



Contents lists available at ScienceDirect

Journal of Econometrics

journal homepage: www.elsevier.com/locate/jeconom

A flexible predictive density combination for large financial data sets in regular and crisis periods[☆]



Roberto Casarin^a, Stefano Grassi^b, Francesco Ravazzolo^{c,d},
Herman K. van Dijk^{e,f,*}

^a Ca' Foscari University of Venice, Italy

^b University of Rome Tor Vergata, Italy

^c BI Norwegian Business School, Norway

^d Free University of Bozen-Bolzano, Italy

^e Erasmus University Rotterdam, Tinbergen Institute, The Netherlands

^f Norges Bank, Norway

ARTICLE INFO

Article history:

Received 24 November 2020

Received in revised form 16 November 2022

Accepted 27 November 2022

Available online 7 January 2023

JEL classification:

C11

C15

C53

E37

Keywords:

Density combination

Large set of predictive densities

Dynamic factor models

Nonlinear state-space

Bayesian inference

ABSTRACT

A flexible predictive density combination is introduced for large financial data sets which allows for model set incompleteness. Dimension reduction procedures that include learning allocate the large sets of predictive densities and combination weights to relatively small subsets. Given the representation of the probability model in extended nonlinear state-space form, efficient simulation-based Bayesian inference is proposed using parallel dynamic clustering as well as nonlinear filtering, implemented on graphics processing units. The approach is applied to combine predictive densities based on a large number of individual US stock returns of daily observations over a period that includes the Covid-19 crisis period. Evidence on dynamic cluster composition, weight patterns and model set incompleteness gives valuable signals for improved modelling. This enables higher predictive accuracy and better assessment of uncertainty and risk for investment fund management.

© 2022 Erasmus University Rotterdam. Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Predicting with large sets of data involving many model structures and explanatory variables is a topic of substantial interest to academic researchers as well as to professional forecasters. It has been studied in several papers (e.g., see [Stock](#)

[☆] This paper should not be reported as representing the views of Norges Bank. The views expressed are those of the authors and do not necessarily reflect those of Norges Bank. Paper was presented at the ISBA2022 World Meeting. The authors are indebted to Mike West, James Mitchell and several participants for very useful discussions and comments. We are also indebted to Jamie Cross, Lennart Hoogerheide and, in particular, to the editor, Torben Andersen, the associate editor and three anonymous referees for valuable comments on earlier versions of this paper, [Casarin et al. \(2020\)](#) which led to a substantial revision and extension. Roberto Casarin's research is supported by the Venice center in Economic and Risk Analytics for public policies. Stefano Grassi gratefully acknowledges financial support from the University of Rome 'Tor Vergata', Italy under Grant "Beyond Borders" (CUP: E84I20000900005). Roberto Casarin and Francesco Ravazzolo acknowledge financial support from Italian Ministry MIUR under the PRIN project Hi-Di NET - Econometric Analysis of High Dimensional Models with Network Structures in Macroeconomics and Finance (grant 2017TA71YC).

* Corresponding author.

E-mail address: hkvandijk@ese.eur.nl (H. K. van Dijk).

and Watson, 1999, 2002, 2005, 2014, and Bañbura et al., 2010). The recent fast growth in big data allows researchers to predict variables of interest more accurately (e.g., see Choi and Varian, 2012; Varian, 2014; Varian and Scott, 2014; Einav and Levin, 2014). Stock and Watson (2005, 2014), Bañbura et al. (2010) and Koop and Korobilis (2013) suggest that there are also potential gains from predicting using a large set of predictors. However, such predictions require new modelling strategies, efficient inference methods and extra computing power possibly resulting from parallel computing. We refer to Granger (1998) for an early discussion of these issues.

Given a large financial micro-data set consisting of US individual stock returns of daily observations over an extended period which includes the Covid-19 crisis period, we propose a flexible predictive density combination in order to approximate such data accurately while allowing for model set incompleteness and combination weight learning. A major motivation for our approach is that in portfolio analysis the process of collecting such a large data set and clustering it in a relatively small number of groups where each group has typical data characteristics is a popular strategy. Next, one replicates an aggregate index through a weighted combination of assets which can be used for a predictive asset allocation strategy in investment funds (e.g., see Corielli and Marcellino, 2006; Kim and Kim, 2020). Our approach provides tools for the quantification of predictive uncertainty and risk for such funds which may be useful management information. A second related motivation is the appearance of fat tail behaviour (and possibly skewness) in most financial data distributions. These data features are usually ignored in point prediction and this may give wrong signals about financial uncertainty and risk.

In terms of methodology, we extend the combination strategies of Billio et al. (2013) and McAlinn and West (2019). McAlinn and West (2019) is founded on a decision theory framework providing a coherent Bayesian interpretation for Bayesian Model Averaging (BMA). However, when the time series tend to become large BMA tends to select one model, supposedly the true model, see Amisano and Geweke (2010). In empirical analysis it is, however, known that the concept of a true model is not very realistic, see Geweke (2010). Our aim is not to pursue the approach of finding the true model, but rather to introduce a flexible predictive density combination that provides an accurate approximation to nonstandard data distributions in finance such as bimodal ones and/or distributions with heavy tails. We refer to Hall and Mitchell (2007), Amisano and Geweke (2010), Billio et al. (2013), Gneiting and Ranjan (2013), Yao et al. (2018) and the discussion therein for further insight on the different combination approaches.

Our extensions to Billio et al. (2013) are threefold. First, replace their normal combination model with a flexible mixture of predictive distributions to account for possible multimodality and heavy tails and specify a measure for model incompleteness or misspecification. Second, make the combination approach operational for large data sets by applying dimension reduction which includes learning about the large sets of predictive densities as well as the combination weights. A major reason for these different dimension reductions is that investment fund managers are interested to obtain not only predictions of the fund but also to learn about the composition over time of the clusters of assets and, also, to learn about the behaviour of weights of these clusters. This may lead to improved modelling, prediction and policy. Third, for efficient inferential purposes make use of a recent parallel Sequential Monte Carlo method.

Extension one is described in the beginning of Section 2, here we explain extensions two and three. Our empirical example contains a large data set of individual US stock returns of daily observations over an extended time period which includes the Covid-19 crisis. As a consequence, we deal with a large set of predictive densities in the combination process, which makes the inference task a substantial challenge. Inference based on the normal combination model from (Billio et al., 2013) is not operational in large panels since the latent space of combination weights is high dimensional and overparameterization and overfitting issues can easily arise. We extend (Billio et al., 2013) with two dimension reductions. The first one is based on dynamic clustering of the large set of predictive densities exploiting common data features in the large set of stock return series such as wide data bands and time-varying volatility. The dynamic clustering maps, at each time t , the large set of predictive densities to a much smaller subset where the cluster composition at time t depends on the past composition. We note that clustering strategies have been successfully used in other models to cope with high dimensional parameter spaces, MacLehose and Dunson (e.g., see 2010), Billio et al. (e.g., see 2019). The second dimension reduction deals with the large number of combination weights. A nonlinear dynamic factor model is specified for these weights which contains learning, possibly, using information about past predictive performance. An alternative is to shrink the combination weights to zero as in the sparse factor model literature, Carvalho et al. (e.g., see 2008), Kaufmann and Schumacher (e.g., see 2017, 2019) but learning is not included. We note that our approach contributes to the literature on time series models with time-varying parameters that take values on a bounded domain, see, e.g., Aitchinson and Shen (1980) and Aitchinson (1982), and applies it to large financial data extending the intuition in Stock and Watson (2014).

With respect to the inference method we show that our mixture model allows for an extended nonlinear state-space representation. This enables us to construct an efficient simulation-based Bayesian inference procedure, where parallel Sequential Monte Carlo is used to filter the set of probabilistic weights and integrate the set of random parameters. Here we make use of the recently developed M-Filter, see Baştürk et al. (2019) and Hoogerheide et al. (2012).

In terms of empirical analysis we provide three contributions: accuracy gains in predictive moments compared to benchmark results; time-varying composition of clusters and cluster weights; diagnostic information on model set incompleteness.

Evidence on substantial accuracy gains in predictive means, volatilities and tail events is presented compared to the no-predictive ability benchmark and predictions from individual models and combination methods as BMA and Equal Weights (EW). The time-varying composition of the set of clusters shows learning. Individual stocks may switch across clusters

or eventually exit them, for example, during and after a crisis like the financial crisis and the Covid-19 crisis. Measures of model set incompleteness and dynamic patterns in the cluster-based weights give valuable diagnostic signals. These empirical results may provide useful information for improved financial modelling and policy analysis by investment fund management.

For a recent survey about the evolution of predictive density combinations we refer to [Aastveit et al. \(2019\)](#), for background to [Billio et al. \(2013\)](#), [McAlinn and West \(2019\)](#) and for a policy application to [Baştürk et al. \(2019\)](#).

The contents of this paper are structured as follows. Section 2 provides details of the methodological contributions of our approach. Section 3 contains an efficient simulation-based Bayesian inference procedure. Section 4 contains results of the empirical application using a large set of US stock return data in regular and crisis periods. Section 5 presents conclusions and suggestions for further research. Some additional results are given in Supplementary Material that serves as an online Appendix.

2. Mixture process with model set incompleteness, dimension reduction and time-varying component weights

We start with extending a standard mixture process for predictive densities to include model set incompleteness. Let the conditional predictive probability distribution of a financial variable of interest, y_t , be specified as a mixture of conditional predictive probability distributions of y_t coming from a large set of n individual financial models with information sets $\mathcal{J}_{i,t-1}$, where the information set $\mathcal{J}_{i,t-1}$ includes data information as well as model structure, denoted by M_i , with $i = 1, \dots, n$. Define weights w_{it} that form a convex combination of the conditional predictive probabilities. In terms of densities this implies a standard mixture process:

$$f(y_t|\mathcal{J}_{t-1}) = \sum_{i=1}^n w_{it}f(y_t|\mathcal{J}_{i,t-1}), \quad 0 \leq w_{it} \leq 1, \quad \sum_{i=1}^n w_{it} = 1, \quad (1)$$

where \mathcal{J}_{t-1} is the joint information set.

A key step is to give specific content to the i th mixture component $f(y_t|\mathcal{J}_{i,t-1})$. Let $y_t = y_{it}^*$ with probability w_{it} , where y_{it}^* is defined for all models M_i as the sum of the following two random variables:

$$y_{it}^* = \tilde{y}_{it} + \varepsilon_{it}, \quad (2)$$

where \tilde{y}_{it} is a generated draw from the predictive distribution with density $f(\tilde{y}_{it}|\mathcal{J}_{i,t-1})$ from model M_i . A new feature is the addition of the disturbance ε_{it} . It points towards two sources of error. There may be misspecification errors due to model set incompleteness and prediction errors due to, for instance, sudden shocks in the series. In this paper we focus on the former, that is, a larger specification error implies a larger error ε_{it} . Investigating the relative importance of a prediction error component is a topic for further research. We note that [Terui and van Dijk \(2002\)](#), [Hoogerheide et al. \(2010\)](#), [Takanashi and McAlinn \(2019\)](#), [McAlinn and West \(2019\)](#) and recently [Aastveit et al., 2022](#) make use of a combination equation which can be interpreted as a linear regression model with time-varying parameters and a time-varying constant and one disturbance term. In contrast, we work with a flexible mixture approach in the combination process where for each component of the mixture there exists a time-varying weight and a disturbance.

The probability density function of ε_{it} is given as:

$$\varepsilon_{it} \sim \mathcal{N}(0, \sigma_{it}^2), \quad (3)$$

where

$$\sigma_{it}^2 = \sigma_i^2 \exp(h_{it}), \quad h_{it} = h_{i,t-1} + \zeta_{it}, \quad \zeta_{it} \sim \mathcal{N}(0, \sigma_{\zeta,i}^2), \quad (4)$$

for each $i = 1, \dots, n$. That is, σ_{it}^2 follows a Stochastic volatility (SV) process. The use of an SV specification is introduced by [Clark \(1973\)](#) and further extended by [Taylor \(1982\)](#), in particular for asset prices, see the survey by [Shephard and Andersen \(2009\)](#). Our random walk assumption is chosen as it makes the specification convenient and parsimonious and the MCMC estimation method becomes straightforward and efficient. For one-period prediction, to which we restrict our attention in this paper, the random walk assumption does not hamper the predictive performance. For long-term forecasting the random walk specification may adversely affect (density) forecast performance by allowing volatility to blow up (becoming either unduly high or low). In such a situation, application of stationary SV processes, see [Shephard and Andersen \(2009\)](#) and [Johannes et al. \(2014\)](#), or stationary GARCH specifications, see [Andersen et al. \(2010\)](#), are interesting alternative specifications to explore. This topic is left for further research.

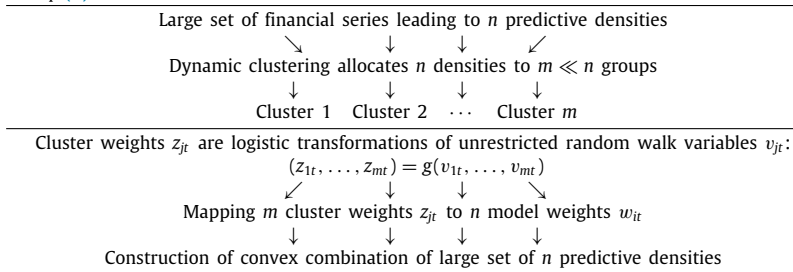
Given Eq. (2), the probability density function of y_{it}^* is the convolution of two densities given as:

$$f(y_{it}^*|\mathcal{J}_{i,t-1}, \sigma_{it}^2) = \int_{\mathbb{R}} \frac{1}{\sigma_{it}} \phi\left(\frac{y_{it}^* - \tilde{y}_{it}}{\sigma_{it}}\right) f(\tilde{y}_{it}|\mathcal{J}_{i,t-1}) d\tilde{y}_{it} \quad (5)$$

where $\phi(\cdot)$ is the standard normal density. So, for our mixture of n models we have

$$f(y_t|\mathcal{J}_{t-1}, \sigma_t^2) = \sum_{i=1}^n w_{it} \int_{\mathbb{R}} \frac{1}{\sigma_{it}} \phi\left(\frac{y_t - \tilde{y}_{it}}{\sigma_{it}}\right) f(\tilde{y}_{it}|\mathcal{J}_{i,t-1}) d\tilde{y}_{it} \quad (6)$$

Table 1
Dimension reduction and computation of latent weights for predictive density combination using a large financial data set. The function $g(\cdot)$ refers to the logistic transformation given in Eq. (8).



where $\sigma_t^2 = \{\sigma_{1t}^2, \dots, \sigma_{nt}^2\}$. One may interpret the information from the predictive densities $f(\tilde{y}_{it} | \mathcal{I}_{t-1})$, $i = 1, \dots, n$ as *prior-predictive* information that is fed into the predictive density combination (6).

However, a large set of predictive densities for each time observation allowing for time-varying combination weights is not easy to handle in terms of econometric inference. Reduction of the large information set is necessary. Dimension reduction techniques are widely used in machine learning to reduce the information of high-dimensional data sets (see e.g., Varian, 2014; Casarin and Veggente, 2020, and references cited therein). We make use of dimension reduction steps for the large number of components in the predictive mixture and the accompanying large number of latent probabilistic weights. The different steps are schematically shown in Table 1.

In our financial case, we start with a preliminary step using diagnostic graphical evidence about typical data features in individual financial series of stock returns like high and low time-varying volatility and wide and narrow data bands. Our motivation for this is that in financial applications large differences across predictions occur in the higher moments and tail behaviour. This leads in our case to the, *a priori*, choice of a Normal density with high and low volatility and a Student- t density with large and small degrees of freedom. We use these data features for dimension reduction of the large number of components in the predictive mixture into four groups using dynamic K-means clustering where the current cluster composition depends on the past composition. This allows for learning about model dependence, cluster grouping and model dimensionality reduction, see e.g. Varian (2014), Casarin and Veggente (2020), and for more general applications of machine learning methods to financial predictions (Gu et al., 2020; Bianchi et al., 2020). This is well documented in data but largely ignored in the predictive density combination literature. We note that Bianchi and McAlinn (2018) and Takanashi and McAlinn (2019) also discuss clustering procedures. Usually, these methods make use of constant clustering. The clustering process is depicted in the second row of Table 1 and empirical results are presented in Section 4. For general background on K-Means clustering we refer to Frühwirth-Schnatter (2006, pp. 97) and Malsiner Walli et al. (2016) and for details about the implementation of our dynamic clustering procedure we refer to Favirar et al. (2008) and to Appendix B.1 in the Supplementary Material.

In order to complete the specification of the probability model given in Eq. (6), we specify a law of motion, which involves dimension reduction and learning, for the latent probabilistic weights w_{it} , $i = 1, \dots, n$, $t = 1, \dots, T$ with large n and large T . The weight model, given in Eqs. (7)–(9), is a nonlinear dynamic factor model of w_{it} . Factor models are well-known as a vehicle for dimension reduction, see Anderson (1984) and Lopes and West (2004) and we make use of a simple nonlinear extension. Let

$$v_{jt} = v_{j,t-1} + \eta_{jt}, \quad \eta_{jt} \sim \mathcal{N}(0, \sigma_{\eta_j}^2), \quad j = 1, \dots, m, \tag{7}$$

$$z_{jt} = \frac{\exp(v_{jt})}{\sum_{i=1}^m \exp(v_{it})}, \quad 0 \leq z_{jt} \leq 1, \quad \sum_{j=1}^m z_{jt} = 1, \tag{8}$$

$$w_{it} = \sum_{j=1}^m b_{ijt} z_{jt}, \quad i = 1, \dots, n. \tag{9}$$

The unrestricted latent variables v_{jt} , $j = 1, \dots, m$ in the $m = 4$ clusters are normal random walk variables, see Eq. (7). This dynamic specification can be generalised, see Billio et al. (2013). Next, a logistic transformation is applied from the unrestricted latent variables v_{jt} , $j = 1, \dots, m$ to an *auxiliary* set of probabilistic cluster weights z_{jt} , $j = 1, \dots, m$, see Eq. (8), which is depicted in the bottom part of Table 1. The purpose of the factor loadings b_{ijt} is to transform the m weights z_{jt} of the convex combination of m clusters to the n weights w_{it} of the convex combination of $n \gg m$ models. Therefore, the factor loadings b_{ijt} are restricted to be nonnegative and sum to 1. We assume in our case, for convenience, that each model i has an equal weight within its cluster j . To allow for learning one may consider a case where the weights are driven by model-specific predictive performance using the log score, see Billio et al. (2013) and also (Mitchell and Hall, 2005). Note that in our case there exists no noise in the connection between w_{it} , z_{jt} and v_{jt} , but this assumption can be

relaxed to account for contemporaneous uncertainty with the addition of logistic-normal noise in the equation for w_{it} . This is left for further research.

Given Eqs. (7)–(9) we can complete the specification of the probability model given in (6) with the specification of the law of motion for the probabilistic weights w_{it} , $i = 1, \dots, n$. Consider the $m \times 1$ vector $\mathbf{v}_t = (v_{1t}, \dots, v_{mt})'$ with the multivariate normal distribution $\mathbf{v}_t \sim \mathcal{N}_m(\mathbf{v}_{t-1}, \Sigma_\eta)$ where Σ_η is a diagonal matrix. Given that the $m \times 1$ vector $\mathbf{z}_t = (z_{1t}, \dots, z_{mt})'$ is given as a logistic transformation, see Eq. (8), the vector $\tilde{\mathbf{z}}_t = (z_{1t}, \dots, z_{m-1,t})'$ follows a logistic-normal distribution given as $\tilde{\mathbf{z}}_t \sim \mathcal{L}_{m-1}(\mathbf{D}\mathbf{v}_{t-1}, \mathbf{D}\Sigma_\eta\mathbf{D}')$ with $z_{mt} = 1 - \sum_{j=1}^{m-1} z_{jt}$, where the expression of the matrix \mathbf{D} is given in the proof of the result which is presented in the Supplementary Material, Appendix A.

Since the large set of n probabilistic weights w_{it} , $i = 1, \dots, n$ consists of linear combinations of the small set of m logistic-normal probabilistic weights z_{jt} , $j = 1, \dots, m$, see Eq. (9), the distribution of the weights w_{it} is logistic-normal. Here, use is made of the class preserving property of the logistic-normal distribution. In matrix notation, consider the vector $\tilde{\mathbf{w}}_t = (w_{1t}, \dots, w_{n-1,t})'$ with $w_{nt} = 1 - \sum_{i=1}^{n-1} w_{it}$. Then we make use of $\tilde{\mathbf{w}}_t = \tilde{\mathbf{B}}_t \tilde{\mathbf{z}}_t$ where $\tilde{\mathbf{B}}_t$ is an $(n-1) \times m$ matrix that contains an appropriate subset of elements b_{ijt} . As a consequence, one can write that the implied logistic-normal distribution of $\tilde{\mathbf{w}}_t$ is given as :

$$\tilde{\mathbf{w}}_t \sim \mathcal{L}_{n-1}(\tilde{\mathbf{B}}_t \mathbf{D} \mathbf{v}_{t-1}, \tilde{\mathbf{B}}_t \mathbf{D} \Sigma_\eta \mathbf{D}' \tilde{\mathbf{B}}_t'), \tag{10}$$

for details on this result, see the Supplementary Material, Appendix A. Note that the density of the complete vector $\mathbf{w}_t = (w_{1t}, \dots, w_{nt})'$ is singular due to the adding-up restriction of the n weights. A second source of degeneracy is intrinsic to our projection strategy which implies rank deficiency of the matrix $\mathbf{B}_t \mathbf{D} \Sigma_\eta \mathbf{D}' \mathbf{B}_t'$.¹

The analytic solution of the probability model of Eqs. (6) and (10) is generally not known but the model can be represented in extended nonlinear state-space form given as:

$$y_t = \sum_{i=1}^n (\tilde{y}_{it} + \varepsilon_{it}) s_{it}, \quad \varepsilon_{it} \sim \mathcal{N}(0, \sigma_{it}^2) \tag{11}$$

$$(s_{1t}, \dots, s_{nt}) \sim \mathcal{M}_n(1, (w_{1t}, \dots, w_{nt})), \tag{12}$$

$$w_{it} = \sum_{j=1}^m b_{ijt} z_{jt}, \quad i = 1, \dots, n, \tag{13}$$

$$z_{jt} = \frac{\exp(v_{jt})}{\sum_{i=1}^m \exp(v_{it})}, \quad j = 1, \dots, m, \tag{14}$$

$$v_{jt} = v_{j,t-1} + \eta_{jt}, \quad \eta_{jt} \sim \mathcal{N}(0, \sigma_{\eta_j}^2), \quad j = 1, \dots, m. \tag{15}$$

The measurement Eq. (11) has as extension that the right-hand side variable \tilde{y}_{it} is not an observation from our data set but refers to a random draw from the predictive distribution of model M_i and ε_{it} gives an indication of possible incompleteness. Using the Multinoulli distribution (also known as the Categorical distribution) of (s_{1t}, \dots, s_{nt}) with parameter vector (w_{1t}, \dots, w_{nt}) given in Eq. (12), one generates draws from the i th component of the mixture distribution with probability w_{it} where (s_{1t}, \dots, s_{nt}) contains $n - 1$ 0's and one element equal to 1. That is, $s_{it} = 1$ means that model i is selected.

A schematic figure of the extended nonlinear state-space representation is given in Fig. 1. The sequence of steps starts with the dynamic clustering step shown at the bottom left, then one proceeds upwards and next to the middle centre where the nonlinear dynamic factor model is shown which is constructed at the bottom right and going upwards. At the top, the measurement equation is shown from the finite mixture process with the n generated predictive draws and disturbances that follow a stochastic volatility process. The complete state-space representation enables us to make use of filtering methods of the nonlinear time series literature to evaluate and update the unobserved weight components in the predictive density combination.

3. Bayesian inference applying the M-filter

The analytic solution of the optimal filtering problem is in most applications not known, except for the case of a Kalman filter where use is made of well-known properties of the multivariate normal distribution. For our large finite mixture process with time-varying weights based on a nonlinear dynamic factor model, we make use of simulation-based numerical methods to tackle the filtering problem and we make use of Bayesian inference to update information about the random model parameters. For a fundamental discussion about a coherent Bayesian framework for evaluation, calibration, and data-informed combination of multiple predictive densities, see McAlinn and West (2019), McAlinn et al. (2020).

¹ Other distributions can be used for weights such as the Dirichlet distribution, but as shown in Aitchinson and Shen (1980) this distributional assumption can be too restrictive in our analysis since the components of a Dirichlet composition have a correlation structure determined solely by the normalisation operation, so that only negative correlations are possible (and given the means, there is only one free parameter for the variances and covariances).




Fig. 1. Directed acyclic graph of the extended nonlinear state-space model. It shows the connections between the variable of interest y_t , the predicted variables \tilde{y}_{it} (rectangles, solid line), the dynamic clustering, the random walk process, the categorical Multinoulli distribution, the stochastic volatility specification (ellipses), the link functions $\frac{\exp(v_{jt} - v_{mt})}{\sum_{k=1}^m \exp(v_{kt} - v_{mt})}$, b_{jit} and the nonlinear dynamic factor model of the latent probabilistic weights w_{it} (rectangles, dashed line). The directed arrows show the dependence structure.

Apart from the fundamental motivation for applying Bayesian inference there exists a practical one which is based on simulation-based Bayesian methods. That is, the generated draws \tilde{y}_{it} from the predictive distributions of the different models are carried forward into the predictive combination density. Thus, the uncertainty in the predictions of the different models carries directly forward into the uncertainty of the combined predictive density. In contrast, frequentist methods like method of moments or maximum likelihood proceed in a two-step fashion by first computing point predictions for the different models and substituting these in a second stage into the combined predictive density. As such the second stage results suffer from the generated regressor problem, that is, uncertainty measures of the second-stage predictions are sensitive to the estimation procedure used for the model predictions, see Pagan (1984) for background details.

In order to filter the latent weights in Eqs. (11)–(14) we make use of the recently developed M-Filter introduced in Baştürk et al. (2019) and based on (Hoogerheide et al., 2012). We restrict ourselves to a summary of the novel, efficient and robust properties of this method and refer to the cited references and Appendix B.2 in the Supplementary Material for technical background.

The M-Filter is a member of the class of Sequential Monte Carlo (SMC) algorithms that are suitable to solve filtering problems in nonlinear state-space models, see Creal (2007) and Herbst and Schorfheide (2014) for background. The basic idea of SMC is that a set of draws, labelled particles, is generated from an approximation to the so-called target density, in our case, the combined predictive density, in a sequential way at each time t . That is, such an SMC method consists in general at each time point t of two steps: a sampling step to generate particles from the approximate density and next a correction step to adjust for the distance between the approximate density and the target. However, the propagation of particles over time points leads to the situation that after many iterations one particle receives all weights which implies degeneracy of the approximate density. In order to avoid this, a so-called resampling step is introduced where the approximate density is updated in such a way that degeneracy does not occur.

The M-filter has two innovations. First, this rather cumbersome resampling procedure in the propagation step is replaced by independent sampling at each time t . This also reduces Monte Carlo variation. Second, the independent sampling step occurs in an online-fashion, contrary to other methods that make use of independent sampling like Efficient Importance Sampling of Richard and Zhang (2007) and Liesenfeld and Richard (2003), or Numerically Accelerated Importance Sampling of Koopman et al. (2015) that are off-line methods.

Table 2

Computing time for number of threads (1, 6, 12, 24), number of draws (5000, 10000, 20000) and maximum number of Student- t components (C). All values are in ratio with the benchmark model given by 5000 draws and one thread. Values higher than 1 means that the computing time is higher than the benchmark. Values lower than 1 means that the computing time is lower than the benchmark.

C	1 Thread			6 Threads			12 Threads			24 Threads		
	5000	10000	20000	5000	10000	20000	5000	10000	20000	5000	10000	20000
2	1.000	1.953	3.773	0.171	0.314	0.725	0.093	0.171	0.377	0.083	0.164	0.325
3	1.067	2.086	3.970	0.178	0.317	0.750	0.098	0.181	0.379	0.083	0.165	0.326
4	1.068	2.097	3.981	0.186	0.322	0.782	0.099	0.185	0.383	0.083	0.168	0.331

The choice of an accurate approximate density is crucial for the performance of any filter method and has received considerable attention in the SMC literature, see Doucet et al. (2001), Liu (2001), Kunsch (2005) and Creal (2012). The M-Filter method approximates a target density using the *Mixture of t by Importance Sampling Weighted Expectation–Maximisation* (MitISEM) algorithm proposed by Hoogerheide et al. (2012) and further developed in Baştürk et al. (2016). MitISEM has been shown to be able to effectively approximate complex, non-elliptical distributions due to two features of this algorithm: the class of importance distributions (mixtures of multivariate Student's t distributions), and their joint optimisation (with the Expectation–Maximisation, EM, algorithm). The former allows to closely track distributions of nonstandard shape (multimodal, skewed). The latter is iteratively carried out with the objective of minimising the Kullback–Leibler (KL) divergence between target density and approximate density. This robustness and flexibility of the M-Filter in constructing approximate densities is particularly important in econometrics where breaks in time series are often observed.

The application of the M-Filter requires to choose the values of a few tuning parameters. First, the number of components of the Student- t mixture to approximate the target density is restricted to be no larger than a maximum of four. For each component the initial mean, variance and degrees of freedom parameter are set to zero, one and three, respectively. Those parameters are then updated with an EM step. Second, the number of draws is fixed to 10000 which is sufficient for good convergence in our case. Third, the measure that is used in order to check the convergence of the algorithm is the relative change in the Coefficient of Variation (CoV) of the Importance Sampling weights, where the Importance Sampling weight is the ratio of target and approximate density. The default convergence criterion is chosen as the change of the CoV being smaller than 0.5%.

The M-Filter is easy to parallelise, this enables our approach to speed up the computations using Multiple CPUs or GPUs. In Table 2, we report results from an experiment about the computing time of the M-Filter with different numbers of threads, numbers of draws and maximum numbers of components of the Student- t mixture. All the values reported are in ratio with the benchmark model given by 5000 draws, one thread and two threads. For a comparison of the computing time between the M-Filter and particle filters we refer to Baştürk et al. (2019), where extensive Monte Carlo studies in benchmark exercises are presented. The results reported in Table 2 show that using multiple threads reduces the computing time considerably in our case. Results are not sensitive to the choice of the maximum number of components in the mixture approximation, while doubling the number of draws also doubles the computing time.

The information on the priors for the SV model in Eq. (4) and for the random walk in Eq. (7) is given as follows. For the σ_{it}^2 we select a log-normal distribution with mean $\log(0.1)$ and standard deviation 0.175; for σ_c^2 , we also select a log-normal distribution with mean -2.3 and standard deviation 0.175. Our prior corresponds to an average incompleteness value that is equal to 10% of the unconditional volatility of our S&P500 data. For the σ_{η}^2 , in Eq. (7), we select a log-normal distribution with mean -0.7 and standard deviation equal to 0.175. Therefore the prior mean for σ_{η}^2 is 5 times higher than for σ_c^2 , assuming bigger changes for the cluster weights than for the incompleteness over time. The standard deviation of 0.175 for both priors implies a relative loose prior in both cases. For the parameters in the four equation models we make use of rather uninformative proper priors. All our priors are default choices, that can be modified by the user, in the MATLAB toolbox that carries out the M-Filter, which is available at <http://www.francescoravazzolo.com/pages/DeCo.html>.

4. Predicting and tracking the S&P500

As discussed in the introduction many investors of mutual funds, hedge funds and exchange-traded funds try to replicate the performance of the S&P500 index by holding a set of stocks, which are not necessarily the exact same stocks included in the index.

We collected 496 individual daily stock prices, components of the S&P500, from Datastream over the sample January 2, 2014, to June 30, 2021, for a total of 1888 daily observations for each series. We computed the time series of log-returns for all stocks, see Fig. 2. Table 3 reports several cross-section average statistics of the individual series together with the same statistics for S&P500 index. Some series have much lower average returns than the index with volatility up to 3 times higher than the index. Heterogeneity in skewness and kurtosis is also evident with the series with the lowest skewness equal to -1.829 and the highest skewness equal to 0.335 and with the lowest kurtosis equal to 9.135 and the highest kurtosis equal to 42.164 .




Fig. 2. Daily (%) log-returns for 496 individual stock components of the S&P500 over the sample January 2, 2014 to June 30, 2021. The red line on January 2, 2019 indicates the beginning of the out-of-sample period.

Table 3

Average cross-section statistics for the 500 daily log-returns of individual stocks for the sample January 2, 2014 to June 30, 2021. The columns “Lower”, “Median” and “Upper” refer to the cross-section 5% lower quantile, median and 95% upper quantile of the 500 statistics in rows, respectively. The rows “Average”, “St dev”, “Skewness”, “Kurtosis”, “Min” and “Max” refer to sample average, sample standard deviation, sample skewness, sample kurtosis, sample minimum and sample maximum statistics, respectively. The column “S&P500” reports the sample statistics for the log-returns of the aggregate S&P500 index.

	Subcomponents			S&P500
	Lower	Median	Upper	
Average	-0.012	0.051	0.116	0.045
St dev	1.328	1.810	2.997	1.111
Skewness	-1.829	-0.469	0.335	-1.045
Kurtosis	9.135	16.707	42.164	24.640
Min	-34.907	-17.013	-10.182	-12.765
Max	9.551	14.662	26.722	8.968

The inclusion in our series of the Covid-19 pandemic explains such high values.² We report results on several features of the combined predictive density of the replication of the S&P500 index, including the economic value of tail events.

Diagnostic determination of four clusters and model estimation. Given the basic statistics about time-series and cross-section patterns, we have determined as typical data features of our financial micro-data set wide and narrow data bands and high and low time-varying volatility.

This led us to specify two clusters of predictive densities based on a Normal GARCH(1,1) model: one cluster with high volatility (labelled $n1$) and one cluster with low volatility (labelled $n2$). Next, two clusters based on a Student- t GARCH(1,1) model: one cluster with low degrees of freedom (labelled $t1$) and one cluster with high degrees of freedom (labelled $t2$).³ Our motivation for this choice is to obtain a mixture of densities which fits both in the centre of the empirical distribution (for mean prediction) and in the tails (for measuring uncertainty and risk).

To ease on the computational workload, we have applied an optimisation method to estimate the posterior modes of the parameters from a Normal GARCH(1,1) model and a Student- t GARCH(1,1) model. Given our rather uninformative proper priors, these mode estimates are equal to approximate Bayes mean estimates.

We use rolling samples of 1258 trading days (about five years) for each stock return in the model:

$$y_{it} = c_i + \kappa_{it}\xi_{it} \tag{16}$$

$$\kappa_{it}^2 = \theta_{i0} + \theta_{i1}(y_{i,t-1} - c_i)^2 + \theta_{i2}\kappa_{i,t-1}^2, \quad i = 1, 2, \dots, n, \tag{17}$$

² It has been suggested to make use of the information about shares outstanding to determine better the time behaviour of weights. We consider only stocks that have survived in the S&P500 basket over the period of our sample, but the methodology can handle series with different sample lengths. We leave these as topics for further research.

³ Low degrees of freedom occur jointly with a large scale and high degrees of freedom occur jointly with a low scale.

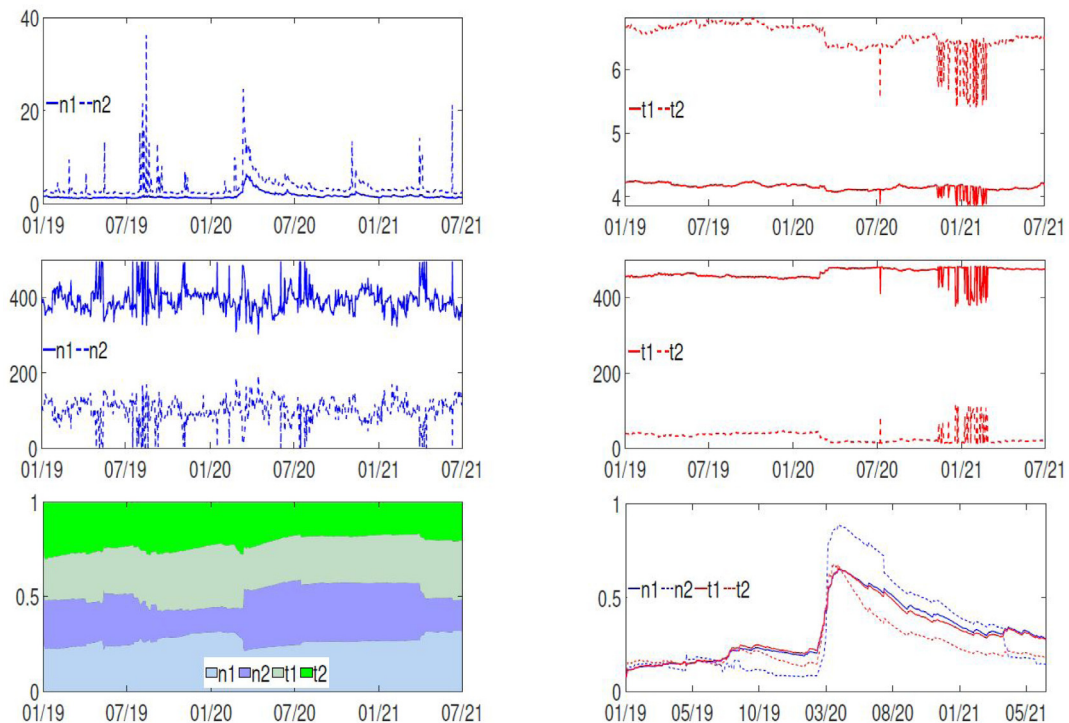


Fig. 3. Top left: the average variance of the predictions from the $n1$ (solid blue line) and $n2$ (dashed blue line) clusters. Top right: the average degree of freedom of the predictions from the $t1$ (solid red line) and $t2$ (dashed red line) clusters. The degrees of freedom are bounded to 30. Middle left: cluster allocation over time between the two clusters for the Normal case: $n1$ (solid blue line) and $n2$ (dashed blue line). Middle right: cluster allocation over time between the two clusters for the Student's t case: $t1$ (solid red line) and $t2$ (dashed red line). Bottom left: the mean logistic-normal weights for the $n1$, $n2$, $t1$ and $t2$ clusters. Bottom right: posterior mean estimates of the incompleteness measures in the four clusters.

where y_{it} is the log return of stock i at day t , $\xi_{it} \sim \mathcal{N}(0, 1)$ and $\xi_{it} \sim \mathcal{T}(v_i)$ for the Normal and Student- t cases, respectively. The number of degrees of freedom v_i is estimated in the latter model. We produce 629 one day ahead predictive densities from January 2, 2019 to June 30, 2021, see red line in Fig. 2 for the sample split. Our out-of-sample period is associated with relatively low volatility in 2019 and high volatility from the end of February 2020 driven by the Covid-19 crisis. In the initial and most dramatic part of the pandemic, the lowest daily return is almost -80%.

Results are reported in Fig. 3. It is seen that the Normal models in cluster $n2$ have a predictive variance (top left plot, dashed line) more than double in size than cluster $n1$ with several spikes over time and this increases further at the beginning of the Covid-19 pandemic. Cluster $n1$ has a relatively constant predicted variance (top left plot, solid line) over the entire period except at the beginning of the Covid-19 pandemic. The Student- t models in cluster $t1$ have a relatively constant thick tail just above 4 over the entire period (top right plot, solid line) while cluster $t2$ has values around 7 for the degrees of freedom (top right plot, dashed line) before the beginning of the Covid-19 which decreases to 6.5 after it. Some instability is also measured in the second wave of the Covid-19 at the end of 2020 and beginning of 2021.

Time-varying cluster composition and weights. The clustering of the predictive densities is repeated at every time t and therefore the cluster composition and weights vary over time. The middle and bottom panels in Fig. 3 present results about these features. The number of stocks varies across consecutive vintages of predictions with the clusters $n1$ and $t1$ being dominant in terms of attracting many individual stocks. The cluster $n1$ contains on average 400 stocks, whereas cluster $n2$ has the remaining 100. The Covid-19 pandemic creates some instability in the stock allocation, but similar patterns existed also in 2019. The cluster $t1$ includes on average 450 stocks, whereas cluster $t2$ has on average 50 stocks. The beginning of the Covid-19 pandemic is associated with a 5% increment in the allocation of stocks to cluster $t1$; the second wave at the end of 2020 and early months of 2021 produces some changes in the stock allocation.

Plots of the estimated cluster weights are shown at the bottom left panel in Fig. 3. These weights are the average weights of w_{it} per cluster. In terms of the importance of the different clusters, we notice that the clusters $n1$ and $t1$ receive large part of the weight over the full sample, see the bottom panel in Fig. 3. At the end of the period, they sum almost to 70% of the total weight. Therefore, the two clusters include most of the stocks and give the larger contribution to the combination. However, there is also evidence of time variations in the weights. The weights of clusters $n2$ and $t2$ are larger at the beginning of sample, but their size reduces from 50% to 30% just before the Covid-19 pandemic. The

Table 4

Predictive results for next day S&P500 log-returns. Root mean square prediction error (RMSPE), logarithmic score (LS) and the continuous rank probability score (CRPS) are reported. Bold numbers indicate the best statistic for each loss function. One or two asterisks indicate that differences in accuracy from the white noise (WN) benchmark are credibly different from zero at 5%, and 1%, respectively, using Bayes estimates of the Diebold–Mariano t -statistic for equal loss. The underlying p -values are based on t -statistics computed with a serial correlation-robust variance, using the pre-whitened quadratic spectral estimator of Andrews and Monahan (1992). The alternative models considered are: GARCH model with Normal error for the aggregate index (Normal GARCH); GARCH model with Student- t error for the aggregate index (Student- t GARCH); Glosten–Jagannathan–Runkle GARCH model for the aggregate index (GJR GARCH, see Glosten et al., 1993); Equal Weights of all disaggregate models (EW); Bayesian Model Averaging of the models of all disaggregate models (BMA); Density Combination with Equal Weights within the clusters and Stochastic Volatility (DCEW-SV). The column “Violation” shows the number of times the realised value exceeds the 5% Value-at-Risk (VaR) predicted by the different models over the sample.

Models	RMSPE	LS	CRPS	avQS-T	avQS-L	Violation
WN	1.518	−2.129	0.689	0.085	0.114	7.15%
Normal GARCH	1.513	−1.532**	0.638**	0.072**	0.104**	5.73%
Student- t GARCH	1.525	−1.420**	0.649**	0.074**	0.106**	3.50%
GJR GARCH	1.512	−1.517**	0.639**	0.072**	0.105**	5.56%
EW	1.522	−14.303	0.804	0.119	0.130	32.11%
BMA	1.525	−21.095	0.822	0.116	0.126	32.47%
DCEW-SV	1.509*	−1.372**	0.557**	0.065**	0.090**	4.97%

crisis increases substantially their weights, in particular for n_2 when volatility increases. The weights of clusters n_2 and t_2 reduce after the first part of the Covid-19 period and in the final part of the sample their contribution is similar to the pre-Covid-19 period.

Measures of incompleteness. We measure incompleteness for the model set Density Combination with Equal Weights and Stochastic Volatility, (DCEW-SV) at the bottom right panel in Fig. 3. We plot the average estimate of the four clusters incompleteness. The estimates are similar over time and they show a 500% increase in February–March 2020, which is due to the Covid-19 pandemic. In particular, the incompleteness for cluster n_2 has the larger increase. Interestingly, the volatilities start to reduce quite fast and they do over the remaining part of the sample. The values for clusters n_1 and t_1 in June are still twice the value pre-Covid-19; whereas the values for clusters n_2 and t_2 are closer to pre-Covid-19 values.

Predictive accuracy of centre and shape of the distribution. We compare the performance of our combination approach with results from five different basic models applied to the S&P500 log-returns: a white noise model in mean, often used as the main benchmark in equity premium predictability; the Normal GARCH(1,1) and the Student- t GARCH(1,1) models described above and applied to the aggregate S&P500 log-returns.

To explore the sensitivity of our results for model set incompleteness in more detail, we include the Normal GJR GARCH(1,1) model in Glosten et al. (1993) that includes leverage effects in the model set. This model is a richer model than the standard GARCH and should fit the data better. The leverage effect is considered among the stylised facts of financial returns and the added feature may become relevant in our analysis. Finally, given that simple combination methods might handle uncertainty accurately, we apply an equal weight combination of the all disaggregated GARCH models, labelled EW; and Bayesian model averaging, labelled BMA. In Section C in the Supplementary Material we also run AR models, AR GARCH models, different EW and BMA combinations of the models in the different clusters, see Table C.1. None of these provides very accurate predictions and are therefore excluded from the main text.

Out-of-sample predictive result are presented in Table 4. The first three columns deal with location and shape features of the predictive densities. It is seen that our combination scheme produces the lowest Root Mean Squared Prediction Error (RMSPE) and Cumulative Rank Probability Score (CRPS) and the highest Log Score (LS). The results indicate that the combination scheme is statistically superior to the no-predictability WN benchmark and it offers the most accurate statistics. The Normal GARCH(1,1) model, the Student- t GARCH(1,1) model and the Normal GJR GARCH(1,1) model fitted on the index also provide more accurate density predictions than the WN in terms of density prediction, but not on point prediction while our DCEW-SV is the only model that is statically superior at 5% level. For all three score criteria, the statistics given by BMA and EW are inferior to our combination scheme. In particular the density performance is very inferior. The lack of time-varying learning weights appears responsible for the poor performance in our data set.

Tail estimates and their economic value. We consider two statistics that refer to tails of the predictive densities. These statistics are the weighted averages of the (Gneiting and Raftery, 2007) quantile scores that are based on quantile predictions that correspond to the predictive densities from the different models. In the Supplementary Material, it is shown that avQS-T emphasises both tails and avQS-L the left tail of the predictive density relative to the realisation 1-step ahead. The fourth and fifth columns of Table 4 show results for tail evaluation. Our scheme provides the lowest avQS-T and avQS-L statistics, confirming the accuracy of our method in the tails of the distribution.

As an economic measure, we apply a Value-at-Risk (VaR) based measure, see Jorion (2006). We compare the accuracy of our models in terms of violations, that is the number of times that negative returns exceed the VaR predictions at time t , with the implication that actual losses on a portfolio are worse than was predicted. Higher accuracy results in

numbers of violations close to the nominal value of 5%.⁴ When looking at VaR violations, reported in the final column of Table 4, the number for all individual models is not very accurate, with the WN higher than 7%, the normal GARCH almost at 6%, the other two GARCH models at 3.5% indicating too large and a conservative density. Our DCEW-SV is the only one having a realised violation close to the 5% nominal value. The dramatic events in particular at the beginning of the Covid-19 pandemic in February/March 2020 drive the results. The property of our combination scheme to increase volatility in both normal clusters, and moreover, allocating more stock series to the fat tail cluster and a larger weight to the high volatility normal one, helps to model more accurately the lower tail of the index returns and covers more adequately risk.

5. Conclusions

We proposed in this paper a flexible Bayesian modelling approach with the construction of a predictive density combination with model set incompleteness and combination weight learning that can deal with large data sets in finance. The approach makes use of dimension reduction by dynamic clustering of the large number of components of the predictive mixture in mutually exclusive small subsets. Using a nonlinear dynamic factor model reduces the dimension of the large number of combination weights to a small set of cluster weights where a learning step is added. A parallel Sequential Monte Carlo algorithm is introduced for efficient Bayesian inference.

We applied the methodology to a large financial data set of individual stock returns which includes the Covid-19 crisis period. Empirical results show that our approach yields substantial accuracy gains in predictive means, volatilities and tail events compared to predictions from individual models and combination methods as Bayesian Model Averaging (BMA) and Equal Weights (EW). Measures of model set incompleteness and dynamic patterns in the cluster weights give valuable signals for improved financial modelling and policy analysis. These empirical results may provide useful information for investment fund management.

The line of research presented in this paper can be extended in several directions. For example, the cluster-based weights contain relevant signals about the importance of the predictive performance and composition of each of the clusters. Some clusters have a substantial weight while others have only little weight and such a pattern may vary over long time periods. This may lead to the construction of alternative model combinations for more accurate out-of-sample prediction and improved policy analysis. Another suggestion is to make use of judgemental information from individual forecasters, see [McAlinn and West \(2019\)](#), and combine this with the predictive information based on our modelling approach. Finally, we emphasise a fruitful connection and possible extension of our approach with work in the field of dynamic portfolio allocation, see [Baştürk et al. \(2019\)](#) for a basic analysis of a momentum strategy with a relatively small set of financial assets, and for a connection with work on machine learning methods for stock and bond market predictions, see [Gu et al. \(2020\)](#), [Bianchi et al. \(2020\)](#).

Appendix A. Supplementary data

Supplementary material related to this article can be found online at <https://doi.org/10.1016/j.jeconom.2022.11.004>.

References

- Aastveit, K.A., Cross, J.L., van Dijk, H.K., 2022. Quantifying time-varying forecast uncertainty and risk for the real price of oil. *J. Bus. Econom. Statist.* 1–15.
- Aastveit, K.A., Mitchell, J., Ravazzolo, F., van Dijk, H.K., 2019. The evolution of forecast density combinations in economics. In: *Oxford Research Encyclopedia of Economics and Finance*. Oxford University Press.
- Aitchinson, J., 1982. The statistical analysis of compositional data. *J. R. Stat. Soc. Series B* 44, 139–177.
- Aitchinson, J., Shen, S.M., 1980. Logistic-normal distributions: Some properties and uses. *Biometrika* 67, 261–272.
- Amisano, G., Geweke, J., 2010. Comparing and evaluating Bayesian predictive distributions of asset returns. *Int. J. Forecast.* 26, 216–230.
- Andersen, T., Bollerslev, T., Hadi, A., 2010. ARCH and GARCH models. In: *Encyclopedia of Statistical Sciences*. John Wiley & Sons, Ltd, pp. 1–11.
- Anderson, T.W., 1984. *An Introduction to Multivariate Statistical Analysis*. John Wiley & Sons, New York.
- Andrews, D.W.K., Monahan, J.C., 1992. An improved heteroskedasticity and autocorrelation consistent covariance matrix estimator. *Econometrica* 60, 953–966.
- Bañbura, M., Giannone, D., Reichlin, L., 2010. Large Bayesian vector auto regressions. *J. Appl. Econometrics* 25, 71–92.
- Baştürk, N., Borowska, A., Grassi, S., Hoogerheide, L., van Dijk, H.K., 2019. Forecast density combinations of dynamic models and data driven portfolio strategies. *J. Econometrics* 210, 170–186.
- Baştürk, N., Grassi, S., Hoogerheide, L., Van Dijk, H.K., 2016. Parallelization experience with four canonical econometric models using ParMitISEM. *Econometrics* 4, 1–11.
- Bianchi, D., Büchner, M., Tamoni, A., 2020. Bond risk premiums with machine learning. *Rev. Financ. Stud.* 34, 1046–1089.
- Bianchi, D., McAlinn, K., 2018. Divide and conquer: Financial ratios and industry returns predictability. Working papers, SSRN.
- Billio, M., Casarin, R., Ravazzolo, F., van Dijk, H.K., 2013. Time-varying combinations of predictive densities using nonlinear filtering. *J. Econometrics* 177, 213–232.
- Billio, M., Casarin, R., Rossini, L., 2019. Bayesian nonparametric sparse VAR models. *J. Econometrics* 212, 97–115.

⁴ We choose a 5% nominal value and not the standard 1% due to a large number of negative returns and the high variability at the beginning of the Covid-19 pandemic. Restricting to 1% will limit the analysis.

- Carvalho, C.M., Chang, J., Lucas, J.E., Nevins, J.R., Wang, Q., West, M., 2008. High-dimensional sparse factor modeling: Applications in gene expression genomics. *J. Amer. Statist. Assoc.* 103, 1438–1456.
- Casarin, R., Grassi, S., Ravazzolo, F., van Dijk, H.K., 2020. A Bayesian dynamic compositional model for large density combinations in finance. Working paper series 20-27, Rimini Centre for Economic Analysis.
- Casarin, R., Veggente, V., 2020. Random projection methods in economics and finance. In: Petr, H., Uddin, M., Abedin, M.Z. (Eds.), *the Essentials of Machine Learning in Finance and Accounting*. Routledge Taylor & Francis, pp. 1–20.
- Choi, H., Varian, H., 2012. Predicting the present with google trends. *Econ. Rec.* 88, 2–9.
- Clark, P.K., 1973. A subordinated stochastic process model with finite variance for speculative prices. *Econometrica* 41, 135–155.
- Corielli, F., Marcellino, M., 2006. Factor based index tracking. *J. Bank. Financ.* 30, 2215–2233.
- Creal, D., 2007. Sequential Monte Carlo samplers for Bayesian DSGE models. Working papers, Vrije Universiteit.
- Creal, D., 2012. A survey of sequential Monte Carlo methods for economics and finance. *Econometric Rev.* 31, 245–296.
- Doucet, A., de Freitas, N., Gordon, N., 2001. *Sequential Monte Carlo Methods in Practice*. Springer Verlag, New York, USA.
- Einav, L., Levin, J., 2014. Economics in the age of big data. *Science* 346, 715–718.
- Favirar, R., Rebolledo, D., Chan, E., Campbell, R., 2008. A parallel implementation of K-means clustering on GPUs. In: *Proceedings of 2008 International Conference on Parallel and Distributed Processing Techniques and Applications*. Vol. 2, pp. 14–17.
- Frühwirth-Schnatter, S., 2006. Finite Mixture and Markov Switching Models: Modeling and Applications to Random Processes. In: *Springer Series in Statistics*, Springer.
- Geweke, J., 2010. Complete and Incomplete Econometric Models. In: *The Econometric and Tinbergen Institutes Lectures*, Princeton University Press.
- Glosten, L.R., Jagannathan, R., Runkle, D.E., 1993. On the relation between the expected value and the volatility of the nominal excess return on stocks. *J. Finance* 48, 1779–1801.
- Gneiting, T., Raftery, A.E., 2007. Strictly proper scoring rules, prediction, and estimation. *J. Amer. Statist. Assoc.* 102, 359–378.
- Gneiting, T., Ranjan, R., 2013. Combining predictive distributions. *Electron. J. Stat.* 7, 1747–1782.
- Granger, C.W.J., 1998. Extracting information from mega-panels and high-frequency data. *Stat. Neerl.* 52, 258–272.
- Gu, S., Kelly, B., Xiu, D., 2020. Empirical asset pricing via machine learning. *Rev. Financ. Stud.* 33, 2223–2273.
- Hall, S.G., Mitchell, J., 2007. Combining density forecasts. *Int. J. Forecast.* 23, 1–13.
- Herbst, E., Schorfheide, F., 2014. Sequential Monte Carlo sampling for DSGE models. *J. Appl. Econometrics* 29, 1073–1098.
- Hoogerheide, L., Kleijn, R., Ravazzolo, R., van Dijk, H.K., Verbeek, M., 2010. Forecast accuracy and economic gains from Bayesian model averaging using time varying weights. *J. Forecast.* 29, 251–269.
- Hoogerheide, L., Opschoor, A., Van Dijk, H.K., 2012. A class of adaptive importance sampling weighted EM algorithms for efficient and robust posterior and predictive simulation. *J. Econometrics* 171, 101–120.
- Johannes, M., Korteweg, A., Polson, N., 2014. Sequential learning, predictability, and optimal portfolio returns. *J. Finance* 69, 611–644.
- Jorion, P., 2006. *Value at Risk: The New Benchmark for Managing Financial Risk*. McGraw-Hill, New York.
- Kaufmann, S., Schumacher, C., 2017. Identifying relevant and irrelevant variables in sparse factor models. *J. Appl. Econometrics* 32, 1123–1144.
- Kaufmann, S., Schumacher, C., 2019. Bayesian estimation of sparse dynamic factor models with order-independent and ex-post mode identification. *J. Econometrics* 210, 116–134.
- Kim, S., Kim, S., 2020. Index tracking through deep latent representation learning. *Quant. Finance* 20, 639–652.
- Koop, G., Korobilis, D., 2013. Large time-varying parameter VARs. *J. Econometrics* 177, 185–198.
- Koopman, S.J., Lucas, A., Scharth, M., 2015. Numerically accelerated importance sampling for nonlinear non-Gaussian state space models. *J. Bus. Econom. Statist.* 33, 114–127.
- Kunsch, H.R., 2005. Recursive Monte Carlo filters: Algorithms and theoretical analysis. *Ann. Statist.* 33, 1983–2021.
- Liesenfeld, R., Richard, J.F., 2003. Univariate and multivariate stochastic volatility models: Estimation and diagnostics. *J. Empir. Financ.* 10, 505–531.
- Liu, J.S., 2001. *Monte Carlo Strategies in Scientific Computing*. Springer Verlag, New York, USA.
- Lopes, H.F., West, M., 2004. Bayesian model assessment in factor analysis. *Statist. Sinica* 1, 41–68.
- MacLehose, R., Dunson, D., 2010. Bayesian semiparametric multiple shrinkage. *Biometrics* 66, 455–462.
- Malsiner Walli, G., Frühwirth-Schnatter, S., Grün, B., 2016. Model-based clustering based on sparse finite Gaussian mixtures. *Stat. Comput.* 26, 303–326.
- McAlinn, K., Aastveit, K.A., Nakajima, J., West, M., 2020. Multivariate Bayesian predictive synthesis in macroeconomic forecasting. *J. Amer. Statist. Assoc.* 115, 1092–1110.
- McAlinn, K., West, M., 2019. Dynamic Bayesian predictive synthesis in time series forecasting. *J. Econometrics* 210, 155–169.
- Mitchell, J., Hall, S.G., 2005. Evaluating, comparing and combining density forecasts using the klic with an application to the bank of England and NIESER “fan” charts of inflation. *Oxf. Bull. Econ. Stat.* 67, 995–1033.
- Pagan, A., 1984. Econometric issues in the analysis of regressions with generated regressors. *Internat. Econom. Rev.* 25, 221–247.
- Richard, J., Zhang, W., 2007. Efficient high-dimensional importance sampling. *J. Econometrics* 141, 1385–1411.
- Shephard, N., Andersen, T.G., 2009. Stochastic volatility: Origins and overview. In: Mikosch, T., Kreiß, J.-P., Davis, R.A., Andersen, T.G. (Eds.), *Handbook of Financial Time Series*. Springer Berlin Heidelberg, Berlin, Heidelberg, pp. 233–254.
- Stock, J.H., Watson, W.M., 1999. Forecasting inflation. *J. Monetary Econ.* 44, 293–335.
- Stock, J.H., Watson, W.M., 2002. Forecasting using principal components from a large number of predictors. *J. Amer. Statist. Assoc.* 97, 1167–1179.
- Stock, J.H., Watson, W.M., 2005. Implications of dynamic factor models for VAR analysis. Technical report, NBER Working Paper No. 11467.
- Stock, J.H., Watson, W.M., 2014. Estimating turning points using large data sets. *J. Econ.* 178, 368–381.
- Takanashi, K., McAlinn, K., 2019. Predictive properties of forecast combination, ensemble methods, and Bayesian predictive synthesis. Working papers, arXiv.
- Taylor, S., 1982. Financial returns modelled by the product of two stochastic processes - a study of daily sugar prices 1961-79. In: Anderson, O.D. (Ed.), *Time Series Analysis: Theory and Practice*. North-Holland, Amsterdam, pp. 203–226.
- Terui, N., van Dijk, H.K., 2002. Combined forecasts from linear and nonlinear time series models. *Int. J. Forecast.* 18, 421–438.
- Varian, H., 2014. Machine learning: New tricks for econometrics. *J. Econ. Perspect.* 28, 3–28.
- Varian, H., Scott, S., 2014. Predicting the present with Bayesian structural time series. *Int. J. Math. Model. Numer. Optimisation* 5, 4–23.
- Yao, Y., Vehtari, A., Simpson, D., Gelman, A., 2018. Using stacking to average Bayesian predictive distributions (with discussion). *Bayesian Anal.* 13, 917–1007.