



Tips and tricks for Bayesian VAR models in `gretl`

Luca Pedini¹ 

Received: 31 October 2023 / Accepted: 23 March 2024 / Published online: 20 April 2024
© The Author(s) 2024

Abstract

Bayesian Vector Autoregressive models have become the natural response to the dense parametrization often required by multivariate time series modeling. However, the Bayesian approach is somehow new to the `gretl` ecosystem: classical analysis of Vector Autoregressions (VARs) is natively supported and supplemented via add-ons and function packages, but a Bayesian counterpart in the form of equally general and advanced functions or, at a lower level, in the form of didactic example scripts is missing. This paper pursues the second route describing, via a replication exercise, how to perform basic Bayesian inference in `gretl` using VARs, with particular reference to structural analysis. The contribution goes in the direction of providing new hints and tools, available to the `gretl` user, for a more complete and up-to-date understanding of modern macroeconometrics.

Keywords `gretl` · Bayesian inference · VAR model · Structural analysis

1 Introduction

The contribution of Vector Autoregressive (VAR) models to modern macroeconometrics can hardly be overstated: starting from the cutting-edge article by Sims (1980), VARs have become an unrivaled tool for macroeconomic analysis (Christiano 2002; Del Negro et al. 2007; Giacomini 2013) with special attention devoted to the monetary policy (Bernanke et al. 1997; Bernanke and Mihov 1998; Sims and Zha 2006a, b).

However the classical or frequentist approach, despite its simplicity in usage, poses some threats to the practitioner in terms of over-parametrization: the moderate size of macroeconomic datasets, coupled with the abundance of series to consider (including the lagged ones), requires some kind of regularization which is naturally addressed within the Bayesian framework via prior distributions (Koop and Korobilis 2010). The celebrated Minnesota prior (Litterman 1986) is the most notorious example, but

✉ Luca Pedini
l.pedini@staff.univpm.it

¹ Department of Economics and Social Sciences, Università Politecnica delle Marche, Piazzale R. Martelli, 8, 60121 Ancona, Marche, Italy

is also the peak of the iceberg: through time, Bayesian analysis of VARs has grown and improved in the direction of *i*) new priors which could cater to peculiar necessities (e.g., lag selection as in George et al. (2008); Korobilis (2013) or hyperparameter tuning, Giannone et al. (2015)); *ii*) forecasting (Karlsson 2013); *iii*) structural analysis (Sims and Zha 1998; Waggoner and Zha 2003; Rubio-Ramirez et al. 2010; Baumeister and Hamilton 2015); *iv*) embodying the time-varying nature of parameters (Cogley and Sargent 2001, 2005; Primiceri 2005).

As a consequence, having dedicated software which could perform Bayesian VAR inference is essential for econometricians: for example, MATLAB with the BEAR toolbox (Dieppe et al. 2016) and the recent Empirical Macro Model toolkit by Canova and Ferroni (2022); however, other alternatives are available such as the new Bayesian functionalities provided in STATA from release 17, Eviews built-in commands or the R packages `bvar` (Kuschnig and Vashold 2021) and `bvartools` (Mohr 2023).

`gretl`, on the other side, dispenses many utilities for VARs either via internal commands (e.g., `var`, `vecm`) and functions (e.g., `vma`, `varsimul`, `irf`) or via the more accessible graphical interface: the routines include, among others, lag selection procedures, Granger-causality tests, forecast, structural analysis via short-run restrictions (Cholesky identification) and cointegration analysis. Additional functionalities can be retrieved via packages: the SVAR addon (Lucchetti and Schreiber 2023), for example, integrates the native structural analysis with various identification schemes such as long-run or sign restrictions; the package MSVAR (Schreiber 2021) provides, instead, Markov-Switching VAR models. The common denominator is the predominant use of frequentist techniques; only recently the BVAR package (Pedini and Schreiber 2024) has brought the attention to Bayesian techniques. This contribution, however, is still preliminary and its applicability in real data context is little documented.

To promote a more rigorous guideline to `gretl` users, this paper proposes a replication exercise of Kilian and Lütkepohl (2017), Chapter 5, where firstly a Bayesian VAR on U.S. main macroeconomic variables (Koop and Korobilis 2010) is estimated under different prior setups and then Cholesky-identified impulse responses are produced. Bayesian computation was originally performed via an independent software implementation, which has now become part of the BVAR package. For this reason, I refer to the package functionalities, but all computing and software technicalities will be minimized favoring, instead, an applied and didactic perspective.¹ The reported exercise wants to offer a first valuable guidance for `gretl` users in understanding and applying the Bayesian paradigm in simple, but often useful, contexts. A generalization to other applications or data is immediate.

The rest of the paper is organized as follows: Sect. 2 describes Bayesian VAR models alongside their main features in terms of prior distributions and posterior simulation; Sect. 3 reports the proposed replication exercise; Sect. 4 concludes.

¹ For more details on BVAR see the package manual available at https://gretl.sourceforge.net/current_fnfiles/unzipped/BVAR.pdf.

2 Statistical background

Define a VAR(p) model as:

$$y_t = a_0 + \sum_{j=1}^p A_j y_{t-j} + \varepsilon_t, \tag{1}$$

where y_t is an $m \times 1$ vector of observable variables at time t ; ε_t the vector of disturbances; a_0 and A_j , respectively a $m \times 1$ vector and a $m \times m$ matrix, are the coefficients of interest. For notational convenience assume from now on $A = [a_0, A_1, \dots, A_p]$. Under the *i.i.d.* assumption of $\varepsilon_t \sim N(\mathbf{0}, \Sigma)$, meaning normally distributed errors with zero mean and variance-covariance matrix Σ , the likelihood $\mathcal{L}(y|A, \Sigma)$ is also normal with maximum \hat{A} as the least squares estimator.

The Bayesian framework replaces the centrality of classical estimators with that of *posterior* distributions, namely $P(A, \Sigma|y)$. On this line, posteriors are obtained by combining the likelihood with *prior* distributions $P(A, \Sigma)$, in symbols

$$P(A, \Sigma|y) \propto \mathcal{L}(y|A, \Sigma)P(A, \Sigma)$$

Priors embody the subjective contribution of the practitioner and should reflect the theoretical beliefs about the long-run properties of the data. However, to make posterior computation feasible, some standardized setups are generally applied with the major discriminant in the error covariance Σ treated either deterministically or as a random variable.

In accordance with the replication exercise in Sect. 3, a brief summary of the involved priors is here reported: Sect. 2.1 considers the case of a deterministic Σ with regressor coefficients as normally distributed with the so-called Minnesota prior; Sect. 2.2, instead, introduces a prior for Σ leading to two major scenarios, the conjugate prior setup and the independent one. Finally, Sect. 2.3 briefly describes the derivation of impulse responses (structural and not) in a Bayesian framework.

2.1 Fixed- Σ case

Assume Σ as deterministic and known *a priori*. In particular, Σ may be substituted by a valid frequentist estimator $\hat{\Sigma}$ which, depending on the application at hand, corresponds to *i*) a diagonal matrix of residual variances from single equation auto-regressive (AR) models; *ii*) the diagonal elements of the VAR residual covariance matrix; *iii*) the full VAR residual covariance matrix.

Given the premise, the only parameter of interest becomes A : to simplify the exposition suppose to work with vectorized parameters, i.e., $\alpha = \text{vec}(A)$. The canonical framework relies on

$$\alpha_{M \times 1} \sim N(\underline{\alpha}, \underline{V}) \tag{2}$$

where $M = m(1 + mp)$ and $N(\mu, V)$ denotes a multivariate Gaussian distribution with mean μ and variance-covariance matrix V . The most popular choice for prior

parameters follows Litterman (1986) who proposes

$$\underline{\alpha}_{[i]} = \begin{cases} 1 & \text{for } E(A_{1[l]l}), \\ 0 & \text{otherwise} \end{cases} \tag{3}$$

$$\underline{V}_{[ii]} = \begin{cases} (\frac{\pi_1}{j\pi_3})^2 & \text{for } V(A_{j[l]l}), \\ (\frac{\pi_1\pi_2\sigma_l}{j\pi_3\sigma_q})^2 & \text{for } V(A_{j[lq]}), \quad l \neq q \\ (\pi_1\pi_4)^2\sigma_l^2 & \text{for } V(a_{0[l]}) \end{cases} \tag{4}$$

with $l = 1, \dots, m$ and $i = 1, \dots, M$. The notation $\mathbf{a}_{[i]}$ and $A_{[lq]}$ identify respectively the i -th element of the vector \mathbf{a} and the lq -th element of the matrix A ; while the operators $E(\cdot)$ and $V(\cdot)$ denote the expected value and the variance. The scalars $\pi_1, \pi_2, \pi_3, \pi_4$ are up to the user, with σ_l^2 commonly replaced by the least squares residual variance of an auto-regressive model for variable l . Clearly, the matrix \underline{V} is diagonal. The above specification is universally known as Minnesota prior and lots of its success derives from a combination of two factors: the interpretation, in line with most macroeconomic stylized facts (unit root series with a declining impact for further lags) and the great computational tractability.

The final implication of the framework pertains to the posteriors: with α as the sole parameter of interest, $P(\alpha|y)$ is normal with easily-derivable moments.

2.2 Priors for Σ

Assume this time that Σ is a random variable, in particular,

$$\Sigma_{m \times m} \sim IW(\underline{S}, \underline{\nu}) \tag{5}$$

where $IW(S, \nu)$ denotes an Inverse Wishart distribution with scale S and ν degrees of freedom. Once again, the prior hyperparameters are up to the user, with some common choices for the scale as the identity matrix or a diagonal matrix holding the residual variances from single-equation AR models; $\underline{\nu}$, instead, is generally set to $m + 2$.

Introducing a prior for Σ leads to defining a joint distribution $P(\alpha, \Sigma)$, which could be further decomposed into the alternatives

- $P(\alpha|\Sigma)P(\Sigma)$;
- $P(\alpha)P(\Sigma)$.

The former case is often associated with the so-called conjugate Normal-Inverse Wishart (NIW) setup; while the latter with the independent NIW. The conjugate NIW postulates the following prior specification:

$$\begin{aligned} \alpha|\Sigma &\sim N(\underline{\alpha}, \Sigma \otimes \underline{V}) \\ \Sigma &\sim IW(\underline{S}, \underline{\nu}) \end{aligned}$$

the paramters $\underline{\alpha}, \underline{S}, \underline{\nu}$ have to be intended as before, while \underline{V} is a $(1 + mp) \times (1 + mp)$ diagonal matrix; \otimes identifies the Kronecker product. Differently from Eq. (2), \underline{V}

represents now the variance of “the single equation” in the VAR, meaning that a simple transposition of the Minnesota structure is not viable. A solution is provided in Karlsson (2013), who actually adapts the Minnesota framework to the conjugate NIW, but an equally popular alternative relies on ridge priors i.e., $\underline{V} = \eta I$ with I as the identity matrix² and η as a shrinking factor. Posteriors are again available in closed-form with $\alpha|\Sigma, y$ as normal and $\Sigma|y$ as an Inverse Wishart.

The independent NIW configuration, instead, relaxes the Kronecker structure on the α -prior by postulating its independence from Σ . In other words,

$$\begin{aligned}
 P(\alpha, \Sigma) &= P(\alpha)P(\Sigma) \\
 \alpha &\sim N(\underline{\alpha}, \underline{V}) \\
 \Sigma &\sim IW(\underline{S}, \nu)
 \end{aligned}$$

The improvement concerns the flexibility in eliciting priors: the parameter α follows *de facto* Eq. (2), allowing for any arbitrary choice in covariance matrices. Deriving the posterior distribution requires the use of Markov Chain Monte Carlo (MCMC) methods: a Gibbs sampler which loops through the conditional posteriors ($\alpha|\Sigma, y$ is normally distributed and $\Sigma|\alpha, y$ is an Inverse Wishart) can be easily devised.

2.3 Impulse response functions

Deriving posteriors for functions of (α, Σ) is almost universally accomplished via simulation regardless of the availability of closed-form distributions: impulse response functions (IRFs) for dynamic analysis are a typical example.

The procedure is summarized in Algorithms 1-2: Algorithm 1 describes the Monte Carlo experiment for drawing (α, Σ) parameters in the previously considered scenarios. Note that a generic draw at iteration i is here expressed with the superscript (i) : $\alpha^{(i)}$, for example, has to be intended as the sampled vector α at iteration i .

Recall that IRFs are defined as the response of y_{t+h} to a one-time shock in ε_t ,

$$C_h = \frac{\partial y_{t+h}}{\partial \varepsilon_t}$$

where C_h contains the IRFs at time h . Note that $C_{h[lq]}$ identifies the response of variable $y_{t+h[l]}$ from a shock in $\varepsilon_{t[q]}$, with $l, q = 1, \dots, m$.

A practical solution for computing $C_{h[lq]}$ is to simulate a VAR process for h period by setting as initial conditions $\varepsilon_{0[q]} = 1, \varepsilon_{0[s]} = 0, s \neq q$ and $y_{-1} = y_{-2} = \dots = y_{-p} = \mathbf{0}$. The detail is reported in Algorithm 2 where the values for A are provided by the previous posterior simulations. Simple IRFs are rarely of any practical interest *per se*: what is widely more predominant is the use of structural IRFs (SIRFs). In the following, I will briefly define one of the possible schemes, i.e., the Cholesky identification, for pursuing structural analysis and, in doing so, I will avoid any unnecessary

² A generic $k \times k$ identity matrix is defined as I_k . Here, for notational convenience, the subscript has been dropped.

Algorithm 1 Sampling from the posterior distributions

```

begin
  Posterior parameter sampling:
  for  $i = 1$  to  $MC$  do
    (where  $MC$  is the number of iterations)
    • Fixed- $\Sigma$  case
      1.  $\Sigma$  is assumed known:  $\Sigma^{(i)} = \hat{\Sigma}$ 
      2. Draw  $\alpha^{(i)}$  from  $P(\alpha|\hat{\Sigma}, y) \sim N$ 
    • Conjugate NIW
      1. Draw  $\Sigma^{(i)}$  from  $P(\Sigma|y) \sim IW$ ;
      2. Draw  $\alpha^{(i)}$  from  $P(\alpha|\Sigma^{(i)}, y) \sim N$ .
    • Independent NIW
      1. Draw  $\Sigma^{(i)}$  from  $P(\Sigma|\alpha^{(i-1)}, y) \sim IW$ ;
      2. Draw  $\alpha^{(i)}$  from  $P(\alpha|\Sigma^{(i)}, y) \sim N$ .
    (Discard a burn-in interval  $b$  at the beginning of the sample)
  end
end

```

detail. For a scrupulous study of Structural VARs or identification schemes I refer to specialized textbooks such as Kilian and Lütkepohl (2017).

Algorithm 2 Sampling impulse response functions

```

begin
  Impulse response function simulation:
  for  $i = 1$  to  $MC$  do
    (where  $MC$  is the number of iterations)
    From one of the previous posterior samplers retain  $\alpha^{(i)}$ .
    Impulse response part:
    From  $\alpha^{(i)}$  recover  $A_1^{(i)}, \dots, A_p^{(i)}$ 
    for  $q = 1$  to  $m$  do
      
$$C_{h[\ast, q]}^{(i)} \leftarrow \tilde{y}_h^{(i)} = \sum_{j=1}^p A_j^{(i)} \tilde{y}_{h-j}^{(i)} \quad h = 1, \dots, H$$

      with  $\tilde{y}_0 = \varepsilon_0$ ;  $\tilde{y}_{-1} = \dots = \tilde{y}_{-p} = \mathbf{0}$ .
      
$$\varepsilon_0 \leftarrow \varepsilon_{0[s]} = \begin{cases} 1 & s = q \\ 0 & s \neq q \end{cases} \quad s = 1, \dots, m$$

      The notation  $C_{h[\ast, q]}$  means the  $q$ -th column of  $C$  (corresponding to the  $q$ -th shock).
    end
  end
end

```

SIRFs with a Cholesky identification scheme are obtained by rewriting the error term as

$$\varepsilon_t = Bu_t, \quad u_t \sim N(\mathbf{0}, \mathbf{I}_m)$$

$$D_h = C_h B = \frac{\partial y_{t+h}}{\partial u_t}$$

where u_t is the vector of structural shocks and B a lower-triangular matrix containing the simultaneous effect of the variables (note that $\Sigma = BB'$, therefore B is the Cholesky factor of Σ); D_h is the structural counterpart of C_h . Clearly, working with u_t instead of ε_t allows to disentangle the shocks from the correlation effect, achieving in this way a more authentic measure of the responses. Algorithm 2 can be revised to account for this change by retrieving the Cholesky decomposition of the sampled $\Sigma^{(i)}$, namely $B^{(i)}$, and post-multiplying each $C_h^{(i)}$ by this term in the IRF step.

3 Replication exercise

In this Section I will replicate the Bayesian VAR examples proposed by Kilian and Lütkepohl (2017), Chapter 5. This exercise has to be intended as a how-to-do “manual” for performing simple Bayesian analyses, with particular attention to structural responses via the Cholesky identification scheme. The data comes from Koop and Korobilis (2010) and includes the quarterly US GDP price deflator (`inflation`), the seasonally adjusted unemployment rate (`unemployment`) and the 3-month Treasury bill yield (`interest`) for the time span 1953q1-2006q3. The series are reported in Fig. 1.

The baseline specification follows a VAR(4); three different prior setups are then introduced: the fixed- Σ case with a Minnesota prior for α , a conjugate and an independent NIW with ridge priors. Each configuration explores various prior-hyperparameter choices whose contribution can be observed by means of SIRFs. In particular, Kilian and Lütkepohl (2017) impose the variable order as `inflation`, `unemployment`, `interest` and inspect the response of inflation to a shock in interest rate (contractionary monetary policy shock). For future comparisons, the frequentist SIRF is reported in Fig. 2.

The replication code exploits the collection of functions provided by the BVAR package (Pedini and Schreiber 2024),³ but a clarification is necessary: even if I refer to its functionalities, I will avoid a comprehensive package showcase. The preliminary state of BVAR prevents any proper package presentation which will be postponed for a future contribution. Conversely, it could be fruitfully used for introducing simple instructions for practitioners new to Bayesian VAR analysis. The scripts, in fact, can be easily extended to different data and applications.

³ BVAR was officially released (version 0.1) on August 2023: the included features were posterior simulators under the diffuse, fixed, conjugate and independent prior setups with the possibility of deriving predictions and Cholesky-identified IRFs. Version 0.2 (January 2024) introduced a substantial modification in the plotting functions in order to accommodate the introduction of new `gretl` commands such as `gridplot`. Version 0.3 will be released in March-April 2024, featuring a major speed-up in the simulation routines. For major details, see the package manual available at https://gretl.sourceforge.net/current_fnfiles/unzipped/BVAR.pdf.




Fig. 1 Quarterly US inflation, unemployment rate and Treasury bill interest rate over the period 1953q1-2006q3

BVAR usage hinges upon four major declarations, namely `BVAR_setup`, `BVAR_sigma_prior`, `BVAR_alpha_prior`, `BVAR_posterior`. In particular, `BVAR_setup(list y, lag p, string prior_type)` is a “setup” function which computes some preliminary quantities given the list of endogenous variables y , the lag order p and the kind of prior `prior_type` (supported

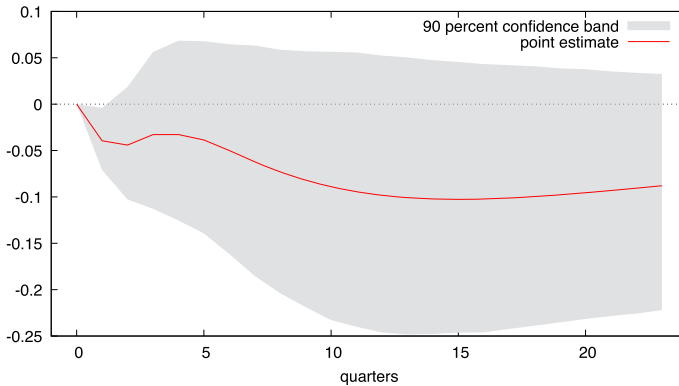


Fig. 2 Structural response of inflation to a shock in interest rate - Classical VAR(4) estimates. The plot is obtained using the gretl GUI, on the path Model->Multivariate Time Series->Vector Autoregression

options: "fixed", "conj", "indep"). Prior customization is then demanded to `BVAR_sigma_prior` and `BVAR_alpha_prior`. Both functions modify the resulting bundle from `BVAR_setup` with the help of dedicated options for prior hyperparameters. Their functioning is directly illustrated in the examples reported below. Finally, `BVAR_posterior` incorporates the routines previously presented in Algorithms 1-2. The signature is `BVAR_posterior(bundle setup, string simul, bundle opt)` where `setup` is the previously defined setup bundle, `simul` governs the simulation procedure: "coeff" will trigger the sole coefficient derivation, "irf" will augment the scheme with IRFS (and Cholesky identified IRFs following the variable order in the list `y`). The bundle `opt` allows the user to finely tune computational details such as the number of replications (bundle key `iter`, default: 10000), the burn-in percentage (bundle key `burn`, default: 10% for MCMC experiments, 0 otherwise), the quantile range to report as "confidence band" for posterior distributions (bundle key `q_range`, default 0.90) or the IRF time horizon (bundle keys `irf_h`, default: 24). If omitted, the default choices will be used. An auxiliary function for plotting the results is also at disposal: `BVAR_irf_plot(bundle out, list from, list to, bundle irfopt)`,⁴ which will display the posterior median of the sampled IRFs with a shaded area around defined via `q_range`; the time horizon follows `irf_h`. The bundle `out` is the returning bundle from `BVAR_posterior`; the lists `from`, `to` act as selectors for the IRFs to plot: in particular, `from` stands for the variables whose error is perturbed and `to` the outcome variables. If omitted, all combinations are shown on the screen. The optional bundle `irfopt` allows to select the saving options (key `save`, default: display on screen), the gretl gridplot integration (key `grid`, default: 1), particularly useful when showing multiple IRFs, the printing of the posterior sample mean (key `no_mean`, default: 1 - meaning suppress the mean plot), and the kind of IRFs to report between structural and not (key `irf_type`, default: `s_irf`).

⁴ For compatibility reasons with older gretl releases than 2023c, BVAR version 0.2 includes an additional function for plotting IRFs, `BVAR_irf_plot_OLD`. Again, the package manual is the reference.

To sum up, the workflow here adopted is the following:

```

1. # Step 1 - load the pkg and open the data
2. include BVAR.gfn
3. open koopkoro.gdt --quiet --frompkg=BVAR # the dataset
4.
5. # Step 2 - define preliminary quantities
6. list dep = inflation unemployment interest
7. scalar lag = 4
8.
9. # Step 3 - setup
10. string type = "fixed" # or "conj", "indep"
11. bundle mod = BVAR_setup(dep, lag, type)
12.
13. # Step 4 - priors, major details below
14. BVAR_sigma_prior(&mod, "ADL") # first sigma prior
15. BVAR_alpha_prior(&mod, "Minn") # then alpha prior
16.
17. # Step 5 - posterior simulation (Bayesian estimation)
18. bundle res = BVAR_posterior(mod,"irf", \
    _(iter=10000, burn = 0, irf_h = 24))
19.
20. # Step 6 - plot
21. BVAR_irf_plot(res, interest, inflation,
    _(no_mean=0))

```

Lines 14-15 report only the prior functions invocation; as previously mentioned, all specific options will be shown in the next examples. Note, however, how the functions in question modify the setup bundle `mod` via the pointer solution (`&mod`).⁵

3.1 Fixed- Σ case

The first scenario assumes a known Σ obtained via single equation Autoregressive Distributed Lag models (ADL) of order 4;⁶ the prior for α is a Minnesota-like type as in Eqs. (2)-(4): Kilian and Lütkepohl (2017) distinguish two major cases depending on how the persistency in the series is perceived. For an integrated nature, they propose a random-walk mean, i.e., $\underline{\alpha}$ is specified as in Eq. (3); conversely, if the long-run tendency is expected to be stationary, the white-noise mean $\underline{\alpha} = \mathbf{0}$ is proposed. \underline{V} follows Eq. (4) where the following combinations of hyperparameters are inspected:

- $\pi_1 = 0.1, \pi_2 = 0.5$;
- $\pi_1 = 0.1, \pi_2 = 1$;

⁵ For major details see Cottrell and Lucchetti (2023b, a).

⁶ Kilian and Lütkepohl (2017) report the use of simple least squares estimates from the VAR(4), however as their replication code clarifies, only the diagonal elements are retrieved. These correspond to simple ADL models of order 4.

- $\pi_1 = 1, \pi_2 = 1$.

The choices $\pi_3 = 1, \pi_4 = 1000$ and σ_I^2 from AR(4) models are adopted in all cases,⁷ while posterior simulation is executed via 10000 Monte Carlo iterations.

Starting from the previous pseudo-code, the prior customization can be performed as follows,

```
14. BVAR_sigma_prior(&mod, "ADL") # as before
15. BVAR_alpha_prior(&mod, "Minn", _(p_mean = 1, \
    p_cov_p1 = 0.1, p_cov_p2 = 0.5))
```

`BVAR_sigma_prior`, in the "fixed" case, is a simply utility function for computing $\hat{\Sigma}$: the second argument can be fixed to "ADL" for computing a diagonal matrix with ADL residual variances; "AR" for a diagonal matrix of AR variances. The lag order is the same specified via `BVAR_setup`.

The invocation `BVAR_alpha_prior(&mod, "Minn")` can be used to recover a Minnesota prior shape; a third and optional input can be supplied in the form of a bundle: this is intended for setting hyperparameter values. The supported keys are `p_mean` for setting $E(A_{1|11})$ in $\underline{\alpha}$ (default = 1); `p_cov_p1` for π_1 (default=0.1), `p_cov_p2` for π_2 (default=0.5), `p_cov_p3` for π_3 (default=1) and `p_cov_inter` for π_4 (default=1000).

Figure 3 reports the BVAR Cholesky-identified IRFs of inflation due to an interest rate shock under the different prior hyperparameter setups: the dashed red line represents the posterior median, while the blue one the posterior sample mean. The shaded area is delimited by the 0.05 and 0.95 quantiles. Figure 4 illustrates the results obtained in MATLAB by means of Kilian and Lütkepohl (2017) replication code⁸ the dashed red lines represent respectively the 0.05, 0.5, 0.95 posterior quantiles; the blue line is the posterior sample mean. As it can be noted, the results are replicated in `gretl`.

3.2 Conjugate NIW

Assume now the conjugate NIW setup: Kilian and Lütkepohl (2017) consider $\Sigma \sim IW(I_3, 4), \alpha|\Sigma \sim N(\underline{\alpha}, \Sigma \otimes \underline{V})$. The expected value $\underline{\alpha}$ follows again the random-walk scenario with ones on the first own lag of the equation endogenous variables or the white-noise case with zeros; \underline{V} is defined via a ridge prior ηI_{13} . Three possible values for η are analyzed:

- $\eta = 0.01$;
- $\eta = 1$;
- $\eta = 100$.

under a common simulation design of 10000 Monte Carlo replications. The prior specification is obtained in `gretl` as follows:

⁷ The value π_4 is highly uninfluential provided $\pi_4 \geq 100$.

⁸ The MATLAB: code is available at the author webpage, <https://sites.google.com/site/ikilian2019/textbook/code>.

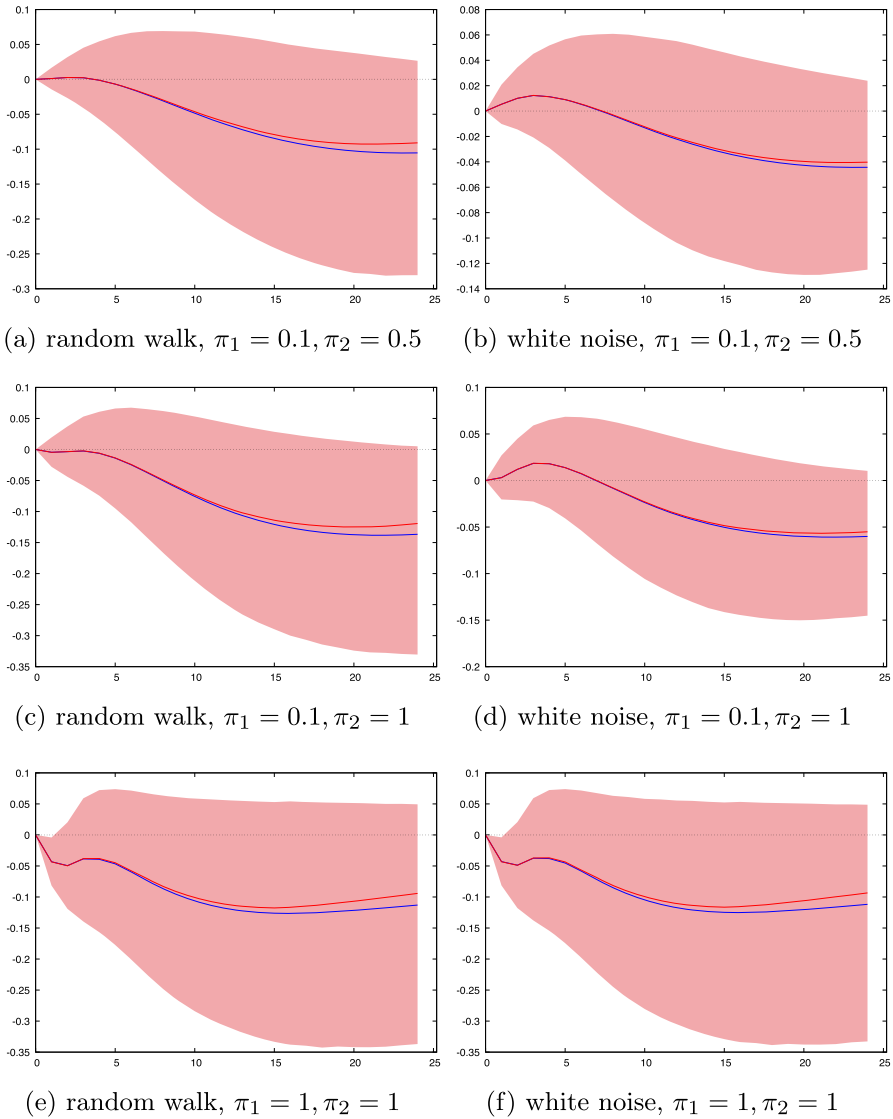


Fig. 3 Structural impulse response of inflation to an interest rate shock. Bayesian VAR with fixed- Σ : gret1 script

```

14. scalar m = nelem(dep) # m = 3
15. scalar param_eq = m*lag + 1 # param per eq = 13
16. scalar param_tot = m*param_eq # param tot = 39
17. BVAR_sigma_prior(&mod, "custom", \
    _(sigma_scale = ones(m), sigma_df = m + 1))
18. BVAR_alpha_prior(&mod, "custom", \

```

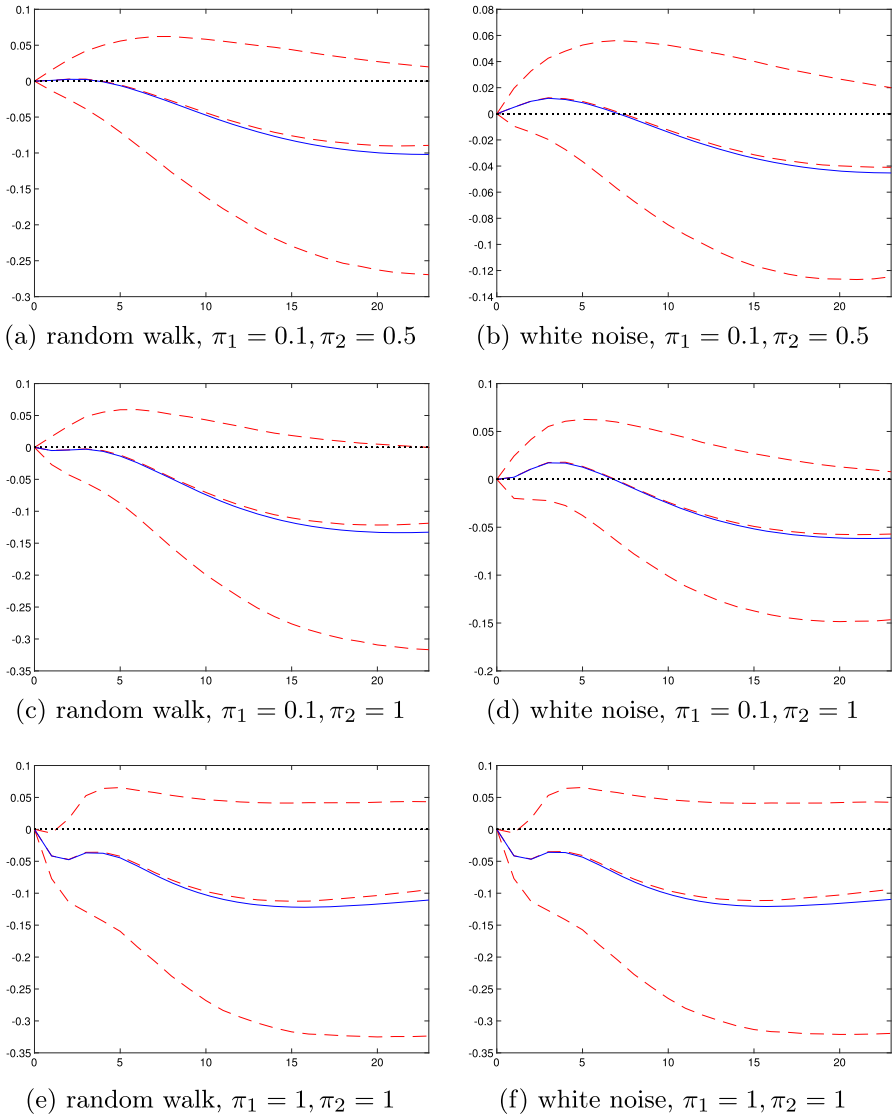


Fig. 4 Structural impulse response of inflation to an interest rate shock. Bayesian VAR with fixed- Σ : Kilian and Lütkepohl (2017) MATLAB script

```
_alpha_mean = zeros(param_tot), \
alpha_vcv = 0.01*ones(param_eq)
```

BVAR_sigma_prior with the "custom" modifier allows to directly express the scale and the degrees of freedom. The same modifier can be applied to α with analogous consequences; note that α has 39 elements corresponding to the total VAR coefficients (param_tot), while \underline{V} is a 13×13 matrix (param_eq=13, m = 3, p

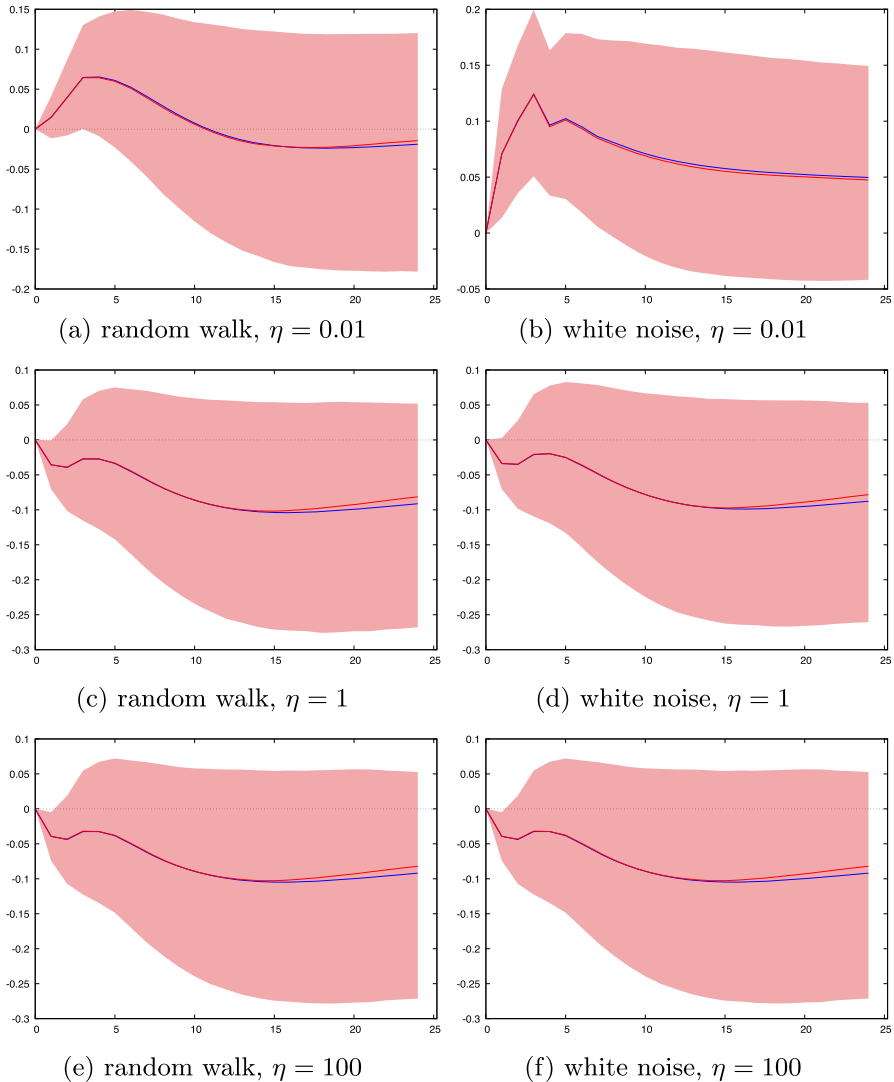


Fig. 5 Structural impulse response of inflation to an interest rate shock. Bayesian VAR with conjugate NIW: `gretl` script

= 4, 1 intercept). Finally, it suffices to specify a vector of elements for both the scale of Σ and the variance of α since diagonal solutions are employed. Figures 5, 6 report the comparison between software: `gretl` replicates exactly the result.




Fig. 6 Structural impulse response of inflation to an interest rate shock. Bayesian VAR with conjugate NIW: Kilian and Lütkepohl (2017) MATLAB script

3.3 Independent NIW

In the independent NIW context, the Σ -specification exactly recalls the one used in the conjugate case, i.e., $\Sigma \sim IW(I_3, 4)$; α is again Normal with the random walk - white noise option for $\underline{\alpha}$. The previous \underline{V} is replaced by $\underline{V} = \eta I_{39}$ with ridge parameter η :

- $\eta = 0.01$;
- $\eta = 1$.

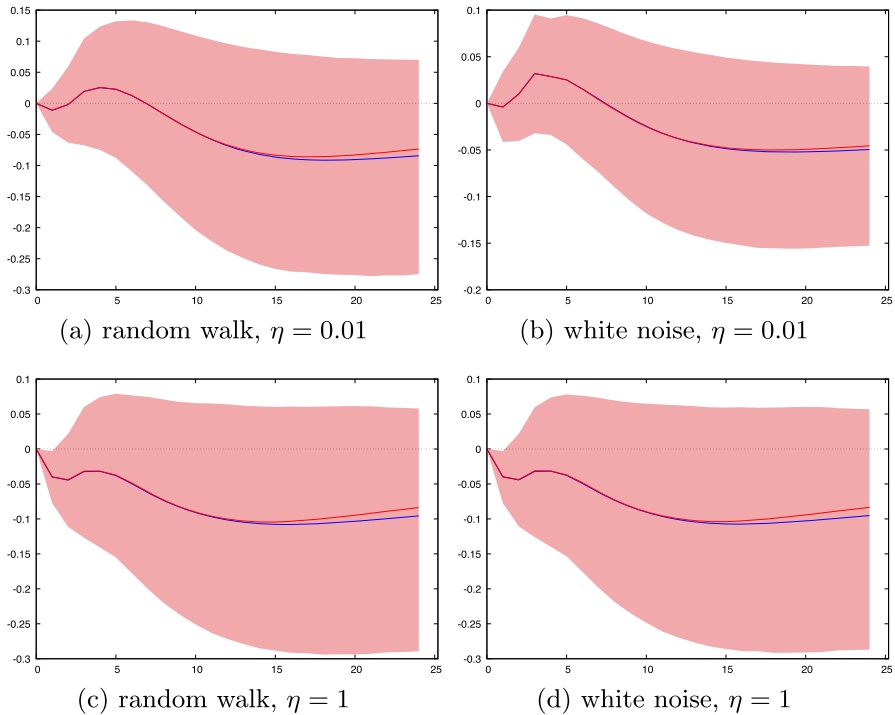


Fig. 7 Structural impulse response of inflation to an interest rate shock. Bayesian VAR with independent NIW: *gretl* script

Since a MCMC is employed, a burn-in period of 10% has been introduced. The *gretl* script is closely related to the previous one, but the dimensional mismatch introduced by \underline{V} requires the following modification

```
18. BVAR_alpha_prior(&mod, "custom", \
    _(alpha_mean = zeros(param_tot), \
    alpha_vcv = 0.01*ones(param_tot))
```

The results of the replication are reported in Fig. 7 for *gretl* and in Fig. 8 for MATLAB. Again, the IRFs are successfully replicated.

4 Conclusion

This short paper has highlighted, via the replication of Kilian and Lütkepohl (2017) - Chapter 5, the use of *gretl* in Bayesian VAR contexts, showing how the software can be effectively counted among those programs that provide such kind of analysis. The utilities and routines are still limited, but can be equally valuable, especially as preliminary analyses or in a didactic context. Needless to say, additional functionalities

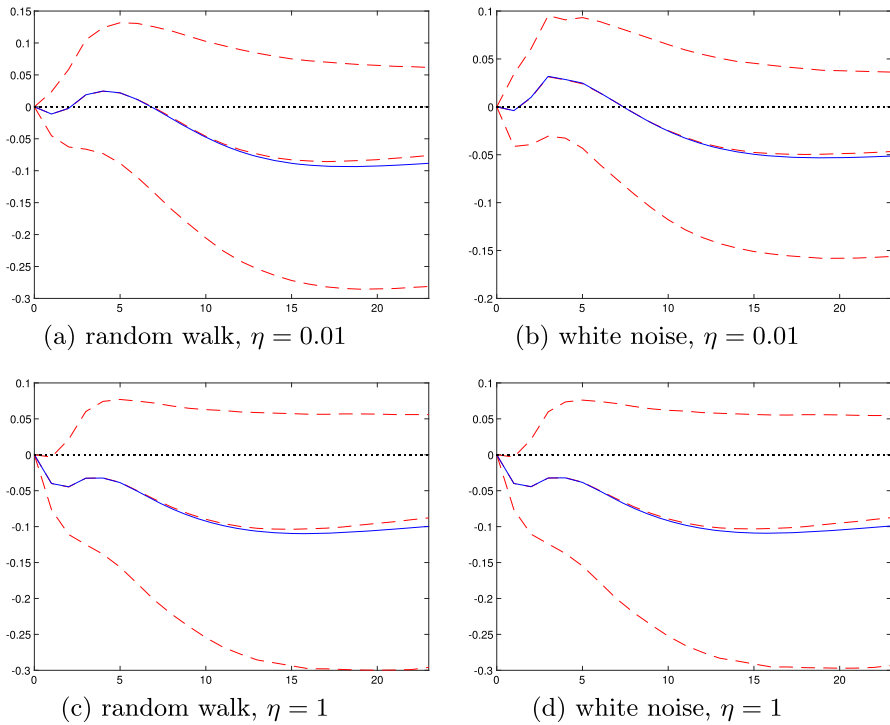


Fig. 8 Structural impulse response of inflation to an interest rate shock. Bayesian VAR with independent NIW: Kilian and Lütkepohl (2017) MATLAB script

and support will be provided in the future: packages that could address different aspects of the Bayesian multivariate time series modeling (e.g., Bayesian Structural VARs or Time Varying VARs) are a must and the interaction with the `gretl` ecosystem could bring to enormous advantages for the whole software community.

Computational details

The replication exercise has been run on `gretl 2023c` and `gretl 2024a` on a macOS laptop, using `BVAR` package, version 0.2.

Acknowledgements I wish to thank Sven Schreiber, all the participants at the 2023 `gretl` Conference and two anonymous referees for their useful remarks and observations.

Funding Open access funding provided by Università Politecnica delle Marche within the CRUI-CARE Agreement. the author declares that no funds, grants or other support were received during the preparation of this manuscript;

Data availability statement the dataset is available as sample data in the package `BVAR`.

Declarations

Conflict of interest the author has no relevant financial or non-financial interest to disclose;

Code availability replication codes are provided as attached files.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Baumeister C, Hamilton JD (2015) Sign restrictions, structural vector autoregressions, and useful prior information. *Econometrica* 83(5):1963–1999
- Bernanke BS, Gertler M, Watson M, Sims CA, Friedman BM (1997) Systematic monetary policy and the effects of oil price shocks. *Brook Pap Econ Activity* 1997(1):91–157
- Bernanke BS, Mihov I (1998) Measuring monetary policy. *Q J Econ* 113(3):869–902
- Canova F, Ferroni F (2022) Empirical macro model toolbox. Retrieved from https://github.com/naffe15/BVAR_
- Christiano LJ (2002) Solving dynamic equilibrium models by a method of undetermined coefficients. *Comput Econ* 20:21–55
- Cogley T, Sargent TJ (2001) Evolving post-world war II US inflation dynamics. *NBER Macroecon Annu* 16:331–373
- Cogley T, Sargent TJ (2005) Drifts and volatilities: monetary policies and outcomes in the post WWII US. *Rev Econ Dyn* 8(2):262–302
- Cottrell A, Lucchetti R (2023a) Gretl function package guide [Computer software manual]. Retrieved from <https://sourceforge.net/projects/gretl/files/manual/pkgbook.pdf/download>
- Cottrell A, Lucchetti R (2023b) Gretl User's Guide [Computer software manual]. Retrieved from <https://sourceforge.net/projects/gretl/files/manual/gretlguide.pdf/download>
- Del Negro M, Schorfheide F, Smets F, Wouters R (2007) On the fit of new Keynesian models. *J Bus Econ Stat* 25(2):123–143
- Dieppe A, Legrand R, Van Roye B (2016) The BEAR toolbox. ECB working paper
- George EI, Sun D, Ni S (2008) Bayesian stochastic search for VAR model restrictions. *J Economet* 142(1):553–580
- Giacomini R (2013) The relationship between DSGE and VAR models. *VAR models in Macroeconomics-New developments and applications: Essays in honor of Christopher A. Sims* 32:1–25
- Giannone D, Lenza M, Primiceri GE (2015) Prior selection for vector autoregressions. *Rev Econ Stat* 97(2):436–451
- Karlsson S (2013) Forecasting with Bayesian vector autoregression. *Handbook Econ Forecast* 2:791–897
- Kilian L, Lütkepohl H (2017) *Structural vector autoregressive analysis*. Cambridge University Press, Cambridge
- Koop G, Korobilis D (2010) Bayesian multivariate time series methods for empirical macroeconomics. *Foundations and Trends@in Econometrics*, 3(4), 267–358
- Korobilis D (2013) VAR forecasting using Bayesian variable selection. *J Appl Economet* 28(2):204–230
- Kuschnig N, Vashold L (2021) Bvar: Bayesian vector autoregressions with hierarchical prior selection in r. *J Stat Softw* 100:1–27
- Litterman RB (1986) Forecasting with Bayesian vector autoregressions—five years of experience. *J Bus Econ Stat* 4(1):25–38
- Lucchetti R, Schreiber S (2023) The SVAR addon for gretl. (addon version 2.1)

- Mohr FX (2023) *bvartools*, Bayesian inference of vector autoregressive and error correction models. Retrieved from <https://cran.rspoproject.org/package=bvartools>
- Pedini L, Schreiber S (2024) BVAR: Bayesian VAR in gretl. Retrieved from https://gretl.sourceforge.net/current_fnfiles/BVAR.zip (packageversion 0.2)
- Primiceri GE (2005) Time varying structural vector autoregressions and monetary policy. *Rev Econ Stud* 72(3):821–852
- Rubio-Ramirez JF, Waggoner DF, Zha T (2010) Structural vector autoregressions: theory of identification and algorithms for inference. *Rev Econ Stud* 77(2):665–696
- Schreiber S (2021) The MSVAR package for gretl. Retrieved from https://gretl.sourceforge.net/current_fnfiles/MSVAR.zip (package version 0.22)
- Sims CA (1980) Macroeconomics and reality. *Econometrica*, 1–48
- Sims CA, Zha T (1998) Bayesian methods for dynamic multivariate models. *Int Econ Rev*, pp 949–968
- Sims CA, Zha T (2006) Does monetary policy generate recessions? *Macroecon Dyn* 10(2):231–272
- Sims CA, Zha T (2006) Were there regime switches in US monetary policy? *Am Econ Rev* 96(1):54–81
- Waggoner DF, Zha T (2003) A Gibbs sampler for structural vector autoregressions. *J Econ Dyn Control* 28(2):349–366

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.