

CHANGE POINT DETECTION IN DYNAMIC GAUSSIAN GRAPHICAL MODELS: THE IMPACT OF COVID-19 PANDEMIC ON THE U.S. STOCK MARKET

BY BEATRICE FRANZOLINI^{1,a}, ALEXANDROS BESKOS^{2,b}, MARIA DE IORIO^{3,d},
WARRICK POKLEWSKI KOZIELL^{2,c} AND KAROLINA GRZESZKIEWICZ^{4,e}

¹*Bocconi Institute for Data Science and Analytics, Bocconi University, beatrice.franzolini@unibocconi.it*

²*Department of Statistical Science, University College London, a.beskos@ucl.ac.uk, warrickpk@gmail.com*

³*Yong Loo Lin School of Medicine, National University of Singapore, mdi@nus.edu.sg*

⁴*Yale-NUS College, karolina.grzeszkiewicz@u.yale-nus.edu.sg*

Reliable estimates of volatility and correlation are fundamental in economics and finance for understanding the impact of macroeconomics events on the market and guiding future investments and policies. Dependence across financial returns is likely to be subject to sudden structural changes, especially in correspondence with major global events, such as the COVID-19 pandemic. In this work we are interested in capturing abrupt changes over time in the conditional dependence across U.S. industry stock portfolios, over a time horizon that covers the COVID-19 pandemic. The selected stocks give a comprehensive picture of the U.S. stock market. To this end, we develop a Bayesian multivariate stochastic volatility model based on a time-varying sequence of graphs capturing the evolution of the dependence structure. The model builds on the Gaussian graphical models and the random change points literature. In particular, we treat the number, the position of change points, and the graphs as object of posterior inference, allowing for sparsity in graph recovery and change point detection. The high dimension of the parameter space poses complex computational challenges. However, the model admits a hidden Markov model formulation. This leads to the development of an efficient computational strategy, based on a combination of sequential Monte-Carlo and Markov chain Monte-Carlo techniques. Model and computational development are widely applicable, beyond the scope of the application of interest in this work.

1. Introduction. Understanding the temporal evolution of the dependence structure among time series is a fundamental topic in many fields, such as psychology (Williams (2021)), speech recognition (Bilmes (2004)), genomics (Yin and Li (2011)), and, in particular, finance (Hlávka, Hušková and Meintanis (2020)). In this latter context, estimates of volatility and correlation of different financial instruments are largely used for portfolio allocation, option-pricing, and to draw conclusions about the impact of macroeconomic events on the markets with the goal of guiding future investments and policies. In particular, estimates of correlation are key to minimise the risk of investment portfolios and define hedging strategies (see, among others, Lien, Tse and Tsui (2002), Lee (2010), Thampanya, Nasir and Huynh (2020), Dutta, Bouri and Noor (2021)). Changes in correlation modify the return/risk profile of the investments and are of interest to both investors and policy makers. To understand how to better prepare for and deal with future major global events, it is important to estimate the impact of the COVID-19 pandemic on the volatility and the dependence structure of financial instruments (Alqaralleh and Canepa (2021), Guidolin, La Cara and Marcellino (2021), Just and Echaust (2020), Sakurai and Kurosaki (2020), Yousfi et al. (2021), Derbali

Received August 2022; revised July 2023.

Key words and phrases. Coronavirus pandemic, graphical models, industry portfolios, particle filter, precision matrix, stochastic volatility.

et al. (2022), Dey et al. (2022)). Global “catastrophic” events, such as financial crises, often lead to sudden changes in the dependence structure across investments. Financial markets’ reaction to the pandemic appears to be no exception: around the end of February 2020 the Dow Jones and S&P 500 fell by 11% and 12%, respectively, marking the biggest weekly decline since the financial crisis of 2008, to the point that the Financial Times described such decline as the “quickest correction since the Great Depression.” Standard statistical approaches assume time-varying dependence to change smoothly over time, which appears to be an unrealistic assumption when investigating financial shocks. In this manuscript we develop statistical machinery to detect abrupt changes in the correlation structure among time series. Such machinery is employed to detect the impact of the COVID-19 pandemic on the U.S. stock market and, in particular, on cross-industry relationships.

There exists a vast literature on models for time-varying second moments. More specifically, there are two main approaches: *conditional volatility models*, as the well-known ARCH and GARCH (Bauwens, Laurent and Rombouts (2006), Engle (1982), Bollerslev, Engle and Nelson (1994), Bollerslev (1986), Engle and Bollerslev (1986), Silvennoinen and Teräsvirta (2009), Boudt et al. (2019)), and *stochastic volatility models* (e.g., Taylor (1982), Wiggins (1987), Hull and White (1987), Asai, McAleer and Yu (2006)). The former class specifies second moments at a certain time t as a deterministic function of past values of observations, volatility, and possibly covariance, given model parameters. The latter assumes second moments to follow a latent stochastic process, typically of Markovian structure, so that, even conditionally on all past information, volatility and correlations are unobservable random variables evolving over time. While stochastic volatility models are often more flexible and may achieve better inferential performances, when compared to conditional volatility approaches (Chan (2013), Clark and Ravazzolo (2015)), they are more difficult to estimate since the likelihood is typically intractable; see, for example, Nilsson (2016).

Within both classes a further distinction may be made between models that explicitly target the covariance matrix Σ_t and those focusing on the precision matrix $\Omega_t = \Sigma_t^{-1}$, specifically allowing for zero entries in Ω_t to favour parsimony. In this work we develop a Bayesian stochastic volatility model for the precision matrix. Specifically, the precision matrix at time t is modelled conditionally on a graph at time t , which describes the dependence structure among time series. As such, our work lies within the literature on Gaussian graphical models (GGMs) (see, e.g., Carvalho and West (2007), Prado and West (2010), Wang (2010), Wang and West (2009), Chandra, Mueller and Sarkar (2021)). This approach presents an important advantage: GGMs target conditional independence instead of marginal, leading to possible identification of macro-components (represented, for instance, by hubs and cliques in the graph) and safeguarding against spurious relationships, in the sense that GGMs aid understanding if pairwise correlations between variables can be fully or partially explained by their relationship with one or more additional variables. The identification of graph substructures is of particular interest in finance, where hubs may be interpreted as risk factors driving the market, while cliques represent financial instruments exposed to the same unobserved risk factor; see Figure 1 for a toy example clarifying the role of graph substructures in financial markets and, in particular, the interpretation of hubs as risk factors. Moreover, marginal dependence and Pearson correlation simply measure the pairwise comovement between two investments but do not provide any indication on whether a risk factor generating the comovement is specific to the two investments or it is common also to other financial instruments of interest. On the contrary, entries of the precision matrix represent comovements conditionally on the effect of all the other instruments considered in the model (Michis (2022)) so that an entry is nonzero if and only if the two returns are dependent conditionally on all other investments.

Changes over time of second moments can be smooth or abrupt. The focus of this work is on changes of the second type. Standard versions of the models cited so far assume variances




FIG. 1. Toy example of graphs substructures in a financial market with three assets. Market n.1: The asset Z acts as common risk factor in the market driving the dependence and is a hub. Market n.2: The return of all three assets X , Y , and Z are driven by an unobservable risk factor W , and the three assets form a clique. Market n.3: X and Y are driven by a common risk factor, not affecting the asset Z . Market n.4: The three assets are independent and the graph does not include any edge. Here $\epsilon_z, \epsilon_x, \epsilon_y$ are white noises.

and covariances changing smoothly over time and, in particular, between any two consecutive time points. For instance, [Carvalho and West \(2007\)](#) propose a Bayesian dynamic stochastic volatility model based on GGMs and conditional independence. In their construction the graph structure is kept constant over time, while the covariance matrix changes smoothly between any two consecutive time points. However, this feature is often in contrast with what is observed in financial markets, where volatility clusters (i.e., periods with a persistent value of volatility, that are interrupted by sudden changes) and correlation breakdowns (i.e., substantial changes in correlations during stressed times and financial crises) are well documented (see, e.g., [Von Furstenberg et al. \(1989\)](#), [Contessi, De Pace and Guidolin \(2014\)](#)).

To detect the possible effect of COVID-19 pandemic on the U.S. stock market, we analyse the correlation structure between nine industry portfolios, considering logarithmic weekly returns in the years 2019, 2020, and 2021. Weekly returns are computed starting from the daily returns available at Kenneth R. French’s Data Library. The Kenneth R. French Data Library also provides portfolio returns corresponding to more fragmented definitions of industries. However, our goal is to detect possible macro effects of the pandemic in the market, and for this reason, we focus on the industry portfolios described in [Table 1](#). Thanks to diversification within the same portfolio, the corresponding returns are less volatile and appropriately represent the tendency of a whole industry. More details about the construction of the portfolios can be found in [Section 5](#).

TABLE 1
Industry portfolios descriptions. SIC codes for each portfolio are available at
https://mba.tuck.dartmouth.edu/pages/faculty/ken.french/Data_Library/det_10_ind_port.html

Portfolio name	Industry	Description
NoDur	Consumer Nondurables	Food, Tobacco, Textiles, Apparel, Leather, Toys
Durbl	Consumer Durables	Cars, TVs, Furniture, Household Appliances
Manuf	Manufacturing	Machinery, Trucks, Planes, Chemicals, Off Furn, Paper
Enrgy	Energy	Oil, Gas, and Coal Extraction and Products
HiTec	Business Equipment	Computers, Software, and Electronic Equipment
Telcm	Telecommunications	Telephone and Television Transmission
Shops	Shops	Wholesale, Retail, and Some Services
Hlth	Health	Healthcare, Medical Equipment, and Drugs
Utils	Utilities	Utilities




FIG. 2. *Moving window: Graph estimated with the R package GGMselect (Bouvier et al. (2022)). The four central graphs are computed on 80 time points, while the first and the last graphs refer, respectively, to the first 70 and last 68 weeks. The window is moved by 20 time points at the time from one graph to the next.*

Figure 2 shows empirical estimates of a graph describing the conditional dependence structure over time. The estimates are obtained using a moving window of 80 weeks, shifting in steps of 20 weeks from one graph to the next. Graph estimates are obtained using an adaptive lasso approach, as implemented in the R package GGMselect (Bouvier et al. (2022)). From this preliminary analysis, changes in dependence are already evident as well as the role of hub of the manufacturing and consumer nondurables industries. However, it is difficult to determine whether the pandemic had an effect on the overall structure and when. Moreover, it is well-known that edge estimation in GGMs is sensitive to sample size and this empirical approach is highly dependent on the arbitrary choice of the window size (and corresponding sample size). To achieve our inferential goals, a sound modelling strategy is needed to be able to effectively infer the existence and location of change points, due to sudden changes, still borrowing information across the entire time horizon.

Sudden changes in volatility and dependence have been modelled generalizing either conditional or stochastic volatility models with the introduction of Markov switching regimes (see, among others, So, Lam and Li (1998), Haas, Mittnik and Paolella (2004), Bianchi et al. (2019), Caporale and Zekokh (2019)). However, in a frequentist framework, Markov switching approaches require an arbitrary choice of the number of different regimes and, consequently, ad hoc criteria for model choice (see also Cribben and Yu (2017)). A full Bayesian model for evolving graphs has been introduced by Warnick et al. (2018); similarly to Markov switching models, here the graph evolves assuming one possible state out of a finite number of exchangeable (not consecutive) states. Still in a Bayesian framework, Schwaller and Robin (2017) develop a strategy for change point detection in graphical models, which, in order to preserve computational tractability, assumes independent consecutive graphs. Their method has the often unrealistic implication that the graphs are estimated independently without borrowing information across the entire time horizon. Recently, Keshavarz, Michailidis and Atchadé (2020) have proposed an accurate algorithmic procedure, which employs multiple frequentist tests to detect abrupt changes in the precision matrix of a GGM; however, the

procedure does not provide estimates of the graph's structure. Alternatively, penalised likelihood techniques have been successfully employed for estimating dynamic GGMs (Bybee and Atchadé (2018), Cribben (2019), Danaher, Wang and Witten (2014), Gibberd and Nelson (2017), Hallac et al. (2017), Kolar and Xing (2012), Roy, Atchadé and Michailidis (2017), Yang and Peng (2020), Yang et al. (2015), Zhou, Lafferty and Wasserman (2010), Liu, Zhang and Liu (2021)); however, such approaches do not allow for uncertainty quantification on the number and temporal location of the abrupt changes and the graph topology. Similar limitations are shared also by algorithmic approaches, as the one proposed by Anastasiou, Cribben and Fryzlewicz (2022). A detailed comparison between our contribution and penalised likelihood approaches is provided in Sections 4 and 5.

In this work we introduce a Bayesian dynamic GGM to detect abrupt changes in the conditional dependence structure between time series. Our proposal is a piecewise constant stochastic volatility model. It favours sparsity at three levels by explicitly penalizing: (i) the number of change points, (ii) the number of edges within each graph, and (iii) the number of edges which are activated (appear) and deactivated (disappear) at each change point. In a Bayesian framework, it is straightforward, at least in principle, to perform posterior inference also on the number of change points and on their location. Finally, we note that our model does not assume global Gaussianity, which would imply the existence of a single Gaussian distribution for the entire temporal span. The assumption of global Gaussianity poses challenges in the analysis of financial returns, which are typically characterised by heavy tails and changes of behaviour. In our setup we assume local Gaussianity between two consecutive change points. Our assumption on the return distribution possibly accommodates the excess of kurtosis typically observed in financial returns' empirical distributions. More precisely, the introduction of change points allows the multivariate Gaussian distribution to change along the overall time horizon so that the observed empirical distribution can be thought of as having been drawn from a mixture of Gaussians, which can accommodate heavy tails (Cui (2012)).

The paper is structured as follows. In Section 2 the dynamic GGM is presented. Section 3 contains a discussion of the computational challenges, the proposed algorithm, and a simulation study to assess the performance of the sequential Monte-Carlo procedure. Results on simulated data and on the U.S. stock market data can be found in Sections 4 and 5, respectively. Section 6 concludes the paper with a discussion about future directions and extensions. In Supplementary Material we provide the dataset, R codes to reproduce all the results in this work, and additional results on the algorithm, simulation studies, and the application.

2. The dynamic Gaussian graphical model. We first introduce some definitions and notation. Let $G = (V, E)$ represent an undirected, simple, and unweighted graph, where $V = \{1, \dots, p\}$, $p \geq 1$, corresponds to the set of labelled nodes and $E \subseteq \{(h, k) \in V \times V : h < k\}$ the set of edges linking pairs of nodes. There is a one-to-one correspondence between G and its $p \times p$ adjacency binary matrix A , which is defined as follows. The element $A[h, k]$ on the h th row and k th column is equal to 1, when an edge exists between nodes h and k , and to 0 otherwise. Note that A is symmetric with zeros on the main diagonal, since G is simple. When each node corresponds to a random variable, the graph structure can be used to encode conditional independence so that an edge is present between vertices h and k if and only if the h th and k th random variables are dependent conditionally on all other variables in the graph (Lauritzen (1996)).

A powerful modelling tool is offered by GGMs, which assume that the distribution of the random variables represented by the nodes in V is multivariate Gaussian. Then the precision matrix Ω can be modelled conditionally on the graph so that the presence of an edge between two nodes in G implies a nonzero entry in the precision matrix between the corresponding

random variables, while the absence of an edge implies a zero entry. Let the cone M^+ be the space of symmetric positive-definite matrices on $\mathbb{R}^{p \times p}$. For graph G and adjacency matrix A , $M^+(G) \subset M^+$ denotes the set of the matrices, M , with $M[h, k] = 0$ if and only if $A[h, k] = 0$, for any $h \neq k$, so that $\Omega \in M^+(G)$.

In a time series setting, let $G_t = (V, E_t)$ describe the (conditional) dependence structure at time t between p time series, each corresponding to one node in V . We propose a prior distribution for the process $\{G_t, t \geq 1\}$, obtained by letting t vary, which lies in the class of stochastic volatility models.

Data are collected at common discrete time points $t = 1, 2, \dots, T$. We denote with Y_t the vector of observations at time t on the p variables (i.e., returns at week t for the considered industry portfolios) and with $Y_{1:T} = [Y_t]_{t=1}^T$ the $T \times p$ data matrix. We assume that, conditionally on a time-indexed collection of precision matrices $\{\Omega_t, t = 1, \dots, T\}$, the vectors of observations are normally distributed and independent over time, that is,

$$(1) \quad Y_t | \Omega_t \stackrel{\text{ind}}{\sim} N_p(0, \Omega_t^{-1}) \quad \text{for } t = 1, \dots, T,$$

where $N_p(\mu, \Sigma)$ denotes a p -variate Gaussian distribution with mean μ and covariance matrix Σ .

We model Ω_t conditionally on a graph at time t , G_t . Then, to allow for time-varying dependence structure among the p variables, we introduce a sequence of random change points. A time point t is said to be a *change point* if the dependence structure among the p observable variables changes between $t - 1$ and t , that is, if $G_t \neq G_{t-1}$ and, consequently, $\Omega_t \neq \Omega_{t-1}$. Let $c_{1:\kappa} = (c_1, c_2, \dots, c_\kappa)$ be the (possibly empty) vector of ordered change points, which, similarly to the precision matrices and the graphs, are unobserved. Here $\kappa \geq 0$ denotes the (random) number of change points. In what follows, we use the conventions $c_{1:0} = \emptyset$, $c_0 = 1$, and $c_{\kappa+1} = T + 1$. Note that between consecutive change points the graph and the corresponding precision matrix are kept constant. Given the sequence of graphs, $G_{1:T} = \{G_t, t = 1, \dots, T\}$, and change points, we assume that

$$(2) \quad \Omega_1 | G_1 \sim W_{G_1}(d, D)$$

and, for $t \geq 2$,

$$(3) \quad \Omega_t | \Omega_{t-1}, G_t, c_{1:\kappa} \sim \begin{cases} W_{G_t}(d, D) & \text{if } t \in \{c_1, \dots, c_\kappa\}, \\ \delta_{\Omega_{t-1}} & \text{otherwise,} \end{cases}$$

where, δ_x denotes the Dirac delta distribution at x and $W_G(d, D)$ the G -Wishart distribution (see [Dobra, Lenkoski and Rodriguez \(2011\)](#), [Roverato \(2002\)](#)), with shape parameter $d > 2$ and inverse scale matrix parameter $D \in M^+$. Its density w.r.t. the Lebesgue measure of dimension equal to the free elements of a matrix in $M^+(G)$ is

$$P(\Omega | G) = \frac{1}{I_G(d, D)} |\Omega|^{(d-2)/2} \exp\left\{-\frac{1}{2} \text{tr}(D\Omega)\right\}, \quad \Omega \in M^+(G).$$

The stated constraints for hyperparameters d and D suffice to ensure the integrability of the above density ([Diaconis and Ylvisaker \(1979\)](#)). The normalizing constant is equal to

$$I_G(d, D) = \int_{M^+(G)} |\Omega|^{(d-2)/2} \exp\left\{-\frac{1}{2} \text{tr}(D\Omega)\right\} d\Omega$$

and will be used later to compute the marginal likelihood of the data conditionally only on the graph structure.

To complete the model, we next describe the graph dynamics. Denote with A_t the adjacency matrix corresponding to G_t , then, for all $h, k \in \{1, \dots, p\}$ with $h < k$, we specify the prior distributions

$$(4) \quad A_1[h, k] | \omega \stackrel{\text{iid}}{\sim} \text{Bernoulli}\left(\frac{2\omega}{p-1}\right)$$

and, for $t \geq 2$,

$$(5) \quad A_t[h, k] | A_{t-1}[h, k], c_{1:\kappa}, z \stackrel{\text{ind}}{\sim} \begin{cases} \left| A_{t-1}[h, k] - \text{Bernoulli}\left(\frac{2z}{p-1}\right) \right| & \text{if } t \in \{c_1, \dots, c_\kappa\}, \\ \delta_{A_{t-1}[h, k]} & \text{otherwise.} \end{cases}$$

Notice that the hyperparameter $\omega \in [0, (p-1)/2]$ controls the graph sparsity so that the expected number of edges for the initial graph a priori equals $p\omega$, while the hyperparameter $z \in [0, (p-1)/2]$ controls the impact of an event on graph structure when a change point is reached, in particular, the (a priori) expected number of edges that will change is equal to pz . Our prior choice is reminiscent of the one proposed in Jones et al. (2005), who recommend setting a prior edge inclusion probability equal to $2/(p-1)$ so that the expected number of edges is p . Alternatively, we suggest to tune ω via a data-driven tuning technique so that a priori the expected number of edges for the graphs is equal to the number of edges detected by estimating one unique graph using all the time points; more details on this procedure are provided later in Section 4.1. Moreover, z controls the a priori abruptness of the changes in the graphs. As a general rule, we suggest to perform sensitivity analysis on this parameter. Nonetheless, in our analysis change point detection does not appear to be sensitive to the choice of the hyperparameters, and the graph recovery performance is limitedly affected by the choice of z (see Section C.5 in the Supplementary Material).

We note that alternative priors can be employed to model the precision matrix and the graph as, for instance, the graphical horseshoe (Li, Craig and Bhadra (2019)) and the prior proposed by Banerjee and Ghosal (2015). The former is a prior used directly on the precision matrix, which, contrarily to our approach, requires (arbitrary) thresholding of its entries in order to recover a sparse graph representation. The latter is more similar to our modelling strategy and consists of three elements: (i) Bernoulli priors for the entries of the adjacency matrix, conditionally on a maximum number of edges, (ii) a Laplace prior on the nonzero off-diagonal elements of the precision matrix, and (iii) an exponential prior for the diagonal elements, still imposing the positive definiteness of the matrix. This construction is still computationally intensive.

Equations (1) and (2)–(5) can be viewed as observation and state dynamics, respectively, of a hidden Markov model with the unobserved signal corresponding to the pair $\{(G_{c_i}, \Omega_{c_i}), i = 0, 1, \dots, \kappa\}$ (see Figure 3 for a graphical representation). For more details, see, for example, West and Harrison (2006).

Finally, the prior distribution for $c_{1:\kappa}$ is chosen as

$$(6) \quad \begin{aligned} c_{1:\kappa} | \kappa &\sim \text{Uniform}(\mathcal{T}_{\kappa\ell}), \\ \kappa | p_0 &\sim \text{Truncated-Geometric}(p_0) \quad \text{for } \kappa = 0, 1, \dots, K_{T\ell} \end{aligned}$$

for hyperparameter $p_0 \in (0, 1)$ so that the (a priori) number of expected change points is

$$(7) \quad \mathbb{E}[\kappa] = \frac{1-p_0}{p_0} \frac{1-(1-p_0)^{K_{T\ell}}(p_0 K_{T\ell} + 1)}{1-(1-p_0)^{K_{T\ell}+1}}.$$




FIG. 3. Graphical representation of the model conditionally on change points. Both graph and precision matrix are constant between change points. Observations (Y_t) are independent conditionally on the model parameters over time, while are i.i.d. between change points. Moreover, the precision matrices $\{\Omega_{c_i}, i = 1, \dots, \kappa\}$ are conditionally independent, given the sequence $\{G_{c_i}, i = 1, \dots, \kappa\}$.

Through appropriate choice of p_0 in (6), we are able to enforce the a priori desired level of sparsity on the number of change points. Nonetheless, in our sensitivity analysis, posterior inference is not sensitive to the choice of p_0 . Here $\mathcal{T}_{\kappa\ell}$ is the space of ordered κ -tuples $c_1 < \dots < c_\kappa$, with c_j in $\{2, \dots, T\}$, under the minimum-span constraint that $c_{j+1} - c_j \geq \ell$, for any $j = 0, 1, \dots, \kappa$, with the convention $c_0 = 1$ and $c_{\kappa+1} = T + 1$. Moreover, $K_{T\ell}$ is the maximum number of change points compatible with the minimum-span constraint. Notice that for $K_{T\ell}$ of moderate size, the second term in the product in (7) is negligible and the Truncated Geometric in (6) approximates a standard Geometric distribution on $\mathbb{N}_0 = \{0, 1, \dots\}$. The imposition of the minimum-span constraint, defined by ℓ , ensures likelihood identifiability between change points, leading to more stable computations and robust inference. Notice that, between any two change points, the sample covariance/precision matrix has $p(p + 1)/2$ entries. To guarantee likelihood identifiability of the model, we need the number of data points between any two change points, that is, $(c_{j+1} - c_j)$, to be greater than p . This poses a trade-off: the lower ℓ , the more flexible the change-point detection procedure, and the higher ℓ , the more stable the estimates of the precision matrices. In the simulation study and in the application, we set ℓ equal to $p + 2$. When $p = 9$, this means that each precision matrix containing 45 entries is estimated with at least 99 single data points.

3. Bayesian inference via sequential Monte-Carlo. The dynamic GGM, proposed herein, is a hierarchical model, with the first layer represented by the change points, $c_{1:\kappa}$, the second by graphs and precision matrices (G_{c_j}, Ω_{c_j}) , and the third by the observations. Markov chain Monte-Carlo (MCMC) methods developed directly on such space of unobserved variables would face major challenges. We ignore the precision matrices in this discussion, as they are later integrated out. Gibbs-type approaches would involve reversible-jump MCMC (Green (1995)), thus requiring the design of a joint update on the “model” space (as determined by the change-points) and model “parameters” of varying dimension (corresponding to the graphs). This joint space is entangled, with very limited space for maneuvering, as updates on the graph space would be heavily constrained by the strong prior Markovian dependencies amongst graphs.

Instead, we perform computationally effective posterior inference for the dynamic GGM through a tailored sequential Monte-Carlo (SMC) algorithm. The proposed Particle MCMC (PMCMC) method is quite appropriate for exploiting the hidden Markov model structure conditionally on the change points, and naturally disentangles the updates on the change points and the latent Markovian signal; see, for example, Karagiannis and Andrieu (2013)

and Persing et al. (2015) for related ideas. The proposed PMCMC algorithm is better understood as comprised of an “outer” cycle and an “inner” cycle. In the former, the change points are updated via a reversible jump Metropolis–Hasting (M–H) algorithm. In the latter, a particle filter of enhanced performance, due to a combination of adaptive tempering, dynamic resampling, and mutation steps, is employed to sample the sequence of graphs and compute the acceptance probability of the outer algorithm. In particular, for each M–H step the inner component provides an unbiased estimate of the conditional likelihood, given the proposed change point sequence together with a corresponding proposed graph sequence. Adaptive tempering and resampling steps are used to improve the robustness of such estimate, while the mutation step is used both to bring particles closer to the modal region of the likelihood and to avoid depletion of the number of unique particles, which can otherwise be a consequence of successive resampling and tempering.

The overall PMCMC algorithm is well-understood as an “exact-approximate” one, in the sense that it targets the correct posterior on the space of graphs and change points, thanks to the unbiasedness (and positivity) of the estimator provided by the inner particle filter.

3.1. *Outer component.* The key component in the development of the M–H step is the choice of proposal distribution, $q(c'_{1:\kappa'}|c_{1:\kappa})$, where $c_{1:\kappa}$ and $c'_{1:\kappa'}$ are the current and proposed collections of ordered change points, respectively. Starting from $c_{1:\kappa}$, one of four alternative events (namely, a birth, a death, a global move, or a local move) generates the proposed new value. With probabilities equal to $P(B|c_{1:\kappa})$, $P(D|c_{1:\kappa})$, $P(M_{\text{glob}}|c_{1:\kappa})$, and $P(M_{\text{loc}}|c_{1:\kappa})$, one of the following four events takes place, respectively: a new change point is added to the current set (birth); a change point is removed from the current set (death); one of the existing change points is moved to another position (global move); one of the existing change points is moved to another position in-between its neighbours (local move). When a new change point, c^* , is created, c^* is chosen uniformly over the set $B(c_{1:\kappa}, \ell) \subset \{2, \dots, T\}$ of allowed positions (i.e., satisfying the minimum-span constraint), of size $|B(c_{1:\kappa}, \ell)| =: n(c_{1:\kappa}, \ell) \geq 0$. Thus, in the birth scenario

$$(8) \quad q(c'_{1:\kappa'}|c_{1:\kappa}) = \frac{P(B|c_{1:\kappa})}{n(c_{1:\kappa}, \ell)}, \quad \kappa' = \kappa + 1, \quad c'_{1:\kappa'} = c_{1:\kappa} \cup c^*, \quad c^* \in B(c_{1:\kappa}, \ell).$$

When a change point, c' , is removed, the change point is chosen uniformly among the current change points, that is, in the death scenario

$$(9) \quad q(c'_{1:\kappa'}|c_{1:\kappa}) = \frac{P(D|c_{1:\kappa})}{\kappa}, \quad \kappa' = \kappa - 1, \quad c'_{1:\kappa'} = c_{1:\kappa} \setminus c', \quad c' \in c_{1:\kappa}.$$

To improve mixing and posterior exploration, we also introduce two move steps. When a change point is moved, first, a change point $c' \in c_{1:\kappa}$ is selected uniformly among the current change points and removed. Then, if the step is a global move, a new change point is selected uniformly in $B(c_{1:\kappa} \setminus c', \ell)$. If instead the step is a local move, a new change point c^* is selected with probability proportional to

$$(10) \quad \exp\{-\lambda|c^* - c'|\} \mathbb{1}_{\{c^* \in [\bar{c}_l, \bar{c}_r]\}}$$

for algorithmic parameter $\lambda > 0$, with

$$\bar{c}_l = c_l + \ell, \quad \bar{c}_r = c_r - \ell,$$

where c_l and c_r denote the left-side and right-side neighbours of c' in $c_{0:\kappa+1}$. Thus, the proposal kernel for a move step is

$$(11) \quad q(c'_{1:\kappa'}|c_{1:\kappa}) = \frac{P(M_{\text{glob}}|c_{1:\kappa})}{\kappa n(c_{1:\kappa} \setminus c', \ell)} + \frac{P(M_{\text{loc}}|c_{1:\kappa})e^{-\lambda|c^* - c'|}}{\kappa \sum_{\chi=\bar{c}_l}^{\bar{c}_r} e^{-\lambda|c^* - \chi|}} \mathbb{1}_{\{c^* \in [\bar{c}_l, \bar{c}_r]\}},$$

$$\kappa' = \kappa, \quad c'_{1:\kappa'} = c_{1:\kappa} \setminus c' \cup c^*, \quad c' \in c_{1:\kappa}, \quad c^* \in B(c_{1:\kappa} \setminus c', \ell).$$

Note that we prefer to write the joint kernel for both the global and local move, as the same configuration $c'_{1:\kappa}$ might be reached by both type of moves. This needs to be accounted for when computing the M–H acceptance probability to ensure detailed balance.

Finally, we choose the event probabilities as

$$(12) \quad P(B|c_{1:\kappa}) = \begin{cases} 1, & \text{if } \kappa = 0 \\ 0, & \text{if } n(c_{1:\kappa}, \ell) = 0 \\ q_B, & \text{otherwise} \end{cases} \quad P(D|c_{1:\kappa}) = \begin{cases} 0, & \text{if } \kappa = 0 \\ q'_D, & \text{if } n(c_{1:\kappa}, \ell) = 0 \\ q_D, & \text{otherwise} \end{cases}$$

$$\text{and } P(M_{\text{glob}}|c_{1:\kappa}) = P(M_{\text{loc}}|c_{1:\kappa}) = [1 - P(B|c_{1:\kappa}) - P(D|c_{1:\kappa})]/2.$$

The Metropolis–Hastings acceptance probability is equal to

$$1 \wedge \frac{P(Y_{1:T}|c'_{1:\kappa'})P(c'_{1:\kappa'})q(c_{1:\kappa}|c'_{1:\kappa'})}{P(Y_{1:T}|c_{1:\kappa})P(c_{1:\kappa})q(c'_{1:\kappa'}|c_{1:\kappa})}$$

where $P(Y_{1:T}|c_{1:\kappa})$ is the marginal likelihood of the data given the change points, that is,

$$P(Y_{1:T}|c_{1:\kappa}) = \int P(Y_{1:T}|G_{1:T})P(G_{1:T}|c_{1:\kappa}) dG_{1:T}$$

with

$$P(Y_{1:T}|G_{1:T}) = \int P(Y_{1:T}|G_{1:T}, \Omega_{1:T})P(\Omega_{1:T}|G_{1:T}) d\Omega_{1:T}.$$

Since $P(Y_{1:T}|c_{1:\kappa})$ is not available in closed form, it needs to be estimated. Algorithm 1 contains the pseudo-code for the outer part of the algorithm described in this section. In the next section, we describe the SMC algorithm to approximate the marginal likelihood.

3.2. Inner component: Particle filter. As already mentioned, the inner component of the algorithm is used to compute $P(Y|c_{1:\kappa})$, that is, the likelihood values required in the acceptance probability of the outer M–H and to provide proposed samples of the graphs to be accepted or rejected by the outer algorithm.

To compute the marginal likelihood given the change point sequence, a standard bootstrap particle filter with multinomial resampling carried out at each change point, samples $N \geq 1$ particles $\{G_{c_j}^{(n)}\}_{n=1}^N$, for $0 \leq j \leq \kappa$, from the joint distribution

$$\prod_{n=1}^N P(G_{c_0}^{(n)}) \times \prod_{j=1}^{\kappa} \left\{ \prod_{n=1}^N \left(\sum_{l=1}^N \frac{w_{j-1}^{(l)}}{\sum_{m=1}^N w_{j-1}^{(m)}} P(G_{c_j}^{(n)}|G_{c_{j-1}}^{(l)}) \right) \right\},$$

where the unnormalised weights are defined as

$$w_j^{(n)} = P(Y_{c_j:c_{j+1}-1}|G_{c_j}^{(n)}), \quad 1 \leq n \leq N, 0 \leq j \leq \kappa.$$

The unbiased estimate $\hat{P}(Y|c_{1:\kappa})$ of $P(Y|c_{1:\kappa})$ could then be obtained as

$$(13) \quad \hat{P}(Y_{1:T}|c_{1:\kappa}) = \prod_{j=0}^{\kappa} \left(\frac{1}{N} \sum_{n=1}^N w_j^{(n)} \right).$$

However, it is often the case that further algorithmic advances must complement the standard particle filter to control the variance of the estimate (13). It is well-understood that such variability is critically linked to the performance of the overall PMCMC algorithm; see, for example, Doucet et al. (2015), Pitt et al. (2012), Sherlock et al. (2015), where in various model settings, standard deviations centred around 1 are proposed for the estimate of the

Algorithm 1 Outer algorithm—reversible jump M-H

Input: change point sequence $c_{1:k}$; $\mathbf{B}(c_{1:k}, \ell)$; likelihood value $P(Y_{1:T}|c_{1:k})$; prior value $P(c_{1:k})$.

Output: new change point sequence $\tilde{c}_{1:\tilde{k}}$; $\mathbf{B}(\tilde{c}_{1:\tilde{k}}, \ell)$; likelihood value $P(Y_{1:T}|\tilde{c}_{1:\tilde{k}})$; prior value $P(\tilde{c}_{1:\tilde{k}})$.

Sample event E from $\{B, D, M_{\text{glob}}, M_{\text{loc}}\}$ according to (12);

if $E = B$ **then**

Sample a new change point uniformly from $\mathbf{B}(c_{1:k}, \ell)$ and propose $c'_{1:k'} = c_{1:k} \cup c'$;

Compute $q(c'_{1:k'}|c_{1:k})$ according to (8);

else

Sample uniformly and remove a change point c' from $c_{1:k}$;

if $E = M_{\text{glob}}$ or M_{loc} **then**

if $E = M_{\text{glob}}$ **then**

Sample a new change point c^* uniformly from $\mathbf{B}(c_{1:k} \setminus c', \ell)$;

else

Sample a new change point c^* from the interval $[\bar{c}_l, \bar{c}_r]$ according to (10);

Propose $c'_{1:k'} = c_{1:k} \setminus c' \cup c^*$ and compute $q(c'_{1:k'}|c_{1:k})$ according to (11);

else

Propose $c'_{1:k'} = c_{1:k} \setminus c'$ and compute $q(c'_{1:k'}|c_{1:k})$ according to (9);

Determine $\mathbf{B}(c'_{1:k'}, \ell)$ and compute $q(c_{1:k}|c'_{1:k'})$ according to (8), (9), or (11), respectively;

Compute prior $P(c'_{1:k'})$ and likelihood $P(Y_{1:T}|c'_{1:k'})$ for proposed configuration;

Sample u from a Uniform(0, 1);

if $u \leq \frac{P(Y_{1:T}|c'_{1:k'})P(c'_{1:k'})q(c_{1:k}|c'_{1:k'})}{P(Y_{1:T}|c_{1:k})P(c_{1:k})q(c_{1:k}|c'_{1:k'})}$ **then**

Return $c'_{1:k'}$, $\mathbf{B}(c'_{1:k'}, \ell)$, $P(Y_{1:T}|c'_{1:k'})$, $P(c'_{1:k'})$.

else

Return $c_{1:k}$, $\mathbf{B}(c_{1:k}, \ell)$, $P(Y_{1:T}|c_{1:k})$, $P(c_{1:k})$.

logarithm of the normalising constant, with exponential decay in performance for PMCMC reported when the standard deviation exceeds a (not too high) threshold.

A standard approach to reduce standard deviation for given number of particles is via the application of tempering, that is, introduction of a sequence of temperatures together with corresponding mutation steps. Such an approach has been shown, in cases, to reduce the required number of particles for a target error from exponential to quadratic in the number T of log-likelihood terms; see, for example, Beskos, Crisan and Jasra (2014), Ruzayqat et al. (2022). Application of tempering will indeed be critical for the class of models we are considering in this work, as shown in Section 3.2.4. The temperatures are determined on-the-fly, thus avoiding the introduction of additional tuning parameters in the algorithm. The complete algorithm can be understood as a particle filter applied on a Feynman–Kac model (Del Moral (2004)) that we extend to include additional Markov iterations and potentials. The overall approach is summarised in Algorithm 2. We stress that the particle filter that includes the tempering and mutation steps will still provide unbiased estimates of $p(Y_{1:T}|c_{1:k})$, and the induced overall PMCMC method will still target the exact posterior $P(c_{1:k}, G_{c_0:c_k}|Y_{1:T})$; see, for example, the original paper on PMCMC (Andrieu, Doucet and Holenstein (2010)) for a detailed justification.

3.2.1. *Preliminary run—determination of temperatures.* Within Algorithm 2 the sequence of temperatures is treated as given. In practice, at each iteration of the outer algorithm, the temperatures are determined by a separate preliminary and independent execution of the

Algorithm 2 Inner algorithm—particle filter with tempering

Input: Data $Y_{1:T}$; change points $c_{1:\kappa}$; number of particles N ; ESS threshold ϵ ; number S_j of temperatures for graph j , $\phi_{0,j} \equiv 0$, $\phi_{0,S_j+1} \equiv 1$, $0 \leq j \leq \kappa$; number of mutation steps $M \geq 1$; temperatures $\{\phi_{1,j}, \dots, \phi_{S_j,j}\}$; M-H kernel $\bar{P}_{j,s}(G_{c_j}|G_{c_j}, G_{c_{j-1}})$, $0 \leq j \leq \kappa$, $1 \leq s \leq S_j + 1$; $\bar{P}_{j,s}^M$ denotes M iterations of such a kernel; hyper-parameters ω, z, d, D .

Output: Unbiased estimate $\hat{P}(Y_{1:T}|c_{1:\kappa}) > 0$ and proposed sequence $(G_{c_0}, \dots, G_{c_\kappa})$.
(Actions over n are understood to be repeated for $1 \leq n \leq N$.)

Set $\hat{P} = 1$;
for j in $0 : \kappa$ **do**
 if $j = 0$ **then**
 Sample $G_{c_0}^{0,n} \stackrel{\text{iid}}{\sim} P(G_{c_0})$ and set $w_0^{0,n} = 1$;
 else
 Initialise particles $G_{c_j}^{0,n} \stackrel{\text{ind.}}{\sim} P(G_{c_j}|G_{c_{j-1}}^{S_{j-1},n})$;
 Initialise weights $w_j^{0,n} = w_{j-1}^{S_{j-1}+1,n}$;
 for s in $1 : S_j + 1$ **do**
 Calculate weights $w_j^{s,n} = w_j^{s-1,n} \cdot [P(Y_{c_j:c_{j+1}-1}|G_{c_j}^{s-1,n})]^{\phi_{j,s}-\phi_{j,s-1}}$;
 if $\text{ESS}(w_j^{s,1:N}) < \epsilon N$ **then**
 $\hat{P} \leftarrow \hat{P} \cdot \frac{1}{N} \sum_{n=1}^N w_j^{s,n}$;
 Resample $\{G_{c_j}^{s-1,n}, G_{c_{j-1}}^{S_{j-1},n}\}$ according to the weights $\{w_j^{s,n}\}$;
 Mutate particles, that is, sample $G_{c_j}^{s,n} \stackrel{\text{ind.}}{\sim} \bar{P}_{j,s}^M(G_{c_j}|G_{c_j}^{s-1,n}, G_{c_{j-1}}^{S_{j-1},n})$;
 Set $w_j^{s,n} = 1$;
 else
 set $G_{c_j}^{s,n} = G_{c_j}^{s-1,n}$;
 Sample a graph G_{c_κ} from $\{G_{c_\kappa}^{S_\kappa+1,n}, w_\kappa^{S_\kappa+1,n}\}_n$ and retrieve its genealogy, $(G_{c_0}, \dots, G_{c_\kappa})$, amongst particles $\{G_{c_0}^{S_0+1,n}\}_n, \dots, \{G_{c_{\kappa-1}}^{S_{\kappa-1}+1,n}\}_n$.
 Return $\hat{P}(Y_{1:T}|G_{1:T}) = \hat{P}$ and the proposed sequence $(G_{c_0}, \dots, G_{c_\kappa})$.

particle filter that identifies and stores the temperatures that are later used within Algorithm 2. Similar ideas have been used in the SMC literature; see, for example, [Jasra et al. \(2011\)](#). That is, we first determine the temperatures according to a target effective sample size (ESS) and then apply the particle filter in Algorithm 2, with the obtained temperatures, to produce a robust unbiased estimator of the likelihood needed to compute the acceptance probability in Algorithm 1.

We describe here how to compute the temperatures $\{\phi_{1,j}, \dots, \phi_{S_j,j}\}$ used in Algorithm 2, for $S_j \geq 0$, where $\phi_{0,j} \equiv 0$, $\phi_{0,S_j+1} \equiv 1$, and $0 \leq j \leq \kappa$. Note that S_j , the number of temperatures, can vary across graphs at different change points; that is, it depends on j . Within this subsection particles and weights $G_{c_j}^{s,n}, w_j^{s,n}$ refer to such a preliminary execution of the particle filter. The temperatures are selected on-the-fly, based on the target ESS, denoted by ESS_0 , with $\text{ESS}_0 = \epsilon N$, $\epsilon \in (0, 1)$. Consider the current collection of particles and weights, $G_{c_j}^{s-1,n}$ and $w_j^{s-1,n}$, generated while filtering data points $Y_{c_j:c_{j+1}-1}$, and the corresponding

likelihood factor $[P(Y_{c_j:c_{j+1}-1}|G_{c_j})]^{\phi_{s-1,j}}$ up to the present step. Then the next temperature ϕ_s is determined so that the ESS equals the target, that is, $ESS(\phi_s) = ESS_0$. More precisely, define the next set of weights as function of the next temperature

$$w_j^{s,n}(\phi) = [P(Y_{c_j:c_{j+1}-1}|G_{c_j}^{s-1,n})]^{\phi-\phi_{s-1,j}},$$

and consider

$$ESS(\phi) := \frac{(\sum_{n=1}^N w_{s,j}^n(\phi))^2}{\sum_{n=1}^N (w_{s,j}^n(\phi))^2} = \epsilon N$$

whose solution—assuming it exists within $(\phi_{s-1,j}, 1]$ —provides the next temperature $\phi_{s,j}$. The solution is obtained with a simple fast bisection method. If $ESS(1) \geq \epsilon N$, we simply select $\phi = 1$. With this procedure we obtain all temperatures related to data $Y_{c_j:c_{j+1}-1}$, and we can then proceed to the next filtering step. We set ϵ to 1/2 to obtain a minimum ESS of $N/2$, which is a common choice (see, e.g., [Chopin and Papaspiliopoulos \(2020\)](#), p.133); see [Algorithm 3](#) for a detailed description.

Algorithm 3 Inner algorithm—temperature tuning (preliminary particle filter)

Input: Data $Y_{1:T}$; change points $c_{1:k}$; number of particles N ;
 hyper-parameters w, z, d, D ; ESS threshold ϵ ; mutation steps $M \geq 1$;
 M-H kernel $\bar{P}_{j,s}(G_{c_j}|G_{c_j}, G_{c_{j-1}})$ that preserves the law
 $[P(Y_{c_j:c_{j+1}-1}|G_{c_j})]^{\phi_{j,s}} \cdot P(G_{c_j}|G_{c_{j-1}})$,
 $\bar{P}_{j,s}^M$ denotes M iterations of such a kernel.
Output: Temperatures $\{\phi_{1,j}, \dots, \phi_{S_j,j}\}$, $S_j \geq 0$, $\phi_{0,j} \equiv 0$, $\phi_{0,S_j+1} \equiv 1$, $0 \leq j \leq \kappa$.
 (Actions over n are understood to be repeated for $1 \leq n \leq N$.)

for j in $0 : \kappa$ **do**
 if $j = 0$ **then**
 Sample $G_{c_0}^{0,n} \stackrel{iid}{\sim} P(G_{c_0})$ and set $w_0^{0,n} = 1$;
 else
 Initialise particles $G_{c_j}^{0,n} \stackrel{ind.}{\sim} P(G_{c_j}|G_{c_{j-1}}^{S_{j-1},n})$;
 Initialise weights $w_j^{0,n} = w_{j-1}^{S_{j-1}+1,n}$;
 $s \leftarrow 1$; $\phi_{0,j} \leftarrow 0$; $\phi_{s,j} \leftarrow 1$;
 while $\phi_{s,j} \neq \emptyset$ **do**
 Find $\phi_{s,j} \in (\phi_{s-1,j}, 1]$ so that $ESS(\phi_{s,j}) \geq \epsilon N$;
 if $\phi_{s,j} = \emptyset$ **then**
 $S_j = s - 1$;
 else
 Set $w_j^{s,n} = [P(Y_{c_j:c_{j+1}-1}|G_{c_j}^{s-1,n})]^{\phi_{s,j}-\phi_{s-1,j}}$;
 Resample $\{G_{c_j:c_{j+1}-1}^{s-1,n}, G_{c_{j-1}}^{S_{j-1},n}\}$ according to the weights $\{w_j^{s,n}\}$;
 Set $w_j^{s,n} = 1$;
 Mutate particles, that is, sample $G_{c_j}^{s,n} \stackrel{ind.}{\sim} \bar{P}_{j,s}^M(G_{c_j}|G_{c_j}^{s-1,n}, G_{c_{j-1}}^{S_{j-1},n})$;
 $s \leftarrow s + 1$
 Set $w_j^{S_j+1,n} = [P(Y_{c_j:c_{j+1}-1}|G_{c_j}^{s-1,n})]^{1-\phi_{s-1,j}}$;
 Return $\{\phi_{1,j}, \dots, \phi_{S_j,j}\}$, $S_j \geq 0$, $0 \leq j \leq \kappa$.

3.2.2. *Determination of mutation kernel \bar{P} .* The mutation kernel \bar{P} is used within the algorithm to jitter particles and move them toward the centre of the support of each filtering distribution under consideration during a full application of the particle filter. The addition of mutation steps has been shown to be, in many cases, critical, both in theoretical and experimental works; see, for example, Beskos, Crisan and Jasra (2014), Ruzayqat et al. (2022) and Llopis et al. (2018), van Leeuwen et al. (2021), respectively. In Section 3.2.4 we illustrate such impact for the specific model at hand through a simulation study.

For the overall algorithm to ensure a correct particle filter on an extended space, the user-specified mutation kernel $\bar{P}_{j,s}(G_{c_j}|G_{c_j}, G_{c_j})$ must have invariant distribution

$$G_{c_j} \mapsto [P(Y_{c_j:c_{j+1}-1}|G_{c_j})]^{\phi_{j,s}} \times P(G_{c_j}|G_{c_{j-1}}),$$

where we use the convention that $G_{c_0-1} = \emptyset$ in which case the rightmost term becomes the prior defined by (4). This is readily achieved via a M–H step. That is, for each current segment $c_j : c_{j+1} - 1$, temperature $\phi_{j,s}$, graph G_{c_j} , with adjacency matrix $A_{c_j} = A_{c_j}[h, k]$, and given $G_{c_{j-1}}$, we define a proposed graph G'_{c_j} , with adjacency matrix $A'_{c_j} = A'_{c_j}[h, k]$, using the symmetric transition

$$(14) \quad A_{c_j}[h, k]'|A_{c_j}[h, k] \stackrel{\text{ind}}{\sim} \left| A_{c_j}[h, k] - \text{Bernoulli}\left(\frac{2s_0}{p-1}\right) \right|$$

for algorithmic tuning parameter $s_0 \in (0, (p-1)/2)$. Thus, under the proposal in (14), the expected number of flips in the edges is $s_0 \cdot p$. The acceptance probability for the mutation step is

$$1 \wedge \frac{[P(Y_{c_j:c_{j+1}-1}|G'_{c_j})]^{\phi_{j,s}} \times P(G'_{c_j}|G_{c_{j-1}})}{[P(Y_{c_j:c_{j+1}-1}|G_{c_j})]^{\phi_{j,s}} \times P(G_{c_j}|G_{c_{j-1}})}.$$

3.2.3. *Likelihood given the graph structure.* An important quantity required within the particle filter is the marginal likelihood

$$P(Y_{c_j:c_{j+1}-1}|G_{c_j}) = \int_{M^+(G_{c_j})} P(Y_{c_j:c_{j+1}-1}|\Omega_{c_j})P(\Omega_{c_j}|G_{c_j}) d\Omega_{c_j}.$$

Since the G -Wishart law is conjugate, we can integrate out the precision matrices $\Omega_{1:T}$. That is, we have (Atay-Kayis and Massam (2005))

$$P(Y_{c_j:c_{j+1}-1}|G_{c_j}) = \frac{1}{(2\pi)^{(c_{j+1}-c_j)p/2}} \frac{I_{G_{c_j}}(d + (c_{j+1} - c_j), D + H_j)}{I_{G_{c_j}}(d, D)},$$

where, for $j = 0, \dots, \kappa$,

$$H_j = \sum_{i=c_j}^{c_{j+1}-1} Y_i Y_i^\top.$$

Notice that, while computing the likelihood of the graphs, we marginalize over Ω_t , and thus, the particles (from the inner algorithm) consist of only the graphs. However, after running the particle filter and thanks to the conjugacy properties of the G -Wishart law, a straightforward independent sampler can be used to get both marginal and conditional posterior of the precision matrices, where with conditional posterior distribution we mean the posterior distribution conditional on the point estimates of the graphs. The normalising constant of the G -Wishart prior can be factorised (Roverato (2002), Uhler, Lenkoski and Richards (2018)), that is, for a given graph G ,

$$(15) \quad I_G(d, D) = \frac{\prod_{m=1}^r I_{G_{P_m}}(d, D_{P_m})}{\prod_{m=2}^r I_{G_{S_m}}(d, D_{L_m})},$$

where $P_1; L_2, P_2; \dots; P_r, L_r$, is a perfect sequence of prime components and corresponding minimal separators of G (see, e.g., Chapter 2 of Lauritzen (1996), for details), and D_{P_m} is the submatrix of D corresponding to the rows and columns in P_m . In the case of a *decomposable* graph G , all prime components are complete graphs. For complete graphs the G -Wishart distribution coincides with the hyper-Wishart distribution (Dawid and Lauritzen (1993)), for which an analytical expression for the normalising constant is available,

$$I_{G_{P_m}}(d, D_{P_m}) = \frac{2^{(d+p_m-1)p_m/2} \Gamma_{p_m} \left(\frac{d+p_m-1}{2} \right)}{|D_{P_m}|^{(d+p_m-1)}}.$$

Here $\Gamma_d(\cdot)$ is the multivariate Gamma function of dimension d , and p_m is the dimension of D_{P_m} . Note that, by construction, the minimal separators are complete subgraphs of G ; thus, the terms in the denominator in (15) are analytically computable. For a general, *nondecomposable* graph G , Atay-Kayis and Massam (2005), Carvalho, Massam and West (2007), Dellaportas, Giudici and Roberts (2003), Roverato (2002) propose Monte-Carlo methods for the approximation of $I_G(d, D)$. Herein to compute the normalizing constant, we employ the method of Atay-Kayis and Massam (2005) implemented in the function `gnorm` of the R package `BDgraph` (Mohammadi, Wit and Dobra (2022)). When dealing with large number of nodes, the implementation of more sophisticated algorithms, such as the exchange algorithm by Murray, Ghahramani and MacKay (2006), is advisable (see, e.g., Cheng and Lenkoski (2012), van den Boom et al. (2022), for an application to GGMs).

3.2.4. Evaluation of SMC approximation. The inner SMC algorithm provides unbiased estimates of the marginal likelihood conditionally on the change points. The variability of such estimates depends on the number of particles, N , and the effect of the tempering and mutation steps, with the number of the latter, M , specified by the user. Thus, a trade-off is posed between accuracy of estimates and computational time.

To assess the effect of the number of particles and the mutation step and, in general, obtain insights into the performance of the SMC component, we perform a series of simulation studies. We simulate data for $p = 10$ nodes and $T = 200$ observation instances. We then fix the change point sequence to its known true value and carry out 30 independent executions of the SMC algorithm, for each different combination of $N \in \{200, 500, 750\}$ and $M \in \{0, 5, 10, 20\}$. Recall that the mutation steps are performed only when the ESS falls below the threshold ϵN , where ϵ is here fixed to be $N/2$. We consider two data generating mechanisms. The first (Scenario A) has no change points, and the p variates are mutually independent (see Figure B.1.1 (a) in the Supplementary Material for the corresponding graph, Franzolini et al. (2024)). In Scenario B we set a change point at $t = 70$, and the two graphs (before and after the change point) encode some nontrivial dependence. The full graph structure of Scenario B is described later in Section 4.1 and displayed in Figure B.1.2 of the Supplementary Material. Figure A.1 of the Supplementary Material (Franzolini et al. (2024)) and Figure 4 show the box plots of the estimates of the log-likelihood, the standard deviation, and the running time of the inner SMC algorithm, coded in R (and run with an Intel Xeon W-1250 processor), under scenarios A and B, respectively.

Under Scenario A the variability of the estimates is limited, as expected, for all pairs (N, M) since data are simulated under the assumption of independence with no change points. However, in real scenarios, as the analysis of financial markets, this is highly unlikely to be the case, and the computational machinery here developed is essential. Figure 4 shows similar box plots obtained under the more realistic and challenging Scenario B. Here we obtain higher variability with values of standard deviation ranging from 2.986 to 29.297. These results highlight the importance of the tempering and mutation steps for the overall




FIG. 4. Log-likelihood estimates for Scenario B obtained with the particle filter by fixing the change points to the truth. Distinct box plots correspond to different numbers of particles N and/or mutations steps M . For each pair (N, M) , we run the algorithm 30 times and obtain the log-likelihood estimates. Each box plot shows the distribution of such estimates. The variability of the estimates is rather limited for all pairs (N, M) , provided that $M \geq 5$.

algorithmic performance (even accounting for the increased computational time) and, in particular, for recovering the complex dependence structures. Their role is essential in reducing variability of the estimates of the normalising constant. Lastly, notice that computational time is increased compared to Scenario A as a consequence of: (i) the presence of a change point, (ii) the computational complexity of the Monte Carlo iterations used to compute $I_G(d, D)$, which increases when particles concentrate on less sparse graphs, and (iii) the increased number of times the ESS threshold is reached.

4. Model performance on simulated data. We investigate model performance through a simulation study considering different scenarios. For each scenario we simulate data for $T = 200$ observation times. There is one change point in Scenario 1 and three change points in Scenario 2. In Scenario 3 the dependence structure, as captured by the precision matrix, presents smooth changes, and thus, our model is misspecified. The number of nodes is $p = 10$ in Scenarios 1 and 3, and $p = 20$ in Scenario 2. In terms of the abruptness of changes in the precision matrix, Scenario 2 presents highly-abrupt changes, Scenario 1 presents mildly-abrupt changes, and Scenario 3 presents smooth changes. For a detailed description of these three scenarios, see Sections 4.1, 4.2, and Section B.4 of the Supplementary Material (Franzolini et al. (2024)).

Moreover, we also consider two additional scenarios (Scenarios A.0 and A.1) with $p = 10$ nodes and no change points to test the baseline performance of the method and perform a sanity check in the case of no signal of the data.

To carry out posterior inference, we run the algorithm to estimate the change point sequence, using $N = 200$ particles, $M = 10$ mutation steps, and performing 10,000 iterations of the outer component, of which the first 2000 are discarded as burn-in. Then we re-run only the inner SMC with $N = 1000$ particles and $M = 20$ mutation steps to obtain the graph estimates conditionally on the maximum-a-posteriori (MAP) estimate of the change point sequence obtained in the first step. When the true sequence of change points is the null set (Scenarios A.0 and A.1), we initialise the MCMC chain at $(c_1 = 51, c_2 = 101, c_3 = 151)$; whereas when the graph and/or the precision matrix changes (Scenarios 1, 2, and 3), we initialise the chain to the state of zero change points. In real applications, when the change points are unknown, we suggest initialising the chain to no change points. The adopted initialization for these simulations better tests the convergence speed of the algorithm.

The inference results for no change points show the expected good performance of the model in terms of both identification of change points and recovery of the dependence structure. The posterior concentrates on the true state of no change points, with posterior probability of no-change points not falling below 0.98 in all replicates. The area under the curve (AUC) for edge detection is approximately 1. We refer to section B of the Supplementary Material (Franzolini et al. (2024)) for a more detailed presentation of the simulation studies with no change points. The following sections describe in details results for Scenarios 1 and 2. Scenario 3 is in section B of the Supplementary Material (Franzolini et al. (2024)).

4.1. Simulation results for one change point. The first simulation scenario (Scenario 1) is obtained by setting one change point at $t = 70$ and generating two precision matrices, as in Peterson, Stingo and Vannucci (2015) and Molinari et al. (2022b). In particular, we first define the Ω_{c_0} and then derive Ω_{c_1} as perturbation of Ω_{c_0} , which defines a mildly-abrupt change. First, Ω_{c_0} is obtained by setting diagonal elements equal to 1, the first off-diagonal elements to 0.5, that is, $\Omega_{c_0}[h, h + 1] = \Omega_{c_0}[h + 1, h] = 0.5$, for $h = 1, \dots, 9$, and the remaining elements to 0. To construct Ω_{c_1} , we randomly remove five edges among the active ones in G_{c_0} and set to 0 the corresponding entries in the precision matrix. Then we add five randomly selected edges drawn from the set of inactive edges in G_{c_0} . Finally, a precision entry equal to 0.2 is assigned to the new edges. The obtained matrix is not necessarily positive-definite, and to this end, we compute the nearest positive-definite approximation through the R function nearPD (Higham (2002)), available in the R package Matrix (Bates et al. (2022)). The resulting graphs are shown in Figure B.1.2 of the Supplementary Material (Franzolini et al. (2024)). We note that the computation of the nearest positive-definite matrix may result in a strong shrinkage of the nonzero elements in the precision matrix, which may cause unrealistic high values in the correspondent covariance matrix. However, this is not the case in our simulation scenario (see Figure B.1.4 and B.1.5 in the Supplementary Material, where the simulated covariance matrices and data are displayed, Franzolini et al. (2024)). We consider 20 replicates of Scenario 1.

The hyperparameter ω in (4) is determined using an approach inspired by empirical Bayes techniques so that a priori the expected number of edges for the graphs is equal to the number of edges detected by estimating one unique graph using all the time points. To this end, we estimate the graph using an adaptive lasso approach, which is a modification of the estimation procedure proposed by Meinshausen and Bühlmann (2006) inspired by the adaptive lasso of (Zou (2006)), as implemented in the R package GGMselect (Bouvier et al. (2022)). For the hyperparameter z in (5), we opt for $z = 0.1$ so that a priori we expect only one edge to change at each change point, favouring graph similarity. This choice also allows us to better understand model performance and hyperparameter sensitivity, as in our simulations we force 10 edges to change across the change point, an event to which our prior associates a probability lower than $4 \cdot 10^{-8}$. The hyperparameter p_0 , which controls the a priori number of the

TABLE 2

Scenario 1: Posterior summaries for change points of the three best and three worst replicates in terms of MAP estimate and MAP probability. MAP estimates, MAP probabilities (for the posterior over all configurations of change points), posterior probability of the number of change points being 1, and credible intervals, mean, and median of the position of the change point (conditionally on having one change point). Credible intervals are obtained computing the smallest credible sets with 90% and 95% credibility, which are not necessary continuous intervals, and then using the minimum and the maximum time points in the credible set as extremes of the provided interval

True change point		MAP est.	MAP prob.	prob $\kappa = 1$	90% C.I.	95% C.I.	Mean	Median	
Scenario 1	(70)	Rep. 10	(70)	0.306	0.971	[70, 74]	[70, 79]	72.21	71
		Rep. 5	(70)	0.294	0.987	[70, 80]	[70, 80]	73.34	72
		Rep. 3	(70)	0.289	0.975	[69, 79]	[68, 79]	71.17	70
		Rep. 11	(68)	0.245	0.966	[68, 76]	[68, 81]	70.60	70
		Rep. 20	(73)	0.179	0.965	[67, 75]	[66, 77]	71.19	71
		Rep. 12	(74)	0.191	0.933	[67, 83]	[67, 83]	73.53	73

change points, is set to $p_0 = 0.1$ to favour sparsity. The hyperparameters of the G-Wishart distribution are set to the common values of $\delta = 3$ and $D = \text{Id}_p$. In Section C.5 of the Supplementary Material (Franzolini et al. (2024)), we provide hyperparameter sensitivity analysis carried on the real dataset. The change point detection procedure appears to be insensitive to the choice of the hyperparameters, and the graph recovery performance is limitedly affected by the choice of z .

Table 2 contains posterior summaries of the three best and three worst replicates based on the accuracy in the recovery of the change point configuration. Summaries for all replicates can be found in Section B.3 of the Supplementary Material (Franzolini et al. (2024)). In all replicates the posterior distribution of the number of change points is concentrated around 1, with a posterior probability greater or equal to 0.91. Moreover, though the space of sequences that satisfy the minimum duration constraint of $\ell = 12$ for $T = 200$ includes more than $4 \cdot 10^{12}$ sequences, in eight of the 20 replicates the MAP estimate (which minimises the 0 – 1 loss function) coincides with the true state $c^* = 70$, and in all replicates the MAP is contained in the interval [68, 74]. Moreover, Table 2 reports also the 90% and 95% credible intervals, the mean, and the median for the position of the change point, confirming that the posterior is concentrated around the true state in all replicates. Figure 5a shows the combined false positive rate (FPR) of edge detection for the two graphs as a function of the threshold used

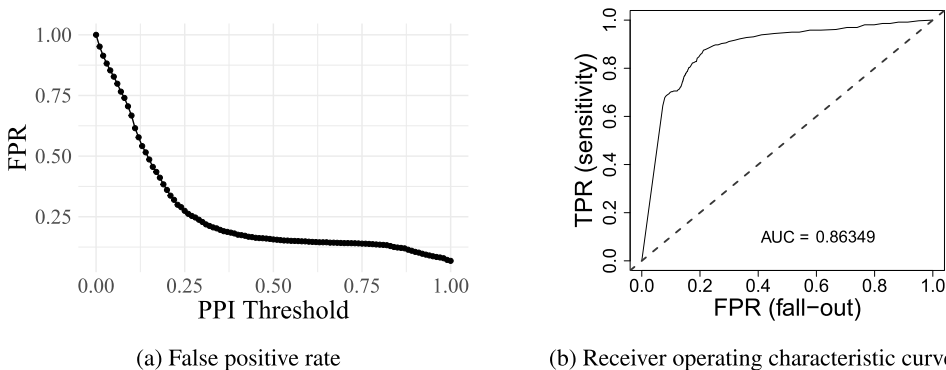


FIG. 5. Panel (a): FPR vs. PPI threshold for Scenario 1, computed via 20 replicates. Panel (b): ROC curve in Scenario 1, computed via 20 replicates.

for the posterior probability of inclusion (PPI). FPRs show a reasonable pattern, and for the 0.5 threshold, the FPR is 0.156. Figure 5b displays the combined ROC curve, with an AUC approximately equal to 0.86. The graph estimates used to evaluate the FPRs are obtained conditionally on the estimated change point sequence, even when it does not coincide with the true one. Indeed, the MAP always identifies a single change point, so we compare the estimated graphs before/after the estimated change point, with the true graph before/after the true change point for fairness of results.

Lastly, we compare the results with those obtained by applying the group-fused graphical lasso (GFGL), introduced by [Gibberd and Nelson \(2017\)](#) (see also [Gibberd and Roy \(2017\)](#)), and the LOcal Group Graphical Lasso Estimation (loggle) of [Yang et al. \(2015\)](#). Similarly to our proposal, GFGL consists of a piecewise constant graphical model. Differently from our model, it is estimated employing a penalized likelihood approach, where two penalties act favouring both sparsity of the graph structure (i.e., shrinkage penalty) and sparsity in the number of change points (i.e., smoothing penalty). The loggle approach is also based on a penalised likelihood approach, but it assumes that the graph topology is gradually changing over time and thus cannot be used to detect change points. We compare our approach with an “oracle version” of loggle, where we estimate the graphs knowing the position of the change points.

From a theoretical point of view, both GFGL and loggle prohibit principled uncertainty quantification on the number and location of change points and graph structure. Contrarily, our strategy allows for straightforward uncertainty quantification, which is one of the main advantages of the Bayesian framework. However, as it is well-known, Bayesian posterior inference typically comes at the expense of the computational time needed to estimate the model. So in this scenario, GFGL and loggle produce estimates in a few seconds or minutes, and our dynamic Gaussian graphical model requires hours to be estimated (see D.2 in the Supplementary Material for details on computational time, [Franzolini et al. \(2024\)](#)). Results from the GFGL model are obtained for different values of the hyperparameters λ_1 and λ_2 . The hyperparameters λ_1 and λ_2 control, respectively, the shrinkage and the smoothness of the solution; for more details, see [Gibberd and Nelson \(2017\)](#). Detailed output summaries for the GFGL are presented in Section B.3 of the Supplementary Material ([Franzolini et al. \(2024\)](#)). As already noticed by [Gibberd and Nelson \(2017\)](#), GFGL’s results can be highly sensitive to the choice of the hyperparameters in terms of both detected change points and recovered graph structure. In our experiment the number of change points estimated by GFGL varies from one to seven, depending on the simulation replicate and on the choice of hyperparameters λ_1 and λ_2 . The location of the change points also varies largely across the different simulation replicates. Contrarily, our model identifies the correct number of change points and their approximate position in all replicates. To evaluate the graph recovery performance of GFGL, we compute the FPR and the TPR at the two middle points $t = 35$ and $t = 135$. For GFGL with $\lambda_1 = 0.25$ and $\lambda_2 = 20$, the FPR is 0.412 and the TPR is 0.550. For GFGL with $\lambda_1 = 0.50$ and $\lambda_2 = 60$, the FPR is 0.175, and the TPR is 0.311. The “oracle version” of the loggle model gives a FPR of 0.242 and a TPR of 0.825. Contrarily, with our approach, if we fix the FPR to 0.242 the corresponding TPR is 0.895, fixing the TPR to 0.825 leads to a FPR of 0.181. (see Table B.3.2 in the Supplementary Material, [Franzolini et al. \(2024\)](#)).

4.2. Simulation results for more changes points and nodes. The second simulation scenario (Scenario 2) is obtained simulating data for $T = 200$ time points and $p = 20$ variables/nodes. The data generating mechanism presents three change points located at $t = 60$, $t = 100$, and $t = 150$ and thus four different graph structures. The true graph structures are displayed in Figure B.1.3 of the Supplementary Material ([Franzolini et al. \(2024\)](#)) and obtained fixing the first graph G_1 , which presents 11 activated edges and, subsequently, randomly changing the graph structure across change points. In each corresponding change




FIG. 6. Panel (a): FPR vs. PPI threshold for Scenario 2. Panel (b): ROC curve in Scenario 2.

point, any active/nonactive edge is deactivated/activated with probability 0.4. The four precision matrices are then generated sampling from a G -Wishart distribution, independently conditional on graph structure. Note that the abruptness in the changes in the dependence structure is higher in this scenario than in Scenario 1, since the precision matrix is generated independently at each change point and not obtained as a perturbation of the previous precision matrix, as in the previous section. The hyperparameters are chosen as described in the previous section.

The posterior distribution for the number of change points assigns probability one to the correct value of 3 (see Figure B.4.1 in the Supplementary Material, Franzolini et al. (2024)). The posterior expectation for the first change point location is 60.68, for the second change point is 98.11, and for the third one is 149.94. More details on posterior inference results are provided in Section B.4 of the Supplementary Material (Franzolini et al. (2024)). Figure 6a shows the combined FPR of edge detection for the four graphs as function of the threshold used for the PPI. Again, the FPRs appear reasonable for any PPI threshold, quickly decaying to zero, and for the 0.5 threshold, the FPR is 0.035. Figure 6b displays the ROC curve, with an AUC approximately equal to 0.89. In the Supplementary Material (Franzolini et al. (2024)), we also report results obtained with the GFGL model of Gibberd and Nelson (2017), which, similarly to what already observed in the previous section, shows high variability of the estimates depending on the value of the hyperparameters. Moreover, in this scenario we note that the GFGL model leads also to poor graph recovery, even when the correct change points are detected (see Figures B.4.3 and B.4.4 of the Supplementary Material, Franzolini et al. (2024)). In particular, for $\lambda_1 = 0.2$ and $\lambda_2 = 60$, GFGL detects the correct location of the change points but has a FPR of 0.176 and TPR of 0.372. The “oracle version” of the loggle model produces a FPR of 0.113 and a TPR of 0.426, presenting significantly worse performance than our approach. For example, for a 0.5 PPI threshold we obtain a FPR of 0.035 and a TPR of 0.611, if we fix the FPR to 0.113 the corresponding TPR is 0.764, fixing the TPR to 0.426 leads to a FPR of 0.021; see Table 3.

5. Industry returns during COVID-19 pandemic. We apply the model to detect changes in the dependence structure of the nine industry portfolios’ weekly returns described in Section 1. We consider weekly data over a time horizon of three years: from January 2019 to December 2021 so that $T = 157$. When choosing which type of returns to include in the analysis, that is, daily, weekly, or monthly, we are faced with a trade-off: higher frequency data may show lower degree of dependence, making harder to detect structure changes in the dependence structure (see, for instance, Ab Razak, Aminuddin and Ismail (2018)); on the other hand, lower frequency data provide a less detailed

TABLE 3

Scenario 2: Comparison on graph recovery between the results obtained with the “oracle version” of the loggle model and our model

	loggle	Bayesian dynamic GGM		
		PPI thres. 0.5	PPI thres. 0.25	PPI thres. 0.78
FPR	0.113	0.035	0.113	0.021
TPR	0.426	0.611	0.764	0.426

representation of markets’ trends. For this reason we consider weekly returns to attain a more detailed level of information, compared to monthly data, and a potentially stronger signal on correlations compared to daily returns. Logarithmic weekly returns are computed starting from the daily returns available from Kenneth R. French’s Data Library at https://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html, where the industry classification used to associate each stock to one of the nine portfolios is defined as follows. Stocks listed in the New York Stock Exchange (NYSE), the American Stock Exchange (AMEX), and National Association of Securities Dealers Automatic Quotation System (NASDAQ) are assigned to an industry at the end of June of year t based on their four-digit standard industrial classification (SIC) code at that time. Then returns are computed from July of year t to June of year $t + 1$. The corresponding standardized time series are represented in Figure C.1 of the Supplementary Material (Franzolini et al. (2024)).

To estimate the dependence structure for the weekly returns, we first run the algorithm to estimate the change point sequence, using 200 particles, 10 mutation steps, performing 32,000 iterations of the outer component, of which the first 2000 are discarded as burn-in, and thinning every 10 iterations. Second, we rerun only the particle filter with 1000 particles and 20 mutation steps to sample the graphs from their posterior distribution conditionally on the MAP estimate of the change point configuration. The algorithm is initialised assuming no change points, and hyperparameters are set as described in Section 4.1.

Figure 7 shows, for each time point, the marginal posterior probability of being a change point, and Table 4 reports the joint posterior distribution of the configurations of change points, which have been accepted by the algorithm. The posterior distribution on the number of change points assigns probability 0.9984 to two change points and the remaining mass 0.0016 to three change points. The posterior distribution for the number and location of




FIG. 7. Marginal posterior probability of every time point to be a change point.

TABLE 4

Posterior distribution of change point configuration. In **bold** we highlight probabilities greater than 0.05 and in underline the MAP estimate

Change points	Post prob.	Change points	Post prob.	Change points	Post prob.	Change points	Post prob.
(57 79)	0.0010	(59 82)	0.0073	(60 99)	0.0003	(61 81)	0.0137
(57 82)	0.0010	(59 84)	0.0017	(61 116)	0.0010	(61 82)	0.1094
(57 84)	0.0007	(60 78)	0.0020	(61 73)	0.0007	(61 83)	0.0650
(57 85)	0.0003	(60 79)	0.0073	(61 74)	0.0033	(61 84)	0.0150
(58 79)	0.0003	(60 80)	0.1010	(61 75)	0.0103	(61 85)	0.0127
(58 82)	0.0033	(60 81)	0.0033	(61 76)	0.0133	(61 86)	0.0043
(59 77)	0.0013	(60 82)	0.0297	(61 77)	0.0447	(61 99)	0.0010
(59 78)	0.0007	(60 83)	0.0100	(61 78)	0.0737	(19 61 83)	0.0010
(59 79)	0.0010	(60 84)	0.0127	<u>(61 79)</u>	<u>0.3735</u>	(39 61 83)	0.0003
(59 81)	0.0003	(60 85)	0.0020	(61 80)	0.0694	(61 80 113)	0.0003

change points are highly concentrated around the posterior mode, showing a low level of uncertainty. Our analysis highlights a first structural change at $t = 61$, that is, during the week starting on February 24, 2020, in correspondence of what appears to be the market's reaction to the first significant world-wide increase in Coronavirus confirmed cases and deaths outside China over the previous weekend. In particular, during the weekend February 21–23, 2020, Italy, the first and hardest-hit country in Europe in 2020, reported the first local cases of COVID-19 (see, e.g., [Just and Echaust \(2020\)](#)). During the week of February 24, 2020, the Dow Jones and S&P 500 fell by 11% and 12%, respectively, marking the biggest weekly declines to occur since the financial crisis of 2008. The identification of a first change point in correspondence of a major shock is coherent with the known stylized fact that, during crisis, dependence among investments typically increases diminishing diversification benefits (see, for instance, [Kotkatvuori-Örnberg, Nikkinen and Äijö \(2013\)](#)). A second change point is detected at $t = 79$, that is, during the last week of June 2020, interpretable as a subsequent and partial restabilization of the financial markets after the initial and most uncertain period of the pandemic. The credible intervals, obtained computing the smallest credible sets with 95% credibility and then using the minimum and the maximum time points in the set as extremes of the interval, are [60, 61] and [76, 83] for the first and second change points, respectively.

Conditionally on the MAP change point configuration, estimates of the three graphs are provided in Figure 8, while the estimated variance and covariance matrices are displayed in Figure C.2.1 of the Supplementary Material ([Franzolini et al. \(2024\)](#)). The graphs are obtained based on the marginal PPI of the edges in order to control the corresponding Bayesian false discovery rate ([Newton et al. \(2004\)](#)). In particular, we set the threshold of inclusion, based on the PPI, to 0.8 in order to guarantee an expected rate of false detection not higher than 0.05, that is, a specificity of at least 0.95 (for more details, see, [Leday and Richardson \(2019\)](#), [Williams \(2021\)](#)). In Section C.1 of the Supplementary Material ([Franzolini et al. \(2024\)](#)), we report the values of degree centrality, betweenness centrality ([Freeman \(1977\)](#)), local clustering and global clustering coefficients ([Watts and Strogatz \(1998\)](#)) for the estimated graphs, which give insights into the role of each node.

A clearly noticeable feature from Figure 8 is the increase in the number of edges from the first change point (eight edges) to the followings (11 and 12 edges, respectively), reflected also in the global measures of clustering of the graphs, which varies from 0 in the pre COVID-19 period, to 0.43 during the first COVID-19 outbreak, to 0.24 to the post COVID-19 outbreak (see Table C.1.3 in the Supplementary Material, [Franzolini et al. \(2024\)](#)). Such increase in the connectivity of the graph is coherent with the hypothesis of the COVID-19 outbreak acting as

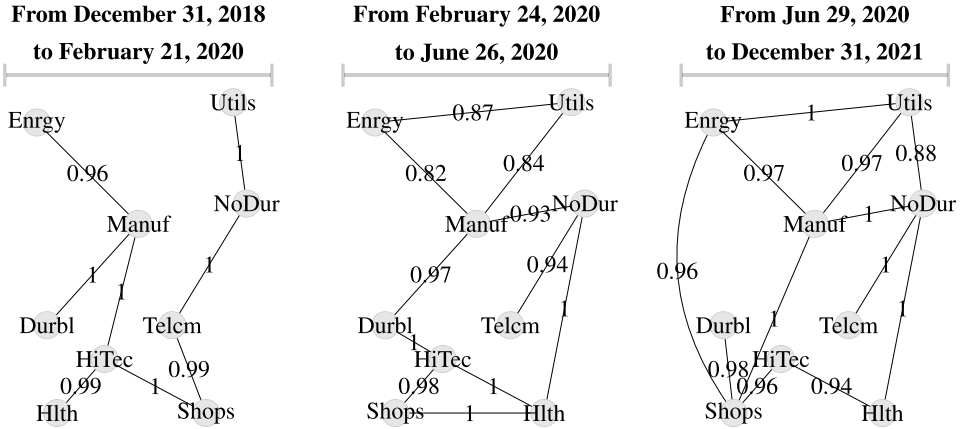


FIG. 8. Posterior estimates of the graphs and PPI for the selected edges. Threshold of inclusion is set to achieve an expected posterior specificity of at least 95%.

a common nonmeasurable risk factor driving the returns in the market. As already noticed in the preliminary analysis summarised by Figure 2, the role of hub of the manufacturing portfolio (i.e., machinery, trucks, planes, chemicals, office furniture, paper) over the three years is confirmed. The corresponding degree centrality (i.e., number of vertices in the neighbourhood of the manufacturing portfolio) is the highest in all three graphs (see Table C.1.1 of the Supplementary Material for more details, Franzolini et al. (2024)). However, contrarily to the conclusions of the initial exploratory analysis, consumer nondurable (i.e., food, tobacco, textiles, apparel, leather, toys) returns appear to play a less central role before the COVID-19 outbreak. Such evidence that was absent in the explanatory analysis is discovered, mainly thanks to the automatic detection of the change points, that allow us to determine the most appropriate time window to estimate the graph and capture differences in structure. Moreover, we identify another hub in the Shops portfolio (i.e., wholesale, retail, and some services, as laundries and repair shops) in the time-interval after the last change point.

In terms of volatility, in all three periods the consumer durables (i.e., cars, TVs, furniture, household appliances) and energy (i.e, oil, gas, and coal extraction and products) are characterised by the highest volatility. Moreover, the analysis confirms that the three periods (pre COVID-19 outbreak, during first COVID-19 global outbreak, and after) coincide with small, high, and medium volatility markets, as already evident from the time-series plot (see Figure C.1 in the Supplementary Material, Franzolini et al. (2024)). In more detail we note that health (i.e., healthcare, medical equipment, and drugs) is the only industry in the market whose portfolio presents a similar volatility before and after the outbreak, while all other portfolios' returns are set to higher levels of variability as consequence of a long-run effect of market uncertainty. Similar conclusions can be drawn also for the pairwise correlations (that can be easily computed from the values in Figure C.2.1 and are reported in Section C.2 of the Supplementary Material, Franzolini et al. (2024)); that is, correlations are higher during COVID-19 outbreak.

Finally, we compare our results with those obtained by applying the group-fused graphical lasso (GFGL) (Gibberd and Nelson (2017)). Output summaries for the GFGL are shown in Table 5, and additional figures can be found in Section C.3 of the Supplementary Material (Franzolini et al. (2024)). We estimate the GFGL model for different values of the hyperparameters λ_1 and λ_2 . We recall that λ_1 and λ_2 control the shrinkage and the smoothness of the solution of the GFGL model, respectively. As already noticed by Gibberd and Nelson (2017) and in the simulation study in Section 4, for GFGL inference results are highly sensitive to the hyperparameters in terms of both detected change points (see Table 5 and Figure C.3.1 in the

TABLE 5
Results obtained using GFGL for different values of the hyperparameters

λ_1	λ_2	# of change points	Change points	Global clustering coef.
0.25	60	0	()	(0.72)
0.35	60	2	(61, 80)	(0, 1, 0.33)
0.50	60	2	(61, 80)	(0, 0, 0)
0.25	55	2	(61, 80)	(0.72, 0.84, 0.81)
0.25	20	3	(61, 80, 98)	(0, 1, 0.92, 0.92)
0.35	20	4	(61, 80, 98, 99)	(0, 1, 0.33, 0.33, 0)
0.50	20	4	(61, 68, 77, 80)	(0, 1, 0.88, 0.87, 0)
0.25	10	6	(57, 61, 80, 98, 99, 116)	(0, 0, 1, 0.89, 0.90, 0.33, 0.33)
0.35	10	4	(61, 80, 98, 116)	(0, 1, 0.33, 0.33, 0)
0.50	10	6	(57, 61, 68, 69, 77, 80)	(0, 0, 1, 0.88, 0.88, 0)

Supplementary Material, Franzolini et al. (2024)) and recovered graph structure (see Figure C.3.2 in the Supplementary Material, Franzolini et al. (2024)). Here we consider values for λ_1 in $\{0.25, 0.35, 0.50\}$ and λ_2 in $\{10, 20, 55, 60\}$, which are in the range of those considered by Gibberd and Nelson (2017) in their work. The number of identified change points ranges widely from zero to six. However, we notice that in all estimated change point configurations, but the one with no estimated change points, GFGL always includes $t = 61$ and at $t = 80$ as change points, which is consistent with the change points identified by our approach. Moreover, even though the graph structure estimated by the GFGL largely varies depending on the hyperparameters, in all settings where change points are detected, the graph structure connectivity increases during the COVID-19 outbreak in February 2020 and diminishes after June 2020 (cf. the global clustering coefficients reported in Table 5 and Figure C.3.2 in the Supplementary Material, Franzolini et al. (2024)). This result is again consistent with the results obtained with our approach. Increased graph connectivity during the COVID-19 outbreak is also found applying loggle; see Section C.3 of the Supplementary Material (Franzolini et al. (2024)).

6. Discussion and conclusion. In this work we study the impact of the COVID-19 pandemic on the U.S. stock market, with a specific focus on changes in dependence structure across stocks related to different industries. To do so, we consider weekly returns recorded for three years starting in January 2019. We identify two structural changes. The first change is in correspondence with the last week of February 2020, a date that for most countries coincided with the beginning of the pandemic. That same week financial markets recorded the weekly biggest losses since the financial crisis of 2008. The second change point is detected after approximately four months, when there is a reduction in market uncertainty, but the dependence structure as well as the volatility are not back to pre COVID-19 levels. Comparing the dependence structure across the three periods (i.e., before February 24th, between February 24th and June 26th, after June 26th) we provide many insights on the impact of the pandemic on the stock market and highlight whose effects appear to be persistent up to the end of 2021, the last year considered.

The main methodological contribution of this work is the development of a dynamic GGM, which allows for abrupt changes in the dependence structure of the random variables represented by the nodes of the graph. Our model builds on existing literature on GGMs as well as random change points. Our model construction allows us to control sparsity in the number of change points and/or in the graph structure. We have designed a tailored SMC algorithm,

arguing for its use in such a complex setup over other M-H based alternatives and demonstrating its performance on simulated data and on our motivating application.

Our work opens up several avenues for future research:

(a) Scalability. We have not made use of the full SMC machinery. We briefly discuss two directions for increase in model dimension, along the number of nodes, p , and along the length of time instances for observations T . In terms of the size of the graph, recent works have developed effective proposals, informed by the observations, for MCMC methods on graph posteriors, instead of previously used random-walk-type blind moves; see, for example, [van den Boom, Beskos and De Iorio \(2022\)](#) and the references therein, for approaches based on Langevin-type analogues for discrete spaces, with parallelisation employed within the specification of the proposal. Such approaches have been seen to be effective for node sizes of $p = \mathcal{O}(10^2)$. In terms of the length T , recent advances on modelling involving change points and accompanying SMC methodologies can permit for recasting models so that change points also become part of the hidden Markov process (see, e.g., [Yildirim, Singh and Doucet \(2013\)](#)). At the same time, SMC methods, based on state-of-art particle Gibbs approaches that incorporate backward steps to improve mixing over the update of the unobserved Markovian states, are shown to provide pseudomarginal methodologies of superior mixing compared to standard PMCMC ([Lindsten, Jordan and Schön \(2014\)](#)). Such new algorithms are supported by strong theoretical results. Indicatively, the number of particles can now be allowed to remain constant, $N = \mathcal{O}(1)$, as a function of T , when PMCMC requires $N = \mathcal{O}(T)$. Thus, costs for the overall SMC algorithm can be brought down to $\mathcal{O}(NT)$ from the previous $\mathcal{O}(T^2)$, for big T —with $\mathcal{O}(NT)$ not taking under consideration the option of parallelisation across particles.

(b) Smooth changes. In this work we have considered abrupt changes in edge inclusion probabilities. Alternatively, we could model edge inclusion probabilities as a function of time, for example, using autoregressive-type models. In this setup shrinkage priors could be specified to link the probability of edge inclusion at time t to the same probability at time $t - 1$ (see, for instance, [Molinari et al. \(2022a\)](#)). This approach is amenable to many generalizations, such as the inclusion of covariates. Moreover, the probability of edge inclusion at time t could be a function of the probabilities of edge inclusions at time $t - 1$ of a neighbourhood of each node.

(c) Graph substructures. Here we have presented changes between graphs, as captured by edge flips before and after a change point. Edge detection is very sensitive to the number of nodes as well as sample size. It has been argued ([van den Boom, De Iorio and Beskos \(2022\)](#), and references therein) that in many applications a more robust approach is to shift the focus of inference to graph sub-structures, such as hubs and communities, with the goal of capturing the evolution over time of such macro-structures, which better describe the underlying phenomenon.

(d) More general response types. The model can be easily extended to accommodate different type of responses, such as binary and count data. An easy solution would be the representation of such data in terms of latent variables ([Albert and Chib \(1993\)](#), [Chib and Greenberg \(1998\)](#)). Moreover, it is straightforward to include time-homogeneous and time-varying covariates to model the mean of the time series as well as a trend and seasonal component.

Acknowledgments. The authors are grateful to the Editor, an Associate Editor, and five anonymous referees for insightful comments and suggestions, which led to a substantial improvement of the manuscript. Most of the paper was completed while B. Franzolini was a Research Fellow at the Singapore Institute for Clinical Sciences of the Agency for Science, Technology and Research, in Singapore, Republic of Singapore. M. De Iorio is also affiliated with the Department of Statistical Sciences of the University College of London and the Singapore Institute for Clinical Sciences of the Agency for Science, Technology and Research.

- CARVALHO, C. M. and WEST, M. (2007). Dynamic matrix-variate graphical models. *Bayesian Anal.* **2** 69–97. MR2289924 <https://doi.org/10.1214/07-BA204>
- CHAN, J. C. C. (2013). Moving average stochastic volatility models with application to inflation forecast. *J. Econometrics* **176** 162–172. MR3084051 <https://doi.org/10.1016/j.jeconom.2013.05.003>
- CHANDRA, N. K., MUELLER, P. and SARKAR, A. (2021). Bayesian scalable precision factor analysis for massive sparse Gaussian graphical models. arXiv preprint, arXiv:2107.11316.
- CHENG, Y. and LENKOSKI, A. (2012). Hierarchical Gaussian graphical models: Beyond reversible jump. *Electron. J. Stat.* **6** 2309–2331. MR3020264 <https://doi.org/10.1214/12-EJS746>
- CHIB, S. and GREENBERG, E. (1998). Analysis of multivariate probit models. *Biometrika* **85** 347–361.
- CHOPIN, N. and PAPASPILIOPOULOS, O. (2020). *An Introduction to Sequential Monte Carlo*. Springer Series in Statistics. Springer, Cham. MR4215639 <https://doi.org/10.1007/978-3-030-47845-2>
- CLARK, T. E. and RAVAZZOLO, F. (2015). Macroeconomic forecasting performance under alternative specifications of time-varying volatility. *J. Appl. Econometrics* **30** 551–575. MR3358636 <https://doi.org/10.1002/jae.2379>
- CONTESSI, S., DE PACE, P. and GUIDOLIN, M. (2014). How did the financial crisis alter the correlations of US yield spreads? *J. Empir. Finance* **28** 362–385.
- CRIBBEN, I. (2019). Change points in heavy-tailed multivariate time series: Methods using precision matrices. *Appl. Stoch. Models Bus. Ind.* **35** 299–320. MR3946128 <https://doi.org/10.1002/asmb.2373>
- CRIBBEN, I. and YU, Y. (2017). Estimating whole-brain dynamics by using spectral clustering. *J. R. Stat. Soc. Ser. C. Appl. Stat.* **66** 607–627. MR3632344 <https://doi.org/10.1111/rssc.12169>
- CUI, K. (2012). Semiparametric Gaussian variance-mean mixtures for heavy-tailed and skewed data. *Int. Sch. Res. Not.* **2012**.
- DANAHER, P., WANG, P. and WITTEN, D. M. (2014). The joint graphical lasso for inverse covariance estimation across multiple classes. *J. R. Stat. Soc. Ser. B. Stat. Methodol.* **76** 373–397. MR3164871 <https://doi.org/10.1111/rssb.12033>
- DAWID, A. P. and LAURITZEN, S. L. (1993). Hyper-Markov laws in the statistical analysis of decomposable graphical models. *Ann. Statist.* **21** 1272–1317. MR1241267 <https://doi.org/10.1214/aos/1176349260>
- DEL MORAL, P. (2004). *Feynman–Kac Formulae: Genealogical and Interacting Particle Systems with Applications. Probability and Its Applications (New York)*. Springer, New York. MR2044973 <https://doi.org/10.1007/978-1-4684-9393-1>
- DELLAPORTAS, P., GIUDICI, P. and ROBERTS, G. (2003). Bayesian inference for nondecomposable graphical Gaussian models. *Sankhyā* **65** 43–55. MR2016776
- DERBALI, A., NAOU, K., SASSI, M. B. and AMIRI, M. M. (2022). Do COVID-19 epidemic explains the dynamic conditional correlation between China’s stock market index and international stock market indices? *Chin. Econ.* **55** 227–242.
- DEY, A. K., HOQUE, G. T., DAS, K. P. and PANOVSKA, I. (2022). Impacts of COVID-19 local spread and Google search trend on the us stock market. *Phys. A, Stat. Mech. Appl.* **589** 126423.
- DIACONIS, P. and YLVIKAKER, D. (1979). Conjugate priors for exponential families. *Ann. Statist.* **7** 269–281. MR0520238
- DOBRA, A., LENKOSKI, A. and RODRIGUEZ, A. (2011). Bayesian inference for general Gaussian graphical models with application to multivariate lattice data. *J. Amer. Statist. Assoc.* **106** 1418–1433. MR2896846 <https://doi.org/10.1198/jasa.2011.tm10465>
- DOUCET, A., PITT, M. K., DELIGIANNIDIS, G. and KOHN, R. (2015). Efficient implementation of Markov chain Monte Carlo when using an unbiased likelihood estimator. *Biometrika* **102** 295–313. MR3371005 <https://doi.org/10.1093/biomet/asu075>
- DUTTA, A., BOURI, E. and NOOR, M. H. (2021). Climate bond, stock, gold, and oil markets: Dynamic correlations and hedging analyses during the COVID-19 outbreak. *Resour. Policy* **74** 102265.
- ENGLE, R. F. (1982). Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation. *Econometrica* **50** 987–1007. MR0666121 <https://doi.org/10.2307/1912773>
- ENGLE, R. F. and BOLLERSLEV, T. (1986). Modelling the persistence of conditional variances. *Econometric Rev.* **5** 1–87. With comments and a reply by the authors. MR0876792 <https://doi.org/10.1080/07474938608800095>
- FRANZOLINI, B., BESKOS, A., DE IORIO, M., POKLEWSKI KOZIELL, W. and GRZESZKIEWICZ, K. (2024). Supplement to “Change point detection in dynamic Gaussian graphical models: The impact of COVID-19 pandemic on the U.S. stock market.” <https://doi.org/10.1214/23-AOAS1801SUPPA>, <https://doi.org/10.1214/23-AOAS1801SUPPB>
- FREEMAN, L. C. (1977). A set of measures of centrality based on betweenness. *Sociometry* 35–41.
- GIBBERD, A. J. and NELSON, J. D. B. (2017). Regularized estimation of piecewise constant Gaussian graphical models: The group-fused graphical Lasso. *J. Comput. Graph. Statist.* **26** 623–634. MR3698672 <https://doi.org/10.1080/10618600.2017.1302340>

- GIBBERD, A. J. and ROY, S. (2017). Multiple changepoint estimation in high-dimensional gaussian graphical models. arXiv preprint, [arXiv:1712.05786](https://arxiv.org/abs/1712.05786).
- GREEN, P. J. (1995). Reversible jump Markov chain Monte Carlo computation and Bayesian model determination. *Biometrika* **82** 711–732. [MR1380810 https://doi.org/10.1093/biomet/82.4.711](https://doi.org/10.1093/biomet/82.4.711)
- GUIDOLIN, M., LA CARA, D. and MARCELLINO, M. G. (2021). Boosting the forecasting power of conditional heteroskedasticity models to account for COVID-19 outbreaks. BAFFI CAREFIN Centre Research Paper (2021-169).
- HAAS, M., MITTNIK, S. and PAOLELLA, M. S. (2004). A new approach to Markov-switching GARCH models. *J. Financ. Econom.* **2** 493–530.
- HALLAC, D., PARK, Y., BOYD, S. and LESKOVEC, J. (2017). Network inference via the time-varying graphical lasso. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* 205–213.
- HIGHAM, N. J. (2002). Computing the nearest correlation matrix—a problem from finance. *IMA J. Numer. Anal.* **22** 329–343. [MR1918653 https://doi.org/10.1093/imanum/22.3.329](https://doi.org/10.1093/imanum/22.3.329)
- HLÁVKA, Z., HUŠKOVÁ, M. and MEINTANIS, S. G. (2020). Change-point methods for multivariate time-series: Paired vectorial observations. *Statist. Papers* **61** 1351–1383. [MR4127478 https://doi.org/10.1007/s00362-020-01175-3](https://doi.org/10.1007/s00362-020-01175-3)
- HULL, J. and WHITE, A. (1987). The pricing of options on assets with stochastic volatilities. *J. Finance* **42** 281–300.
- JASRA, A., STEPHENS, D. A., DOUCET, A. and TSAGARIS, T. (2011). Inference for Lévy-driven stochastic volatility models via adaptive sequential Monte Carlo. *Scand. J. Stat.* **38** 1–22. [MR2760137 https://doi.org/10.1111/j.1467-9469.2010.00723.x](https://doi.org/10.1111/j.1467-9469.2010.00723.x)
- JUST, M. and ECHAUST, K. (2020). Stock market returns, volatility, correlation and liquidity during the COVID-19 crisis: Evidence from the Markov switching approach. *Finance Res. Lett.* **37** 101775.
- KARAGIANNIS, G. and ANDRIEU, C. (2013). Annealed importance sampling reversible jump MCMC algorithms. *J. Comput. Graph. Statist.* **22** 623–648. [MR3173734 https://doi.org/10.1080/10618600.2013.805651](https://doi.org/10.1080/10618600.2013.805651)
- KESHAVARZ, H., MICHAILIDIS, G. and ATCHADÉ, Y. (2020). Sequential change-point detection in high-dimensional Gaussian graphical models. *J. Mach. Learn. Res.* **21** Paper No. 82, 57. [MR4119150](https://doi.org/10.48550/jmlr.2020.21.19150)
- KOLAR, M. and XING, E. P. (2012). Estimating networks with jumps. *Electron. J. Stat.* **6** 2069–2106. [MR3020257 https://doi.org/10.1214/12-EJS739](https://doi.org/10.1214/12-EJS739)
- KOTKATVUORI-ÖRNBERG, J., NIKKINEN, J. and ÄIJÖ, J. (2013). Stock market correlations during the financial crisis of 2008–2009: Evidence from 50 equity markets. *Int. Rev. Financ. Anal.* **28** 70–78.
- LAURITZEN, S. L. (1996). *Graphical Models. Oxford Statistical Science Series 17*. The Clarendon Press, Oxford Univ. Press, New York. Oxford Science Publications. [MR1419991](https://doi.org/10.1017/CBO9780511526908)
- LEDAY, G. G. R. and RICHARDSON, S. (2019). Fast Bayesian inference in large Gaussian graphical models. *Biometrics* **75** 1288–1298. [MR4041830 https://doi.org/10.1111/biom.13064](https://doi.org/10.1111/biom.13064)
- LEE, H.-T. (2010). Regime switching correlation hedging. *J. Bank. Financ.* **34** 2728–2741.
- LI, Y., CRAIG, B. A. and BHADRA, A. (2019). The graphical horseshoe estimator for inverse covariance matrices. *J. Comput. Graph. Statist.* **28** 747–757. [MR4007755 https://doi.org/10.1080/10618600.2019.1575744](https://doi.org/10.1080/10618600.2019.1575744)
- LIEN, D., TSE, Y. K. and TSUI, A. K. (2002). Evaluating the hedging performance of the constant-correlation GARCH model. *Appl. Financ. Econ.* **12** 791–798.
- LINDSTEN, F., JORDAN, M. I. and SCHÖN, T. B. (2014). Particle Gibbs with ancestor sampling. *J. Mach. Learn. Res.* **15** 2145–2184. [MR3231604](https://doi.org/10.48550/jmlr.2014.15.1604)
- LIU, B., ZHANG, X. and LIU, Y. (2021). Simultaneous change point inference and structure recovery for high dimensional Gaussian graphical models. *J. Mach. Learn. Res.* **22** Paper No. [274], 62. [MR4353053](https://doi.org/10.48550/jmlr.2021.22.13053)
- LLOPIS, F. P., KANTAS, N., BESKOS, A. and JASRA, A. (2018). Particle filtering for stochastic Navier–Stokes signal observed with linear additive noise. *SIAM J. Sci. Comput.* **40** A1544–A1565. [MR3807951 https://doi.org/10.1137/17M1151900](https://doi.org/10.1137/17M1151900)
- MEINSHAUSEN, N. and BÜHLMANN, P. (2006). High-dimensional graphs and variable selection with the lasso. *Ann. Statist.* **34** 1436–1462. [MR2278363 https://doi.org/10.1214/009053606000000281](https://doi.org/10.1214/009053606000000281)
- MICHIS, A. A. (2022). Multiscale partial correlation clustering of stock market returns. *J. Risk Financ. Manag.* **15** 24.
- MOHAMMADI, R., WIT, E. and DOBRA, A. (2022). BDgraph: Bayesian structure learning in graphical models using birth-death MCMC. R package version 2.72.
- MOLINARI, M., CREMASCHI, A., DE IORIO, M., CHATURVEDI, N., HUGHES, A. and TILLIN, T. (2022a). Bayesian dynamic network modelling: An application to metabolic associations in cardiovascular disease. arXiv preprint, [arXiv:2207.13418](https://arxiv.org/abs/2207.13418).
- MOLINARI, M., CREMASCHI, A., DE IORIO, M., CHATURVEDI, N., HUGHES, A. D. and TILLIN, T. (2022b). Bayesian nonparametric modelling of multiple graphs with an application to ethnic metabolic differences. *J. R. Stat. Soc. Ser. C. Appl. Stat.* **71** 1181–1204. [MR4511106 https://doi.org/10.1111/rssc.12570](https://doi.org/10.1111/rssc.12570)

- MURRAY, I., GHAHRAMANI, Z. and MACKAY, D. J. C. (2006). MCMC for doubly-intractable distributions. In *Proceedings of the Twenty-Second Conference on Uncertainty in Artificial Intelligence, UAI'06*.
- NEWTON, M. A., NOUEIRY, A., SARKAR, D. and AHLQUIST, P. (2004). Detecting differential gene expression with a semiparametric hierarchical mixture method. *Biostatistics* **5** 155–176.
- NILSSON, O. (2016). On stochastic volatility models as an alternative to GARCH type models. Dept. Statistics, Uppsala Univ., Sweden.
- PERSING, A., JASRA, A., BESKOS, A., BALDING, D. and DE IORIO, M. (2015). A simulation approach for change-points on phylogenetic trees. *J. Comput. Biol.* **22** 10–24.
- PETERSON, C., STINGO, F. C. and VANNUCCI, M. (2015). Bayesian inference of multiple Gaussian graphical models. *J. Amer. Statist. Assoc.* **110** 159–174. MR3338494 <https://doi.org/10.1080/01621459.2014.896806>
- PITT, M. K., SILVA, R. D. S., GIORDANI, P. and KOHN, R. (2012). On some properties of Markov chain Monte Carlo simulation methods based on the particle filter. *J. Econometrics* **171** 134–151. MR2991856 <https://doi.org/10.1016/j.jeconom.2012.06.004>
- PRADO, R. and WEST, M. (2010). *Time Series: Modeling, Computation, and Inference*. CRC Press/CRC.
- ROVERATO, A. (2002). Hyper inverse Wishart distribution for non-decomposable graphs and its application to Bayesian inference for Gaussian graphical models. *Scand. J. Stat.* **29** 391–411. MR1925566 <https://doi.org/10.1111/1467-9469.00297>
- ROY, S., ATCHADÉ, Y. and MICHAILIDIS, G. (2017). Change point estimation in high dimensional Markov random-field models. *J. R. Stat. Soc. Ser. B. Stat. Methodol.* **79** 1187–1206. MR3689314 <https://doi.org/10.1111/rssb.12205>
- RUZAYQAT, H., ER-RAIY, A., BESKOS, A., CRISAN, D., JASRA, A. and KANTAS, N. (2022). A lagged particle filter for stable filtering of certain high-dimensional state-space models. *SIAM/ASA J. Uncertain. Quantificat.* **10** 1130–1161. MR4490319 <https://doi.org/10.1137/21M1450392>
- SAKURAI, Y. and KUROSAKI, T. (2020). How has the relationship between oil and the us stock market changed after the Covid-19 crisis? *Finance Res. Lett.* **37** 101773.
- SCHWALLER, L. and ROBIN, S. (2017). Exact Bayesian inference for off-line change-point detection in tree-structured graphical models. *Stat. Comput.* **27** 1331–1345. MR3647100 <https://doi.org/10.1007/s11222-016-9689-3>
- SHERLOCK, C., THIERY, A. H., ROBERTS, G. O. and ROSENTHAL, J. S. (2015). On the efficiency of pseudo-marginal random walk Metropolis algorithms. *Ann. Statist.* **43** 238–275. MR3285606 <https://doi.org/10.1214/14-AOS1278>
- SILVENNOINEN, A. and TERÄSVIRTA, T. (2009). Multivariate GARCH models. In *Handbook of Financial Time Series* 201–229. Springer.
- SO, M. K. P., LAM, K. and LI, W. K. (1998). A stochastic volatility model with Markov switching. *J. Bus. Econom. Statist.* **16** 244–253. MR1615375 <https://doi.org/10.2307/1392580>
- TAYLOR, S. J. (1982). Financial returns modelled by the product of two stochastic processes—a study of the daily sugar prices 1961–75. *Time Ser. Anal. Theory Pract.* **1** 203–226.
- THAMPANYA, N., NASIR, M. A. and HUYNH, T. L. D. (2020). Asymmetric correlation and hedging effectiveness of gold & cryptocurrencies: From pre-industrial to the 4th industrial revolution. *Technol. Forecast. Soc. Change* **159** 120195.
- UHLER, C., LENKOSKI, A. and RICHARDS, D. (2018). Exact formulas for the normalizing constants of Wishart distributions for graphical models. *Ann. Statist.* **46** 90–118. MR3766947 <https://doi.org/10.1214/17-AOS1543>
- VAN LEEUWEN, P. J., CRISAN, D., LANG, O. and POTTHAST, R. (2021). Bayesian inference for fluid dynamics: A case study for the stochastic rotating shallow water model. arXiv preprint, arXiv:2112.15216.
- VAN DEN BOOM, W., BESKOS, A. and DE IORIO, M. (2022). The G -Wishart weighted proposal algorithm: Efficient posterior computation for Gaussian graphical models. *J. Comput. Graph. Statist.* **31** 1215–1224. MR4513382 <https://doi.org/10.1080/10618600.2022.2050250>
- VAN DEN BOOM, W., DE IORIO, M. and BESKOS, A. (2022). Bayesian learning of graph substructures. arXiv preprint, arXiv:2203.11664.
- VAN DEN BOOM, W., JASRA, A., DE IORIO, M., BESKOS, A. and ERIKSSON, J. G. (2022). Unbiased approximation of posteriors via coupled particle Markov chain Monte Carlo. *Stat. Comput.* **32** Paper No. 36, 19. MR4414445 <https://doi.org/10.1007/s11222-022-10093-3>
- VON FURSTENBERG, G. M., JEON, B. N., MANKIW, N. G. and SHILLER, R. J. (1989). International stock price movements: Links and messages. *Brookings Pap. Econ. Act.* **1989** 125–179.
- WANG, H. (2010). Sparse seemingly unrelated regression modelling: Applications in finance and econometrics. *Comput. Statist. Data Anal.* **54** 2866–2877. MR2720481 <https://doi.org/10.1016/j.csda.2010.03.028>
- WANG, H. and WEST, M. (2009). Bayesian analysis of matrix normal graphical models. *Biometrika* **96** 821–834. MR2564493 <https://doi.org/10.1093/biomet/asp049>

- WARNICK, R., GUINDANI, M., ERHARDT, E., ALLEN, E., CALHOUN, V. and VANNUCCI, M. (2018). A Bayesian approach for estimating dynamic functional network connectivity in fMRI data. *J. Amer. Statist. Assoc.* **113** 134–151. MR3803445 <https://doi.org/10.1080/01621459.2017.1379404>
- WATTS, D. J. and STROGATZ, S. H. (1998). Collective dynamics of ‘small-world’ networks. *Nature* **393** 440–442.
- WEST, M. and HARRISON, J. (2006). *Bayesian Forecasting and Dynamic Models*. Springer Science & Business Media. MR1020301 <https://doi.org/10.1007/978-1-4757-9365-9>
- WIGGINS, J. B. (1987). Option values under stochastic volatility: Theory and empirical estimates. *J. Financ. Econ.* **19** 351–372.
- WILLIAMS, D. R. (2021). Bayesian estimation for Gaussian graphical models: Structure learning, predictability, and network comparisons. *Multivar. Behav. Res.* **56** 336–352. <https://doi.org/10.1080/00273171.2021.1894412>
- YANG, J. and PENG, J. (2020). Estimating time-varying graphical models. *J. Comput. Graph. Statist.* **29** 191–202. MR4085874 <https://doi.org/10.1080/10618600.2019.1647848>
- YANG, S., LU, Z., SHEN, X., WONKA, P. and YE, J. (2015). Fused multiple graphical lasso. *SIAM J. Optim.* **25** 916–943. MR3343365 <https://doi.org/10.1137/130936397>
- YILDIRIM, S., SINGH, S. S. and DOUCET, A. (2013). An online expectation-maximization algorithm for change-point models. *J. Comput. Graph. Statist.* **22** 906–926. MR3173749 <https://doi.org/10.1080/10618600.2012.674653>
- YIN, J. and LI, H. (2011). A sparse conditional Gaussian graphical model for analysis of genetical genomics data. *Ann. Appl. Stat.* **5** 2630–2650. MR2907129 <https://doi.org/10.1214/11-AOAS494>
- YOUSFI, M., ZAIED, Y. B., CHEIKH, N. B., LAHOUEL, B. B. and BOUZGARROU, H. (2021). Effects of the COVID-19 pandemic on the us stock market and uncertainty: A comparative assessment between the first and second waves. *Technol. Forecast. Soc. Change* **167** 120710.
- ZHOU, S., LAFFERTY, J. and WASSERMAN, L. (2010). Time varying undirected graphs. *Mach. Learn.* **80** 295–319. MR3108169 <https://doi.org/10.1007/s10994-010-5180-0>
- ZOU, H. (2006). The adaptive lasso and its oracle properties. *J. Amer. Statist. Assoc.* **101** 1418–1429. MR2279469 <https://doi.org/10.1198/016214506000000735>