



Australian
National
University

Crawford School of Public Policy

CAMA

Centre for Applied Macroeconomic Analysis

Will the real eigensystem VAR please stand up? A univariate primer

CAMA Working Paper 1/2019
January 2019

Leo Krippner

Reserve Bank of New Zealand

University of Waikato

Centre for Applied Macroeconomic Analysis, ANU

Abstract

I introduce the essential aspects of the eigensystem vector autoregression (EVAR), which allows VARs to be specified and estimated directly in terms of their eigensystem, using univariate examples for clarity. The EVAR guarantees non-explosive dynamics and, if included, non-redundant moving-average components. In the empirical application, constraining the EVAR eigenvalues to be real and positive leads to “desirable” impulse response functions and improved out-of-sample forecasts.

Keywords

vector autoregression, moving average, lag polynomial

JEL Classification

C22, C32, C53

Address for correspondence:

(E) cama.admin@anu.edu.au

ISSN 2206-0332

[The Centre for Applied Macroeconomic Analysis](#) in the Crawford School of Public Policy has been established to build strong links between professional macroeconomists. It provides a forum for quality macroeconomic research and discussion of policy issues between academia, government and the private sector.

The Crawford School of Public Policy is the Australian National University's public policy school, serving and influencing Australia, Asia and the Pacific through advanced policy research, graduate and executive education, and policy impact.

Will the real eigensystem VAR please stand up? A univariate primer

Leo Krippner*

16 January 2019

Abstract

I introduce the essential aspects of the eigensystem vector autoregression (EVAR), which allows VARs to be specified and estimated directly in terms of their eigensystem, using univariate examples for clarity. The EVAR guarantees non-explosive dynamics and, if included, non-redundant moving-average components. In the empirical application, constraining the EVAR eigenvalues to be real and positive leads to “desirable” impulse response functions and improved out-of-sample forecasts.

JEL classification: C22, C32, C53

Keywords: vector autoregression; moving average, lag polynomial

1 Introduction

In this note I introduce the essential aspects of the eigensystem vector autoregression (EVAR), where VARs are specified and estimated directly in terms of their eigensystem. Relative to the typical ordinary least squares (OLS) estimation of VARs, I show that EVARs have three distinct advantages: (1) stable (i.e. non-explosive) dynamics can be enforced by constraining the eigenvector magnitudes; (2) the eigenvalues may be additionally constrained to be real and positive, which produces “desirable” dynamics; and (3) an EVAR may be extended to include moving-average (MA) components that are guaranteed to avoid redundant lag polynomial factors in the resulting VARMA.

I use univariate models in this paper to make the points above most clearly.¹ Hence, in section 2, I use the AR(2) exposition to illustrate the nature of autoregressive dynamics, and also to provide the basis for generalizing to higher-order models. Section 3 develops the generalization. Section 4 contains the results of an empirical application to forecasting 3-month United States Treasury bill data. Section 5 concludes with the main points and potential extensions.

*Reserve Bank of New Zealand, University of Waikato, and Centre for Applied Macroeconomic Analysis. Email: leo.krippner@rbnz.govt.nz. I thank Iris Claus, James Morley, Christie Smith, and Benjamin Wong for useful comments on this note and related work.

¹A univariate model is, of course, simply the special case of a VAR with one variable. Hence, the names and abbreviations introduced in this note apply in the univariate and multivariate contexts.

2 The AR(2) model and VAR dynamics

The AR(2) model is typically represented in OLS regression form as follows:

$$y_t = \alpha + \phi_1 y_{t-1} + \phi_2 y_{t-2} + \varepsilon_t \quad (1)$$

where α is the constant parameter, and ϕ_1 and ϕ_2 are the AR parameters. An alternative representation is the lag polynomial form:

$$(1 - \phi_1 L - \phi_2 L^2)(y_t - \mu) = \varepsilon_t \quad (2)$$

where L is the lag (or backshift) operator so that $Lx_t = x_{t-1}$, and $\mu = \alpha / (1 - \phi_1 - \phi_2)$ is the series mean (if the model is stationary). Equation 2 may be factored as:

$$(1 - \lambda_1 L)(1 - \lambda_2 L)(y_t - \mu) = \varepsilon_t \quad (3)$$

where:

$$(\lambda_1, \lambda_2) = \frac{\phi_1 \pm \sqrt{\phi_1^2 + 4\phi_2}}{2} \quad (4)$$

are the two eigenvalues for the AR(2) model. For reference later in section 3, note that:

$$(\phi_1, \phi_2) = (\lambda_1 + \lambda_2, -\lambda_1 \lambda_2) \quad (5)$$

The first thing a researcher will hope for in an estimated VAR is a stable model, which is typically tested by calculating the model's eigenvalues to ensure they have an absolute value of less than 1. For the AR(2) model, both eigenvalues need to be inside the complex unit circle, as illustrated in the top-right subplot of figure 1. The equivalent is for the (ϕ_1, ϕ_2) point to be within the stability triangle plotted in the top-left subplot of figure 1.²

Depending on the particular application, the second thing a researcher may hope for is “desirable” impulse response functions (IRFs). While this criteria is subjective, smooth and either monotonic or “hump-shaped” decays, like those in the bottom-right subplot of figure 1, are generally preferred.³ Such IRFs are guaranteed when both eigenvalues (λ_1, λ_2) are real and positive, which corresponds to (ϕ_1, ϕ_2) being within the shaded region of the stability triangle. For example, the two eigenvalues indicated with a “+” (i.e. 0.8 and 0.6) produce $(\phi_1, \phi_2) = (1.40, -0.48)$. A pure monotonic decay is the $(\phi_1, 0)$ special case, which is an AR(1) model with the eigenvalue ϕ_1 . The remaining regions of the stability triangle result in oscillatory IRFs, which are associated with complex or negative real eigenvalues.⁴ The “+” cloud in figure 1 illustrates that AR(2) models estimated from simulated data (20 samples each of 30 data points) using OLS often produce oscillatory results even when the true data-generating process is the non-oscillatory 0.8 and 0.6 eigenvalue example.

²See Hamilton (1994) pp. 7–18 and Sargent (1987) pp. 183–191 for related discussion, and on the aspects discussed earlier in this section.

³For example, such IRFs feature in the Smets and Wouters (2007) model, and Barnichon and Brownlees (2018) impose those properties empirically using B-spline functions. The IRFs from empirical VARs, e.g. Stock and Watson (2001) for the United States, are typically hump-shaped with some superimposed oscillatory dynamics.

⁴The region in the stability triangle below the parabola $\phi_2 = -\frac{1}{4}\phi_1^2$ corresponds to a complex conjugate pair of eigenvalues, the region above $\phi_2 = 0$ corresponds to one positive and one negative eigenvalue, and the mirror of the shaded region in the stability triangle corresponds to two real negative eigenvalues.

The third aspect a researcher should hope for is that the model provides “good” forecasts. Unlike IRFs, this criterion may be tested objectively using a quantitative assessment, like the application in section 4.⁵

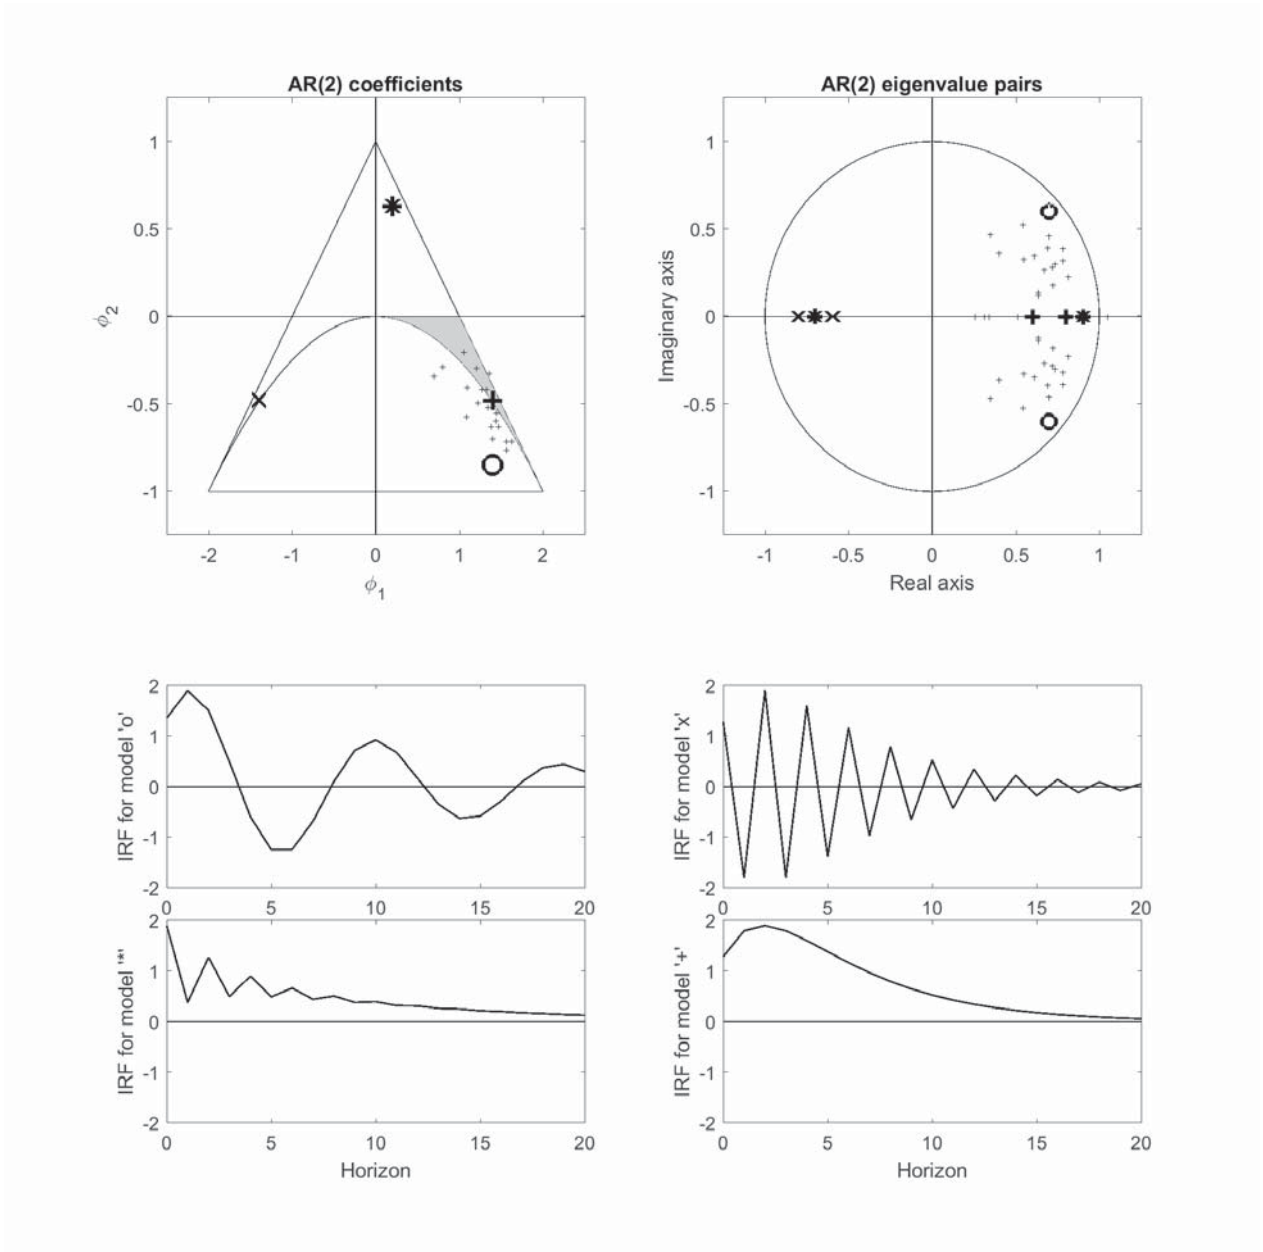


Figure 1: Four AR(2) model examples, denoted ‘+’, ‘o’, ‘x’, ‘*’, and their IRFs. The top-left figure plots the AR(2) coefficients as the points (ϕ_1, ϕ_2) , the top-right figure plots the pairs of eigenvalues associated with each model, and the remaining four figures plot the IRFs from each model (all standardized to a maximum of 1.9 units).

⁵As a separate issue, institutional forecasts appear to have a qualitative dimension analogous to IRFs, in that monotonic decays or hump-shaped dynamics are typically published.

3 EVAR specification and estimation

Generalizing the AR(2) representation in equation 3, any univariate VAR with P lags may be represented by P eigenvalues as follows:

$$\left[\prod_{k=1}^P (1 - \lambda_k L) \right] (y_t - \mu) = \varepsilon_t \quad (6)$$

where $k = 1, \dots, P$ are the index numbers for the eigenvalues. Analogous to equation 2, equation 6 multiples out to the following lag polynomial:

$$\left[1 - \sum_{p=1}^P \phi_p L^p \right] (y_t - \mu) = \varepsilon_t \quad (7)$$

where ϕ_p is the coefficient associated with the lag exponent L^p . Applying the operators L^p to $(y_t - \mu)$ thereby produces:

$$y_t = \mu + \left[\sum_{p=1}^P \phi_p (y_{t-1} - \mu) \right] + \varepsilon_t \quad (8)$$

which is analogous to equation 1.

Maximum likelihood estimation (MLE) of the EVAR is easiest when conditioned on the initial P observations of the $1 \times T$ dataset $\{y\}_1^T$, as for the conditional MLE that underlies the OLS estimation of VARs.⁶ The MLE of the EVAR is therefore the following non-linear constrained minimization:

$$\text{minimize } (\varepsilon\varepsilon') \text{ subject to } |\lambda_k| < 1 \quad (9)$$

where $\varepsilon = \varepsilon \left([\lambda_1, \dots, \lambda_P, \mu], \{y\}_1^T \right)$ is the $1 \times T$ vector of residuals associated with the parameters $[\lambda_1, \dots, \lambda_P, \mu]$ and $\{y\}_1^T$, and λ_k can be real or complex. The EVAR with real (and positive) eigenvalues, which I hereafter abbreviate to REVAR, is:

$$\text{minimize } (\varepsilon\varepsilon') \text{ subject to } 0 < \lambda_k < 1 \quad (10)$$

It is convenient to transform the constrained estimation to an unconstrained estimation by expressing the eigenvalues λ_k in terms of unconstrained parameters x_k . This is straightforward for the REVAR, by simply using the logistic function as specified in appendix A. For the EVAR, (x_k, x_{k+1}) pairs are required so the eigenvalues can be positive, negative, or complex conjugate pairs while ensuring $|\lambda_k| < 1$. I achieve this via the AR(2) stability triangle, as detailed in appendix A.

The eigenvalues resulting from the unconstrained parameters $[x_1, \dots, x_P]$ are used to create the lag polynomial equation in equation 6, which is undertaken in practice using vector convolution.⁷ For example, with the AR(2) model in section 2, $(1 - \lambda_1 L)$ is

⁶I have also undertaken exact MLE using the Kalman filter. The results are immaterially different, but the estimations take longer.

⁷If u and v are vectors of polynomial coefficients, the convolution of u and v is equivalent to multiplying the two polynomials. The function is “conv(u, v)” in MatLab, and appendix B provides the algorithm and sample code applicable to other programming languages.

represented as the two-element vector $u = [1, -\lambda_1]$ and $(1 - \lambda_2 L)$ as $v = [1, -\lambda_2]$. The convolution of those two vectors is:

$$\begin{aligned} w_2 &= \text{conv}(u, v) \\ &= \text{conv}([1, -\lambda_1], [1, -\lambda_2]) \\ &= [1, -(\lambda_1 + \lambda_2), \lambda_1 \lambda_2] \end{aligned} \quad (11)$$

where w_2 is a three-element vector representing the lag polynomial for the AR(2) model.⁸ Generalizing to P eigenvalues, iterating on $w_3 = \text{conv}(w_2, [1, -\lambda_3])$, etc., produces the convolution of all $[1, -\lambda_k]$ vectors. The resulting vector w_P will contain $1 + P$ elements, with the first element being 1. Negating the remaining P elements obtains the coefficients ϕ_p for equation 8. Together with μ , which is an unconstrained parameter to be estimated (or it can be fixed as discussed in section 4), the residual vector ε may be calculated from the data $\{y\}_1^T$ for the current set of parameters $[x_1, \dots, x_P, \mu]$.

In summary, the estimation is now simply the unconstrained non-linear minimization:

$$\text{minimize } \left[\varepsilon \left([x_1, \dots, x_P, \mu], \{y\}_1^T \right) \varepsilon \left([x_1, \dots, x_P, \mu], \{y\}_1^T \right)' \right] \quad (12)$$

which is estimated using standard algorithms (e.g. I use the MatLab function “lsqnonlin”). I have found fast convergence with arbitrary starting values for $[x_1, \dots, x_P, \mu]$, but appendix A discusses how better starting values may be obtained.

There are two ready extensions to the EVAR and REVAR. First, it is straightforward to allow for unit root dynamics, as with standard OLS VARs, by setting $\lambda_k = 1$ for $k = 1, \dots, D$, where D is the imposed order of integration. For example, in the case of a single unit root, setting $\lambda_1 = 1$ gives $(1 - L)(y_t - \mu) = (y_t - \mu) - (y_{t-1} - \mu) = \Delta y_t$, and therefore:

$$\left[\prod_{k=2}^P (1 - \lambda_k L) \right] \Delta y_t = \varepsilon_t \quad (13)$$

Hence, the OLS VAR, EVAR, and REVAR estimations would simply be undertaken on the first difference of y_t , i.e. Δy_t , with the constant α or mean μ parameters set to zero.

The second, and I believe more important, extension is to use a constrained EVAR to avoid the issue of parameter redundancy that arises in VARMA models from the cancellation (or near-cancellation in practice) of factors in the VAR and MA lag polynomials.⁹ For example, a given ARMA(P, Q) may be written as:

$$\left[\prod_{k=1}^P (1 - \lambda_k L) \right] (y_t - \mu) = \left[\prod_{j=1}^Q (1 - \gamma_j L) \right] \varepsilon_t \quad (14)$$

where γ_j are the eigenvalues from the MA lag polynomial. Any instance of $\lambda_k = \gamma_j$ (or $\lambda_k \simeq \gamma_j$ in practice) means that an equivalent ARMA($P - 1, Q - 1$) may be obtained by cancelling the $(1 - \lambda_k L)$ factor from both sides of equation 14. But specifying the eigenvalues of an EVAR to come from a particular region, and the eigenvalues of the MA component to come from a non-overlapping region will guarantee that no factor

⁸This result may be verified using the coefficients $1, -(\lambda_1 + \lambda_2) = -\phi_1$, and $\lambda_1 \lambda_2 = -\phi_2$ that respectively apply to $L^0 = 1$ (i.e. no lag), $L^1 = L$, and L^2 from equations 5 and 2.

⁹See, for example, Lütkepohl (2006) section 12 for further discussion on this issue and its implications.

cancellations in the ARMA can occur. At the same time, the MA eigenvalues may be constrained to ensure that the MA component is invertible.

Regarding the lag lengths P and Q , any methods typically used for the OLS estimation of VARs could be applied (e.g. imposing 12 lags for monthly data, or 4 for quarterly data). In section 4, I select them using information criteria.

4 Empirical application

In this section, I compare the EVAR and the REVAR models to standard OLS VAR models using an out-of-sample forecasting exercise using a cumulating sample. Specifically, I forecast the monthly 3-month United States Treasury bill rate, which is convenient because it is not subject to revision (hence avoiding real-time considerations that would otherwise need to be accounted for), and it also presents non-trivial and challenging dynamics to model in real time (i.e. inspection of the data plotted below in figure 2 indicates that it could be represented as a persistent but stationary process or a unit root process).¹⁰ Using a cumulating sample avoids any possibility of coincidentally selecting a rolling window favorable to any particular model, and it makes full use of all information from past data.

The data are obtained from the Federal Reserve Economic Database (FRED) on the St. Louis Federal Reserve Bank website. The sample period is July 1948 to October 2008, which represents the longest period over which the data could move freely (the rate was fixed at 0.38 percent for the five years prior to July 1948, and it was constrained by the lower bound from November 2008 to December 2015).

The forecasting models I have tested are combinations of the following three aspects: (1) OLS, EVAR, or REVAR estimation; (2) lag selection using the Akaike or Schwarz Information Criterion (AIC or SIC, respectively); and (3) using y_t in levels or first differences. As a proof of concept, I have also estimated a single VARMA model that uses a REVAR specification in conjunction with a MA process specified to have real and “Negative” eigenvalues (the REVARMAN), and I also select the P and Q using the AIC and SIC, with $\min(P) = \min(Q) = 1$. I use the random walk (RW) model as a benchmark. For all models estimated in levels, I set the mean parameter μ equal to the mean of the current sample used for estimation.¹¹

The out-of-sample-forecasting exercise begins in June 1967 using the sample of 240 data points (20 years). For each model, the lag lengths are selected using the AIC and SIC, the model is estimated, and the estimated model is then used to obtain forecasts for horizons of 1 to 60 months (5 years). I then expand the sample to 241 data points by including the July 1967 data point, and repeat the lag selection, estimation, and forecasts. That process continues until the final forecasts are made using the 738 data points for the full sample.

¹⁰The Dickey-Fuller and Phillips-Perron tests do not reject a unit root, and the Kwiatkowski-Phillips-Schmidt-Shin test rejects stationarity.

¹¹The forecast results from models when μ is estimated are very similar, but for all models (particularly the OLS estimations) there are many occasions where the estimated μ has extreme values, i.e. it is well outside the range of the sample. The models with Δy_t have $\mu = 0$.





Figure 2: The data used for the cumulative out-of-sample forecast exercise (top), and examples of the forecasts from models estimated using the data from July 1948 to October 1980 (O = OLS VAR, E = EVAR, R = REVAR, and M = REVARMAN).

The bottom subplots of figure 2 provide examples of the out-of-sample forecasts for five years using the models estimated with data up to October 1980. I have selected this set of forecasts because they provide the clearest empirical illustrations of the phenomena raised in the introduction and discussed in section 2 for the AR(2) model. That is, when the model is estimated with y_t in levels, the OLS forecasts are both explosive and markedly oscillating. The EVAR estimation results in a stable model, consistent with all eigenvalues being constrained to have absolute values of less than 1, but the oscillations remain because the eigenvalues can still be complex (or real and negative). The REVAR and REVARMAN estimations produce forecasts with a minor hump shape, i.e. a higher forecast for the first horizon and then a monotonic reversion to the mean. For the models

estimated using Δy_t , the OLS and EVAR estimations produce identical forecasts that still have marked oscillations as they revert to their implicit steady-state level as at October 1980. The REVAR forecasts contain a small monotonic increase before settling at the implied steady-state level, while the RW forecast simply maintains the October 1980 level.

I evaluate the forecasts from each model by subtracting the forecasts from the associated future realized values, and then using the series of forecast errors for each horizon to calculate the mean absolute forecast error (MAFE) associated with that horizon. To facilitate comparisons across the different models, table 1 expresses the MAFEs for each model and horizon as a percentage relative to the MAFEs for the RW model, i.e.:

$$\text{Table Entry (Model } j, \text{ Horizon } h) = \frac{\text{MAFE}(j, h)}{\text{MAFE}(\text{RW}, h)} - 1 \quad (15)$$

Hence, a positive (negative) entry shows a higher (lower) MAFE than the RW model, and the actual RW MAFEs are displayed in the end column. I use a one-sided Diebold-Mariano test to obtain the statistical significance of $\text{MAFE}(j, h) - \text{MAFE}(\text{RW}, h)$.

Table 1: MAFEs relative to RW for models with SIC-selected lags

Model	y_t models					Δy_t models			RW
	O	O ^E	E	R	M	O	E	R	% pts
$h = 1$	-2.0	-2.0	-2.1	-4.7 ^v	-5.2 ^v	-2.1	-2.1	-6.1 ⁱ	0.29
2	3.1	3.2	2.8	-1.3	-1.6	2.7	2.7	-3.2 ^v	0.47
3	7.6	8.0	7.5	2.0	1.1	7.1	7.1	-1.0	0.62
6	9.4	8.8	8.5	2.5	1.2	7.8	7.8	-1.0	0.97
9	5.4	4.8	4.1	-1.7	-2.0	3.5	3.5	-1.7 ^v	1.23
12	2.9	1.0	-0.1	-3.9	-3.9	0.5	0.5	-0.9	1.50
18	2.1	-1.1	-1.9	-7.9 ^x	-7.1 ^x	-0.4	-0.4	-0.5	1.98
24	3.6	-3.3	-4.0	-11.2 ^v	-9.9 ^v	-0.5	-0.5	0.1	2.30
36	6.9	-4.5 ^x	-5.2	-14.8 ^v	-13.8 ^v	0.3	0.3	0.6	2.62
48	9.0	-2.9	-2.8	-14.5	-15.3 ^x	-0.5	-0.5	1.2	2.58
60	11.8	6.6	7.2	-2.3	-5.1	-0.0	-0.0	0.9	2.44

Notes: O = OLS VAR, E = EVAR, R = REVAR, M = REVARMAN, and O^E = O with E override if O produces explosive dynamics. Superscripts “x” = 10%, “v” = 5%, and “i” = 1% indicate statistical significance.

The main point from table 1 is that the REVAR and REVARMAN models applied to y_t produce better forecasts than the competing y_t and Δy_t models. Indeed, the REVAR and REVARMAN forecasts are better than the RW for all horizons, except the 2 to 10 or 9 month range, and the results for $h = 1$ and $17 \leq h \leq 47$ or 49 are statistically significant.

Other notable points from table 1 are: (1) the OLS VAR (O) produces the highest MAFEs, consistent with about 27 percent of the estimated models being explosive; (2) using the EVAR as the override in those cases (O^E) improves the longer-horizon forecasts, consistent with the EVARs always being stable; (3) the EVAR forecast results (E) are very similar to O^E, consistent with stable OLS VARs being close or identical to the EVAR.

Appendix C contains additional results, from using AIC lag selection, and forecasts assessments using root mean squared forecast errors (RMSFEs). None of those results overturn the comments already made, except the statistical significance varies. In addition, forecast comparisons over the so-called Great Moderation period (from 1985) show that the REVAR and REVARMAN models outperform the RW at all horizons tested, and the lower MAFEs and RMSFEs are all statistically significant.

5 Conclusion

The eigensystem VAR (EVAR) introduced in this note allows VARs to be specified and estimated directly in terms of their eigensystem. Hence, rather than having to accept the dynamics implied by an OLS estimation, which are sometimes explosive and often oscillatory, one can impose the dynamics desired. For example, imposing positive real eigenvalues to create the REVAR produces models with smooth IRFs that can be monotonic or hump-shaped, and it also improves the forecasts in the empirical application to United States 3-month Treasury bill data.

There are several potential extensions to the univariate illustrations and applications in this note. The most straightforward is the application to other univariate series, including polynomial time trends and/or unit roots with drift as might be required, and then to multivariate time series. The early results from such applications are encouraging, in particular generally favoring the REVAR as in this note, even against Bayesian VAR models (which is a standard method of constraining/influencing VAR estimations).

In summary, the EVAR and REVAR, plus the inclusion of MA components as illustrated in this note, offer interesting extensions to standard VAR applications, and further investigation is warranted.

References

- Barnichon, R. and C. Brownlees (2018). Impulse response estimation by smooth local projections. *The Review of Economics and Statistics* (forthcoming).
- Hamilton, J. (1994). *Time Series Analysis*. Princeton University Press.
- Krippner, L. (2019). New tricks for old VARs: An eigensystem approach. *Manuscript*.
- Lütkepohl, H. (2006). *New Introduction to Multiple Time Series*. Springer.
- Morley, J. (1999). A note on constraining AR(2) parameters in estimation. *Working Paper*.
- Sargent, T. (1987). *Macroeconomic Theory, Second Edition*. Academic Press.
- Smets, F. and R. Wouters (2007). Shocks and frictions in US business cycles: a Bayesian DSGE approach. *American Economic Review* 97(3), 586–606.
- Stock, J. and M. Watson (2001). Vector autoregressions. *Journal of Economic Perspectives* 15(4), 101–115.

A Unconstrained estimation set-up

Section A.1 details the transformations that produce constrained eigenvalues from unconstrained parameters, therefore allowing the EVAR and REVAR to be estimated as an unconstrained non-linear minimization. Section A.2 details how suitable starting values may be obtained for EVAR and REVAR estimation.

A.1 Producing constrained eigenvalues

Using an unconstrained value x_k to produce $0 < \lambda_k < 1$ for the REVAR is straightforward, by simply using the logistic function:

$$\lambda_k = \frac{1}{1 + \exp(-x_k)} \quad (16)$$

For the EVAR, eigenvalue pairs $(\lambda_k, \lambda_{k+1})$ are required to allow for real or complex values, at the same time as ensuring $|\lambda_k| < 1$. Such eigenvalue pairs are obtainable indirectly using a geometric method to choose AR(2) coefficient pairs (ϕ_k, ϕ_{k+1}) within the stability triangle from figure 1.¹² Hence, I use the scaled shifted logistic function of the unconstrained value x_k :

$$\phi_k = \frac{4}{1 + \exp(-x_k)} - 2 \quad (17)$$

to first obtain a value for ϕ_k such that $-2 < \phi_k < 2$ (which is the range of allowable values for the base of the stability triangle). Given ϕ_k , the upper constraint for ϕ_{k+1} , which I denote $\bar{\phi}_{k+1}$, to ensure that (ϕ_k, ϕ_{k+1}) falls within the stability triangle is:

$$\bar{\phi}_{k+1} = \begin{cases} 1 + \phi_k & ; \text{if } \phi_k < 0 \\ 1 - \phi_k & ; \text{if } \phi_k \geq 0 \end{cases} \quad (18)$$

and the lower constraint for ϕ_{k+1} is -1 . Hence, given an unconstrained value x_{k+1} , the following scaled shifted logistic function:

$$\phi_{k+1} = \frac{\bar{\phi}_{k+1} + 1}{1 + \exp(-x_{k+1})} - 1 \quad (19)$$

will ensure a value of ϕ_{k+1} so that (ϕ_k, ϕ_{k+1}) falls within the stability triangle. Equation 4 is then used to obtain $(\lambda_k, \lambda_{k+1})$ from (ϕ_k, ϕ_{k+1}) .

Note that if there is an odd number of lags (and hence eigenvalues) for the EVAR, then the single remaining unpaired eigenvalue must be real. However, it could be negative or positive so long as $-1 < \lambda_k < 1$. This result is readily obtained from an unconstrained value x_k within the following shifted scaled logistic function:

$$\lambda_k = \frac{2}{1 + \exp(-x_k)} - 1 \quad (20)$$

¹²Morley (1999) provides an alternative method based on the algebraic eigenvalues of the AR(2) lag polynomial and a constraining device similar to the logistic function.

A.2 EVAR and REVAR starting values

For both the EVAR and the REVAR, if μ is to be estimated then a suitable starting value is $\mu_0 = \text{mean}(y)$.

For the EVAR, one could use arbitrary starting values for $[x_1, \dots, x_P]$. However, very good starting values, which make for convergence within a few iterations, may be obtained from the OLS VAR estimates as follows: (1) use the ϕ_p estimates in the AR(P) companion matrix to calculate the eigenvalues λ_k ; (2) if any eigenvalues $|\lambda_k| > 1$, scale them to $|\lambda_k| < 1$ (I scale them to $|\lambda_k| = 0.99$); (3) convert $(\lambda_k, \lambda_{k+1})$ pairs to AR(2) (ϕ_k, ϕ_{k+1}) pairs using equation 5; and (4) invert AR(2) (ϕ_k, ϕ_{k+1}) pairs to unconstrained (x_k, x_{k+1}) pairs using equations 17 and 19.

For the REVAR, OLS estimates will not necessarily be available. Hence, I simply use arbitrary starting values $[x_1, \dots, x_P]$ obtained with linear spacing between -2.197 and 2.944 (which are the values that would produce 0.1 and 0.95 from equation 16). If values closely related to the current sample are available, e.g. the REVAR parameters from the previous sample, then they could be used as starting values that would converge more quickly. However, for the empirical exercise in this paper, I use the arbitrary starting values noted earlier to ensure that all REVAR estimations are fully independent from each other.

B Vector convolution code

The MatLab function “conv(u, v)” is the simplest way within MatLab to undertake the vector convolution evaluations noted in the main text. Otherwise, as detailed below, the algorithm is a straightforward double summation that may readily be coded in other programming languages.

MatLab documentation for the function “conv” provides the following expression for the convolution of an $m \times 1$ vector u and an $n \times 1$ vector v :

$$w(k) = \sum_j u(j) v(k - j + 1) \quad (21)$$

where $w(k)$ is element k of the $(m + n - 1) \times 1$ vector w . The sum is over all the values of j that lead to legal subscripts for $u(j)$ and $v(k - j + 1)$, which is a minimum index of $j = \max(1, k + 1 - n)$ and a maximum index of $j = \min(k, m)$.

The full expression for w is therefore:

$$w = \sum_{k=1}^{m+n-1} \sum_{j=\max(1, k+1-n)}^{\min(k, m)} u(j) v(k - j + 1) \quad (22)$$

and I have written MatLab code for this algorithm in table B.1. Note, however, that the MatLab function and the code in table B.1 is only for vectors, which applies to the univariate EVARs as outlined in the main text. Krippner (2019) generalizes the convolution technique and also shows how the VAR companion form may be used for the application to multivariate EVARs.

Table B.1: MatLab vector convolution code

```
function w=VectorConvolution(u,v)
% Matrix polynomial convolution.
% u is an m-element vector.
% v is an n-element vector.
[~,m]=size(u);
[~,n]=size(v);
w=zeros(1,m+n-1);
for k=1:m+n-1
    for j=max(1,k+1-n):1:min(k,m)
        w(:,k)=w(:,k)+u(:,j)*v(:,k-j+1);
    end
end
end
```

C All forecast comparison results

This appendix contains the forecast comparison results for all the models mentioned in the main text. To make the set of results self-contained, I have first repeated table 1 from section 4 of the main text, which summarizes the MAFE results for the full sample from the models using the SIC to select the lags. Table B.2 uses the AIC to select the lags. Tables B.3 and B.4 contain the RMSFE results for the full sample, using SIC and AIC respectively. Note that the insignificance of the full-sample RMSFE comparisons is due to all models having some very large forecast errors in the middle of the sample, but the relative improvements are otherwise all similar to those using MAFEs.

Tables B.5 to B.8 contain the results as in tables B.1 to B.4, but for the Great Moderation period (i.e. from 1985 to 2008). I am happy to provide any further results on request.

Table B.1: Full-sample MAFEs relative to RW (SIC-selected lags)

Model	y_t models					Δy_t models			RW
	O	O ^E	E	R	M	O	E	R	% pts
$h = 1$	-2.0	-2.0	-2.1	-4.7 ^v	-5.2 ^v	-2.1	-2.1	-6.1 ⁱ	0.29
2	3.1	3.2	2.8	-1.3	-1.6	2.7	2.7	-3.2 ^v	0.47
3	7.6	8.0	7.5	2.0	1.1	7.1	7.1	-1.0	0.62
6	9.4	8.8	8.5	2.5	1.2	7.8	7.8	-1.0	0.97
9	5.4	4.8	4.1	-1.7	-2.0	3.5	3.5	-1.7 ^v	1.23
12	2.9	1.0	-0.1	-3.9	-3.9	0.5	0.5	-0.9	1.50
18	2.1	-1.1	-1.9	-7.9 ^x	-7.1 ^x	-0.4	-0.4	-0.5	1.98
24	3.6	-3.3	-4.0	-11.2 ^v	-9.9 ^v	-0.5	-0.5	0.1	2.30
36	6.9	-4.5 ^x	-5.2	-14.8 ^v	-13.8 ^v	0.3	0.3	0.6	2.62
48	9.0	-2.9	-2.8	-14.5	-15.3 ^x	-0.5	-0.5	1.2	2.58
60	11.8	6.6	7.2	-2.3	-5.1	-0.0	-0.0	0.9	2.44

Notes: O = OLS VAR, E = EVAR, R = REVAR, M = REVARMAN, and O^E = O with E override if O produces explosive dynamics. Superscripts “x” = 10%, “v” = 5%, and “i” = 1% indicate statistical significance.

Table B.2: Full-sample MAFEs relative to RW (AIC-selected lags)

Model	y_t models					Δy_t models			RW
	O	O ^E	E	R	M	O	E	R	% pts
$h = 1$	3.5	3.2	3.5	-4.6 ^v	-5.0 ^v	3.3	3.3	-6.3 ⁱ	0.29
2	9.2	8.8	9.1	-1.3	-1.1	8.8	8.8	-3.3 ^v	0.47
3	11.1	11.3	12.7	2.0	1.7	11.3	11.3	-1.1	0.62
6	9.1	8.3	9.7	2.3	1.7	8.2	8.2	-1.0	0.97
9	5.4	4.2	4.9	-1.7	-1.4	4.1	4.1	-1.7 ^v	1.23
12	5.5	1.3	1.9	-3.9	-3.5	2.8	2.8	-0.9	1.50
18	3.8	-1.1	0.3	-7.9 ^x	-6.8 ^x	0.8	0.8	-0.5	1.98
24	4.2	-3.7 ^x	-3.0	-11.1 ^v	-9.9 ^v	-0.0	-0.0	0.1	2.30
36	3.4	-8.1 ^v	-7.5	-14.7 ^v	-14.2 ^v	-2.3	-2.3	0.6	2.62
48	6.3	-7.6 ^v	-7.8	-14.3	-14.3 ^x	-1.5	-1.5	1.2	2.58
60	13.5	1.2	0.8	-2.0	-3.1	2.6	2.6	0.8	2.44

Notes: O = OLS VAR, E = EVAR, R = REVAR, M = REVARMAN, and O^E = O with E override if O produces explosive dynamics. Superscripts “x” = 10%, “v” = 5%, and “i” = 1% indicate statistical significance.

Table B.3: Full-sample RMSFEs relative to RW (SIC-selected lags)

Model	y_t models					Δy_t models			RW
	O	O ^E	E	R	M	O	E	R	% pts
$h = 1$	-2.4	-2.7	-2.7	-3.2	-4.6	-2.9	-2.9	-3.5	0.50
2	3.8	3.3	3.4	2.5	0.1	2.9	2.9	1.3	0.82
3	5.9	5.2	5.3	3.7	0.9	4.5	4.5	2.2	1.02
6	7.3	5.9	6.3	4.2	1.2	4.9	4.9	2.4	1.43
9	5.6	4.1	4.6	1.1	-0.4	2.6	2.6	0.6	1.64
12	3.8	1.4	1.9	-0.1	-1.3	0.1	0.1	1.1	1.95
18	4.8	0.3	0.8	-4.0	-3.9	-0.2	-0.2	0.1	2.44
24	5.1	-1.4	-0.6	-5.3	-5.0	-1.1	-1.1	1.0	2.80
36	9.5	-1.9	-0.3	-6.3	-6.2	-0.1	-0.1	1.0	3.21
48	14.1	-1.7	1.2	-6.3	-6.8	-0.5	-0.5	1.2	3.38
60	22.1	3.3	7.1	-1.9	-3.2	-0.5	-0.5	1.1	3.38

Notes: O = OLS VAR, E = EVAR, R = REVAR, M = REVARMAN, and O^E = O with E override if O produces explosive dynamics. Superscripts “x” = 10%, “v” = 5%, and “i” = 1% indicate statistical significance.

Table B.4: Full-sample RMSFEs relative to RW (AIC-selected lags)

Model	y_t models					Δy_t models			RW
	O	O ^E	E	R	M	O	E	R	% pts
$h = 1$	4.0	3.6	4.1	-3.2	-4.2	3.5	3.5	-3.9	0.50
2	14.0	13.6	14.7	2.4	1.2	13.3	13.3	1.1	0.82
3	14.6	14.0	16.0	3.5	2.0	13.8	13.8	1.9	1.02
6	7.4	5.7	8.4	3.9	1.9	6.1	6.1	2.2	1.43
9	6.2	3.6	5.9	1.0	-0.0	3.6	3.6	0.5	1.64
12	6.9	2.7	5.5	-0.1	-1.0	3.2	3.2	0.9	1.95
18	6.1	1.0	3.6	-3.7	-3.8	1.3	1.3	0.1	2.44
24	5.8	-0.3	1.3	-5.0	-4.8	0.1	0.1	1.0	2.80
36	9.1	-2.3	-0.2	-5.9	-6.2	0.4	0.4	1.0	3.21
48	14.8	-2.6	0.9	-5.8	-6.4	1.3	1.3	1.2	3.38
60	23.8	1.4	6.0	-1.1	-2.0	2.6	2.6	1.1	3.38

Notes: O = OLS VAR, E = EVAR, R = REVAR, M = REVARMAN, and O^E = O with E override if O produces explosive dynamics. Superscripts “x” = 10%, “v” = 5%, and “i” = 1% indicate statistical significance.

Table B.5: 1985–2008 MAFEs relative to RW (SIC-selected lags)

Model	y_t models					Δy_t models			RW
	O	O ^E	E	R	M	O	E	R	% pts
$h = 1$	2.0	2.0	1.0	-9.9 ⁱ	-10.4 ⁱ	1.9	1.9	-10.8 ⁱ	0.16
2	10.7	10.7	8.7	-6.7 ⁱ	-4.1 ^x	10.9	10.9	-9.1 ⁱ	0.27
3	13.7	13.7	10.5	-6.9 ^v	-4.3 ^x	14.1	14.1	-8.4 ⁱ	0.37
6	12.9	12.9	8.5	-3.6	-2.3	13.5	13.5	-4.8 ⁱ	0.65
9	6.1	6.1	1.1	-5.6	-4.8	6.9	6.9	-3.6 ⁱ	0.92
12	3.1	3.1	-2.4	-8.5 ^x	-7.7 ^x	4.1	4.1	-2.6 ⁱ	1.18
18	1.1	1.1	-5.9	-12.9 ^v	-11.7 ^v	2.6	2.6	-1.4 ⁱ	1.56
24	-0.5	-0.5	-9.3 ^x	-16.7 ^v	-15.3 ^v	1.5	1.5	-0.9 ^v	1.85
36	-3.3 ^v	-3.3 ^v	-14.4 ^v	-20.8 ^v	-19.5 ^v	-0.3	-0.3	-0.1	2.15
48	-6.1 ⁱ	-6.1 ⁱ	-18.5 ⁱ	-24.7 ^v	-23.4 ^v	-2.1 ^v	-2.1 ^v	0.6	2.16
60	-6.3 ⁱ	-6.3 ⁱ	-20.9 ⁱ	-27.9 ^v	-27.3 ^v	-1.5 ^x	-1.5 ^x	0.7	2.01

Notes: O = OLS VAR, E = EVAR, R = REVAR, M = REVARMAN, and O^E = O with E override if O produces explosive dynamics. Superscripts “x” = 10%, “v” = 5%, and “i” = 1% indicate statistical significance.

Table B.6: 1985–2008 MAFEs relative to RW (AIC-selected lags)

Model	y_t models					Δy_t models			RW
	O	O ^E	E	R	M	O	E	R	% pts
$h = 1$	4.2	4.2	3.9	-10.1 ⁱ	-10.4 ⁱ	4.4	4.4	-10.8 ⁱ	0.16
2	9.9	9.9	8.2	-6.7 ⁱ	-4.1 ^x	10.5	10.5	-8.6 ⁱ	0.27
3	8.7	8.7	6.1	-6.9 ^v	-4.3 ^x	9.5	9.5	-8.2 ⁱ	0.37
6	6.2	6.2	3.2	-3.6	-2.3	7.1	7.1	-4.4 ⁱ	0.65
9	-0.3	-0.3	-4.5	-5.6	-4.8	0.9	0.9	-3.4 ⁱ	0.92
12	-1.7	-1.7	-7.4 ^x	-8.5 ^x	-7.7 ^x	-0.3	-0.3	-2.5 ⁱ	1.18
18	-1.3	-1.3	-8.6 ^x	-12.9 ^v	-11.7 ^v	0.3	0.3	-1.4 ⁱ	1.56
24	-2.4	-2.4	-11.3 ^v	-16.7 ^v	-15.3 ^v	-0.4	-0.4	-0.9 ^v	1.85
36	-9.8 ⁱ	-9.8 ⁱ	-18.9 ⁱ	-20.7 ^v	-19.5 ^v	-7.2 ⁱ	-7.2 ⁱ	-0.1	2.15
48	-10.8 ⁱ	-10.8 ⁱ	-22.0 ⁱ	-24.7 ^v	-23.4 ^v	-7.0 ⁱ	-7.0 ⁱ	0.6	2.16
60	-7.4 ^v	-7.4 ^v	-23.9 ⁱ	-27.9 ^v	-27.3 ^v	-2.5	-2.5	0.6	2.01

Notes: O = OLS VAR, E = EVAR, R = REVAR, M = REVARMAN, and O^E = O with E override if O produces explosive dynamics. Superscripts “x” = 10%, “v” = 5%, and “i” = 1% indicate statistical significance.

Table B.7: 1985–2008 RMSFEs relative to RW (SIC-selected lags)

Model	y_t models					Δy_t models			RW % pts
	O	O ^E	E	R	M	O	E	R	
$h = 1$	0.5	0.5	-0.1	-10.8 ⁱ	-11.0 ⁱ	0.7	0.7	-10.9 ⁱ	0.22
2	10.0	10.0	8.9	-6.6 ⁱ	-4.9 ⁱ	10.3	10.3	-7.6 ⁱ	0.38
3	10.9	10.9	9.4	-6.6 ⁱ	-4.7 ⁱ	11.3	11.3	-7.7 ⁱ	0.50
6	11.4	11.4	8.7	-5.2 ^v	-4.1 ^x	12.0	12.0	-4.8 ⁱ	0.86
9	4.8	4.8	1.4	-6.1 ^x	-5.3 ^x	5.6	5.6	-3.4 ⁱ	1.20
12	2.2	2.2	-2.2	-7.8 ^v	-7.1 ^v	3.3	3.3	-2.2 ⁱ	1.49
18	0.8	0.8	-5.5 ^x	-11.3 ^v	-10.5 ^v	2.3	2.3	-1.3 ⁱ	1.92
24	-0.9	-0.9	-8.8 ^v	-14.6 ^v	-13.7 ^v	1.1	1.1	-0.7 ^x	2.25
36	-3.2 ^v	-3.2 ^v	-13.8 ^v	-19.6 ^v	-18.6 ^v	-0.4	-0.4	-0.1	2.60
48	-5.3 ⁱ	-5.3 ⁱ	-17.6 ⁱ	-22.3 ^v	-21.6 ^v	-1.9 ⁱ	-1.9 ⁱ	0.6	2.63
60	-5.3 ⁱ	-5.3 ⁱ	-16.6 ^v	-19.9 ^v	-19.9 ^v	-2.1 ^v	-2.1 ^v	0.7	2.41

Notes: O = OLS VAR, E = EVAR, R = REVAR, M = REVARMAN, and O^E = O with E override if O produces explosive dynamics. Superscripts “x” = 10%, “v” = 5%, and “i” = 1% indicate statistical significance.

Table B.8: 1985–2008 RMSFEs relative to RW (AIC-selected lags)

Model	y_t models					Δy_t models			RW % pts
	O	O ^E	E	R	M	O	E	R	
$h = 1$	2.4	2.4	1.8	-10.8 ⁱ	-11.0 ⁱ	2.6	2.6	-11.0 ⁱ	0.22
2	8.1	8.1	7.0	-6.5 ⁱ	-4.9 ⁱ	8.5	8.5	-7.3 ⁱ	0.38
3	5.7	5.7	4.2	-6.6 ⁱ	-4.7 ⁱ	6.2	6.2	-7.5 ⁱ	0.50
6	4.4	4.4	1.7	-5.2 ^v	-4.1 ^x	5.1	5.1	-4.5 ⁱ	0.86
9	-2.1	-2.1	-5.6 ^v	-6.1 ^x	-5.3 ^x	-1.1	-1.1	-3.2 ⁱ	1.20
12	-3.0	-3.0	-7.7 ^v	-7.8 ^v	-7.1 ^v	-1.8	-1.8	-2.1 ⁱ	1.49
18	-2.9	-2.9	-9.5 ^v	-11.3 ^v	-10.5 ^v	-1.3	-1.3	-1.3 ⁱ	1.92
24	-3.9 ^v	-3.9 ^v	-11.8 ^v	-14.6 ^v	-13.7 ^v	-1.9	-1.9	-0.7 ^x	2.25
36	-7.4 ⁱ	-7.4 ⁱ	-17.4 ⁱ	-19.6 ^v	-18.6 ^v	-4.8 ⁱ	-4.8 ⁱ	-0.0	2.60
48	-8.8 ⁱ	-8.8 ⁱ	-20.2 ⁱ	-22.3 ^v	-21.6 ^v	-5.7 ⁱ	-5.7 ⁱ	0.6	2.63
60	-6.5 ^v	-6.5 ^v	-18.0 ^v	-19.9 ^v	-19.9 ^v	-3.4 ^v	-3.4 ^v	0.7	2.41

Notes: O = OLS VAR, E = EVAR, R = REVAR, M = REVARMAN, and O^E = O with E override if O produces explosive dynamics. Superscripts “x” = 10%, “v” = 5%, and “i” = 1% indicate statistical significance.