Vine Copula VAR*

Hunter Ng[a]

[a] *Zicklin School of Business, Baruch College, City University of New York, United States*

May 4, 2026

Abstract

This paper introduces a scalable and order-invariant framework for modeling high-dimensional macro-financial dynamics. We develop a *Time-Varying Parameter VAR with Vine Copula Dependence* (VCVAR), which couples a lightweight TVP backbone with a flexible copula representation of cross-sectional dependence. The approach separates marginal dynamics from joint dependence, allowing the model to capture structural drift together with asymmetric, nonlinear, and tail-dependent co-movements that standard Gaussian or Cholesky-based VARs cannot represent. Despite this flexibility, the VCVAR remains computationally tractable through discounted recursive estimation and simplified vine construction. Applications to empirical macroeconomic data and controlled simulation designs show that the VCVAR offers consistent improvements in medium- and long-horizon forecasting, demonstrating the value of combining time variation with non-Gaussian dependence in large systems.

Keywords: Copula; Dynamic Dependence

JEL Classification: C32; C53; C58.

*

1 Introduction

Large-scale vector autoregressions (VARs) are standard tools for forecasting and structural analysis in macroeconomics (Sims, 1980). As the dimension grows, however, the number of autoregressive coefficients and the elements of the error covariance matrix expand quadratically, so some form of regularization or structure is required. Recent work uses Bayesian shrinkage priors and efficient covariance factorizations to make large VARs feasible in practice (Bańbura, Giannone, and Reichlin, 2010; Carriero, Clark, and Marcellino, 2019; Wu and Koop, 2025). Most of these approaches, however, retain a Gaussian or elliptical structure for the joint distribution of shocks, so that all cross-series dependence is summarized by a covariance or factor matrix. This is restrictive for macro financial data, where co-movements often become nonlinear and asymmetric in stressed periods, and where tail dependence can change over time.

At the same time, the literature has developed two distinct ways to introduce time variation into VARs. (1) Stochastic volatility models (SV VARs) allow the variances and covariances of the innovations to evolve over time, while keeping the autoregressive coefficients fixed (e.g. Cogley and Sargent, 2005; Carriero, Clark, and Marcellino, 2019). (2) Time varying parameter VARs (TVP VARs), in contrast, allow the coefficients themselves to drift, capturing gradual changes in monetary transmission, pricing relationships, or risk sensitivities (e.g. Koop, Korobilis, and Pettenuzzo, 2019; Chan, Yu, and Zhang, 2024). A full TVP SV VAR combines both features, but even then dependence is typically modeled through a Gaussian covariance structure or a factor stochastic volatility block (e.g. Kastner, 2019). Volatility dynamics can generate heavy tails and regime shifts in individual series, but they do not by themselves deliver non Gaussian or edge specific dependence across variables.

Copula based VARs offer a different route. Tsionas, Izzeldin, and Trapani (2022) use Sklar's theorem to factor the joint distribution into univariate TVP marginals and a Gaussian mixture copula, estimated with MCMC and dimension reduction in the covariance matrices. This breaks the problem into many univariate TVP equations plus a separate copula stage, and allows non Gaussian cross equation dependence. However, the Gaussian mixture copula still relies on high dimensional covariance objects and requires nontrivial MCMC machinery. Other recent contributions address the ordering problem in stochastic volatility VARs by working with eigendecompositions or random compression (Koop, Korobilis, and Pettenuzzo, 2019; Wu and Koop, 2025), but they also maintain Gaussian dependence.

In this paper, we propose a simple copula VAR that combines a TVP backbone with a vine copula layer. Each series is modeled by a univariate TVP autoregression estimated with discounted recursive least squares and variance discounting, which yields Gaussian predictive marginals and

time varying volatility in a computationally light way. Cross series dependence is then introduced via a simplified regular vine copula (Bedford and Cooke, 2001; Nagler, 2025), fitted to the probability integral transforms of the TVP residuals. The vine copula is selected using the greedy Kendall's τ algorithm of Dissmann, Brechmann, Czado, and Kurowicka (2013). This construction is order invariant, allows each pair of variables to follow its own copula family, and delivers asymmetric, tail dependent co movements that cannot be represented by a Gaussian covariance matrix.

The vine copula is best viewed as an extension to the standard TVP VAR, rather than a substitute for stochastic volatility. The TVP component captures slow drift in the conditional mean dynamics, while the copula captures nonlinear and potentially asymmetric dependence in the innovations. In principle, one could combine TVP, stochastic volatility, and a vine copula in a single model, but this would substantially increase computational cost and complicate estimation. Our aim is to strike a balance between flexibility and tractability by pairing a simple TVP marginal specification with a flexible copula layer, which offers superior performance in tail tests.

The paper makes three contributions. First, we implement a TVP VAR with a modular copula dependence in which univariate dynamics are handled by a Bayesian estimator and cross-series dependence is handled by a vine copula. Second, we compare our vine copula VAR (VCVAR) to alternative large VAR specifications, including basic TVP-VAR and TVP-SV-VAR benchmarks, and document that it performs competitively and offers superior performance in longer-horizon forecasting. Third, using U.S. macroeconomic data and simulated data, we show that the VCVAR matches or outperforms standard models in both real time forecasting and simulation experiments, especially in long-horizon environments with structural drift and nonlinear co movements.

This paper links to two streams of literature. First, it contributes to the large-scale VAR literature that develops shrinkage, factor structures, and efficient covariance representations to make high-dimensional forecasting feasible (Bańbura, Giannone, and Reichlin, 2010; Carriero, Clark, and Marcellino, 2019; Wu and Koop, 2025). Our approach complements this work by replacing the usual Gaussian covariance block with a modular copula layer, allowing more flexible cross-series dependence without sacrificing computational tractability. Second, the paper relates to the literature introducing time variation into VARs. Stochastic-volatility VARs model evolving variances and covariances but retain fixed autoregressive coefficients (Cogley and Sargent, 2005; Carriero, Clark, and Marcellino, 2019), while TVP-VARs allow the coefficients themselves to drift (Koop, Korobilis, and Pettenuzzo, 2019; Chan, Yu, and Zhang, 2024). Our framework fits into this second class, using discounted recursive least squares for efficient updating of time-varying parameters, but augments it with a flexible copula structure rather than a Gaussian covariance or factor-SV representation (e.g. Kastner, 2019). Finally, we build on the emerging work on copula-based VARs and non-Gaussian dependence. Closest to us is Tsionas, Izzeldin, and Trapani (2022), who factor the joint distribu-

tion via Sklar's theorem and estimate a Gaussian-mixture copula, but their approach still requires manipulating large covariance matrices and MCMC. By contrast, our simplified vine construction (Bedford and Cooke, 2001; Nagler, 2025; Dissmann et al., 2013) delivers pairwise, family-specific, and order-invariant dependence in a computationally light way that scales well to higher dimensions.

2 Methodology

2.1 The Workhorse SV-VAR

Stochastic volatility and time-varying parameters are two orthogonal concepts: the former governs how the variances of structural shocks change over time, while the latter governs how the autoregressive coefficients change over time. We begin by describing stochastic volatility in an otherwise constant-parameter VAR.

The reduced-form VAR with stochastic volatility (SV-VAR) is a widely used specification in empirical macroeconomics. Let $\mathbf{y}_t = (y_{1t}, \dots, y_{nt})'$ denote an $n \times 1$ vector of observed variables for $t = 1, \dots, T$. Consider the reduced-form VAR(p):

$$\mathbf{y}_t = \mathbf{a}_0 + \mathbf{A}_1 \mathbf{y}_{t-1} + \cdots + \mathbf{A}_p \mathbf{y}_{t-p} + \boldsymbol{\varepsilon}_t, \qquad \boldsymbol{\varepsilon}_t \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}_t), \tag{2.1}$$

where $\mathbf{a}_0$ is an $n \times 1$ vector of intercepts and the coefficient matrices $\{\mathbf{A}_j\}_{j=1}^p$ are fixed over time, so the propagation mechanism does not change. Following Cogley and Sargent (2005), the innovation covariance matrix evolves according to the Cholesky representation

$$\boldsymbol{\Sigma}_t^{-1} = \mathbf{B}_0' \mathbf{D}_t^{-1} \mathbf{B}_0, \tag{2.2}$$

where $\mathbf{D}_t = \operatorname{diag}\big(e^{h_{1t}}, \dots, e^{h_{nt}}\big)$ and $\mathbf{B}_0$ is lower triangular with ones on the diagonal. Each log-variance $h_{i,t}$ follows a random-walk process:

$$h_{i,t} = h_{i,t-1} + u_{i,t}^h, \qquad u_{i,t}^h \sim \mathcal{N}(0, \sigma_i^2). \tag{2.3}$$

The joint density of $\{\mathbf{y}_t\}_{t=1}^T$ and the latent log-volatilities $\{h_{i,t}\}$ factorizes as

$$p(\mathbf{y}_{1:T}, h_{1:T}) = \prod_{t=1}^{T} \mathcal{N}\big(\mathbf{y}_t \mid \mathbf{A}_1 \mathbf{y}_{t-1} + \cdots + \mathbf{A}_p \mathbf{y}_{t-p}, \boldsymbol{\Sigma}_t\big) \prod_{i=1}^{n} \mathcal{N}\big(h_{i,t} \mid h_{i,t-1}, \sigma_i^2\big).$$

This specification allows the covariance matrix to evolve stochastically over time while keeping the autoregressive coefficients constant. When the eigenvalues of the VAR companion matrix have absolute value less than one, the process is stable and weakly dependent (Lütkepohl, 2005). The

latent log-volatility processes require simulation-based inference, typically through Bayesian MCMC. See Carriero, Clark, and Marcellino (2019) for computationally efficient implementations and Smith and Klein (2021) for variational approximations.

2.2 TVP-VAR

We now turn to time-varying parameters. While the SV-VAR introduces time variation only in the shock variances, the time-varying parameter VAR (TVP-VAR) allows the autoregressive coefficients themselves to evolve. A TVP-VAR can be written as

$$\mathbf{y}_t = \sum_{j=1}^{p} \mathbf{A}_{j,t}\mathbf{y}_{t-j} + \boldsymbol{\varepsilon}_t, \qquad \boldsymbol{\varepsilon}_t \sim F_t, \tag{2.4}$$

where the stacked coefficient vector $\boldsymbol{\beta}_{i,t}$ evolves according to a random walk,

$$\boldsymbol{\beta}_{i,t} = \boldsymbol{\beta}_{i,t-1} + \mathbf{u}_{i,t}, \qquad \mathbf{u}_{i,t} \sim \mathcal{N}\big(\mathbf{0},\, \delta_\beta^{-1}\mathbf{P}_{i,t-1}\big), \tag{2.5}$$

The SV-VAR in (2.1)-(2.3) corresponds to a specification in which the autoregressive coefficients are held fixed over time, which can be represented by setting $\delta_\beta = 1$. In that case, all time variation in the model arises from the stochastic volatility process that governs the shock variances. More broadly, time-varying parameters and stochastic volatility are conceptually distinct modeling choices. Time variation in the coefficients determines how the autoregressive relationships evolve, while stochastic volatility determines how the variances of the innovations change. These features may be introduced independently or combined within a unified VAR framework.

A full TVP-SV-VAR (e.g., Koop, Korobilis, and Pettenuzzo, 2019) incorporates both mechanisms and has been actively studied in the literature. In applied finance, each of these VAR extensions aligns naturally with a different form of time variation. A TVP-VAR could be preferable in settings where economic relationships drift gradually while volatility plays a secondary role, such as long-horizon monetary transmission where the Phillips curve slope and interest rate pass-through evolve over decades, or in credit spread models where sensitivities to leverage and profitability shift with market structure. An SV-VAR could be preferable when the structural coefficients remain stable, but the volatility of shocks varies substantially, as seen in exchange rate movements, stock-bond return dynamics, or uncertainty shocks, where crises generate large swings in volatility without altering underlying relations. A TVP-SV-VAR could be preferable when both types of time variation operate simultaneously, as during financial crises when monetary policy regimes shift and credit or equity market volatility spikes, or in sovereign spread dynamics during the Eurozone crisis, where both the responsiveness to fundamentals and the level of volatility change over time.

In this paper, we focus on TVP-VAR. A TVP-VAR pairs more naturally with a copula than a VAR-SV because the two components address distinct and complementary parts of the model. The TVP structure governs slow-moving changes in the autoregressive coefficients, capturing structural drift in macro-financial relationships, while the copula governs the cross-equation dependence and allows for asymmetric, nonlinear, and tail-driven co-movements in the shocks. In contrast, pairing a copula with an SV-VAR introduces redundancy: stochastic volatility already handles time-varying scale and marginal heavy tails but leaves the dependence structure Gaussian and elliptical. As a result, the copula and SV may compete to explain similar non-Gaussian features. A TVP-VAR with a copula cleanly separates the roles of mean drift and dependence flexibility, making it a better combination in settings where structural change and nonlinear dependence are the primary empirical features.

2.3 Time-Varying Parameter VAR with Copula Dependence

Following Tsionas, Izzeldin, and Trapani (2022), we now focus on a TVP-VAR specification enriched with a flexible vine-copula layer for cross-series dependence. The copula acts as an extension to the standard TVP framework, but we introduce it here because, as discussed above, the copula structure is complementary to time variation in the autoregressive coefficients. The TVP component captures slow drift in the conditional mean dynamics, while the copula captures nonlinear and potentially asymmetric dependence in the innovations. In principle, a researcher could combine TVP, stochastic volatility, and copula dependence in a single specification, although doing so would substantially increase computational cost.

2.3.1 How this improves on existing approaches

As documented by Arias, Rubio-Ramirez, and Shin (2023), the TVP-SV-VAR model in Cogley and Sargent (2005) is not order-invariant due to the lower triangular assumption. If the order of the variables in estimation are changed, the model is unstable and changes. Therefore, recent literature has been focusing on order-invariant methods which typically focus on Bayesian MCMC (Koop, Korobilis, and Pettenuzzo, 2019; Chan, Koop, and Yu, 2024).

Another way to bypass the ordering issue is the use of copula VARs. We extend the work of Tsionas, Izzeldin, and Trapani (2022). While they use a Gaussian-mixture copula, we use a vine copula, which allows for richer forms of dependence. Instead of assuming that all pairs of variables share the same elliptical dependence, each pair in the vine can take its own copula family, such as Gaussian, Student-t, Clayton, Gumbel, Joe, etc. This flexibility matters for macro-financial data, where relationships often strengthen asymmetrically in periods of market stress. For example, equity

returns and credit spreads may move together more tightly in downturns than in booms, creating left-tail dependence that a single Gaussian copula cannot capture. The vine can assign stronger tail dependence to these specific pairs while leaving others nearly independent.

Although using a vine copula increases computational demands, several features help keep it manageable. The model can be truncated to limit the depth of dependence, and the simplified vine assumption greatly reduces complexity. Because the joint likelihood decomposes into univariate and pairwise components, estimation can proceed in parallel and scales only mildly with the number of variables. This avoids the explosion of parameters and heavy matrix operations that arise with high-dimensional covariance matrices or stochastic-volatility systems. By holding pair-copula parameters fixed with respect to conditioning variables, the simplified vine remains tractable while still capturing non-Gaussian and time-varying dependence.

Overall, combining a TVP-VAR with a vine-copula layer delivers several advantages over existing methods. It removes the ordering bias inherent in Cholesky-based SV-VARs without imposing a major computational penalty, captures edge-specific asymmetric and heavy-tailed co-movements that Gaussian copulas miss, and remains easy to implement. Moreover, recent advances in fast vine fitting (Vatter and Nagler, 2025) can further reduce runtime, helping the approach remain scalable in larger systems.

2.4 Vine Copula VAR Model

2.4.1 TVP VAR

We refer to the time-varying parameter VAR in (2.4). Let $\{y_{i,t}\}_{t=1}^{T}$ denote the ith component of $\mathbf{y}_t$. For each series, we construct the lagged regression design for an AR(p) representation,

$$y_{i,t} = \beta_{i,0,t} + \beta_{i,1,t} y_{i,t-1} + \cdots + \beta_{i,p,t} y_{i,t-p} + u_{i,t}, \qquad t = p+1, \ldots, T. \tag{2.6}$$

Define the regressor vector

$$\mathbf{z}_{i,t} = \big(1,\, y_{i,t-1},\, \ldots,\, y_{i,t-p}\big)', \qquad y_{i,t} = \mathbf{z}_{i,t}' \boldsymbol{\beta}_{i,t} + u_{i,t},$$

and collect the stacked regressors as

$$Z_i = \begin{pmatrix} \mathbf{z}_{i,p+1}' \\ \vdots \\ \mathbf{z}_{i,T}' \end{pmatrix}, \qquad \mathbf{y}_i = (y_{i,p+1}, \ldots, y_{i,T})'.$$

For multivariate regressors $\mathbf{Z}_t \in \mathbb{R}^K$, a random compression matrix $\Phi \in \mathbb{R}^{k \times K}$, with $k \ll K$, is introduced to reduce dimensionality. Each entry is drawn as

$$\Phi_{ij} = \begin{cases} \pm\phi^{-1/2}, & \text{with probability } \phi, \\ 0, & \text{with probability } 1-\phi, \end{cases}$$

so that the compressed regressors become

$$\tilde{\mathbf{Z}}_t = \Phi \mathbf{Z}_t.$$

The resulting TVP regression takes the form

$$y_t = \tilde{\mathbf{Z}}_t' \boldsymbol{\beta}_t + u_t, \qquad \boldsymbol{\beta}_t = \boldsymbol{\beta}_{t-1} + \boldsymbol{\eta}_t, \quad \boldsymbol{\eta}_t \sim N(0, \delta_\beta^{-1} P_{t-1}),$$

and the parameters are updated through discounted recursive least squares. This gives us time-varying coefficient trajectories $\boldsymbol{\beta}_t$ together with the associated variance sequence σ_t^2.

2.4.2 Copula Layer for Cross-Series Dependence

In our proposed framework, we first estimate univariate TVP autoregressions for each series and then introduce joint dependence through a copula. This produces a modular "margins-then-copula" structure: the univariate TVP dynamics determine the Gaussian predictive marginals, while the copula introduces flexible Gaussian/non-Gaussian, cross-series dependence.

Copulas and Sklar's Theorem Let $X = (X_1, \ldots, X_d)$ be a continuous random vector with marginal CDFs F_j and joint distribution F. Sklar's theorem (Sklar, 1959) states that there exists a unique copula C such that

$$F(x_1, \ldots, x_d) = C(F_1(x_1), \ldots, F_d(x_d)),$$

and the transformed variables

$$U_j = F_j(X_j) \sim \text{Unif}(0, 1).$$

The associated density factorizes as

$$f(x_1, \ldots, x_d) = \left(\prod_{j=1}^{d} f_j(x_j) \right) c(F_1(x_1), \ldots, F_d(x_d)),$$

where c is the copula density. Thus, marginal behavior and dependence can be modeled entirely

separately. In particular, even if X_j is Gaussian conditional on its TVP history, the joint dependence encoded by c can be non-Gaussian, asymmetric, or heavy-tailed.

Constructing Copula Inputs from TVP Marginals For each series $j = 1, \ldots, n$, the univariate TVP regression yields a Gaussian predictive marginal:

$$y_{j,t} \mid \mathcal{F}_{t-1} \sim \mathcal{N}\big(\mu_{j,t}, \sigma_{j,t}^2\big).$$

Let $e_{j,t}$ be the one-step innovation. We obtain PIT residuals

$$u_{j,t} = \Phi\left(\frac{e_{j,t}}{\sigma_{j,t}}\right), \qquad t = 1, \ldots, T.$$

Stacking these gives the sequence $\mathbf{u}_t = (u_{1,t}, \ldots, u_{n,t})$, which is used to estimate the copula.

Vine Copulas for Non-Gaussian Dependence To capture rich cross-series dependence, we employ a simplified regular vine (R-vine) copula. A vine decomposes the joint copula density as a product of bivariate conditional copulas:

$$c(u_1, \ldots, u_n) = \prod_{e \in V} c_{a_e, b_e; D_e}\big(u_{a_e|D_e},\, u_{b_e|D_e}\big),$$

where each edge e connects variables a_e and b_e conditioned on a set D_e. Under the simplifying assumption, each pair-copula does not depend on the actual values of the conditioning set, greatly reducing computational cost.

We allow each pair-copula to be chosen from a flexible family set (Clayton, Joe, BB1, BB6, Frank, etc.), and selection is performed using the BIC criterion. Because pair-copulas can differ across edges, the vine can represent asymmetric tail dependence and nonlinear dependence patterns that a Gaussian copula cannot capture.

Importantly, the vine affects only the dependence structure–the univariate TVP marginals remain Gaussian. Thus, the overall model combines

- Gaussian predictive marginals from the TVP regressions, and
- a non-Gaussian joint dependence structure from the vine copula.

Simulation from the fitted vine copula produces joint innovations $(u_{1,t}, \ldots, u_{n,t})$ that are then transformed into VAR shocks using the inverse-PIT of the TVP marginals:

$$\varepsilon_{j,t} = \mu_{j,t} + \sigma_{j,t}\, \Phi^{-1}(u_{j,t}).$$

This yields a multivariate innovation process with flexible, edge-specific, time-varying dependence

while preserving the interpretable Gaussian TVP forecasts for each individual series.

Vine Copula VAR Algorithm

1. **Recursively expanding-window origins:** for each annual origin t_k, form $Y^{(k)} = \{Y_t^{\text{med}} : t \leq t_k\}$.
2. **Random-compression VAR design:** build VAR(p) lags (X_{full}, Y_t); standardize X_{full} and apply a sparse random matrix Φ to obtain compressed regressors $Z_t = [1,\, \Phi X_{\text{full},t}]$.
3. **TVP regressions:** for each series j, fit discounted RLS $y_{j,t} = Z_t \beta_{j,t} + \varepsilon_{j,t}$ to obtain paths $(\beta_{j,t}, \sigma^2_{j,t})$ and final states $(\beta_{j,T_k}, \sigma^2_{j,T_k})$.
4. **Rank Transformation and Copula:** compute one-step TVP predictive means/variances and PITs $u_{j,t} = \Phi(y_{j,t}; \mu_{j,t}, \sigma^2_{j,t})$; fit dependence via either (i) Gaussian copula or (ii) BIC-optimized R-vine copula on $\{u_{j,t}\}$.
5. **Forecast using Copula Shock** for $s = 1, \ldots, S$ and $h = 1, \ldots, 12$: (i) build compressed regressors from the last p simulated observations; (ii) compute marginal means $\mu_{j,t_k+h} = Z_{t_k+h} \beta_{j,T_k}$; (iii) draw a joint copula shock, convert to Gaussian, scale by $\sqrt{\sigma^2_{j,T_k}}$, and set $y^{(s)}_{j,t_k+h} = \mu_{j,t_k+h} + \varepsilon^{(s)}_{j,t_k+h}$.

2.5 Model Overview

Our framework separates univariate dynamics from cross-series dependence. Each series is first modeled with a time-varying parameter autoregression (TVP-AR). While we use discounted recursive least squares (RLS), this choice is not essential. Recent research offers several alternatives to Kalman-filter/RLS estimation. For example, (Tsionas, Izzeldin, and Trapani, 2022) use MCMC to estimate each univariate TVP equation and then combine the resulting posteriors through a copula. (Carriero, Clark, and Marcellino, 2019) use Bayesian methods for large VARs with stochastic volatility and flexible priors. (Kastner, 2019) use factor stochastic volatility with global-local shrinkage and Gibbs sampling to estimate high-dimensional, time-varying covariance structures. (Wu and Koop, 2025) use an eigendecomposition of the error covariance matrix to obtain order-invariant Bayesian inference with stochastic volatility. (Wu and Koop, 2023) use a Plackett-Luce prior and MCMC to learn variable orderings in Bayesian VARs. (Chan, Yu, and Zhang, 2024) use variational Bayes and variational importance sampling as fast approximate alternatives to MCMC. (Koop, Korobilis, and Pettenuzzo, 2019) use random compression of regressors and Bayesian model averaging across compressions. These approaches provide fully Bayesian or approximate Bayesian alternatives for the univariate estimation step.

After computing the marginal predictive distributions (which remain Gaussian in our implementation), we introduce cross-series dependence using a copula. Sklar's theorem states that any

joint distribution can be decomposed into its marginals and a copula. Unlike (Tsionas, Izzeldin, and Trapani, 2022), who use a Gaussian-mixture copula with compressed covariance matrices, we use a simplified regular vine copula selected via the greedy Kendall's τ algorithm of (Dissmann et al., 2013). This construction is order-invariant, in contrast to dependence structures based on Cholesky factorizations or triangular decompositions (e.g. Carriero, Clark, and Marcellino, 2019) and also differs from Gaussian elliptical dependence used in factor stochastic-volatility models (e.g. Kastner, 2019) or the mixture-Gaussian copula of (Tsionas, Izzeldin, and Trapani, 2022). Each variable pair can follow a distinct copula family (Clayton, Joe, Frank, Mixtures, etc.), enabling non-Gaussian, asymmetric, and tail-dependent co-movements. Vine truncation keeps the number of pair-copulas moderate, so estimation scales almost linearly with the number of series. The copula layer therefore provides a flexible, data-driven, and order-invariant representation of cross-series dependence without altering the marginal TVP forecasts, and it is simpler to implement than high-dimensional Gaussian-mixture or SV-based dependence systems.

3 Forecasting Evaluation

To assess the out-of-sample performance of the vine copula VAR, we compare them against other classes of forecasting models for large multivariate systems: Univariate TVP VAR (UTVP), dynamic factor models (DFMs), factor-augmented VARs (FAVARs), Bayesian VARs with Minnesota priors (BVARs) and Bayesian Random Compression VAR (BCVARs). As a simple baseline, we also include AR(1) forecasts. We largely follow Koop, Korobilis, and Pettenuzzo (2019) in benchmarking the forecasting.

3.0.1 Univariate TVP VAR UTVP

To properly compare against our VCVAR, we implement a TVP VAR constructed from univariate time varying AR models, without any copula addition. Let $Y_t = (y_{1,t}, \ldots, y_{n,t})'$ denote the n observed series. For each series $j = 1, \ldots, n$ and at each forecast origin, we estimate a TVP AR(p)

$$y_{j,t} = \beta_{j,t}' z_{j,t} + \varepsilon_{j,t}, \qquad \varepsilon_{j,t} \sim N(0, \sigma_{j,t}^2),$$

where $z_{j,t} = (1, y_{j,t-1}, \ldots, y_{j,t-p})'$ collects the intercept and p lags. The coefficients $\beta_{j,t}$ follow a random walk evolution implemented by discounted recursive least squares, and the innovation variance $\sigma_{j,t}^2$ is updated through exponential discounting of squared residuals. This yields filtered paths $\{\beta_{j,t}, \sigma_{j,t}^2\}$ and final states $\beta_{j,T_k}, \sigma_{j,T_k}^2$ at each expanding-window origin t_k.

Given these marginal TVP fits, we approximate one-step predictive distributions within the

estimation sample and construct probability integral transforms

$$u_{j,t} = F_{j,t}(y_{j,t}),$$

where $F_{j,t}$ is the Gaussian cdf with mean $z'_{j,t}\beta_{j,t}$ and variance $\sigma^2_{j,t}$. Stacking the PITs across series and time, we apply a normal-scores transformation and estimate a shrunk correlation matrix using Ledoit and Wolf style shrinkage. The resulting correlation estimate summarizes the contemporaneous dependence implied by the marginal models, while the point forecasts and marginal variances are determined entirely by the univariate TVP AR components.

3.0.2 Dynamic Factor Model - DFM

We implement a general dynamic factor model

$$Y_t = \lambda_0 + \lambda_1 F_t + \varepsilon_t, \qquad F_t = \Phi_1 F_{t-1} + \cdots + \Phi_p F_{t-p} + \varepsilon_t^F,$$

where F_t is a $q \times 1$ vector of common factors ($q < \sqrt{n}$), $\varepsilon_t \sim N(0, \Sigma_Y)$ with Σ_Y diagonal, and $\varepsilon_t^F \sim N(0, \Sigma_F)$. Factors are estimated by PCA . At each forecast origin we select the lag order p and number of factors q using BIC over the grids $q \leq \sqrt{n}, p \leq 13$.

3.0.3 Factor-Augmented VAR - FAVAR

We implement FAVAR following Bernanke, Boivin, and Eliasz (2005), which partitions Y_t into the m^* key variables of interest Y_t^* and the remaining variables $\tilde{Y}_t$. The measurement and state equations are

$$\tilde{Y}_t = \Lambda F_t + \varepsilon_t^{\tilde{Y}},$$

$$\begin{bmatrix} F_t \\ Y_t^* \end{bmatrix} = B_0 + B_1 \begin{bmatrix} F_{t-1} \\ Y_{t-1}^* \end{bmatrix} + \cdots + B_p \begin{bmatrix} F_{t-p} \\ Y_{t-p}^* \end{bmatrix} + \varepsilon_t^*,$$

with mutually independent Gaussian disturbances. Factors are extracted with PCA and (q, p) are chosen by BIC.

3.0.4 Bayesian VAR with Minnesota Prior - BVAR

We implement a Bayesian VAR Following on the forecast target variables Y_t^* using the Minnesota prior Bańbura, Giannone, and Reichlin (2010). For each expanding-window sample, a VAR(p) is estimated by a posterior ridge estimator, and both the lag order p and the Minnesota tightness ω are selected by BIC. To ensure a quicker implementation, we use increasing steps of 1 for the

tightness parameter and a support of 0.5 to 5.5

The overall tightness parameter is chosen from the data-driven grid

$$\omega^{-1} \in \Big\{0.5\sqrt{mp},\, 1.5\sqrt{mp},\, \ldots,\, 5.5\sqrt{mp}\Big\},$$

where m is the number of forecasted series. For each candidate (p, ω), we compute the Minnesota prior variances equation-by-equation using the AR(1) innovation scales of each series: $B_{\cdot j} \sim \mathcal{N}\big(0,\, V_j(\omega)\big)$ with $V_j(\omega)$ diagonal. The optimal (p^*, ω^*) minimizes the BIC.

$$\mathrm{BIC}(p, \omega) = -2\ell(\widehat{B}, \widehat{\Sigma}_U) + k \log(T - p), \qquad k = m^2 p + m + \tfrac{1}{2} m(m+1).$$

3.0.5 Bayesian Random Compression VAR - BCVAR

We implement a VAR following the random compression approach of Koop, Korobilis, and Pettenuzzo (2019), but in the reduced-form space of each equation's lagged regressors. Let $Y_t \in \mathbb{R}^n$ denote the medium-dimensional panel. At each forecast origin, we build the full lagged design matrix for a VAR(p), which contains $K = np$ regressors per equation.

For each random compression draw $r = 1, \ldots, R$, and for each equation i, we generate a compression matrix $\Phi_i^{(r)} \in \mathbb{R}^{m_i \times K}$ with $m_i \ll K$ and project the regressors into a low-dimensional subspace. An OLS regression of $y_{i,t}$ on the compressed regressors produces a coefficient vector, which is then lifted back to the original K-dimensional space using $\Phi_i^{(r)}$. Stacking across equations yields a full VAR coefficient matrix $B^{(r)}$ and residual covariance $\Sigma_U^{(r)}$ obtained by Gaussian MLE. A Gaussian log likelihood and corresponding BIC_r are computed for each draw.

For each draw we generate a deterministic H-step mean forecast path by iterating the VAR forward without shocks. Bayesian model averaging is performed using BIC weights

$$w_r \propto \exp\Big(-\tfrac{1}{2}(\mathrm{BIC}_r - \min_s \mathrm{BIC}_s)\Big), \qquad \sum_r w_r = 1.$$

The BCVAR forecast is the weighted average of these deterministic paths, and model-uncertainty variance is computed as the weighted second moment minus squared mean. This procedure avoids full Minnesota-prior estimation and instead averages over many randomly compressed VAR specifications, yielding a fast and scalable approximation to Bayesian model uncertainty.

3.1 Forecasting FRED Data

We perform two separate forecasts. Firstly, we use empirical macroeconomics data. We use the FRED-MD database of monthly US variables from January 1960 to December 2014 (McCracken

and Ng, 2016). We use the 2016-03 version and follow their instructions for transforming, imputing missing variables, and demeaning. After these steps, we obtain a 134-variable panel. We remove variables with excessive missing variables and document the steps in the appendix. We then adopt a recursive expanding-window forecasting scheme. The initial estimation sample runs from 1960:1 to 1987:6. Using these data, we estimate each model and generate forecasts for horizons $h = 1, 2, \ldots, 12$ for the period 1987:7 to 1988:6. At the next forecast origin, 1987:8, the sample is expanded to include the new past year's observation. All model components are re-estimated using the enlarged dataset. This procedure continues until the final origin in 2013:6, yielding a full history of real-time forecasts that mimic the information set available to an empirical forecaster.

We perform two forecasting exercises. One panel uses the full 128 variables while the other uses a reduced panel of 19 variables, to forecast seven variables of interest following Koop, Korobilis, and Pettenuzzo (2019): industrial production growth (INDPRO), the unemployment rate (UNRATE), total nonfarm employment (PAYEMS), the change in the Fed funds rate (FEDFUNDS), the change in the 10 year T-bill rate (GS10), the finished good producer price inflation (PPIFGS) and consumer price inflation (CPIAUCSL). To standardize the settings across models and reduce computational time, we choose a lag length of (p=4) with 500 runs. A full set of instructions to replicate and code are provided in Appendix B.

3.2 Forecast Accuracy Measures

We evaluate forecast accuracy using weighted forecast error differentials relative to an AR(1) benchmark. For each model i, variable j, horizon h, and forecast date t, define

$$d_{i,j,h,t} = w_{j,h}\left(e^2_{\mathrm{AR}(1),j,h,t} - e^2_{i,j,h,t}\right),$$

where $e_{i,j,h,t} = y_{j,t} - \widehat{y}_{i,j,h,t}$ is the h-step forecast error and

$$w_{j,h} = \frac{1}{\widehat{\sigma}^2_{j,h}}, \qquad \widehat{\sigma}^2_{j,h} = \widehat{\mathrm{Var}}\big(e_{\mathrm{AR}(1),j,h,\tau}\big)$$

is the inverse of the AR(1) error variance computed over the out-of-sample period for variable j and horizon h (Christoffersen and Diebold, 1998).

The cumulative sum of weighted forecast error differentials (CSWFED) for variable j is

$$\mathrm{CSWFED}_{i,j,h}(t) = \sum_{\tau \le t} d_{i,j,h,\tau}.$$

To evaluate multivariate accuracy across a set of variables J, we aggregate the *per-variable cumu-*

lative sums:

$$\mathrm{CSWFED}_{i,h}^{(\mathrm{sum})}(t) = \sum_{j \in J} \mathrm{CSWFED}_{i,j,h}(t),$$

The aggregated CSWFED curve measures, at each forecast date t, the total (or average) cumulative advantage of model i over the AR(1) benchmark across all variables and all past forecast origins. Positive values indicate that model i has produced lower weighted MSFEs than AR(1) up to date t. Density forecast accuracy is evaluated using the average log predictive likelihood differential of Geweke and Amisano (2010), which is a variance-type comparison:

$$\mathrm{ALPL}_{ijh} = \frac{1}{t - t - h + 1} \sum_{\tau=t}^{t-h} \big(\mathrm{LPL}_{i,j,\tau+h} - \mathrm{LPL}_{\mathrm{AR}(1),j,\tau+h}\big),$$

and its multivariate analogue

$$\mathrm{MVALPL}_{ih} = \frac{1}{t - t - h + 1} \sum_{\tau=t}^{t-h} \big(\mathrm{MVLPL}_{i,\tau+h} - \mathrm{MVLPL}_{\mathrm{AR}(1),\tau+h}\big).$$

Positive values indicate superior density forecasting relative to the AR(1) benchmark.

Finally, we conduct Diebold-Mariano (Diebold and Mariano, 2002) tests for equal predictive accuracy using serial-correlation-robust standard errors and the finite-sample correction of Harvey, Leybourne, and Newbold (1997). Significance at the 1%, 5%, and 10% levels is denoted by ***, **, and * respectively.

3.3 Out-of-sample forecast performance with FRED-MD

We assess the out-of-sample performance of the competing models using both point and density forecast criteria. Tables 1 and 2 report relative mean squared forecast errors (MSFEs) with respect to an AR(1) benchmark for the 19-variable and 128-variable FRED-MD panels, respectively, while Tables 3 and 4 report corresponding average log predictive likelihood (ALPL) differentials. Figures 1 and 2 complement these tables by plotting cumulative weighted forecast-error differentials (CSWFED) over time, providing a dynamic view of how forecast gains accumulate across origins and horizons.

For the 19-variable FRED-MD panel, the traditional benchmarks–DFM, FAVAR, and BVAR–already deliver sizable improvements over the AR(1) at short horizons ($h = 1, 2, 3$) for most real activity and inflation targets. However, the copula-based specifications, and in particular the VC-VAR, become increasingly dominant at medium and long horizons. In Table 1, the VCVAR typically attains the lowest or near-lowest relative MSFE for key real variables such as industrial production

(INDPRO), unemployment (UNRATE), and payroll employment (PAYEMS) once $h \geq 6$, while maintaining competitive performance at shorter horizons. The BCVAR and UTVP models also yield nontrivial gains for selected series, but their improvements are less systematic across variables and horizons than those of the VCVAR. The ALPL results in Table 3 reinforce this pattern: the VCVAR and, to a lesser extent, the UTVP model often exhibit positive and statistically significant ALPL differentials relative to AR(1), indicating that the copula-based dependence structure translates into sharper density forecasts as well as smaller squared errors.

The richer 128-variable panel amplifies these patterns. In Table 2, the relative MSFEs indicate that simply adding more predictors does not uniformly benefit all models: some linear benchmarks exhibit mixed or even deteriorating performance for certain variables and horizons. By contrast, the VCVAR continues to perform well and, for many target series at horizons $h \geq 6$, emerges as the most robust specification in terms of relative MSFE. This suggests that the flexible dependence structure embedded in the VCVAR is particularly well-suited to exploiting the additional cross-sectional information in the larger FRED-MD panel. The density-based comparisons in Table 4 are consistent with this interpretation: copula-based models, especially VCVAR, tend to deliver positive ALPL differentials at medium and long horizons, indicating that they not only reduce forecast errors on average but also assign higher probability mass to eventual realizations.

The CSWFED plots in Figures 1 and 2 summarize these out-of-sample gains over time. For both the 19-variable and 128-variable panels, the cumulative weighted forecast-error differential of the VCVAR lies above that of competing models at medium and long horizons, and increases steadily with the forecast origin once $h \geq 6$. Because CSWFED is defined relative to AR(1), positive and upward-sloping paths indicate persistent outperformance in terms of weighted mean squared error. The fact that the VCVAR paths remain positive and trend upward over multiple decades of forecast origins implies that the gains are not driven by a small number of episodes, but instead reflect a robust improvement in out-of-sample forecasting performance across macroeconomic regimes and information sets.

The ALPL results for the 128 variable FRED-MD panel in Table 4 show that density forecasting performance becomes more heterogeneous as dimensionality increases. The VCVAR model does not deliver uniform gains across all variables and forecast horizons, but it does retain several clear strengths. For real activity indicators such as `PAYEMS`, the VCVAR produces positive and often statistically significant ALPL improvements at short and medium horizons. These gains indicate that the flexible dependence structure helps capture nonlinear comovements that arise in high dimensional macroeconomic panels. The VCVAR also performs relatively well for `FEDFUNDS` at longer horizons, where time variation and tail dependence play a more important role in shaping predictive densities.

In contrast, variables related to interest rates such as GS10 and inflation measures such as CPIAUCSL exhibit more modest or mixed ALPL patterns. These series are known to be difficult to improve upon using nonlinear multivariate methods because they display strong persistence and relatively stable conditional distributions. As a result, linear benchmarks often remain competitive. Overall, while the VCVAR model does not dominate across the full 128 variable system, it delivers competitive and in several cases superior predictive densities for economically important macroeconomic series. This pattern indicates that nonlinear cross sectional dependence can yield measurable forecasting benefits even in very large information sets.

Table 5 reports relative MSFE ratios together with Diebold-Mariano significance markers for the 19-variable FRED-MD panel. Several patterns emerge. At short horizons, no single model dominates, although UTVP, DFM, and FAVAR each achieve pockets of improvement for activity variables such as INDUSTRIAL PRODUCTION and PAYEMS. The DM statistics indicate that these gains are only occasionally statistically significant, consistent with the limited cross-series information available in the small panel. For monetary policy and financial variables, the gains are more concentrated. VCVAR delivers large and statistically significant improvements for FEDFUNDS at horizons $h = 1$, $h = 2$, and $h = 3$, and BVAR and BCVAR obtain modest but sometimes significant gains for GS10 at intermediate horizons. Overall, many MSFE ratios remain close to one, and the DM tests confirm that the 19-variable environment yields only selective forecasting gains. This highlights the modest but nontrivial advantages associated with models incorporating smooth time variation or flexible dependence structures.

Expanding the information set to the full 128-variable FRED-MD panel substantially changes the relative ranking of the models. Table 6 shows that the richer environment amplifies the benefits of cross-series pooling, with many more MSFE differences now flagged as statistically significant by the DM test. For real activity indicators, DFM and BCVAR generate statistically significant gains at short and medium horizons, while VCVAR consistently improves long-horizon forecasts. These gains are markedly stronger than in the 19-variable case, confirming that the larger panel provides exploitable covariance structure that the AR(1) benchmark cannot capture. Monetary and financial variables show an even sharper improvement. UTVP remains competitive at $h = 1$ for FEDFUNDS, but from $h = 3$ onward, DFM, BCVAR, and VCVAR deliver substantial and statistically significant accuracy gains, dominating the remaining alternatives. Long-term rates such as GS10 also benefit from the data-rich setting, with DFM performing well at $h = 2$ and $h = 3$, and BCVAR and VCVAR achieving the strongest performance for horizons $h \geq 6$. While price indices remain comparatively difficult to forecast, even here the DM results identify incremental improvements not visible in the smaller panel. Taken together, the 128-variable results show that data-rich environments meaningfully strengthen the relative advantage of models such as DFM,

BCVAR, and VCVAR, especially at medium and long horizons where pooling high-dimensional macroeconomic information produces the greatest forecasting gains.

4 Forecasting Simulated Data

4.1 Simulated TVP-AR(1) Panel with Stochastic Volatility and Regime-Switching Dependence

To complement the empirical FRED-MD results, we also evaluate all forecasting models in a controlled high-dimensional environment. We simulate a panel of 100 monthly variables over $T = 1000$ periods from a rich data-generating process that incorporates three key features commonly observed in macro-financial datasets: (i) time-varying persistence, (ii) stochastic volatility, and (iii) regime-dependent cross-sectional dependence. Each series follows a time-varying AR(1) process, where the autoregressive coefficient evolves smoothly over time and is restricted to the stationary region. Residual volatilities are driven by persistent log-variance processes, generating stochastic volatility and heteroskedasticity.

Cross-sectional dependence is introduced through a two-state Markov regime-switching structure. In the "calm" regime, shocks are drawn from a multivariate normal distribution with a moderate correlation matrix. In the "crisis" regime, innovations become both more strongly correlated and heavy-tailed through a scale-mixture construction, producing tail dependence and joint extreme realizations. The combination of drifting persistence, stochastic volatility, and episodic increases in correlation generates a flexible synthetic environment resembling macro-financial stress periods.

We evaluate forecasting performance for the first ten simulated variables. To mimic the empirical expanding-window setting, we generate real-time forecasts using a grid of estimation endpoints $\{500, 600, 700, 800, 900\}$ and horizons $\{1, 3, 5, 10, 25, 50\}$. At each endpoint, models are estimated using only data available up to that date, and multi-step-ahead forecasts are produced for all ten variables. This design allows us to assess how forecasting accuracy varies with sample size, forecast horizon, and underlying dependence conditions. Full implementation details and the complete simulation code appear in Appendix B.

4.2 Out-of-sample forecast performance with simulation data

Figure 3 summarizes the cumulative weighted forecast-error differentials (CSWFED) for the 100-variable simulation across horizons. The patterns closely mirror the structure of the data-generating process. At very short horizons ($h = 1$), none of the models delivers systematic gains, although the copula-based VCVAR exhibits pronounced underperformance due to the difficulty of exploiting

nonlinear dependence at such a short forecasting window. Beginning at $h = 3$ and $h = 5$, however, the relative ranking shifts: the UTVP model accumulates increasingly positive differentials, indicating that time-varying marginal dynamics provide meaningful short- and medium-horizon improvements when the underlying coefficients drift. At longer horizons ($h = 10$ and beyond), the VCVAR becomes the most robust performer. Its CSWFED paths turn positive and trend upward, surpassing all linear benchmarks as the horizon increases. This reflects the fact that nonlinear and tail dependence–features explicitly encoded in the simulated shocks–contribute more heavily to forecast performance at longer horizons, and the vine-copula layer of the VCVAR exploits this structure. By contrast, the linear models accumulate negative or near-zero differentials as h grows, underscoring their inability to capture the multivariate dependence embedded in the DGP. Taken together, the CSWFED curves demonstrate that in a high-dimensional setting with time variation and non-Gaussian dependence, the copula-based VAR yields persistent long-horizon gains, while UTVP remains competitive primarily at shorter horizons.

For the 100-variable simulation in Table 7, the MSFE results align with the structure of the data-generating process. At short horizons, all models perform similarly because forecast errors are dominated by local AR dynamics. Differences emerge as the horizon increases. The UTVP model improves over the linear benchmarks at $h = 3$ and $h = 5$, reflecting its ability to track coefficient drift. At longer horizons ($h \geq 10$), the VCVAR becomes the strongest performer for most series. These gains arise because the copula layer captures the nonlinear cross-sectional dependence built into the simulated shocks, which the linear models cannot represent. Across the system, MSFE levels show that UTVP is competitive only at short horizons, while the VCVAR consistently reduces forecast errors at medium and long horizons in a high-dimensional, non-Gaussian environment.

Table 8 reports ALPL differentials for the 100-variable simulation. At short horizons ($h = 1$ and $h = 3$), the UTVP model delivers the strongest density forecasts because it adapts quickly to time-varying marginal dynamics. The linear benchmarks (DFM, FAVAR, and BVAR) show mixed results, with small gains for some series but many differentials near zero or negative.

At medium horizons ($h = 5$ and $h = 10$), VCVAR produces the highest ALPL values for many variables, with several significant improvements (such as `X5`, `X7`, and `X9`). These gains reflect the ability of the vine copula to capture nonlinear and tail dependence in the simulated shocks. At longer horizons ($h = 25$ and $h = 50$), forecast accuracy declines for all models, but VCVAR remains competitive for variables with strong nonlinear dependence, while the linear benchmarks often deteriorate. Overall, UTVP dominates at short horizons, whereas VCVAR becomes the most effective density forecaster once cross-sectional dependence becomes more important.

Table 9 reports MSFE ratios with Diebold-Mariano significance for the 100-variable simulation. At short horizons ($h = 1$ and $h = 3$), UTVP performs best, with several ratios below one and a few

significant improvements (for example, `X1` at $h = 3$). The linear benchmarks show mixed results: they achieve isolated significant gains but many ratios remain close to or above one.

At medium horizons ($h = 5$ and $h = 10$), VCVAR begins to outperform the linear models for many series, although not always at a statistically significant level. Its advantages are most visible for variables driven by nonlinear dependence, such as `X5`, `X7`, and `X9`. Variables dominated by near-unit-root behavior or weak multivariate dependence (for example, `X2` and `X4`) show less consistent patterns across models. At long horizons ($h = 25$ and $h = 50$), performance across models converges, and DM significance becomes rare. Overall, UTVP is strongest at short horizons, and VCVAR provides the most gains at medium horizons, especially for series affected by nonlinear dependence, but no single model dominates in every case.

5 Conclusion

This paper develops a simple and scalable way to introduce non-Gaussian, asymmetric dependence into large time-varying VARs. We combine a univariate TVP backbone, estimated by discounted recursive least squares, with a modular vine copula layer that captures nonlinear and tail dependence across series while remaining order invariant and computationally tractable. This design separates the roles of mean dynamics and cross-sectional dependence. The TVP component tracks gradual changes in macro-financial relationships, and the copula component shapes the joint distribution of shocks without relying on high-dimensional covariance matrices or heavy MCMC machinery. In this sense, the VCVAR framework sits between traditional Gaussian large VARs and fully Bayesian TVP-SV-VARs with rich covariance structures and provides a practical tool for applied forecasting when both structural drift and non-Gaussian dependence matter.

The empirical and simulation results show that this additional dependence flexibility has practical value. In FRED-MD data, copula-based specifications, and VCVAR in particular, match or improve on standard benchmarks for point forecasts and density forecasts, with gains that are most pronounced at medium and long horizons and in larger panels where cross-series information is richer. In the simulated 100-variable system, where the data-generating process features time-varying persistence, stochastic volatility, and regime-dependent tail dependence, VCVAR delivers lower MSFEs and higher ALPLs at longer horizons, while UTVP remains competitive only at short horizons. These findings suggest that practitioners who face data with evolving relationships and episodic stress can benefit from adding a copula layer on top of relatively simple TVP marginals. Future work could extend the framework to structural identification, alternative copula families, or joint modeling of time variation in both volatilities and copula parameters, as well as explore fast approximate inference schemes that preserve scalability in even higher-dimensional settings.

References

Arias, J. E., Rubio-Ramirez, J. F., Shin, M., 2023. Macroeconomic forecasting and variable ordering in multivariate stochastic volatility models. Journal of Econometrics 235, 1054–1086, publisher: Elsevier.

Bańbura, M., Giannone, D., Reichlin, L., 2010. Large Bayesian vector auto regressions. Journal of Applied Econometrics 25, 71–92.

Bedford, T., Cooke, R. M., 2001. Probability Density Decomposition for Conditionally Dependent Random Variables Modeled by Vines. Annals of Mathematics and Artificial Intelligence 32, 245–268.

Bernanke, B. S., Boivin, J., Eliasz, P., 2005. Measuring the effects of monetary policy: a factor-augmented vector autoregressive (FAVAR) approach. Quarterly Journal of Economics 120, 387–422, publisher: MIT Press.

Carriero, A., Clark, T. E., Marcellino, M., 2019. Large Bayesian vector autoregressions with stochastic volatility and non-conjugate priors. Journal of Econometrics 212, 137–154, publisher: Elsevier.

Chan, J. C. C., Koop, G., Yu, X., 2024. Large Order-Invariant Bayesian VARs with Stochastic Volatility. Journal of Business & Economic Statistics 42, 825–837, publisher: ASA Website _eprint: https://doi.org/10.1080/07350015.2023.2252039.

Chan, J. C. C., Yu, X., Zhang, W., 2024. Bayesian Model Comparison for Large Bayesian Vars after the Covid-19 Pandemic.

Christoffersen, P. F., Diebold, F. X., 1998. Cointegration and Long-Horizon Forecasting. Journal of Business & Economic Statistics 16, 450–456.

Cogley, T., Sargent, T. J., 2005. Drifts and volatilities: monetary policies and outcomes in the post WWII US. Review of Economic Dynamics 8, 262–302, publisher: Elsevier.

Diebold, F. X., Mariano, R. S., 2002. Comparing Predictive Accuracy. Journal of Business & Economic Statistics 20, 134–144.

Dissmann, J., Brechmann, E. C., Czado, C., Kurowicka, D., 2013. Selecting and estimating regular vine copulae and application to financial returns. Computational Statistics & Data Analysis 59, 52–69, publisher: Elsevier.

Geweke, J., Amisano, G., 2010. Comparing and evaluating Bayesian predictive distributions of asset returns. International Journal of Forecasting 26, 216–230, publisher: Elsevier.

Harvey, D., Leybourne, S., Newbold, P., 1997. Testing the equality of prediction mean squared errors. International Journal of Forecasting 13, 281–291, publisher: Elsevier.

Kastner, G., 2019. Sparse Bayesian time-varying covariance estimation in many dimensions. Journal of Econometrics 210, 98–115.

Koop, G., Korobilis, D., Pettenuzzo, D., 2019. Bayesian compressed vector autoregressions. Journal of Econometrics 210, 135–154.

Lütkepohl, H., 2005. New Introduction to Multiple Time Series Analysis. Springer.

McCracken, M. W., Ng, S., 2016. FRED-MD: A Monthly Database for Macroeconomic Research. Journal of Business & Economic Statistics 34, 574–589.

Nagler, T., 2025. Simplified vine copula models: state of science and affairs. Risk Sciences p. 100022, publisher: Elsevier.

Sims, C. A., 1980. Macroeconomics and reality. Econometrica: Journal of the Econometric Society pp. 1–48, publisher: JSTOR.

Sklar, M., 1959. Fonctions de répartition à n dimensions et leurs marges. In: *Annales de l'ISUP*, vol. 8, pp. 229–231.

Smith, M. S., Klein, N., 2021. Bayesian inference for regression copulas. Journal of Business & Economic Statistics 39, 712–728.

Tsionas, M. G., Izzeldin, M., Trapani, L., 2022. Estimation of large dimensional time varying VARs using copulas. European Economic Review 141, 103952.

Vatter, T., Nagler, T., 2025. Throwing Vines at the Wall: Structure Learning via Random Search. ArXiv:2510.20035.

Wu, P., Koop, G., 2023. Estimating the ordering of variables in a VAR using a Plackett-Luce prior. Economics Letters 230, 111247.

Wu, P., Koop, G., 2025. Fast, order-invariant bayesian inference in vars using the eigendecomposition of the error covariance matrix. Journal of Business & Economic Statistics pp. 1–21.

Figure 1: **CSWFED for the 19-Variable FRED-MD Panel**

This figure reports cumulative weighted forecastâ€“error differentials (CSWFED) relative to the AR(1) benchmark for the 19-variable FRED-MD panel. Each panel corresponds to a forecast horizon $h \in \{1, 2, 3, 6, 9, 12\}$ and shows the cumulative sum of weighted squaredâ€“error differences across the seven target variables. Positive values indicate superior forecasting performance relative to AR(1). The VCVAR model delivers the most pronounced gains at medium- and long-horizon forecasts ($h \geq 6$).

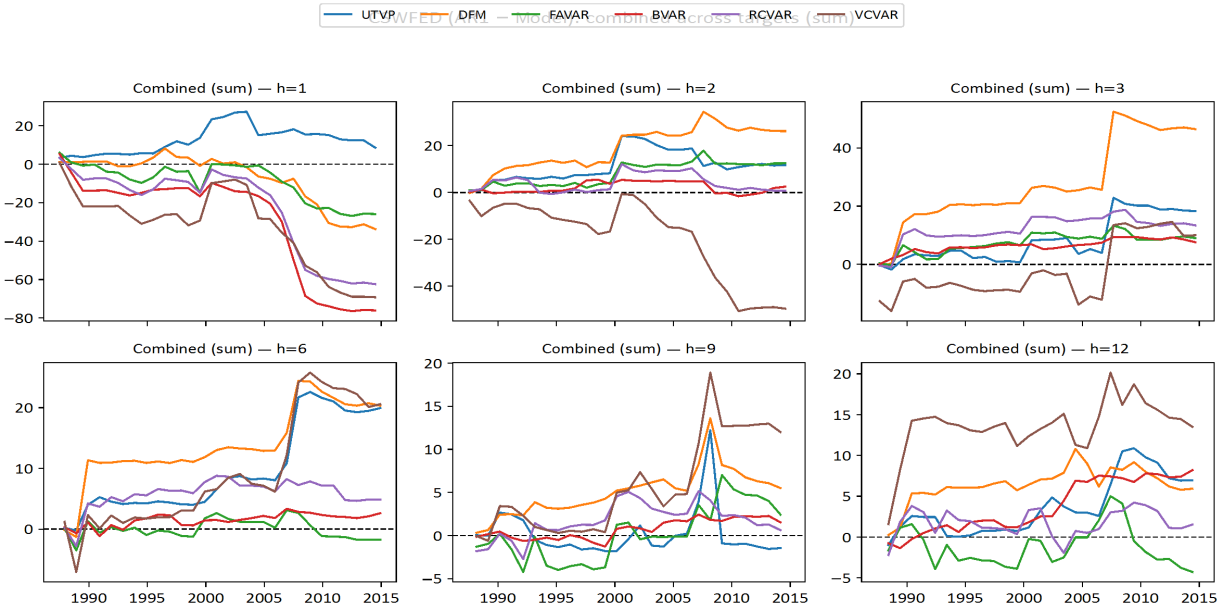

Figure 2: **CSWFED for the 128-Variable FRED-MD Panel**

This figure reports cumulative weighted forecastâ€“error differentials (CSWFED) relative to the AR(1) benchmark for the full 128-variable FRED-MD panel. Each panel corresponds to a forecast horizon $h \in \{1, 2, 3, 6, 9, 12\}$ and shows the cumulative sum of weighted squaredâ€“error differences across the seven target variables. Positive values indicate superior forecasting performance relative to AR(1). The VCVAR model continues to deliver the most pronounced gains at long-horizon forecasts ($h \geq 6$).

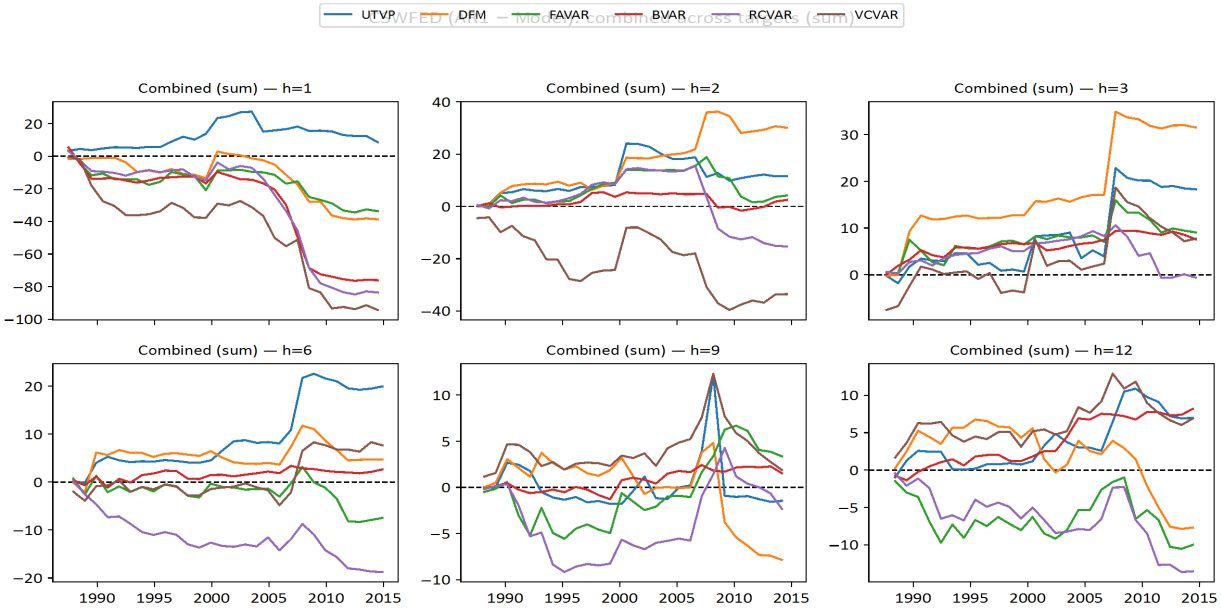

Table 1: Relative MSFE by Variable, Horizon, and Model (19-Variable FRED-MD)

	UTVP	DFM	FAVAR	BVAR	BCVAR	VCVAR
INDPRO						
h=1	0.959	0.939	1.032	1.145	1.171	1.255
h=2	1.049	0.858	0.775*	0.979	0.676**	1.271
h=3	0.900	0.619	0.915	1	0.805	0.786
h=6	0.949	0.948	1.155	0.998	1.063	0.955
h=9	0.936	1.118	1.051	0.999	1.094	0.961
h=12	1.010	1.013	1.041	1	1.057	1.189
UNRATE						
h=1	0.958	0.981	1.113	1.008	1.051	1.259
h=2	0.794	0.708	0.885	0.893	1.060	1.107
h=3	0.928	0.775	0.895	0.913	0.954	0.886
h=6	0.944	0.842	0.989	0.980	0.979	0.842
h=9	0.999	0.878	0.953	1.014	0.965	0.850
h=12	1.020	0.903***	0.873*	0.927	0.859***	0.949
PAYEMS						
h=1	0.550**	0.939	1.374	2.077	1.362	0.604
h=2	0.592	0.597	1.008	1.479	0.860	0.759
h=3	0.486	0.622	1.090	1.206	0.998	0.541
h=6	0.486*	0.645	1.121	1.026	0.975	0.499
h=9	0.999	0.832	1.048	0.996*	1.069	0.734
h=12	0.607**	0.876	1.151	0.990***	1.055	0.630
FEDFUNDS						
h=1	0.787	1.429	1.202	1.420	1.490	1.793
h=2	0.901	0.677	0.797	0.546	0.941	0.779
h=3	0.901	0.604*	0.758*	0.631**	0.692**	0.967
h=6	0.972	0.834	0.801*	0.849	0.866*	0.805
h=9	1.120	0.951	0.820	0.894	0.862	1.013
h=12	1.417	1.076	0.984	0.849	0.956	1.250
GS10						
h=1	0.991	1.004	1.059	1.088	1.091	1.548
h=2	0.938	1.094	1.121	0.989	1.343	2.086
h=3	1.034	0.835*	1.039	1.042	1.218	1.203
h=6	1.013	1.072	1.002	1.048	0.990	1.264
h=9	1.016	1.022	1.016	1.041	0.973	0.952
h=12	0.977	1.005	1.003	0.947	0.990	0.877
WPSFD49207						
h=1	1.342	1.678	1.040	1.783	1.833	1.388
h=2	1.128	1.145	1.003	1.041	1.002	1.373
h=3	1.067	0.976	1.010	0.960	0.961	1.035
h=6	1.004	1.003	0.998	1	0.992	0.994
h=9	0.995	1.009	1.001	1	0.993	1.065
h=12	1.001	1.003	0.999	1	1.012	0.933**
CPIAUCSL						
h=1	1.093	1.201	1.083	1.104	1.144	1.440
h=2	1.227	1.078	1.001	1.023	1.092	1.376
h=3	1.070	0.960	0.980	0.985	0.917	1.254
h=6	1.000	0.987	0.979*	1.001	0.968**	0.987
h=9	0.991	1.012	1.011	1	1.007	1.025
h=12	1.010	0.999	0.995	1	0.977*	0.973

Table 2: Relative MSFE by Variable, Horizon, and Model (128-Variable FRED-MD)

	UTVP	DFM	FAVAR	BVAR	BCVAR	VCVAR
INDPRO						
h=1	0.959	0.988	0.914	1.145	0.946	0.976
h=2	1.049	0.775*	0.939	0.979	0.847	1.022
h=3	0.900	0.747*	0.920	1	0.991	0.884
h=6	0.949	0.964	1.043	0.998	1.240	0.930
h=9	0.936	1.083	0.861	0.999	0.995	0.997
h=12	1.010	1.065	1.059	1	1.068	1.086
UNRATE						
h=1	0.958	1.096	1.292	1.008	1.059	1.252
h=2	0.794	0.809	0.926	0.893	0.941	0.957
h=3	0.928	0.852	0.880	0.913	0.980	1.176
h=6	0.944	1.005	1.059	0.980	1.051	0.852
h=9	0.999	1.012	1.003	1.014	0.954	0.936
h=12	1.020	0.914**	0.919	0.927	0.906	0.962
PAYEMS						
h=1	0.550**	0.828	1.346	2.077	1.756	0.839
h=2	0.592	0.670*	1.167	1.479	1.267	0.758
h=3	0.486	0.731	1.075	1.206	1.074	0.619
h=6	0.486*	0.859	1.179	1.026	1.194	0.615
h=9	0.999	1.109	0.993	0.996*	1.007	0.751
h=12	0.607**	1.050	1.239	0.990***	1.154	0.740*
FEDFUNDS						
h=1	0.787	1.577	1.500	1.420	1.717	2.267
h=2	0.901	0.608	0.829	0.546	1.131	0.801
h=3	0.901	0.679	0.802	0.631**	0.887	0.677
h=6	0.972	0.905	0.825	0.849	1.111	1.103
h=9	1.120	1.087	0.917	0.894	1.005	1.141
h=12	1.417	1.250	0.995	0.849	1.143	1.105
GS10						
h=1	0.991	1.058	1.066	1.088	1.335	1.435
h=2	0.938	1.109	1.009	0.989	1.088	1.775
h=3	1.034	0.954	1.014	1.042	0.946	1.262
h=6	1.013	1.122	1.133	1.048	1.160	1.410
h=9	1.016	0.982	1.087	1.041	1.092	1.055
h=12	0.977	0.964	0.993	0.947	1.084	1.012
WPSFD49207						
h=1	1.342	1.733	1.004	1.783	1.816	1.865
h=2	1.128	1.039	1.008	1.041	1.078	1.321
h=3	1.067	0.977	1.025	0.960	0.999	1.040
h=6	1.004	1.002	0.997	1	0.971	0.921
h=9	0.995	0.998	0.999	1	1.034	1.060
h=12	1.001	1	0.999	1	1.013	1.056
CPIAUCSL						
h=1	1.093	1.048	1.070	1.104	1.219	1.573
h=2	1.227	1.020	0.992	1.023	1.127	1.538
h=3	1.070	0.972	0.969	0.985	1.142	1.110
h=6	1.000	0.998	1.004	1.001	0.915**	0.964
h=9	0.991	1.002	1.017	1	0.997	1.032
h=12	1.010	1.008	0.986	1	1.005	0.978

Table 3: Relative ALPL by Variable, Horizon, and Model (19-Variable FRED-MD)

	UTVP	DFM	FAVAR	BVAR	BCVAR	VCVAR
INDPRO						
h=1	-0.27	-0.071	-0.109	-0.073	-0.382	-0.476
h=2	-0.026	0.118	0.121***	0.027*	0.17***	-0.12
h=3	0.089	0.225***	0.103**	0.004	0.099**	0.003
h=6	0.046	0.082	-0.026	0	-0.003	-0.098
h=9	0.155***	0.023	0.021	-0.01	-0.019	-0.044
h=12	-0.08	0.012	-0.006	-0.015	-0.015	-0.145
UNRATE						
h=1	0.095**	0.072	0.033	0.091	0.049	-0.041
h=2	0.168***	0.194**	0.082*	0.085*	-0.002	-0.025
h=3	0.094	0.121	0.051	0.048	0.026	0.06
h=6	0.066	0.093	0.004	0.035	0.017	-0.016
h=9	-0.09	0.061	0.026	-0.04	0.013	-0.025
h=12	-0.03	0.085**	0.064**	0.067*	0.053***	-0.134
PAYEMS						
h=1	0.412***	0.09	-0.025	-0.258	-0.02	0.311***
h=2	0.358***	0.239***	0.091	-0.129	0.158*	0.23**
h=3	0.27	0.268***	-0.002	-0.073	0.042	0.247***
h=6	0.199	0.225***	-0.051	-0.009	0.025	0.206
h=9	-0.202	0.171**	0.032	-0.014	0.026	0.039
h=12	0.243	0.132**	-0.02	-0.023	0.006	-0.011
FEDFUNDS						
h=1	0.654***	-0.03	0.049	0.108***	0.012	0.432***
h=2	0.303	0.131**	0.094***	0.181***	0.072***	0.412***
h=3	0.418**	0.141***	0.062***	0.161***	0.094***	0.37***
h=6	-0.68	0.1***	0.086***	0.101***	0.048***	0.191*
h=9	0.325	0.11***	0.003	0.102***	0.034***	0.189**
h=12	0.602***	0.107***	-0.031	0.11***	0.023***	0.182***
GS10						
h=1	0.049	-0.006	0.03	0.018	-0.002	-0.166
h=2	0.132***	-0.012	0.002	0.084***	-0.032	-0.193
h=3	0.048	0.078***	-0.016	0.035	-0.038	-0.113
h=6	0.084**	0.008	0.014	0.041*	0.001	-0.158
h=9	-0.023	0.007	0	0.002	0.01	-0.083
h=12	0.054	0.016	-0.015	0.04**	0.005	-0.06
WPSFD49207						
h=1	-0.032	-0.323	-0.036	-0.334	-0.367	-0.027
h=2	0.071	-0.165	0.027	0.025	0.033	0.13
h=3	-0.043	-0.094	-0.016	0.022	0.035	0.062
h=6	-0.235	-0.092	0.015	0.003	0.007	0.109
h=9	0.037	-0.061	-0.002	0.009**	0.004	-0.1
h=12	-0.024	-0.11	-0.01	0.009	-0.013	0.21
CPIAUCSL						
h=1	0.027	-0.135	-0.065	-0.062	-0.066	-0.151
h=2	-0.181	-0.15	-0.003	0.014	-0.097	-0.007
h=3	0.063	-0.025	-0.002	0.025	0.058	-0.006
h=6	-0.103	-0.007	0.015	0.003	0.021**	-0.049
h=9	-0.307	-0.021	-0.009	-0.002	-0.004	-0.247
h=12	0.108	0.001	0	0.001	0.008	-0.001

Table 4: Relative ALPL by Variable, Horizon, and Model (128-Variable FRED-MD)

	UTVP	DFM	FAVAR	BVAR	BCVAR	VCVAR
INDPRO						
h=1	-0.27	-0.127	-0.071	-0.073	-0.029	-0.145
h=2	-0.026	0.144***	0.061	0.027*	0.11*	-0.021
h=3	0.089	0.167***	0.069*	0.004	0.052	0.039
h=6	0.046	0.069*	0.019	0	-0.051	-0.047
h=9	0.155***	0.046	0.071**	-0.01	0.019	-0.048
h=12	-0.08	-0.004	-0.029	-0.015	-0.013	-0.054
UNRATE						
h=1	0.095**	0.04	-0.029	0.091	0.041	-0.036
h=2	0.168***	0.125**	0.048	0.085*	0.056	0.033
h=3	0.094	0.086	0.063	0.048	0.006	-0.074
h=6	0.066	0.031	-0.016	0.035	-0.012	0.015
h=9	-0.09	-0.009	0.007	-0.04	0.027	-0.008
h=12	-0.03	0.069*	0.046*	0.067*	0.018	-0.104
PAYEMS						
h=1	0.412***	0.15***	-0.01	-0.258	-0.113	0.206**
h=2	0.358***	0.265***	0.035	-0.129	0.017	0.285***
h=3	0.27	0.196***	0.003	-0.073	0.008	0.326***
h=6	0.199	0.073	-0.069	-0.009	-0.094	0.218
h=9	-0.202	0.021	0.04	-0.014	0.028	0.033
h=12	0.243	0.014	-0.06	-0.023	-0.043	0.019
FEDFUNDS						
h=1	0.654***	-0.006	0.075**	0.108***	-0.031	0.241*
h=2	0.303	0.143**	0.037	0.181***	0.024	0.366***
h=3	0.418**	0.163***	0.029	0.161***	0.061**	0.344***
h=6	-0.68	0.081**	0.057**	0.101***	0.002	0.019
h=9	0.325	0.146***	0.027*	0.102***	-0.002	0.051
h=12	0.602***	0.097***	-0.006	0.11***	-0.015	0.105**
GS10						
h=1	0.049	-0.001	0.017	0.018	-0.078	-0.137
h=2	0.132***	0.036	-0.024	0.084***	-0.002	-0.118
h=3	0.048	0.07**	-0.024	0.035	0.01	-0.086
h=6	0.084**	0.022	-0.007	0.041*	-0.04	-0.171
h=9	-0.023	0.021	-0.032	0.002	-0.036	-0.101
h=12	0.054	0.037**	-0.005	0.04**	-0.038	-0.057
WPSFD49207						
h=1	-0.032	-0.317	-0.009	-0.334	-0.383	-0.26
h=2	0.071	-0.038	-0.013	0.025	-0.047	0.102
h=3	-0.043	0.033	0.012	0.022	0.026	0.206
h=6	-0.235	-0.11	0.068	0.003	0.09	0.142
h=9	0.037	0.026***	-0.001	0.009**	0.008	-0.13
h=12	-0.024	-0.023	-0.032	0.009	0.02	0.284
CPIAUCSL						
h=1	0.027	0.033	-0.072	-0.062	-0.123	-0.157
h=2	-0.181	-0.053	-0.016	0.014	-0.134	-0.252
h=3	0.063	0.021	0.019	0.025	-0.088	0.061
h=6	-0.103	-0.024	0.007	0.003	0.058	-0.033
h=9	-0.307	-0.001	-0.01	-0.002	-0.026	-0.211
h=12	0.108	0.001	0.002	0.001	-0.031	-0.041

Table 5: Diebold–Mariano Test Statistics for Relative MSFE, by Variable and Horizon (19-Variable FRED-MD Panel)

	UTVP	DFM	FAVAR	BVAR	BCVAR	VCVAR
INDPRO						
h=1	0.959	0.939	1.032	1.145	1.171	1.255
h=2	1.049	0.858	0.775	0.979	0.676*	1.271
h=3	0.900	0.619	0.915	1.000	0.805	0.786
h=6	0.949	0.948	1.155	0.998	1.063	0.955
h=9	0.936	1.118	1.051	0.999	1.094	0.961
h=12	1.010	1.013	1.041	1.000	1.057	1.189
UNRATE						
h=1	0.958	0.981	1.113	1.008	1.051	1.259
h=2	0.794	0.708	0.885	0.893	1.060	1.107
h=3	0.928	0.775*	0.895	0.913	0.954	0.886
h=6	0.944	0.842	0.989	0.980	0.979	0.842
h=9	0.999	0.878	0.953	1.014	0.965	0.850
h=12	1.020	0.903	0.873**	0.927	0.859***	0.949
PAYEMS						
h=1	0.550**	0.939	1.374	2.077	1.362	0.604
h=2	0.592	0.597	1.008	1.479	0.860	0.759
h=3	0.486	0.622	1.090	1.206	0.998	0.541
h=6	0.486*	0.645	1.121	1.026	0.975	0.499
h=9	0.999	0.832	1.048	0.996*	1.069	0.734
h=12	0.607*	0.876	1.151	0.990***	1.055	0.630
FEDFUNDS						
h=1	0.787	1.429	1.202	1.420	1.490	1.793
h=2	0.901	0.677	0.797	0.546	0.941	0.779
h=3	0.901	0.604*	0.758**	0.631**	0.692*	0.967
h=6	0.972	0.834	0.801**	0.849	0.866*	0.805
h=9	1.120	0.951	0.820	0.894	0.862	1.013
h=12	1.417	1.076	0.984	0.849	0.956	1.250
GS10						
h=1	0.991	1.004	1.059	1.088	1.091	1.548
h=2	0.938	1.094	1.121	0.989	1.343	2.086
h=3	1.034	0.835	1.039	1.042	1.218	1.203
h=6	1.013	1.072	1.002	1.048	0.990	1.264
h=9	1.016	1.022	1.016	1.041	0.973	0.952
h=12	0.977	1.005	1.003	0.947	0.990	0.877
WPSFD49207						
h=1	1.342	1.678	1.040	1.783	1.833	1.388
h=2	1.128	1.145	1.003	1.041	1.002	1.373
h=3	1.067	0.976	1.010	0.960	0.961	1.035
h=6	1.004	1.003	0.998	1.000	0.992	0.994
h=9	0.995	1.009	1.001	1.000	0.993	1.065
h=12	1.001	1.003	0.999	1.000	1.012	0.933
CPIAUCSL						
h=1	1.093	1.201	1.083	1.104	1.144	1.440
h=2	1.227	1.078	1.001	1.023	1.092	1.376
h=3	1.070	0.960	0.980	0.985	0.917	1.254
h=6	1.000	0.987	0.979*	1.001	0.968**	0.987
h=9	0.991	1.012	1.011	1.000	1.007	1.025
h=12	1.010	0.999	0.995	1.000	0.977	0.973

Table 6: Diebold–Mariano Test Statistics for Relative MSFE, by Variable and Horizon (128-Variable FRED-MD Panel)

	UTVP	DFM	FAVAR	BVAR	BCVAR	VCVAR
INDPRO						
h=1	0.959	0.988	0.914	1.145	0.946	0.976
h=2	1.049	0.775	0.939	0.979	0.847	1.022
h=3	0.900	0.747**	0.920	1.000	0.991	0.884
h=6	0.949	0.964	1.043	0.998	1.240	0.930
h=9	0.936	1.083	0.861	0.999	0.995	0.997
h=12	1.010	1.065	1.059	1.000	1.068	1.086
UNRATE						
h=1	0.958	1.096	1.292	1.008	1.059	1.252
h=2	0.794	0.809	0.926	0.893	0.941	0.957
h=3	0.928	0.852*	0.880	0.913	0.980	1.176
h=6	0.944	1.005	1.059	0.980	1.051	0.852
h=9	0.999	1.012	1.003	1.014	0.954	0.936
h=12	1.020	0.914	0.919	0.927	0.906	0.962
PAYEMS						
h=1	0.550**	0.828	1.346	2.077	1.756	0.839
h=2	0.592	0.670*	1.167	1.479	1.267	0.758
h=3	0.486	0.731	1.075	1.206	1.074	0.619
h=6	0.486*	0.859	1.179	1.026	1.194	0.615
h=9	0.999	1.109	0.993	0.996*	1.007	0.751
h=12	0.607*	1.050	1.239	0.990***	1.154	0.740
FEDFUNDS						
h=1	0.787	1.577	1.500	1.420	1.717	2.267
h=2	0.901	0.608	0.829	0.546	1.131	0.801
h=3	0.901	0.679	0.802	0.631**	0.887	0.677
h=6	0.972	0.905	0.825	0.849*	1.111	1.103
h=9	1.120	1.087	0.917	0.894	1.005	1.141
h=12	1.417	1.250	0.995	0.849	1.143	1.105
GS10						
h=1	0.991	1.058	1.066	1.088	1.335	1.435
h=2	0.938	1.109	1.009	0.989	1.088	1.775
h=3	1.034	0.954	1.014	1.042	0.946	1.262
h=6	1.013	1.122	1.133	1.048	1.160	1.410
h=9	1.016	0.982	1.087	1.041	1.092	1.055
h=12	0.977	0.964	0.993	0.947	1.084	1.012
WPSFD49207						
h=1	1.342	1.733	1.004	1.783	1.816	1.865
h=2	1.128	1.039	1.008	1.041	1.078	1.321
h=3	1.067	0.977	1.025	0.960	0.999	1.040
h=6	1.004	1.002	0.997	1.000	0.971	0.921
h=9	0.995	0.998	0.999	1.000	1.034	1.060
h=12	1.001	1.000	0.999	1.000	1.013	1.056
CPIAUCSL						
h=1	1.093	1.048	1.070	1.104	1.219	1.573
h=2	1.227	1.020	0.992	1.023	1.127	1.538
h=3	1.070	0.972	0.969	0.985	1.142	1.110
h=6	1.000	0.998	1.004	1.001	0.915**	0.964
h=9	0.991	1.002	1.017	1.000	0.997	1.032
h=12	1.010	1.008	0.986	1.000	1.005	0.978

Figure 3: **Cumulative Weighted Forecastâ€“Error Differentials (CSWFED) for the 100-Variable Simulation Panel**

This figure reports cumulative weighted forecastâ€“error differentials relative to an AR(1) benchmark for the 100-variable simulated panel. Each panel corresponds to a forecast horizon $h \in \{1, 3, 5, 10, 25, 50\}$ and shows the cumulative sum of weighted squaredâ€“error differentials across the forecast target variables. Positive values indicate superior forecasting performance relative to the AR(1) benchmark. The VCVAR model delivers the most pronounced gains at h=10 and h=50 forecasts ($h \geq 6$).

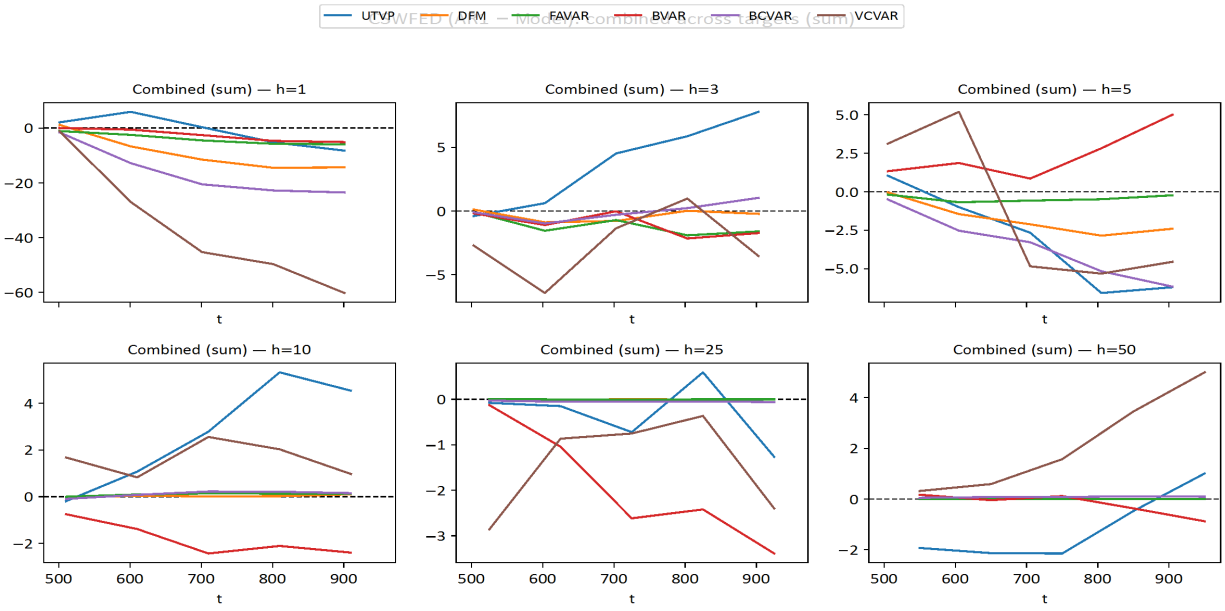

Table 7: Relative MSFE by Variable, Horizon, and Model (100-Variable Simulation Panel)

	UTVP	DFM	FAVAR	BVAR	BCVAR	VCVAR
X1						
h=1	0.81	1.972	0.987	0.961	2.5	1.976
h=3	0.638	1.03	0.985	0.972	1.123	1.788
h=5	0.758	0.988	0.982	1.02	1.01	1.016
h=10	0.886	0.999	0.998	1.032	1.001	0.848
h=25	1.037	1.000*	1	1.009	1	0.95
h=50	0.868	1.001	1	1.021	0.999	0.929
X2						
h=1	0.906	0.96	0.892	0.899	1.267	3.221
h=3	0.979	1.06	1.065	1.083	0.925	1.669
h=5	0.942	0.984	1.01	1.019	0.962	0.967
h=10	1.035	1.007	0.998	1.056	1.015	1.148
h=25	0.997	0.999	0.999	0.999	1.001	0.986
h=50	1.027	1.001	1	0.997	1	1.052
X3						
h=1	2.868	1.527	1.577	1.525	2.791	2.208
h=3	1.287	0.668*	0.967	0.898	0.577	1.154
h=5	1.238	1.032	0.987	0.804*	1.11	1.056
h=10	0.848	0.997	1.003	1.029	0.997	0.792
h=25	1.13	1	1	0.957	1.004	1.346
h=50	0.908	1	1	0.945	0.998	0.829
X4						
h=1	1.611	0.442	1.118	1.115	0.367	0.252
h=3	0.936	1.068	1.024	1.103	1.153	1.025
h=5	1.102	0.989	0.999	0.889	1.068	0.978
h=10	1.063	0.995	1.001	1.035	1.006	1.132
h=25	1.069	1.001	1	1.015	1.002	1.054
h=50	0.958	1	1	1.013	0.999	0.988
X5						
h=1	1.663	0.68	0.734	0.74	0.617	2.041
h=3	0.899	0.789	0.879	0.892	0.876	0.94
h=5	1.851	1.359	1.069	0.976	1.69	2.511
h=10	0.946	1.005	1.001	1.17	1.016	1.005
h=25	1.051	0.999	1	1.237	1.001	0.928
h=50	0.974	0.997	0.999	1.169	0.996*	0.882
X6						
h=1	0.934	0.705	0.814*	0.816*	0.87	3.257
h=3	0.493	1.053	0.997	0.927	0.941	0.321
h=5	1.057	1.009	0.987	0.896	0.997	0.983
h=10	0.994	1.001	1.001	1.106	1.009	1.095
h=25	0.992	1	1	1.013	1.001	0.992
h=50	1.029	0.998	1	1	0.992	0.85
X7						
h=1	0.701	1.511	1.173	1.174	1.482	1.005
h=3	0.819	0.993	0.97	0.972	0.998	0.862
h=5	0.745	0.987	0.989	0.923	1.038	0.803
h=10	0.544	1	1	0.986	0.994	0.828
h=25	0.966	1.002	1.002	0.988	1.001	1.052
h=50	0.783	1	1	1.003	1.001	0.819
X8						
h=1	0.71	0.693	0.951	0.938	0.669	1.145
h=3	1.017	1.09	1.002	1.012	0.924	0.792
h=5	1.145	1.02	0.995	1.041	1.004	1.144
h=10	0.942	0.996	1.001	0.866	1.002	0.975
h=25	1.038	1	1	1.008	1.001	0.972
h=50	1.021	1.001	1	0.989	1	0.885
X9						
h=1	0.794	0.95	1.539	1.565	1.005	0.81
h=3	0.965	1.031	1.046	1.036	0.94	0.681
h=5	1.042	0.999	1.004	0.973	1.011	0.99
h=10	1.009	1.001	1.004	1.121	0.995	0.94
h=25	1.114	1	1	1.067	0.999	1.041
h=50	0.992	1	1.001	0.994	1	0.953
X10						
h=1	0.815	2.796	1.297	1.214	2.515	4.866
h=3	0.826	1.216	1.404	1.405	1.233	1.872
h=5	0.967	1.051	1.057	1.054	0.987	0.458
h=10	0.88	0.979	0.968	0.899	0.918	0.878
h=25	0.785	1.001	1	1.119	0.999	1.081
h=50	1.4	1.001	1.001	1.017	1	0.923

Table 8: Relative ALPL by Variable, Horizon, and Model (100-Variable Simulation Panel)

	UTVP	DFM	FAVAR	BVAR	BCVAR	VCVAR
X1						
h=1	0.319	-0.448	0.014	0.030	-0.768	-0.334
h=3	0.870	-0.142	0.016	0.008	-0.204	0.202
h=5	0.385	-0.010	0.092	0.000	-0.018	0.442
h=10	0.657	-0.025	0.051	-0.019	-0.006	0.615
h=25	-0.169	0.022	-0.022	-0.007	-0.006	0.201
h=50	-0.019	0.048***	0.013***	0.026*	-0.002	-0.278
X2						
h=1	0.059	-0.031	0.054	0.056	-0.144	-0.672
h=3	0.079	-0.006	-0.014	-0.017	0.012	-0.178
h=5	0.176	0.004	-0.007	-0.004	0.021	0.039
h=10	-0.149	-0.001	0.012**	0.005	0.000	-0.224
h=25	0.017	0.003	-0.005	0.001	0.000	-0.053
h=50	0.015	0.010	0.000	0.000	0.000	0.009
X3						
h=1	-0.012	-0.143	-0.052	-0.046	-0.241	-0.216
h=3	-0.176	0.073**	-0.005	-0.010	0.078	-0.180
h=5	0.025	-0.011	-0.003	0.081**	-0.040	-0.101
h=10	0.116	0.022	0.010	-0.018	0.002	-0.104
h=25	0.083	-0.005	0.003	-0.016	-0.001	-0.201
h=50	0.083	-0.008	0.006	0.018	-0.001	0.098
X4						
h=1	-0.084	0.228	-0.048	-0.049	0.234	0.122
h=3	-0.167	-0.017	-0.017	-0.037	-0.060	-0.118
h=5	-0.255	-0.003	-0.002	0.056***	-0.034	-0.153
h=10	0.123	0.000	0.002	-0.012	-0.004	-0.029
h=25	0.134	-0.033	0.005	-0.006	0.001	0.064
h=50	-0.019	0.009	0.000	0.001	0.001	-0.122
X5						
h=1	0.165	-0.067	0.025	0.023	-0.066	0.018
h=3	0.285	0.059	0.043	0.044	0.040	0.335*
h=5	0.538**	-0.033	-0.008	0.022	-0.018	0.324
h=10	0.571**	0.030***	0.009***	-0.014	-0.001	0.350*
h=25	0.368	0.028*	-0.027	-0.039	0.002	0.296
h=50	0.449*	0.019***	-0.059	-0.006	0.002	0.310*
X6						
h=1	-0.153	0.051	0.057**	0.033	0.074	-0.348
h=3	0.309	-0.054	0.015	0.056	0.038	0.307
h=5	0.285	-0.002	0.008	0.087	0.006	0.238
h=10	0.064	-0.031	-0.003	-0.053	-0.008	-0.054
h=25	-0.059	-0.009	-0.004	-0.013	-0.001	-0.115
h=50	-0.151	0.015***	0.025**	-0.032	0.004*	-0.285
X7						
h=1	0.052	-0.499	-0.291	-0.237	-0.435	-0.165
h=3	1.395	-0.056	-0.123	0.029	-0.034	1.773
h=5	0.551	0.009	0.055	0.042	-0.042	0.409
h=10	0.108	0.003	0.005	0.010	0.002	-0.102
h=25	0.100	-0.001	-0.029	0.001	-0.001	-0.138
h=50	-0.327	-0.027	-0.026	0.002	-0.001	0.002
X8						
h=1	0.167	0.083	0.025	-0.008	0.047	-0.037
h=3	0.085	0.006	-0.023	-0.017	0.023	-0.029
h=5	-0.260	0.026*	0.003	-0.010	0.000	-0.228
h=10	-0.089	-0.007	0.013	-0.015	0.000	-0.220
h=25	-0.225	-0.004	-0.013	-0.042	0.000	-0.255
h=50	0.097	0.015	-0.002	0.006	0.000	-0.042
X9						
h=1	-0.168	0.012	-0.093	-0.081	0.006	-0.252
h=3	0.439	0.031	-0.057	-0.223	0.090	0.709
h=5	1.685	-0.088	0.051	-0.010	-0.031	1.828
h=10	-0.142	0.021	-0.040	-0.015	0.001	-0.280
h=25	-0.310	-0.005	0.007	-0.042	0.001	-0.226
h=50	0.078	0.015	0.000	-0.010	0.001	0.047
X10						
h=1	0.262*	-0.386	-0.073	-0.050	-0.325	-0.137
h=3	0.091	0.023	-0.028	-0.009	-0.009	-0.208
h=5	0.350**	-0.002	-0.020	-0.016	0.003	0.238*
h=10	0.279	-0.020	0.010***	0.091***	-0.001	-0.010
h=25	0.281	0.074**	-0.012	-0.007	-0.002	0.004
h=50	0.231	0.026**	0.032**	0.020	-0.002	0.047

Table 9: Diebold–Mariano Test Statistics for Relative MSFE, by Variable and Horizon (100-Variable Simulation Panel)

	UTVP	DFM	FAVAR	BVAR	BCVAR	VCVAR
X1						
h=1	0.81	1.972	0.987	0.961	2.5	1.976
h=3	0.638*	1.03	0.985	0.972	1.123	1.788
h=5	0.758	0.988***	0.982***	1.02	1.01	1.016
h=10	0.886	0.999	0.998	1.032	1.001	0.848
h=25	1.037	1	1	1.009	1	0.95
h=50	0.868	1.001	1	1.021	0.999	0.929
X2						
h=1	0.906	0.96	0.892	0.899	1.267	3.221
h=3	0.979	1.06	1.065	1.083	0.925	1.669
h=5	0.942	0.984***	1.01	1.019	0.962	0.967
h=10	1.035	1.007	0.998	1.056	1.015	1.148
h=25	0.997	0.999	0.999	0.999	1.001	0.986
h=50	1.027	1.001	1	0.997	1	1.052
X3						
h=1	2.868	1.527	1.577	1.525	2.791	2.208
h=3	1.287	0.668	0.967	0.898	0.577	1.154
h=5	1.238	1.032	0.987	0.804	1.11	1.056
h=10	0.848	0.997	1.003	1.029	0.997	0.792
h=25	1.13	1	1	0.957	1.004	1.346
h=50	0.908	1	1	0.945	0.998	0.829
X4						
h=1	1.611	0.442	1.118	1.115	0.367	0.252
h=3	0.936	1.068	1.024	1.103	1.153	1.025
h=5	1.102	0.989***	0.999	0.889***	1.068	0.978***
h=10	1.063	0.995	1.001	1.035	1.006	1.132
h=25	1.069	1.001	1	1.015	1.002	1.054
h=50	0.958	1	1	1.013	0.999	0.988
X5						
h=1	1.663	0.68	0.734	0.74	0.617	2.041
h=3	0.899	0.789	0.879	0.892	0.876	0.94
h=5	1.851	1.359	1.069	0.976	1.69	2.511
h=10	0.946	1.005	1.001	1.17	1.016	1.005
h=25	1.051	0.999	1	1.237	1.001	0.928
h=50	0.974	0.997	0.999	1.169	0.996	0.882
X6						
h=1	0.934	0.705	0.814*	0.816*	0.87	3.257
h=3	0.493	1.053	0.997	0.927**	0.941	0.321
h=5	1.057	1.009	0.987	0.896***	0.997	0.983
h=10	0.994	1.001	1.001	1.106	1.009	1.095
h=25	0.992	1	1	1.013	1.001	0.992
h=50	1.029	0.998	1	1	0.992	0.85
X7						
h=1	0.701	1.511	1.173	1.174	1.482	1.005
h=3	0.819	0.993	0.97	0.972	0.998	0.862
h=5	0.745	0.987	0.989	0.923	1.038	0.803***
h=10	0.544	1	1	0.986	0.994	0.828
h=25	0.966	1.002	1.002	0.988	1.001	1.052
h=50	0.783	1	1	1.003	1.001	0.819
X8						
h=1	0.71	0.693	0.951	0.938	0.669	1.145
h=3	1.017	1.09	1.002	1.012	0.924	0.792
h=5	1.145	1.02	0.995***	1.041	1.004	1.144
h=10	0.942	0.996	1.001	0.866	1.002	0.975
h=25	1.038	1	1	1.008	1.001	0.972
h=50	1.021	1.001	1	0.989	1	0.885
X9						
h=1	0.794	0.95	1.539	1.565	1.005	0.81
h=3	0.965	1.031	1.046	1.036	0.94	0.681
h=5	1.042	0.999	1.004	0.973	1.011	0.990***
h=10	1.009	1.001	1.004	1.121	0.995	0.94
h=25	1.114	1	1	1.067	0.999	1.041
h=50	0.992	1	1.001	0.994	1	0.953
X10						
h=1	0.815	2.796	1.297	1.214	2.515	4.866
h=3	0.826	1.216	1.404	1.405	1.233	1.872
h=5	0.967	1.051	1.057	1.054	0.987***	0.458***
h=10	0.88	0.979	0.968	0.899	0.918	0.878
h=25	0.785	1.001	1	1.119	0.999	1.081
h=50	1.4	1.001	1.001	1.017	1	0.923

Table 10: Time of Runs (seconds)

Table 10 reports the wall-clock time required to estimate each forecasting model across the three datasets examined in the paper. Each entry corresponds to a single recursive expanding-window forecasting run, covering the full sequence of real-time forecast origins. All models are estimated with a uniform lag length of $p = 3$ and 500 Monte Carlo runs, ensuring fully comparable computational requirements across specifications.

	TVP-VAR Benchmark	DFM	FAVAR	BVAR	BCVAR	VCVAR
19 Variable FRED	2.2	1.4	2.9	56.3	51.7	249
128 Variable FRED	14.6	15.3	14.5	63.5	997.7	8540.2
100 Variable Simulated	3.0	3.7	3.9	0.9	188.4	3107.4

A Algorithms for Model Comparisons

This section shows the algorithms for the models used for ease of comparison.

A.1 DFM

DFM Algorithm

1. **Recursively expanding-window origins:** for each annual origin t_k, form the training sample up to t_k and split variables into remainder $Y^{(k)} = \{Y_t : t \le t_k\}$.
2. **Select** (q, p) **by BIC:** q is the number of static factors and p is the VAR lag order. (i) extract $F^{(k)}(q) = \mathrm{PCA}(Y^{(k)})$; (ii) for each $p \le 13$, fit VAR(p) on $F^{(k)}(q)$ and compute its BIC; (iii) fit the diagonal measurement block; (iv) choose (q_k^*, p_k^*) that minimizes the total BIC.
3. **Forecast factors:** for a chosen forecast horizon H, compute the H-step mean forecasts $\widehat{F}_{t_k+h|t_k}$ for $h = 1, \dots, H$.
4. **Forecast observables:** compute $\widehat{Y}_{t_k+h|t_k} = \lambda_0 + \lambda_1 \widehat{F}_{t_k+h|t_k}$ for each $h = 1, \dots, H$.
5. **Predictive densities:** simulate draws using $\varepsilon^F \sim N(0, \Sigma_F)$ and $\varepsilon^Y \sim N(0, \Sigma_Y)$;

A.2 FAVAR

FAVAR Algorithm

1. **Recursively expanding-window origins:** for each annual origin t_k, form the training sample up to t_k and split variables into remainder $Y_t^{\tilde{}}$ and primary Y_t^*.
2. **Select** (q, p) **by BIC:** q is the number of static factors and p is the VAR lag order. (i) for each $q \le \lfloor \sqrt{n_{\mathrm{rem}}} \rfloor$, where n_{rem} is the number of remainder non-predicted variables, extract $F_t(q) = \mathrm{PCA}(Y_t^{\tilde{}})$; (ii) for each $p \le 13$, fit VAR(p) on $Z_t = [F_t, Y_t^*]$, (iii) fit the diagonal measurement block $Y_t^{\tilde{}} = \lambda_0 + \lambda_1 F_t + \varepsilon_t$, (iv) choose (q_k^*, p_k^*) that minimizes the total BIC.
3. **Joint VAR forecasts:** from the selected VAR on $Z_t = [F_t, Y_t^*]$, for a chosen forecast horizon H, compute the H-step mean forecasts $(\widehat{F}_{t_k+h|t_k}, \widehat{Y}^*_{t_k+h|t_k})$.
4. **Predictive densities:** simulate draws using $\varepsilon^Z \sim N(0, \Sigma_U)$

A.3 BVAR (Minnesota)

BVAR Algorithm

1. **Expanding-window origins:** for each annual origin t_k, form $Y^{*(k)} = \{Y_t^* : t \le t_k\}$ using only the forecast target variables.
2. **Select** (p, ω) **by BIC:** for each $p \le p_{\max}$ and grid value $\omega^{-1} \in [0.5\sqrt{mp},\ 10\sqrt{mp}]$ with increments of $0.1\sqrt{mp}$ (Koop, Korobilis, and Pettenuzzo, 2019) and $\lambda_2 = 1, \lambda_3 = 1$ and we only tune ω, (i) estimate AR(1) residuals σ_i for each variable; (ii) build Minnesota prior variances $V_j(\omega)$; (iii) compute the posterior mode $B_j = (X'X + V_j^{-1})^{-1} X' Y_{\cdot j}$ using ridge regression loss functions; (iv) form $\widehat{\Sigma}_U$ from residuals and compute $\mathrm{BIC}(p, \omega) = -2\ell(B, \widehat{\Sigma}_U) + k \log T$; (v) choose (p_k^*, ω_k^*) with minimal BIC.
3. **Posterior-mode VAR:** with (p_k^*, ω_k^*), retain $\widehat{B}^{(k)}$ and $\widehat{\Sigma}_U^{(k)}$ as the BVAR approximation on Y_t^*.
4. **Multi-step forecasts:** recursively iterate the VAR(p_k^*) with coefficients $\widehat{B}^{(k)}$ to obtain 12-step mean forecasts $\widehat{Y}^*_{t_k+h|t_k}$, $h = 1, \dots, 12$.

A.4 BCVAR

RC-VAR+BMA Algorithm

1. **Expanding-window origins:** for each annual origin t_k, form the full panel $Y^{(k)} = \{Y_t^{\text{med}} : t \le t_k\}$ on `medium_vars`.
2. **Random compression VAR draws:** fix lag order p; for $r = 1, \dots, R$:
   (i) build the VAR design $(Y_{\text{tar}}, X_{\text{full}})$ with p lags; (ii) for each equation i, draw a random projection $\Phi_i \in \mathbb{R}^{m_i \times K}$ with $m_i \sim \text{Unif}\{1, \dots, m_{\max}\}$ and $K = np$; (iii) regress y_i on the compressed regressors $X_c = X_{\text{lags}}\Phi_i'$ (ridge OLS) and lift coefficients back to full dimension via $\beta_i = \Phi_i'\beta_{c,i}$; (iv) stack intercepts and lifted coefficients into $B^{(r)}$ and compute residuals, Gaussian log-likelihood, $\widehat{\Sigma}_U^{(r)}$, and $\text{BIC}^{(r)}$.
3. **BIC weighting (BMA):** compute weights

$$w_r \propto \exp\Big(-\tfrac{1}{2}(\text{BIC}^{(r)} - \min_s \text{BIC}^{(s)})\Big), \qquad \sum_{r=1}^{R} w_r = 1.$$

4. **BMA mean path:** for each draw r, convert $B^{(r)}$ into $(c^{(r)}, \{A_\ell^{(r)}\}_{\ell=1}^p)$ and iteratively generate a 12-step deterministic mean path $\widehat{Y}_{t_k+h|t_k}^{(r)}$; form the BMA forecast

$$\widehat{Y}_{t_k+h|t_k} = \sum_{r=1}^{R} w_r \, \widehat{Y}_{t_k+h|t_k}^{(r)}, \qquad h = 1, \dots, 12.$$

5. **Targets and output:** extract the forecasted subvector for `forecast_targets` from $\widehat{Y}_{t_k+h|t_k}$ and stack all (t_k, h) into a tidy panel with columns `origin`, `forecast_date`, and the target series.

A.5 Univariate TVP VAR

Vine Copula TVP-VAR Algorithm

1. **Expanding-window origins:** for each annual origin t_k, form the balanced panel $Y^{(k)} = \{Y_t : t \le t_k\}$ on `var_list`.
2. **Univariate TVP marginals:** for each series $j = 1, \dots, n$, fit a discounted RLS TVP-AR(p) to $Y_{\cdot j}^{(k)}$, obtaining filtered states $\beta_{j,t}$, variance path $\sigma_{j,t}^2$, and final $(\beta_{j,T}, \sigma_{j,T}^2)$.
3. **Output:** stack point forecasts and, if simulated, bands for all (t_k, h) in tidy form with columns `origin`, `forecast_date`, and the variables in `var_list` (later restricted to `forecast_targets`).

A.6 VCVAR

Vine Copula TVP-VAR Algorithm

1. **Expanding-window origins:** for each annual origin t_k, form the balanced panel $Y^{(k)} = \{Y_t : t \le t_k\}$ on `var_list`.
2. **Univariate TVP marginals:** for each series $j = 1, \dots, n$, fit a discounted RLS TVP-AR(p) to $Y_{\cdot j}^{(k)}$, obtaining filtered states $\beta_{j,t}$, variance path $\sigma_{j,t}^2$, and final $(\beta_{j,T}, \sigma_{j,T}^2)$.
3. **Empirical pseudo-observations:** apply the empirical CDF columnwise to $Y^{(k)}$ to obtain $U^{(k)} \in (0,1)^{T_{\text{eff}} \times n}$ with entries $u_{t,j} = \text{rank}(Y_{t,j})/(T_{\text{eff}} + 1)$.

4. **Vine copula fit:** estimate an R-vine copula on $U^{(k)}$ using `pyvinecopulib` (`Vinecop.from_data`) with data-driven family and structure selection; this yields a dependence model $\widehat{C}^{(k)}$.
5. **12-step mean forecasts:** starting from the last p observations of $Y^{(k)}$, recursively apply the TVP-AR(p) one-step predictors to construct marginal means $\widehat{Y}_{t_k+h|t_k}$ for $h = 1, \ldots, 12$.
6. **Output:** stack point forecasts and, if simulated, bands for all (t_k, h) in tidy form with columns `origin`, `forecast_date`, and the variables in `var_list` (later restricted to `forecast_targets`).

B Forecasting

B.1 Construction of the FRED-MD Panels and Variable Mapping

We start from the March 2016 vintage of the FRED-MD database (McCracken and Ng, 2016), which contains 134 monthly U.S. macroeconomic and financial time series from 1959:1 to 2014:2. We follow the authors' transformation codes and imputation procedures to obtain a balanced, demeaned panel. The resulting dataset is stored as `2016-03_imputed.csv`, where the first column contains calendar dates and all remaining columns correspond to transformed series.

For the forecasting exercises in Section 3.1, we work with two nested sets of variables:

1. a smaller macro panel of 19 series, denoted `medium_vars`, which is used to forecast the seven target variables in Koop, Korobilis, and Pettenuzzo (2019).
2. a "full" panel of FRED-MD style variables, denoted `fred_vars`;

Forecast Target Variables. Following Koop, Korobilis, and Pettenuzzo (2019), we forecast the following seven key U.S. macroeconomic and financial variables:

- **INDPRO**: Industrial production growth
- **UNRATE**: Unemployment rate
- **PAYEMS**: Total nonfarm payroll employment
- **FEDFUNDS**: Change in the effective Federal Funds rate
- **GS10**: Change in the 10-year Treasury yield
- **WPSFD49207**: Finished goods producer price inflation (PPI)
- **CPIAUCSL**: Consumer price inflation (CPI)

The exact composition of `medium_vars` and `fred_vars` is given below. If the variable is not in the medium panel, it is in the full panel.

Table 11: Variables in the Medium Panel and Full FRED-MD Panel

Medium Panel?	FRED-MD Category	Variable Code
Yes	Income & Output	RPI
Yes	Industrial Production	INDPRO

Continued on next page

(continued)

Medium Panel?	FRED-MD Category	Variable Code
Yes	Labor Market	HWI
Yes	Labor Market	UNRATE
Yes	Labor Market	PAYEMS
Yes	Housing	HOUST
Yes	Housing	PERMIT
Yes	Money & Credit	M1SL
Yes	Money & Credit	BUSLOANS
Yes	Business Investment	INVEST
Yes	Interest Rates	FEDFUNDS
Yes	Interest Rates	GS10
Yes	Interest Rates	T10YFFM
Yes	Exchange Rates	EXUSUKx
Yes	Prices	WPSFD49207
Yes	Commodity Prices	OILPRICEx
Yes	Prices (CPI)	CPIAUCSL
Yes	Stock Market	S&P 500
Yes	Surveys/PMI	NAPM
No	Income & Output	W875RX1
No	Industrial Production	IPFPNSS
No	Industrial Production	IPFINAL
No	Industrial Production	IPCONGD
No	Industrial Production	IPDCONGD
No	Industrial Production	IPNCONGD
No	Industrial Production	IPBUSEQ
No	Industrial Production	IPMAT
No	Industrial Production	IPDMAT
No	Industrial Production	IPNMAT
No	Industrial Production	IPMANSICS
No	Industrial Production	IPB51222S
No	Industrial Production	IPFUELS
No	Labor Market	HWIURATIO
No	Labor Market	CLF16OV
No	Labor Market	CE16OV
No	Labor Market	UEMPMEAN
No	Labor Market	UEMPLT5
No	Labor Market	UEMP5TO14
No	Labor Market	UEMP15OV
No	Labor Market	UEMP15T26

Continued on next page

(continued)

Medium Panel?	FRED-MD Category	Variable Code
No	Labor Market	UEMP27OV
No	Labor Market	CLAIMSx
No	Labor Market	USGOOD
No	Labor Market	CES1021000001
No	Labor Market	USCONS
No	Labor Market	MANEMP
No	Labor Market	DMANEMP
No	Labor Market	NDMANEMP
No	Labor Market	SRVPRD
No	Labor Market	USTPU
No	Labor Market	USWTRADE
No	Labor Market	USTRADE
No	Labor Market	USFIRE
No	Labor Market	USGOVT
No	Labor Market	CES0600000007
No	Labor Market	AWOTMAN
No	Labor Market	AWHMAN
No	Labor Market	CES0600000008
No	Labor Market	CES2000000008
No	Labor Market	CES3000000008
No	Housing	HOUSTNE
No	Housing	HOUSTMW
No	Housing	HOUSTS
No	Housing	HOUSTW
No	Housing	PERMITNE
No	Housing	PERMITMW
No	Housing	PERMITS
No	Housing	PERMITW
No	Consumption	DPCERA3M086SBEA
No	Consumption	CMRMTSPLx
No	Consumption	RETAILx
No	Surveys/PMI	NAPMNOI
No	Surveys/PMI	NAPMSDI
No	Surveys/PMI	NAPMII
No	Durable Goods Orders	AMDMNOx
No	Durable Goods Orders	AMDMUOx
No	Inventories	BUSINVx
No	Inventories	ISRATIOx
No	Money & Credit	M2SL

Continued on next page

(continued)

Medium Panel?	FRED-MD Category	Variable Code
No	Money & Credit	M2REAL
No	Money & Credit	AMBSL
No	Money & Credit	TOTRESNS
No	Money & Credit	REALLN
No	Money & Credit	NONREVSL
No	Money & Credit	CONSPI
No	Money & Credit	MZMSL
No	Money & Credit	DTCOLNVHFNM
No	Money & Credit	DTCTHFNM
No	Rates & Spreads	CP3Mx
No	Rates & Spreads	TB3MS
No	Rates & Spreads	TB6MS
No	Rates & Spreads	GS1
No	Rates & Spreads	GS5
No	Rates & Spreads	AAA
No	Rates & Spreads	BAA
No	Rates & Spreads	COMPAPFFx
No	Rates & Spreads	TB3SMFFM
No	Rates & Spreads	TB6SMFFM
No	Rates & Spreads	T1YFFM
No	Rates & Spreads	T5YFFM
No	Rates & Spreads	AAAFFM
No	Rates & Spreads	BAAFFM
No	Exchange Rates	EXSZUSx
No	Exchange Rates	EXJPUSx
No	Exchange Rates	EXCAUSx
No	Prices	WPSFD49502
No	Prices	WPSID61
No	Prices	WPSID62
No	Prices	PPICMM
No	Prices	NAPMPRI
No	Prices	CPIAPPSL
No	Prices	CPITRNSL
No	Prices	CPIMEDSL
No	Prices	CUSR0000SAC
No	Prices	CUUR0000SAD
No	Prices	CUSR0000SAS
No	Prices	CPIULFSL
No	Prices	CUUR0000SA0L2

Continued on next page

(continued)

Medium Panel?	FRED-MD Category	Variable Code
No	Prices	CUSR0000SA0L5
No	Prices	PCEPI
No	Prices	DDURRG3M086SBEA
No	Prices	DNDGRG3M086SBEA
No	Prices	DSERRG3M086SBEA
No	Stock Market	S&P div yield
No	Stock Market	S&P PE ratio

Forecasting Metrics Details

Inverse-Variance Weights: For each variable $j \in \mathcal{J}$ and forecast horizon h, let $\mathcal{T}_{j,h}$ denote the set of out-of-sample forecast dates for which an AR(1) h-step forecast is available.

The AR(1) forecast error is

$$e_{\mathrm{AR}(1),j,h,\tau} = y_{j,\tau} - \widehat{y}_{\mathrm{AR}(1),j,h,\tau}, \qquad \tau \in \mathcal{T}_{j,h}.$$

The error variance used in the weighting scheme is estimated as

$$\widehat{\sigma}^2_{j,h} = \frac{1}{|\mathcal{T}_{j,h}| - 1} \sum_{\tau \in \mathcal{T}_{j,h}} \left(e_{\mathrm{AR}(1),j,h,\tau} - \overline{e}_{\mathrm{AR}(1),j,h}\right)^2, \qquad \overline{e}_{\mathrm{AR}(1),j,h} = \frac{1}{|\mathcal{T}_{j,h}|} \sum_{\tau \in \mathcal{T}_{j,h}} e_{\mathrm{AR}(1),j,h,\tau}.$$

The inverse-variance weight is then

$$w_{j,h} = \frac{1}{\widehat{\sigma}^2_{j,h}}.$$

Equivalently, the weight matrix for horizon h is diagonal:

$$W_h = \mathrm{diag}\big(w_{1,h},\, w_{2,h}, \ldots, w_{|\mathcal{J}|,h}\big)\,.$$

CSWFED for a Single Variable

For a model $i \neq \mathrm{AR}(1)$, define its h-step forecast error as

$$e_{i,j,h,\tau} = y_{j,\tau} - \widehat{y}_{i,j,h,\tau}.$$

The weighted squared error differential used in the CSWFED statistic is

$$d_{i,j,h,\tau} = w_{j,h}\left(e^2_{\mathrm{AR}(1),j,h,\tau} - e^2_{i,j,h,\tau}\right).$$

Order the forecast dates as

$$\tau_1 < \tau_2 < \cdots < \tau_{T_{j,h}}.$$

Then the cumulative sum of weighted forecast error differentials (CSWFED) for variable j is

$$\mathrm{CSWFED}_{i,j,h}(\tau_m) = \sum_{k=1}^{m} d_{i,j,h,\tau_k} = \sum_{k=1}^{m} w_{j,h}\left(e^2_{\mathrm{AR}(1),j,h,\tau_k} - e^2_{i,j,h,\tau_k}\right), \qquad m = 1, \ldots, T_{j,h}.$$

CSWFED Aggregated Across All Medium-VAR Variables

Let $\mathcal{J}$ denote the set of medium-VAR variables (e.g., seven series in our application).

The combined CSWFED curve plotted in the paper is

$$\mathrm{CSWFED}^{\mathrm{sum}}_{i,h}(\tau_m) = \sum_{j\in\mathcal{J}} \mathrm{CSWFED}_{i,j,h}(\tau_m) = \sum_{j\in\mathcal{J}} \sum_{k=1}^{m} w_{j,h} \left(e^2_{\mathrm{AR}(1),j,h,\tau_k} - e^2_{i,j,h,\tau_k}\right).$$

If instead one averages across variables (i.e. `reduce_op="mean"`), then

$$\mathrm{CSWFED}^{\mathrm{mean}}_{i,h}(\tau_m) = \frac{1}{|\mathcal{J}|} \sum_{j\in\mathcal{J}} \mathrm{CSWFED}_{i,j,h}(\tau_m).$$

For each origin and each model, the h-step point forecasts are produced by iterating the estimated state equation forward:

$$\widehat{y}_{t+h|t} = \mathbb{E}(y_{t+h} \mid \mathcal{F}_t).$$

When $h > 1$, these forecasts require repeated substitution of model-implied predictions into future lags. The resulting (t, h) forecast panel allows direct computation of mean squared forecast errors for evaluating point forecast accuracy.

When predictive densities are required, we follow the model-specific simulation procedure. For linear Gaussian models, densities are iterated analytically; for Bayesian VARs, MC draws from the posterior predictive distribution are used; and for copula-based models, paths are generated by first sampling marginal innovations and then imposing the estimated copula dependence before drawing from a MC. The forecasts at horizon h are thus based on the density

$$p(y_{t+h} \mid \mathcal{F}_t)$$

either obtained in closed form or approximated by Monte Carlo sampling. This unified forecasting design ensures comparability across models in both point and density forecast evaluations.

B.2 Construction of the Simulation Data

In addition to the empirical FRED-MD application, we construct a controlled Monte Carlo environment to study how the forecasting models behave under rich, time-varying dependence. We simulate a high-dimensional panel $\{X_t\}_{t=1}^{T}$ with $T = 1000$ monthly observations and $d = 100$ variables, where each series follows a time-varying AR(1) with stochastic volatility and correlated shocks.

Baseline cross-sectional dependence and regimes. We first draw a random positive-definite covariance matrix Σ_{base} from a Gaussian ensemble and convert it to a correlation matrix R by scaling with the implied standard deviations. This provides a realistic but generic cross-sectional dependence structure for the d variables. We then construct two regimes: a "calm" regime with correlation matrix $R_{\mathrm{calm}} = R$, and a "crisis" regime with higher average correlation

$$R_{\mathrm{crisis}} = \rho R + (1-\rho)\mathbf{1}\mathbf{1}^\top, \quad \rho = 0.7,$$

with unit diagonals enforced ex post. The regime indicator $S_t \in \{0, 1\}$ follows a two-state Markov chain: with probability 0.97 the process remains in the calm regime and switches to crisis otherwise; once in crisis, it remains there with probability 0.95 and reverts with the complementary probability. This construction generates persistent episodes of high correlation that mimic financial crises.

Time-varying autoregressive coefficients. For each variable $j = 1, \ldots, d$, we model the AR(1) coefficient as a smooth, bounded process

$$X_{t,j} = \phi_{j,t} X_{t-1,j} + u_{t,j},$$

where $\phi_{j,t} = \tanh(g_{t,j})$ and $g_{t,j}$ follows a stationary AR(1) process

$$g_{t,j} = \mu_g + \phi_g(g_{t-1,j} - \mu_g) + \sigma_g \varepsilon_{t,j}^{(g)},$$

with $\mu_g = 0.3$, $\phi_g = 0.98$, and $\sigma_g = 0.1$. The $\tanh(\cdot)$ mapping ensures $|\phi_{j,t}| < 1$ for all t, so the conditional dynamics remain stable while allowing for gradual drifts in persistence over time and across variables.

Stochastic volatility. Residual volatilities are driven by log-variance processes $h_{t,j}$,

$$h_{t,j} = \mu_h + \phi_h(h_{t-1,j} - \mu_h) + \sigma_h \varepsilon_{t,j}^{(h)}, \qquad \sigma_{t,j} = \exp\big(\tfrac{1}{2} h_{t,j}\big),$$

with $\mu_h = -0.5$, $\phi_h = 0.98$, and $\sigma_h = 0.2$. This specification produces highly persistent stochastic volatility with episodes of elevated uncertainty, again aligned with standard SV formulations used in the macro-finance literature.

Correlated, heavy-tailed shocks. Conditional on the regime S_t and volatilities $\{\sigma_{t,j}\}$, the innovations $u_t = (u_{t,1}, \ldots, u_{t,d})'$ are generated in two steps. First, we draw a standardized shock z_t:

- in the calm regime, $z_t \sim \mathcal{N}(0, R_{\text{calm}})$;
- in the crisis regime, we construct a heavy-tailed, strongly correlated vector by drawing $z_t^{(g)} \sim \mathcal{N}(0, R_{\text{crisis}})$ and then applying a scale mixture
$$z_t = \frac{z_t^{(g)}}{\sqrt{\chi^2_\nu / \nu}}, \quad \nu = 3,$$
which yields t-like tails and stronger tail dependence.

Second, we scale by the contemporaneous volatilities:

$$u_{t,j} = \sigma_{t,j} z_{t,j}, \qquad j = 1, \ldots, d.$$

This delivers a panel with time-varying persistence, stochastic volatility, and correlated heavy-tailed shocks that become more pronounced in crisis regimes.

Output and forecast targets. The simulated series are collected in a data frame with columns `X1`, ..., `X100` and a time index `t`. For the forecasting experiments, we focus on the first ten variables,

$$\texttt{forecast\_vars} = \{\texttt{X1}, \ldots, \texttt{X10}\},$$

which serve as the pseudo-observed macro variables of interest.

B.3 Forecasting Design in the Simulated Panel

The forecasting design in the simulated environment mirrors the real-time exercise used for the FRED-MD data, but with full control over the data-generating process. We define a grid of estimation endpoints and forecast horizons,

$$\mathcal{T}_{\text{end}} = \{500, 600, 700, 800, 900\}, \qquad \mathcal{H} = \{1, 3, 5, 10, 25, 50\}.$$

For each combination $(T_{\text{end}}, h) \in \mathcal{T}_{\text{end}} \times \mathcal{H}$, the following steps are carried out:

1. We treat all observations up to T_{end} as the information set available to the forecaster and estimate the competing models on $\{X_t\}_{t=1}^{T_{\text{end}}}$.
2. Using the estimated model, we generate h-step-ahead forecasts for the ten target variables in `forecast_vars`, i.e. for $(X_{T_{\text{end}}+h,1}, \ldots, X_{T_{\text{end}}+h,10})'$.
3. The procedure is repeated over all (T_{end}, h) combinations, producing a structured collection of pseudo-real-time forecasts that can be directly compared across models and horizons.

Programmatically, we construct three aligned lists: (i) `train_end_vec`, which stores the estimation endpoints in $\mathcal{T}_{\text{end}}$; (ii) `horizon_vec`, which stores the corresponding forecast horizons in $\mathcal{H}$; and (iii) `pred_vars_vec`, which records the forecast target set `forecast_vars` for each task. These lists define the full grid of forecasting tasks and ensure that each model is evaluated on an identical set of (T_{end}, h) combinations. This design allows us to assess how the models perform as the effective sample size grows and as the forecast horizon moves from short-term (e.g. $h = 1$) to medium- and long-term horizons (e.g. $h = 25$ or 50) in a DGP that explicitly features time variation in both persistence and volatility.

B.4 System Specifications

Table 12: System Specifications

Category	Specification
CPU	
CPU Model	12th Gen Intel(R) Core(TM) i7-12700KF
Architecture	x86_64
Bits	64
Logical Cores	20
Physical Cores	12
Base Frequency	3.6100 GHz
Measured Frequency	3.6100 GHz
L2 Cache	12,582,912 bytes
L3 Cache	26,214,400 bytes
RAM Total	34.21 GB
Operating System / Python	
OS	Windows 10 (Build 10.0.26200, SP0)
Python Version	3.11.14

C Other Forms of Copula VARs

Tsionas, Izzeldin, and Trapani (2022) describe a copula VAR in which cross equation dependence is modeled through a Gaussian mixture copula, while each marginal equation is a time varying parameter regression with stochastic volatility. For each variable $i = 1, \ldots, n$, they estimate

$$y_{i,t} = \beta_{i,t}' z_{i,t} + u_{i,t}, \qquad u_{i,t} = \exp(h_{i,t}/2)\, u_{i,t}^*, \qquad u_{i,t}^* \sim N(0,1),$$

where both the time varying coefficients $\beta_{i,t}$ and log variances $h_{i,t}$ follow separate state space evolutions. These marginal states are sampled jointly using the Kim Shephard Chib MCMC algorithm, which makes the marginal step

computationally intensive.

After obtaining marginal predictive densities $f_i(y_{i,t} \mid z_{i,t})$, the probability integral transforms

$$v_{i,t} = F_i(y_{i,t} \mid z_{i,t})$$

are combined through a Gaussian mixture copula,

$$c(v_{1,t}, \ldots, v_{n,t}; \alpha) = \sum_{g=1}^{G} p_g \, \varphi_n\big(\Psi^{-1}(v_t) \mid \mu_g, \Omega_g\big) \prod_{i=1}^{n} \phi_i\big(\Psi_i^{-1}(v_{i,t})\big) \,,$$

where p_g, μ_g, and Ω_g are mixture weights, component means, and structured covariance matrices. The complete likelihood is

$$\ell = \sum_{t} \Big[\sum_{i=1}^{n} \log f_i(y_{i,t} \mid z_{i,t}) + \log c(v_{1,t}, \ldots, v_{n,t}; \alpha) \Big],$$

so that the marginal densities and copula term contribute separately. Even with dimension reduction, the estimator requires joint sampling of all TVP states, volatilities, mixture weights, and copula covariance structures inside a high dimensional MCMC scheme. This tight coupling makes replication difficult.

In contrast, the approach used here follows the same two step structure but replaces the full MCMC with a simpler procedure. Marginal TVP dynamics are estimated by discounted recursive least squares, and dependence is captured by a Gaussian or vine copula fitted directly to rank based pseudo observations. This produces a model aligned with the structure of Tsionas, Izzeldin, and Trapani (2022) while being much easier to implement.