Stochastic Dynamic Correlations with Exogenous Shifts: Connecting Macroeconomic Events and Financial Risk

Igor Ferreira Batista Martins*, Sune Karlsson, Tamás Kiss, Stepan Mazur

School of Business, Örebro University, Sweden

February 20, 2026

Abstract

We propose a novel dynamic correlation model that allows for systematic shifts in correlations driven by exogenous variables while ensuring positive definiteness. Using matrix logarithm transformation, we obtain unbounded eigencorrelations that uniquely recover the correlation matrix. Our MCMC algorithm not only handles the non-linearity induced by the transformation but also easily incorporates information from realized correlations. Our approach enables separate treatment of correlations and volatilities, and we leverage this feature in our empirical application. Using data on the Australian Dollar and Japanese Yen exchange rate against the US Dollar, we find that correlation and volatilities generally increase on days when events occur, with announcements in the United States—and in particular FOMC meetings—being most impactful. Our findings highlight the dominant role of the US within the global financial markets.

Keywords: Dynamic correlation, macroeconomic announcements, currency returns

*Igor Ferreira Batista Martins gratefully acknowledges the funding from Jan Wallanders och Tom Hedelius stifelse samt Tore Browaldhs stiftelse - grant BFv22-0005.

1 Introduction

Multivariate modeling of variances and correlations of asset returns is fundamental to financial decision-making, playing a crucial role in portfolio optimization, risk management, and asset pricing. While univariate models of time-varying volatility such as stochastic volatility, GARCH, and realized volatility models are well-established in the literature, modelling of dynamic correlations remains more challenging, primarily due to the difficulty of ensuring a positive definite conditional correlation matrix. The Dynamic Conditional Correlation (DCC) model of Engle (2002), the factor stochastic volatility (FSV) models, as in Aguilar and West (2000), Lopes and Carvalho (2007), and Kastner et al. (2017) and the Autoregressive Inverse Wishart model of Philipov and Glickman (2006), are among the leading approaches for capturing correlation dynamics. All of them rely on quadratic forms to construct covariance matrices, thereby guaranteeing positive definiteness. However, while this structure allows for exogenous shifts in variances, it is not well suited for capturing exogenous shifts in correlations. This limitation is particularly relevant in applications where correlation dynamics respond to external economic or financial conditions.

This paper proposes a novel dynamic correlation model that allows systematic shifts in correlations driven by exogenous variables while ensuring positive definiteness. Central to our approach is the vector representation of the correlation matrix using the matrix logarithmic transformation proposed by Archakov and Hansen (2021). This non-linear transformation yields a vector, which we call *eigencorrela-*

tion vector, composed of unbounded elements that uniquely recover the correlation matrix. Our approach leverages two key properties of this transformation. First, the unboundness of the eigencorrelations allow us to capture shifts using a standard Gaussian linear model. Second, the uniqueness of the transformation ensures that shifts in the eigencorrelation vector translate directly into changes in the correlation matrix. The benefits of the transformation comes with the cost of a more complex estimation due to the non-linearities of the matrix logarithmic transformation, which do not allow for the use of filtering techniques as in Carter and Kohn (1994), Frühwirth-Schnatter (1994) and Chan and Jeliazkov (2009). To address this issue, we propose an MCMC scheme that not only handles the non-linearity induced by the transformation but also easily incorporates information from realized correlations by extending the MCMC of Jacquier et al. (1994). Another advantage of directly transforming the correlation matrix is that it allows for separate treatment of correlation and volatility dynamics, enabling more flexible model specifications.

We leverage the features of our model when analyzing the effects of macroeconomic announcements on exchange rate correlations and volatilities. While the impact of announcements on exchange rate volatility has been recently studied in Martins and Lopes (2024), to our knowledge this is the first paper to analyze the impact of macroeconomic events on the dynamics of the correlation between assets. We employ our model in a bivariate setting using daily returns of the Australian Dollar and Japanese Yen against the US Dollar. Based on estimates of latent time-varying correlation and volatility of the currency returns, we document several empirical regularities. First, volatility is higher on days when macroeconomic announcements occur in either the country of the currency or the US, with the FOMC meetings being most impactful for both currencies. Second, the correlation between the two currencies also increases following US announcements but shows little reaction to those from Japan or Australia. These findings align with previous literature highlighting the dominant role of US macroeconomic news in global financial markets (Savor and Wilson, 2014; Brusa et al., 2020), usually motivated by the central role the US plays in the global economy.

Other papers have also tackled the issue of guaranteeing positive definiteness. Lopes et al. (2010) and Shirota et al. (2017) rely on the Cholesky decomposition to enforce this constraint. However, a well-known limitation of the Cholesky-based approach is that it is not order-invariant, meaning that estimated parameters depend on the arbitrary ordering of assets. In contrast, our proposal is order invariant and does not suffer from this limitation. Moreover, Chiu and Tsui (1996), Kawakatsu (2006) and Bauer and Vorkink (2011) have also employed matrix logarithmic but they do so in the context of covariance matrices rather than correlations. By not explicitly disentangling the correlation and volatility components, these approaches lose flexibility when compared to our proposal. The closest paper to ours is Archakov et al. (2025) where the authors also benefit from matrix logarithmic in a realized GARCH framework. Nonetheless, none of the previous papers investigate the effect of exogenous variables shifting correlation matrices.

Our paper also contributes to the broader literature on the effects of macroeconomic announcements on financial markets. Savor and Wilson (2013), Savor and Wilson (2014), Brusa et al. (2020), and Lucca and Moench (2015) focus on how macroeconomic news affects stock returns, while Andersen et al. (2007), Stroud and Johannes (2014), and Martins and Lopes (2024) investigate their impact on volatility dynamics. However, to the best of our knowledge, this is the first paper to analyze correlation shifts due to macroeconomic announcements.

The remainder of the paper is organized as follows. In Section 2, we present our econometric model. Section 3 discuss plausible priors and our MCMC scheme capable of tackling the non-linearities induced by the matrix logarithm transformation. In Section 4, we estimate our model using the proposed MCMC scheme and quantify the effect of macroeconomic announcements on both correlation and volatilities for the Australian Dollar and Japanese Yen. Section 5 concludes.

2 Model Specification

Let $y_t = (y_{t,1}, \ldots, y_{t,n})$ denote an n-dimensional vector of returns observed at time t. We assume that that y_t has zero mean and a dynamic covariance matrix:

$$y_t = \Omega_t^{1/2} \varepsilon_t \text{ with } \varepsilon_t \sim N(0, \Sigma_{\varepsilon,t}),$$

where $\Omega_t^{1/2}$ is a diagonal matrix capturing time-varying volatilities and $\Sigma_{\varepsilon,t}$ is a time-varying correlation matrix. The zero-mean assumption is appropriate for returns at a daily or higher frequency and also allows us to keep our focus on modeling the second moments. However, adding a mean equation is straightforward and can be handled by standard techniques. An element $h_{i,t}$ in the matrix Ω_t

$$\Omega_t^{1/2} = diag\Big(exp\Big(\frac{h_{i,t}}{2}\Big)\Big)$$

captures the time-varying volatility of asset i, with

$$h_{i,t} = x_{i,t} + \delta_i^v A_t.$$

The first component, $x_{i,t}$, is a persistent component following a standard stochastic volatility specification

$$x_{i,t} = \alpha_i^v + \beta_i^v x_{i,t-1} + \sigma_i^v \varepsilon_{i,t}^v \text{ with } \varepsilon_{i,t}^v \sim N(0,1),$$

where α_i^v and β_i^v are the intercept and the persistence parameter of the persistent component and σ_i^v is the volatility of the volatility innovation. The second component is a transitory component containing exogenous variables that impact volatility contemporaneously. In our application, we focus on the impact of macroeconomic announcements, therefore we specify the exogenous variables A_t as a set of announcement dummies, common for all assets. The loading of the volatility on the announcements δ_i^v is asset specific and constant over time. In line with the previous literature (e.g. Archakov et al., 2025) we also use realized variance as an additional source of information through the equation

$$logrvar_{i,t} = \xi_i^v + h_{i,t} + \sigma_i^{v,rl} \varepsilon_{i,t}^{v,rl} \text{ with } \varepsilon_{i,t}^{v,rl} \sim N(0,1),$$

where ξ_i^v captures the asset-specific difference between the level of realized volatility and the latent volatility process and $\sigma_i^{v,rl}$ is the volatility of the innovation in the realized-variance equation.

We impose a similar structure on the elements of the correlation matrix $\Sigma_{\varepsilon,t}$. However, in order to preserve positive definiteness without imposing further restrictions, we use the matrix logarithmic transformation of $\Sigma_{\varepsilon,t}$. Starting from the spectral decomposition, we have $\Sigma_{\varepsilon,t} = Q\Lambda_{\varepsilon,t}Q'$. Taking the log of the eigenvalues we get

$$log\Sigma_{\varepsilon,t} = Qlog\Lambda_{\varepsilon,t}Q'$$

with $log\Lambda_{\varepsilon,t} = diag(log\lambda_{1,t}, \ldots, log\lambda_{I,t})$. Archakov and Hansen (2021) show that the elements below the diagonal of $log\Sigma_{\varepsilon,t}$ are sufficient to uniquely recover $\Sigma_{\varepsilon,t}$. We denote these values, which we refer to as eigencorrelations, by $\gamma_{j,t}^{\Sigma_\varepsilon}$. The domain of eigencorrelations is unrestricted, they span the whole real line, which then uniquely transform to a well-defined correlation matrix. Therefore, we can use a specification similar to the time-varying volatility for $\gamma_{j,t}^{\Sigma_\varepsilon}$. Specifically, each eigencorrelation $\gamma_{j,t}^{\Sigma_\varepsilon}$ has a persistent and a transitory component

$$\gamma_{j,t}^{\Sigma_\varepsilon} = z_{j,t} + \delta_j^c A_t.$$

The persistent component $z_{j,t}$ follows an equation similar to stochastic volatility

$$z_{j,t} = \alpha_j^c + \beta_j^c z_{j,t-1} + \sigma_j^c \varepsilon_{j,t}^c \text{ with } \varepsilon_{j,t}^c \sim N(0,1),$$

where α_j^c and β_j^c are the intercept and the persistence parameters of the persistent component and σ_j^c is the volatility of the innovation term. Regarding the transitory component, we use the same set of announcements (A_t) for the eigencorrelations as for volatility, with eigencorrelation-specific loadings δ_j^c. Furthermore, we use the matrix logarithmic transformation of the realized correlation matrix to convey additional information to the latent correlations by the equation

$$\gamma_{j,t}^{rc} = \xi_j^c + \gamma_{j,t}^{\Sigma_\varepsilon} + \sigma_j^{c,rl} \varepsilon_{j,t}^{c,rl} \text{ with } \varepsilon_{j,t}^{c,rl} \sim N(0,1).$$

Here ξ_j^c captures the level difference between the realized and the latent eigencorrelations and $\sigma_j^{c,rl}$ is the innovation variance.

3 Bayesian Inference

We adopt a Bayesian framework and use MCMC methods to sample from the joint posterior distribution of the parameters and latent states i.e. to sampled from

$$\{\alpha_i^v, \beta_i^v, \sigma_i^{2,v}, \xi_i^v, \sigma_i^{2,v,rl}, \delta_i^v, x_i\}_i^I, \{\alpha_j^c, \beta_j^c, \sigma_j^{2,c}, \xi_j^c, \sigma_j^{2,c,rl}, \delta_j^c, z_j\}_j^J | y, rvar, rcor$$

where I is the number of assets and $J = I(I-1)/2$ the number of eigencorrelations. Subsection 3.1 describes the choice for priors. Subsection 3.2 details the MCMC approach.

3.1 Priors

Our choice of priors reflects several empirical regularities documented in the literature and allows for great simplification of the MCMC scheme. We opt to break down the choice of priors by assuming prior independence across volatilities and eigencorrelations. Therefore, we split the description based on the unknown quantities related to volatilities $p(\{\alpha_i^v, \beta_i^v, \sigma_i^{2,v}, \xi_i^v, \sigma_i^{2,v,rl}, \delta_i^v, x_i\}_i^I)$ and associated with eigencorrelations $p(\{\alpha_j^c, \beta_j^c, \sigma_j^{2,c}, \xi_j^c, \sigma_j^{2,c,rl}, \delta_j^c, z_j\}_j^J)$

Our prior specification for the volatility component assumes independence among its components and is designed to incorporate several empirical regularities and model constraints. First, volatility levels can vary substantially across asset classes. Second, volatility exhibits strong persistence, with higher persistence typically observed in high-frequency data and lower persistence in low-frequency settings. We capture the first two empirical regularities by specifying $p(\alpha_i^v, \beta_i^v)$ as a multivariate normal with zero correlation. For our empirical application, we select

$$\begin{pmatrix} \alpha_i^v \\ \beta_i^v \end{pmatrix} \sim N\left(\begin{pmatrix} -0.10 \\ 0.95 \end{pmatrix}, \begin{pmatrix} 0.1 & 0 \\ 0 & 0.1 \end{pmatrix} \right).$$

This prior implies that annualized volatility levels are predominantly concentrated between 2% and 40%, while the half-life of log-volatility is on average two weeks. This choice is quite reasonable for exchange rates. For example, Martins and Lopes (2024) recover volatility levels for exchange rate rates mostly concentrated around 10%.

Third, while some studies acknowledge the effect of events on volatility, most papers do not account for this component. Thus, we take a conservative approach to modeling events, centering our prior beliefs on the assumption that events have no impact on volatilities but we express our uncertainty about this claim by selecting $\delta_i^v \sim N(0,1)$.

Fourth, realized log-variance and the latent stochastic volatility estimates should capture the same quantity but they can differ due to, among other reasons, non-traded hours and microstructure noise. We account for this by selecting $\xi_i^v \sim N(0,1)$.

Fifth, volatility is inherently non-negative, needing a prior that ensures strict positivity. Thus, we choose $\sigma_i^{2,v}$ and $\sigma_i^{2,v,rl}$ as $IG(1,1)$.

Our prior specification for the eigencorrelation component follows the same rational. Since eigencorrelations span the entire real line, we model their levels and persistence using Gaussian priors. Specifically, we assume

$$\begin{pmatrix} \alpha_j^c \\ \beta_j^c \end{pmatrix} \sim N\left(\begin{pmatrix} 0 \\ 0.95 \end{pmatrix}, \begin{pmatrix} 0.1 & 0 \\ 0 & 0.1 \end{pmatrix}\right).$$

This prior centers our beliefs around the uncorrelated case while allowing correlation levels to take both positive and negative values while accounting for the persistence. We keep our conservative approaches to modeling events and deviations of realized and latent quantities by specifying $\delta_j^c \sim N(0,1)$ and $\xi_j^c \sim N(0,1)$. We complete our prior specification by assigning $\sigma_j^{2,c}$ and $\sigma_j^{2,c,rl}$ a $IG(1,1)$ prior allowing them to be primarily determined by the likelihood.

3.2 Posterior Inference

We combine the likelihood described in Section 2 with the priors described in Subsection 3.1 to obtain the joint posterior distribution of the parameters and latent states represented by

$$(\{\alpha_i^v, \beta_i^v, \sigma_i^{2,v}, \xi_i^v, \sigma_i^{2,v,rl}, \delta_i^v, x_i\}_i^I, \{\alpha_j^c, \beta_j^c, \sigma_j^{2,c}, \xi_j^c, \sigma_j^{2,c,rl}, \delta_j^c, z_j\}_j^J)|y, rvar, rcor$$

Since no analytical solution is available, we rely on a Metropolis within Gibbs approach to sample from the posterior of the model.

In addition to reflecting our beliefs, our prior choice greatly simplifies the steps of our Markov Chain Monte Carlo (MCMC) scheme by providing full-conjugacy for all unknown quantities with the exception of $\{z_{j,t}\}_{t=1}^T$ and $\{x_{i,t}\}_{t=1}^T$. Appendix A describes the MCMC scheme in detail.

Sampling from $\{z_{j,t}\}_{t=1}^T|\cdot$ is not a trivial task. Since both the measurement equation and the matrix logarithm transformation are non-linear, we cannot rely on samplers such as Carter and Kohn (1994), Frühwirth-Schnatter (1994) and Chan and Jeliazkov (2009) which are well suited for conditionally Gaussian dynamic linear models. To tackle this issue, we extend the original Metropolis approach by Jacquier et al. (1994) to our eigencorrelation framework where we have realized (eigen)correlations as an additional source of information.

Our model specification implies a multivariate normal for $\{z_{j,t}, z_{j,t+1}, \gamma_{j,t}^{rc}\}$

$$\begin{pmatrix} z_{j,t} \\ z_{j,t+1} \\ \gamma_{j,t}^{rc} \end{pmatrix} \sim N\left(\begin{pmatrix} \alpha_j^c + \beta_j^c z_{j,t-1} \\ (1+\beta_j^c)\alpha_j^c + \beta_j^{c,2} z_{j,t-1} \\ \xi_j^c + \delta_j^c A_t + \alpha_j^c + \beta_j^c z_{j,t-1} \end{pmatrix}, \begin{pmatrix} \sigma_j^{c,2} & \beta_j^c \sigma_j^{c,2} & \sigma_j^{c,2} \\ \beta_j^c \sigma_j^{c,2} & (1+\beta_j^{c,2})\sigma_j^{c,2} & \beta_j^c \sigma_j^{c,2} \\ \sigma_j^{c,2} & \beta_j^c \sigma_j^{c,2} & \sigma_j^{c,2} + \sigma_j^{c,rl,2} \end{pmatrix}\right)$$

Using the properties of the multivariate normal, we obtain expressions for $E(z_{j,t}|\cdot)$ and $V(z_{j,t}|\cdot)$ detailed on Appendix A. If there was no information from realized (eigen)correlation, $\gamma_{j,t}^{rc}$, our approach would be a direct adaptation of the Metropolis algorithm in Jacquier et al. (1994) for latent volatilities. However, realized (eigen)correlations act as an additional source of information helping pin down the latent quantities and our algorithm accounts for this during sampling. The use of realized information is particularly helpful in diminishing the estimation uncertainty around the latent quantities.

Our approach for sampling $\{z_{j,t}\}_{t=1}^{T}|\cdot$ can be summarized as follows:

1. Given a current $z_{j,t}^{(iter)}$, sample a random walk candidate $z_{j,t}^{*}$

2. Compute $\Sigma_{\varepsilon,t}^{(iter)}$ and $\Sigma_{\varepsilon,t}^{*}$ using the inverse Archakov-Hansen transformation for $z_{j,t}^{(iter)}$ and $z_{j,t}^{*}$.

3. Compute the Metropolis acceptance probability

$$p = min(1, \frac{N(z_{j,t}^{*}|E(z_{j,t}|\cdot), V(z_{j,t}|\cdot))N(y_t|0, \Omega^{1/2}\Sigma_{\varepsilon,t}^{*}\Omega^{1/2})}{N(z_{j,t}^{(iter)}|E(z_{j,t}|\cdot), V(z_{j,t}|\cdot))N(y_t|0, \Omega^{1/2}\Sigma_{\varepsilon,t}^{(iter)}\Omega^{1/2})})$$

4. Accept the candidate with probability p and reject it with probability 1 - p.

An analogous approach is used to sample from $\{x_{i,t}\}_{t=1}^{T}|\cdot$. Again, relying on a multivariate normal for $\{x_{i,t}, x_{i,t+1}, logrvar_{i,t}\}$ followed by computing $E(x_{i,t}|\cdot)$ and $V(x_{i,t}|\cdot)$, computing the Metropolis acceptance probability and concluding acceptance/rejection step.

Our algorithm provides a simple way to tackle the non-linearities while also accounting for realized eigencorrelations and realized variances.

4 Empirical analysis

We estimate the model outlined in Section 2 using the prior specifications and MCMC scheme detailed in Section 3. Subsection 4.1 describes our dataset, which consists of Australian Dollar and Japanese Yen futures, two currencies frequently used in carry trade strategies, serving as y_t in our analysis. Additionally, we provide an overview of the macroeconomic announcements, A_t, that may influence exchange rate volatilities and correlations. Subsection 4.2 presents our in-sample estimation results.

4.1 Data

Our returns data is based on prices of Australian Dollar and Japanese Yen futures (quoted in US dollars) at the Chicago Mercantile Exchange (CME). We have price data on a one minute frequency between January 4, 2016 to June 26, 2024, hence our dataset covers approximately 8.5 years. We use data up to December 30, 2022 for estimation and reserve the remainder observations for out of sample analysis (to be added later). The futures contracts are traded almost uninterrupted from Sunday 5pm (Central Time, CT) to Friday 4pm (CT). The only exception is the one-hour break between 4pm and 5pm (CT) on each trading day. For each currency, we use the price quotes of the most liquid futures contract.

We consider three types of events for our analysis. First, we have the day of the release of the policy decisions from the Reserve Bank of Australia (RBA) and the Bank of Japan (BOJ), as well as the meetings of the Federal Open Market Committee (FOMC) of the Federal Reserve in the US. We also use the days when headline inflation (Consumer Price Index, CPI) is published in Australia, Japan and the United States. Lastly, to capture information related to the real economy, we use release dates of the Australian and Japanese unemployment rates, and the release dates of the non-farm payrolls in the United States. The CPI, unemployment rate and non-farm payroll releases have monthly frequency, while monetary policy decisions are made approximately eight times a year.

Our choice of events is guided by macroeconomic principles. Covered and uncovered interest rate parity are textbook connections between exchange rates and interest rates. All events considered are either the interest rate decision itself or tightly connect to interest rates via the Taylor Rule (Taylor, 1993) and the dual mandate of price stability and maximum sustainable employment.

We note that all calculations are carried out based on the New-York time zone, basically taking the perspective of an investor sitting in New York. That is, the time of the events are all calculated according to the New-York time zone, and event dummies are based on the whether the event happened on a certain day in New-York time.

4.2 Estimation results

Given our focus on event induced volatility and correlation, we first discuss estimates related to the events. Our main results are summarized in Table 1, which presents estimates of the volatility of the two currencies and their correlation. The baseline, presented in the first column in each panel, are the posterior mean and 90 percent credible interval of (annualized) volatility and correlation over days with no events while the rest of the table contain estimates of the same quantities over days where a specific type of event (specified by the column header) occurs.

[TABLE 1 ABOUT HERE]

Starting with the events in Australia (Panel A in Table 1), we see that the volatility of the Australian dollar is increased on average with approximately one to 1.5 percentage points. The effects are strongest for the release of the CPI and the policy decisions of the RBA, with the latter difference being significant, based on the 90 percent credible interval of the baseline. In contrast, events in Australia have no impact on the volatility of Japanese Yen. Regarding the correlation between the currencies, Australian events have a positive effect on the correlation, with the RBA-related event exerting the strongest impact, raising correlation from 0.18 to 0.25 (a substantial difference). Events in Japan (Panel B in Table 1) are analogous to the Australian events. In particular, the volatility of the Yen increases on days with an event related to Japan, while the volatility of the Australian Dollar is unaffected. The effect on the correlation is qualitatively similar, but more dampened: events in Japan are associated with an increase in the correlation between the Yen and the Australian dollar, but the increase is not statistically significant. Panel C in Table 1 presents the results for events that occur in the United States. Overall, we see that these events are the most impactful: volatilities of both currencies, as well as their correlation increase on all days with events in the US. It is also clear from the results that the most pronounced is the effect of the FOMC meetings. On FOMC days, the average volatility of the Australian dollar increases to 11.1 percent (from a baseline of 7.1) and the volatility of the Japanese Yen increases to 8.4 percent (from a baseline of 5.5). These are substantial changes, well beyond the credible interval of the baseline. Correlation raises from 0.18 to 0.35, also a significant increase.

Examining the estimates of the latent variables provides additional insights into the model's behavior. Figure 1 displays the posterior mean, black line, and the 90% credible intervals, red dashed lines, for the persistent component, z_t, and eigencorrelation, $\gamma_t^{\Sigma_\epsilon}$, in Panels A and B, respectively. As a reference, we also include the realized eigencorrelation subtracting ξ^c shown as a gray shaded line.

[FIGURE 1 ABOUT HERE]

The persistent component serves as a de-noised representation of the realized eigencorrelation, effectively filtering out short-term fluctuations. The figure suggests that incorporating dynamic correlation behavior is well justified. We observe an initial phase, up to the second half of 2017, where the (eigen)correlation between the currency pairs was negative, followed by a stable period until 2020, and a subsequent increase after the COVID-19 pandemic. Incorporating the events component reveals that some spikes removed from the realized measure should be retained since they are associated with macroeconomic announcements, others appear to be noise and can be smoothed out.

Figures 2 and 3 display the posterior mean, black line, and the 90% credible intervals, red dashed lines, for the persistent component, x_t shown in Panel A, and the log-variance, h_t shown in Panel B, for the Australian Dollar ($h_{AUD,t}$) and Japanese Yen ($h_{JPY,t}$), respectively. As a reference, we also include the realized log-variance subtracting ξ^v shown as a gray shaded line.

[FIGURE 2 ABOUT HERE]

[FIGURE 3 ABOUT HERE]

Similarly to eigencorrelation, the persistent component represents a de-noised version of the realized measure by filtering out short-term fluctuations. In both cases, volatility jumps during the financial crisis associated with the Covid-19 outbreak and is fairly persistent. In addition to the relatively slow movement, there are regularly spaced spikes in volatility associated with the events.

The estimates of parameters of the model not directly related to the events are presented in Table 2. In particular, the parameters of the persistent component (α, β and innovation variance σ^2) and the parameters for the equation connecting the latent and realized measures (ξ and $\sigma^{2,rl}$) are presented, both for (log-)volatilities and for the (eigen)correlation.

[TABLE 2 ABOUT HERE]

The results are overall fairly standard: both (log-)volatility and (eigen)correlation appears to have high persistence. The results also show that there is a level difference between the latent processes and their realized counterparts, with realized (log-)volatility being higher and realized (eigen)correlation being lower than their latent counterparts. These differences are usually attributed to the fact that the realized volatility and the latent volatility process capture slightly different objects, for example, realized volatility is affected by market microstructure noise. Therefore, it is important to account for the level difference when using the realized measures as an additional source of information to the latent second moments.

The general empirical conclusions of our results are intuitive and in line with previous research. Of particular interest are the event-related results, which can be summarized in three empirical regularities. First, events tend to increase both volatility and correlation when they have an impact. Second, events in the US have the strongest impact, in line with the US playing a central role in the global financial system. Third, the FOMC meetings have by far the strongest impact on the currency returns, which is reasonable given that exchange rates are directly affected by changing interest rates through the interest rate parity.

5 Conclusion

This paper introduces a novel dynamic correlation model that accommodates systematic shifts in correlations induced by exogenous variables while ensuring positive definiteness. By leveraging the matrix logarithm transformation, we obtain unbounded eigencorrelations that uniquely recover the correlation matrix, allowing for a more flexible representation of correlation dynamics. Our MCMC scheme handles the non-linearities introduced by the transformation and seamlessly integrates realized correlations and realized variances.

Our empirical analysis, using Australian Dollar and Japanese Yen exchange rates, demonstrates the effectiveness of the proposed model in capturing correlation dynamics. We find that both exchange rate volatilities and correlations increase on days with macroeconomic announcements, with US events—particularly FOMC meetings—exerting the strongest influence. These results reinforce the central role of US macroeconomic news in shaping global financial markets and suggest that accounting for exogenous drivers is relevant when modelling dynamic correlations.

References

Aguilar, O. and M. West (2000). Bayesian dynamic factor models and portfolio allocation. *Journal of Business & Economic Statistics 18*(3), 338–357.

Andersen, T. G., T. Bollerslev, F. X. Diebold, and C. Vega (2007). Real-time price discovery in global stock, bond and foreign exchange markets. *Journal of international Economics 73*(2), 251–277.

Archakov, I. and P. R. Hansen (2021). A new parametrization of correlation matrices. *Econometrica 89*(4), 1699–1715.

Archakov, I., P. R. Hansen, and A. Lunde (2025). A multivariate realized garch model.

Bauer, G. H. and K. Vorkink (2011). Forecasting multivariate realized stock market volatility. *Journal of Econometrics 160*(1), 93–101.

Brusa, F., P. Savor, and M. Wilson (2020). One central bank to rule them all. *Review of Finance 24*(2), 263–304.

Carter, C. K. and R. Kohn (1994). On gibbs sampling for state space models. *Biometrika 81*(3), 541–553.

Chan, J. C. and I. Jeliazkov (2009). Efficient simulation and integrated likelihood estimation in state space models. *International Journal of Mathematical Modelling and Numerical Optimisation 1*(1-2), 101–120.

Chiu, T. L. and K.-W. Tsui (1996). The matrix-logarithmic covariance model. *Journal of the American Statistical Association 91*(433), 198–210.

Engle, R. F. (2002). Dynamic conditional correlation: A simple class of multivariate generalized autoregressive conditional heteroskedasticity models. *Journal of Business & Economic Statistics 20*, 339–350.

Frühwirth-Schnatter, S. (1994). Data augmentation and dynamic linear models. *Journal of time series analysis 15*(2), 183–202.

Jacquier, E., N. G. Polson, and P. E. Rossi (1994). Bayesian analysis of stochastic volatility models. *Journal of Business & Economic Statistics 12*(4), 371–87.

Kastner, G., S. Frühwirth-Schnatter, and H. F. Lopes (2017). Efficient bayesian inference for multivariate factor stochastic volatility models. *Journal of Computational and Graphical Statistics 26*(4), 905–917.

Kawakatsu, H. (2006). Matrix exponential garch. *Journal of Econometrics 134*(1), 95–128.

Lopes, H. F. and C. M. Carvalho (2007). Factor stochastic volatility with time varying loadings and markov switching regimes. *Journal of Statistical Planning and Inference 137*(10), 3082–3091.

Lopes, H. F., R. McCulloch, and R. Tsay (2010). Cholesky stochastic volatility. *Unpublished Technical Report, University of Chicago, Booth Business School 2.*

Lucca, D. O. and E. Moench (2015). The pre-fomc announcement drift. *The Journal of finance 70*(1), 329–371.

Martins, I. and H. F. Lopes (2024). What events matter for exchange rate volatility ?

Philipov, A. and M. E. Glickman (2006). Multivariate stochastic volatility via wishart processes. *Journal of Business & Economic Statistics 24*(3), 313–328.

Savor, P. and M. Wilson (2013). How much do investors care about macroeconomic risk? evidence from scheduled economic announcements. *Journal of Financial and Quantitative Analysis 48*(2), 343–375.

Savor, P. and M. Wilson (2014). Asset pricing: A tale of two days. *Journal of Financial Economics 113*(2), 171–201.

Shirota, S., Y. Omori, H. F. Lopes, and H. Piao (2017). Cholesky realized stochastic volatility model. *Econometrics and Statistics 3*, 34–59.

Stroud, J. R. and M. S. Johannes (2014). Bayesian modeling and forecasting of 24-hour high-frequency volatility. *Journal of the American Statistical Association 109*(508), 1368–1384.

Taylor, J. B. (1993). Discretion versus policy rules in practice. In *Carnegie-Rochester conference series on public policy*, Volume 39, pp. 195–214. Elsevier.

Appendix A: MCMC scheme

The Markov Chain Monte Carlo (MCMC) algorithm for the model is outlined as follows:

For every asset i,

1. Sample $(\alpha_i^v, \beta_i^v)|\cdot$ from the Gaussian conjugate posterior.
2. Sample $\sigma_i^{2,v}|\cdot$ from the IG conjugate posterior.
3. Sample $\xi_i^v|\cdot$ from the Gaussian conjugate posterior.
4. Sample $\sigma_i^{2,v,rl}|\cdot$ from the IG conjugate posterior.
5. Sample $\delta_i^v|\cdot$ from the Gaussian conjugate posterior.
6. Sample $\{x_{i,t}\}_{t=1}^T|\cdot$ using expanded JPR accounting for log-realized variances.

For every eigencorrelation j

7. Sample $(\alpha_j^c, \beta_j^c)|\cdot$ from the Gaussian conjugate posterior.
8. Sample $\sigma_j^{2,c}|\cdot$ from the IG conjugate posterior.
9. Sample $\xi_j^c|\cdot$ from the Gaussian conjugate posterior.
10. Sample $\sigma_j^{2,c,rl}|\cdot$ from the IG conjugate posterior.
11. Sample $\delta_j^c|\cdot$ from the Gaussian conjugate posterior.

12. Sample $\{z_{j,t}\}_{t=1}^{T}|\cdot$ using expanded JPR accounting for realized eigencorrelations.

The only non straightforward step in our scheme are the sampling of $\{z_{j,t}\}_{t=1}^{T}|\cdot$ and $\{x_{i,t}\}_{t=1}^{T}|\cdot$. Bellow we detailed the sampling for $\{z_{j,t}\}_{t=1}^{T}|\cdot$. The sampling for $\{x_{i,t}\}_{t=1}^{T}|\cdot$ is analogous.

Our model specification implies a multivariate normal for $\{z_{j,t}, z_{j,t+1}, \gamma_{j,t}^{rc}\}$

$$\begin{pmatrix} z_{j,t} \\ z_{j,t+1} \\ \gamma_{j,t}^{rc} \end{pmatrix} \sim N\left(\begin{pmatrix} \alpha_j^c + \beta_j^c z_{j,t-1} \\ (1+\beta_j^c)\alpha_j^c + \beta_j^{c,2} z_{j,t-1} \\ \xi_j^c + \delta_j^c A_t + \alpha_j^c + \beta_j^c z_{j,t-1} \end{pmatrix}, \begin{pmatrix} \sigma_j^{c,2} & \beta_j^c \sigma_j^{c,2} & \sigma_j^{c,2} \\ \beta_j^c \sigma_j^{c,2} & (1+\beta_j^{c,2})\sigma_j^{c,2} & \beta_j^c \sigma_j^{c,2} \\ \sigma_j^{c,2} & \beta_j^c \sigma_j^{c,2} & \sigma_j^{c,2} + \sigma_j^{c,rl,2} \end{pmatrix}\right)$$

Writing the conditional expectation for the MVN, we obtain

$$E(z_{j,t}|\cdot) = \lambda_0 \alpha_j^c + \lambda_1 z_{j,t-1} + \lambda_2 z_{j,t+1} + \lambda_3 (\gamma_{j,t}^{rc} - \xi_j^c - \delta_j^c A_t)$$

with

$det\Sigma_{22} = \sigma_j^{c,2}[(1+\beta_j^{c,2})(\sigma_j^{c,2} + \sigma_j^{c,rl,2}) - \beta_j^{c,2}\sigma_j^{c,2}]$

$\lambda_0 = 1 - \frac{\beta_j^c \sigma_j^{c,2} \sigma_j^{c,rl,2}(1+\beta_j^{c,2}) + \sigma_j^{c,4}}{det\Sigma_{22}}$

$\lambda_1 = \beta_j^c - \frac{\sigma_j^{c,2}\sigma_j^{c,rl,2}\beta_j^{c,3} + \sigma_j^{c,4}\beta_j^c}{det\Sigma_{22}}$

$\lambda_2 = \frac{\beta_j^c \sigma_j^{c,2}\sigma_j^{c,rl,2}}{det\Sigma_{22}}$

$\lambda_3 = \frac{\sigma_j^{c,4}}{det\Sigma_{22}}$

Similarly, the conditional variance for the MVN implies

$$V(z_t|\cdot) = \sigma_j^{c,2}\left(1 - \frac{\beta_j^{c,2}\sigma_j^{c,2}\sigma_j^{c,rl,2} + \sigma_j^{c,4}}{det\Sigma_{22}}\right)$$

Table 1: The impact of events on volatility and correlations

This table presents model estimates of the volatility and correlation. The first column shows the posterior mean and the 90 percent credible interval of the annualized volatility of the Australian Dollar and Japanese Yen, as well as their correlation averaged over days without events, computed based on the level $\frac{\alpha}{1-\beta}$, while columns two to four show the analogous values for events, computed as the level $\frac{\alpha}{1-\beta}$ added of the event effect δ, specified in the column header. The events related to Australia, Japan and the United States are presented in the first, second and third panels of the table, respectively.

Panel A: Events in Australia

	Baseline	RBA	CPI	Unemp
vol_{AUD}	7.144	8.431	8.692	8.124
	(6.674–7.649)	(7.421–9.643)	(7.115–10.594)	(7.105–9.296)
vol_{JPY}	5.500	5.453	5.243	5.681
	(5.074–5.967)	(4.783–6.241)	(4.266–6.399)	(4.950–6.499)
$\text{corr}_{AUD,JPY}$	0.180	0.253	0.216	0.211
	(0.143–0.219)	(0.190–0.320)	(0.113–0.314)	(0.142–0.277)

Panel B: Events in Japan

	Baseline	BOJ	CPI	Unemp
vol_{AUD}	7.144	7.349	7.802	6.988
	(6.674–7.649)	(6.248–8.609)	(6.687–9.139)	(5.994–8.180)
vol_{JPY}	5.500	6.789	6.292	5.679
	(5.074–5.967)	(5.824–7.921)	(5.385–7.368)	(4.839–6.666)
$\text{corr}_{AUD,JPY}$	0.180	0.217	0.182	0.214
	(0.143–0.219)	(0.138–0.296)	(0.100–0.260)	(0.135–0.292)

Panel C: Events in the United States

	Baseline	FOMC	CPI	Payroll
vol_{AUD}	7.144	11.092	8.031	8.210
	(6.674–7.649)	(9.415–13.177)	(7.036–9.163)	(7.245–9.231)
vol_{JPY}	5.500	8.388	6.528	6.617
	(5.074–5.967)	(7.101–9.865)	(5.757–7.415)	(5.802–7.507)
$\text{corr}_{AUD,JPY}$	0.180	0.354	0.266	0.215
	(0.143–0.219)	(0.278–0.426)	(0.201–0.330)	(0.148–0.279)

Table 2: Parameter estimates

This table presents the parameter estimates of the model (except for the parameters for the events). Posterior means, along with 90 percent credible intervals are shown for each parameter. The following parameters are presented: the intercept (α) and persistence parameter (β) of the persistent components, the variance of the innovation to the persistent component (σ^2) and the intercept and innovation volatility of the equations for the realized measures (ξ and $\sigma^{2,rl}$). The first and second rows are related to the volatility of the Australian dollar and the Yen, respectively, while the last row shows estimates related to the eigencorrelation.

	α	β	σ^2	ξ	$\sigma^{2,rl}$
$h_{AUD,t}$	-0.055	0.965	0.014	1.106	0.082
	(-0.073, -0.038)	(0.955, 0.976)	(0.012, 0.017)	(1.044, 1.154)	(0.072, 0.095)
$h_{JPY,t}$	-0.097	0.954	0.039	0.821	0.130
	(-0.128, -0.067)	(0.940, 0.968)	(0.032, 0.047)	(0.770, 0.884)	(0.115, 0.146)
$\gamma_t^{\Sigma_\varepsilon}$	0.020	0.893	0.005	-0.086	0.007
	(0.014, 0.026)	(0.871, 0.914)	(0.004, 0.006)	(-0.118, -0.058)	(0.007, 0.008)

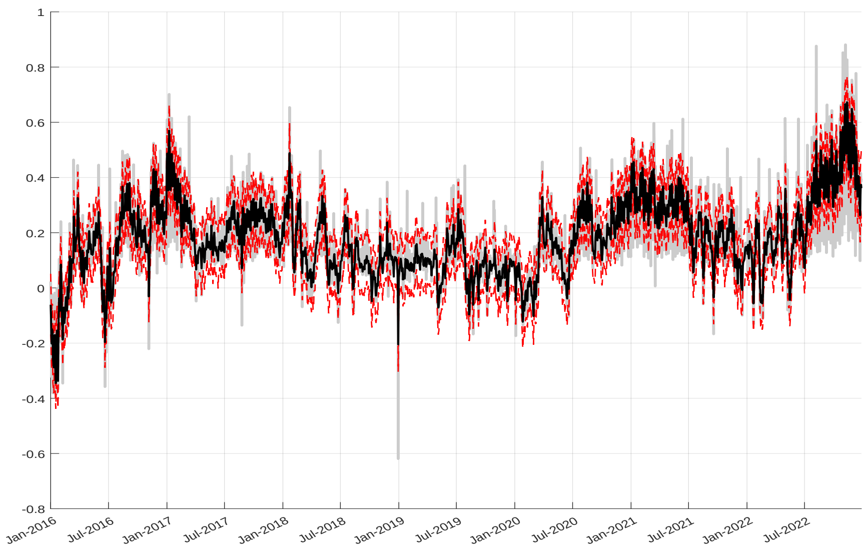

Panel A: z_t

Panel B: $\gamma_t^{\Sigma_\epsilon}$

Figure 1: **Estimates of the latent components of the eigencorrelation.** This figure shows the estimates of the persistent component of the eigencorrelation, z_t, in Panel A and the sum of persistent and announcement component yielding the full eigencorrelation, $\gamma_t^{\Sigma_\epsilon}$, in Panel B. The solid black line and the red dashed lines show the posterior mean and the 90 percent credible interval, respectively. For reference, we also add the realized eigencorrelation after subtracting the estimated intercepts ξ^c with a gray shaded line.

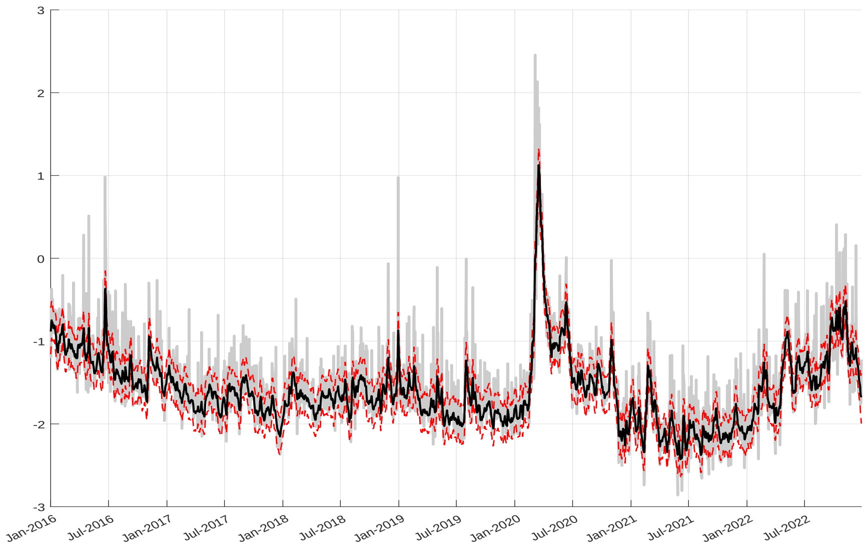

Panel A: $x_{AUD,t}$

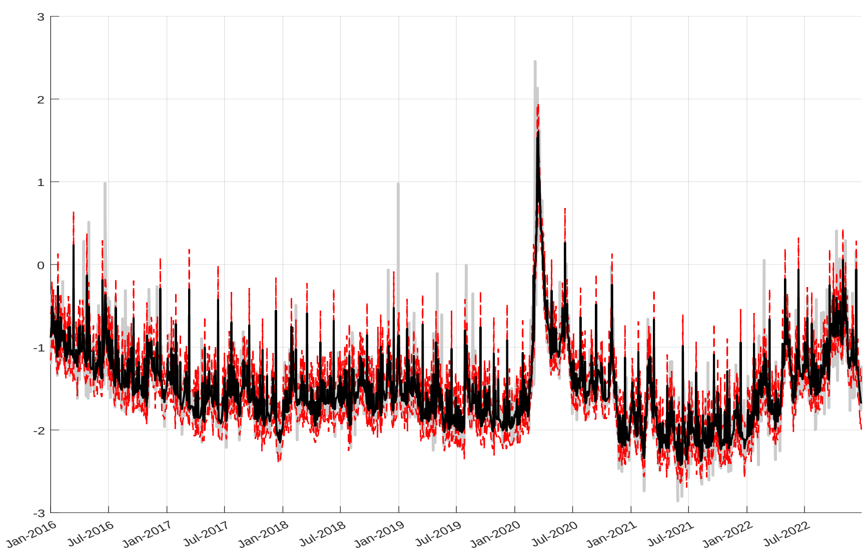

Panel B: $h_{AUD,t}$

Figure 2: **Estimates of the latent components of the Australian Dollar volatility.** This figure shows the estimates of the persistent component of the volatility, $x_{AUD,t}$, in Panel A and the sum of persistent and announcement component yielding the log-volatility, $h_{AUD,t}$, in Panel B. The solid black line and the red dashed lines show the posterior mean and the 90 percent credible interval, respectively. For reference, we also add the realized log-variance after subtracting the estimated intercepts ξ^v with a gray shaded line.

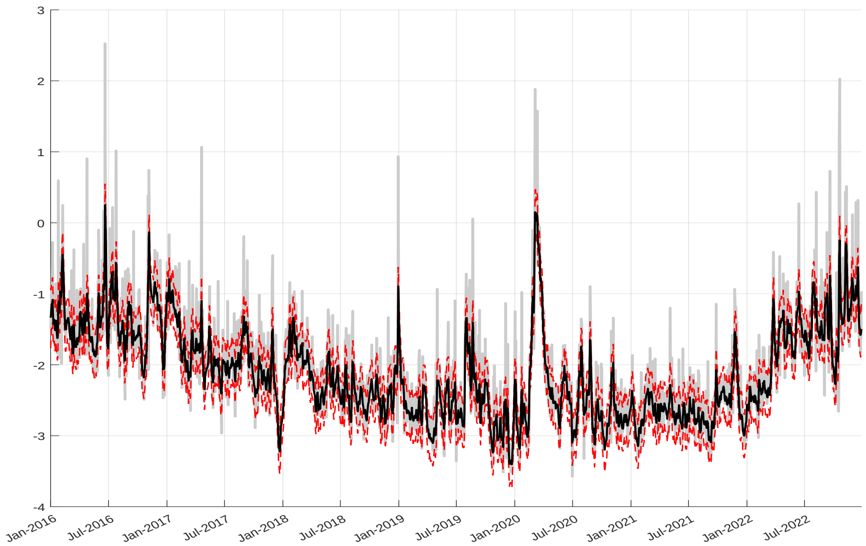

Panel A: $x_{JPY,t}$

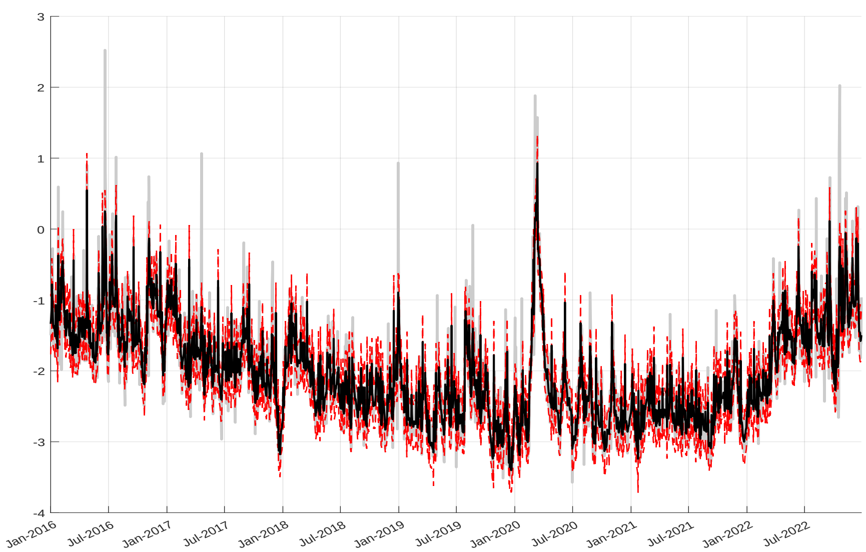

Panel B: $h_{JPY,t}$

Figure 3: **Estimates of the latent components of the Japanese Yen volatility.** This figure shows the estimates of the persistent component of the volatility, $x_{JPY,t}$, in Panel A and the sum of persistent and announcement component yielding the log-volatility, $h_{JPY,t}$, in Panel B. The solid black line and the red dashed lines show the posterior mean and the 90 percent credible interval, respectively. For reference, we also add the realized log-variance after subtracting the estimated intercepts ξ^v with a gray shaded line.