Switching Macroeconomic Growth and Volatility: Evidence from a Mean-Variance Markov-Switching Dynamic Factor Model

Catherine Doz[1] Laurent Ferrara[2*] Pierre-Alain Pionnier[3]

July 15, 2025

Abstract

As illustrated by the Great Recession, the COVID-19 pandemic and the global decline in GDP growth since the mid-2000s, economists need to account for sudden and deep recessions, shifts in macroeconomic volatility, and longer-term fluctuations in GDP growth. This paper puts forward a Mean-Variance Markov-Switching Dynamic Factor Model (MV-MS-DFM) that accounts for these stylised facts by allowing the mean and the volatility of macroeconomic variables to switch abruptly, and trend GDP growth to vary smoothly over time. We show that allowing for different volatility regimes improves the detection of turning points in the U.S. business cycle, that the Great Recession and the COVID-19 pandemic only led to temporary increases in volatility, and that the U.S. trend GDP growth has declined by around 1 percentage point since the early 2000s. Information criteria and marginal likelihood comparisons support our model specification. The model provides a unified framework connecting the literature on turning-point detection to the more recent literature on Growth-at-Risk in macroeconomic forecasting. While tightening financial conditions are shown to increase the probability of falling in a recession, the model can generate left-skewed density forecasts without including any financial variable in the information set. The paper finally discusses how to adjust the model estimation strategy to deal with the COVID-19 period.

Keywords: Markov-Switching, Bayesian estimation, Turning-point detection, Business cycles, Real-time, Growth-at-Risk, COVID-19

JEL classification: C22, C51, E32, E37

This paper is an improved and updated version of the paper that was circulated under the title "Business Cycle Dynamics after the Great Recession: An Extended Markov-Switching Dynamic Factor Model". We would like to thank Marcelle Chauvet, Guillaume Chevillon, Arnaud Dufays, Domenico Giannone, Sylvia Kaufmann, Siem-Jan Koopman, Josefine Quast, Luca Sala, Paul Schreyer, Dave Turner, and participants in the following conferences: OxMetrics User Conference, French Econometrics Conference, Symposium of the Society for Nonlinear Dynamics and Econometrics (SNDE), Conference of the International Association for Applied Econometrics (IAAE), Conference on Forecasting at Central Banks, Conference on Real-Time Data Analysis, Methods and Applications, CIRET/KOF/OECD/INSEE Workshop, International Symposium on Forecasting, CFE-CMS Conference. We also thank participants in internal seminars at OECD, Namur University, Bocconi University, Nanterre University, Piraeus University and KOF for helpful comments and suggestions. All remaining errors are ours.

[1]Paris School of Econonomics and University Paris 1 Panthéon Sorbonne. Email: *catherine.doz@univ-paris1.fr*
[2]*Corresponding author. Skema Business School - University Côte d'Azur. Email: *laurent.ferrara@skema.edu*
[3]OECD. Email: *pierre-alain.pionnier@oecd.org*

1 Introduction

The sequence of major shocks that has affected the global economy since the early 2000s, including the Great Recession of 2008-09, the COVID-19 pandemic of 2020-21 and the energy crisis in 2022, contrasts with the two previous decades of limited macroeconomic volatility known as the Great Moderation (McConnell and Perez-Quiros, 2000). These shocks occurred against a background of declining long-term growth in advanced economies, often referred to as secular stagnation (Summers, 2005). Accounting for these stylised facts is increasingly important for tracking business cycles in real time (Antolin-Diaz et al., 2024).

Our starting point is a small-scale Markov-Switching Dynamic Factor Model (MS-DFM). This type of models is increasingly used for nowcasting, short-term forecasting, and turning point detection (Chauvet and Potter 2013, Camacho et al. 2014, 2018). The MS-DFM was put forward by Diebold and Rudebusch (1996) to simultaneously account for the co-movement in macroeconomic time series and the different dynamics during expansions and recessions, as originally suggested by Burns and Mitchell (1946). Our main contribution to the literature is to extend this MS-DFM to account for declining long-term growth and changing macroeconomic volatility. In the following, we refer to this new model as a Mean-Variance Markov-Switching Dynamic Factor Model (MV-MS-DFM). As argued by Antolin-Diaz et al. (2017), accounting for declining long-term growth improves GDP forecasts even at short horizons because standard DFM forecasts tend to revert quickly to the unconditional mean of GDP. We also provide evidence that the introduction of Markov-Switching volatility in a standard MS-DFM is supported by statistical criteria and improves the detection of turning points in the US economy after the mid-1980s.

As a by-product of the estimation of our model, we quantify the decline in U.S. long-run GDP growth to be about one percentage point per year compared to the early 2000s, half of it occurring before the Great Recession. Fernald et al. (2017) highlighted the role of total factor productivity and participation in the labour force to explain this decline. We also show that the Great Recession did not put an end to the Great Moderation. According to our model, the Great Recession was characterized by a temporary increase in macroeconomic volatility, but we observe a return to a low-volatility environment after 2010, in line with Charles et al. (2018) and Gadea-Rivas et al. (2018). This result is important for the calibration of macroeconomic models and for anticipating future U.S. recessions.

As the model is estimated using Bayesian techniques, it can generate density forecasts that include three sources of uncertainty: future shocks, parameter estimation, and the identification of the underlying business cycle and volatility regimes. While accounting for two business cycle regimes (expansions and recessions) is key to capture abrupt changes in GDP growth, accounting for the larger uncertainty in average GDP growth during recessions than during expansions and for two volatility regimes allows the model to generate left-skewed density forecasts, as in the Growth-at-Risk literature initiated by Adrian et al. (2019), and recenlty extended by Loria et al. (2024). Even though our model can generate left-skewed density forecasts without including financial variables in the information set, we show that tighter financial conditions increase the probability of falling in a recession.

Finally, we discuss how our model estimation strategy can be adjusted to account for the COVID-19 period. In line with other authors who suggested to remove this period from the estimation sample of macroeconomic models (Lenza and Primiceri, 2022; Baumeister and Hamilton, 2023), we show that our model can be used during the pandemic by freezing its parameters in 2019 and letting it identify a recession and a high-volatility regime based on the parameters estimated over the previous period. Adding post-COVID data to the estimation sample leaves the model parameters and the recession and volatility regimes identified over 1970-2019 hardly modified.

The paper is organized as follows. Section 2 provides a review of the related literature. Section 3 introduces the MV-MS-DFM. Section 4 describes the data and the estimation methodology. Section 5 provides in-sample estimation results based on U.S. data from 1970 to 2019. Section 6 evaluates the real-time out-of-sample properties of the model, with a focus on density forecasts and Growth-at-Risk features. Section 7 discusses the sensitivity of GDP growth to financial conditions and how our model is able to capture it. Section 8 reviews possible adjustments of the estimation strategy to deal with the COVID-19 period, and Section 9 concludes.

2 Related literature

Initially introduced by Diebold and Rudebusch (1996), MS-DFMs have two attractive features: they account for co-movement in macroeconomic time series and for different dynamics during expansion and recession phases. In this paper, we put forward an extended MS-DFM that also incorporates a time-varying long-run GDP growth rate and Markov-Switching volatility. Our model can be seen as a non-linear extension of the model proposed by Antolin-Diaz et al. (2017), who include a time-varying long-run GDP growth rate in a linear DFM model. Our main difference with them is the inclusion of Markov-Switching features to distinguish expansion and recession phases, as well as high- and low-volatility regimes.

Like Marcellino et al. (2016) and Antolin-Diaz et al. (2017) , we account for time variation in macroeconomic volatility, but whereas they rely on a stochastic volatility specification, we rely on a Markov-Switching volatility specification. The evidence provided by Stock and Watson (2002) and Sensier and van Dijk (2004) supports this specification, since they both conclude that the Great Moderation episode is better accounted for by a sharp break in volatility rather than a gradual decline.

Our model is also closely connected to the one proposed by Eo and Kim (2016) who relax the assumption that all recessions and expansions are alike in a univariate Markov-Switching model of the business cycle. However, whereas Eo and Kim (2016) allow for a fully flexible evolution of the regime-specific mean growth rates over time, we link regime-specific mean growth rates to macroeconomic volatility, as suggested by McConnell and Perez-Quiros (2000), Giordani et al. (2007) and Bai and Wang (2011). Our MV-MS-DFM is also connected to the model proposed by Chauvet and Su (2014), who include independent Markov-Switching processes for the mean and the variance of GDP growth, and add a structural break in GDP growth.

In this paper, we consider a larger information set than Eo and Kim (2016) or Chauvet and Su (2014) who only used quarterly GDP. In addition to quarterly GDP, we also use the four monthly variables advocated by the NBER Business Cycle Dating Committee. This multivariate framework also distinguishes our model from the univariate models of McConnell and Perez-Quiros (2000), Giordani et al. (2007) and Bai and Wang (2011). Note that our information set could be further expanded to include additional monthly or quarterly indicators, which makes our model suitable for nowcasting GDP growth.

Table 1 summarises how our modelling strategy relates to the existing literature.

Table 1: Comparison with the existing literature

	Multivariate	MS mean	MS volatility	SV	TV-Trend
Antolin-Diaz et al. (2017)	×			×	×
Bai and Wang (2011)		×	×		
Chauvet and Su (2014)		×	×		
Diebold and Rudebusch (1996)	×	×			
Eo and Kim (2016)		×		×	×
Giordani et al. (2007)		×	×		×
Marcellino et al. (2016)	×			×	
McConnell and Perez-Quiros (2000)		×	×		
This paper: MV-MS-DFM	×	×	×		×

Note: MS = Markov switching; SV = Stochastic Volatility; TV = Time-Varying

This paper finally relates to the Growth-at-Risk literature initiated by Adrian et al. (2019). While these authors relied on quantile regressions to capture downside risks to future GDP growth, other econometric strategies have been proposed since then. For example, Lhuissier (2022) relied on Markov-Switching skewnormal models, and Carriero et al. (2024) relied on Bayesian VARs with stochastic volatility in mean. Together with Caldara et al. (2021), our paper demonstrates that MS-DFMs can also generate left-skewed density forecasts and capture downside risks to future GDP growth. Nevertheless, it improves over Caldara et al. (2021) in several important dimensions. First, it does not impose that the same Markov process drives the business cycle regime (expansion/recession) and the volatility regime. This allows capturing recessions that are not associated with increased volatility, e.g. the 1990-91 and 2001 U.S. recessions which took place during the Great Moderation. Second, MS features in our model impact all macroeconomic variables simultaneously because they drive the dynamics of the underlying factor. In contrast, Caldara et al. (2021) first extract macroeconomic and financial factors using linear DFMs and then assume that the link between those factors and GDP growth depends on a Markov variable, leaving it unclear why GDP growth and other macroeconomic variables are not treated in a symmetric way.

While our model can generate left-skewed density forecasts even without resorting to financial variables, it offers two different ways to capture the information brought by financial conditions: as variables contributing to the estimation of the underlying economic activity factor, or as variables influencing transition probabilities between expansions and recessions. In line with Adrian et al. (2019) but using a different model, we show that tightening financial conditions make future GDP growth more fragile by increasing the probability of falling in a recession.

3 Econometric specification

This section describes the econometric specification of the MV-MS-DFM. We assume that both the intercept and the volatility of the shocks in the state equation (factor dynamics) are governed by two independent Markov-Switching processes. This specification can be thought of as a generalization of the work by McConnell and Perez-Quiros (2000) to a multivariate framework. While they used US quarterly GDP growth as their only observable variable, we rely on the four monthly series mentioned by the NBER Business Cycle Dating Committee (BCDC) in addition to quarterly GDP.

We first consider a general state-space model where, for the time being, all n observable variables y_{it}, including GDP, are available at a monthly frequency. The model is given by the following equations.

First, $\forall i \in 1, \ldots, n$, a measurement equation relates the observable variables to latent variables:

$$\Delta y_{it} = a_{it} + \gamma_i(L)c_t + u_{it} \tag{1}$$

Since all observable variables are demeaned prior to model estimation, there is no need to add constant terms in the measurement equations.

Then, the state equations specify the dynamics of the latent variables, as follows:

$$a_{it} = a_{i,t-1} + \sigma_{a_i} \cdot \eta_t^{a_i}, \tag{2}$$

$$\phi(L)c_t = \mu_{S_t,V_t} + \sqrt{1 + h \cdot V_t} \cdot \sigma_c \cdot \eta_t^c, \tag{3}$$

$$\psi_i(L)u_{it} = \sigma_i \cdot \varepsilon_{it}, \tag{4}$$

where $\eta_t^{a,i}$ iid $\mathrm{N}(0,1)$, η_t^c iid $\mathrm{N}(0,1)$ and ε_{it} iid $\mathrm{N}(0,1)$.

In the measurement equation (1), c_t captures common fluctuations across observed variables, and a_{it} captures low-frequency fluctuations in the growth rate of series i.

In the state equations (2)-(4), a_{it} are independent random walks. The common component c_t follows an autoregressive process (AR) with Markov-Switching intercept and Markov-Switching variance. The error terms u_{it} are modelled as independent AR processes.

S_t and V_t are two independent first-order Markov-Switching processes, each with only two possible states (0 or 1). Transition probabilities from one regime to the other are given by:

$$\begin{aligned} P(S_t = j | S_{t-1} = i) &= p_{ij}, \\ P(V_t = j | V_{t-1} = i) &= q_{ij}. \end{aligned}$$

S_t impacts the intercept of the state equation and is equal to 0 during recessions and to 1 during expansions. V_t is equal to 0 when volatility is low, and equal to 1 when volatility is high. As a result, the variance of $\phi(L)c_t$ is σ_c^2 in the low-volatility regime and $(1+h)\sigma_c^2$ in the high-volatility regime.

The intercept of state equation (3) depends on both S_t and V_t, as follows:

$$\mu_{S_t,V_t} = \mu_{00} + \mu_{01}V_t + \mu_{10}S_t + \mu_{11}S_t, \tag{5}$$

where:

- μ_{00} is the intercept in the low-volatility recession regime
- $\mu_{00} + \mu_{01}$ is the intercept in the high-volatility recession regime
- $\mu_{00} + \mu_{10}$ is the intercept in the low-volatility expansion regime
- $\mu_{00} + \mu_{01} + \mu_{10} + \mu_{11}$ is the intercept in the high-volatility expansion regime.

Thus μ_{01} is likely to be negative, μ_{10} is likely to be positive, and μ_{11} is also likely to be positive. Leaving V_t influence μ_{S_t,V_t}, we follow McConnell and Perez-Quiros (2000), Bai and Wang (2011), and Chauvet et al. (2015).

In practice, we use a mix of quarterly and monthly frequencies in our information set: GDP growth is measured at quarterly frequency, and the growth rate of all other series is measured at monthly frequency. The quarterly GDP growth Δy_{1t}^q is observed once every three months, while the growth rates of the monthly variables Δy_{jt}^m are observed every month. This is a situation that the Kalman filter can effectively accommodate. The standard approach to relate quarterly GDP to the underlying monthly state variables is to approximate the quarterly GDP growth rate as a weighted average of current and past monthly GDP growth rates, as proposed by Mariano and Murasawa (2003). Therefore, the measurement equation for quarterly GDP growth can be rewritten as follows.

$$\begin{aligned} \Delta y_{1t}^q =& \left(\frac{1}{3}a_{1,t}^q + \frac{2}{3}a_{1,t-1}^q + a_{1,t-2}^q + \frac{2}{3}a_{1,t-3}^q + \frac{1}{3}a_{1,t-4}^q\right) \\ &+ \gamma_1^q(L)\left(\frac{1}{3}c_t + \frac{2}{3}c_{t-1} + c_{t-2} + \frac{2}{3}c_{t-3} + \frac{1}{3}c_{t-4}\right) \\ &+ \left(\frac{1}{3}u_{1,t}^q + \frac{2}{3}u_{1,t-1}^q + u_{1,t-2}^q + \frac{2}{3}u_{1,t-3}^q + \frac{1}{3}u_{1,t-4}^q\right) \end{aligned} \tag{6}$$

The measurement equation of each monthly variable takes the following form:

$$\Delta y_{it}^m = a_{it}^m + \gamma_i^m c_t + u_{it}^m \tag{7}$$

We assume an AR(1) dynamics for the idiosyncratic shocks and the underlying factor, as in Chauvet and Hamilton (2005):

$$\Psi_1^q(L) = 1 - \psi_{11}^q L$$

$$\Psi_1^m(L) = 1 - \psi_{j1}^m L \text{ for all } j = 1 \dots n$$
$$\phi(L) = 1 - \phi_1 L$$

In practice, GDP growth is the only observed series for which we allow $a_{it} \neq 0$, and this should only be considered as a convenient modelling trick to capture the long-term decline in U.S. GDP growth. Only accounting for the low-frequency fluctuations of the main variable of interest (GDP growth) usefully limits the number of state variables in the model. Antolin-Diaz et al. (2017) follow a similar modelling strategy and conduct simulation experiments showing that not including a time-varying trend in the dynamics of the other variables is harmless for the estimation of the GDP trend, as long as persistence is allowed for in the idiosyncratic components of all variables. This is precisely the role of the $\Psi_i^m(L)$ polynomials in our model, for which we do not rule out the possibility to have unit roots. Alternatively, Stock and Watson (2012) subtract a local mean to U.S. GDP growth before including it in their model. In our case, this adjustment is made within the model.

We finally assume that quarterly GDP growth and the first three monthly variables (industrial production, real manufacturing and trade sales, and real personal income excluding transfer payments) are linked only contemporaneously to the business cycle (i.e. to the underlying factor). In other words, we assume $\gamma_1^q(L) = \gamma_{10}^q$ and $\gamma_j^m(L) = \gamma_{j0}^m$ for $j = 1, 2, 3$. For the fourth monthly variable (non-agricultural civilian employment), we include a richer lag structure to take into account that employment may be lagging, as in Stock and Watson (1989), that is: $\gamma_4^m(L) = \gamma_{40}^m + \gamma_{41}^m L + \gamma_{42}^m L^2 + \gamma_{43}^m L^3$. As in any factor model, we also have to impose an identifying restriction (otherwise, the factor would be defined up to a multiplicative constant). Here, we impose the following restriction: $\gamma_{10}^q = 1$.

Additional details on the state-space representation of the model can be found in Appendix A.

4 Data and estimation strategy

Our information set contains five variables: quarterly real GDP and four monthly variables (industrial production, real manufacturing and trade sales, real personal income excluding transfer payments, and non-agricultural civilian employment). These variables are advocated by the NBER Business Cycle Dating Committee. Following Chauvet and Hamilton (2005), we prefer to rely on civilian employment rather than payroll employment.[1] The U.S. quarterly, real and seasonally adjusted GDP is extracted from the ALFRED database maintained by the St.Louis FED. The corresponding quarterly vintages are available from December 1991 onwards. The four monthly series in the information set are extracted from the FRED-MD database, also maintained by the St.Louis FED (McCracken and Ng, 2016). All of these monthly series are seasonally adjusted. The corresponding monthly vintages are available from August 1999 onwards. In the FRED-MD database available for month M, industrial production and non-agricultural civilian employment are typically available up to month (M-1), real personal income excluding transfer payments up to month (M-1) or (M-2), and real manufacturing trade and sales up to month (M-2) or (M-3).

Following Kim and Nelson (1998), we estimate the model using a Bayesian approach and rely on Gibbs sampling. The Bayesian estimation strategy has several important advantages. First, it is a modular estimation technique, allowing to add or remove building blocks in the model and facilitating comparisons between different model specifications. Second, it simplifies inference on the Markov-Switching variables S_t and V_t by allowing to condition on the underlying factor and treating it as if it were an observed monthly variable. In this way, the inclusion of both quarterly and monthly variables in the information set does not complicate inference on S_t and V_t. This is a key advantage compared to Camacho et al. (2018) who directly rely on the

[1]Payroll employment is based on a survey of business establishments whereas civilian employment is based on a household survey. As noted by Chauvet and Hamilton (2005), payroll employment only includes job creation and destruction with a lag and does not include self-employment or off-the-book employment, which can delay the detection of a recovery. Moreover, it double counts jobs if a person changes job within a payroll survey reference period, which can overestimate employment around peaks. Lastly, payroll employment undergoes substantial revisions over time whereas civilian employment does not get revised.

variables in the information set to estimate the underlying Markov-Switching variables. When the variables in the information set have different frequencies, their distribution potentially depends on many lags of the Markov-Switching variables, thus generating a curse of dimensionality problem that they solve at the cost of approximations. None of these approximations is needed here. Finally, the Bayesian estimation strategy allows accounting for three sources of uncertainty at the same time (future shocks, parameter estimation, and the identification of the underlying business cycle and volatility regimes), which is key for producing meaningful density forecasts (Section 7).

Our Gibbs sampling algorithm consists of four main blocks, which are sequentially iterated until convergence. Subsequently, we use the following notation for all time series: $\widetilde{\alpha_{1...T}} = (\alpha_1, \ldots, \alpha_T)$.

Block 1: Draw the state vector $\widetilde{\alpha_{1...T}}$ conditional on $\widetilde{\Delta y^*_{1...T}}$, $\widetilde{S_{1...T}}$ and $\widetilde{V_{1...T}}$ and the model parameters, based on the sequential Kalman filter with diffuse initialization of Koopman and Durbin (2000, 2003) and the simulation smoother of Durbin and Koopman (2002). The sequential Kalman filter considers the series in the information set one by one when updating estimates of the state vector. Beyond saving some computing time, it simplifies diffuse initialization compared to the multivariate Kalman filter. Moreover, both the Kalman filter and the simulation smoother efficiently deal with missing observations, which is crucial in the present context where the variables in the information set are released at different dates (ragged-edge sample) and have different frequencies.

Block 2: Draw $\widetilde{S_{1...T}}$ conditional on $\widetilde{\alpha_{1...T}}$, $\widetilde{V_{1...T}}$ and model parameters, based on Hamilton's (1989) filter.

Block 3: Draw $\widetilde{V_{1...T}}$ conditional on $\widetilde{\alpha_{1...T}}$, $\widetilde{S_{1...T}}$ and model parameters, based on Hamilton's (1989) filter.

Block 4: Sequentially draw the various model parameters, conditional on $\widetilde{\Delta y^*_{1...T}}$, $\widetilde{\alpha_{1...T}}$, $\widetilde{S_{1...T}}$, $\widetilde{V_{1...T}}$, and the other model parameters.

Additional details on the estimation strategy are available in Appendix A.

5 In-sample results (1970-2019)

This section presents our in-sample results based on data from January 1970 to November 2019. It focuses on the two features that we have added to the standard MS-DFM: high- and low-volatility regimes, and a time-varying long-term GDP growth rate.

5.1 Identifying volatility regimes

The inclusion of different volatility regimes in the state equations of the model allows distinguishing different phases since 1970 (Figure 1): high volatility from 1970 to 1984, during which oil price shocks impacted the U.S. economy; low volatility from 1984 to 2007 (the period known as the *Great Moderation*); a volatility spike during the Great Recession (2007-09); and low volatility from 2009 to 2019. Examining the probability of being in a high-volatility regime, the transitions between the four periods are quite sharp. Based on these results, the increase in volatility during the Great Recession was only temporary and did not signal the end of the Great Moderation, which aligns with other empirical findings such as those reported by Gadea Rivas et al. (2018) and Charles et al. (2018). The model also identifies minor fluctuations in the probability of being in a high-volatility state during the first period (1970-1984), something also pointed out by Antolin-Diaz et al. (2017) who rely on a stochastic volatility specification.

Figure 1: Probability of being in a high-volatility regime

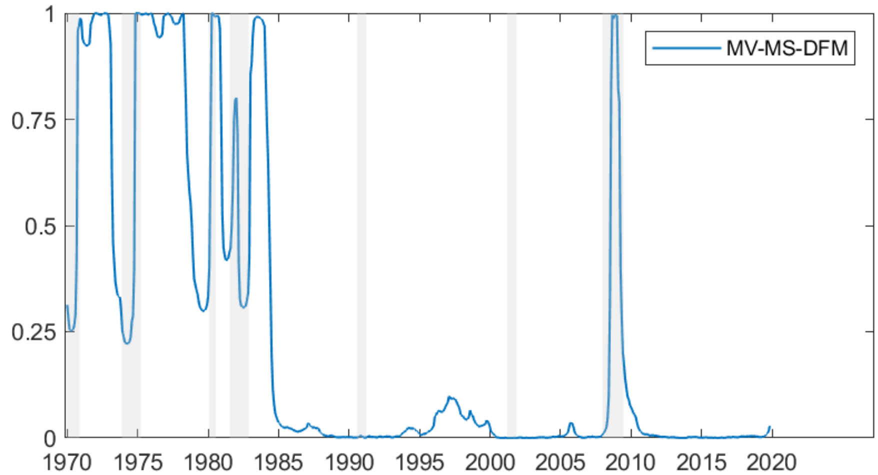

Note: Estimation sample: 1970M01-2019M11. Data vintage: 2019M12. Estimation based on 15000 draws of the Gibbs sampler, from which the first 5000 are discarded. Shaded areas correspond to NBER recessions.

5.2 Dating recessions

Figure 2 shows that allowing for two volatility regimes, along with different expansion and recession characteristics depending on volatility (blue line), helps to identify U.S. recessions compared to a standard MS-DFM where only the mean growth rate can switch (red line). In particular, this specification enhances the ability to capture the two recessions that occurred during the Great Moderation, in 1990-91 and following the burst of the Internet bubble in 2001. This is in line with the results of McConnell and Perez-Quiros (2000) based on a univariate MS model with switching volatility and mean growth rate. Here, we extend their findings by allowing for a multivariate framework and updating their sample.

For the MV-MS-DFM, Table 2 indicates that volatility is associated with deeper recessions but matters less for growth during expansions. More specifically, the discrepancy between high- and low-volatility recessions is twice larger than between high- and low-volatility expansions. This asymmetry in the estimated impact of volatility on economic growth is consistent with the vulnerable growth dynamics put forward by Adrian et al. (2019) where deteriorating financial conditions lead to increased volatility and a decline in the lower quantiles of future GDP growth, while the upper quantiles remain broadly stable. Furthermore, Figure 2 in their paper illustrates that financial conditions were tighter during periods when our model identifies a high-volatility regime, i.e. prior to the Great Moderation and during the Great Recession.

At the same time, Table 2 demonstrates that, for the MS-DFM with constant volatility, the average growth rate during recessions is estimated with greater uncertainty than the average growth rate during expansions. This discrepancy contributes to larger forecasting uncertainty during recessions than during expansions and aligns with the concept of Growth-at-Risk (GaR). Section 6.3 on density forecasts further elaborates on this issue.

Table 2: Intercept of the state equation driving factor dynamics, depending on growth and volatility regimes

		Intercept
MV-MS-DFM	High-volatility recession	-0.77 [-1.07 ; -0.42]
	Low-volatility recession	-0.25 [-0.39 ; -0.09]
	Low-volatility expansion	0.03 [0.00 ; 0.06]
	High-volatility expansion	0.23 [0.08 ; 0.37]
MS-DFM	Recession	-0.50 [-0.76 ; -0.33]
	Expansion	0.03 [0.01 ; 0.05]

Note: Values in squared brackets represent 95% probability intervals. Estimation sample: 1970M01-2019M11. Data vintage: 2019M12.

Figure 2: Probability of being in a recession: MV-MS-DFM and MS-DFM with constant volatility

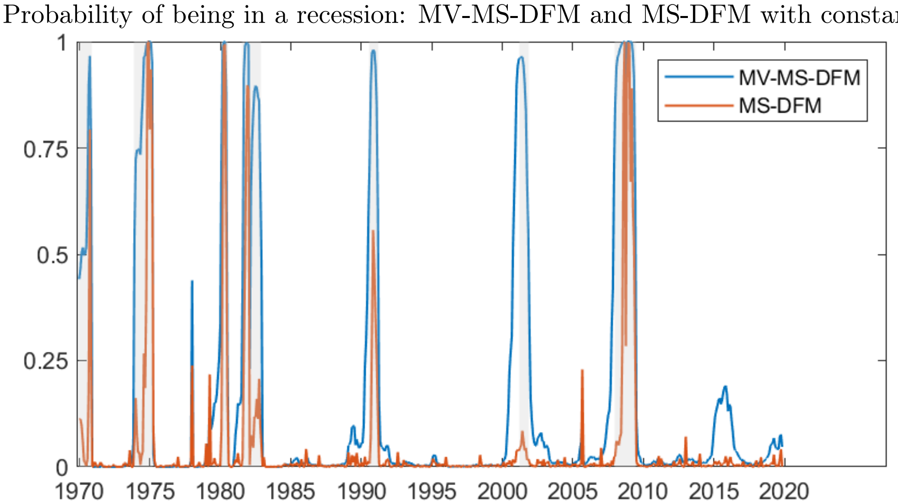

Note: Estimation sample: 1970M01-2019M11. Data vintage: 2019M12. Estimation based on 15000 draws of the Gibbs sampler, from which the first 5000 are discarded. Shaded areas correspond to NBER recessions.

5.3 Declining long-term GDP growth

We identify a rebound in U.S. long-term GDP growth in the second half of the 1990s, which coincides with a well-documented surge in productivity related to an information and communication technology (ICT) shock (Jorgenson et al. 2008). Long-term GDP growth then began to decline in the early 2000s, prior to the Great Recession, and continued to do so thereafter (Figure 3). This observation is consistent with the view that the Great Recession exerted a negative impact on U.S. potential growth by reinforcing pre-existing factors. According to our model estimates, U.S. long-term GDP growth was approximately one percentage point (pp) lower in the early 2010s compared to 2000. A subsequent rebound reduced this gap to 0.7 pp in the late

2010s. These findings align with those obtained by Antolin-Diaz et al. (2017) based on a linear DFM with a time-varying trend.

Figure 3: Estimated U.S. long-run GDP growth

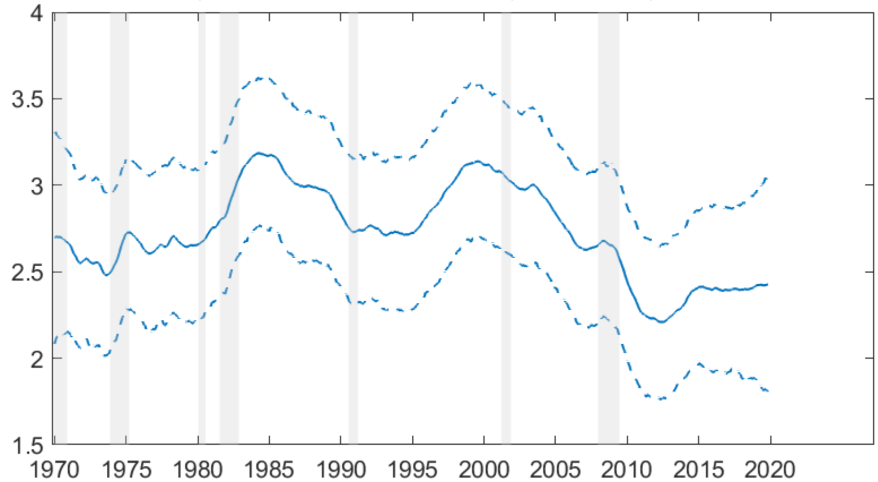

Note: Estimation sample: 1970M01-2019M11. Data vintage: 2019M12. Estimation based on 15000 draws of the Gibbs sampler, from which the first 5000 are discarded. Annualised quarter-on-quarter GDP growth in percentage points. Dotted lines correspond to the 68% credibility interval. Shaded areas correspond to NBER recessions.

5.4 Comparison of specifications based on DICs and marginal likelihoods

The two refinements to the standard MS-DFM introduced in this paper are supported not only by their ability to capture changes in long-term GDP growth and to enhance turning point detection, but also by statistical criteria.

We first rely on the Deviance Information Criterion (DIC) to compare model specifications. This criterion has been specifically designed to compare complex hierarchical models, such as those presented in our paper (Spiegelhalter et al., 2002). It is also used by Eo and Kim (2016). Like other information criteria, the DIC includes two parts: one captures model fit, while the other is a penalty for model complexity:

$$\begin{aligned} DIC =& E_{\tilde{\theta},\tilde{S},\tilde{V}|Y}[-2\log f(\widetilde{\Delta y_{1\ldots T}}|\tilde{\theta},\widetilde{S_{1\ldots T}},\widetilde{V_{1\ldots T}})] \\ &+E_{\tilde{\theta},\tilde{S},\tilde{V}|Y}[-2\log f(\widetilde{\Delta y_{1\ldots T}}|\tilde{\theta},\widetilde{S_{1\ldots T}},\widetilde{V_{1\ldots T}})] + 2\log f(\widetilde{\Delta y_{1\ldots T}}|\tilde{\theta}^{*},\widetilde{S^{*}_{1\ldots T}},\widetilde{V^{*}_{1\ldots T}}) \end{aligned} \tag{8}$$

In equation (8), $\tilde{\theta}$ corresponds to a vector of model parameters (e.g. autoregressive parameters), $\tilde{\theta}^{*}$ to their posterior mean, and $f(\widetilde{\Delta y_{1\ldots T}}|\tilde{\theta},\widetilde{S_{1\ldots T}},\widetilde{V_{1\ldots T}})$ to the model conditional likelihood. The first term increases when the model fit improves, whereas the other two capture the penalty for model complexity. The best model is the one that minimizes the DIC. This information criterion is relatively easy to compute when model estimation is based on Bayesian Markov Chain Monte Carlo (MCMC) techniques, like Gibbs Sampling. Nevertheless, the literature indicates that the DIC is all the more accurate that the number of parameters in the conditioning set of the log-likelihood is limited (Chan and Grant, 2016). Here, we compute the conditional log-likelihood from the Kalman filter step of the Gibbs Sampler. In summary, the Kalman filter estimates the expectation and variance of the model's state vector, conditional on past observations and other model variables and parameters, which then allows estimating $\log f(\Delta y_t|\widetilde{\Delta y_{1\ldots t-1}},\tilde{\theta},\widetilde{S_{1\ldots T}},\widetilde{V_{1\ldots T}})$ for each date t.[2] In this case, the

[2] Given the structure of the model, the following equality applies: $f(\Delta y_t|\widetilde{\Delta y_{1\ldots t-1}},\tilde{\theta},\widetilde{S_{1\ldots T}},\widetilde{V_{1\ldots T}}) =$

conditioning set includes the two Markov variables $\widetilde{S_{1...T}}$ and $\widetilde{V_{1...T}}$ and all constant model parameters, but not the large state vector $\widetilde{\alpha_{1...T}}$. Summing these log-predictive likelihoods over time and averaging over posterior draws of the Gibbs Sampler ultimately provides an estimate of $E_{\tilde{\theta},\tilde{S},\tilde{V}|Y}[-2\log f(\widetilde{\Delta y_{1...T}}|\tilde{\theta},\widetilde{S_{1...T}},\widetilde{V_{1...T}})]$. First averaging over the posterior draws of the Gibbs Sampler and then computing the predictive log-likelihoods with the Kalman filter, allows estimating $\log f(\widetilde{\Delta y_{1...T}}|\tilde{\theta}^*,\widetilde{S^*_{1...T}},\widetilde{V^*_{1...T}})$.

As an alternative to DICs, which may be sensitive to the inclusion of high-dimensional variables in the conditioning set, we also compute marginal log-likelihoods for different model specifications. In this case, the preferred specification is the one with the highest marginal log-likelihood. As before, we take advantage of the fact that, conditional on the two Markov variables, the MS-DFM is a linear Gaussian state-space model. Therefore, we can use the Kalman filter to compute $f(\Delta y_t|\widetilde{\Delta y_{1...t-1}},\tilde{\theta}^*,\widetilde{S_{1...t}},\widetilde{V_{1...t}})$ for each date. The most challenging part of the marginal log-likelihood computation for this class of models is the numerical integration over $\widetilde{S_{1...t}}$ and $\widetilde{V_{1...t}}$ in $f(\Delta y_t|\widetilde{\Delta y_{1...t-1}},\tilde{\theta}^*,\widetilde{S_{1...t}},\widetilde{V_{1...t}})$. To accomplish this, we rely on a particle filter, which provides a sample of N independent particles $\widetilde{S_{1...T}}^n$ and $\widetilde{V_{1...T}}^n$ $(n=1,\ldots,N)$. The idea is to draw $\widetilde{S_{1...t}}^n$ and $\widetilde{V_{1...t}}^n$, conditional on $\widetilde{\Delta y_{1...t-1}}$ and $\tilde{\theta}^*$ using an importance sampling distribution that approximates the true but unknown distribution of these variables. The particle filter draws $\widetilde{S_{1...t}}^n$ and $\widetilde{V_{1...t}}^n$ recursively, based on $\widetilde{S_{1...t-1}}^n$ and $\widetilde{V_{1...t-1}}^n$. The functioning of the particle filter and key issues such as the choice of the importance sampling distribution and the weighting of the particles are explained in detail in Appendix B.

This previous step gives $f(\Delta y_t|\widetilde{\Delta y_{1...t-1}},\tilde{\theta}^*)$ for each date, from which we compute $\log f(\widetilde{\Delta y_{1...T}}|\tilde{\theta}^*)$. We ultimately rely on Chib's (1995) formula to estimate the marginal log-likelihood $\log f(\widetilde{\Delta y_{1...T}})$ of each specification, as follows:

$$\log f(\widetilde{\Delta y_{1...T}}) = \log f(\widetilde{\Delta y_{1...T}}|\tilde{\theta}^*) + \log\phi(\tilde{\theta}^*) - \log\phi(\tilde{\theta}^*|\widetilde{\Delta y_{1...T}})$$

$\tilde{\theta}^*$ is the vector of constant model parameters evaluated at their posterior mean. $\log\phi(\tilde{\theta}^*)$ and $\log\phi(\tilde{\theta}^*|\widetilde{\Delta y_{1...T}})$ are the log-prior and log-posterior densities of these parameters, respectively. Note that the use of a particle filter to reduce the dimension of the conditioning set prior to the use of Chib's (1995) formula is key to ensure accurate results, as shown by Frühwirth-Schnatter and Wagner (2008).

To the best of our knowledge, Kaufmann (2000) is the only other paper in the literature that computes the marginal likelihood of MS-DFMs. However, our framework is more complex because we need to account for Markov-Switching volatility and time-varying growth rates. Therefore, the marginal likelihood computation method described in Appendix B can be considered as an additional contribution of the present paper to the literature.

Table 3 compares the DICs and marginal log-likelihoods of four model specifications: a linear DFM, a MS-DFM in which only the intercept of the state equation governing factor dynamics can switch, a MS-DFM in which both the intercept and the volatility can switch in the state equation, and the MV-MS-DFM. All models use the same five observable variables (see Section 4). Both Markov-Switching features reduce the DIC compared to the linear DFM. This reduction is particularly pronounced for Markov-Switching volatility. Allowing for a time-varying long-term GDP growth rate slightly deteriorates the DIC and the marginal log-likelihood, although not to a point that allows for a clear-cut conclusion. It seems prudent to consider that the last two model specifications cannot be distinguished based on these statistical criteria. The fact that quarterly GDP is only one of the five observed variables, and observed only once every three months, may help to explain this result. In any case, the time-varying long-term GDP growth rate is an important feature to account for the limited rebound in economic activity after the Great Recession.

$f(\Delta y_t|\widetilde{\Delta y_{1...t-1}},\tilde{\theta},\widetilde{S_{1...t}},\widetilde{V_{1...t}})$.

Table 3: Comparison of model specifications based on DICs and marginal log-likelihoods

Specification	DIC	Marginal Log-Lik
Linear DFM	3883.1	-1978.3
MS-DFM: MS on intercept only	3822.2	-1963.4
MS-DFM: MS on intercept and volatility	**3720.9**	**-1934.7**
MS-DFM: MS on intercept and volatility + TV trend	3728.6	-1937.4

Note: The posterior mean of model parameters is estimated based on 15000 draws of the Gibbs sampler, from which the first 5000 are discarded. In addition, 500 particles are used for marginal log-likelihood computations (see Appendix B for details). Sample: 1970M01-2019M11. Data vintage: 2019M12. The preferred specification is the one minimising the DIC and maximising the marginal log-likelihood. In each case, this specification is indicated in bold.

6 Real-time performance of the MV-MS-DFM

In this section, we conduct a real-time assessment of the MV-MS-DFM using data vintages available from January 2007 to December 2019. This assessment includes the real-time estimation of long-run GDP growth, the real-time identification of the Great Recession, and real-time density forecasts of GDP growth at a six-month horizon.

6.1 Real-time estimation of U.S. long-term GDP growth

Figure 4 illustrates how the MV-MS-DFM uncovers the decline in U.S. long-term GDP growth during the 2000s in real time. As expected, the model requires some time to determine that the decline in GDP growth during this period is not attributable to business cycle developments but to factors influencing long-term growth. Until 2010, estimated long-run growth remains above 3%, but it progressively decreases as subsequent data vintages become available (dashed curves). It is noteworthy that the professional forecasters surveyed by the Philadelphia FED revised their long-term growth expectations by a similar amount over the same period, and that both estimates converge at the end of the sample.[3] Nevertheless, the MV-MS-DFM identifies the decline in long-term GDP growth more rapidly than the Survey of Professional Forecasters (SPF).

[3]We rely on variable RGDP10 in the SPF, which corresponds to average GDP growth projections over the next 10 years. This question is only asked once a year in January, which is why we compare the SPF with our model using data vintages that are made available in January each year.

Figure 4: U.S. long-run GDP growth estimated in real time

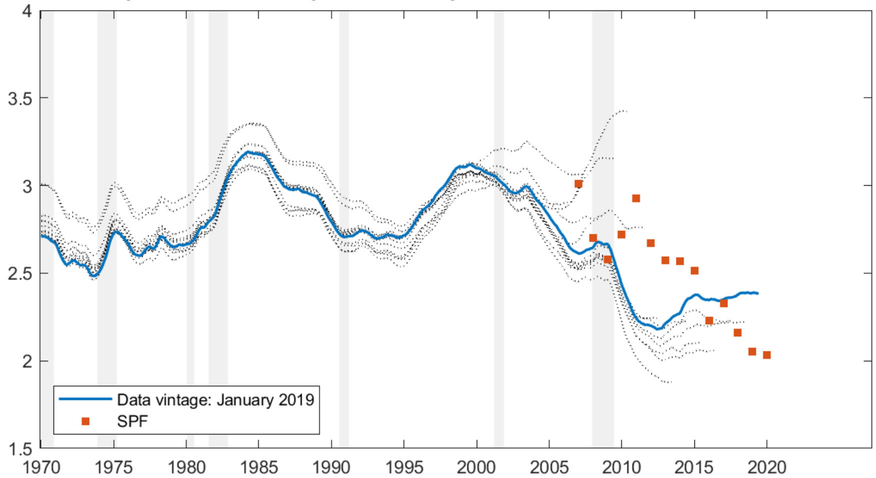

Note: Data vintages: 2007M01 to 2019M01. Estimation based on 15000 draws of the Gibbs sampler, from which the first 5000 are discarded. Shaded areas correspond to NBER recessions.

6.2 Real-time detection of the Great Recession

We now rely on real-time data vintages to assess when the model detects the onset and the end of the Great Recession.

There is no consensus in the literature regarding the timing of announcements for the onset and the end of a recession, based on recession probabilities estimated using Markov-Switching models. Here, we rely on a decision rule that is similar to those used by Chauvet and Piger (2008) and Hamilton (2011).[4] Concretely, we consider that a recession has begun when the recession probability exceeds 0.7 and stays above this threshold for three consecutive months, and we identify the onset of the recession as the first month for which the probability of recession exceeds 0.5. Conversely, we announce that a recession is over when the probability of recession declines below 0.3 and remains below this threshold for three consecutive months, and we identify the end of the recession as the last month for which the probability is above 0.5.

Results are reported in Table 4. The dates of the Great Recession identified by the MV-MS-DFM are in close agreement with those published by the NBER. Nevertheless, the model identifies the start and the end of the recession respectively two and eleven months prior to the NBER's official announcement. This demonstrates that the model is capable of providing reliable forecasts of business cycle turning points in a timely manner.

[4] Hamilton (2011) suggests the following decision rule when relying on quarterly GDP data (with his notations $S_t = 1$ during recessions and $S_t = 0$ during expansions): "When the one-quarter ahead smoothed inference $P(S_t = 1|y_{t+1}, y_t, \ldots, y_1; \hat{\theta}_{t+1})$ first exceeds 0.65, declare that a recession has started, and at that time assign a probable starting point for the recession as the beginning of the most recent set of observations for which $P(S_{t-j} = 1|y_{t+1}, y_t, \ldots, y_1; \hat{\theta}_{t+1})$ exceeds 0.5." Using only the four NBER monthly series to infer recession probabilities, Chauvet and Piger (2008) require that the probability of recession moves from below to above 0.8 and remains above 0.8 for three consecutive months before announcing a recession. Similarly to Hamilton (2011), they identify the start of the recession as the first month for which the probability exceeds 0.5.

Table 4: Real-time dating of the Great Recession by the NBER and the MV-MS-DFM

Business cycle peak			
Date		**Announcement**	
NBER	MV-MS-DFM	NBER	MV-MS-DFM
2008 M01	2007 M11	2008 M12	2008 M10
Business cycle trough			
Date		**Announcement**	
NBER	MV-MS-DFM	NBER	MV-MS-DFM
2009 M06	2009 M06	2010 M09	2009 M10

6.3 Real-time density forecasts and Growth-at-Risk

We now focus on the real-time density forecasts produced by linear DFMs, two different specifications of the MS-DFM (with and without Markov-Switching volatility), and we discuss how our results relate to the Growth-at-Risk literature.

Since the seminal work of Adrian et al. (2019) who relied on quantile regressions to capture macroeconomic tail risks, other econometric specifications, including Bayesian VARs (Carriero et al., 2024) and Markov-Switching models (Doz et al., 2020 ; Caldara et al., 2021) have been proposed to account for asymmetric predictive densities. While Caldara et al. (2021) first extract macroeconomic and financial factors using linear DFMs and then assert that the link between those factors and GDP growth depends on a Markov-Switching state variable, our approach treats GDP and the other macroeconomic variables in a similar way. As both GDP and macroeconomic variables, such as industrial production and sales, ultimately measure economic activity, albeit at different frequencies and from different perspectives, their underlying dynamics are unlikely to differ. For example, Loria et al. (2024) show that quarterly real GDP and monthly industrial production are both subject to similar tail risks. This finding supports our econometric specification, wherein Markov Switching underlies the dynamics of all variables.

GDP forecasts for the current and upcoming quarters can be generated at the end of each month. However, the analysis in this paper concentrates on forecasts produced six months ahead of the quarterly GDP release date by the US Bureau of Economic Analysis (BEA). Given that quarterly GDP for quarter Q is released at the end of the first month of quarter $(Q + 1)$, these forecasts are produced at the end of the first month of quarter $(Q-1)$. Since all models presented in this paper are estimated using a Bayesian approach, uncertainty can arise from three sources: future shocks, parameter estimation, and the identification of the underlying business cycle and volatility regimes.

Both MS-DFM specifications, with and without Markov-Switching volatility, generate left-skewed predictive densities, indicating an increase in the probability of low GDP growth during recessions while the probability of high GDP growth remains broadly constant over time, in line with the findings of Adrian et al. (2019). This is a key difference with linear DFM specifications, with or without TV-Trend, where uncertainty is centered around long-term GDP growth and similar in all phases of the business cycle (Figure 5). It should be stressed that the Growth-at-Risk features of the MS-DFM specifications are obtained by only including the four real variables advocated by the NBER Business Cycle Dating Committee in the information set, i.e. without any financial variable. Possible ways to account for financial conditions will be explored in more detail in the next section.

Figure 5: Real-time density forecast (6-month horizon): Linear DFM

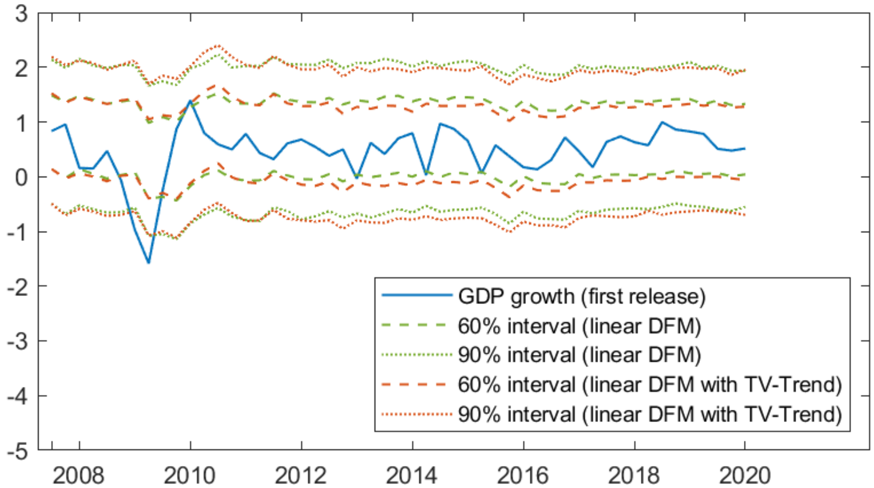

Note: Expanding estimation sample starting in 1970M01. 6-month ahead density forecasts are produced for each quarter from 2007Q2 to 2019Q4. All density forecasts are based on 7500 draws of the Gibbs sampler, from which the first 2500 are discarded. GDP growth is the first release of the non-annualised quarter-on-quarter U.S. GDP growth rate.

With the MS-DFM specification, where only the mean growth rate can switch (Figure 6), the asymmetry of the predictive distribution is related to the fact that the average GDP growth during recessions is estimated with greater uncertainty than the average GDP growth during expansions (Table 2). This phenomenon can be attributed to the limited number of recessions and the substantial variation in the depth of recessions over the estimation sample. Without increased uncertainty surrounding average GDP growth during recessions, all quantiles would shift in parallel, resulting in a downward shift of the entire distribution during recessions, without any asymmetry.

Figure 6: Real-time density forecast (6-month horizon): MS-DFM with constant volatility

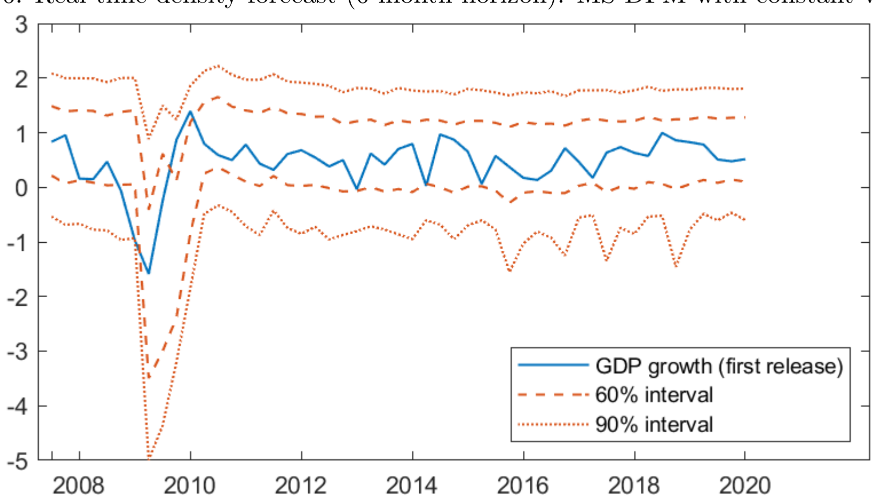

Note: Expanding estimation sample starting in 1970M01. 6-month ahead density forecasts are produced for each quarter from 2007Q2 to 2019Q4. All density forecasts are based on 7500 draws of the Gibbs sampler, from which the first 2500 are discarded. GDP growth is the first release of the non-annualised quarter-on-quarter U.S. GDP growth rate.

With the MV-MS-DFM, time-varying volatility also contributes to the asymmetry of the predictive distribution (Figure 7). Even if average GDP growth in expansion and recession states were estimated without any uncertainty, the combination of lower average growth and higher volatility during certain recessions would generate asymmetric predictive distributions, as highlighted by Carriero et al. (2024). Indeed, the effects of increased volatility and lower average GDP growth tend to compensate each other in the upper part of the distribution, while they cumulate in the lower part of the distribution.

Figure 7: Real-time density forecast (6-month horizon): MV-MS-DFM

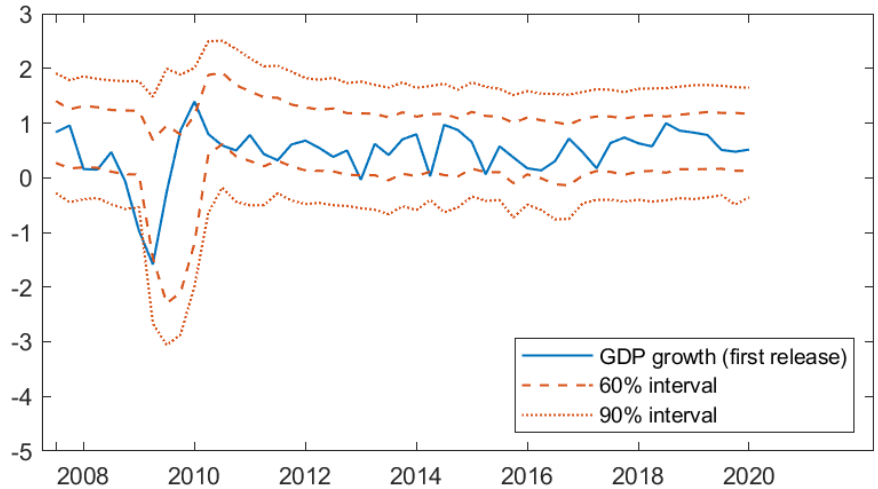

Note: Expanding estimation sample starting in 1970M01. 6-month ahead density forecasts are produced for each quarter from 2007Q2 to 2019Q4. All density forecasts are based on 7500 draws of the Gibbs sampler, from which the first 2500 are discarded. GDP growth is the first release of the non-annualised quarter-on-quarter U.S. GDP growth rate.

We consider three metrics to assess the quality of the predictive densities of the linear DFMs and of the MS-DFM where only the mean growth rate can switch, compared to the MV-MS-DFM, with a focus on downside risks: the 5% and 10% quantile scores (Giacomini and Komunjer, 2005) and the left-tail focused quantile-weighted continuous ranked probability score (qw-CRPS-L) (Gneiting and Ranjan, 2011).[5]

From 2007Q2 to 2019Q4, the MV-MS-DFM outperforms the alternative models (Table 5). While the MS-DFM where only the mean growth rate can switch generates left-skewed predictive densities during the Great Recession, it exaggerates the downside risk to the same extent that linear DFMs underestimate it, which is why their lower quantile scores are comparable in the end.

[5]The α% quantile score of the predictive density at horizon h is calculated as:

$$\mathrm{QS}_{\alpha,h} = \frac{1}{N+1} \sum_{t=t_0}^{t_0+N} \left(y_{t+h} - Q_{\alpha,t+h}\right) \cdot \left(\alpha - \mathbf{1}_{y_{t+h} \leq Q_{\alpha,t+h}}\right)$$

where y_{t+h} is the first release of the quarterly GDP growth at date $(t+h)$ and $Q_{\alpha,t+h}$ is the α% quantile of the predictive density at date $(t+h)$. t_0 and (t_0+N) are the first and the last dates at which forecasts at horizon h are produced. $N = 50$ for the 2007Q2–2019Q4 sample.

The quantile-weighted continuous ranked probability score is a weighted average across quantile scores, with larger weights for lower quantiles:

$$qwCRPSL_h = \frac{2}{J-1} \sum_{j=1}^{J-1} [\nu_{\alpha_j} \cdot QS_{\alpha_j,h}]$$

where $\alpha_j = j \cdot 10\%$ and $\nu_{\alpha_j} = (1-\alpha_j)^2$ is the corresponding weight.

Table 5: Relative predictive density accuracy over 2007Q2–2019Q4, compared to the MV-MS-DFM

	5% quantile score	10% quantile score	qw-CRPS-L
Linear DFM	0.86*	0.84	0.98
Linear DFM with TV-Trend	0.82*	0.79**	0.95
MS-DFM where only the mean growth rate can switch	0.87	0.83***	0.95**

*Note: A ratio below 1 indicates that the predictive density of the MV-MS-DFM is more accurate than the alternative. *, ** and *** indicate the statistical significance of differences in quantile scores and continuous ranked probability scores at the 10%, 5% and 1% levels, respectively. While limiting distributions in Gneiting and Ranjan (2011) are obtained for fixed-length estimation samples, their relative accuracy test is here applied to estimation samples expanding from 1970-2007 to 1970-2019.*

7 Vulnerability of GDP growth to financial conditions

The way financial conditions influence future GDP growth and how to take it into account for macroeconomic forecasting is still being debated. While Adrian et al. (2019) and Adrian et al. (2022) explicitly link downside risks to U.S. GDP growth to tighter financial conditions, Huang et al. (2024) highlight that financial conditions comove with macroeconomic uncertainty, and that the way uncertainty varies over time plays a key role in shaping the distribution of future growth by itself. In a similar vein, Plagborg-Møller et al. (2020) show that the information provided by financial variables for forecasting the distribution of future GDP growth is significantly reduced once past and current macroeconomic conditions are already accounted for. This does not imply that financial shocks are irrelevant in explaining macroeconomic developments and GDP growth. Actually, Huang et al. (2024) show that structural financial shocks tightening financial conditions shift the entire distribution of future GDP growth downwards, and increase macroeconomic uncertainty, which in turn widens the distribution of future GDP growth. As a result, the upper quantiles of future GDP growth remain broadly unchanged while the lower quantiles decrease significantly, which is consistent with the findings of Adrian et al. (2019). The main question is more about the marginal informational content of financial variables and how to use it in practice for macroeconomic forecasting.

We now rely on a standard MS-DFM where only the mean growth rate can switch. In this model, a variable summarising financial conditions can be used in two distinct ways: (1) as a variable akin to industrial production or employment, from which a useful signal on economic activity can be extracted, similarly to what Adrian et al. (2019) do in a quantile regression framework, and (2) as a variable indicating that GDP growth is becoming more or less vulnerable over time, along the lines of Adrian et al. (2022). We assess both options simultaneously in the following, with the MS-DFM specification where only the mean growth rate can switch.[6]

Following Adrian et al. (2019), we use the Chicago FED's National Financial Condition Index (NFCI) as our measure of financial conditions in the U.S.. We first include it in the model information set, along with the macroeconomic variables used by the NBER Business Cycle Dating Committee. In addition, the model is extended to allow for time-varying transition probabilities between expansions and recessions. Following the seminal work of Filardo and Gordon (1998), this is done with a Probit specification where the probability of switching from a Markov state to another depends on the current state and financial conditions.

The Probit model is specified as follows:

$$P\left(S_t = 0 \mid S_{t-1}, NFCI_{t-1}\right) = P\left(S_t^* < 0 \mid S_{t-1}, NFCI_{t-1}\right) \tag{9}$$

[6]With simulated data where the number and length of recessions and high-volatility episodes are similar to the U.S. evidence, all attempts to estimate an MV-MS-DFM where financial conditions can influence both $P(S_t = j|S_{t-1} = i)$ and $P(V_t = j|V_{t-1} = i)$ proved unsuccessful. This may be due to the fact that the sample does not include enough transitions between high- and low-volatility regimes to properly identify the impact of financial conditions on $P(V_t = j|V_{t-1} = i)$.

where

$$S_t^* = \beta_0^S + \beta_1^S S_{t-1} + \beta_2^S \text{ NFCI}_{t-1} + u_t \quad ; \quad u_t \sim \text{i.i.d. } N(0,1) \tag{10}$$

While the NFCI is only weakly related to the underlying economic activity factor (the posterior distribution of γ_5^m does not significantly deviate from zero), tightening financial conditions significantly increase the probability of entering and staying in a recession, as shown by the posterior distribution of β_2^S (Table 6, Figure 8). This confirms that tighter financial conditions make growth more fragile, in line with Adrian et al. (2019) but relying on a different model. Assessing which specific financial variables are the most relevant and whether their information content is sufficient to significantly improve macroeconomic forecasts in real time is left for future research.

Table 6: Estimated Factor loadings and Probit parameters in a MS-DFM with financial conditions

Parameter	Prior distribution	Posterior distribution - 95% Credibility Interval
γ_1^m	$N(0,1)$	$[1.81; 2.36]$
γ_2^m	$N(0,1)$	$[1.72; 2.23]$
γ_3^m	$N(0,1)$	$[0.60; 0.96]$
$\gamma_{4,0}^m$	$N(0,1)$	$[0.17; 0.38]$
$\gamma_{4,1}^m$	$N(0,1)$	$[-0.05; 0.22]$
$\gamma_{4,2}^m$	$N(0,1)$	$[-0.05; 0.21]$
$\gamma_{4,3}^m$	$N(0,1)$	$[0.06; 0.27]$
γ_5^m	$N(0,1)$	$[-0.11; 0.06]$
β_0^S	$N(\beta_0^{S,prior}, 0.01)$ with $\Phi[\beta_0^{S,prior}] = \frac{7}{83}$	$[-1.95; -1.63]$
β_1^S	$N(\beta_1^{S,prior}, 0.01)$ with $\Phi[\beta_0^{S,prior} + \beta_1^{S,prior}] = \frac{511}{517}$	$[3.27; 3.63]$
β_2^S	$N(0, 0.01)$	$[-0.30; -0.06]$

Note: The means of the prior distributions for β_0^S and β_1^S reflect the number of recessions identified by the NBER and their average length over 1970-2019. There were 7 recessions spanning 83 months during this period. Assuming that financial conditions do not play a role, the prior mean of β_0^S is set so that the probability of exiting a recession is $\frac{7}{83}$. As already noticed by Filardo and Gordon (1998), the priors of the Probit specification need to be relatively tight to achieve convergence. Although the prior distribution of β_0^S is tightly centered around zero, its posterior distribution significantly deviates from zero, which means that the data is informative about the role played by financial conditions. Φ is the cumulative distribution function of the standard normal distribution.

Figure 8: Probability of falling in a recession

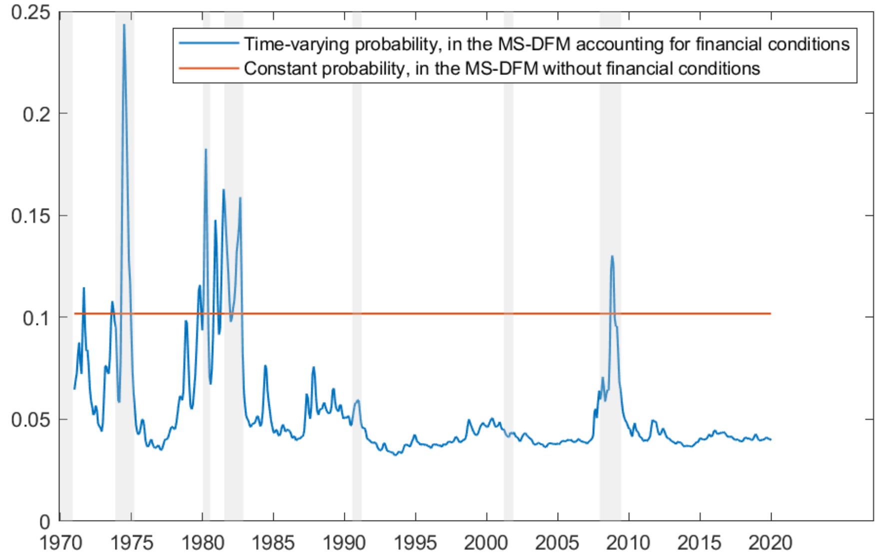

Note: Estimation sample: 1970M01-2019M12. Data vintage: 2024M09. Estimation based on 15000 draws of the Gibbs sampler, from which the first 5000 are discarded. The above displayed probability of falling in a recession corresponds to $P(S_t = 0|S_{t-1} = 1, NFCI_{t-1})$*. It should not be confused with the probability of being in a recession* $P(S_t = 0|\widetilde{\Delta y_{1...T}})$ *that is estimated by the Kalman smoother and displayed in Figure 2. Shaded areas correspond to NBER recessions.*

8 Dealing with the COVID-19 period

The economic lockdown during the COVID-19 pandemic triggered a deep recession in most advanced countries, which, according to the NBER Business Cycle Dating Committee, lasted from February 2020 (peak) to April 2020 (trough) in the U.S. This event was exceptional in many aspects and led to a truly unusual behaviour of economic time series, characterized by substantial fluctuations and volatility over a relatively short period. Although our MV-MS-DFM has been designed to handle high volatility periods, it cannot adequately address this very specific period without a targeted treatment for the COVID-19 period.

Two main approaches have been proposed in the literature to address the specific features of the COVID-19 period. First, some authors explicitly model time-varying volatility and outliers. For instance, Carriero et al. (2024) rely on different BVARs with stochastic volatility (SV) and outliers, while Antolin-Diaz et al. (2024) rely on a DFM that integrates stochastic volatility, fat tails and outliers. Second, other authors suggest to simply drop the COVID-19 period from the estimation sample, especially when the focus is not on this specific period. This is what Lenza and Primiceri (2022) do for estimating a BVAR with time-varying volatility. Maroz, Stock and Watson (2021) add a specific "COVID factor" in an otherwise standard DFM but also recommend excluding the COVID-19 period when the focus is not on this specific period. Schorfheide and Song (2021) and Baumeister and Hamilton (2023) similarly exclude the COVID-19 period from their econometric estimates.

Although the MV-MS-DFM introduced in the present paper includes two volatility regimes and accounts for deeper recessions when volatility is high, it is not suitably equipped to address the unusually large recession and volatility observed during the COVID-19 pandemic. Including this period in the sample without any further refinement would result in a situation where the only high-volatility episode identified by the model would be the COVID-19 period, leading most previous recessions to go unnoticed by the model.

In this paper, we consider two approaches that ultimately very similar results for the post-COVID-19 period (Figures 10 and 9), suggesting that both approaches are relevant.

The first approach consists in skipping the COVID-19 period, as proposed in the aforementioned literature. Here, we choose to omit the years 2020 and 2021 and to estimate the model from 1970 to 2024, without these two years. The resulting probability of recession (red line in Figure 9) captures all past recessions, but is not available during the COVID-19 period, by construction. When excluding 2020 and 2021, the recession probability remains low since 2015, despite some minor fluctuations. Similarly, the probability of being in a high-volatility regime (red line in Figure 10) displays a clear historical pattern, consistent with the historical narrative we previously discussed, but remains at zero since the end of the Great Recession.

In the second approach, we allow the model to go through the COVID-19 period, but without re-estimating its parameters. More precisely, the parameters are estimated over 1970-2019 and then frozen. After 2019, we simply update the estimation of the underlying factor and of the probabilities of being in a recession and a high-volatility regime.[7] In this case, the model identifies a high-volatility regime from January 2020 to August 2021 (blue line in Figure 10) and a recession from January to April 2020 (blue line in Figure 9). The recession probability peaks again in February 2021, which seems related to a temporary decline in real manufacturing and trade sales and real personal income, but this peak is too short-lived to consider that a new recession has started, according to the criteria we put forward in section 6.2.

Figure 9: Probability of being in a recession

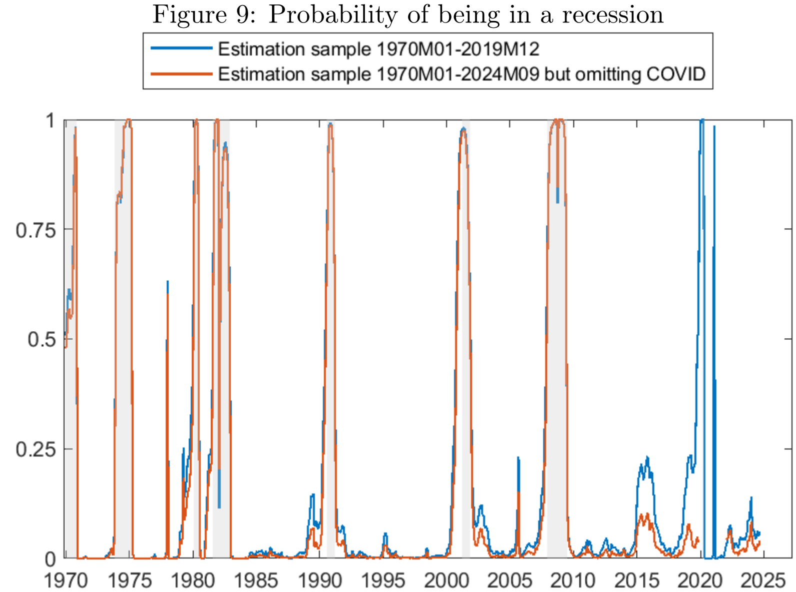

Note: Data vintage: 2024M09. Estimation based on 15000 draws of the Gibbs sampler, from which the first 5000 are discarded. Shaded areas correspond to NBER recessions.

[7]More precisely, Kalman and Hamilton smoothers are run over the full sample to estimate the underlying factor tracking economic activity, long-run GDP growth, and probabilities of being in a recession and a high-volatility regime over 1970M01-2024M09, but all other parameters are drawn in posterior distributions estimated over 1970M01-2019M12 to avoid distortions related to the COVID-19 period.

Figure 10: Probability of being in a high-volatility regime

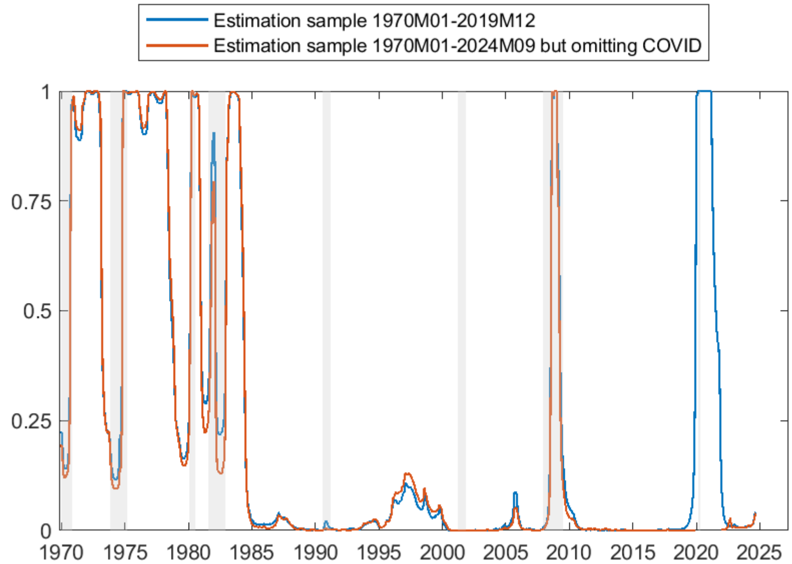

Note: Data vintage: 2024M09. Estimation based on 15000 draws of the Gibbs sampler, from which the first 5000 are discarded. Shaded areas correspond to NBER recessions.

While further research would help refine the specific period to be removed from the estimation sample in order to prevent the COVID-19 period from permanently distorting model estimates, removing the years 2020 and 2021 entirely appears to be a transparent and straightforward solution to this issue.

9 Conclusion

This paper extends the standard Markov-Switching Dynamic Factor Model (MS-DFM) to account for abrupt changes in macroeconomic volatility and longer-term fluctuations in long-run GDP growth, two important characteristics of the U.S. economy that are also prevalent in other advanced economies. We refer to this model as the Mean-Variance Markov-Switching Dynamic Factor Model (MV-MS-DFM).

The MV-MS-DFM identifies a continuous decline in U.S. long-term GDP growth that began in the early 2000s, consistent with the existing literature, and quantifies this decline at 0.7 percentage points (pp) between 2000 and 2019. While long-term growth declined by 0.5 pp before 2008, this decline persisted after the Great Recession, reaching 1 pp by 2012, and has slightly narrowed since then.

The introduction of Markov-Switching volatility into the standard MS-DFM framework is supported by statistical criteria (DICs and marginal likelihoods) and improves the detection of turning points in the U.S. business cycle during the Great Moderation. The model provides evidence indicating that the Great Moderation, which began in the mid-1980s, was only temporarily interrupted by the Great Recession of 2008-09 and the COVID-19 pandemic of 2020-21. Therefore, the ability to detect recessions in both high- and low-volatility environments is relevant for the detection of future recessions.

As the model is estimated using Bayesian techniques, it can generate density forecasts that encompass three sources of uncertainty: future shocks, parameter estimation, and the identification of the underlying business cycle and volatility regimes. While accounting for two business cycle regimes (expansions and recessions) is key to capture abrupt changes in GDP growth, accounting for the larger uncertainty in average GDP growth during recessions than expansions, as well as for two volatility regimes, allows the model to generate left-skewed density forecasts, as advocated in the Growth-at-Risk literature initiated by Adrian et al. (2019). Our model can generate left-skewed density forecasts without including financial variables in the information set, but we show that tighter financial conditions make future GDP growth more fragile and increase the probability of falling in a recession.

While the model may not adequately capture the exceptionally large recession and the huge increase in volatility during the COVID-19 pandemic without further refinements, freezing its parameters during this period provides a transparent and straightforward solution to this issue. In this case, the model identifies a short-lived recession in early 2020 and an increase in volatility until mid-2021.

References

[1] Adrian, T., Boyarchenko, N. and Giannone, D. (2019), "Vulnerable Growth", American Economic Review, 109(4), 1263–1289.

[2] Adrian, T., Grinberg, F., Yang, N., Malik, S. and Yu, J. (2022), "The Term Structure of Growth-at-Risk", American Economic Journal: Macroeconomics 14(3), 283-323.

[3] Amburger, A.J. and McCracken, M.W. (2023), "On the Real-Time Predictive Content of Financial Condition Indices for Growth", Journal of Applied Econometrics 38(2), 137-163.

[4] Antolin Diaz, J., Drechsel, T. and Petrella, I. (2017), "Tracking the Slowdown in Long-Run GDP Growth", The Review of Economics and Statistics 99(2), 343-356.

[5] Antolin Diaz, J., Drechsel, T. and Petrella, I. (2024), "Advances in Nowcasting Economic Activity: The Role of Heterogeneous Dynamics and Fat Tails, Journal of Econometrics, 238.

[6] Bai, J. and Wang, P. (2011), "Conditional Markov Chain and its Application in Economic Time Series Analysis", Journal of Applied Econometrics 26, 715-734.

[7] Baumeister, C. and J. Hamilton (2024), "A Full Information Approach to Granular Instrumental Variables", *Mimeo.*

[8] Burns, A.F. and Mitchell, W.C. (1946), "Measuring Business Cycles", National Bureau of Economic Research.

[9] Caldara, D., Cascaldi-Garcia, D., Cuba-Borda, P. and Loria F.(2021), Understanding Growth-at-Risk: A Markov SwitcTing Approach, SSRN Electronic Journal, DOI:10.2139/ssrn.3992793

[10] Camacho, M., Perez-Quiros, G. and Poncela, P. (2018), "Markov Switching Dynamic Factor Models in Real Time", International Journal of Forecasting 34, 598-611.

[11] Camacho, M., Perez-Quiros, G. and Poncela, P. (2014), "Green Shoots and Double Dips in the Euro Area: A Real-Time Measure", International Journal of Forecasting 30, 520-535.

[12] Carriero, A., Clark, T.E., Marcellino, M. and Mertens, E. (2024), "Addressing COVID-19 Outliers in BVARs with Stochastic Volatility", The Review of Economics and Statistics 106(5), 1403-1417.

[13] Carriero, A., Clark, T.E. and Marcellino, M. (2024), "Capturing Macroeconomic Tail Risks with Bayesian Vector Autoregressions", Journal of Money, Credit and Banking 56, 1099-1127.

[14] Charles, A., Darné, O. and Ferrara, L. (2018), "Does the Great Recession Imply the End of the Great Moderation? International evidence", Economic Inquiry 56(2), 745-760.

[15] Chauvet, M., Senyuz, Z. and Yoldas, E. (2015), "What does Financial Volatility Tell us about Macroeconomic Fluctuations?" Journal of Economic Dynamics and Control 52, 340-360.

[16] Chauvet, M. and Su, Y. (2014), "Nonstationarities and Markov Switching Models", in Recent Advances in Estimating Nonlinear Models With Applications in Economics and Finance, eds Ma and M. Wohar, Springer Science Business Media, New York.

[17] Chauvet, M. and Potter, S. (2013), "Forecasting Output", Handbook of Economic Forecasting, Chapter 3, 141-194.

[18] Chauvet, M. and Piger, J. (2008), "A Comparison of the Real-Time Performance of two Business Cycle Dating Methods", Journal of Business and Economics Statistics 26(1), 42-49.

[19] Chauvet, M. and Hamilton, J.D. (2005), "Dating Business Cycle Turning Points", NBER Working Paper 11422.

[20] Chib, S. (1995), "Marginal Likelihood from the Gibbs Output", Journal of the American Statistical Association, 90(432), 1313-1321.

[21] Clark, T.E. and McCracken, M.W. (2013), "Advances in Forecast Evaluation", Handbook of Economic Forecasting, Chapter 20, 1107-1201.

[22] Diebold, F.X. and Rudebusch, G.D. (1996), "Measuring Business Cycles: A Modern Perspective", The Review of Economics and Statistics 78(1), 67-77.

[23] Diebold, F.X. and Mariano, R.S. (1995), "Comparing Predictive Accuracy", Journal of Business and Economic Statistics 13, 253-263.

[24] Doz, C., Ferrara, L. and Pionnier P-A. (2020), "Business cycle dynamics after the Great Recession: An Extended Markov-Switching Dynamic Factor Model", Paris School of Economics Working Paper 2020-03.

[25] Eo, Y. and Kim, C.J. (2016), "Markov Switching Models with Evolving Regime-Specific Parameters: Are Post-War Booms or Recessions all Alike?", The Review of Economics and Statistics 98(5), 940-949.

[26] Fernald, J., Hall, R., Stock, J. and Watson, M. (2017), "The Disappointing Recovery of Output after 2009", Brooking Papers on Economic Activity 48(1), 1-81.

[27] Filardo, A.J. and Gordon, S.J. (1998), "Business Cycle Duration", Journal of Econometrics 85, 99-123.

[28] Gadea Rivas, M.D., Gomez Loscos, A. and Perez Quiros, G. (2018), "Great Moderation and Great Recession: From Plain Sailing to Stormy Seas?", International Economic Review 59(4), 2297-2321.

[29] Giacomini, R. and Komunjer, I. (2005), "Evaluation and Combination of Conditional Quantile Forecasts", Journal of Business and Economic Statistics 23(4), 416-431.

[30] Giordani, P., Kohn R. and van Dijk D. (2007), "A Unified Approach to Nonlinearity, Structural Change and Outliers", Journal of Econometrics 137, 112-133.

[31] Gneiting, T. and Ranjan, R. (2011), "Comparing Density Forecasts using Threshold- and Quantile-Weighted Scoring Rules", Journal of Business and Economic Statistics 29(3), 411-422.

[32] Hamilton, J. (2011), "Calling Recessions in Real Time", International Journal of Forecasting 27, 1006-1026.

[33] Harvey, D., Leybourne, S. and Newbold, P. (1997), "Testing the Equality of Prediction Mean Square Errors", International Journal of Forecasting 13, 281-291.

[34] Huang, Y.F., Liao, W., Luo, S. and Ma, J. (2024), "Financial Conditions, Macroeconomic Uncertainty and Macroeconomic Tail Risks", Journal of Economic Dynamics and Control 163.

[35] Jorgenson, D.W., Ho, M.S. and Stiroh K.J. (2008), "A Retrospective Look at the US Productivity Growth Resurgence", Journal of Economic Perspectives 22(1), 3-24.

[36] Kim, C.J. and Nelson, C.R. (1998), "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.

[37] Lenza, M. and G. Primiceri (2022), "How to Estimate a Vector AutoRegression after March 2020?", Journal of Applied Econometrics, 37(4), 688-699.

[38] Lhuisser, S. (2022), "Financial Consitions and Macroeconomic Downside Risks in the Euro Area", European Economic Review, 143.

[39] Loria F., Matthes, C. and Zhang, D. (2024), "Assessing Macroeconomic Tail Risk", The Economic Journal 135(665), 264-284.

[40] Marcellino, M., Porqueddu, M. and Venditti, F. (2016), “Short-Term Forecasting with a Mixed-Frequency Dynamic Factor Model with Stochastic Volatility”, Journal of Business and Economic Statistics 34(1), 118-127.

[41] Mariano, R. and Murasawa, Y. (2003), “A New Coincident Index of Business Cycles Based on Monthly and Quarterly Series”, Journal of Applied Econometrics 18, 427-443.

[42] Maroz, D., Stock, J.H. and Watson, M.W. (2021), “Comovement of Economic Activity During the COVID Recession”, *Mimeo.*

[43] McConnell, M.M. and Perez Quiros, G. (2000), “Output Fluctuations in the United States: What has Changed since the Early 1980s?”, American Economic Review 90(5), 1464-1476.

[44] McCracken, M.W. and Ng, S. (2016), “FRED-MD: A Monthly Database for Macroeconomic Research”, Journal of Business and Economic Statistics 34(4), 574-589.

[45] Plagborg-Møller, M., Reichlin, L., Ricco G. and Hasenzagl, T. (2020), “When is Growth at Risk?”, Brookings Papers on Economic Activity, Spring 2020.

[46] Sensier, M. and van Dijk, D. (2004), “Testing for Volatility Changes in US Macroeconomic Time Series”, The Review of Economics and Statistics 86(3), 833-839.

[47] Schorfheide, F. and Song, D. (2021), “Real-Time Forecasting with a (Standard) Mixed-Frequency VAR During a Pandemic, National Bureau of Economic Research, Working Paper 29535.

[48] Stock, J.H. and Watson, M.W. (1989), “New Indexes of Coincident and Leading Economic Indicators”, NBER Macroeconomics Annual, 351-394.

[49] Stock, J.H. and Watson, M.W. (2002), “Has the Business Cycle Changed and Why?”, NBER Macroeconomics Annual, 159-230.

[50] Stock, J.H. and Watson, M.W. (2012), “Disentangling the Channels of the 2007-2009 Recession”, Brookings Papers on Economic Activity, Spring 2012, 81-156.

[51] Summers, L. H. (2015), "Demand-Side Secular Stagnation" American Economic Review, 105 (5), 60-65.