 5 are met. It holds*

$$\begin{bmatrix} \Lambda_{1,\diamond} & 0 \\ 0 & \Lambda_{2,\diamond} \end{bmatrix} := \arg \min_{\Lambda \in \mathcal{I}_\lambda} \text{Err}(\hat{\beta}_\diamond^R(\Lambda)) = \begin{bmatrix} o_p(T^{-1/2}) & 0 \\ 0 & o_P(1) \end{bmatrix}$$

Moreover, any $\Lambda_{2,\diamond}$ such that $0 \preceq \Lambda_{2,\diamond} \preceq \lambda I$ is asymptotically valid.

In theory, one would like to be able to quantify the gains obtained in the asymptotic shrinkage setup of Theorem 6 compared to the standard setting of Theorems 1 and 2, particularly when using cross-validation. Unfortunately, it is in general hard to study the cross-validation error loss even in setups without dependence. Stephenson et al. (2021) in fact show that the ridge leave-one-out CV loss is not generally convex. This suggests that studying the behavior of CV when penalizing with a diagonal anisotropic Λ can be a very complex task in a finite sample setup.

7 Simulations

To study the performance of ridge-regularized estimators, I now perform simulation exercises focused on impulse response functions (IRFs). Throughout the experiments I will consider structural impulse responses, and I assume that identification can be obtained in a recursive way (Kilian and Lütkepohl, 2017), which is a widely used approach for structural shock identification in macroeconometrics.

I consider two setups:

1. The three-variable VARMA(1,1) design of Kilian and Kim (2011), representing a small-scale macro model. I term this setup “A”.

Table 1: List of Estimation Methods

Type	Name	Description
Frequentist	LS	Least squares estimator
	RIDGE	Ridge estimator, CV penalty
	RIDGE-GLS	GLS ridge estimator, CV penalty
	RIDGE-AS	Ridge estimator with asymptotic shrinkage, CV penalty
	LP	Local projections with Newey-West covariance estimate
Bayesian	BVAR-CV	Litterman-Minnesota Bayesian VAR, CV tightness prior
	H-BVAR	Hierarchical Bayesian VAR of Giannone et al. (2015)

2. A VAR(5) model in levels, using the model specification of [Giannone et al. \(2015\)](#) with the dataset of [Hansen \(2016b\)](#) consisting of $K = 7$ variables in levels.⁵ I term this setup “B”. For the ease of exposition, in the discussion I will tabulate results only for three variables – real GDP, investment and federal funds rate – but complete tables can be found in Appendix D.5.

The specification of [Kilian and Kim \(2011\)](#) has already been extensively used in the literature as a benchmark to gauge the basic properties of inference methods. On the other hand, the estimation task of [Giannone et al. \(2015\)](#) involves more variables and a higher degree of persistence. This setting is useful to evaluate the effects of ridge shrinkage when applied to realistic macroeconomic questions. It is also a suitable test bench to compare Bayesian methods with frequentist ridge.

Estimators. For frequentist methods, I include both $\hat{\beta}^R$ and $\hat{\beta}^{RGLS}$ ridge estimators as well as the local projection estimator of [Jordà \(2005\)](#). For Bayesian methods, I implement both the Minnesota prior approach of [Bańbura et al. \(2010\)](#) with stationary prior and the hierarchical prior BVAR of [Giannone et al. \(2015\)](#).⁶ The full list of method I consider is given in Table 1. To make methods comparable, I have extended the ridge estimators to include an intercept in the regression. A precise discussion regarding the tuning of penalties and hyperparameters of all methods can be found in Supplementary Appendix D.

⁵The dataset is supplied by the author at <https://users.ssc.wisc.edu/~bhansen/progs/var.html>. While the data provided by [Hansen \(2016b\)](#) includes releases until 2016. I do not include more recent quarterly data since this is a simulation exercise. Moreover, due to the effects of the COVID-19 global pandemic, an extended sample would likely only add data released until Q4 2019 due to overwhelming concerns of a break point.

⁶To estimate hierarchical prior BVARs I rely on the original MATLAB implementation provided by [Giannone et al. \(2015\)](#) on the authors’ website at <http://faculty.wcas.northwestern.edu/gep575/GLPreplicationWeb.zip>.

7.1 Pointwise MSE

The first two simulation designs explore the MSE performance of ridge-type estimators versus alternatives. Let $\theta_{km}(h)$ be the horizon h structural IRF for variable k given a unit shock from variable m . To compute the MSE for each k , define

$$\text{MSE}_k(h) := \sum_{m=1}^K \mathbb{E} \left[\left(\hat{\theta}_{km}(h) - \theta_{km}(h) \right)^2 \right],$$

which is the total MSE for variable k over all possible structural shocks. In simulations, I use B replications to estimate the expectation. All MSEs are normalized by the mean squared error of the least squares estimator.

Setup A. A time series of length $T = 200$ is generated a number $B = 10\,000$ of times for replication. All VAR estimators are computed using $p = 10$ lags, while LPs include $q = 10$ regression lags. Table 2 shows relative MSEs for this design. It is important to notice that in this situation GLS ridge has remarkably low performance at horizon $h = 1$ compared to other methods. The primary issue is that Σ_u features strong correlation between components, and thus the diagonal lag-adapted structure does not shrink along the appropriate directions. This is much less prominent as the horizon increases due to the fact that impulse responses eventually decay to zero, since the underlying VARMA DGP is stationary. While there is no clear ranking, the MSE of the baseline ridge VAR estimator is in between those of the BVAR and hierarchical BVAR approaches. The degrading quality of local projection estimates are mainly due to the smaller samples available in regressions at each increasing horizon (Kilian and Kim, 2011). This behavior is one of the prime reasons behind the development of LP shrinkage estimators, like that proposed in Plagborg-Møller (2016) or the SLP estimator of Barnichon and Brownlees (2019).

Setup B. Using the data of Hansen (2016b), I estimate and simulate a stationary but highly persistent VAR(5) model using the same sample size and number of replications of Setup A. For all methods, $p = 5$ lags are used, so that VAR estimators are correctly specified. In this setup, unlike in the previous experiment, one can clearly notice that impulse responses computed via cross-validated ridge show increasing MSE as horizon h grows. There are two main reasons behind this behavior. First, the chosen setup features a very persistent data generating process, as the largest root of the underlying VAR model is 0.9945. This means that the true IRFs revert to zero only over long horizons, while lag-adapted ridge estimates yields models with lower persistence and thus flatter impulse responses. Secondly, the dataset from Hansen (2016b) is not normalized, and the included series have markedly heterogenous

Table 2: MSE Relative to OLS – Setup A

Variable	Method	$h = 1$	$h = 4$	$h = 8$	$h = 12$	$h = 16$	$h = 20$	$h = 24$
Investment Growth	RIDGE	0.97	0.74	0.64	0.64	0.65	0.63	0.60
	RIDGE-GLS	5.16	0.89	0.55	0.47	0.44	0.41	0.38
	LP	1.00	1.05	1.13	1.52	2.15	3.20	4.87
	BVAR-CV	1.55	0.84	0.70	0.70	0.71	0.70	0.66
	H-BVAR	1.80	0.66	0.53	0.52	0.54	0.53	0.50
Deflator	RIDGE	0.93	0.78	0.69	0.68	0.67	0.64	0.59
	RIDGE-GLS	2.43	0.83	0.59	0.52	0.48	0.44	0.40
	LP	1.00	1.05	1.13	1.44	1.99	2.90	4.47
	BVAR-CV	1.03	0.89	0.74	0.73	0.73	0.70	0.66
	H-BVAR	1.01	0.70	0.58	0.56	0.55	0.53	0.50
Paper Rate	RIDGE	0.94	0.76	0.66	0.66	0.66	0.64	0.60
	RIDGE-GLS	1.80	0.87	0.59	0.52	0.47	0.43	0.39
	LP	1.00	1.05	1.13	1.46	1.99	2.86	4.31
	BVAR-CV	0.87	0.87	0.74	0.73	0.73	0.71	0.66
	H-BVAR	0.81	0.69	0.57	0.55	0.56	0.54	0.51

Table 3: MSE Relative to OLS – Setup B

Variable	Method	$h = 1$	$h = 4$	$h = 8$	$h = 12$	$h = 16$	$h = 20$	$h = 24$
Real GDP	RIDGE	1.11	1.08	1.16	1.06	0.90	0.89	0.94
	RIDGE-GLS	1.16	1.00	0.99	1.00	0.93	0.93	0.95
	LP	1.00	1.14	1.37	1.52	1.72	1.98	2.24
	BVAR-CV	0.90	0.87	1.04	1.01	0.92	0.92	0.98
	H-BVAR	0.83	0.62	0.78	0.73	0.62	0.62	0.68
Investment	RIDGE	1.49	1.27	1.17	0.99	0.70	0.73	1.61
	RIDGE-GLS	1.34	1.14	1.02	1.02	0.86	0.82	0.86
	LP	1.00	1.15	1.40	1.63	2.03	2.76	3.59
	BVAR-CV	1.51	1.01	0.97	0.97	0.93	1.08	1.24
	H-BVAR	1.06	0.68	0.69	0.66	0.63	0.87	1.14
Fed Funds Rate	RIDGE	2.17	1.21	0.96	0.93	1.03	4.00	53.18
	RIDGE-GLS	1.21	1.04	0.90	0.93	0.90	0.88	0.91
	LP	1.00	1.18	1.51	1.71	1.97	2.44	2.99
	BVAR-CV	0.92	0.94	0.91	0.90	0.86	0.87	0.92
	H-BVAR	0.75	0.77	1.32	1.38	1.25	1.15	1.20

variances. Since GLS ridge shrinks along covariance-rotated data, shrinkage is adjusted according to each series variance, unlike that baseline ridge estimator $\hat{\beta}^R$. The MSE for the Fed Fund Rate impulse responses shows that the pointwise difference between baseline and GLS ridge can be severe for long horizon IRFs when the DGP is highly persistent. On short horizons, Bayesian estimators perform on par or better than baseline least squares estimates, while at longer horizons differences are less stark. It is, however, clear that the hierarchical prior BVAR of Giannone et al. (2015) shows the overall best results. As in the previous setup, local projections show degrading performance at larger horizons.

7.2 Confidence Intervals

I now try and evaluate whether ridge shrinkage has a negative impact on inference. There have also been recent contributions directly aimed at studying shrinkage effects. Li et al. (2022) give an extensive treatment of the issue in terms of bias-variance trade-off, showing that, to a large extent, shrinkage is desirable unless bias is a primary and sensitive concern. Using the same simulation setups as in the previous section, I investigate coverage and size properties of pointwise CIs constructed using the methods in Table 1. All confidence intervals are constructed with nominal 90% level coverage.

In this set of simulations, I swap GLS ridge for the asymptotic shrinkage ridge estimator, $\hat{\beta}_{as}^R$, see Section 6, since the latter allows for a partially non-negligible penalization in the limit. To implement $\hat{\beta}_{as}^R$, one needs to choose a partition of β which identifies asymptotically negligible coefficient. To do this, I split β by lag and penalize all coefficients with lag orders greater than a given threshold \bar{p} , such that $1 < \bar{p} < p$. In setup A, I choose $\bar{p} = 6$, while in setup B I set $\bar{p} = 3$. In Bayesian methods, including the cross-validated Minnesota BVAR, I construct high-probability intervals by drawing from the posterior. Comparison between frequentist CIs and Bayesian posterior densities is not generally valid, because they are not analogous concepts. Therefore, the discussion below is intended to highlight differences in *structure* between ridge approaches.

Setup A. Simulations with the DGP of Kilian and Kim (2011), presented in Tables 4 and 5, highlight some of the advantages of applying ridge when performing inference. Focusing on estimator $\hat{\beta}^R$, it is clear that CI coverage is in fact higher than the intervals obtained by least squares estimation in all situations. At impact, ridge CIs are larger than the LS baseline, but they shrink as horizons increase. Thus, as IRFs revert relatively quickly to zero, ridge can effectively reduce length while preserving coverage. As discussed in Section 3, these gains are inherently local to the DGP – shrinkage to zero at deep lags embodies correct prior knowledge of a weakly persistent process. For Bayesian estimators, one can

Table 4: Impulse Response Inference – Setup A – CI Coverage

Variable	Method	$h = 1$	$h = 4$	$h = 8$	$h = 12$	$h = 16$	$h = 20$	$h = 24$
Investment Growth	LS	0.88	0.88	0.87	0.88	0.91	0.93	0.94
	RIDGE	0.90	0.92	0.94	0.93	0.94	0.95	0.95
	RIDGE-AS	0.90	0.92	0.88	0.88	0.88	0.89	0.89
	LP	0.88	0.97	0.99	0.99	0.99	0.99	0.99
	BVAR-CV	0.77	0.88	0.88	0.90	0.92	0.94	0.96
	H-BVAR	0.72	0.89	0.89	0.92	0.93	0.95	0.96
Deflator	LS	0.88	0.87	0.86	0.88	0.91	0.92	0.94
	RIDGE	0.91	0.92	0.93	0.92	0.93	0.94	0.95
	RIDGE-AS	0.91	0.91	0.88	0.88	0.87	0.87	0.88
	LP	0.88	0.97	0.99	0.99	0.99	0.99	1.00
	BVAR-CV	0.80	0.86	0.88	0.91	0.93	0.94	0.96
	H-BVAR	0.84	0.88	0.90	0.92	0.94	0.95	0.97
Paper Rate	LS	0.87	0.86	0.86	0.88	0.90	0.92	0.94
	RIDGE	0.90	0.91	0.93	0.93	0.93	0.94	0.95
	RIDGE-AS	0.89	0.90	0.89	0.88	0.88	0.88	0.88
	LP	0.87	0.97	0.99	0.99	0.99	0.99	0.99
	BVAR-CV	0.82	0.84	0.87	0.90	0.92	0.93	0.95
	H-BVAR	0.88	0.88	0.90	0.92	0.93	0.95	0.96

Table 5: Impulse Response Inference – Setup A – CI Length

Variable	Method	$h = 1$	$h = 4$	$h = 8$	$h = 12$	$h = 16$	$h = 20$	$h = 24$
Investment Growth	LS	2.99	5.11	5.78	5.35	4.79	4.17	3.56
	RIDGE	3.13	5.20	5.82	5.17	4.48	3.78	3.09
	RIDGE-AS	3.11	5.15	4.84	4.33	3.70	3.06	2.48
	LP	2.99	7.50	10.97	12.89	13.99	14.55	14.70
	BVAR-CV	2.84	4.48	4.70	4.38	3.99	3.56	3.11
	H-BVAR	2.71	4.20	4.50	4.29	3.96	3.56	3.13
Deflator	LS	1.19	1.92	2.23	2.14	1.94	1.71	1.46
	RIDGE	1.24	1.97	2.25	2.09	1.84	1.54	1.26
	RIDGE-AS	1.24	1.95	1.95	1.78	1.52	1.25	1.01
	LP	1.19	3.03	4.56	5.42	5.90	6.14	6.21
	BVAR-CV	1.03	1.69	1.87	1.80	1.67	1.50	1.31
	H-BVAR	1.01	1.64	1.83	1.79	1.67	1.51	1.33
Paper Rate	LS	0.97	1.42	1.64	1.57	1.44	1.27	1.09
	RIDGE	1.01	1.44	1.65	1.53	1.36	1.16	0.95
	RIDGE-AS	1.01	1.43	1.42	1.31	1.13	0.94	0.77
	LP	0.97	2.19	3.28	3.90	4.26	4.43	4.48
	BVAR-CV	0.84	1.22	1.35	1.30	1.21	1.09	0.96
	H-BVAR	0.85	1.21	1.34	1.31	1.22	1.10	0.97

note that quantile intervals at small horizons tend to be shorter compared to least squares and ridge methods.

Setup B. The effects of ridge shrinkage on a DGP with high persistence are much more severe, as shown in Tables 6 and 7. Focusing on frequentist ridge, one can observe that close to impact ($h = 1$) ridge has similar or even higher coverage than other methods for real GDP⁷ However, as the IRF horizon grows, shrinkage often leads to severe undercoverage, with asymptotic shrinkage estimator $\hat{\beta}_{as}^R$ giving the worst results. In comparison, Bayesian methods are much more reliable at all horizons, although the only estimator that can consistently improve upon the benchmark least squares VAR CIs is the hierarchical prior BVAR of Giannone et al. (2015). The reason behind this is simple enough: the implementation of the Minnesota-prior BVAR I have used here has a white noise prior on all variables, which in this case is far from the truth. Indeed, Bańbura et al. (2010) implement the same BVAR by tuning the prior to a random walk for very persistent variables in their applications. In this sense the cross-validated BVAR considered – which is assumed centered at zero – is really the flip-side of ridge estimators. Therefore, the addition of a prior on the mean of the autoregressive parameters as done by Giannone et al. (2015) is a key element to perform shrinkage in high persistence setups in a way that does not systematically undermine asymptotic inference on impulse responses.

8 Conclusion

In this paper, I have studied ridge regression and its application to vector autoregressive model estimation in detail. This appears to be the first work that provides a thorough analysis of ridge penalization in the context of time series data, including geometric as well as asymptotic properties. I have also derived results on the validity of cross-validation as a method to select the penalty intensity in practice, and I have shown that CV produces asymptotically valid penalization rates. Finally, I have compared both frequentist and Bayesian ridge formulation in simulations aimed at quantifying the applicability of ridge for impulse response inference.

The key takeaway of this work is that ridge penalization is a useful approach to VAR estimation as long as the chosen penalty structure is well-adapted to the models' structure. Bayesian ridge posteriors are especially flexible, with hierarchical priors also allowing shrinkage towards non-zero coefficient vectors. However, it is important to note that the Bayesian approach also permits the researcher to specify uninformative priors, so that the influence of

⁷This also is the case with consumption and compensation, c.f. Tables 9 and 10 in Appendix D.5.

Table 6: Impulse Response Inference – Setup B: CI Coverage

Variable	Method	$h = 1$	$h = 4$	$h = 8$	$h = 12$	$h = 16$	$h = 20$	$h = 24$
Real GDP	LS	0.87	0.81	0.75	0.72	0.71	0.72	0.73
	RIDGE	0.90	0.79	0.66	0.62	0.65	0.68	0.68
	RIDGE-AS	0.89	0.72	0.61	0.58	0.61	0.65	0.65
	LP	0.87	0.93	0.94	0.94	0.93	0.93	0.91
	BVAR-CV	0.70	0.71	0.63	0.64	0.71	0.75	0.76
	H-BVAR	0.84	0.86	0.76	0.76	0.83	0.88	0.88
Investment	LS	0.87	0.82	0.76	0.73	0.75	0.82	0.87
	RIDGE	0.85	0.79	0.65	0.62	0.73	0.80	0.81
	RIDGE-AS	0.82	0.69	0.59	0.57	0.68	0.77	0.77
	LP	0.87	0.94	0.94	0.95	0.94	0.94	0.94
	BVAR-CV	0.70	0.73	0.67	0.71	0.77	0.81	0.83
	H-BVAR	0.80	0.86	0.81	0.82	0.87	0.88	0.88
Fed Funds Rate	LS	0.85	0.83	0.80	0.78	0.77	0.79	0.80
	RIDGE	0.79	0.77	0.74	0.68	0.68	0.72	0.72
	RIDGE-AS	0.78	0.66	0.68	0.64	0.64	0.68	0.69
	LP	0.85	0.94	0.96	0.96	0.95	0.94	0.93
	BVAR-CV	0.76	0.72	0.76	0.77	0.77	0.81	0.83
	H-BVAR	0.87	0.86	0.74	0.73	0.78	0.84	0.87

Table 7: Impulse Response Inference – Setup B: CI Length (rescaled $\times 100$)

Variable	Method	$h = 1$	$h = 4$	$h = 8$	$h = 12$	$h = 16$	$h = 20$	$h = 24$
Real GDP	LS	0.71	1.56	2.07	2.31	2.32	2.24	2.15
	RIDGE	0.79	1.56	1.85	1.95	1.92	1.85	1.77
	RIDGE-AS	0.74	1.31	1.65	1.76	1.75	1.70	1.64
	LP	0.71	2.42	4.21	5.40	5.90	5.91	5.70
	BVAR-CV	0.53	1.23	1.74	2.00	2.10	2.13	2.15
	H-BVAR	0.58	1.36	1.87	2.16	2.32	2.44	2.55
Investment	LS	3.38	6.65	7.89	7.89	7.31	6.69	6.18
	RIDGE	3.79	6.81	6.93	6.46	5.79	5.19	4.73
	RIDGE-AS	3.59	5.57	6.11	5.77	5.21	4.72	4.34
	LP	3.37	10.16	16.00	18.85	19.06	18.22	17.23
	BVAR-CV	2.64	5.26	6.59	6.91	6.78	6.57	6.38
	H-BVAR	2.89	5.74	7.08	7.54	7.63	7.60	7.58
Fed Funds Rate	LS	0.25	0.39	0.43	0.43	0.41	0.38	0.35
	RIDGE	0.29	0.39	0.37	0.36	0.33	0.30	0.29
	RIDGE-AS	0.27	0.31	0.33	0.32	0.30	0.28	0.27
	LP	0.25	0.59	0.88	1.01	1.05	1.03	0.98
	BVAR-CV	0.21	0.31	0.36	0.37	0.36	0.35	0.34
	H-BVAR	0.23	0.36	0.42	0.44	0.45	0.45	0.46

the priors' hyperparameters is less pronounced. This is not the case in frequentist ridge, c.f. including an explicit non-zero centering vector. However, prior knowledge or a pre-estimation procedure may be available to the researcher, so that ridge can be effectively implemented without the need to implement a BVAR.

To conclude, there are still avenues of research regarding ridge which would be interesting to develop. First and foremost, the high-dimensional setup, for which, however, it seems non-trivial to find a domain of applicability, as discussed in the introduction. Secondly, a more in-depth analysis of cross-validation, especially in the multivariate case, would be extremely valuable. Moreover, both the latter and former topics should be jointly addressed in the context of mild cross-sectional dimension growth, i.e. $K \rightarrow \infty$ such that $K/T \rightarrow \rho \in (0, 1)$, which is comparable to factor model setups.

Declaration of Interest Statement. The author reports there are no competing interests to declare.

References

- Babii, A., Ghysels, E., and Striaukas, J. (2021). Machine Learning Time Series Regressions With an Application to Nowcasting. *Journal of Business & Economic Statistics*, 40(3):1–23.
- Bañbura, M., Giannone, D., and Reichlin, L. (2010). Large Bayesian vector auto regressions. *Journal of Applied Econometrics*, 25(1):71–92.
- Barnichon, R. and Brownlees, C. (2019). Impulse Response Estimation by Smooth Local Projections. *The Review of Economics and Statistics*, 101(3):522–530.
- Bergmeir, C., Hyndman, R. J., and Koo, B. (2018). A note on the validity of cross-validation for evaluating autoregressive time series prediction. *Computational Statistics & Data Analysis*, 120:70–83.
- Bhatia, R. (1997). *Matrix Analysis*. Springer, New York, NY, New York, NY, USA.
- Boubacar Mainassara, Y. and Francq, C. (2011). Estimating structural VARMA models with uncorrelated but non-independent error terms. *Journal of Multivariate Analysis*, 102(3):496–505.
- Brüggemann, R., Jentsch, C., and Trenkler, C. (2016). Inference in VARs with conditional heteroskedasticity of unknown form. *Journal of Econometrics*, 191(1):69–85.
- Burman, P., Chow, E., and Nolan, D. (1994). A Cross-Validatory Method for Dependent Data. *Biometrika*, 81(2):351–358.
- Cavaliere, G., Gonçalves, S., and Nielsen, M. Ø. (2022). Bootstrap inference in the presence of bias. *Working Paper*.
- Davidson, J. (1994). *Stochastic limit theory: An introduction for econometricians*. OUP Oxford.
- 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.
- Dobriban, E. and Wager, S. (2018). High-dimensional asymptotics of prediction: Ridge regression and classification. *Annals of Statistics*, 46(1):247–279.
- Fu, W. and Knight, K. (2000). Asymptotics for lasso-type estimators. *The Annals of Statistics*, 28(5):1356–1378.
- Fuleky, P., editor (2020). *Macroeconomic Forecasting in the Era of Big Data: Theory and Practice*, volume 52 of *Advanced Studies in Theoretical and Applied Econometrics*. Springer International Publishing, Cham.

- Ghosh, S., Khare, K., and Michailidis, G. (2019). High-Dimensional Posterior Consistency in Bayesian Vector Autoregressive Models. *Journal of the American Statistical Association*, 114(526):735–748.
- Giannone, D., Lenza, M., and Primiceri, G. E. (2015). Prior Selection for Vector Autoregressions. *The Review of Economics and Statistics*, 97(2):436–451.
- Goulet Coulombe, P. (2023). Time-Varying Parameters as Ridge Regressions. *Working Paper*.
- Goulet Coulombe, P., Leroux, M., Stevanovic, D., and Surprenant, S. (2022). How is machine learning useful for macroeconomic forecasting? *J. Appl. Econ.*, 37(5):920–964.
- Hamilton, J. D. (1994). *Time Series Analysis*. Princeton University Press.
- Hansen, B. E. (2016a). Efficient shrinkage in parametric models. *Journal of Econometrics*, 190(1):115–132.
- Hansen, B. E. (2016b). Stein Combination Shrinkage for Vector Autoregressions. *Working Paper*.
- Hastie, T. (2020). Ridge Regularization: An Essential Concept in Data Science. *Technometrics*, 62(4):426–433.
- Hastie, T., Montanari, A., Rosset, S., and Tibshirani, R. J. (2022). Surprises in high-dimensional ridgeless least squares interpolation. *The Annals of Statistics*, 50(2):949–986.
- Hastie, T., Tibshirani, R., Friedman, J. H., and Friedman, J. H. (2009). *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. Springer, second edition.
- Hoerl, A. E. and Kennard, R. W. (1970). Ridge Regression: Applications to Nonorthogonal Problems. *Technometrics*, 12(1):69–82.
- Inoue, A. and Kilian, L. (2008). How Useful Is Bagging in Forecasting Economic Time Series? A Case Study of U.S. Consumer Price Inflation. *Journal of the American Statistical Association*, 103(482):511–522.
- Jordà, Ò. (2005). Estimation and Inference of Impulse Responses by Local Projections. *American Economic Review*, 95(1):161–182.
- Kilian, L. and Kim, Y. J. (2011). How Reliable Are Local Projection Estimators of Impulse Responses? *Review of Economics and Statistics*, 93(4):1460–1466.
- Kilian, L. and Lütkepohl, H. (2017). *Structural vector autoregressive analysis*. Cambridge University Press.

- Li, D., Plagborg-Møller, M., and Wolf, C. K. (2022). Local Projections vs. VARs: Lessons From Thousands of DGPs. *Working Paper*.
- Litterman, R. B. (1986). Forecasting with Bayesian vector autoregressions five years of experience. Working Papers 274, Federal Reserve Bank of Minneapolis.
- Liu, S. and Dobriban, E. (2020). Ridge Regression: Structure, Cross-Validation, and Sketching. In *International Conference on Learning Representations*.
- Lütkepohl, H. (1990). Asymptotic Distributions of Impulse Response Functions and Forecast Error Variance Decompositions of Vector Autoregressive Models. *The Review of Economics and Statistics*, 72(1):116–125.
- Lütkepohl, H. (2005). *New Introduction to Multiple Time Series Analysis*. Springer.
- Medeiros, M. C., Vasconcelos, G. F. R., Veiga, Á., and Zilberman, E. (2021). Forecasting Inflation in a Data-Rich Environment: The Benefits of Machine Learning Methods. *Journal of Business & Economic Statistics*, 39(1):98–119.
- Mikusheva, A. (2007). Uniform Inference in Autoregressive Models. *Econometrica*, 75(5):1411–1452.
- Mikusheva, A. (2012). One-Dimensional Inference in Autoregressive Models With the Potential Presence of a Unit Root. *Econometrica*, 80(1):173–212.
- Park, J. Y. and Phillips, P. C. (1988). Statistical inference in regressions with integrated processes: Part 1. *Econometric Theory*, 4(3):468–497.
- Park, J. Y. and Phillips, P. C. (1989). Statistical inference in regressions with integrated processes: Part 2. *Econometric Theory*, pages 95–131.
- Patil, P., Wei, Y., Rinaldo, A., and Tibshirani, R. (2021). Uniform Consistency of Cross-Validation Estimators for High-Dimensional Ridge Regression. In *Proceedings of The 24th International Conference on Artificial Intelligence and Statistics*, pages 3178–3186. PMLR.
- Pesavento, E. and Rossi, B. (2006). Small-sample confidence intervals for multivariate impulse response functions at long horizons. *Journal of Applied Econometrics*, 21(8):1135–1155.
- Phillips, P. C. B. (1988). Regression Theory for Near-Integrated Time Series. *Econometrica*, 56(5):1021–1043.
- Plagborg-Møller, M. (2016). *Essays in Macroeconometrics*. PhD thesis, Harvard University.
- Pratt, J. W. (1961). Length of Confidence Intervals. *Journal of the American Statistical Association*, 56(295):549–567.

- Rio, E. (2017). *Asymptotic Theory of Weakly Dependent Random Processes*, volume 80 of *Probability Theory and Stochastic Modelling*. Springer Berlin Heidelberg, Berlin, Heidelberg.
- Sims, C. A., Stock, J. H., and Watson, M. W. (1990). Inference in linear time series models with some unit roots. *Econometrica: Journal of the Econometric Society*, pages 113–144.
- Smeeke, S. and Wijler, E. (2018). Macroeconomic forecasting using penalized regression methods. *International Journal of Forecasting*, 34(3):408–430.
- Stephenson, W., Frangella, Z., Udell, M., and Broderick, T. (2021). Can we globally optimize cross-validation loss? Quasiconvexity in ridge regression. In *Advances in Neural Information Processing Systems*, volume 34, pages 24352–24364. Curran Associates, Inc.
- Whitney, H. (1972). *Complex Analytic Varieties*, volume 131. Addison-Wesley Reading.

Appendix

A Preliminaries

A.1 LS and RLS Estimators.

Lütkepohl (2005), Chapter 3, shows that the multivariate least squares and GLS estimator of parameter vector $\boldsymbol{\beta}$ is given by

$$\hat{\boldsymbol{\beta}} = ((Z'Z)^{-1}Z \otimes I_K)\mathbf{y}$$

as the minimizer of $S(\boldsymbol{\beta}) = T^{-1}\text{tr}[(Y - BZ)' \Sigma_u(Y - BZ)]$. The multivariate ridge-regularized least squares (RLS) – or, simply, ridge – estimator considered in this paper is defined to be the minimizer of the regularized problem,

$$\begin{aligned} S^R(\boldsymbol{\beta}; \Lambda) &= T^{-1}\text{tr}[(Y - BZ)'(Y - BZ)] + \text{tr}[B'\Lambda B] \\ &= \frac{\mathbf{y}'\mathbf{y}}{T} + \boldsymbol{\beta}' \left(\frac{ZZ'}{T} \otimes I_K \right) \boldsymbol{\beta} - 2\boldsymbol{\beta}' \frac{(Z \otimes I_K)\mathbf{y}}{T} + \boldsymbol{\beta}'\Lambda\boldsymbol{\beta} \end{aligned}$$

The first partial derivative,

$$\frac{\partial S^R(\boldsymbol{\beta}; \Lambda)}{\partial \boldsymbol{\beta}} = 2 \left(\frac{ZZ'}{T} \otimes I_K \right) \boldsymbol{\beta} - 2 \frac{(Z \otimes I_K)\mathbf{y}}{T} + 2\Lambda\boldsymbol{\beta}$$

gives the normal equations $(T^{-1}ZZ' \otimes I_K + \Lambda)\boldsymbol{\beta} = T^{-1}(Z \otimes I_K)\mathbf{y}$. The Hessian $\partial^2 S^R(\boldsymbol{\beta})/\partial^2 \boldsymbol{\beta} = 2(T^{-1}ZZ' \otimes I_K + \Lambda)\boldsymbol{\beta}$ is positive definite when $\Lambda > 0$, thus indeed the minimum is achieved by

$$\hat{\boldsymbol{\beta}}^R(\Lambda) = \left(\frac{ZZ'}{T} \otimes I_K + \Lambda \right)^{-1} \frac{(Z \otimes I_K)\mathbf{y}}{T}$$

Identical derivations prove that re-centering the ridge penalty at $\boldsymbol{\beta}_0 \in \mathbb{R}^{K^2p}$ produces the estimator

$$\hat{\boldsymbol{\beta}}^R(\Lambda, \boldsymbol{\beta}_0) = \left(\frac{ZZ'}{T} \otimes I_K + \Lambda \right)^{-1} \left(\frac{(Z \otimes I_K)\mathbf{y}}{T} + \Lambda\boldsymbol{\beta}_0 \right)$$

A.2 Structure of the Regularization Matrix

The vectorized RLS estimator $\hat{\boldsymbol{\beta}}^R(\Lambda)$ has maximal flexibility in terms of the regularization structure that matrix $\Lambda = \text{diag}\{\lambda_{1,1}, \dots, \lambda_{K,p}\}$ ($K^2p \times K^2p$) imposes. Since $\boldsymbol{\beta}$ contains all the coefficients of (A_1, \dots, A_p) it is indeed possible to individually penalize each lag of each series differently. In fact, by relaxing the assumption that Λ be a diagonal matrix, even more

general penalization structures are possible, although I do not consider them in this paper.

An interesting special case arises if the RLS estimator is instead written in its matrix form⁸,

$$\hat{B}_{\text{mat}}^R(\Lambda_{Kp}) = \frac{YZ'}{T} \left(\frac{ZZ'}{T} + \Lambda_{Kp} \right)^{-1}$$

where here it is of note that $\Lambda_{Kp} > 0$ has size $(Kp \times Kp)$. The regularization structure imposed is different in general than that in $\hat{\beta}^R(\Lambda)$: Λ_{Kp} induces *column-wise* ridge regularization, which penalizes coefficient estimates uniformly over each of the Kp columns of B . The associated vectorized estimator then simplifies:

$$\begin{aligned} \hat{\beta}^R(\Lambda_{Kp}) &= \left(\left(\frac{ZZ'}{T} + \Lambda_{Kp} \right) \otimes I_K \right)^{-1} (Z \otimes I_K) \mathbf{y} \\ &= \left(\left(\frac{ZZ'}{T} + \Lambda_{Kp} \right)^{-1} \frac{Z \otimes I_K}{T} \right) \mathbf{y} \end{aligned}$$

On the other hand, the *devectorized* RLS estimator is given by

$$\hat{B}^R(\Lambda_{K^2p}) = \text{reshape}(\beta^R(\Lambda), K, Kp)$$

that is, \hat{B}^R is simply a restructuring of the vectorized estimator into a matrix with identical dimensions to B . Importantly then, $\hat{B}^R(\Lambda_{K^2p})$ is equivalent to $\hat{B}^R(\Lambda_{Kp})$ if $\Lambda_{K^2p} = \Lambda_{Kp} \otimes I_K$. Because $\beta^R(\Lambda_{K^2p})$ and $\hat{B}^R(\Lambda_{K^2p})$ allow for the most generality in penalization structure, I will consider them to be the reference RLS estimators, so the dimension subscript to Λ will be dropped unless explicitly required.

A.3 Autocovariance and Asymptotic Conditioning

The conditioning of the autocovariance $\Sigma_y = \mathbb{E}[y_t y_t']$ is an important measure for the role that the regularization in the RLS estimator should be playing. This in turn depends on the eigenvalues of $\hat{\Sigma}_y$ with respect to those of Σ_y . [Hoerl and Kennard \(1970\)](#) showed in the linear regression setting that, when the sample covariance matrix deviates significantly from the identity matrix, its small eigenvalues excessively inflate the variance of least squares estimates, even though the regression problem itself is well-posed. This fragility is inherently a byproduct of finite sampling, and partially due to numerical procedures. Nowadays, numerical precision is virtually not a concern anymore, as robust linear algebra procedures are implicitly implemented in most scientific languages and toolboxes. Yet estimation issues

⁸For details in the least squares case, see again [Lütkepohl \(2005\)](#), Chapter 3. The derivations for the ridge estimator are identical.

tied to small or unfavorable data samples remain extremely relevant from both theoretical and practical viewpoints.

In the spirit of ridge as a regularization procedure, the following Lemma establishes convergence in probability of the ordered eigenvalues of the sample autocovariance matrix.

Lemma A.1. *If $\hat{\Sigma}_y = T^{-1} \sum_{t=1}^{T-1} y_t y_t' \xrightarrow{P} \Sigma_y$ where $\Gamma \in \mathbb{R}^{K \times K}$ is positive definite, then*

$$\omega_j(\hat{\Sigma}_y) \xrightarrow{P} \omega_j(\Sigma_y)$$

where $\omega_j(A)$ is the j largest eigenvalue of A .

Proof. First, recall that for all matrices $A \in \mathbb{R}^{K \times K}$, the determinant $\det(A)$ is clearly a continuous mapping⁹. Furthermore, for any polynomial $g(z) = z^n + a_1 z^{n-1} + \dots + a_n$, $a_i \in \mathbb{C}$ factored as $g(z) = (z - w_1) \cdots (z - w_n)$, $w_i \in \mathbb{C}$, where the ordering of roots w_i is arbitrary, it holds that for any $\epsilon > 0$ there exists $\delta > 0$ such that for every polynomial $h(z) = z^n + b_1 z^{n-1} + \dots + b_n$ with $|a_i - b_i| < \delta$ decomposed as $g(z) = (z - \bar{w}_1) \cdots (z - \bar{w}_n)$, $|w_i - \bar{w}_i| < \epsilon$, $i = 1, \dots, n$, see [Whitney \(1972\)](#), Appendix V.4. This in particular implies that the roots of the characteristic polynomial of matrix A are continuous functions of its coefficients.

Let $\varrho_{\hat{\Sigma}_y}(z) = z^K + a_1 z^{K-1} + \dots + a_K = (z - \hat{\omega}_1) \cdots (z - \hat{\omega}_K)$ and $\varrho_{\Sigma_y}(z) = z^K + b_1 z^{K-1} + \dots + b_K = (z - \omega_1) \cdots (z - \omega_K)$ be the (real) characteristic polynomials of $\hat{\Sigma}_y$ and Σ_y respectively. Because of the continuity arguments above, for every $\epsilon > 0$ there exist $\delta_1, \delta_2 > 0$ such that

$$\begin{aligned} \mathbb{P}(|\hat{\omega}_i - \omega_i| > \epsilon) &\leq \mathbb{P}(|a_i - b_i| > \delta_1) \\ &\leq \mathbb{P}(\|\hat{\Sigma}_y - \Sigma_y\| > \delta_2) \end{aligned}$$

for $i \in \{1, \dots, K\}$. Since by assumption $\hat{\Sigma}_y \xrightarrow{P} \Sigma_y$, the RHS of the above converges to zero as $T \rightarrow \infty$, thus $\hat{\omega}_i \xrightarrow{P} \omega_i$. \square

B Proofs

B.1 Shrinkage

Proof of Proposition 2

Proof. Notice that, by introducing $\Lambda_p := \text{diag}\{\lambda_1, \dots, \lambda_p\}$, any lag-adapted regularization

⁹This follows from $\det(A_{i,j}) = \sum_{\varsigma} \text{sign}(\varsigma) \prod_{i=1}^K A_{\varsigma(i),i}$ for permutation ς over $\{1, \dots, K\}$

matrix can be written as $\Lambda^{(p)} = \Lambda_p \otimes I_{K^2} = (\Lambda_p \otimes I_K) \otimes I_K$, so that

$$\begin{aligned}\hat{\boldsymbol{\beta}}^R(\Lambda_i^{(p)}) &= [(ZZ' + \Lambda_{p,i} \otimes I_K) \otimes I_K]^{-1} (Z \otimes I_K) \mathbf{y} \\ &= [(ZZ' + \Lambda_{p,i} \otimes I_K)^{-1} \otimes I_K] (Z \otimes I_K) \mathbf{y}\end{aligned}$$

by the properties of Kronecker product. It is now possible to derive the statements of the proposition as follows:

- (a) The result regarding isotropic regularizer $\Lambda^{(p)} = \lambda I_{K^2}$ is trivial given Proposition 1.
- (b) Without loss of generality due to the ordering of lags in Z , one may write the Gram matrix ZZ' in a block fashion,

$$ZZ' + \Lambda_p = \begin{bmatrix} (ZZ')_{[\mathcal{S}]} + \Lambda_{[\mathcal{S}]} & D \\ D' & (ZZ')_{[\mathcal{S}^c]} + \Lambda_{[\mathcal{S}^c]} \end{bmatrix}$$

where $(ZZ')_{[\mathcal{S}^c]}$ is the sub-matrix containing all the components *not* indexed by subset \mathcal{S} , and the subscript has been dropped from Λ_p for ease of notation.

Define $A_{[\mathcal{S}]} = (ZZ')_{[\mathcal{S}]} + \Lambda_{[\mathcal{S}]}$, $B_{[\mathcal{S}^c]} = (ZZ')_{[\mathcal{S}^c]} + \Lambda_{[\mathcal{S}^c]}$ and $\Delta = (B_{[\mathcal{S}^c]} - D'A_{[\mathcal{S}]}^{-1}D)$. The matrix block-inversion formula yields

$$(ZZ' + \Lambda_p)^{-1} = \begin{bmatrix} A_{[\mathcal{S}]}^{-1} + A_{[\mathcal{S}]}^{-1}D\Delta^{-1}D'A_{[\mathcal{S}]}^{-1} & A_{[\mathcal{S}]}^{-1}D\Delta^{-1} \\ -\Delta^{-1}D'A_{[\mathcal{S}]}^{-1} & \Delta^{-1} \end{bmatrix}.$$

If $\Lambda_{[\mathcal{S}]} \rightarrow 0$ and $\Lambda_{[\mathcal{S}^c]} \rightarrow \infty$, then $A_{[\mathcal{S}]} \rightarrow (ZZ')_{[\mathcal{S}]}$, $B_{[\mathcal{S}^c]} \rightarrow \infty$. Therefore $\Delta^{-1} \rightarrow 0$, since for $\Lambda_{[\mathcal{S}^c]}$ sufficiently large $\|B_{[\mathcal{S}^c]}^{-1}D'A_{[\mathcal{S}]}^{-1}D\| < 1$ and thus the Sherman-Morrison-Woodbury formula implies

$$\|(B_{[\mathcal{S}^c]} - D'A_{[\mathcal{S}]}^{-1}D)^{-1}\| \leq \frac{\|B_{[\mathcal{S}]}^{-1}\|}{1 - \|B_{[\mathcal{S}^c]}^{-1}D'A_{[\mathcal{S}]}^{-1}D\|} \rightarrow 0.$$

The above results finally yield

$$[(ZZ' + \Lambda_p)^{-1}Z \otimes I_K] \mathbf{y} \rightarrow \begin{bmatrix} (ZZ')_{[\mathcal{S}]} & 0 \\ 0 & 0 \end{bmatrix} (Z \otimes I_K) \mathbf{y} = \hat{\boldsymbol{\beta}}_{[\mathcal{S}]}^{LS}$$

as required. □

B.2 Ridge Asymptotic Theory

Proof of Theorem 1

Proof. (a) Assumptions A-B imply directly that $\hat{\Gamma}_T$ is a consistent estimator for Γ : in particular, y_t is a stationary, stable and ergodic VAR process.

(b) Rewriting $\hat{\beta}^R(\Lambda, \beta_0)$ yields

$$\begin{aligned}\hat{\beta}^R(\Lambda, \beta_0) &= \left(\frac{ZZ'}{T} \otimes I_K + \Lambda \right)^{-1} \left[T^{-1}(Z \otimes I_K)((Z' \otimes I_K)\beta + \mathbf{u}) + \Lambda\beta_0 \right] \\ &= \left(\frac{ZZ'}{T} \otimes I_K + \Lambda \right)^{-1} \left[\left(\frac{ZZ'}{T} \otimes I_K \right) \beta + \frac{(Z \otimes I_K)\mathbf{u}}{T} + \Lambda\beta_0 \right] \\ &= \left(\frac{ZZ'}{T} \otimes I_K + \Lambda \right)^{-1} \left[\left(\frac{ZZ'}{T} \otimes I_K + \Lambda \right) \beta + \frac{(Z \otimes I_K)\mathbf{u}}{T} + \Lambda(\beta_0 - \beta) \right] \\ &= \beta + \left(\frac{ZZ'}{T} \otimes I_K + \Lambda \right)^{-1} \Lambda(\beta_0 - \beta) + \left(\frac{ZZ'}{T} \otimes I_K + \Lambda \right)^{-1} \frac{(Z \otimes I_K)\mathbf{u}}{T}.\end{aligned}$$

I study the last two terms of the last equality separately. The first term is $o_p(1)$,

$$\left(\left(\frac{ZZ'}{T} \right) \otimes I_K + \Lambda \right)^{-1} \Lambda(\beta_0 - \beta) = \left(\left(\frac{ZZ'}{T} \right) \otimes I_K + o_p(1) \right)^{-1} o_p(1)(\beta_0 - \beta) \xrightarrow{P} 0$$

since $(\beta - \beta_0) = (\beta - \underline{\beta}_0) + (\underline{\beta}_0 - \beta_0) = (\beta - \underline{\beta}_0) + o_p(1)$. Considering the matrix sequence

$$\zeta_T = [T^{-1}(ZZ'), T^{-1}\Lambda],$$

which under Assumptions B and D.(ii) converges in probability to $[\Gamma, 0]$, by the continuous mapping theorem (Davidson, 1994) the second term gives

$$\left(\left(\frac{ZZ'}{T} \right) \otimes I_K + o_p(1) \right)^{-1} \left(\frac{1}{T}(Z \otimes I_K)\mathbf{u} \right) \xrightarrow{P} \Gamma^{-1} \mathbb{E}[(Z \otimes I_K)\mathbf{u}] = 0$$

under Assumption A.

(c) The residuals \hat{U} can be written as

$$\hat{U} = Y - \hat{B}^R Z = BZ + U - \hat{B}^R Z = U + (B - \hat{B}^R)Z$$

Thus

$$\frac{\hat{U}\hat{U}'}{T} = \frac{UU'}{T} + (B - \hat{B}^R) \left(\frac{ZZ'}{T} \right) (B - \hat{B}^R)' + (B - \hat{B}^R) \left(\frac{ZU'}{T} \right) + \left(\frac{UZ'}{T} \right) (B - \hat{B}^R)' \quad (9)$$

From (a) one finds that $\text{vec}(B) - \text{vec}(\hat{B}^R) = \beta - \hat{\beta}^R = o_p(1)$, so $(B - \hat{B}^R) \xrightarrow{P} 0$, while $T^{-1}(ZZ') \xrightarrow{P} \mathbb{E}[ZZ']$ and $T^{-1}(ZU') \xrightarrow{P} \mathbb{E}[ZU'] = 0$: the terms involving these

quantities then vanish asymptotically. Lastly, the first term of the sum gives

$$\frac{UU'}{T} = \frac{1}{T} \sum_{t=1}^T u_t u_t' \xrightarrow{P} \mathbb{E}[u_t u_t'] = \Sigma_u$$

for $T \rightarrow \infty$ under Assumptions A and B.

(d) With the same expansion used in (b),

$$\sqrt{T}(\hat{\boldsymbol{\beta}}^R(\Lambda, \boldsymbol{\beta}_0) - \boldsymbol{\beta}) = Q_T^{-1}(\sqrt{T}\Lambda)(\boldsymbol{\beta}_0 - \boldsymbol{\beta}) + Q_T^{-1}\left(\frac{1}{\sqrt{T}}(Z \otimes I_K)\mathbf{u}\right)$$

where $Q_T = (T^{-1}ZZ' + \Lambda) \xrightarrow{P} \Gamma$. Following the arguments above, the first term in the sum converges in probability,

$$Q_T^{-1}(\sqrt{T}\Lambda)(\boldsymbol{\beta}_0 - \boldsymbol{\beta} + o_p(1)) \xrightarrow{P} \Gamma^{-1}\Lambda_0(\boldsymbol{\beta} - \boldsymbol{\beta}_0)$$

The second term has normal limiting distribution,

$$Q_T^{-1}\left(\frac{1}{\sqrt{T}}(Z \otimes I_K)\mathbf{u}\right) \xrightarrow{d} \mathcal{N}(0, \Gamma^{-1} \otimes \Sigma_u)$$

see [Lütkepohl \(2005\)](#), Proposition 3.1. By Slutsky's theorem claim (d) follows. □

Proof of Theorem 2

Proof. (1) Since condition (i) implies that $\sqrt{T}\Lambda \xrightarrow{P} 0$, results (a)-(c) are unchanged, while (d) now involves the limit

$$Q_T^{-1}(\sqrt{T}\Lambda)(\boldsymbol{\beta}_0 - \boldsymbol{\beta}) = Q_T^{-1} \cdot o_P(1) \cdot (\boldsymbol{\beta}_0 - \boldsymbol{\beta} + o_P(1)) \xrightarrow{P} 0$$

yielding (d').

(2) Assuming $\boldsymbol{\beta}_0 \xrightarrow{P} \boldsymbol{\beta}$ simplifies the terms in the proof of Theorem 1 since now $\boldsymbol{\beta} - \boldsymbol{\beta}_0 = o_P(1)$. The weaker rate imposed on Λ does not influence results (a)-(c). Moreover,

$$Q_T^{-1} \cdot (\sqrt{T}\Lambda) \cdot (\boldsymbol{\beta}_0 - \boldsymbol{\beta}) = Q_T^{-1} \cdot O_P(1) \cdot o_P(1) \xrightarrow{P} 0$$

so that (d') follows. □

Proof of Theorem 3

Proof. I make a straightforward adaptation of the proof found in [Hamilton \(1994\)](#), Proposition 11.2. Define $\hat{\Sigma}_u^* = T^{-1}(UU')$, which is expanded to

$$\begin{aligned}\hat{\Sigma}_u^* &= \frac{1}{T}(Y - BZ)(Y - BZ)' \\ &= \frac{1}{T} \left(Y - \hat{B}^R Z + (\hat{B}^R - B)Z \right) \left(Y - \hat{B}^R Z + (\hat{B}^R - B)Z \right)' \\ &= \hat{\Sigma}_u^R + (\hat{B}^R - B) \left(\frac{ZZ'}{T} \right) (\hat{B}^R - B)' + \\ &\quad + \frac{1}{T} \left((Y - \hat{B}^R Z)Z'(\hat{B}^R - B)' + (\hat{B}^R - B)Z(Y - \hat{B}^R Z)' \right)\end{aligned}$$

Contrary to the least squares estimator, cross-terms do not cancel out since for $\Lambda \succeq 0$ the residuals $(Y - \hat{B}^R Z)$ are not in the orthogonal space of Z . From the consistency results of [Theorem 1](#),

$$(\hat{B}^R - B) \left(\frac{ZZ'}{T} \right) (\hat{B}^R - B)' = o_p(1) \left(\frac{ZZ'}{T} \right) o_p(1) \xrightarrow{P} 0$$

and

$$\sqrt{T}(\hat{B}^R - B) \left(\frac{ZZ'}{T} \right) (\hat{B}^R - B)' = O_p(1) \left(\frac{ZZ'}{T} \right) o_p(1) \xrightarrow{P} 0$$

Further,

$$\sqrt{T} \left[\frac{1}{T}(Y - \hat{B}^R Z)Z'(\hat{B}^R - B)' \right] = \left(\frac{\hat{U}Z'}{T} \right) \sqrt{T}(\hat{B}^R - B)' \xrightarrow{P} 0$$

since again $\sqrt{T}\hat{B}^R$ is asymptotically normal, and $T^{-1}(\hat{U}Z') = T^{-1}(UZ') + (B - \hat{B}^R) \cdot T^{-1}(ZZ') = T^{-1}(UZ') + o_p(1) \xrightarrow{P} \mathbb{E}[UZ'] = 0$. The same holds for the remaining transpose term, too.

It is thus proven that $\sqrt{T}(\hat{\Sigma}_u^* - \hat{\Sigma}_u^R) \xrightarrow{P} 0$, meaning $\sqrt{T}(\hat{\Sigma}_u^* - \Sigma_u) \xrightarrow{P} \sqrt{T}(\hat{\Sigma}_u^R - \Sigma_u)$ so that the two terms may be exchanged in computing the joint asymptotic distribution. [Theorem 1](#) accordingly yields

$$\begin{bmatrix} \hat{\beta}_0^R - \beta \\ \text{vec}(\hat{\Sigma}_u^R) - \text{vec}(\Sigma_u) \end{bmatrix} \xrightarrow{P} \begin{bmatrix} Q_T^{-1} (\sqrt{T}\Lambda) (\beta_0 - \beta) + Q_T^{-1} \frac{1}{\sqrt{T}} \boldsymbol{\xi} \\ \frac{1}{\sqrt{T}} \boldsymbol{\eta} \end{bmatrix}$$

where $\boldsymbol{\xi} = (Z \otimes I_K)\mathbf{u}$ and $\boldsymbol{\eta} = \text{vec}(UU' - \Sigma_u)$. As in [Hamilton \(1994\)](#), Proof of Proposition 11.2, $(\boldsymbol{\xi}', \boldsymbol{\eta})'$ is a martingale difference sequence, thus the claim

$$\sqrt{T} \begin{bmatrix} \hat{\beta}_0^R - \beta \\ \text{vec}(\hat{\Sigma}_u^R) - \text{vec}(\Sigma_u) \end{bmatrix} \xrightarrow{d} \mathcal{N} \left(\begin{bmatrix} \Gamma^{-1} \Lambda_0 (\beta_0 - \beta) \\ 0 \end{bmatrix}, \begin{bmatrix} \Gamma^{-1} \otimes \Sigma_u & 0 \\ 0 & \Omega \end{bmatrix} \right)$$

as $T \rightarrow \infty$ follows. When the strengthened assumptions (1) or (2) of [Theorem 2](#) are used

instead, the non-zero limiting mean vanishes

$$Q_T^{-1} \cdot \Lambda \cdot \sqrt{T}(\beta_0 - \beta) \xrightarrow{P} 0$$

proving that the joint asymptotic distribution is mean-zero Gaussian.

Finally, to compute the explicit expression of the asymptotic variance Ω , one must take care and note that u_t is not assumed to be normally distributed, contrary to the remainder of the proof in [Hamilton \(1994\)](#), pp. 342-343. A correct expression for i.i.d. non-Gaussian u_t can be found in Remark 2.1, [Brüggemann et al. \(2016\)](#), yielding

$$\Omega = \text{Var}[\text{vec}(u_t u_t')] = \mathbb{E}[\text{vec}(u_t u_t') \text{vec}(u_t u_t)'] - \boldsymbol{\sigma} \boldsymbol{\sigma}'$$

where $\boldsymbol{\sigma} = \text{vec}(\Sigma_u)$. □

Proof of Theorem 6

Proof. The stated results reduce to studying the behavior of two components used in the proof of Theorem 1 and Theorem 2, under the additional simplification of $\beta_0 = 0$.

- (a) Identical to result (a) in Theorem 1.
- (b) Consistency follows immediately by the fact that $\Lambda \beta \xrightarrow{P} 0$.
- (c) Follows from (c), Theorem 1 and (b) above.
- (d'') The bias term in the expression of $\sqrt{T}(\hat{\beta}^R - \beta)$ is driven by

$$\sqrt{T} \Lambda \beta = \begin{bmatrix} \sqrt{T} \Lambda_1 \cdot \beta_1 \\ \sqrt{T} \Lambda_2 \cdot T^{-(1/2+\delta)} \mathbf{b}_2 \end{bmatrix} = \begin{bmatrix} o_P(1) \cdot \beta_1 \\ \Lambda_2 \cdot T^{-\delta} \mathbf{b}_2 \end{bmatrix} \xrightarrow{P} 0$$

meaning there is no asymptotic bias. On the other hand,

$$\left((T^{-1} Z Z') \otimes I_K + \Lambda \right)^{-1} \xrightarrow{P} (\Gamma + \bar{\Lambda})^{-1} \otimes I_K.$$

Setting $(\Gamma + \bar{\Lambda}) = \Gamma_{\bar{\Lambda}}$ yields the claim since there are no further simplifications in the asymptotic variance formula, cf. proof of (d), Theorem 1. □

C Cross-validation

Later in this section, the following lemma will be useful.

Lemma 1. *Let y_t a stationary and stable mean-zero $AR(p)$ process with companion form matrix $\mathbb{A} \in \mathbb{R}^{p \times p}$. Then, the associated $MA(\infty)$ coefficients, $\{\vartheta_k\}_{k \in \mathbb{N}_0}$, decay exponentially for k sufficiently large, that is,*

$$|\vartheta_k| = O\left(\exp(-C_{\mathbb{A}} k)\right)$$

for some constant $C_{\mathbb{A}} > 0$.

Proof. Recall from e.g. [Lütkepohl \(2005\)](#) that if \mathbb{A} is the companion matrix of the $AR(p)$ model, then $\vartheta_k = \iota \mathbb{A}^j \iota'$, where $\iota := (1, 0, \dots, 0)' \in \mathbb{R}^p$. Note that $\|\iota\| = 1$ by construction and that the spectral radius of \mathbb{A} is less than one under the assumption of stability. Let $\mathbb{A} = V D V^{-1}$ be the Jordan canonical form of the companion matrix and ω_1 the dominant eigenvalue. Under the assumption of stability $|\omega_1| < 1$. Now observe that, supposing D has $l \leq p$ diagonal blocks, for $k \geq 0$

$$\frac{|\iota \mathbb{A}^k \iota'|}{|\omega_1^k|} = \left| \iota V \left(\frac{D^k}{\omega_1^k} \right) V^{-1} \iota' \right| = \left| \iota V \begin{bmatrix} [1] & & & \\ & \left[\frac{1}{\omega_1^k} D_2^k \right] & & \\ & & \ddots & \\ & & & \left[\frac{1}{\omega_1^k} D_l^k \right] \end{bmatrix} V^{-1} \iota' \right|,$$

where $[1]$ is the dominant Jordan block, while D_2, \dots, D_l are the non-dominant blocks. Then, letting $k \rightarrow \infty$, one gets

$$\frac{|\iota \mathbb{A}^k \iota'|}{|\omega_1^k|} \rightarrow \left| \iota V \begin{bmatrix} [1] & & & \\ & 0 & & \\ & & \ddots & \\ & & & 0 \end{bmatrix} V^{-1} \iota' \right| = C_D$$

for some constant $C_D < \infty$, as for $\ell \in \{2, \dots, l\}$ one can see that $D_\ell^k / \omega_1^k \rightarrow 0$. Since for complex ω_1 it holds $|\omega_1^k| = |\omega_1|^k$, letting $C_{\mathbb{A}} = -\log(|\omega_1|)$ the proof is complete. \square

C.1 Two-fold CV

Proof of Theorem 4

Proof. Write $\text{cv}2_m(\Lambda)$ as

$$\begin{aligned} \text{cv}2_m(\Lambda) &= \tilde{T}^{-1} \left(\tilde{\mathbf{y}} - \tilde{Z}' \hat{\beta}^R(\Lambda) \right)' \left(\tilde{\mathbf{y}} - \tilde{Z}' \hat{\beta}^R(\Lambda) \right) \\ &= \tilde{T}^{-1} \left[\tilde{\mathbf{u}} + \tilde{Z}' \left(\beta - \hat{\beta}^R(\Lambda) \right) \right]' \left[\tilde{\mathbf{u}} + \tilde{Z}' \left(\beta - \hat{\beta}^R(\Lambda) \right) \right] \end{aligned}$$

$$= (\boldsymbol{\beta} - \hat{\boldsymbol{\beta}}^R(\Lambda))' \left(\frac{\tilde{Z}\tilde{Z}'}{\tilde{T}} \right) (\boldsymbol{\beta} - \hat{\boldsymbol{\beta}}^R(\Lambda)) + 2 (\boldsymbol{\beta} - \hat{\boldsymbol{\beta}}^R(\Lambda))' \left(\frac{\tilde{Z}\tilde{\mathbf{u}}}{\tilde{T}} \right) + \frac{\tilde{\mathbf{u}}'\tilde{\mathbf{u}}}{\tilde{T}}.$$

By a strong LLN for weakly dependent processes (see, for example, [Rio 2017](#)), it holds that $\tilde{T}^{-1}\tilde{Z}\tilde{Z}' \xrightarrow{a.s.} \Gamma_z$, $\tilde{T}^{-1}\tilde{Z}\tilde{\mathbf{u}} \xrightarrow{a.s.} \mathbf{0}$ and $\tilde{T}^{-1}\tilde{\mathbf{u}}'\tilde{\mathbf{u}} \xrightarrow{a.s.} \Sigma_u$. Since, by a similar decomposition, it holds that

$$\text{err}(\Lambda) = (\boldsymbol{\beta} - \hat{\boldsymbol{\beta}}^R(\Lambda))' \Gamma (\boldsymbol{\beta} - \hat{\boldsymbol{\beta}}^R(\Lambda)) + \Sigma_u,$$

where $\Gamma = \mathbb{E}[z_t z_t']$ is positive definite, almost sure convergence is proven.

To prove uniform convergence over compact subsets, I follow the proof of [Patil et al. \(2021\)](#), Theorem 4.1, which relies on verifying the conditions of the Arzelà-Ascoli theorem. That is, one must prove that function $\text{cv}2_m(\Lambda)$ as well as its first derivatives are bounded over compact sets. As the Arzelà-Ascoli theorem readily generalizes to Euclidean spaces of arbitrary (fixed) dimension, I will directly consider the matrix derivative when checking boundedness.

Assume that $\Lambda \in \mathcal{I}$, where \mathcal{I} is a compact set of positive semidefinite penalization matrices Λ such that $\|\Lambda\|_{\max} < \infty$ and $\Lambda \succ \lambda_{\min} I$. Note $\boldsymbol{\beta} - \hat{\boldsymbol{\beta}}^R(\Lambda) = -(\hat{\Gamma} + \Lambda)^{-1} \Lambda \boldsymbol{\beta} + (\hat{\Gamma} + \Lambda)^{-1} (T^{-1} Z \mathbf{u})$ where $\hat{\Gamma} = T^{-1} Z Z'$. Using this decomposition, one gets first

$$\begin{aligned} |\text{err}(\Lambda)| &\leq \Sigma_u + \|\Gamma\|_2 \|\boldsymbol{\beta} - \hat{\boldsymbol{\beta}}^R(\Lambda)\|_2^2 \\ &\leq \Sigma_u + \|\Gamma\|_2 \left[\left\| (\hat{\Gamma} + \Lambda)^{-1} \Lambda \right\|_2^2 \|\boldsymbol{\beta}\|_2^2 + \left\| (\hat{\Gamma} + \Lambda)^{-1} \right\|_2^2 \left\| T^{-1} Z \mathbf{u} \right\|_2^2 \right. \\ &\quad \left. + 2 \left\| (\hat{\Gamma} + \Lambda)^{-1} \right\|_2^2 \left\| T^{-1} Z \mathbf{u} \right\|_2 \|\Lambda \boldsymbol{\beta}\|_2 \right] \\ &\stackrel{a.s.}{\leq} \Sigma_u + \|\Gamma\|_2 \frac{\omega_{\max}(\Lambda)^2 \|\boldsymbol{\beta}\|_2^2 + C_{zu}^2 + 2C_{zu} \omega_{\max}(\Lambda) \|\boldsymbol{\beta}\|_2}{(\omega_{\min}(\hat{\Gamma}) + \omega_{\min}(\Lambda))^2} \end{aligned}$$

where the last line follows from applying Weil's eigenvalue inequalities ([Bhatia, 1997](#)) to $(\hat{\Gamma} + \Lambda)^{-1}$ and the fact that $T^{-1} Z \mathbf{u} \xrightarrow{a.s.} \mathbf{0}$ by a strong LLN, so that there exists a constant $C_{zu} > 0$ bounding $\|T^{-1} Z \mathbf{u}\|_2$ for T large enough.

Additionally, the matrix derivative of $\boldsymbol{\beta} - \hat{\boldsymbol{\beta}}^R(\Lambda)$ with respect to Λ is

$$\frac{\partial(\boldsymbol{\beta} - \hat{\boldsymbol{\beta}}^R(\Lambda))}{\partial \Lambda} = -(\hat{\Gamma} + \Lambda)^{-1} \boldsymbol{\beta} + (\hat{\Gamma} + \Lambda)^{-2} \Lambda \boldsymbol{\beta} + (\hat{\Gamma} + \Lambda)^{-2} (T^{-1} Z \mathbf{u}),$$

so that, by using similar argument as the one used above, one gets

$$\begin{aligned} \left| \frac{\partial \text{err}(\Lambda)}{\partial \Lambda} \right| &\stackrel{a.s.}{\leq} \Sigma_u + \|\Gamma\|_2 \left[\frac{\|\boldsymbol{\beta}\|_2^2}{(\omega_{\min}(\hat{\Gamma}) + \omega_{\min}(\Lambda))^2} + \frac{2 \omega_{\max}(\Lambda) \|\boldsymbol{\beta}\|_2^2 + 2C_{zu} \|\boldsymbol{\beta}\|_2}{(\omega_{\min}(\hat{\Gamma}) + \omega_{\min}(\Lambda))^3} \right. \\ &\quad \left. + \frac{\omega_{\max}(\Lambda)^2 \|\boldsymbol{\beta}\|_2^2 + C_{zu}^2 + 2C_{zu} \omega_{\max}(\Lambda) \|\boldsymbol{\beta}\|_2}{(\omega_{\min}(\hat{\Gamma}) + \omega_{\min}(\Lambda))^4} \right]. \end{aligned}$$

The almost sure bound in the last display is also clearly finite for any $\Lambda \in \mathcal{I}_\lambda$, as required.

One can easily bound $\text{cv}2_m(\Lambda)$ and its first derivative as $\text{err}(\Lambda)$, with only addition of an extra term depending on $(\boldsymbol{\beta} - \hat{\boldsymbol{\beta}}^R(\Lambda))'(\tilde{T}^{-1}\tilde{Z}\tilde{\boldsymbol{u}})$. This means that $\text{err}(\Lambda) - \text{cv}2_m(\Lambda)$ forms an equicontinuous family of functions with respect to Λ over any \mathcal{I}_λ . Therefore, Arzelà-Ascoli yields uniform convergence of a subsequence, and since the difference converges to zero pointwise, too, the entire sequence converges uniformly. \square

C.2 Cross-validation under Dependence

The result of Theorem 4 may be only partially informative in practice, as it does not give information on how dependence, in terms of the buffer block of size m , impacts $\text{cv}2_m(\Lambda)$. Indeed, due to averaging, the effects of time dependence between the estimation and evaluation folds are washed out in the limit $\tilde{T} \rightarrow \infty$ even when m is fixed. Therefore, Theorem 4 is not useful in finite samples, where one would preferably set m to be as small as possible.

To address dependence, in the same setup as above, consider an alternative predictive error measure, the m -dependence prediction error,

$$\text{Err}_m(\hat{\boldsymbol{\beta}}^R(\Lambda)) := \mathbb{E}_{y_{T+m+1}, z_{T+m+1}} \left[\left(y_{T+m+1} - z'_{T+m+1} \hat{\boldsymbol{\beta}}^R(\Lambda) \right)^2 \middle| Z, \mathbf{y} \right],$$

and the associated error curve, $\text{err}_m(\Lambda) := \text{Err}_m(\hat{\boldsymbol{\beta}}^R(\Lambda))$. The empirical counterpart to this quantity is given by $\text{cv}2_m(\Lambda)$ for $\tilde{T} = 1$. The next theorem shows that in the case of a purely autoregressive data generating process, the error one commits by choosing a finite buffer size is exponentially small.

Theorem 7. *Under Assumptions A-C, for every Λ in the cone of diagonal positive definite penalty matrices with diagonal entries in (λ_{\min}, ∞) , it holds that*

$$\text{err}_m(\Lambda) - \text{err}(\Lambda) = O(\exp(-C_\beta m))$$

where C_β is a constant that does not depend on Λ .

Proof. In line with the definition of $\text{cv}2_m(\Lambda)$, I set $\tilde{y}_1 = y_{T+m+1}$, $\tilde{z} = z_{T+m+1}$ and $\tilde{u}_1 = y_{T+m+1} - z'_{T+m+1}\boldsymbol{\beta}$. With the same approach as in the proof of Theorem 4, here one finds

$$\begin{aligned} \text{err}_m(\Lambda) &= (\boldsymbol{\beta} - \hat{\boldsymbol{\beta}}^R(\Lambda))' \mathbb{E} [\tilde{z}_1 \tilde{z}'_1 | Z, \mathbf{y}] (\boldsymbol{\beta} - \hat{\boldsymbol{\beta}}^R(\Lambda)) \\ &\quad + 2 (\boldsymbol{\beta} - \hat{\boldsymbol{\beta}}^R(\Lambda))' \mathbb{E} [\tilde{z}_1 \tilde{u}_1 | Z, \mathbf{y}] + \mathbb{E} [\tilde{u}_1^2 | Z, \mathbf{y}], \end{aligned}$$

where I have removed the subscript from expectation \mathbb{E} to make notation clearer. Since \tilde{u}_1 is independent of \tilde{z}_1 , the cross term reduces to zero, while $\mathbb{E} [\tilde{u}_1^2 | Z, \mathbf{y}] = \Sigma_u$. Thus, it is the

first term in the last display that is effected by dependence.

To see this, let \tilde{z}_{i1} for $1 \leq i \leq p$ be the i th entry of \tilde{z}_1 . Then, using the MA(∞) decomposition of y_t , i.e. $y_t = \sum_{\ell=0}^{\infty} \phi_{\ell} u_{t-\ell}$, one can write

$$\begin{aligned} \tilde{z}_{i1} = z_{T+m+1-i} &= \sum_{\ell=0}^{m-i} \phi_{\ell} u_{T+m+1-i-\ell} + \sum_{\ell=M+1-i}^{\infty} \phi_{\ell} u_{T+m+1-i-\ell} \\ &= \sum_{\ell=0}^{m-i} \phi_{\ell} u_{T+m+1-i-\ell} + \sum_{s=0}^{\infty} \phi_{m+1-i+s} u_{T-s} \\ &= \eta_i + \zeta_i. \end{aligned}$$

Note that η_i is independent of ζ_i , and ζ_i is belongs with the σ -algebra generated by Z and \mathbf{y} . Therefore,

$$\begin{aligned} \mathbb{E} [\tilde{z}_1 \tilde{z}'_1 | Z, \mathbf{y}] &= \mathbb{E} [(\eta_i + \zeta_i)(\eta_i + \zeta_i)' | Z, \mathbf{y}] \\ &= \Gamma_{\eta} + \zeta_i \zeta'_i, \end{aligned}$$

as $\mathbb{E} [\eta_i \zeta'_i | Z, \mathbf{y}] = \mathbb{E} [\eta_i | Z, \mathbf{y}] \zeta'_i = 0$.

Now, I prove that $\Gamma_{\eta} \rightarrow \Gamma_z$ and $\zeta_i \zeta'_i \rightarrow 0$ at an exponential rate. First, let

$$H_{\phi} := \begin{bmatrix} \sum_{s=0}^{\infty} \phi_{m+1+s}^2 & \sum_{s=0}^{\infty} \phi_{m+1+s} \phi_{m+1+s-1} & & & \\ \sum_{s=0}^{\infty} \phi_{m+1+s-1} \phi_{m+1+s} & \sum_{s=1}^{\infty} \phi_{m+1+s}^2 & & & \\ & & \ddots & & \\ & & & \ddots & \\ & & & & \sum_{s=p}^{\infty} \phi_{m+1+s}^2 \end{bmatrix}$$

and observe that

$$\|\Gamma_{\eta} - \Gamma_z\|_2 \leq \Sigma_u p \|H_{\phi}\|_{\max} \leq C \exp(-C_{\eta} m),$$

since $\Gamma_{\eta} = \Sigma_u \otimes H_{\phi}$, lag order p is fixed and the maximal entry of H_{ϕ} decays exponentially for m sufficiently large following Lemma 1. Secondly, much in the same vein

$$\|\zeta_i \zeta'_i\|_2 = \zeta'_i \zeta_i \leq p \left(\sum_{s=0}^{\infty} \phi_{m+1-i+s} u_{T-s} \right)^2 \leq C' \exp(-C_{\zeta} m).$$

The proof concludes by setting $C_{\beta} = \max(C_{\eta}, C_{\zeta})$. \square

Remark 4. Theorem 7 is reassuring because it suggests that, in practice, if the AR(p) model is correctly specified, one may keep m small and still get a valid prediction error estimate in sense of Theorem 4. In simulations, I set $m = 0$, which is a common simplification to more effectively exploit the entire sample and does not, as discussed above, effect consistency (Bergmeir et al., 2018). Moreover, note that Theorem 7 intuitively gives a *worst-case* rate:

the dependence of between \tilde{z}_t and data in the estimation set gets milder, on average, as \tilde{T} grows. Thus, if CV aspect ratio \tilde{T}/T is balanced, dependence only plays a negligible role.

C.3 Asymptotically Valid CV

Proof of Theorem 5

Proof. First, recall that

$$\begin{aligned}\hat{\beta}_{\blacklozenge}^R(\Lambda) &= \left(\frac{ZZ'}{T} + \sqrt{T}\Lambda \right)^{-1} \frac{Z\mathbf{y}}{T} \\ &= \beta - \left(\frac{ZZ'}{T} + \sqrt{T}\Lambda \right)^{-1} (\sqrt{T}\Lambda) \beta + \left(\frac{ZZ'}{T} + \sqrt{T}\Lambda \right)^{-1} \frac{Z\mathbf{u}}{T}.\end{aligned}$$

It also holds $\text{Err}(\hat{\beta}_{\blacklozenge}^R(\Lambda)) = (\beta - \hat{\beta}_{\blacklozenge}^R(\Lambda))' \Gamma (\beta - \hat{\beta}_{\blacklozenge}^R(\Lambda)) + \Sigma_u$. Now, notice that

$$(T^{-1}ZZ' + \sqrt{T}\Lambda)^{-1} = O_P(1) \quad \text{and} \quad (T^{-1}ZZ' + \sqrt{T}\Lambda)^{-1} (\sqrt{T}\Lambda) = O_P(1),$$

since $\Lambda \in \mathcal{I}_\lambda$. It follows

$$\text{Err}(\hat{\beta}_{\blacklozenge}^R(\Lambda)) = \beta' (\sqrt{T}\Lambda) \left(\frac{ZZ'}{T} + \sqrt{T}\Lambda \right)^{-1} \Gamma \left(\frac{ZZ'}{T} + \sqrt{T}\Lambda \right)^{-1} (\sqrt{T}\Lambda) \beta + \Sigma_u + O_P(\sqrt{T}).$$

One can now consider a sequence $\tilde{\Lambda} = O_p(T^{-1/2})$ of regularizers in \mathcal{I}_λ . By taking the limit, one gets that

$$\begin{aligned}\lim_{T \rightarrow \infty} \text{Err}(\hat{\beta}_{\blacklozenge}^R(\tilde{\Lambda})) &= \beta' \cdot O_p(1) \cdot (\Gamma + O_p(1))^{-1} \Gamma (\Gamma + O_p(1))^{-1} \cdot O_p(1) \cdot \beta + \Sigma_u \\ &\geq \Sigma_u = \lim_{T \rightarrow \infty} \text{Err}(\hat{\beta}_{\blacklozenge}^R(0)),\end{aligned}$$

meaning that $\tilde{\Lambda}$ can not be optimal asymptotically, since the least squares solution at $\Lambda = 0$ achieves a lower predictive error. Additionally, any sequence with lower convergence order is also asymptotically invalid. Therefore, by contradiction, it must hold that $\Lambda_{\blacklozenge} = o_p(T^{-1/2})$. \square

Proof of Corollary 2

Proof. The first results follows directly from Theorem 5. Further, the fact that $\Lambda_{2,\blacklozenge} = O_P(1)$ is trivial because it is assumed that $\Lambda \in \mathcal{I}_\lambda$.

To prove the second part of the theorem, one can simply notice that, given the assump-

tion on the coefficients β in Theorem 6,

$$\left(\frac{ZZ'}{T} + \sqrt{T}\Lambda\right)^{-1} (\sqrt{T}\Lambda) \beta = \left(\frac{ZZ'}{T} + \sqrt{T}\Lambda\right)^{-1} \begin{bmatrix} \sqrt{T}\Lambda_1 \cdot \beta_1 \\ \Lambda_2 \cdot T^{-\delta} \mathbf{b}_2 \end{bmatrix} \xrightarrow{P} 0,$$

where Λ was partitioned into two diagonal blocks, Λ_1 and Λ_2 , as done previously. Block Λ_1 must be $O_P(T^{-1/2})$ following the proof of Theorem 5. Finally, the fact that $T^{-\delta} \mathbf{b}_2 \rightarrow 0$ as $T \rightarrow \infty$ means that, in the limit, a nonzero Λ_2 does not yield a sub-optimal cross-validation loss value. \square

D Monte Carlo Simulations

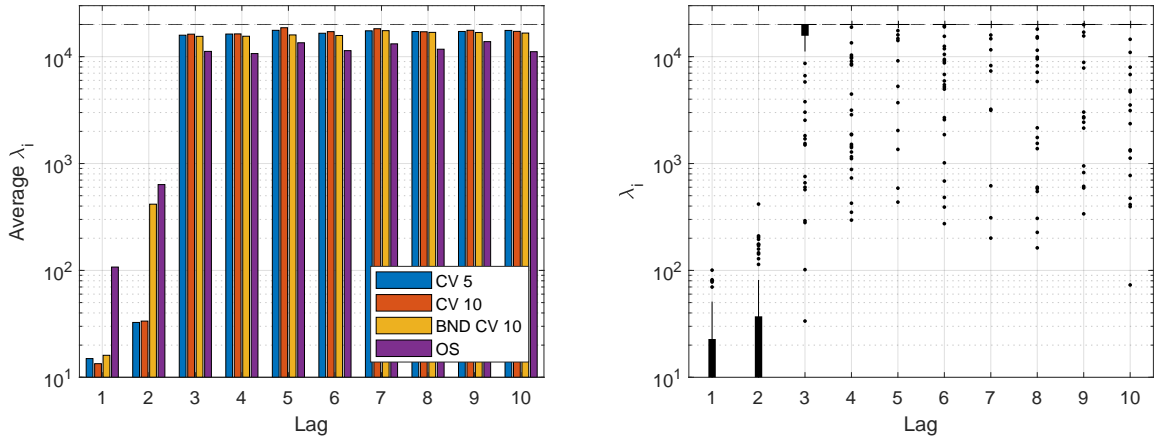
D.1 Cross-validation Details

To select the ridge penalty, I implement the lag-adapted structure and choose the relevant λ_i 's using block non-dependent cross-validation, c.f. [Burman et al. \(1994\)](#), [Bergmeir et al. \(2018\)](#). I constraint the optimization domain of λ_i to be $[0, 10^2]$, without discretization. An issue with cross-validation regards the GLS ridge estimator: the matrices involved can quickly become prohibitively large due to Kronecker products, making CV optimization impractical. To avoid this, I set the penalty for lag-adapted $\hat{\beta}^{RGLS}(\Lambda)$ to be the same as that obtained for $\hat{\beta}^R$ via CV, which means the regularizer is tuned sub-optimally. Nonetheless, an identical choice of Λ for both methods can help shed light on the difference in structure between the two estimators.

In contrast to [Bańbura et al. \(2010\)](#), I do not tune the shrinkage parameter of the Minnesota BVAR using a mean squared forecasting error (MSFE) criterion: instead, I again use block CV. Since a Minnesota prior can be easily implemented with the use of augmented regression matrices, cross-validation can be much more efficiently implemented than for $\hat{\beta}^{RGLS}$. The resulting choice of prior tightness λ^2 is reasonable because CV, too, estimates the (one step ahead) forecasting risk. Since in this context only the mean of the posterior is used to compute pointwise impulse responses, one can even directly interpret the cross-validated Minnesota BVAR estimator as a refinement of GLS ridge.

D.2 Penalty Selection in Simulations

In the simulations of Section 7, an interesting aspect to study is how data-driven penalty selection methods behave. Both average and individual behavior are important, because the former generally gives intuition for the kind of regularization structure that is selected for the model, while the latter is relevant for empirical modeling where estimation can be done



(a) Average penalty selected by lag.

(b) Distribution of 10-fold CV penalties.

Figure 2: Comparison of penalty selection methods in Setup A, 1000 replications.

only on one sample.

I use Setup A from the Monte Carlo experiments (with 1000 replications) and instead just focus on the behavior of a number of penalty validation techniques. The rationale behind the choice is straightforward: Setup A involves models with more lags. Figure 2a shows the mean selected penalty parameter λ_i for $i \in \{1, \dots, 10\}$. The methods I compare are: out-of-sample validation (OOS) with a split of 80% of sample for estimation and 20% for testing; block cross-validation with 5 (CV 5) or 10 (CV 10) folds; block non-dependent cross validation with 10 folds (BND CV 10). The differences between 5 and 10-fold block CV methods are small, and both largely agree with BND cross-validation apart at lag 2. In contrast, out-of-sample validation appears to select on average much higher penalties at early lags and slightly lower ones at higher lags.

Regarding the distribution of selected λ_i over all replications, one can notice from Figure 2b that there is indeed important variation in the individual penalty choices. The evidence is for the specific case of 10-fold block CV, but it appears as a common pattern with other techniques, too. The implications of such variability in lag-adapted penalties are hard to gauge because in any given sample it is not possible to say whether the choice of $\{\lambda_i\}$ is good or bad outside of speculation. A guiding principle might be to compare CV with any "hyperpriors" one might have on Λ itself – like in Bayesian paradigm – but then parameteric penalty matrices like the one used with the Minnesota prior should be preferred. Indeed, the question of whether a more robust but still general method other than cross-validation can be applied to the time series context is highly relevant.

D.3 Penalty Selection with Many Lags

The minimization problem involved in OOS and CV penalty selection in its most general form suffers from the curse of dimensionality. This is somehow mitigated when using a lag-adapted regularizer since the loss with Λ^ℓ depends only on p non-negative parameters, rather than K^2p with non-block-diagonal Λ . But whenever p is chosen large (e.g. $p > 20$) the problem resurfaces.

I suggest a basic shortcut to make computation easier. Such simplification stems from the following observation. If one is willing to believe the assumptions of Section 6, then, because deep lag coefficients are small, after the first few lags penalization can be equally strong on all remaining lags with negligible additional bias. The shortcut, then, is to estimate only $\{\lambda_1, \dots, \lambda_r\}$ for $r < p$, then extrapolate and use $\{\lambda_1, \dots, \lambda_{r-1}, \lambda_r, \dots, \lambda_r\}$, where λ_r is repeated $p - r$ times, as lag-adapted penalty parameters. The idea is supported by the results in Figure 2a. However, this strategy is not generally appropriate, because it could be that even at relatively deep lags some coefficients are large, while on the other hand the early coefficients are small. Therefore, in applications where the ridge penalty needs to be estimated only once or a handful of times I would suggest to avoid this shortcut altogether.

D.4 Numerical Optimization

For a VAR(p) and a lag-adapted Λ^ℓ , a collection $\{\lambda_1, \dots, \lambda_p\}$ must be chosen. To implement OS and CV for the ridge estimators, I rely on MATLAB optimization routines, in both cases using the optimization function `patternsearch` from the MATLAB Optimization Toolbox. The domain of optimization is chosen to be the hypercube $[0, 10^2]^p$, where T is the sample size. The choice of a bounded domain is asymptotically valid, c.f. Theorem 5 and Corollary 2.

In applications, since the CV loss needs not be convex (Stephenson et al., 2021), it is appropriate to employ advanced optimization routines, e.g. genetic or pattern-based optimizers like `patternsearch`, if possible. When one only requires to estimate the VAR model once, then the selection of Λ is a one-time cost. The gains of better optimization solutions therefore are often superior to the higher computational costs one incurs in when using more sophisticated routines.

D.5 Additional Tables

Table 8: MSE Relative to OLS – Setup B

Variable	Method	$h = 1$	$h = 4$	$h = 8$	$h = 12$	$h = 16$	$h = 20$	$h = 24$
Real GDP	RIDGE	1.11	1.08	1.16	1.06	0.90	0.89	0.94
	RIDGE-GLS	1.16	1.00	0.99	1.00	0.93	0.93	0.95
	LP	1.00	1.14	1.37	1.52	1.72	1.98	2.24
	BVAR-CV	0.90	0.87	1.04	1.01	0.92	0.92	0.98
	H-BVAR	0.83	0.62	0.78	0.73	0.62	0.62	0.68
GDP Deflator	RIDGE	2.25	2.10	1.81	1.54	1.39	1.47	3.37
	RIDGE-GLS	1.17	1.09	1.10	1.10	1.08	1.05	1.04
	LP	1.00	1.13	1.25	1.33	1.40	1.47	1.54
	BVAR-CV	1.06	1.03	1.00	0.95	0.92	0.92	0.93
	H-BVAR	0.81	0.91	1.10	1.11	1.04	0.98	0.94
Consumption	RIDGE	0.99	1.29	1.31	1.08	0.96	1.05	2.31
	RIDGE-GLS	0.94	0.96	1.04	0.99	0.95	0.96	0.99
	LP	1.00	1.13	1.32	1.44	1.63	1.83	2.03
	BVAR-CV	1.04	1.07	1.15	1.00	0.92	0.94	0.99
	H-BVAR	0.94	0.83	0.98	0.83	0.78	0.83	0.91
Investment	RIDGE	1.49	1.27	1.17	0.99	0.70	0.73	1.61
	RIDGE-GLS	1.34	1.14	1.02	1.02	0.86	0.82	0.86
	LP	1.00	1.15	1.40	1.63	2.03	2.76	3.59
	BVAR-CV	1.51	1.01	0.97	0.97	0.93	1.08	1.24
	H-BVAR	1.06	0.68	0.69	0.66	0.63	0.87	1.14
Hours	RIDGE	1.22	1.24	1.18	1.03	0.77	0.76	1.27
	RIDGE-GLS	1.07	1.05	1.01	1.03	0.90	0.85	0.90
	LP	1.00	1.14	1.33	1.53	1.81	2.35	2.92
	BVAR-CV	0.89	0.88	1.03	1.02	0.91	0.95	1.05
	H-BVAR	0.77	0.71	0.97	0.96	0.81	0.85	0.98
Compensation	RIDGE	0.85	0.99	0.85	0.94	1.13	1.40	4.70
	RIDGE-GLS	0.93	0.97	0.89	0.94	1.04	1.04	1.00
	LP	1.00	1.18	1.52	1.78	1.90	1.93	1.99
	BVAR-CV	1.07	0.92	0.94	0.93	0.99	1.01	0.98
	H-BVAR	0.86	0.80	1.12	1.31	1.35	1.27	1.22
Fed Funds Rate	RIDGE	2.17	1.21	0.96	0.93	1.03	4.00	53.18
	RIDGE-GLS	1.21	1.04	0.90	0.93	0.90	0.88	0.91
	LP	1.00	1.18	1.51	1.71	1.97	2.44	2.99
	BVAR-CV	0.92	0.94	0.91	0.90	0.86	0.87	0.92
	H-BVAR	0.75	0.77	1.32	1.38	1.25	1.15	1.20

Table 9: Impulse Response Inference – Setup B: CI Coverage

Variable	Method	$h = 1$	$h = 4$	$h = 8$	$h = 12$	$h = 16$	$h = 20$	$h = 24$
Real GDP	LS	0.87	0.81	0.75	0.72	0.71	0.72	0.73
	RIDGE	0.90	0.79	0.66	0.62	0.65	0.68	0.68
	RIDGE-AS	0.89	0.72	0.61	0.58	0.61	0.65	0.65
	LP	0.87	0.93	0.94	0.94	0.93	0.93	0.91
	BVAR-CV	0.70	0.71	0.63	0.64	0.71	0.75	0.76
	H-BVAR	0.84	0.86	0.76	0.76	0.83	0.88	0.88
GDP Deflator	LS	0.86	0.83	0.80	0.76	0.73	0.72	0.70
	RIDGE	0.85	0.76	0.66	0.62	0.61	0.60	0.58
	RIDGE-AS	0.83	0.70	0.61	0.58	0.57	0.57	0.55
	LP	0.86	0.93	0.94	0.94	0.93	0.91	0.89
	BVAR-CV	0.76	0.72	0.70	0.72	0.72	0.71	0.70
	H-BVAR	0.84	0.83	0.79	0.78	0.78	0.78	0.77
Consumption	LS	0.87	0.80	0.75	0.72	0.70	0.70	0.70
	RIDGE	0.90	0.74	0.60	0.60	0.64	0.66	0.65
	RIDGE-AS	0.89	0.67	0.55	0.55	0.60	0.62	0.62
	LP	0.86	0.93	0.94	0.94	0.94	0.92	0.90
	BVAR-CV	0.73	0.66	0.60	0.63	0.70	0.73	0.74
	H-BVAR	0.84	0.79	0.70	0.74	0.79	0.82	0.84
Investment	LS	0.87	0.82	0.76	0.73	0.75	0.82	0.87
	RIDGE	0.85	0.79	0.65	0.62	0.73	0.80	0.81
	RIDGE-AS	0.82	0.69	0.59	0.57	0.68	0.77	0.77
	LP	0.87	0.94	0.94	0.95	0.94	0.94	0.94
	BVAR-CV	0.70	0.73	0.67	0.71	0.77	0.81	0.83
	H-BVAR	0.80	0.86	0.81	0.82	0.87	0.88	0.88
Hours	LS	0.86	0.81	0.76	0.74	0.74	0.79	0.81
	RIDGE	0.88	0.80	0.68	0.64	0.70	0.72	0.66
	RIDGE-AS	0.87	0.74	0.62	0.59	0.65	0.67	0.62
	LP	0.86	0.93	0.94	0.94	0.94	0.94	0.93
	BVAR-CV	0.73	0.72	0.64	0.66	0.74	0.78	0.77
	H-BVAR	0.88	0.86	0.73	0.72	0.80	0.85	0.86
Compensation	LS	0.86	0.82	0.76	0.75	0.72	0.68	0.67
	RIDGE	0.93	0.82	0.75	0.72	0.66	0.60	0.58
	RIDGE-AS	0.91	0.71	0.69	0.67	0.61	0.56	0.55
	LP	0.86	0.93	0.95	0.95	0.94	0.92	0.90
	BVAR-CV	0.78	0.71	0.69	0.71	0.69	0.67	0.69
	H-BVAR	0.85	0.83	0.80	0.81	0.82	0.82	0.83
Fed Funds Rate	LS	0.85	0.83	0.80	0.78	0.77	0.79	0.80
	RIDGE	0.79	0.77	0.74	0.68	0.68	0.72	0.72
	RIDGE-AS	0.78	0.66	0.68	0.64	0.64	0.68	0.69
	LP	0.85	0.94	0.96	0.96	0.95	0.94	0.93
	BVAR-CV	0.76	0.72	0.76	0.77	0.77	0.81	0.83
	H-BVAR	0.87	0.86	0.74	0.73	0.78	0.84	0.87

Table 10: Impulse Response Inference – Setup B: CI Length (rescaled $\times 100$)

Variable	Method	$h = 1$	$h = 4$	$h = 8$	$h = 12$	$h = 16$	$h = 20$	$h = 24$
Real GDP	LS	0.71	1.56	2.07	2.31	2.32	2.24	2.15
	RIDGE	0.79	1.56	1.85	1.95	1.92	1.85	1.77
	RIDGE-AS	0.74	1.31	1.65	1.76	1.75	1.70	1.64
	LP	0.71	2.42	4.21	5.40	5.90	5.91	5.70
	BVAR-CV	0.53	1.23	1.74	2.00	2.10	2.13	2.15
	H-BVAR	0.58	1.36	1.87	2.16	2.32	2.44	2.55
GDP Deflator	LS	0.26	0.74	1.47	2.14	2.69	3.12	3.46
	RIDGE	0.32	0.81	1.41	1.96	2.43	2.82	3.13
	RIDGE-AS	0.31	0.71	1.28	1.79	2.23	2.59	2.88
	LP	0.26	1.09	2.68	4.39	5.95	7.23	8.13
	BVAR-CV	0.21	0.61	1.26	1.88	2.43	2.91	3.32
	H-BVAR	0.23	0.71	1.42	2.09	2.71	3.27	3.80
Consumption	LS	0.63	1.35	1.97	2.29	2.34	2.30	2.24
	RIDGE	0.71	1.35	1.74	1.93	1.96	1.93	1.88
	RIDGE-AS	0.67	1.15	1.57	1.75	1.79	1.77	1.74
	LP	0.63	2.04	3.92	5.27	5.89	6.00	5.88
	BVAR-CV	0.49	1.09	1.65	1.97	2.12	2.20	2.24
	H-BVAR	0.53	1.21	1.79	2.16	2.41	2.60	2.77
Investment	LS	3.38	6.65	7.89	7.89	7.31	6.69	6.18
	RIDGE	3.79	6.81	6.93	6.46	5.79	5.19	4.73
	RIDGE-AS	3.59	5.57	6.11	5.77	5.21	4.72	4.34
	LP	3.37	10.16	16.00	18.85	19.06	18.22	17.23
	BVAR-CV	2.64	5.26	6.59	6.91	6.78	6.57	6.38
	H-BVAR	2.89	5.74	7.08	7.54	7.63	7.60	7.58
Hours	LS	0.70	1.64	2.27	2.42	2.29	2.11	1.99
	RIDGE	0.82	1.73	2.10	2.06	1.86	1.66	1.52
	RIDGE-AS	0.79	1.49	1.87	1.85	1.68	1.51	1.40
	LP	0.70	2.49	4.52	5.62	5.83	5.57	5.26
	BVAR-CV	0.57	1.30	1.90	2.07	2.03	1.94	1.88
	H-BVAR	0.62	1.49	2.15	2.35	2.38	2.37	2.37
Compensation	LS	0.86	1.17	1.18	1.18	1.21	1.24	1.25
	RIDGE	0.97	1.21	1.07	1.05	1.06	1.06	1.06
	RIDGE-AS	0.93	0.95	0.95	0.96	0.97	0.98	0.98
	LP	0.86	1.80	2.53	2.87	3.11	3.26	3.32
	BVAR-CV	0.69	0.94	1.00	1.05	1.11	1.18	1.23
	H-BVAR	0.78	1.10	1.26	1.40	1.54	1.67	1.78
Fed Funds Rate	LS	0.25	0.39	0.43	0.43	0.41	0.38	0.35
	RIDGE	0.29	0.39	0.37	0.36	0.33	0.30	0.29
	RIDGE-AS	0.27	0.31	0.33	0.32	0.30	0.28	0.27
	LP	0.25	0.59	0.88	1.01	1.05	1.03	0.98
	BVAR-CV	0.21	0.31	0.36	0.37	0.36	0.35	0.34
	H-BVAR	0.23	0.36	0.42	0.44	0.45	0.45	0.46