

Improving the robustness of Markov-switching dynamic factor models with time-varying volatility

Romain Aumond*, Julien Royer†

March 2024

Abstract

Tracking macroeconomic data at a high frequency is difficult as most time series are only available at a low frequency. Recently, the development of macroeconomic nowcasters to infer the current position of the economic cycle has attracted the attention of both academics and practitioners, with most of the central banks having developed statistical tools to track their economic situation. The specifications usually rely on a Markov-switching dynamic factor model with mixed-frequency data whose states allow for the identification of recession and expansion periods. However, such models are notoriously not robust to the occurrence of extreme shocks such as Covid-19. In this paper, we show how the addition of time-varying volatilities in the dynamics of the model alleviates the effect of extreme observations and renders the dating of recessions more robust. Both stochastic and conditional volatility models are considered and we adapt recent Bayesian estimation techniques to infer the competing models parameters. We illustrate the good behavior of our estimation procedure as well as the robustness of our proposed model to various misspecifications through simulations. Additionally, in a real data exercise, it is shown how, both insample and in an out-of-sample exercise, the inclusion of a dynamic volatility component is beneficial for the identification of phases of the US economy.

Keywords: Nowcasting; Bayesian Inference; Dynamic Factor Models; Markov Switching

*CREST, ENSAE and Institut Polytechnique de Paris, 5 Avenue Henri Le Chatelier, 91120 Palaiseau, France; E-mail: romain.aumond@ensae.fr

†CREST and Institut Polytechnique de Paris, 5 Avenue Henri Le Chatelier, 91120 Palaiseau, France; E-mail: julien.royer@ensae.fr

We thank Christian Francq, Anna Simoni, and Jean-Michel Zakoïan for their insightful comments, as well as participants at the 17th International Conference on Computational and Financial Econometrics (CFE 2023).

1 Introduction

Dating economic recession and expansion periods is paramount for policy makers and asset managers. However, recessions are often identified after the publication of lagged low-frequency macroeconomic variables, resulting in an identification process that may be severely delayed. For example, the end of the Great Recession in June 2009 was announced by the National Bureau of Economic Research (NBER) more than a year later. In order to track the state of the economy in real-time and at a higher frequency, building upon nowcasting techniques popularized by [Evans \(2005\)](#), [Giannone et al. \(2008\)](#) and [Bańbura et al. \(2013\)](#), the use of Markov-switching dynamic factor models (MS-DFM) has recently gained popularity ([Camacho et al. \(2014, 2018\)](#); [Doz et al. \(2020\)](#)). Such models were initially introduced by [Diebold and Rudebusch \(1996\)](#) to capture the co-movements of multiple time series while allowing for the dynamic to be specific to different regimes. [Chauvet and Piger \(2008\)](#) and [Hamilton \(2011\)](#) emphasized the benefits of this specification to time economic recessions in the US.

The performance of these models has however been tremendously challenged by the occurrence of extreme values observed during the Covid-19 pandemic. [Figure 1](#) provides a compelling and motivating example. Using the five variables recommended by the NBER and pursuant to [Chauvet and Piger \(2008\)](#), we fit a standard MS-DFM model to obtain the in-sample probability of the US economy being in a recession. Using a sample ranging from February 1947 to December 2019¹, we observe that the model is very accurate to date the recessions. This accuracy however sharply decreases after the introduction of Covid data in the sample. Indeed, the same model fitted on a sample from February 1947 to June 2023 fails to identify five recessions, only capturing very large shocks (the oil shock of 1973, the interest rates shock of 1980 and the Great Financial Crisis of 2008).

The treatment of Covid data when modeling macroeconomic phenomena has puzzled econometricians as well as practitioners. While simply 'dumming out' data linked to the pandemic shock may seem appealing, these extreme values are not void of economic content as noted by [Ng \(2021\)](#). Focusing on linear VAR models, the author introduces Covid indicators to act as control variables and disentangle the pandemic and economic effects in the data. [Carriero et al. \(2022\)](#) propose a different approach, and specify a VAR with time-varying volatility to filter the extreme values through the volatility process. The benefits of including dynamic volatility components in VAR models are well documented in the literature (see for example [Clark \(2011\)](#); [Clark and Ravazzolo \(2015\)](#); [Chan and Eisenstat \(2018\)](#)) and have influenced models with latent variables,

¹Details about the data are presented in Section 4.




Figure 1: *Smoothed probability of being in a recession (Chauvet and Piger (2008)). The black line represents a sample from February 1947 to December 2019 while the red dashed line represents a sample from February 1947 to June 2023. Blue shades show the recessions as dated by the NBER.*

Antolin-Diaz et al. (2017) introducing a linear DFM with stochastic volatility, as well as Markov-switching models (see for example Eo and Kim (2016)). To the best of our knowledge, the development of MS-DFM with time-varying volatility has been limited and mostly focused around dynamic but piecewise-constant volatility processes as presented in Doz et al. (2020).

In this paper, we introduce new MS-DFM specifications to allow for continuous dynamic volatility processes. In particular, we will consider two well-known competing models: conditional volatility, where the dynamic variance is a measurable function of the past observations, and stochastic volatility processes, where the volatility is a random variable not directly linked to the σ -field of the data. We will show that, both in-sample and out-of-sample, models with dynamic volatilities outperform standard MS-DFM model in which homoskedasticity is assumed, in particular when large shocks occur. The remainder of the paper is organized as follows. Section 2 presents the general form of the model and discusses its Bayesian inference. Section 3 presents Monte Carlo experiments, illustrating the good-behavior of our estimation procedure. Additionally, a careful attention is dedicated to show the ability of our proposed model to remain robust even under misspecification and jumps in the simulated data generating process. Section 4 presents an application on real data, emphasizing the out-performance of our model for timing US recessions when Covid-data are considered in-sample. Section 5 provides an out-of-sample exercise where we compare the ability to date recessions in a real-time nowcasting exercise. Section 6 concludes. Technical details about the Bayesian

estimation are relegated to the appendix.

2 Markov-Switching Dynamic Factor Models with continuous time-varying volatility

The modelling of potentially large systems of economic time series via a small number of latent factors has been a workhorse of the economic literature. Initially introduced by [Diebold and Rudebusch \(1996\)](#), MS-DFM aim at capturing both the co-movements of multiple macroeconomic data - possibly sampled at different frequencies - and the changes in time series dynamics induced by latent regimes usually linked to the business cycle, in the spirit of [Hamilton \(1989\)](#). [Chauvet \(1998\)](#) and [Kim and Nelson \(1998\)](#) were the first to propose estimation procedures, the former considering a frequentist approach and the latter a Bayesian approach. Recently, the inclusion of time-varying volatility has proved useful in modeling linear DFM ([Antolin-Diaz et al. \(2017\)](#)) as well as Markov-Switching models where no factor structure is assumed ([Eo and Kim \(2016\)](#)). The aim of our specification is to bridge the gap between the latter and the MS-DFM literature.

2.1 Model specification

Let \mathbf{y}_t a vector of q quarterly and m monthly observable time series and let \mathbf{f}_t a set of k latent common factors. We have the standard DFM given by

$$\mathbf{y}_t = \mathbf{\Lambda} \mathbf{f}_t + \mathbf{u}_t, \quad (1)$$

where $\mathbf{\Lambda}$ denotes the loadings matrix, \mathbf{u}_t is orthogonal to \mathbf{f}_t and for all $j = 1, \dots, m + q$

$$\psi_j(L)u_{j,t} = e_{j,t}, \quad e_{j,t} \sim \mathcal{N}(0, \sigma_{e,j}^2), \quad (2)$$

and the set of factors follows

$$\mathbf{f}_t = \boldsymbol{\mu}_{S_t} + \mathbf{\Phi} \mathbf{f}_{t-1} + \boldsymbol{\Sigma}_t^{1/2} \boldsymbol{\eta}_t \quad (3)$$

where $\boldsymbol{\eta}_t$ is iid $(0, \mathbf{I}_k)$ and S_t is an independent first order Markov chain.

In most applications, a two-state ($S_t = 0$ or 1) regime-switching model is considered with a unique factor ($k = 1$) capturing the co-movements of the economic time series. For the sake of simplicity, in the remainder of the paper we will focus on this simplified

form. We thus have

$$\begin{aligned} \mathbf{f}_t &= f_t = \mu_{S_t} + \phi f_{t-1} + \varepsilon_t \\ \varepsilon_t &= \sigma_t \eta_t \end{aligned} \quad (4)$$

with $\mu_{S_t} = \mu_1 S_t + \mu_0(1 - S_t)$.

Among the most famous specifications, [Kim and Nelson \(1998\)](#) assume a constant volatility for the factor residual ($\sigma_t = \sigma$). The inclusion of a dynamic volatility component was first proposed by [Chauvet \(1998\)](#) by letting the constant volatility switch with the latent regimes, yielding $\sigma_t = \sigma_{S_t}$. More recently, [Doz et al. \(2020\)](#) extend this model by letting the volatility process be influenced by an additional two-state first-order Markov chain V_t , independent from S_t , yielding $\sigma_t = \sigma_{V_t}$. It is noteworthy that, while time-varying, the proposed volatility processes are piecewise-constant which might be difficult to justify empirically. Alternatively, we consider models where the volatility is time-varying but not piecewise constant. In particular, we introduce well-known competitors: conditional volatility models and stochastic volatility (SV) models. In the former, the volatility is a measurable function with regard to the σ -field \mathcal{F}_t generated by $\{f_u, t < u\}$. The simplest form of this model was introduced by [Engle \(1982\)](#) through the ARCH(1) equation

$$\sigma_t^2 = \omega + \alpha \varepsilon_{t-1}^2 \quad (5)$$

which was extended by [Bollerslev \(1986\)](#) to yield the famous GARCH(1,1) model

$$\sigma_t^2 = \omega + \alpha \varepsilon_{t-1}^2 + \beta \sigma_{t-1}^2 \quad (6)$$

where ω , α and β are positive parameters to estimate. On the contrary, stochastic volatility models do not assume that volatility is a function of the past but a random variable (in particular $\sigma_t^2 \neq \mathbb{E}[\varepsilon_t^2 | \mathcal{F}_t]$). The standard SV model is given by

$$\sigma_t^2 = e^{h_t}, \quad h_t = \mu_h + \phi_h(h_{t-1} - \mu_h) + \varepsilon_{h,t}, \quad \varepsilon_{h,t} \sim \mathcal{N}(0, \omega_h^2). \quad (7)$$

where the log-volatility h_t follows a stationary AR(1) process with μ_h , $|\phi_h| < 1$ and ω_h being parameters to estimate.

Tracking real-time economic conditions require to integrate time series measured at different frequencies, such as quarterly GDP and monthly employment data. Following [Mariano and Murasawa \(2003\)](#), the model is therefore specified at a monthly frequency where the observed quarterly data $\mathbf{y}_t^{(q)}$ can be related to unobserved synthetic monthly data $\mathbf{y}_t^{(m)}$

$$\mathbf{y}_t^{(q)} = \frac{1}{3} \mathbf{y}_t^{(m)} + \frac{2}{3} \mathbf{y}_{t-1}^{(m)} + \mathbf{y}_{t-2}^{(m)} + \frac{2}{3} \mathbf{y}_{t-3}^{(m)} + \frac{1}{3} \mathbf{y}_{t-4}^{(m)}. \quad (8)$$

Substituting quarterly series in (1) with the synthetic higher-frequency data in (8) allows us to obtain a dynamic factor model at a monthly frequency where missing data can be inferred from our Bayesian estimation procedure.

2.2 Bayesian estimation

It is useful to rewrite the MS-DFM model defined by equations (1)-(2)-(4) into a state-space equation. For any $j = 1, \dots, m + q$, let us denote $\boldsymbol{\psi}_j = (\psi_{j,1}, \dots, \psi_{j,l})'$ the coefficients of the lag polynomial $\psi_j(L)$ assumed of order l and $\boldsymbol{\Psi} = (\boldsymbol{\psi}'_1, \dots, \boldsymbol{\psi}'_m)'$. Additionally, let \mathbf{H} the $(m + q) \times (ml + 5 + 5q)$ matrix such that

$$\mathbf{H} = \begin{bmatrix} \lambda_1 \mathbf{h}_q & \mathbf{h}_q & \dots & 0 & 0 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ \lambda_q \mathbf{h}_q & 0 & \dots & \mathbf{h}_q & 0 & \dots & 0 \\ \lambda_{q+1} \mathbf{h}_m^5 & 0 & \dots & 0 & \mathbf{h}_m^l & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ \lambda_{q+m} \mathbf{h}_m^5 & 0 & \dots & 0 & 0 & \dots & \mathbf{h}_m^l \end{bmatrix}$$

where $\mathbf{h}_q = \left[\frac{1}{3} \quad \frac{2}{3} \quad 1 \quad \frac{2}{3} \quad \frac{1}{3} \right]$, $\mathbf{h}_m^5 = \left[1 \quad 0 \quad 0 \quad 0 \quad 0 \right]$, $\mathbf{h}_m^l = \left[1 \quad 0 \quad \dots \quad 0 \right]$. \mathbf{h}_m^l is a $1 \times l$ vector with the only first element equal to one. We define the lag vectors $\mathbf{L}^4 = (1, L, \dots, L^4)$. The vector of unobserved variables \mathbf{z}_t is given by

$$\mathbf{z}_t = (\mathbf{L}^4 f_t, \mathbf{L}^4 u_{1,t}, \dots, \mathbf{L}^4 u_{q,t}, (1, L, \dots, L^{l-1}) u_{q+1,t}, \dots, (1, L, \dots, L^{l-1}) u_{m+q,t})'$$

We can rewrite the factor model into a state-space equation as follows

$$\begin{aligned} \mathbf{y}_t &= \mathbf{H} \mathbf{z}_t + \boldsymbol{\varsigma}_t & \boldsymbol{\varsigma}_t &\sim \mathcal{N}(0, \mathbf{R}) \\ \mathbf{z}_t &= \boldsymbol{\delta}_{S_t} + \boldsymbol{\Xi} \mathbf{z}_{t-1} + \boldsymbol{\zeta}_t & \boldsymbol{\zeta}_t &\sim \mathcal{N}(0, \mathbf{Q}_t) \end{aligned} \quad (9)$$

with $\boldsymbol{\delta}_{S_t} = (\mu_{S_t}, 0, \dots, 0)'$, $\text{diag}(\mathbf{Q}_t) = (\sigma_t^2 \mathbf{h}_m^5, \sigma_{e,1}^2 \mathbf{h}_m^5, \dots, \sigma_{e,q}^2 \mathbf{h}_m^5, \sigma_{e,q+1}^2 \mathbf{h}_m^l, \dots, \sigma_{e,m+q}^2 \mathbf{h}_m^l)$ where $\boldsymbol{\Xi}$ is a $ml + 5 + 5q$ square block diagonal matrix given by

$$\boldsymbol{\Xi} = \begin{bmatrix} \boldsymbol{\phi} & & & & \\ & \boldsymbol{\Xi}_1 & & & \\ & & \ddots & & \\ & & & \ddots & \\ & & & & \boldsymbol{\Xi}_{m+q} \end{bmatrix} \quad \text{where} \quad \boldsymbol{\phi} = \begin{bmatrix} \phi & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \end{bmatrix}$$

such that, for $j_q = 1, \dots, q$ and $j_m = q + 1, \dots, m + q$, Ξ_{j_q} is a 5×5 matrix and Ξ_{j_m} a $l \times l$ matrix given by

$$\Xi_{j_q} = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \end{bmatrix} \quad \text{and} \quad \Xi_{j_m} = \begin{bmatrix} \psi_{j_m,1} & \psi_{j_m,2} & \dots & \psi_{j_m,l-1} & \psi_{j_m,l} \\ 1 & 0 & \dots & 0 & 0 \\ 0 & 1 & \dots & 0 & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \dots & 1 & 0 \end{bmatrix}.$$

Additionally, we follow [Leiva-Leon et al. \(2020\)](#) and restrict the individual components of the quarterly observations to be white noises.

Denoting the transition probabilities

$$q = \mathbb{P}(S_t = 0 | S_{t-1} = 0) \quad \text{and} \quad p = \mathbb{P}(S_t = 1 | S_{t-1} = 1),$$

the vector of parameters to estimate is given by

$$\boldsymbol{\vartheta} = (p, q, \boldsymbol{\Psi}', \sigma_{e,1}, \dots, \sigma_{e,m}, \boldsymbol{\Lambda}', \mu_0, \mu_1, \phi, \boldsymbol{\theta}^{(\cdot)'})'$$

where $\boldsymbol{\theta}^{(\cdot)}$ is the vector of parameters driving the dynamic volatility equation σ_t

$$\boldsymbol{\theta}^{(\text{ARCH})} = (\omega, \alpha), \quad \boldsymbol{\theta}^{(\text{GARCH})} = (\omega, \alpha, \beta), \quad \text{and} \quad \boldsymbol{\theta}^{(\text{SV})} = (\mu_h, \phi_h, \omega_h).$$

Let us denote $\mathbf{z}^{(T)} = \{\mathbf{z}_1, \dots, \mathbf{z}_T\}$ the unobserved state vector in equation (9), $\mathbf{y}^{(T)} = \{\mathbf{y}_1, \dots, \mathbf{y}_T\}$ the observed data, and $S^{(T)} = \{S_1, \dots, S_T\}$ the unobserved Markov Chain. The model is estimated using a Markov Chain Monte Carlo (MCMC) Gibbs sampling algorithm in the spirit of [Kim and Nelson \(1999\)](#) and [Bai and Wang \(2015\)](#) where conditional draws of the state vector, the Markov Chain, and the parameters vector $\boldsymbol{\vartheta}$ are obtained sequentially. In particular, we adapt the Metropolis Hastings procedure presented in [Chan and Grant \(2016\)](#) and [Chan \(2023\)](#) to sample stochastic and conditional volatilities. Details on the priors and a complete description of our sampling algorithm are presented in Appendices A and B but the main steps can be summarized as follows:

1. We generate conditional draws of the state vector from $p(\mathbf{z}^{(T)} | \mathbf{y}^{(T)}, S^{(T)}, \boldsymbol{\vartheta})$ using the forward-filtering backward-smoothing algorithm of [Carter and Kohn \(1994\)](#).
2. We generate conditional draws of the Markov chain from $p(S^{(T)} | \mathbf{y}^{(T)}, \mathbf{z}^{(T)}, \boldsymbol{\vartheta})$ based on the Hamilton filter ([Hamilton \(1989\)](#)).
3. We generate conditional draws for the parameters vector from $p(\boldsymbol{\vartheta} | \mathbf{y}^{(T)}, \mathbf{z}^{(T)}, S^{(T)})$

by sequentially drawing in the conditional distribution of components of ϑ as follows:

- We obtain conditional draws for p and q following [Albert and Chib \(1993\)](#).
- We obtain conditional draws for $(\Psi', \sigma_{e,1}, \dots, \sigma_{e,m}, \Lambda')$ from usual priors in the literature (see for example [Bai and Wang \(2015\)](#)).
- To obtain conditional draws for $(\mu_0, \mu_1, \phi, \theta^{(\cdot)'})$, we build upon the Metropolis Hasting algorithm presented in [Chan and Grant \(2016\)](#). In particular, in the case of stochastic volatility, we use the precision sampler of [Chan and Jeliazkov \(2009\)](#) for efficiency gains.

3 Monte Carlo experiments

In order to assess the performance of our estimation procedure as well as the robustness of our model to potential misspecifications, we conduct some Monte Carlo experiments.

In all cases, we model a five-variable system, with one quarterly data and four monthly data, driven by a single latent regime-switching factor. On all simulations, we use our MCMC procedure to fit five competing MS-DFM models as described by (1)-(2)-(4): a standard model with constant volatility (simply denoted MS-DFM), a model where the factor volatility follows an ARCH(1) dynamic (denoted MS-DFM-ARCH), a model where the factor volatility follows a GARCH(1,1) dynamic (denoted MS-DFM-GARCH), a model where the factor volatility is stochastic with dynamic (7) (denoted MS-DFM-SV), and a model where the factor volatility is piece-wise constant and driven by an additional Markov- chain independent of S_t as presented by [Doz et al. \(2020\)](#) (denoted 2MS-DFM). All estimation results are obtained from 1600 draws of the Gibbs sampler.

To conduct our experiments, we first simulate a two-state first order Markov chain with $p = q = 0.97$ and $T = 1000$. From the obtained Markov chain, we simulate f_t under different specifications (that we will develop in the reminder of this section), which yields the four monthly variables and the quarterly variable from (8). In particular, we use $\beta_1 = \dots = \beta_5 = 0.1$ and an AR(1) structure for the $u_{j,t}$ sequence with $\psi_j = 0.7$ and $\sigma_{e,j} = 1$ for all j .

We first simulate the factor assuming the true data generating process (DGP) follows the standard MS-DFM with constant volatility of [Kim and Nelson \(1998\)](#). In particular, we set $\mu_0 = -2$, $\mu_1 = 2$, $\phi = 0,7$ and $\varepsilon_t \sim \mathcal{N}(0,1)$. [Figure 2](#) presents the insample probability of being in regime 1 when fitting a standard MS-DFM (dotted red

line) and a MS-DFM-GARCH (plain black line). Shaded areas correspond to the simulated regimes sequence. In this case, the standard MS-DFM model is apt at identifying regime switches and performs very well insample. Interestingly, although misspecified, the MS-DFM-GARCH also appears to be well behaved even if the true DGP has constant volatility.

We then repeat the same experience, with the exact same setting, but proceed to shock the innovation sequence by forcing ε_{460} to take an extreme value ($\varepsilon_{460} = -35$). This shock will act as a jump, impacting the factor conditional mean through the AR(1) process. Figure 3 presents the insample probabilities when a shock occurs at time $t = 460$. The occurrence of an extreme value clearly derails the standard MS-DFM, failing at identifying most of the regime switches insample and reminding us of the deterioration of the performance of this model to date NBER recessions after the Covid-19 pandemic as emphasized by Figure 1. The insample probabilities derived from a MS-DFM-GARCH contrast sharply. The model is able to identify most regime switches even in the presence of an extreme shock, emphasizing the gain in robustness brought by the inclusion of a time-varying volatility process.

To rank the performance of the competing models, we consider two standard metrics based on the errors between the simulated state sequence $S_{0,t}$ (assumed observed) and the filtered probability of being in regime 1 at time t ($S_t = 1$). More precisely, we consider the Quadratic Probability Score (QPS) given by

$$\frac{1}{T} \sum_{t=1}^T (S_{0,t} - p(S_t = 1 | \mathcal{F}_{t-1}))^2 \quad (10)$$

and the False Probability Score (FPS) given by

$$\frac{1}{T} \sum_{t=1}^T (S_{0,t} - 1_{p(S_t=1|\mathcal{F}_{t-1})>0.5})^2. \quad (11)$$

Table 1 presents FPS and QPS metrics for our five competing models in the two presented cases (no jump and one jump), as well as an additional case where two shocks of lower intensities occur ($\varepsilon_{460} = -30$ and $\varepsilon_{800} = -20$). Although misspecified, the four competing time-varying volatility models appear well suited both without shocks and in the presence of extreme values.

We then simulate the factor assuming the true (DGP) follows the standard MS-DFM-GARCH with dynamic volatility following (6). In particular, we let $\mu_0 = -2$, $\mu_1 = 2$,




Figure 2: *Simulated recession regimes and smoothed recession probabilities from the MS-DFM model (in red) and MS-DFM-GARCH model (in black) when the DGP is a MS-DFM model*




Figure 3: *Simulated recession regimes and smoothed recession probabilities from the MS-DFM model (in red) and MS-DFM-GARCH model (in black) when the DGP is a MS-DFM model but a large jump occurs at $t = 460$*

Models	No jump		One jump		Two jumps	
	FPS	QPS	FPS	QPS	FPS	QPS
MS-DFM	0.071	0.053	0.362	0.252	0.354	0.329
MS-DFM-ARCH	0.090	0.072	0.072	0.053	0.073	0.053
MS-DFM-GARCH	0.114	0.093	0.079	0.059	0.072	0.054
MS-DFM-SV	0.079	0.055	0.072	0.054	0.084	0.059
2MS-DFM	0.075	0.055	0.081	0.055	0.082	0.060

Table 1: *Simulated regime dating under jumps and misspecifications when the true DGP is a standard MS-DFM*

$\phi = 0.7$ and set $\omega = 1$, $\alpha = \beta = 0.4$ with $\eta_t \sim \mathcal{N}(0, 1)$. Figure 4 presents the insample probability of being in regime 1 when fitting a standard MS-DFM (dotted red line) and a MS-DFM-GARCH (plain black line). Surprisingly, spuriously assuming constancy of the volatility process does not appear to penalize the inference of the regimes as the MS-DFM proves able to accurately detect regime switches. This relative irrelevance of heteroskedasticity on the Bayesian estimation of such MS-DFM models may explain the only recent attention to MS-DFM with dynamic volatility. The behavior of the MS-DFM-GARCH on this exercise illustrate the good performance of our Bayesian estimation procedure detailed in Appendices A and B.

We again repeat the experience and shock the innovation sequence with the same amplitude as the previous simulation. This extreme value will, this time, impact both the factor conditional mean through the AR(1) process and the conditional volatility through the GARCH process. Figure 5 presents the insample probabilities when a shock occurs at time $t = 460$. The occurrence of an extreme value once again strongly deteriorates the performance of the standard MS-DFM, while the insample probabilities derived from a MS-DFM-GARCH remain relatively unchanged.

Results in Table 2 confirm the previous findings and emphasize the robustness to extreme values stemming from the inclusion of time-varying volatility. Interestingly, in this case, although better than the standard MS-DFM, the MS-DFM-ARCH and the 2MS-DFM of Doz et al. (2020) appear less apt than the MS-DFM-GARCH to identify regime switches. This could be due to the slow decay of the volatility path after the shocks, induced by the GARCH(1,1) equation, that is incompatible with short memory feature of the ARCH(1) model and the piecewise-constant nature of the volatility in the 2MS-DFM. The MS-DFM-SV and MS-DFM-GARCH specifications appear difficult to discriminate in this simulation exercise.

4 A Covid-robust timing of US recessions

The main application of MS-DFM is the timing of recession and expansion periods underlying economic data. Tracking these recurring cycles is paramount for policy makers and asset managers, but Covid-data have tremendously complicated the identification process as emphasized by Figure 1. In a recent exercise, Doz et al. (2020) show that the inclusion of time-varying volatility improves the detection of recessions in the US. However, their sample stops before the occurrence of the Covid-19 pandemic. Simulations presented in the previous section yield promising results on the ability of MS-DFM models, when extended with continuous volatility processes, to be robust to extreme




Figure 4: *Simulated recession regimes and smoothed recession probabilities from the MS-DFM model (in red) and MS-DFM-GARCH model (in black) when the DGP is a MS-DFM-GARCH model*




Figure 5: *Simulated recession regimes and smoothed recession probabilities from the MS-DFM model (in red) and MS-DFM-GARCH model (in black) when the DGP is a MS-DFM-GARCH model but a large jump occurs at $t = 460$*

Models	No jump		One jump		Two jumps	
	FPS	QPS	FPS	QPS	FPS	QPS
MS-DFM	0.074	0.050	0.364	0.359	0.357	0.348
MS-DFM-ARCH	0.084	0.063	0.086	0.065	0.085	0.068
MS-DFM-GARCH	0.087	0.072	0.097	0.077	0.082	0.062
MS-DFM-SV	0.078	0.055	0.075	0.059	0.102	0.072
2MS-DFM	0.069	0.050	0.092	0.067	0.095	0.067

Table 2: *Simulated regime dating under jumps and misspecifications when the true DGP is a MS-DFM-GARCH*

values. In this section, we propose to confront this assumption to a real-data exercise.

Following the NBER Business Cycle Dating Committee recommendations, we consider a five-variable system, as presented in Table 3. The US quarterly GDP is obtained from the ALFRED database while the four monthly series are extracted from the FRED-MD database, both maintained by the Federal Reserve Bank of St Louis (see [McCracken and Ng \(2016\)](#)). All series are seasonally adjusted. The sample ranges from January 1947 to June 2023.

Data	Frequency	Start date	Transformation
Real GDP	Quarterly	Q1 1947	Diff Log
Industrial production	Monthly	Jan. 1947	Diff Log
Real personal income excluding transfer payments	Monthly	Jan. 1959	Diff Log
Real manufacturing trade and sales	Monthly	Jan. 1959	Diff Log
Non-agricultural civilian employment	Monthly	Jan. 1948	Diff Log

Table 3: *Data description*

Similarly to [Kim and Nelson \(1998\)](#), [Chauvet and Piger \(2008\)](#) and [Doz et al. \(2020\)](#), we assume a single-factor structure driven by a two-state Markov chain corresponding to recessions ($S_t = 1$) and expansion periods ($S_t = 0$). On our data, we fit five competing models as presented in the previous section. In addition, we also include constrained specifications where the autoregressive term in the latent factor dynamic is constrained to 0 ($\phi = 0$ in (4)). Figures 6, 7 and 8 present the smoothed recession probabilities for the unconstrained MS-DFM-GARCH, MS-DFM-SV and 2MS-DFM respectively (black lines). It is remarkable that the inclusion of time-varying volatility proves very effective for a Covid-robust timing of US recessions. However, the 2MS-DFM, clearly underperforms the MS-DFM-GARCH and MS-DFM-SV insample, emphasizing the need for a continuous volatility process. The MS-DFM-SV, although outperforming the standard MS-DFM, fails at identifying the shallow early-1970s recession.

To quantify the performance of the competing models, we evaluate the QPS and FPS by replacing $S_{0,t}$ in (10) and (11) with the recessions regimes as provided by the NBER. Results are reported in Table 4. In addition, we report the Portmanteau test statistics of [Li and Mak \(1994\)](#), an extended version of the standard goodness-of-fit test for conditional volatility models. Interestingly, the MS-DFM-GARCH is the only unconstrained specification to pass this test of at all considered lags. All other models reject the goodness-of-fit hypothesis, which emphasizes the heteroskedasticity of the latent factor, often neglected in the literature. Moreover, volatility exhibits some persistence, the




Figure 6: *Smoothed recession probability based on the MS-DFM-GARCH model*




Figure 7: *Smoothed recession probability based on the MS-DFM-SV model*




Figure 8: *Smoothed recession probability based on the 2MS-DFM model*

short-memory feature of volatility induced by the ARCH(1) equation yielding autocorrelated squared residuals. It is noteworthy that the constrained MS-DFM-GARCH with $\phi = 0$ both present the smallest FPS and QPS and does not reject the goodness-of-fit test.

Models	QPS	FPS	Portmanteau test		
			$q_{\max} = 3$	$q_{\max} = 5$	$q_{\max} = 10$
MS-DFM	0.045	0.045	129.2	129.2	130.8
MS-DFM ($\phi = 0$)	0.054	0.052	49.7	49.8	50.8
MS-DFM-ARCH	0.031	0.019	169.1	169.6	170.6
MS-DFM-ARCH ($\phi = 0$)	0.035	0.041	28.0	28.8	30.1
MS-DFM-GARCH	0.022	0.026	3.7	3.8	4.0
MS-DFM-GARCH ($\phi = 0$)	0.019	0.026	6.5	6.8	7.1
MS-DFM-SV	0.041	0.041	17.8	20.5	21.6
MS-DFM-SV ($\phi = 0$)	0.036	0.033	19.8	22.5	23.0
2MS-DFM	0.041	0.048	88.2	90.2	91.9
2MS-DFM ($\phi = 0$)	0.059	0.067	52.4	53.4	53.6
<i>threshold</i>			7.81	11.07	18.31

Table 4: *Empirical results on US recessions dating; Li and Mak (1994) Portmanteau test on factor residuals*

5 Real-time assessment: nowcasting US recessions

We proceed to a real-time downturn assessment of the five competing models. The monthly seasonally adjusted macroeconomic aggregates building up our information sample come from the FRED-MD database provided by the Federal Reserve Bank of St Louis [McCracken and Ng \(2016\)](#). We carry the exercise on vintages available from January 2001 onwards. In the FRED-MD database, for a given month M, industrial production and non farm payroll employment are displayed up until month M-1, real personal income excluding transfer payments up until month M-1 or M-2, and real manufacturing and trade sales up until month M-2 or M-3. We also use real seasonally adjusted quarterly GDP vintages from the ALFRED database. The first estimate of a quarterly GDP figure for a given quarter is usually available at the end of the first month following the reference period. It is then revised up to two times with the third estimate available at the end of the following quarter. We run the competing models on a monthly basis from January 2000 to June 2023 adding the new information available this given month (it means the only data available in March 2009 is employment and industrial production reports from February 2009). To fit best the day to day monitoring of macroeconomic conditions by practitioners we attribute the probability extracted from a given month data vintage to this month as the macroeconomic information availability is known to be lagging and asynchronous. The models filtered probabilities associated to downturns from the competing models are reported in [Figure 9](#).




Figure 9: *Real-time probabilities of being in a recession.* The black line represents the MS-DFM-GARCH probabilities, the red line displays the MS-DFM-SV probabilities. The blue line represents the Constant variance model probabilities. Blue shades show the recessions as dated by the NBER.

Given the fact that both amplitude and heterogeneity of downturn episodes have been increasing over the three recession episodes available in our vintages, a constant-volatility model seems to be performing well in capturing the entry date and the exit date of the downturns. This fact is depicted in the QPS and FPS identification measures reported in the first columns of Table 4. However, one could argue that this good performance of the standard MS-DFM model on this out-of-sample exercise is misleading, as the real-time datasets were not polluted by abnormal episodes. Indeed, Covid-data may have deteriorated the ability of the MD-DFM to identify softer low growth regimes. Because no recession has occurred in the recent post-Covid era, it is difficult to outperform standard MS-DFM in nowcasting the US economy from 2001 to 2023.

In order to assess the benefits of our model, we thus propose to simulate the real-time occurrence of a crisis. Focusing on a GFC-type crisis and a dotcom-type crisis, we investigate the ability of the competing models to capture, in a timely manner, the entry in these of recessions, should they repeat after the shock produced by the Covid crisis. Therefore, two datasets are built. We first extend the observation sample from July 2023 by adding an expanding information sample composed of the observations from January 2008 up until May 2009 — the *artificial* GFC. A second observation sample is composed of data observed up until July 2023 to which we add observations from March to November 2001 — the *artificial* dotcom bubble. Figures 10a and 10b focus on the real GFC crisis and the *artificial* GFC as well as the corresponding filtered probabilities of being in a recession regime in real-time. Figures 10c and 10d focus on the dotcom crisis.




Figure 10: *Real-time probabilities of being in a recession obtained from a MS-DFM-GARCH (black line), MS-DFM-SV (red dashed line) and standard MS-DFM (blue dotted line). Blue shades show the GFC recession as dated by the NBER.*

In both cases, extended MS-DFM with dynamic volatility do not seem to outperform the standard MS-DFM when the crisis occur prior to the Covid-19 shock. However, when the crisis repeat after the Covid-induced recession of 2020, the MS-DFM-GARCH is clearly outperforming its competitors at timing the artificially created recessions, as emphasized in the last columns of Table 5. This feature is a further evidence to the observations brought by [Camacho et al. \(2018\)](#) and [Leiva-Leon et al. \(2020\)](#) regarding the weakness of constant-variance model to identify in a timely fashion the occurrence of a downturn episode. The results defer from the above-mentioned papers to the extent that the exit of the recession episode as given by the real-time filtered probabilities intervenes around 6 months after the official dating given by the NBER.

Models	All sample		True GFC		True .com		Art. GFC		Art. .com	
	FPS	QPS	FPS	QPS	FPS	QPS	FPS	QPS	FPS	QPS
MS-DFM-ARCH	0.11	0.08	0.50	0.30	0.88	0.53	<i>0.47</i>	<i>0.31</i>	<i>1.00</i>	<i>0.78</i>
MS-DFM-GARCH	0.12	0.09	0.28	0.20	0.88	0.48	<i>0.18</i>	<i>0.09</i>	<i>0.75</i>	<i>0.49</i>
MS-DFM-SV	0.10	0.08	0.28	0.22	0.50	0.34	<i>0.41</i>	<i>0.26</i>	<i>1.00</i>	<i>0.79</i>
MS-DFM	0.10	0.07	0.56	0.31	0.50	0.26	<i>0.65</i>	<i>0.54</i>	<i>1.00</i>	<i>0.91</i>
2MS-DFM	0.10	0.08	0.44	0.23	0.50	0.31	<i>0.59</i>	<i>0.39</i>	<i>1.00</i>	<i>0.81</i>

Table 5: *Performance of the competing models on the whole out-of-sample exercise, the GFC and dotcom recessions, and their artificial counterparts. Bold numbers display the minimum statistics while italic indicate the artificially created samples.*

6 Conclusion

Coping with the Covid shock in macroeconomic aggregates is a new challenge for econometricians and practitioners. Given the increasing heterogeneity of downturn phases - and the amplitude of the last recessionary episode - this challenge is even more essential for dating business cycles. In this paper, we introduce a novel Markov-switching dynamic factor model that proves highly robust to extreme shocks. This model extends the existing literature by allowing the latent factor to have a continuous-path dynamic volatility process. We present a detailed MCMC Gibbs sampling algorithm and show its good performance on simulated data. Additionally, we establish the robustness of this framework in its capacity to properly identify regimes under misspecified data generating processes and artificial jumps through Monte Carlo experiments. We compare the performances of the new framework to text-book multi-frequency MS-DFM models on an insample turning point detection exercise. In particular, our model yields a precise, Covid-robust, dating of the NBER recessions. Finally, a real-time exercise is carried out, showing the ability of our model to consistently capture the entry date into recession, as well as its readiness to detect new downturns in the future. The addition of alternative data, supposed to be contemporaneous with the business cycle and observable at a higher frequency (such as market prices) could however improve the timing of recessions ends. We leave this problem for future research.

References

- J. H. Albert and S. Chib. Bayesian analysis of binary and polychotomous response data. *Journal of the American Statistical Association*, 88(422):669–679, 1993.
- J. Antolin-Diaz, T. Drechsel, and I. Petrella. Tracking the Slowdown in Long-Run GDP Growth. *The Review of Economics and Statistics*, 99(2):343–356, 2017.
- J. Bai and P. Wang. Identification and bayesian estimation of dynamic factor models. *Journal of Business & Economic Statistics*, 33(2):221–240, 2015.
- M. Bańbura, D. Giannone, M. Modugno, and L. Reichlin. Now-casting and the real-time data flow. In G. Elliott and A. Timmermann, editors, *Handbook of Economic Forecasting*, volume 2, pages 195–237. 2013.
- T. Bollerslev. Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics*, 31(3):307–327, 1986.
- M. Camacho, G. Perez-Quiros, and P. Poncela. Green shoots and double dips in the euro area: A real time measure. *International Journal of Forecasting*, 30(3):520–535, 2014.
- M. Camacho, G. Perez-Quiros, and P. Poncela. Markov-switching dynamic factor models in real time. *International Journal of Forecasting*, 34(4):598–611, 2018.
- A. Carriero, T. E. Clark, M. Marcellino, and E. Mertens. Addressing COVID-19 Outliers in BVARs with Stochastic Volatility. *The Review of Economics and Statistics*, pages 1–38, 06 2022.
- C. K. Carter and R. Kohn. On Gibbs sampling for state space models. *Biometrika*, 81(3):541–53, 1994.
- J. C. Chan. The stochastic volatility in mean model with time-varying parameters: An application to inflation modeling. *Journal of Business & Economic Statistics*, 35(1):17–28, 2017.
- J. C. Chan. Comparing stochastic volatility specifications for large Bayesian VARs. *Journal of Econometrics*, 235(2):1419–1446, 2023.
- J. C. Chan and E. Eisenstat. Bayesian model comparison for time-varying parameter VARs with stochastic volatility. *Journal of Applied Econometrics*, 33(4):509–532, 2018.
- J. C. Chan and A. L. Grant. Modeling energy price dynamics: GARCH versus stochastic volatility. *Energy Economics*, 54:182–189, 2016.

- J. C. Chan and I. Jeliazkov. Efficient simulation and integrated likelihood estimation in state space models. *International Journal of Mathematical Modelling and Numerical Optimisation*, 1(1-2):101–120, 2009.
- M. Chauvet. An econometric characterization of business cycle dynamics with factor structure and regime switches. *International Economic Review*, 39(4):969–996, 1998.
- M. Chauvet and J. Piger. A comparison of the real-time performance of business cycle dating methods. *Journal of Business & Economic Statistics*, 26(1):42–49, 2008.
- T. E. Clark. Real-time density forecasts from Bayesian Vector Autoregressions with stochastic volatility. *Journal of Business & Economic Statistics*, 29(3):327–341, 2011.
- T. E. Clark and F. Ravazzolo. Macroeconomic forecasting performance under alternative specifications of time-varying volatility. *Journal of Applied Econometrics*, 30(4):551–575, 2015.
- F. X. Diebold and G. D. Rudebusch. Measuring business cycles: A modern perspective. *The Review of Economics and Statistics*, 78(1):67–77, 1996.
- C. Doz, L. Ferrara, and P.-A. Pionnier. Business cycle dynamics after the Great Recession: An extended Markov-switching dynamic factor model. Working Paper 02443364, Paris School of Economics, 2020.
- R. F. Engle. Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation. *Econometrica*, 50(4):987–1007, 1982.
- Y. Eo and C.-J. Kim. Markov-switching models with evolving regime-specific parameters: are postwar booms or recessions all alike? *The Review of Economics and Statistics*, 98(5):940–949, 2016.
- M. Evans. Where are we now? Real-time estimates of the macroeconomy. *International Journal of Central Banking*, 2(1):127–175, 2005.
- D. Giannone, L. Reichlin, and D. Small. Nowcasting: The real-time informational content of macroeconomic data. *Journal of Monetary Economics*, 55(4):665–676, 2008.
- J. D. Hamilton. A new approach to the economic analysis of nonstationary time series and the business cycle. *Econometrica*, 57(2):357–384, 1989.
- J. D. Hamilton. Calling recessions in real time. *International Journal of Forecasting*, 27(4):1006–1026, 2011.

- C.-J. Kim and C. R. Nelson. Business cycle turning points, a new coincident index and tests of duration dependence based on a dynamic factor model with regime switching. *The Review of Economics and Statistics*, 80(2):188–201, 1998.
- C.-J. Kim and C. R. Nelson. *State space models with regime switching*. The MIT Press, Cambridge, Massachusetts, 1999.
- D. Leiva-Leon, G. Pérez-Quirós, and E. Rots. Real-time weakness of the global economy: A first assessment of the coronavirus crisis. *Working Paper Series, European Central Bank*, 2381, 2020.
- W. K. Li and T. Mak. On the squared residual autocorrelations in non-linear time series with conditional heteroskedasticity. *Journal of Time Series Analysis*, 15(6):627–636, 1994.
- R. S. Mariano and Y. Murasawa. A new coincident index of business cycles based on monthly and quarterly series. *Journal of Applied Econometrics*, 18(4):427–443, 2003.
- M. W. McCracken and S. Ng. FRED-MD: A monthly database for macroeconomic research. *Journal of Business & Economic Statistics*, 34(4):574–589, 2016.
- S. Ng. Modeling macroeconomic variations after COVID-19. Working Paper 29060, National Bureau of Economic Research, July 2021.

Appendix A Priors

This section describes the priors used for the distributions of the parameter vector $\boldsymbol{\vartheta}$. λ_1 is set to one for identification purposes. For all $j = 2, \dots, m + q$, we use the following prior to sample λ_j the j -th element of the factor loading matrix $\mathbf{\Lambda}$ in (1)

$$\lambda_j \sim \mathcal{N}(a_j, A_j) \quad (12)$$

where hyperparameters are set to $a_j = 0$ and $A_j = 0.1$. To sample the parameters linked to the residuals $u_{j,t}$ in (2), we use the following priors, for $l = 1, 2$,

$$\begin{aligned} \psi_{j,l} &\sim \mathcal{N}(\pi, \Pi) \quad \pi = 0, \Pi = 0.1 \\ \sigma_{e,j}^2 &\sim IG(\nu_i, Z_i) \quad \nu_i = 10, Z_i = 2 \end{aligned} \quad (13)$$

where IG denotes the inverse-gamma distribution. Additionally, independent beta distributions can be used as conjugate prior for each transition probability

$$\pi(q, p) \propto q^{u_{00}}(1 - q)^{u_{01}}p^{u_{11}}(1 - p)^{u_{10}} \quad (14)$$

Note that, contrary to [Doz et al. \(2020\)](#), we do not put an informative prior and set $u_{00} = 470, u_{01} = 9, u_{10} = 9, u_{11} = 90$ in order to take into account the relative persistence of each of the regimes as observed on macroeconomic data. The prior for the Markov-switching intercept in equation (4) is given by :

$$\boldsymbol{\mu} = (\mu_0, \mu_1)' \sim \mathcal{N}(\boldsymbol{\alpha}^*, A^*) \quad (15)$$

with $\boldsymbol{\alpha}^* = (4, -2)'$ and $A^* = \text{diag}(0.02, 0.02)$. We acknowledge that, in the spirit of [Leiva-Leon et al. \(2020\)](#), relatively tight priors are used for identification purposes. The informativeness brought by the first moment is indeed needed to discriminate between the regimes over the parameters space. The prior for the autoregressive parameter ϕ in equation (4) is given by

$$\phi \sim \mathcal{N}(\alpha, A) \quad (16)$$

where $\alpha = 0, A = 0.1$. In the case of a MS-DFM-ARCH, we use the following prior for the vector $\boldsymbol{\theta}^{(\text{ARCH})} = (\omega, \alpha)$

$$\log \boldsymbol{\theta}^{(\text{ARCH})} \sim \mathcal{N}(\boldsymbol{\theta}_0^{(\text{ARCH})}, V_\theta) \mathbf{1}(\alpha < 1).$$

$\boldsymbol{\theta}^{(\text{ARCH})}$ thus follows a truncated log-normal distribution with the stationarity restriction that $\alpha < 1$. We set the hyperparameters to $\boldsymbol{\theta}_0^{(\text{ARCH})} = \log(1, 0.5)$ and $V_\theta = \text{diag}(1, 1)$. In the case of MS-DFM-GARCH, we use the following prior for the vector $\boldsymbol{\theta}^{(\text{GARCH})} =$

(ω, α, β)

$$\log \boldsymbol{\theta}^{(\text{GARCH})} \sim \mathcal{N}(\boldsymbol{\theta}_0^{(\text{GARCH})}, V_\theta) \mathbb{1}(\alpha + \beta < 1).$$

Similarly to the ARCH(1) specification, $\boldsymbol{\theta}^{(\text{GARCH})}$ follows a truncated log-normal distribution with the adapted stationarity restriction $\alpha + \beta < 1$. Hyperparameters are set to $\boldsymbol{\theta}_0^{(\text{GARCH})} = \log(1, 0.5, 0.4)$ and $V_\theta = \text{diag}(1, 1, 1)$. Note that in both cases, the priors are relatively non-informative. Finally, in the MS-DFM-SV, we use the following prior for the vector $\boldsymbol{\theta}^{(\text{SV})} = (\mu_h, \phi_h, \omega_h)$

$$\mu_h \sim \mathcal{N}(\mu_{h0}, V_{\mu_h}) \quad \phi_h \sim \mathcal{N}(\phi_{h0}, V_{\phi_h}) \quad \omega_h \sim IG(\nu_h, S_h)$$

where $\mu_{h0} = 1$, $V_{\mu_h} = 50$, $\phi_{h0} = 0.9$, $V_{\phi_h} = 1$, $\nu_h = 1$, and $S_h = 1$. These priors are intended to make the stochastic volatility process exhibits persistence in a similar fashion as the conditional variance in the GARCH framework.

Appendix B Bayesian Estimation

Let $\mathbf{z}^{(T)} = \{\mathbf{z}_1, \dots, \mathbf{z}_T\}$ the unobserved state, $\mathbf{y}^{(T)} = \{\mathbf{y}_1, \dots, \mathbf{y}_T\}$ the observed data and $S^{(T)} = \{S_1, \dots, S_T\}$ the first order Markov-Chain. We describe the Gibbs sampler steps based on [Kim and Nelson \(1999\)](#) and follow their notations. The Gibbs sampler consists of iterating between the three following steps sequentially.

B.1 Generation of the state vector

The joint distribution of $\mathbf{z}^{(T)}$, given $\mathbf{y}^{(T)}$, $S^{(T)}$ and $\boldsymbol{\vartheta}$ can be defined as

$$p(\mathbf{z}^{(T)} | \mathbf{y}^{(T)}, S^{(T)}, \boldsymbol{\vartheta}) = p(\mathbf{z}_T | \mathbf{y}^{(T)}, S^{(T)}, \boldsymbol{\vartheta}) \prod_{t=1}^{T-1} p(\mathbf{z}_t | \mathbf{y}^{(t)}, S^{(t)}, \boldsymbol{\vartheta}, \mathbf{z}_{t+1})$$

which boils down to generating \mathbf{z}_t for $t = T, T-1, \dots, 1$ from

$$\begin{aligned} \mathbf{z}_T | \mathbf{y}^{(T)}, S^{(T)}, \boldsymbol{\vartheta} &\sim \mathcal{N}(\mathbf{z}_{T|T}, \mathbf{V}_{T|T}) \\ \mathbf{z}_t | \mathbf{y}^{(t)}, S^{(t)}, \mathbf{z}_{t+1}, \boldsymbol{\vartheta} &\sim \mathcal{N}(\mathbf{z}_{t|t, \mathbf{z}_{t+1}}, \mathbf{V}_{t|t, \mathbf{z}_{t+1}}) \end{aligned} \tag{17}$$

where $\mathbf{z}_{t|t} = E(\mathbf{z}_t | \mathbf{y}^{(t)})$ and $\mathbf{V}_{t|t} = \text{Var}(\mathbf{z}_t | \mathbf{y}^{(t)})$ for $t = 1, \dots, T$. In equation (17), $\mathbf{z}_T | \mathbf{y}^{(T)}, S^{(T)}, \boldsymbol{\vartheta}$ can be generated using the Multi-move Gibbs sampling introduced by [Carter and Kohn \(1994\)](#) as follows

1. We use the Kalman filter to obtain $\mathbf{z}_{t|t}$ and $\mathbf{V}_{t|t}$ for $t = 1, \dots, T$. The last iteration of the filter gives $\mathbf{z}_{T|T}$ and $\mathbf{V}_{T|T}$ which are then used to generate \mathbf{z}_T .
2. For $t = T-1, T-2, \dots, 1$, $\mathbf{z}_{t|t}$ and $\mathbf{V}_{t|t}$, \mathbf{z}_{t+1} can be considered as an incremental

vector of observations in the system. The distribution $p(\mathbf{z}_t | \mathbf{y}^{(T)}, S^{(t)}, \boldsymbol{\vartheta}, \mathbf{z}_{t+1})$ is then deduced from the Kalman smoother. From equation (9), updating equation are then given by

$$\begin{aligned}\mathbf{z}_{t|t, \mathbf{z}_{t+1}} &= \mathbf{z}_{t|t} + \mathbf{V}_{t|t} \boldsymbol{\Xi} \tilde{\boldsymbol{\zeta}}_t / R_t \\ \mathbf{V}_{t|t, \mathbf{z}_{t+1}} &= \mathbf{V}_{t|t} - \mathbf{V}_{t|t} \boldsymbol{\Xi}' \boldsymbol{\Xi} \mathbf{V}_{t|t}' / R_t\end{aligned}$$

where $\tilde{\boldsymbol{\zeta}}_t = \mathbf{z}_{t+1} - \boldsymbol{\delta}_{S_{t+1}} - \boldsymbol{\Xi} \mathbf{z}_{t|t}$ and $R_t = \boldsymbol{\Xi} \mathbf{V}_{t|t} \boldsymbol{\Xi}' + \sigma_{t+1}^2$.

B.2 Generation of the Markov Chain

Once $\mathbf{z}^{(T)}$ has been simulated, given $\boldsymbol{\vartheta}$, the Markov Chain $S^{(T)}$ can be generated from the following distribution

$$\begin{aligned}p(S^{(T)} | \mathbf{y}^{(T)}, \mathbf{z}^{(T)}, \boldsymbol{\vartheta}) &= p(S_T | \mathbf{y}^{(T)}, \mathbf{z}^{(T)}, \boldsymbol{\vartheta}) \prod_{t=1}^{T-1} p(S_t | \mathbf{y}^{(t)}, \mathbf{z}^{(t)}, S_{t+1}, \boldsymbol{\vartheta}) \\ &= p(S_T | \mathbf{z}^{(T)}, \boldsymbol{\vartheta}) \prod_{t=1}^{T-1} p(S_t | \mathbf{z}^{(t)}, S_{t+1}, \boldsymbol{\vartheta})\end{aligned}$$

as the distribution of $S^{(T)}$ is orthogonal to $\mathbf{y}^{(T)}$ given $\mathbf{z}^{(T)}$. We can thus obtain conditional draws for $S^{(T)}$ as follows

1. We use the [Hamilton \(1989\)](#) filter on (1) to generate $p(S_t | \mathbf{z}^{(t)}, \boldsymbol{\vartheta})$ for $t = 1, 2, \dots, T$ and save them. The last iteration gives $p(S_T | \mathbf{z}^{(T)}, \boldsymbol{\vartheta})$ from which we get S_T .
2. To draw S_t given $\mathbf{z}^{(T)}$ and S_{t+1} , for $t = T - 1, T - 2, \dots, 1$ the following result is used

$$p(S_t | \mathbf{z}^{(t)}, S_{t+1}, \boldsymbol{\vartheta}) = \frac{p(S_{t+1} | S_t) p(S_t | \mathbf{z}^{(t)}, \boldsymbol{\vartheta})}{p(S_{t+1} | \mathbf{z}^{(t)}, \boldsymbol{\vartheta})} \propto p(S_{t+1} | S_t) p(S_t | \mathbf{z}^{(t)}, \boldsymbol{\vartheta})$$

where $p(S_{t+1} | S_t)$ is the transition probability in $\boldsymbol{\vartheta}$ and $p(S_t | \mathbf{z}^{(t)}, \boldsymbol{\vartheta})$ is obtained from the values saved in the previous step.

3. The last step consist of drawing from

$$Pr(S_t = 1 | \mathbf{z}^{(t)}, S_{t+1}, \boldsymbol{\vartheta}) = \frac{p(S_{t+1} | S_t = 1) p(S_t = 1 | \mathbf{z}^{(t)}, \boldsymbol{\vartheta})}{\sum_{j=0}^1 p(S_{t+1} | S_t = j) p(S_t = j | \mathbf{z}^{(t)}, \boldsymbol{\vartheta})}$$

where S_t is drawn from a uniform distribution $S_t \sim \mathcal{U}(0, 1)$. If the generated number is smaller than $Pr(S_t = 1 | S_{t+1}, \mathbf{z}^{(t)}, \boldsymbol{\vartheta})$, $S_t = 1$, otherwise $S_t = 0$.

B.3 Generation of the parameters vector

We now turn to the generation of draws for the vector of parameters. To do so, we will sequentially draw components of the $\boldsymbol{\vartheta}$ vector as follows.

We obtain conditional draws for the transition probabilities p and q following [Albert and Chib \(1993\)](#). In particular, given $S^{(T)}$ and the initial state, we denote the transition from the state $S_{t-1} = i$ to $S_t = j$ by n_{ij} , the log-likelihood is given by

$$L(q, p) = q^{n_{00}}(1 - q)^{n_{01}}p^{n_{11}}(1 - p)^{n_{10}}.$$

By combining the likelihood function and the conjugate priors presented in the previous section, from equation (14), we get the conditional distributions of (p, q) as the product of the independent beta distributions from which we generate p and q as

$$\begin{aligned} q \mid S^{(T)} &\sim \text{Beta}(u_{00} + n_{00}, u_{01} + n_{01}) \\ p \mid S^{(T)} &\sim \text{Beta}(u_{11} + n_{11}, u_{10} + n_{10}). \end{aligned}$$

Given $\mathbf{y}^{(T)}$ and $f^{(T)}$, we can rewrite equation-by-equation equation (1) with

$$\mathbf{y}_{j,t}^* = \lambda_j f_{j,t}^* + e_{j,t}$$

for $j = 2, \dots, m + q$, where $y_{j,t}^*$ and $f_{j,t}^*$ are the j -th respective components of

$$\begin{aligned} \mathbf{y}_t^* &= \mathbf{y}_t - \bar{\psi}_1 \circ \mathbf{y}_{t-1} + \bar{\psi}_2 \circ \mathbf{y}_{t-2} \\ \mathbf{f}_t^* &= \mathbf{e}_m f_t - \bar{\psi}_1 f_{t-1} + \bar{\psi}_2 f_{t-2} \end{aligned} \quad (18)$$

with \mathbf{e}_m denoting a vector of 1 of length m and $\bar{\psi}_l = (\psi_{1,l}, \dots, \psi_{m,l})$, $l = 1, \dots, L$; $L = 2$ being the order of the AR specification in equation (2). From (12) and (18), we obtain conditional draws for λ_j from the posterior distribution

$$\mathcal{N} \left[\left(A_j^{-1} + \sigma_{e,j}^{-2} f_j^{*(T)'} f_j^{*(T)} \right)^{-1} \left(A_j a_j + \sigma_{e,j}^{-2} f_j^{*(T)'} \mathbf{y}_j^{*(T)} \right), \left(A_j^{-1} + \sigma_{e,j}^{-2} f_j^{*(T)'} f_j^{*(T)} \right)^{-1} \right].$$

Given $\mathbf{y}^{(T)}$ and $f^{(T)}$, from (1) we can measure $\mathbf{u}^{(T)}$ and from equation (2) and the prior distribution (13), for all $j = 1, \dots, m$, we can draw $\boldsymbol{\psi}_j$ from the posterior distribution

$$\mathcal{N} \left[\left(\boldsymbol{\Pi}_j^{-1} + \sigma_{e,j}^{-2} \mathbf{w}_j^{(T)'} \mathbf{w}_j^{(T)} \right)^{-1} \left(\boldsymbol{\Pi}_j^{-1} \boldsymbol{\pi}_j + \sigma_{e,j}^{-2} \mathbf{w}_j^{(T)'} \mathbf{u}_j^{(T)} \right), \left(\boldsymbol{\Pi}_j^{-1} + \sigma_{e,j}^{-2} \mathbf{w}_j^{(T)'} \mathbf{w}_j^{(T)} \right)^{-1} \right]$$

where $\mathbf{w}_{j,t} = (u_{j,t-1}, u_{j,t-2})'$. Similarly, from the generated $\boldsymbol{\psi}_j$ and from (13), we can

draw $\sigma_{e,j}$ from the posterior distribution

$$IG \left(\frac{\nu_j + T}{2}, \frac{Z_j}{2} + \frac{\left(u_j^{(T)} - \boldsymbol{\psi}'_j \mathbf{w}_j^{(T)} \right)' \left(u_j^{(T)} - \boldsymbol{\psi}'_j \mathbf{w}_j^{(T)} \right)}{2} \right).$$

Finally, we turn to the generation of $(\mu_0, \mu_1, \phi, \boldsymbol{\theta}^{(\cdot)'})$. As the Gibbs sampling methods differ between conditional and stochastic volatility, we will detail our algorithm for the different models, starting with the latter.

For the MS-DFM-SV, we draw the individual parameters in $(\mu_0, \mu_1, \phi, \boldsymbol{\theta}^{(SV)'})$ sequentially. Rewriting equation (4), we have

$$\frac{f_t - \phi f_{t-1}}{\sigma_t} = \frac{\mu_0(1 - S_t) + \mu_1 S_t}{\sigma_t} + \eta_t$$

Let us denote G_t^* the left-hand side of the above equation and Q_t^* the right-hand side. From the prior distribution (15), $\boldsymbol{\mu}$ can be drawn from the posterior distribution

$$\boldsymbol{\mu} \sim \mathcal{N}((A^{*-1} + Q^{*(T)' } Q^{*(T)})^{-1} (A^{*-1} \boldsymbol{\alpha} + Q^{*(T)' } G^{*(T)}), (A^{*-1} + Q^{*(T)' } Q^{*(T)})^{-1}),$$

and only draws verifying the condition $\mu_0 > \mu_1$ are kept. Rewriting again equation (4) yields

$$\frac{f_t - \mu_0(1 - S_t) - \mu_1 S_t}{\sigma_t} = \phi \frac{f_{t-1}}{\sigma_t} + \eta_t.$$

Denoting \tilde{G}_t the left-hand side of the above equation and \tilde{Q}_t the right-hand side, from (16) ϕ can be drawn from the following posterior distribution

$$\phi \sim \mathcal{N}((A^{-1} + \tilde{Q}' \tilde{Q})^{-1} (A^{-1} \boldsymbol{\alpha} + \tilde{Q}' \tilde{G}), (A^{-1} + \tilde{Q}' \tilde{Q})^{-1}).$$

Only draws satisfying the stationarity condition $|\phi| < 1$ are kept. We then jointly sample the log-volatility from the conditional density

$$p(\mathbf{h}^{(T)} \mid \mathbf{z}^{(T)}, S^{(T)}, \mu_0, \mu_1, \phi, \boldsymbol{\theta}^{(SV)'})$$

based on the acceptance-rejection Metropolis Hastings algorithm described in Chan (2017) using the precision sampler of Chan and Jeliazkov (2009). To that end, we compute the mode of $p(\mathbf{h}^{(T)} \mid \mathbf{z}^{(T)}, S^{(T)}, \mu_0, \mu_1, \phi, \boldsymbol{\theta}^{(SV)'})$ and the Hessian of the log-density evaluated at this mode denoted \hat{h} and K_h . We then use $\mathcal{N}(\hat{h}, K_h^{-1})$ as a proposal distribution in the acceptance-rejection Metropolis Hastings step from which we can directly sample $\mathbf{h}^{(T)}$.

For the MS-DFM-GARCH, we draw the parameters from the three full conditional distributions $p(\boldsymbol{\mu} \mid \mathbf{z}^{(T)}, \mathcal{S}^{(T)}, \phi, \boldsymbol{\theta}^{(\text{GARCH})})$, $p(\phi \mid \mathbf{z}^{(T)}, \mathcal{S}^{(T)}, \boldsymbol{\theta}^{(\text{GARCH})})$ and $p(\boldsymbol{\theta}^{(\text{GARCH})} \mid \mathbf{z}^{(T)}, \mathcal{S}^{(T)}, \phi)$ sequentially. Since μ_{S_t} and ϕ appear in the conditional variance equation, those distributions are non-standard, as noted by [Chan and Grant \(2016\)](#), and Metropolis Hastings algorithms are required. To sample ϕ we use a Gaussian proposal with mean $\bar{\phi}$ and variance V_ϕ given by

$$\begin{aligned}\bar{\phi} &= (A^{-1} + \tilde{Q}'\tilde{Q})^{-1}(A^{-1}\alpha + \tilde{Q}'\tilde{G}) \\ V_\phi &= (A^{-1} + \tilde{Q}'\tilde{Q})^{-1}.\end{aligned}$$

Only draws satisfying the stationarity condition $|\phi| < 1$ are kept. To sample $bm\mu$, we use a multivariate Gaussian proposal

$$\mathcal{N}\left[(A^{*-1} + Q^{*(T)'}Q^{*(T)})^{-1}(A^{*-1}\alpha + Q^{*(T)'}G^{*(T)}), (A^{*-1} + Q^{*(T)'}Q^{*(T)})^{-1}\right]$$

and only keep draws verifying $\mu_0 > \mu_1$. Finally to sample $\boldsymbol{\theta}^{(\text{GARCH})}$, we use a Gaussian proposal centered at the mode of $p(\boldsymbol{\theta}^{(\text{GARCH})} \mid \mathbf{z}^{(T)}, \mathcal{S}_t, \phi)$ with covariance matrix set to be the outer product of the scores.