Empirical Economics (2024) 66:1443–1489
https://doi.org/10.1007/s00181-023-02514-7

Holding the economy by the tail: analysis of short- and long-run macroeconomic risks

Michal Franta[1] **· Jan Libich**[2,3]

Received: 29 May 2023 / Accepted: 7 October 2023 / Published online: 3 November 2023

Abstract

We put forward a macro-financial empirical modeling framework that can examine the tails of distributions of macroeconomic variables and the implied risks. It does so without quantile regression, also allowing for non-normal distributions. The framework offers a number of relevant insights into higher moments of the US output growth distribution, as well as the effects of monetary policy and financial (risk premia) shocks on downside macroeconomic risk. This is not only from the short-run perspective but also from the long-run perspective, which has remained largely unexamined in the existing Macro-at-Risk literature. In particular, we estimate the short-run (conditional) and long-run US output growth distributions and study their evolution. The short-run analysis finds that monetary policy and financial shocks render the conditional output growth distribution asymmetric. As such, they affect downside risk over and above their impact on the conditional mean that policymakers routinely focus on. The long-run analysis indicates that US output growth left-tail risk showed a general downward trend in the two decades preceding the Global Financial Crisis (GFC), but this trend got reversed post-2008. Our examination strongly points to the adopted unconventional monetary policy framework featuring quantitative easing as a potential source of elevated long-run downside tail risk in the post-GFC period.

We would like to thank Christian Matthes, Martin Kábrt and seminar participants at the Czech National Bank and at the IAAE 2021 for their comments. The views expressed are those of the authors and do not necessarily represent those of the Czech National Bank. Both authors declare no conflict of interest and compliance with the journal's data availability policy.

✉ Jan Libich
j.libich@latrobe.edu.au

Michal Franta
michal.franta@cnb.cz

1 Czech National Bank, Prague 1, Czechia

2 La Trobe University, Bundoora, VIC, Australia

3 VŠB-TU Ostrava, Ostrava, Czech Republic

Keywords Growth-at-risk · Threshold VAR · Non-normal distribution · US output growth · Macroeconomic policy · Macro-financial modeling

JEL Classification C53 · C54 · E32

1 Introduction

Mark Twain once argued that "if you hold a cat by the tail you learn things you cannot learn any other way." This paper shows that Twain's argument can be extended from cats to economies. Our proposed nonlinear macro-financial empirical modeling framework focuses on distribution tails and empirically examines the distributions' higher moments in order to broaden our understanding of the downside threats to the macroeconomy. The analysis can provide answers to relevant policy questions related to identifying and alleviating macroeconomic risks.

Historically, researchers have focused on the mean (the first moment) when modeling macroeconomic relationships. The period described as the Great Moderation brought the issue of volatility (the second moment) to the fore, but it was generally taken to support the earlier focus on the mean as a useful tool for policymakers. The sentiment started to change with the 2007–2009 Global Financial Crisis (GFC). Since then, academic economists and policymakers have attempted to understand the likelihood and dynamics of extreme scenarios that occur at the tails, far from the mean of the distribution of macroeconomic variables (see, e.g., Cecchetti 2008).

Discussions of the associated risk have been gaining prominence over time. A rapidly growing body of Macro-at-Risk research has gone beyond mean and volatility, empirically investigating specific distribution quantiles in order to shed light on downside (left-tail) risks to economic performance.[1] We attempt to advance this line of work in several directions, offering both methodological and macroeconomic policy contributions.

On the methodological front, we expand the macro-financial empirical modeling framework of Balke (2000) toward simulating predictive distributions. We do not rely on a nonparametric approach based on quantile regressions used in the majority of Macro-at-Risk studies. Instead, our framework makes progress within the newly emerging stream of this literature that simulates predictive distributions using fully parametric models. This is because it tends to provide superior out-of-sample predictive ability for distribution tails of macroeconomic variables compared to quantile-regression-based approaches.[2]

We enrich this literature by employing a mean-adjusted threshold vector autoregression that permits asymmetric predictive distributions with fat tails. This means that

[1] The eminent topic in the Macro-at-Risk literature has been Growth-at-Risk focusing on lower percentiles (usually the 5th or the 10th) of the conditional output growth distribution several quarters ahead (see Adrian et al. 2019). Studies in the Macro-at-Risk literature have also examined inflation (Banerjee et al. 2020) and capital flows (Eguren-Martin et al. 2020).

[2] This body of research has focused on Growth-at-Risk and employed Bayesian vector autoregressions with stochastic volatility (Carriero et al. 2020), Markov-switching models (Caldara et al. 2020), GARCH-type models (Brownlees and Souza 2021), and skewed t-distributions with time-varying parameters (Delle Monache et al. 2020; Plagborg-Møller et al. 2020).

we can formally explore higher moments of the distributions. Let us highlight three aspects of the analysis. First, the distributions are allowed to be multimodal, which the recent literature shows to be important (see Adrian et al. 2021; Kiss et al. 2023).[3] Second, our approach allows for nonlinearities between financial conditions and the real economy. Third, we are able to examine macroeconomic tail risk in a mutually consistent manner from both the short-run perspective, which has been the focus of the literature, and from the long-run perspective.

The long-run viewpoint has been virtually absent in the related literature; we are aware of only one other study that has made a step in this direction so far. Delle Monache et al. (2020) modeled explicitly the trend and cyclical deviations in skewed t-distribution parameters (location, scale, shape and heavy-tailedness). In contrast to their purely statistical approach, we stay within a framework of a standard (reduced-form) multivariate representation of the macroeconomic relationships. While it is still a statistical approach, it allows us to provide an economic narrative of the changes in long-run macroeconomic tail risk, as well as to formulate policy implications. Some medium-run considerations (up to 12 quarters) are present also in Adrian et al. (2022), who discuss intertemporal trade-offs in Growth-at-Risk related to loose financial conditions. Our focus goes beyond the business cycle horizon and can thus provide an additional perspective.

As for the macroeconomic policy contribution, our framework extends the basic question of the Growth-at-Risk literature: "How low could a country's output growth fall in one year's time in the worst 1, 5 or 10% of adverse economic scenarios?" By looking at the distribution moments instead of a particular lower quantile of the conditional distribution, the policymaker can weigh the severity of economic recessions against the opposite prospect of a strong boom in the economy. The policymaker can also prudently evaluate the probabilities of extreme events captured by fat-tailedness, which could be hidden under a particular value of lower quantiles. Furthermore, the long-run perspective offers results over and above the short-term risk fluctuations. It can therefore be used to assess the tail risk arising from longer-term factors such as the structure of the economy, economic trends and alternative policy frameworks.

In order to report both the short-run (conditional) and long-run output growth distribution moments and estimates of US Growth-at-Risk at both horizons, we estimate predictive distributions using 30-year data windows from 1964Q2-1994Q1 to 1989Q2-2019Q1. We then examine the distributional effects of two types of structural shocks. One is financial in nature, namely a change in the excess bond premium. The other is a monetary policy shock.[4] Finally, we investigate trends in downside output growth tail

[3] In line with our arguments for approaches based on predictive distributions, Geweke and Keane (2007) point out that quantile regressions provide only a finite number of quantile estimates, whereas the research question can be related to the entire distribution. Furthermore, there is no guarantee that the estimated quantiles follow the ordering from lower to higher quantiles, implying a potential loss of efficiency in a finite sample due to the so-called quantile crossing.

[4] This aspect of our analysis contributes to the literature on the impact of economic characteristics and shocks on Growth-at-Risk. Existing studies have focused on financial variables (Adrian et al. 2019), monetary policy and financial shocks (Duprey and Ueberfeldt 2020; Loria et al. 2023; Jung and Lee 2019; Kim et al. 2019), as well as on macroprudential policy shocks and indicators (Franta and Gambacorta 2020; Aikman et al. 2019).

risk and the effect of policy frameworks by focusing on changes in model parameters and the role of economic shocks.

Let us summarize the main empirical findings and key policy insights. The short-run perspective demonstrates that downside US output growth tail risk is strongly influenced by economic and policy shocks, and the effect is different from the effect on the conditional means. This is in line with the results reported in the related literature. We extend these findings by showing the effects of distribution moments. In particular, we find that an unexpected monetary policy shock makes the output growth distribution asymmetric in the short term, with a marked increase in the probability of very low growth. We further show that an unexpected easing of financial conditions (a fall in the risk premium) influences the distribution's symmetry toward a longer left tail. In addition, there is an increase in the variance. The implication is that policymakers acting as risk managers should weigh potential future output growth realizations above and below the conditional mean differently. Our results on the asymmetry of the distribution similarly highlight the fact that policymakers need to adjust their decisions relative to scenarios based on the conditional mean alone.

In terms of the long-run perspective, the key finding is that the post-GFC unconventional monetary policy framework featuring large-scale asset purchases has led to a sizeable increase in output growth risk over the long run. This is in contrast to the pre-GFC trend of decreasing output growth risk, which was mainly due to the decline of variance of output growth [for updated empirical evidence on this Great Moderation see Perron and Yamamoto (2022)]. Our analysis demonstrates that this 'declining economic volatility effect' has no longer been operational in the post-GFC period. This is why long-run downside risk stopped declining around 2008, and has been on the rise since 2014.

We uncover changes in short-term macroeconomic dynamics—namely reduced-form monetary policy rules—that underlie this shift in the long-run risk trend. The takeaway for policymakers is that the choice of the policy framework matters for long-run downside macroeconomic tail risk, and that they should not solely focus on the short-run effects. Frameworks that imply an asymmetric distribution should be complemented by policies that can address the resulting longer-term problems.

2 Model for simulation of predictive distributions

Our approach to the estimation of output growth distributions is based on the simulation of predictive distributions.[5] We use a mean-adjusted threshold vector autoregression with the threshold variable potentially dependent on (the lags of) all endogenous variables. The regime switch is thus linked explicitly to the macroeconomic variables, allowing for the regime to be path-dependent. An additional benefit is that the parameters of the predictive distribution of output growth (the mean, variance, skewness

[5] Note that for stationary time series featuring the Markovian property the conditional predictive distribution fully characterizes the probabilistic aspect of the time series.

and kurtosis) can be directly linked to macroeconomic entities. These include regime-specific steady states and regime-specific shock volatilities together with the regime probabilities dependent on macroeconomic variables.

The model is formulated as follows. For some regime r we have

$$y_t = A_1^{(r)} y_{t-1} + \cdots + A_p^{(r)} y_{t-p} + F^{(r)} - A_1^{(r)} F^{(r)} \cdots - A_p^{(r)} F^{(r)} + \varepsilon_t^{(r)}, \qquad (1)$$

where the $n \times 1$ vector y_t denotes the vector of endogenous variables. The coefficient matrices at lag l, denoted $A_l^{(r)}$, are of dimension $n \times n$. The $n \times 1$ vector $F^{(r)}$ contains the unconditional means of endogenous variables. Finally, the vector of error terms $\varepsilon_t^{(r)}$ is distributed independently and normally with a zero mean and a regime-dependent covariance matrix $\Sigma^{(r)}$, i.e., $\varepsilon_t^{(r)} \sim N\big(0, \Sigma^{(r)}\big)$. The vector autoregressions (VARs) commonly used to estimate predictive distributions allow for stochastic volatility [they were introduced by Cogley et al. (2005), and in the context of Growth-at-Risk employed by Carriero et al. (2020)]. In order to examine long-run risk, our framework requires bounded variance. Volatility is therefore allowed to change only across the regimes, not within each regime.

Two regimes are considered. They are determined by the threshold variable y_t^{TR} and the threshold value R. The threshold variable is defined as a weighted average of normalized and smoothed endogenous variables:

$$y_t^{TR} \equiv \sum_{i=1}^{n} w_i NMA^q(y_{it}), \qquad (2)$$

where the weights w_i are estimated. The operator $NMA^q(\cdot)$ denotes the normalized moving average of the last q observations, whereby we set $q = 4$. A smoothing of the threshold variable components allows for a certain persistence of the regimes. It reflects the commonly held view that a single shock should not generally lead to a regime change, i.e., regime changes are better described by gradual underlying processes. To facilitate the interpretation of the estimated weights, smoothed endogenous variables are normalized by first subtracting the respective mean and then dividing them by the standard deviation of the smoothed series.

The threshold value R is also estimated. Regimes 1 and 2 are defined for periods when $y_t^{TR} < R$ and $y_t^{TR} \geq R$, respectively. In order for the threshold and weights to be uniquely determined we impose $0 \leq w_i \leq 1$ and $\sum_{i=1}^{n} w_i = 1$. The definition of the threshold variable generalizes a standard approach within the threshold VARs literature, which employs a particular (lagged) endogenous variable to drive the regime. For instance, if we are interested in credit regimes and feedback loops between credit and the real economy, we would work with the credit spread variable solely in the definition of the threshold variable (similarly to Balke 2000). However, if our focus is on regimes of low and high interest rates or on changes in the interest rate rule, it would solely be the interest rate as the threshold variable. In our analysis of the distribution moments, we take the agnostic approach and let the estimation procedure to pick the (combination) of endogenous variable(s) and weights that maximize the likelihood of data given the priors.

The standard VAR is reformulated into the mean-adjusted form as in (1) to work directly with parameters interpretable as steady states. For a stationary process it follows from (1) that in regime r we have

$$E(y_t) = F^{(r)}. \tag{3}$$

The mean-adjustment disciplines the forecasts in the medium run and long run. Such forecasts are usually very poor when based solely on standard VARs. Furthermore, the mean-adjustment implies that both the short-run and the long-run perspectives draw on the same model, being mutually consistent.

3 Data and estimation

The data are quarterly and cover the period 1964Q2-2019Q1. The vector of endogenous variables y_t consists of one financial and three macroeconomic variables (all series have been downloaded from the FRED database). The latter consist of real GDP growth (annualized quarter-on-quarter change in real GDP), consumer price inflation (annualized quarter-on-quarter change in the CPI), and the short-term interest rate (the 3-month interbank rate). They capture the key macroeconomic features of a closed economy. The financial sector is represented by the spread between the 10-year bond yield and the Federal Funds rate. It has been included because the slope of the yield curve is often a predictor of future economic activity, see, e.g., Estrella and Trubin (2006). In addition, the spread (together with the short-term interest rate) can account for post-2008 unconventional monetary policy measures implemented through large-scale asset purchases, see for example Baumeister and Benati (2013).[6]

The model in (1) is estimated via the Gibbs sampler. The estimation procedure follows a standard Bayesian estimation of threshold vector autoregressions (Chen and Lee 1995), extended for the estimation of weights in the threshold variable (Chen and So 2006) and mean-adjustment (Villani 2009). Threshold VARs were first introduced into macro-financial modeling by Balke (2000). One difference with respect to the standard approach involves sampling the threshold R and weights w_i, i.e., the parameters whose conditional posteriors are not analytically tractable. The key is to sample from the multinomial distribution instead of employing the usual Metropolis step. We exploit the fact that the domains of the threshold value, as well as the weights, are bounded and discretize the domain. Then all possible values of the parameters can be assessed by their respective likelihood and a multinomial distribution can be employed. Details of the estimation procedure and informative priors are discussed in Appendices A.1–A.3.

[6] Note that the spread enters our vector of endogenous variables as its negative – the reasoning for this normalization can be found in Appendix A.5. The ability to capture unconventional measures results in our preference for the spread even though other suitable predictors of recessions have also been explored in the literature (e.g., Hwang 2019, Borio et al. 2020).

The model is estimated on windows of 30 years' length. The first estimation window covers the period 1964Q2-1994Q1, and the last estimation window covers the 1989Q2-2019Q1 period. A fixed window length is essential because parameter uncertainty is part of the simulated predictive distributions, hence changing the window length could spuriously suggest trends in output growth tail risks. For inference, we use 40,000 iterations of the Gibbs sampler that are generated after 10,000 'burn-in' iterations. Convergence diagnostics are presented in Appendix A.4. The out-of-sample validation of the model can be found in Appendix B.5, where the predictive distribution performance is compared with a standard (constant-parameter) VAR model and nonlinearity tests are reported.

The recursive estimation serves to shed some light on the changes in model parameters over time. An alternative approach would be to use a modeling framework, which directly allows time-variation in model parameters. In the context of predictive distribution estimations, VARs with stochastic volatility (Carriero et al. 2020), Markov-switching models (Caldara et al. 2020) and time-varying parameter VARs could be employed.

While the primary purpose of the model is to estimate predictive distributions of output growth and the implied tail-risk measures, the estimation results themselves offer several interesting insights into the evolution of US macroeconomic dynamics. Moreover, they provide an empirical validation of our model in terms of its in-sample properties, and help us to interpret changes in long-run tail risk over time. The estimation results of the model in (1), as well as their detailed discussion and comparison with the existing literature, can be found in Appendix B.1.

During the Gibbs sampler procedure, each drawn vector of the model parameters is used to simulate the model forward. For each vector of the model parameters drawn, one simulated path of endogenous variables is obtained. All such simulated paths yield an estimate of the predictive distributions at various horizons. Finally, we assume that our model describes a stable system, i.e., explosive paths of macroeconomic variables are not considered when the model parameters are sampled and predictive distributions simulated. The stability of the system is important if the long-run output growth distribution and tail risk are to be examined. Appendix A.3 provides further details.

The simulated predictive distributions are then used to estimate distribution moments and measures of output growth tail risk. The third and fourth moments are represented by the measures of skewness and kurtosis less sensitive to outliers; as introduced in Groeneveld and Meeden (1984) and Hogg (1974), respectively. Regarding the output growth tail risk we focus on downside tail risk as represented by the 10th percentile of the output growth distribution.[7] As usual in the literature, the short-term measure of macroeconomic risk is conditional on data in period $t-4$: the 10th percentiles are taken from the predictive output growth distribution four periods ahead. More specifically, we take all simulated paths of output growth and look at the 10th percentile at a distance of four quarters.

[7] As a robustness check, a comparison of the 1st and the 10th percentile of the output growth distribution is reported in Appendix C.4. It shows that the specific choice of the percentile to summarize the tail of the output growth distribution is not crucial for our findings.

Our second indicator of macroeconomic risk is linked to the long-run output growth distribution. It can be considered unconditional because we utilize all simulated paths of output growth from the simulation of predictive distribution starting with the 48th quarter to eliminate initial conditions. The length of the simulated path is 1000 periods. We take each 20th simulated value to get independent draws from the long-run output growth distribution, and use the 10th percentile from the resulting distribution. In its construction, the indicator is similar to existing measures of likelihood of zero lower bound events, which are also based on long-run model simulations (e.g., Chung et al. 2012).

Such long-run risk measure is not directly related to a specific sequence of shocks occurring during the estimation period of the underlying model. Instead, it encompasses long-term risks related to both the structure of the economy and the policy framework prevailing at the time. In terms of the model, it captures all estimated macroeconomic relationships, shock volatilities and steady states in both regimes. It should be noted that the zero lower bound (ZLB) is imposed during the simulation, i.e., a zero value is forced whenever the simulated interest rate path tends to get into negative territory. In Appendix C.3, we perform a formal check in which the ZLB is not imposed, and show that the selected approach of accounting for the ZLB in the simulation of predictive distributions does not affect our findings.

4 Estimates of risk measures

Having obtained the estimates of the output growth distribution, which are discussed in Appendix B.2, we can focus on their 10th percentile that serves as a common left-tail risk measure. Naturally, whenever the 10th percentile falls downside risk rises, and vice versa. Figure 1 presents two variants of the estimated percentile, representing downside risk in the short run (the red dashed line) and in the long run (the blue dotted line). As a reference point, ex post observed output growth is also plotted in Fig. 1 (the black solid line).

In terms of the short-run perspective, it is reaffirming that the observed output growth falls below the 10th percentile during recessions (in 2001 and 2007–2009). This means that a realization of output growth from the left tail of the conditional distribution occurred.

In terms of long-run risk some trends can be observed. The initial two thirds of the estimation windows until the GFC can be broadly characterized by increasing 10th percentile of the long-run output growth distribution, i.e., long-run downside tail risk was trending downward (the blue line). This trend stopped when the data from the GFC period enter the estimation windows. Importantly, the trend got reversed around 2014, after which long-run risk started to increase steadily. This reversal accounts approximately for 0.5 p.p. reduction in the 10th percentile of long-run output growth.[8]

Figure 1 emphasizes the importance of studying macroeconomic risk both from the short-run and the long-run perspectives. The long-run and short-run trend of the 10th

[8] Note that in terms of the confidence bands around the smoothed trend of long-run risk, the trend reversal is statistically significant. The confidence bands are very narrow due to the fact that the long-run output growth distribution is simulated using many periods.

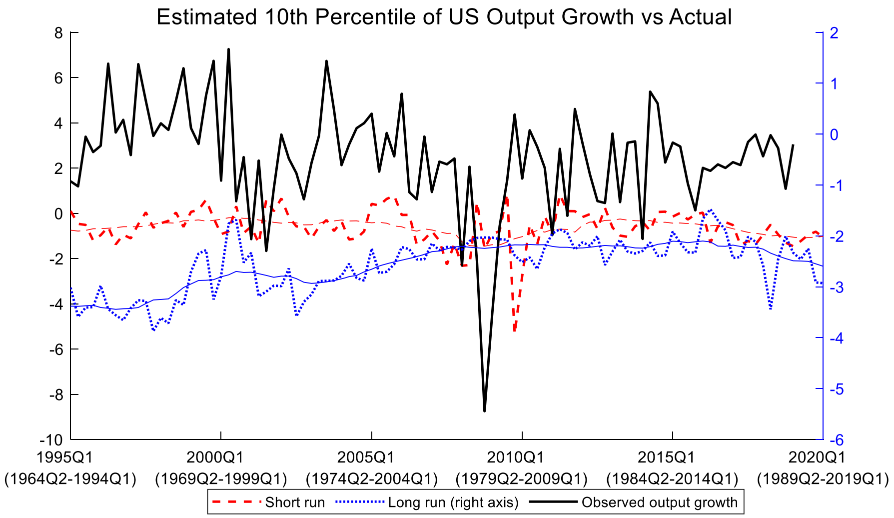

Fig. 1 Estimates of the 10th percentile of the short-run (four-quarters-ahead) and the long-run US output growth distribution, together with their underlying trends and ex post observed output growth. Real output growth is in the form of annualized quarter-on-quarter percentage changes. The x-axis indicates the quarter for which the 10th percentile of the short-run output growth distribution is estimated and ex post output growth observed. The periods in parenthesis indicate the windows used to estimate the long-run output growth percentile. Moving averages of the respective 10th percentiles are reported using windows of 16 quarters and truncating the size of the window at the beginning and end of the sample from left and right, respectively

percentile of output growth differ. Most notably, in the first half of the sample the 10th percentile of the long-run output growth distribution was generally increasing, whereas its short-run analog remained largely unchanged. In addition, the 10th percentiles of the conditional and long-run distributions are uncorrelated, i.e., short-run developments do not systematically relate to the long-run distribution.

5 Analysis of short-run risk

We start with the short-run risk analysis by estimating the effects of monetary policy and financial shocks on the moments of the conditional output growth distribution. The results are then interpreted with respect to the monetary policy framework and long-run macroeconomic risk.

The analysis is built on the following linear regression, estimated by OLS:

$$\mu^{(m)}_{t+h+4|t+h} = \alpha^{(m)}_h + \beta^{(m)}_h Shock_t + \gamma^{(m)}_h \mu^{(m)}_{t+4|t} + \varepsilon^{(m)}_t,\ h = 0, \ldots, H \qquad (4)$$

where m denotes the respective moment, i.e., $\mu^{(m)}_{t+h+4|t+h}$ expresses for the mean ($m =$ 1), variance ($m = 2$), skewness ($m = 3$) and kurtosis ($m = 4$) in period $t + h + 4$. These moments are estimated in the first stage by employing the model in (1) and using the data available in period $t + h$ (conditional on the information set Ω_{t+h}).

Even though we assume in model (1) that the shock propagation can differ across the regimes and is thus history-dependent, and the regimes can also differ across the estimation windows, here we take the estimated conditional distribution as based purely on a statistical model. Specifically, we take the conditional distribution as coming from a mixture of two normals that best fits the data (together with the priors), assuming a constant effect of the shock over time. The estimated effect can then be interpreted as the average effect over historical episodes and also over regimes.[9] As such Eq. (4) can provide useful information on the effect of structural shocks on short-run risk rather than directly on effects of structural shocks on a particular macroeconomic variable. Finally, note that to account for the fact that the moments are estimated quantities, the confidence intervals for the model parameters are bootstrapped; the procedure is described in Appendix B.4.

The variable $Shock_t$ denotes the exogenous structural disturbance taken from the existing literature. The identification is done in one of three distinct ways. The first is a narrative identification approach (monetary policy shocks in the baseline analysis). The second is high frequency identification (monetary policy shocks used in a robustness check). The third employs the timing restrictions imposed on indices derived from micro-level data (excess bond premium shocks).[10] Similarly to moments the employed structural shocks in (4) are also estimated quantities. Unfortunately, the confidence bands of shock estimates are not available, which could in principle result in biased estimates in (4).

Note that the linear regression in (4) is reminiscent of the local projections introduced in Jordà (2005), and often applied in the analysis of the effect of shocks on distribution quantiles. Local projections represent a way of estimating the impulse responses if a measure of exogenous variation is available. Importantly, local projections allow for a simple examination of nonlinear specifications and asymmetric effects of shocks, which is the focus of our analysis. For example, variance could be affected similarly by positive and negative shocks; both types of shocks could increase volatility. Thus a linear specification could provide insignificant results, whereas a specification with shocks separated according to their direction would provide valuable insights.

The Growth-at-Risk literature has typically studied the effect of a shock on a particular quantile. We focus directly on the effects of the distribution moments, and are thus able to interpret changes in risk from the perspective relevant to the policymaker.[11] As an example, consider a frequently analyzed monetary policy tightening scenario (e.g., Duprey and Ueberfeldt 2020; Loria et al. 2023; Jung and Lee 2019). Research

[9] Note that similar approach is used in Loria et al. (2023). They employ local projections on quantiles estimated by quantile regression to examine the effect of various structural shocks. Interestingly, they conclude that a threshold VAR is able to match their findings.

[10] If sufficiently long series of shocks were available, an alternative approach would be to incorporate a shock measure directly into the threshold VAR from the first stage of our analysis (along the lines of proxy VARs). However, such an approach is associated with several methodological questions that are yet to be resolved in the literature. For example, it is not clear how the equivalence between local projections and impulse responses from proxy VARs (Plagborg-Møller and Wolf 2021) would apply for nonlinear models and for quantile responses instead of standard conditional mean responses.

[11] Another advantage of the focus on moments rather than quantiles is that it allows us to report statistical significance of the difference between the conditional mean effect and the tail response. This difference is generally not tested statistically in the literature.

studies commonly find that tighter monetary conditions shift the left tail of the output growth distribution leftward; in a way that exceeds the shift of the mean. However, this could be a consequence of several distinct factors: (i) increased conditional variance, (ii) a more negative conditional skewness, or (iii) fatter tails of the conditional output growth distribution.

Finding out which of these three explanations applies is of great relevance to policymakers. In their decision making, changes in volatility are viewed, and responded to, differently from changes in the asymmetry of the distribution. Each explanation (i)–(iii) above would imply a different probability that in a period of realization from the left tail of the output growth distribution there will be an offsetting positive realization from the distribution's right tail.

5.1 Monetary policy shocks

Given the importance of the monetary policy shock for the dynamics of US macroeconomic variables (see Liu et al. 2019), we make it our starting point. Monetary policy shocks are taken from Wieland and Yang (2020), who extend the series of narratively identified shocks by Romer and Romer (2004) until 2007Q4. Given that the first estimate of the moments from the model in (1) is available for 1994Q2, the regression in (4) is estimated for the period 1994Q2-2007Q4. A robustness check with shocks identified by high-frequency techniques, as well as shocks covering longer periods with both conventional and unconventional monetary policies, can be found in Appendix C.1.

In regard to the conditional mean, the effect of a monetary policy shock as depicted in panel a) of Fig. 2 is consistent with the literature in terms of both its direction and profile. There is a fall in output growth with a trough after a year and half, and a subsequent compensation in the form of higher output growth. Skewness decreases within a year after the shock and then again at the horizon of four years. This means that short-run left-tail risk (the probability of very low output growth) temporarily increases and then does so in the medium term again. The change in the asymmetry of the distribution is statistically significant implying that the left tail temporarily diverts from the conditional mean (and from the right tail too). This effect is economically significant, which can be seen from the estimated skewness of the conditional output growth distribution reported in Appendix B.2.

The observed asymmetry of the output growth distribution after the monetary policy shock provides a message to risk-managing monetary policymakers.[12] Depending on the specific form of the policymakers' loss function, broadening the focus from the mean to the entire output growth distribution implies different monetary policy actions. Regarding the even moments (variance and kurtosis), the specification in (4) with separated easing and tightening shocks reported in Appendix C.1 implies that both the tightening shock and (even more so) the easing shock result in a decrease of kurtosis in

[12] Alessi et al. (2014) describe the risk-management approach to monetary policy at the Federal Reserve Bank of New York. Kilian and Manganelli (2008) show what a central banker's loss function can look like, and how it can be linked to macroeconomic risk represented by (a subset of) distributions of macroeconomic variables. A simple general equilibrium framework where the policymaker considers GDP tail risk can be found in, e.g., Duprey and Ueberfeldt (2020).

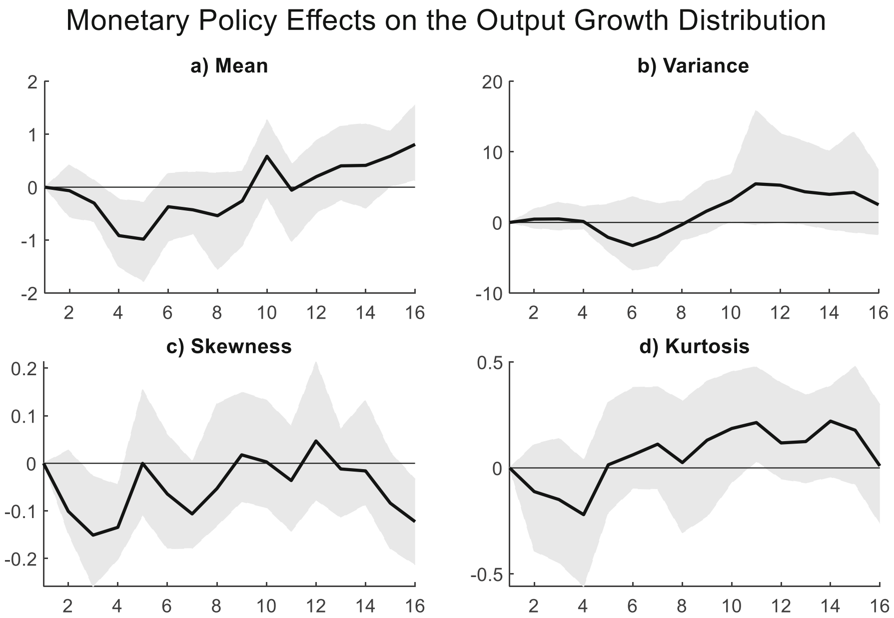

Fig. 2 Response of the conditional mean, variance, skewness and kurtosis of the four-quarters-ahead output growth distribution to an unexpected 100-basis-points monetary policy shock. The shaded areas indicate the 90% confidence intervals

the short run. The benchmark specification implies only a minor impact as the effect is not symmetric. This points to a potentially important policy message that runs counter the arguments for interest rate smoothing. If monetary policy reacts immediately to shocks without excessive smoothing, fat-tailedness of the output growth distribution is reduced. This is beneficial as it reduces the probability of periods featuring very low output growth.

The literature usually finds that the monetary policy shock has a more profound effect on the left tail of the output growth distribution represented by the 5th or 10th percentiles in comparison with its effect on the conditional mean/median (Duprey and Ueberfeldt 2020; Loria et al. 2023; Jung and Lee 2019). Our analysis of the moment responses reveals that the temporary increase in output growth tail risk after a monetary policy shock is due to both the conditional mean effect and an increased asymmetry of the output growth distribution.

5.2 Financial shocks

The second type of shock employed in our analysis is financial in nature, represented by an unexpected change in the excess bond premium. We start with the credit spread index from Gilchrist and Zakrajšek (2012), regularly updated to cover the period 1973Q1-2019Q1 (for details see Favara et al. 2016). In contrast to the narratively identified monetary policy shocks described in the previous section, excess bond premium changes cannot be viewed as a source of exogenous variation. We follow Loria et al. (2023) and extend the regression equation in (4) to the contemporaneous Federal Funds rate, contemporaneous output growth, contemporaneous inflation and four lags of the excess bond premium in order to control for the premium's endogenous movement. We estimate this extended model for the 1994Q2-2019Q1 period.

Figure 3 presents separately the effects of a tightening of financial conditions (an increase in the risk premium) and financial easing (a decrease in the risk premium). The asymmetric effects represent a stylized fact caused by strong nonlinearities between the financial conditions and the real economy (Akinci and Queralto 2017).

Figure 3 shows an increase in the conditional variance primarily after an unexpected easing of financial conditions, implying elevated short-run tail risk. This second moment effect is further strengthened by the shift of conditional skewness toward a more negative skew, which is economically significant and more profound under financial easing. The economic mechanisms behind the rise in short-run risk after easing of financial conditions includes, for example, accumulation of risk in the financial sector due to an increase in risk-taking.[13]

Comparing the effects of monetary policy and financial shocks suggests an important difference between conventional and unconventional monetary policy measures. The latter, based on quantitative or credit easing (i.e., lowering of credit spread taking the form of a financial shock), are shown to weaken the resilience of the economy in the short-term—conditional skewness decreases. The opposite is true when easing is delivered by conventional monetary policy actions.

Our estimates go hand in hand with the findings of the literature, providing a new context for them. Adrian et al. (2019) found that deterioration of financial conditions represented by the National Financial Conditions Index relates to a fall in the conditional mean and an increase in the conditional volatility. This is consistent with a fall in the lower quantiles of the conditional output growth distribution while the upper quantiles remain stable. Our analysis based on effects of financial shocks, not general index, shows an important role of changes in skewness.

Next, Aikman et al. (2019) find that credit booms increase downside risk over the 3 to 5-year horizon. Considering our estimated impact of loosening financial conditions in Fig. 3 and assuming a link between a fall in the credit spread and a credit boom, our analysis implies that the increase in downside risk is due to two factors. One is an increase in volatility and the other is greater negative skewness. It thus follows that the fall in Growth-at-Risk is not fully compensated on the upside. This provides an empirical argument for policymakers to impose more stringent financial market regulatory measures during exuberant credit booms.

Finally, Loria et al. (2023) report that one year after a credit spread shock the conditional distribution narrows, i.e., the 10th and 90th quantiles get closer to each other. Our analysis offers an explanation for this finding. It is likely to be a consequence of the fall in conditional volatility, which arises due to aggregation of the effects of

[13] It is worth noting the counter-intuitive profile of the conditional mean effect in the short run in Fig. 4. The increase of the mean could be due to the reverse causality issue not sufficiently accounted for by the extended regression (4). Our argument is, however, based on the effects on higher moments of the conditional distribution.

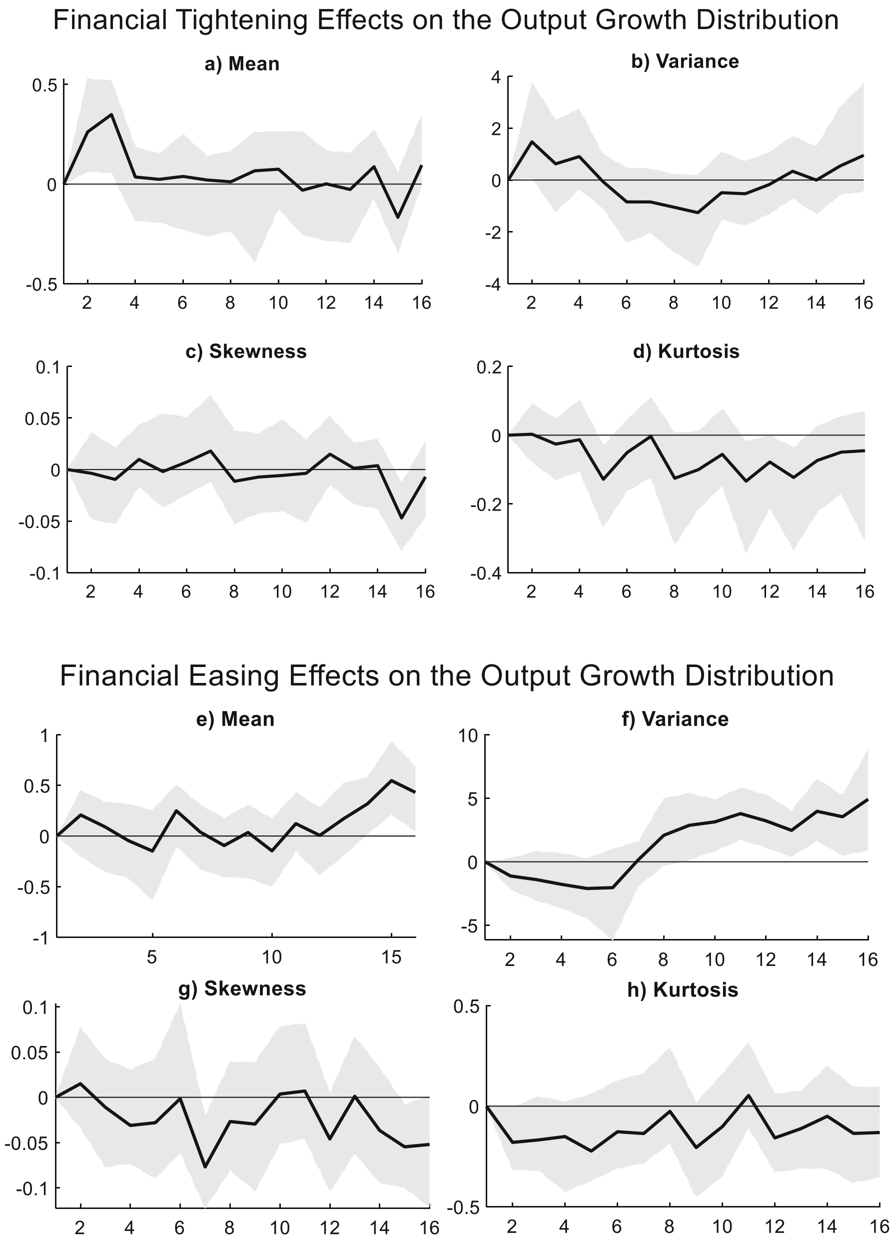

Fig. 3 Response of the conditional mean, variance, skewness and kurtosis of the four-quarters-ahead output growth distribution to an unexpected 100-basis-points change in the excess bond premium. **a**–**d** Show a tightening of financial conditions, **e**–**h** show their easing. The shaded areas indicate the 90% confidence intervals. The separate policy shocks are employed in a way that if only tightening shocks are considered unexpected easing is represented by 0, and vice versa

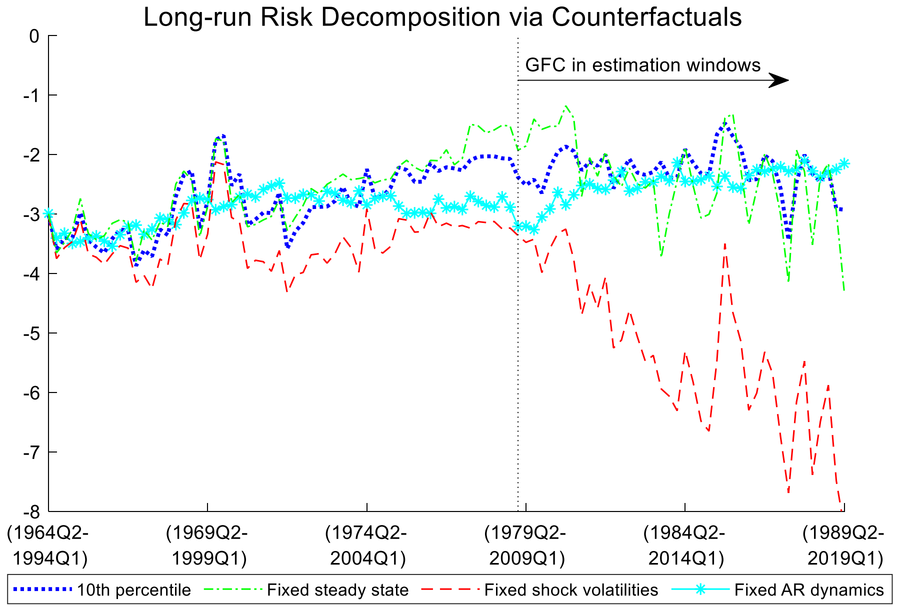

Fig. 4 The 10th percentile of the long-run output growth distribution and the counterfactuals with a fixed steady state, fixed shock volatilities and fixed AR parameters, respectively. The *x*-axis indicates the period used to estimate the percentile in the long run. The dotted vertical line indicates the estimation window in which the GFC enters the estimation (1979Q1-2008Q4)

credit tightening and easing within one linear specification. In this specification credit easing enters with a minus sign, whereas tightening enters with a positive sign. Given that the effect on easing on volatility dominates, a fall in volatility is observed in Loria et al. (2023). The takeaway from this aspect of our analysis is that distinguishing between credit easing and credit tightening is essential for designing appropriate risk-management policies.

5.3 Shocks and long-run risk

In terms of the monetary policy shock, unexpected policy easing and tightening have the same short-run distributional effects with an opposite sign (with the exception of kurtosis). The nature of the financial shock is however different. Credit spread easing results in a negatively skewed distribution, and, in addition, leads to a higher variance while tightening implies an insignificant effect both statistically and economically.

It follows that policy frameworks based predominantly on the interest rate as an instrument tend to imply a symmetric long-run output growth distribution, whereas those based on the credit spread lead to an asymmetric and more volatile distribution. These findings are highly relevant for the post-2008 period. In the aftermath of the GFC the monetary policy framework has been largely based on the credit spread. Quantitative easing influences credit spreads through the portfolio rebalancing channel, whereby its direct impact on yield spreads spills over to close substitutes of safe assets bought and sold by the central bank. Our findings therefore imply that the post-GFC policy in the US has made the long-run output growth distribution asymmetric and increased its variance.[14] This argument draws on structural shocks and is thus more informative in regard of nonsystematic component of the policies. A supplementary argument based on the systematic component of policies is presented in the next section.

6 Analysis of long-run risk

In this section, we examine directly the profile of the 10th percentile of the long-run output growth distribution and the implied long-run downside tail risk. In contrast to the above short-term perspective, which related to shocks and specific policy actions within a given policy framework, the long-run insights relate to the choice of the policy framework as such. The argument here draws strictly on the model in (1) and should thus be interpreted as being of a reduced-form nature. It is based on estimated (reduced-form) policy rules, so it relates to the systematic component of the observed policies.

The first step in analyzing the 10th percentile of the long-run output growth distribution is to construct counterfactuals that help uncover the role of macro dynamics, shock volatility and steady states, respectively. We focus on the 10th percentile rather than the distribution moments because the exposition is more straightforward, but a similar exercise is also conducted for the moments of the output growth distribution (see Appendix B.3).

The counterfactuals are constructed by keeping various subsets of the whole parameter vector in model (1) constant, one at a time.[15] Figure 4 plots the three counterfactuals we examine: fixed steady states, fixed reduced-form shock volatilities and fixed autoregressive parameters. For example, the first counterfactual shows what the 10th percentile would be if the steady states in both regimes were kept unchanged over time. The second and third counterfactuals conduct the same exercise, fixing the regime-specific shock volatilities and the regime-specific short-run dynamics, respectively. Comparing the counterfactual with the factual then indicates whether the specific feature (subset of parameters) contributed to the observed profile of tail risk, and in what way.[16]

[14] There is evidence supporting our argument that the nonsystematic component of monetary policy based on the interest rate implies the dominance of interest rate shocks hitting the economy, whereas monetary policy based on the credit spread results in the dominance of credit spread shocks. Most notably, focusing on the US Liu et al. (2019) found the late 1990s and mid-2000s as periods with the interest rate shocks dominating, whereas the post-GFC period featured predominantly yield spread shocks.

[15] The counterfactuals draw on a reduced-form model, and hence the regime-specific error covariance matrices include the contemporaneous effects of structural shocks on endogenous variables. As such, imposing reduced-form shock volatilities imposes part of the short-run dynamics. In what follows, short-run dynamics thus means reduced-form AR parameters, and shock volatilities mean reduced-form shock volatilities.

[16] More specifically, for a given estimation window we use the respective output from the Gibbs sampler, i.e., the draws of all subsets of model parameters (except the one that is fixed over all estimation windows whereby the first estimation window is used). Then, we simulate the model again 1,048 periods ahead to

Figure 4 implies that the increasing trend in the 10th percentile, i.e., the decline in downside risk in the pre-GFC period (the first half of the estimation windows) is mainly due to shock volatilities (the red dashed line). Intuitively, keeping shock volatilities unchanged would lead to a lower and more stable 10th percentile, i.e., higher downside risk. Moreover, the role of shock volatilities increases over time as the counterfactual diverts from the actual 10th percentile more profoundly over time. The impact of volatility on risk is partially offset by a fall in the steady state of output growth (see Appendix B.1). This is especially the case in the middle of the series of estimation windows (the green dot-dashed line, which shows how the 10th percentile would evolve if the steady state was fixed).

Focusing on the post-GFC evolution in the 10th percentile, Fig. 4 suggests that its switch from a rising to a falling trend is due to short-term macroeconomic dynamics. The counterfactual with fixed AR parameters (the blue line with stars) exhibits an increase in the 10th percentile, while fixed shock volatilities and the steady states result in a fall of the 10th percentile.

To shed some light on the nature of the change in the short-run dynamics after 2008, Fig. 5 presents the evolution of the AR coefficients averaged over the regimes in the interest rate equation and in the spread equation of the model in (1).[17] We focus on the two equations because the interest rate served as an instrument of conventional monetary policy, whereas the spread served as an unconventional monetary policy instrument after the interest rate hit the zero lower bound.

Figure 5 shows that remarkable qualitative changes in the estimated coefficients occurred in the post-GFC period coinciding with a shift from a declining to a rising trend in long-term tail risk to its rise. The change in coefficients exhibits some statistical significance, especially for the AR coefficients related to output growth (panel a) and inflation (panel b) in the interest rate equation.

The changes can be interpreted as a transition to a different monetary policy regime. After the GFC enters the estimation windows (2008Q4) the monetary policy regime can be characterized by the monetary authority placing a low weight on both output growth and inflation within its interest rate policy rule. This reflects the ZLB. The post-2014Q1 that witnessed the rebound of the interest rate from zero can be characterized by the monetary authority assigning a low weight on inflation in combination with a nontrivial weight on output growth (in fact the weight on inflation was zero for the whole post-2008 period). This contrasts the pre-GFC windows featuring a high weight on inflation with a low weight on output growth, and hence means a significant change in monetary policy objectives.[18]

Footnote 16 continued
obtain the long-run risk measure as in the benchmark case. This procedure is conducted for all estimation windows.

[17] AR coefficients in other equations of model (1) together with regime-specific AR coefficients are presented and discussed in Appendix B.1.

[18] The changes of our AR coefficients are consistent with the regimes of monetary policy from Bianchi (2013). His estimates based on a micro-funded model indicate that the 1960s, the 1970s and the GFC period can be characterized as the "Dove" regime, involving an interest rate policy rule with a low weight on inflation, whereas the rest of the periods can be described as the "Hawk" regime with a high weight on inflation. Note that Bianchi worked with a structural model that included expectations of a regime change. Our results should hence be interpreted as an alternative (reduced-form) evidence supporting his findings.

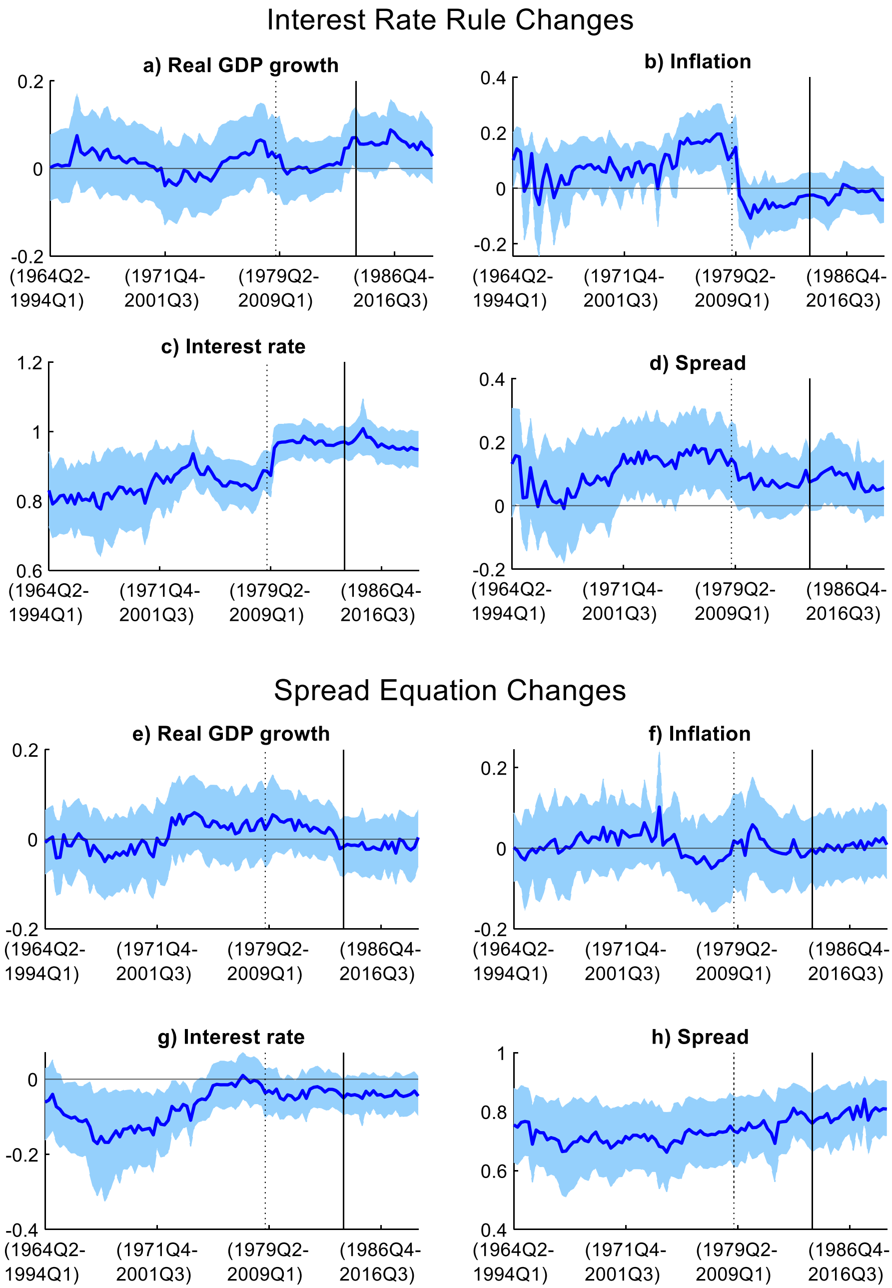

Fig. 5 The weighted sum of the estimated coefficients at lags of the respective variable in the interest rate equation and in the spread equation. We report the weighted sum over the two regimes estimated in model (1) with weights equal to the probability of the regime in the steady state together with 68% confidence intervals. The x-axis indicates the range of quarters of the estimation window. The dotted vertical line indicates the estimation window in which the GFC first enters the estimation (1979Q1-2008Q4). The solid vertical line indicates the estimation window 1984Q2-2014Q1, for which the risk profile changes its direction

The relationship between the spread and lagged output growth also undergoes a qualitative change. It turns from positive to zero, whereby the switch coincides with the 1984Q2-2014Q1 estimation window, but the change is not statistically significant. Post-GFC quantitative easing was conducted in the form of credit spread shocks rather than in the form of a systematic policy reaction to inflation and output growth. Therefore, statistical insignificance of AR parameters is observed.

The changes in the parameters of model (1) suggest that the post-GFC period represents a major departure from earlier policy frameworks. The departure reflects the change in the monetary policy instrument from the interest rate to the interest rate spread, and very different weights policymakers assigned to output and inflation in their objective function. Our analysis demonstrates that such departure from past practice has been accompanied by an increase in downside macroeconomic risk.

7 Summary and conclusions

This paper proposes an econometric approach to estimating and analyzing output growth tail risks. It can provide valuable insights into past macroeconomic developments as well as into efficacy of various policy measures and frameworks. Our exploration is divided into two stages. In the first stage, we put forward an empirical framework to estimate the evolution of the entire US output growth distribution. The framework allows for asymmetry and fat tails of the distribution, i.e., it relaxes the conventional Gaussian assumption. Importantly, the framework features a mechanism that disciplines its long-run behavior. Therefore, the long-run output growth distribution and risks can be analyzed over and above the standard short-run tail risk measures built on quantiles of the conditional output growth distribution. Adding the long-run perspective—in a way that is consistent with the short-run perspective—is one of the methodological contributions of this paper.

The second stage of the analysis consists of an empirical examination of the derived US output growth distributions. We exploit the fact that the entire distributions are simulated and investigate the effect of monetary policy and financial shocks on the moments of the conditional output growth distribution. This extends the literature that has generally only considered a handful of selected quantiles of the distribution. Being able to pin down which distribution moment is behind the effect of a certain shock on various quantiles of the output growth distribution (most importantly its left tail) can help policymakers to better tailor their actions.

We then take advantage of our nonlinear model from the first stage to interpret the changes in the simulated distributions from a long-run perspective. It is not conditional on any short-term shocks and thus encapsulates the trends and structural developments. In particular, we examine the time profile of the long-run output growth distribution, and the long-run output growth tail risk, by decomposing the distinct effects of three types of changes. They are changes in the steady states of macroeconomic variables, changes in their shock volatilities, and changes in short-run macroeconomic dynamics. Such decomposition sheds light on the general role of alternative economic policy frameworks and structural factors for output growth tail risk.

The main broad message is that the entire distribution of key macroeconomic variables should be taken into account, because various policy and economic/financial shocks affect the conditional mean differently from the tails. Furthermore, both the short- and long-run perspectives on downside macroeconomic risk need to be explored, as both offer valuable (and sometimes conflicting) information. Most relevantly, our analysis shows that US monetary policy responses during the 2007–2009 period, and the subsequent policy normalization between 2014 and 2019, represented important regime changes that can be linked to a rise in long-run output growth tail risk over the past decade.

While our focus was on monetary policy, let us note that macroprudential policy measures could, in principle, also adversely affect the left-tail risk of output growth (over and above the conditional mean). This is because macroprudential changes in bank capital requirements or risk weights work through credit spreads, and we have demonstrated above that these can manifest through an increase in conditional variance or more negative conditional skewness.

However, two relevant issues must be stressed. First, the linkage between our financial shocks (credit spread) and macroprudential policy is not two-sided. It is apparent that changes in the macroprudential policy setting affect the excess bond premium, but not all unexpected changes in this premium can be attributed to macroprudential policy measures. For instance, the rise in the bond premium around 2008 was undoubtedly driven by financial market turmoil, not macroprudential policy actions at the time. Second, the transmission via credit spreads only captures one of the potential channels of macroprudential policy, so the overall effect of the policy cannot be unambiguously established. These two reasons imply that more evidence would be required to make the claim that macroprudential policy easing may have possible negative consequences for downside short-run output growth risk.

Funding Open Access funding enabled and organized by CAUL and its Member Institutions.

Appendix A: Bayesian estimation of the TVAR model

This appendix discusses the estimation procedure of the model introduced in Sect. 2. We present the model in the vector form (Sect. A.1), describe the likelihood and priors (Sect. A.2), and examine the resulting conditional posteriors and the Gibbs sampler

(Sect. A.3). Finally, we discuss the convergence of the sampler and correlation of endogenous variables in the threshold variable (Sects. A.4, A.5).

Model (A.1)

The model in (1) can be reformulated into a vector form as follows:

$$\overline{y}^{(r)} = \overline{X}^{(r)}\beta^{(r)} + \overline{Z}^{(r)}\delta^{(r)} + \varepsilon^{(r)} \quad \text{for} \quad r = 1, 2, \tag{A1}$$

where $\varepsilon \sim N\left(0, \overline{\Sigma}^{(r)}\right)$. Defining $Y^{(r)}$ as a $t^{(r)} \times n$ matrix of observations related to regime r, then the $nt^{(r)} \times 1$ vector $\overline{y}^{(r)}$ is defined as $\overline{y}^{(r)} \equiv vec\left(Y^{(r)}\right)$. Next, defining $X^{(r)}$ as a $t^{(r)} \times np$ regime-specific matrix that comprises up to p lags of the vector of endogenous variables, the $nt^{(r)} \times n^2 p$ matrix $\overline{X}^{(r)}$ is defined by the Kronecker product $\overline{X}^{(r)} \equiv I_n \otimes X^{(r)}$. The $n^2 p \times 1$ vector $\beta^{(r)}$ is defined as follows:

$$\beta^{(r)} = vec\begin{pmatrix} \left(A_1^{(r)}\right)' \\ \vdots \\ \left(A_p^{(r)}\right)' \end{pmatrix}. \tag{A2}$$

The $nt^{(r)} \times n(p+1)$ matrix $\overline{Z}^{(r)}$ equals to $I_n \otimes Z^{(r)}$, where $Z^{(r)}$ is a $t^{(r)} \times (p+1)$ matrix of the following constant terms:

$$Z^{(r)} = \begin{bmatrix} 1 & -1 & \cdots & -1 \\ 1 & -1 & \cdots & -1 \\ \vdots & \vdots & \ddots & \vdots \\ 1 & -1 & \cdots & -1 \end{bmatrix}. \tag{A3}$$

Finally,

$$\delta^{(r)} = vec\begin{pmatrix} F^{(r)'} \\ F^{(r)'}\left(A_1^{(r)}\right)\prime \\ \vdots \\ F^{(r)'}\left(A_p^{(r)}\right)\prime \end{pmatrix} \tag{A4}$$

is an $n(p+1) \times 1$ vector including the steady states of the endogenous variables. The vector of the model parameters is then the following:

$$\boldsymbol{\theta} \equiv \left[\beta^{(1)}, \beta^{(2)}, F^{(1)}, F^{(2)}, R, w, \overline{\Sigma}^{(1)}, \overline{\Sigma}^{(2)}\right].$$

Likelihood and Priors (A.2)

The estimation procedure draws on the Bayesian approach, i.e., a posterior distribution is formulated based on the likelihood of data and a joint prior distribution. The likelihood function of the model in (A1) is:

$$\pi(\overline{y}|\theta) = \left|\overline{\Sigma}^{(1)}\right|^{-\frac{1}{2}}\left|\overline{\Sigma}^{(2)}\right|^{-\frac{1}{2}}$$
$$* \, exp\left\{-\frac{1}{2}\sum_{r=1}^{2}\left(\overline{y}^{(r)} - \overline{X}^{(r)}\beta^{(r)} - \overline{Z}^{(r)}\delta^{(r)}\right)\left(\overline{\Sigma}^{(r)}\right)^{-1}\left(\overline{y}^{(r)} - \overline{X}^{(r)}\beta^{(r)} - \overline{Z}^{(r)}\delta^{(r)}\right)\right\}. \tag{A5}$$

The prior distribution is set as follows: the vectors of the autoregressive coefficients $\beta^{(1)}$ and $\beta^{(2)}$ are distributed normally around zero and 0.9, respectively, for the first lag of the interest rate in the interest rate equation. The interest rate enters the vector of endogenous variables in levels. The variance of the prior distribution is of the Minnesota-type prior, that is, a diagonal matrix. The element corresponding to the *i*-th equation and (each lag of) the*j*-th variable equals to

$$\sigma_i^2/\sigma_j^2, \tag{A6}$$

where σ_i^2 is the estimated standard error from the univariate *AR(4)* model of *i*-th variable.

The prior on the error covariance matrix $\overline{\Sigma}^{(r)}$ is distributed as the inverse Wishart distribution. The scale matrix is such that its diagonal elements are equal to the estimates of the error variances in the *AR(4)* models of the respective endogenous variables. The degrees of freedom are equal to $n + 1$, suggesting a rather uninformative prior capable of dealing with the fact that the variances of the endogenous variables differ by orders of magnitude.

The prior on the unconditional mean of endogenous variables $F^{(r)}$ is distributed normally. In the case of the interest rate and output growth, the means of the prior distribution are based on the estimates from Holston et al. (2017). More precisely, we take the authors' average estimate over the period for which our model is estimated. For inflation, we assume the prior mean equal to two, and the historical average of 0.71 is taken for the spread. The variances of the prior on the steady state values are equal to one. Note that the priors are the same for the steady states in both regimes. The prior distribution on the weights w_i in the definition of the threshold variable in (2) is assumed to be uniform.

Finally, the threshold R is the last parameter for which the prior is formulated, whereby there are two a priori restrictions. First, it is the minimum number of observations in a regime. Second, it is the requirement that the steady state in Regime 1 belongs to Regime 1, i.e., $Y^{TR}\left(F^{(1)}\right) < R$, and/or the steady state in Regime 2 belongs to Regime 2, i.e., $Y^{TR}\left(F^{(2)}\right) \geq R$. These two restrictions are also imposed when taking draws of the threshold during the Gibbs sampling (a detailed discussion follows below). As such, the conditional prior $p\left(R|F^{(1)},\ F^{(2)},\ w\right)$ is uniform on the set of all values R, which satisfy:

(a) for a certain threshold variable (given by w), R implies regimes with at least a given minimum number of observations, and
(b) for a given unconditional mean $F^{(1)}$ and $F^{(2)}$ and for the threshold variable defined by w, R is such that $Y^{TR}\left(F^{(1)}\right) < R$ or $Y^{TR}\left(F^{(2)}\right) \geq R$.

The joint prior on the vector of parameters $\boldsymbol{\theta}$ is formulated as follows:

$$p(\boldsymbol{\theta}) \propto \left[\prod_{r=1}^{2} p\left(\beta^{(r)}\right)p\left(\overline{\Sigma}^{(r)}\right)p\left(F^{(r)}\right)\right]p\left(R|F^{(1)},\ F^{(2)},\ w\right)p(w). \tag{A7}$$

Conditional posteriors and the Gibbs sampler (A.3)

Following the Bayes theorem, combining the likelihood in (A5) and the joint prior in (A7) yields a joint posterior. The joint posterior distribution is then simulated by taking draws from the conditional posteriors.

Let $\boldsymbol{\theta}_{-parameter}$ denote the parameter vector excluding a particular parameter. For example, $\boldsymbol{\theta}_{-R}$ contains all model parameters other than R, namely $\boldsymbol{\theta}_{-R} \equiv \left[\beta^{(1)},\ \beta^{(2)},\ F^{(1)},\ F^{(2)},\ w,\ \overline{\Sigma}^{(1)},\ \overline{\Sigma}^{(2)}\right]$. Let us now describe the steps in the simulation procedure and discuss the conditional posteriors.

(1) The conditional posterior of the vectors of AR parameters $\beta^{(1)}$ and $\beta^{(2)}$ is distributed normally:

$$p\left(\beta^{(r)}|\overline{y},\ \boldsymbol{\theta}_{-\beta^{(r)}}\right) \propto N\left(\beta^{(r)}_{POST},\ \Omega^{(r)}_{POST}\right), \tag{A8}$$

where

$$\beta^{(r)}_{POST} = \Omega^{(r)}_{POST}\left[\left(\Omega^{(r)}_{PRIOR}\right)^{-1}\beta^{(r)}_{PRIOR} + \overline{X}^{(r)}_{dem}{}'\left(\overline{\Sigma}^{(r)}\right)^{-1}\overline{X}^{(r)}_{dem}\beta^{(r)}_{OLS}\right], \tag{A9}$$

and

$$\Omega^{(r)}_{POST} = \left[\left(\Omega^{(r)}_{PRIOR}\right)^{-1} + \overline{X}^{(r)}_{dem}{}'\left(\overline{\Sigma}^{(r)}\right)^{-1}\overline{X}^{(r)}_{dem}\right]^{-1}. \tag{A10}$$

The variable $\overline{X}^{(r)}_{dem}$ denotes demeaned data matrices and $\beta^{(r)}_{OLS}$ is a standard OLS estimate of the vector $\beta^{(r)}$ estimated on the demeaned data. The demeaning of the data is due to the fact that the unconditional means of endogenous variables are treated separately in the estimation procedure.

For each draw from the conditional posterior a stability check is carried out, i.e., the eigenvalues related to the regime-specific VAR are compared to unity. In the case of an eigenvalue larger than one, the system is unstable and another draw is taken from the conditional posterior. The maximum of tries is set to 10,000. If neither $\beta^{(1)}$ nor $\beta^{(2)}$ ensuring a stable VAR is found, the draws from other conditionals based on this 'unstable' draw are discarded. It turns out that in all estimations the ratio of such $\beta^{(1)}$ or $\beta^{(2)}$ does not exceed 7%, and is below 1% for the majority of the estimation windows.

(2) The conditional posterior of the error covariance matrix $\overline{\Sigma}^{(1)}$ and $\overline{\Sigma}^{(r)}$ is distributed as inverse Wishart:

$$p\left(\Sigma^{(r)}|\overline{y}, \boldsymbol{\theta}_{-\Sigma^{(r)}}\right) \propto iW\left(S_{POST}, DoF_{PRIOR} + t^{(r)}\right), \tag{A11}$$

where

$$S_{POST} = S_{PRIOR} + \left(\overline{y}_{dem}^{(r)} - \overline{X}_{dem}^{(r)}\beta^{(r)}\right)'\left(\overline{y}_{dem}^{(r)} - \overline{X}_{dem}^{(r)}\beta^{(r)}\right). \tag{A12}$$

(3) The conditional posterior of the vector of unconditional means $F^{(r)}$ is distributed normally:

$$p\left(F^{(r)}|\overline{y}, \boldsymbol{\theta}_{-F^{(r)}}\right) \propto N\left(\varphi_{POST}^{(r)}, \Phi_{POST}^{(r)}\right), \tag{A13}$$

where

$$\begin{aligned}\varphi_{POST}^{(r)} &= \Phi_{POST}^{(r)} \\ &\times \left\{\Phi_{POST}^{(r)}{}^{-1}\varphi_{PRIOR} + U^{(r)'} vec\left[\left(\overline{\Sigma}^{(r)}\right)^{-1}\left(\overline{y}^{(r)} - \overline{X}^{(r)}\beta^{(r)}\right)'\overline{Z}^{(r)}\right]\right\},\end{aligned} \tag{A14}$$

and

$$\Phi_{POST}^{(r)} = \left[\Phi_{POST}^{(r)}{}^{-1} + U^{(r)'}\left(\overline{Z}^{(r)'}\overline{Z}^{(r)} \otimes \left(\overline{\Sigma}^{(r)}\right)^{-1}\right)U^{(r)}\right]^{-1}, \tag{A15}$$

such that

$$U^{(r)} \equiv \begin{pmatrix} I_4 \\ A_1^{(r)} \\ \vdots \\ A_p^{(r)} \end{pmatrix}.$$

The data matrices $\overline{y}^{(r)}$ and $\overline{X}^{(r)}$ are not demeaned. The last subset of the parameter vector $\boldsymbol{\theta}$ consists of the weights constituting the threshold variable and the threshold value r. Both types of parameters are treated as discrete variables.

(4) The space for weights $[w_1, \ldots, w_4]$ is discretized such that we work with all 4-tuples, with each element taking values from $\{0, 0.1, \ldots, 0.9, 1\}$ and adding up to one. The likelihood of data is then computed for each 4-tuple and a random draw from the multinomial distribution with probabilities given by the likelihood. The implementation of the random draw from the multinomial distribution reflects the fact that the likelihood of data for the majority of 4-tuples of weights is almost zero. Therefore, we work with log-likelihoods and keep the sampling procedure in log-space. We proceed as follows:

  (4a) Index all 4-tuples and evaluate the log of likelihood defined in (A5) for each 4-tuple of weights.
  (4b) Define bins by cumulative log-likelihoods.
  (4c) Take a random draw from the negative of the exponential distribution with mean 1. Here we exploit the equivalence between a log of the uniform random value and a negative of an exponential distribution random draw, which enables us to adhere to the log space.
  (4d) The bin defined in (4b), which includes the random draw from (4c), indicates the index of the drawn 4-tuple of weights.

(5) Finally, the conditional posterior of the threshold is affected by the prior conditional on the vector of weights w and by the unconditional means $F^{(1)}$ and $F^{(2)}$. The set of possible values of the threshold is restricted such that at least a given minimum number of observations is in each regime.

From an econometric point of view, in imposing a minimal number of observations we face a trade-off. On the one hand, rare events may be represented by a few observations only and their solid modeling is important for an accurate probabilistic assessment of the output growth distribution tails. On the other hand, some minimal number of observations is necessary to estimate our relatively heavily parameterized model in order to avoid (i) the estimation results being driven by the assumed priors and (ii) the situation of one additional observation changing the regimes completely. We assume that a regime contains at least 28 observations out of the 120 total observations, i.e., 7 out of the 30 years. The reason is that this seven-year minimum is close to the average length of the US business cycle.

All eligible values of the threshold are evaluated in terms of the likelihood, and a random draw is taken from the multinomial distribution with probabilities related to the likelihoods analogously to the procedure for weights described in (4).

The stability checks on $\beta^{(1)}$, $\beta^{(2)}$ are not sufficient conditions for the stability of a threshold VAR. We therefore check each simulated path of output growth and exclude those suggesting exploding patterns that exceed 400 in absolute value, i.e., those that describe the fall of the economy's production between two quarters as being equal to what the economy produced in the previous period. In this way, less than 4% of the simulated paths (and the corresponding TVARs) are excluded.

The Gibbs sampler involves repeating steps (1)–(5). The initial values include (i) the OLS estimates for the AR parameters, (ii) the degrees-of-freedom-adjusted estimate of the error covariance matrices in the regimes for the error covariance matrix, (iii) the means of the prior distributions for the unconditional means, (iv) a random draw

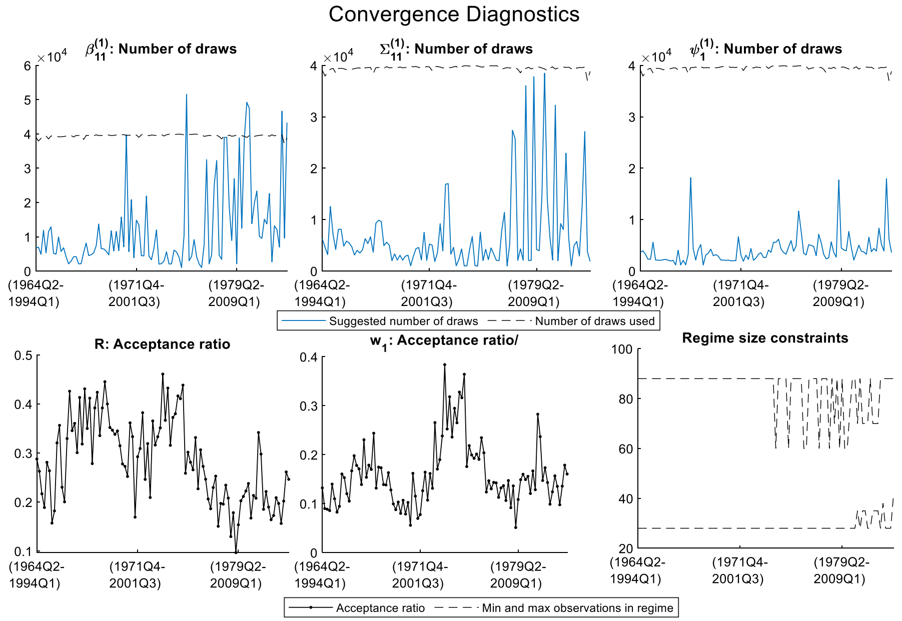

Fig. 6 Convergence diagnostics of the selected model parameters based on Raftery and Lewis (1992), acceptance ratios and imposed constraints on the number of observations in Regime 1

from the uniform distribution for the weights, and finally (v) a random draw from the uniform distribution over all threshold values satisfying the two above-mentioned restrictions.

Convergence of the Gibbs sampler (A.4)

This subsection presents the standard convergence diagnostics for all model parameters. For AR parameters, shock volatilities and steady states, the convergence measure is based on Raftery and Lewis (1992). It suggests how many draws should be taken from the conditional posteriors within the Gibbs sampler to obtain a stationary joint distribution.[19] For the threshold value and the weights in the threshold variable, the standard acceptance ratios are presented.

Figure 6 suggests that a sufficient number of draws is employed in the estimation procedure. Similarly, acceptance ratios belong into the usual interval between 0.1 and 0.5 suggesting a good mixing of the algorithm. For some estimation windows, the mixing was inferior with the sampler getting stuck in a subspace of the parameters' space. In such case, we restricted the number of observations in a regime (from below or from above) more than in the benchmark specification, where at least 28 observations in a regime are imposed. The minimum and maximum numbers of observations in Regime 1 imposed in the estimation window are reported in the last panel of Fig. 6.

[19] The usual diagnostics parameters are used: for the 0.025th and 0.975th quantiles of a marginal posterior distribution, an accuracy of 0.025 with a probability of 0.95 must be achieved.

Estimation with correlated variables in the threshold variable (A.5)

One of the estimation issues relates to the correlation of endogenous variables constituting the threshold variable y_t^{TR}. Negatively correlated endogenous variables could lead to switching between the regime labels across two adjacent estimation windows. In such case, the regimes could be almost the same, but in one estimation window one could in principle be denoted in the opposite way than in the other.

In the type of modeling exercise we perform, such a problem can occur in principle because the spread is strongly negatively correlated with inflation and the interest rate for some estimation windows (Fig. 7). Our solution is to simply work with the *negative* of the spread in the vector of the endogenous variables.

Appendix B: Additional results and procedures

In this appendix, we present additional results and procedures related to: (i) the estimation of the threshold VAR model (Sect. B.1), (ii) the estimated moments of the output growth distribution (Sect. B.2), (iii) the counterfactuals of the long-run output growth distribution moments (Sect. B.3), (iv) the bootstrap procedure for the confidence intervals in the equation for the conditional moments (Sect. B.4), and (v) validation of the model in (1) and the normality tests of the short-run and long-run output growth distribution (Sect. B.5).

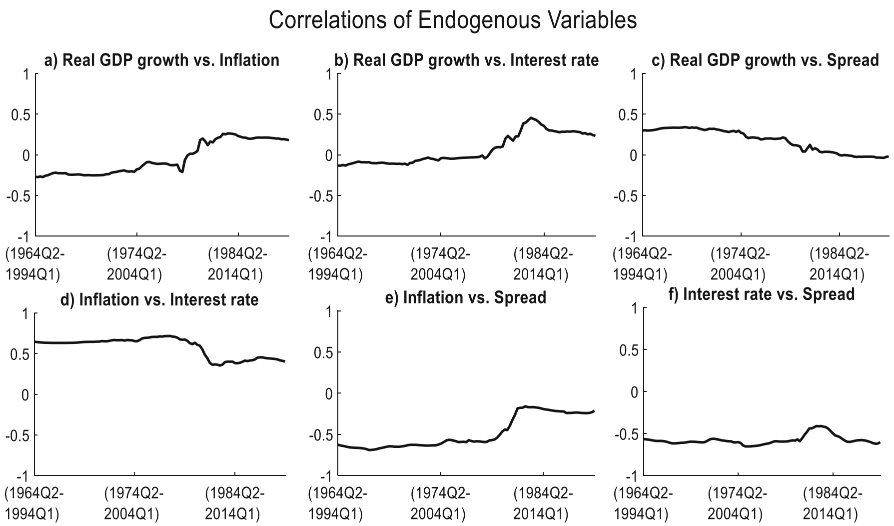

Fig. 7 Correlation coefficients of the endogenous variables. The x-axis indicates the range of quarters of the estimation window

Threshold VAR model (B.1)

Employing a model with two regimes raises the question of how the two regimes differ. Figures 8 and 9 plot the posterior means of the regime-specific steady states and the regime-specific shock volatilities. Figure 9 shows that the estimation windows preceding the GFC are characterized mainly by different shock volatilities across the regimes. In contrast, Fig. 8 indicates that the windows containing the 2001 recession and the GFC periods are characterized by different steady states, implying bimodality of the long-run output growth distribution. This is in line with Adrian et al. (2021), who found the bimodality of the output growth distribution to be a feature related to financially turbulent times, whereas 'normal' times are characterized by a Gaussian output growth distribution.

In addition to the distinct steady states or shock volatilities, the regimes can be interpreted based on the estimated weights in the threshold variable defined in (2). To provide details, Fig. 10 shows the most frequent weights drawn during the Gibbs sampling for each estimation window, whereas Fig. 11 reports the portion of explained variation of the threshold variable by a given endogenous variable. The dominant role of the interest rate for the first third of the sample is apparent. After that, however, inflation and the spread start to drive the regime change. Some role for real GDP growth is found at the end of the sample.

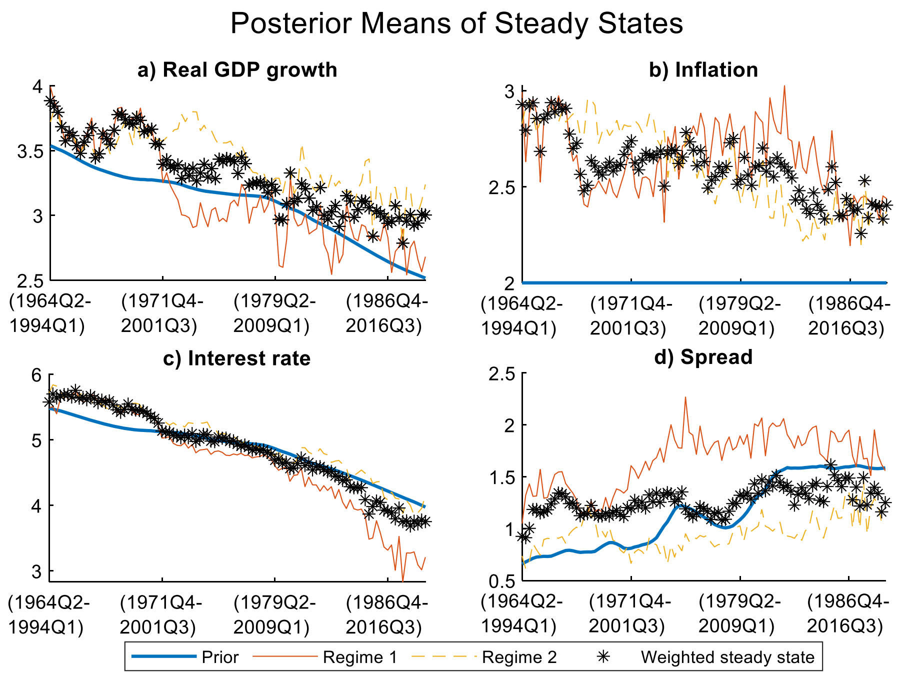

Fig. 8 The mean of the prior distribution of the steady states, the means of the posterior distribution in the two regimes, and the weighted posterior mean of long-run output growth with weights equal to the probability of the regime in the long run. The x-axis indicates the range of quarters of the estimation window

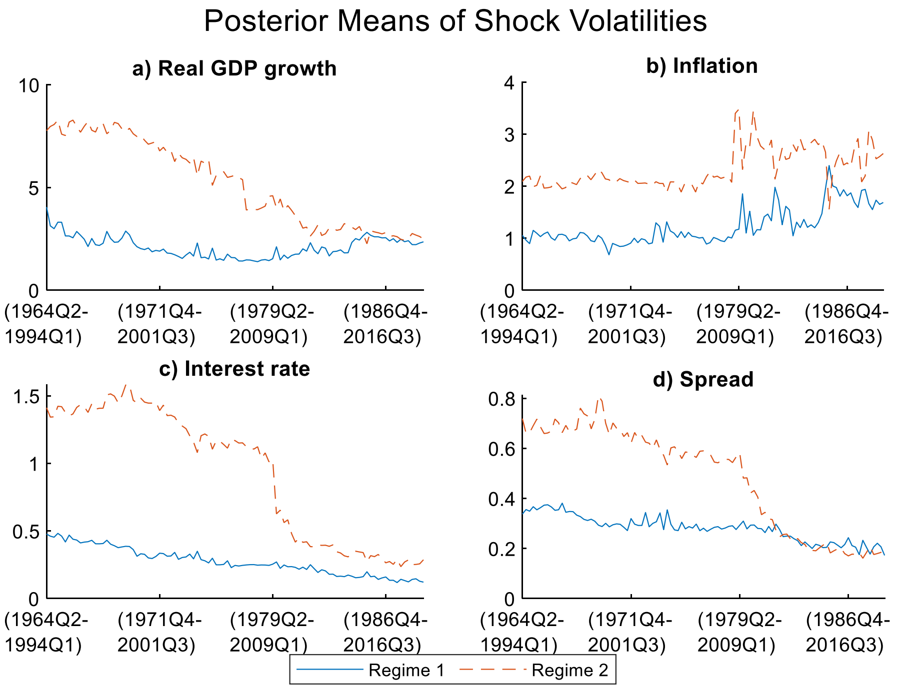

Fig. 9 The posterior mean of the shock volatilities in the two regimes. The x-axis indicates the range of quarters of the estimation window

Figure 12 shows the number of observations in Regime 1. The purpose of the figure is to demonstrate that there are no abrupt changes in the size of the regimes in adjacent estimation windows that would indicate the issue of sample instabilities and outliers. In case of such instabilities, there would be a risk of confusing them with changes in the output growth distribution properties.

Figures 13, 14, and 15 present the estimated (reduced-form) coefficients at the lagged endogenous variables in the interest rate equation, the spread equation and the inflation equation. More precisely, we take the sum of the estimates at all lags of a given endogenous variable to characterize the overall effect (our interest does not lie in the exact profile). Figures 13 and 14 demonstrate changes in the behavior of the monetary authority after the GFC (the dotted vertical line), i.e., when the GFC enters the estimation windows.

Regarding the change in the behavior of the monetary authority during and after the GFC in the form of an adjustment to its interest rate rule, it turns out that there is a break in the AR parameters at lagged inflation from positive values to values close to zero starting with the window of 1979Q2-2009Q1 (Fig. 13b). Such change can be observed in both regimes and also if we summarize the information across the regimes by a weighted sum of the AR parameters with weights equal to the probability of each regime in the long run (the solid curve). Furthermore, the weight of output growth in the interest rate equation declined to zero when the GFC entered the estimation

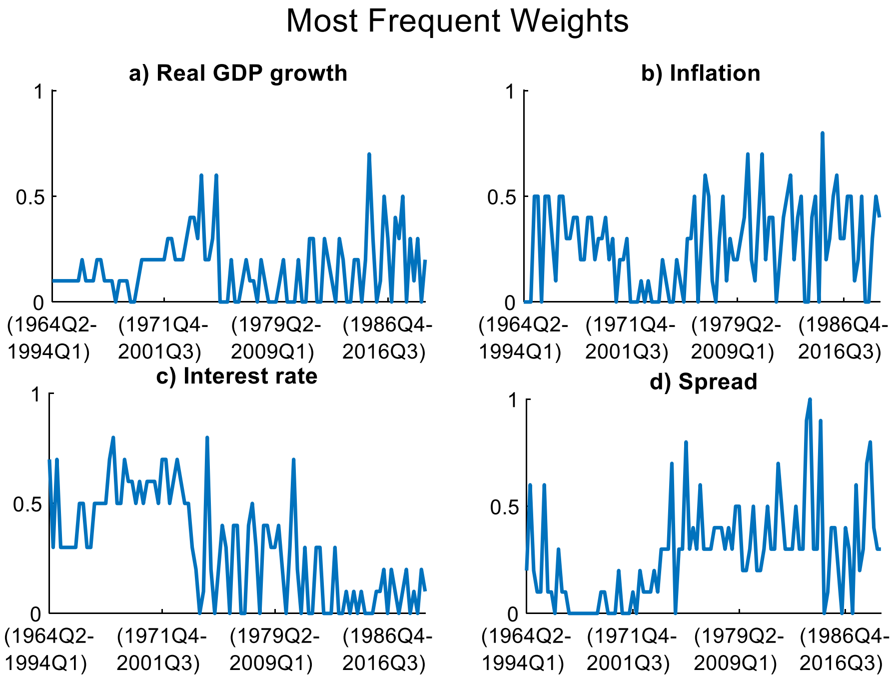

Fig. 10 The most frequent weight of endogenous variables in the threshold variable in the Gibbs chain. The weights can attain values from the set $\{0, 0.1, \ldots, 0.9, 1\}$ only. For details see Appendix A.3. The x-axis indicates the range of quarters of the estimation window

windows (due to the ZLB), and then increased when the post-2014Q1 entered the estimation windows. The zero effect of lagged output growth on the interest rate can also be observed at the very beginning of the series of estimation windows when the mid-1960s and 1970s are included. This is in line with, e.g., Bianchi (2013) discussed in the main text.

The link between the spread and lagged output growth changed qualitatively as well (see Fig. 14a). When focusing on weighted sum of coefficient, it turns out that the positive reduced-form relationship estimated for the windows containing the GFC switched to zero when the post-2014Q1 period entered the windows.

Estimated distribution moments (B.2)

Figures 16 and 17 present the first four moments of the conditional and long-run output growth distributions, respectively. Both figures also report the skewness and kurtosis of several standard distributions (dashed and dotted lines) to provide some guidance regarding the extent of the asymmetry and fat tails.

In the short-run perspective, panels (c) a (d) of Fig. 16 suggest that the conditional output growth distribution is generally asymmetric and does not feature fat tails. In particular, panel c) shows that the size of the asymmetry exceeds the difference in

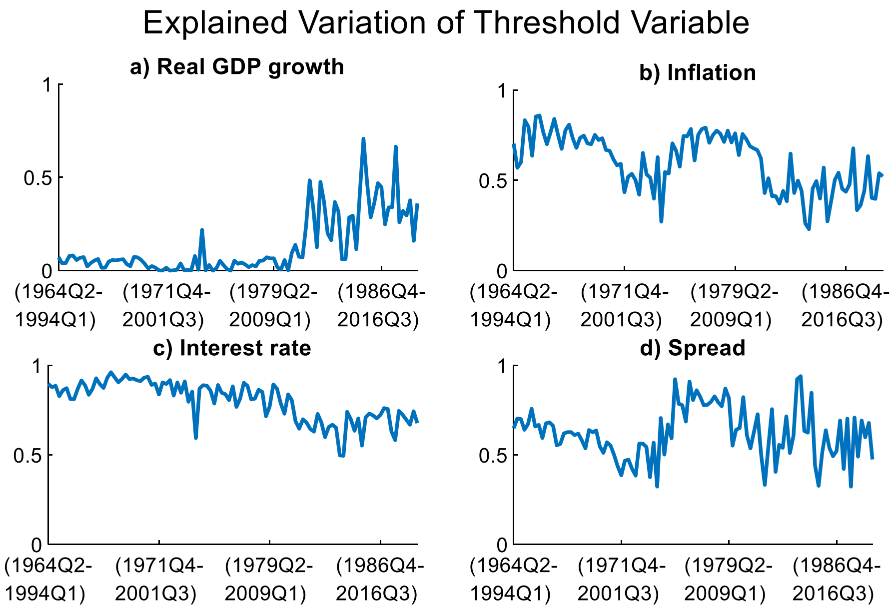

Fig. 11 The portion of the variation of the threshold variable explained by the endogenous variable. The explained variation is represented by R2 from regression of mean threshold variable on each smoothed (MA(4)) endogenous variable in turn. The *x*-axis indicates the range of quarters of the estimation window

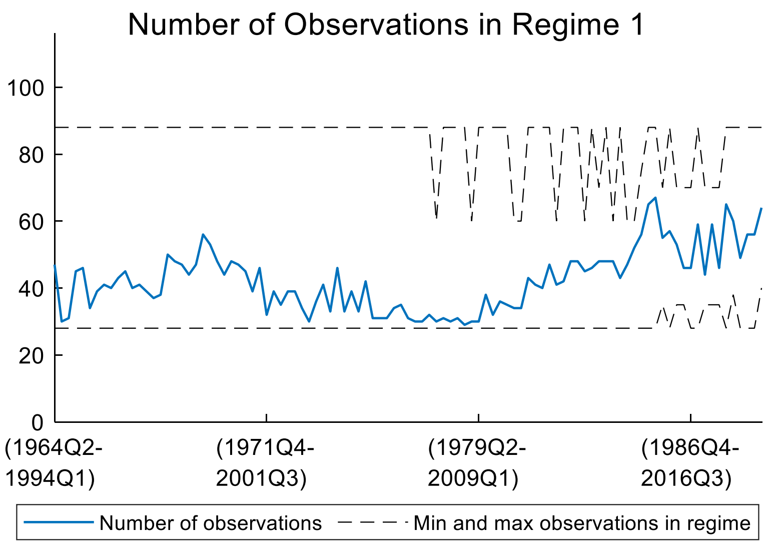

Fig. 12 The number of observations in Regime 1 and the imposed minimum and maximum of observations in Regime 1. The *x*-axis indicates the range of quarters of the estimation window

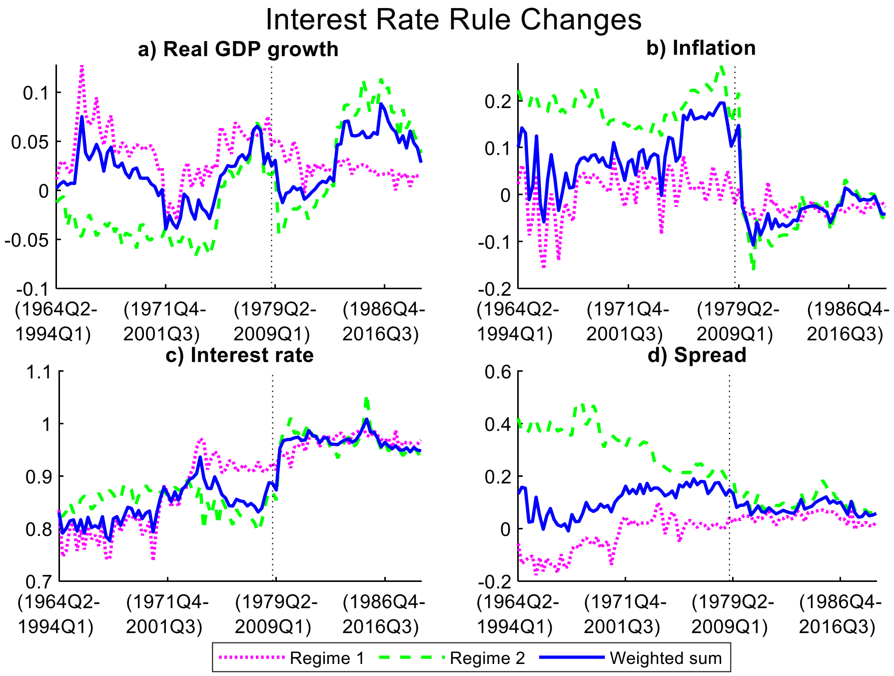

Fig. 13 The sum of the estimated coefficients at lags of the respective variable in the interest rate equation, together with a weighted sum over the two regimes (with weights equal to the probability of the regime in the steady state). The x-axis indicates the range of quarters of the estimation window. The dotted vertical line indicates the estimation window in which the GFC first enters the estimation (1979Q1-2008Q4)

skewness of the symmetric distribution and the lognormal distribution with $\mu = 0$ and $\sigma = 0.1$.

Moving from the short-run to the long-run perspective, Fig. 17 shows how the unconditional moments have evolved over time. Panels (c) and (d) suggest that the long-run distribution is closer to being symmetric and it is heavy-tailed, which contrasts the conditional distribution in Fig. 16 that is asymmetric with almost no fat tails.

Our ability to reproduce important findings in the literature serves as a further justification for our modeling approach. The lack of asymmetry of the long-run distribution is in line with, e.g., Carriero et al. (2020), who show this for an unconditional distribution. In terms of the fourth moment, the fat tails of the unconditional output growth distribution have been detected by, e.g., Fagiolo et al. (2008). We can also replicate two well-known long-term phenomena of the past several decades related to the first two moments. They are a gradual decrease in potential output growth over the past several decades [see, e.g., Antolin-Diaz et al. 2017)] and a decrease in the shock volatility of macroeconomic variables [known as the Great Moderation, see, e.g., Stock and Watson (2003)].

Figure 17 also reports sample moments based on observations in the respective estimation window. The simulated long-run moments need not coincide with the sample moments for two reasons. First, we impose a mixture of two normals, while the empirical distribution can be, for example, with three modes. Second, we impose priors on

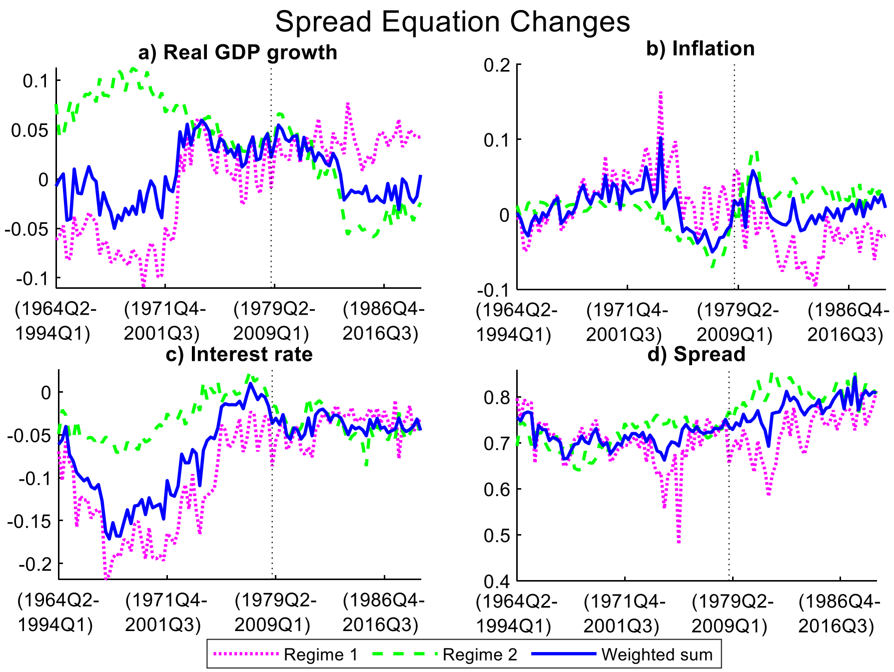

Fig. 14 The sum of the estimated coefficients at lags of the respective variable in the spread equation, together with a weighted sum over the two regimes with weights equal to the probability of the regime in the steady state. The x-axis indicates the range of quarters of the estimation window. The dotted vertical line indicates the estimation window in which the GFC first enters the estimation (1979Q1-2008Q4)

model parameters that affect posteriors, whereas the unconditional sample moments are purely data driven. An alternative would be to work directly with the empirical distribution of observed output growth rates, but it would have several shortcomings. For example, we could not construct the counterfactuals presented in Sect. 6.

Distribution moments of counterfactuals (B.3)

Figures 18, 19 and 20 below present the mean, variance, skewness and kurtosis of the long-run output growth distribution for our three counterfactuals from Sect. 6, respectively. These were discussed in detail in Sect. 6.

Confidence intervals: model for conditional distribution moments (B.4)

The model in (4) is estimated by the OLS. Confidence intervals for the parameters of interest $\beta_h^{(m)}$ are constructed employing a bootstrap procedure. To deal with the potential inefficiency of estimates implied by the fact that the dependent variable $\mu_{t+h+4|t+h}^{(m)}$ is generated in the first stage, we take 500 bootstrap samples of $\mu_{t+h+4|t+h}^{(m)}$ by sampling from the set of all simulated forecast paths from the first stage with

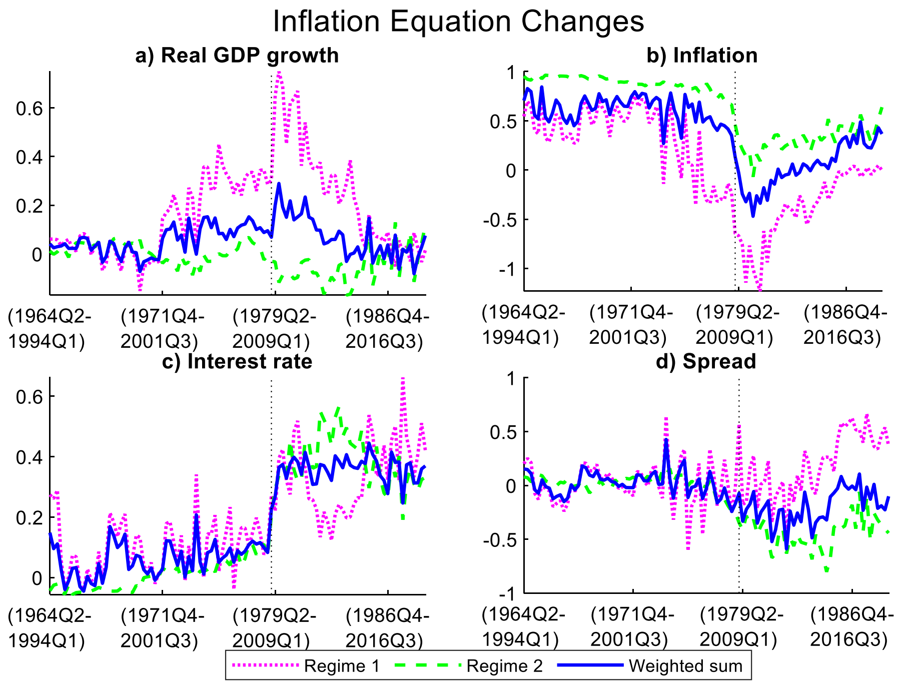

Fig. 15 The sum of the estimated coefficients at lags of the respective variable in the inflation equation, together with a weighted sum over the two regimes with weights equal to the probability of the regime in the steady state. The x-axis indicates the range of quarters of the estimation window. The dotted vertical line indicates the estimation window in which the GFC first enters the estimation (1979Q1-2008Q4)

replacement. Similarly, bootstrap samples are generated for $\mu_{t+4|t}^{(m)}$ to account for the generated regressor. To get confidence intervals for the coefficient $\beta_h^{(m)}$ moving block bootstrap is used. More specifically, for each resampled set of forecast paths, we compute a particular quantile for each t. Then, we take blocks of the data matrix (containing both left- and right-hand side variables) of length 4 and a construct a new data matrix of the length of the original data matrix to estimate the model in (4). Then, the OLS estimates based on the bootstrap samples are summarized to get the confidence intervals.

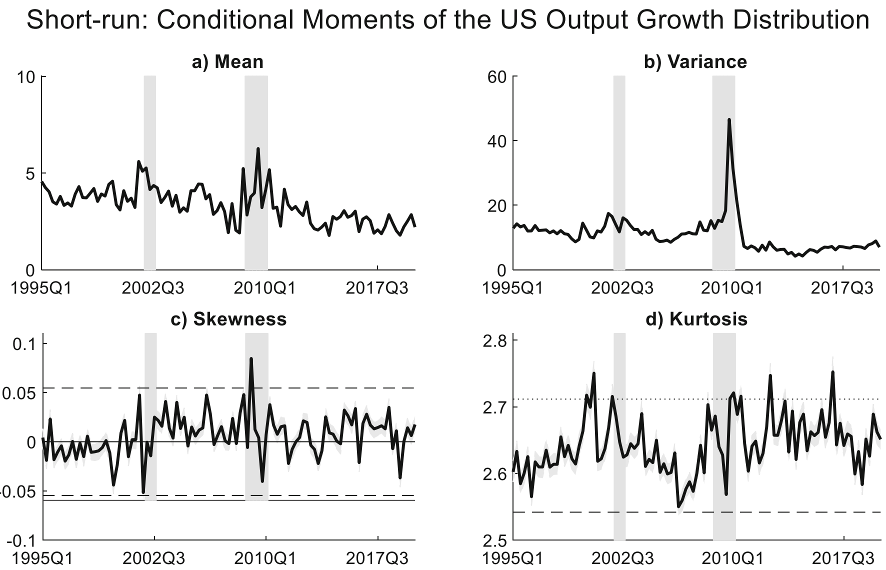

Fig. 16 The conditional mean, variance, skewness and kurtosis of the four-quarters-ahead output growth distribution. The x-axis indicates the quarter of the estimated conditional output growth distribution. The distribution is conditional on the data available four quarters earlier. The shaded areas denote the 90% confidence intervals obtained by bootstrapping from 500 bootstrap samples. The dashed line in **c** indicates skewness of the simulated lognormal distribution (the negative of) with parameters $\mu = 0$, $\sigma = 0.1$. The dashed line in **d** indicates the kurtosis of the simulated standard normal distribution. The dotted line in **d** indicates the kurtosis of the simulated Student's t-distribution with 10 degrees of freedom. The number of simulated values from the respective distributions equals the number of simulated values of the output growth distribution reported in the respective panel. The vertical shaded areas in all panels represent official NBER recessions. The first two moments are represented by the mean and variance of the simulated values of output growth. The third and fourth moments are represented by the measures of skewness and kurtosis less sensitive to outliers as introduced in Groeneveld and Meeden (1984) and Hogg (1974), respectively

Out-of-sample fit (B.5)

To validate the use of the model in (1) from the out-of-sample perspective, we take two approaches. The first approach is to directly measure the density forecasting ability of the model. To that end, the expected logarithmic score for the threshold VAR and constant parameter VAR are computed and compared. The expected logarithmic score is estimated based on the average logarithmic score—the average log of simulated density at a horizon of 4 quarters for a realization of the endogenous variable y_t^i:

$$\frac{1}{N}\sum_t \ln f_{t+4,t}\left(y_t^i\right), \tag{B1}$$

where N denotes the number of available realizations of the variable at the given horizon. It turns out that the threshold VAR is preferred to the constant parameter VAR as their average logarithmic scores equal − 1.12 and − 1.19, respectively.

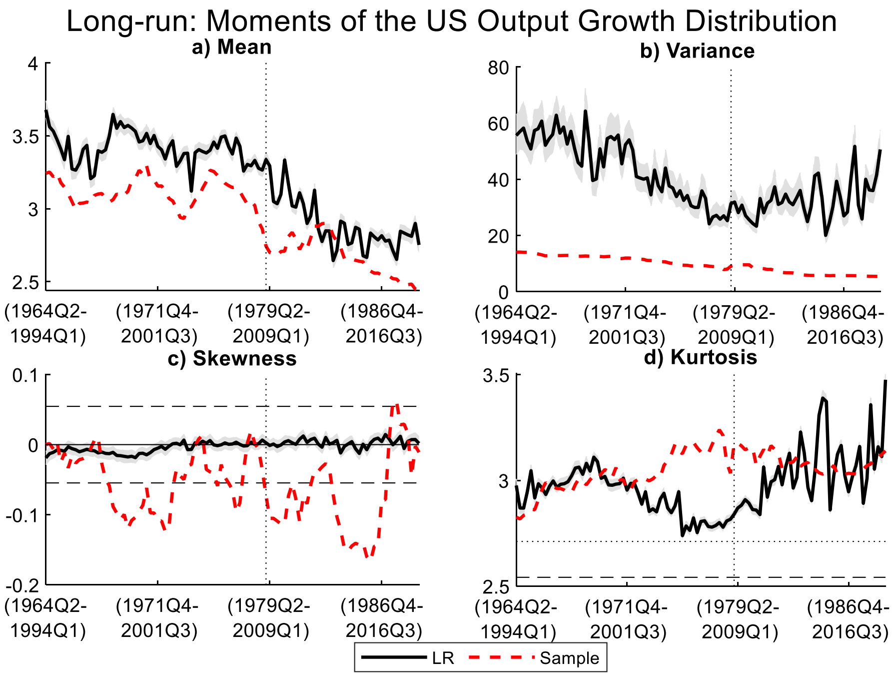

Fig. 17 The mean, variance, skewness and kurtosis of the long-run output growth distribution. The x-axis indicates the period used to estimate the model employed for simulation of the long-run output growth distribution. The shaded areas denote 90% confidence intervals obtained by bootstrapping from 500 bootstrap samples. The dashed line in **c** indicates skewness of the simulated the lognormal distribution (the negative of) with parameters $\mu = 0, \sigma = 0.1$. The dashed line in **d** indicates the kurtosis of the simulated standard normal distribution. The dotted line in **d** indicates the kurtosis of the Student's t distribution with 10 degrees of freedom. The dotted vertical line indicates the estimation window when the GFC enters the estimation (1979Q1-2008Q4). The first two moments are represented by the mean and variance of the simulated values of output growth. The third and fourth moments are represented by the measures of skewness and kurtosis less sensitive to outliers as introduced in Groeneveld and Meeden (1984) and Hogg (1974), respectively

The second approach consists of two normality tests of the simulated output growth distributions. The first is a one-sample Kolmogorov–Smirnov test with the null hypothesis of the sample being from the standard normal distribution applied on the standardized simulated sample. The second is the Lilliefors test with the null of the sample being generated by a distribution from the family of normal distributions. For the long-run output growth distribution, both tests reject the null of normality for all estimation windows. For the output growth distributions four quarters ahead, both tests reject normality in the vast majority of cases (98 out of 101 cases for the Kolmogorov Smirnov test and all cases for the Lilliefors test).

The normality tests suggest that a model that allows for non-normal predictive distributions is informative. To illustrate the non-normality of the simulated distributions, Fig. 21 shows a histogram of the observed interest rate in the last estimation window 1989Q2-2019Q1 and the corresponding simulated long-run interest rate distribution. Note that zero lower bound is imposed.

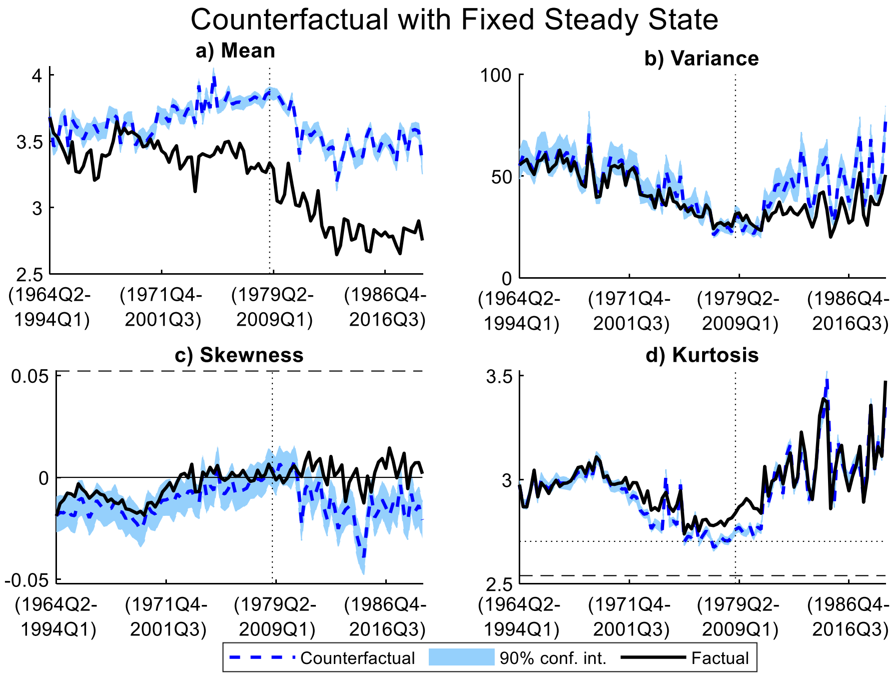

Fig. 18 Moments of the long-run output growth distribution of the counterfactual with a fixed steady state. The x-axis indicates the range of quarters of the estimation window. The dotted vertical line indicates the estimation window when the GFC enters the estimation (1979Q1-2008Q4)

Appendix C: Robustness checks

This appendix includes several robustness checks related to: (i) the analysis of the distributional effects of structural shocks on the conditional output growth distribution (Sect. C.1), (ii) the choice of the quantile that defines our measure of tail risk (Sect. C.2), and (iii) the way in which the zero lower bound is imposed in the simulation of predictive distributions (Sect. C.3).

Monetary policy shock: alternative specifications and shock measure (C.1)

Several specifications of the model for the conditional output growth distribution moments are estimated here to provide further insights and robustness checks. Figure 22 presents the nonlinear variant of the model in (4) separating policy tightening and easing.

An important aspect of the analysis relates to the exogeneity of monetary policy shocks. The exogeneity of shocks with respect to the conditional mean is guaranteed by the procedure in Romer and Romer (2004). The authors use historical documents and construct a series of intended monetary policy actions, which they confront with real-time macroeconomic data and forecasts to render policy changes orthogonal to

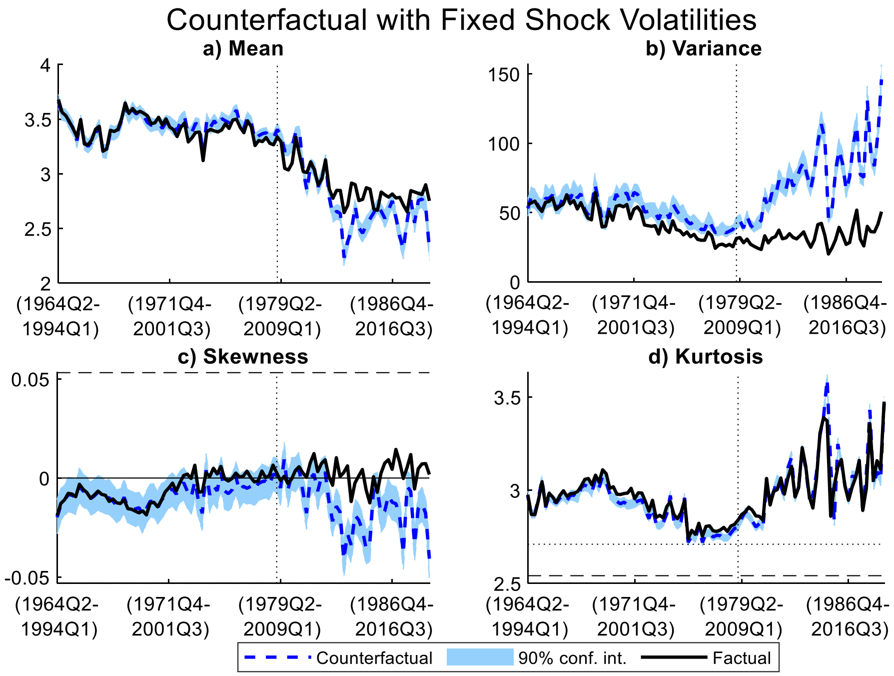

Fig. 19 Moments of the long-run output growth distribution of the counterfactual with fixed shock volatilities. The x-axis indicates the range of quarters of the estimation window. The dotted vertical line indicates the estimation window when the GFC enters the estimation (1979Q1-2008Q4)

the policymakers' information set. They thus filter out the change in the Federal Funds rate that is not related to current and future inflation and real economic activity.

Importantly, Romer and Romer (2004) control for the conditional mean forecasts when estimating unexpected changes in the Federal Funds rate. Despite these measures, endogeneity can still arise in two ways. First, the unexpected change relates to a higher moment of the output growth distribution while, at the same time, the moment affects future output growth moments. Second, higher moments of the future output growth distribution can be affected by a macroeconomic variable not included in the policymakers' information set. If this were to occur, the estimated effect of shocks on variance, skewness and kurtosis would be biased. To shed some light on this question, we conduct a robustness check with differently identified monetary policy shocks. The shocks are taken over from Gertler and Karadi (2015), who look at monetary policy surprises that are orthogonal to both economic and financial variables. The exogeneity with respect to financial variables is crucial because the literature shows financial conditions affect future output growth distributions (Adrian et al. 2019).

Gertler and Karadi's (2015) identification scheme draws on Federal Funds futures observed on FOMC meeting days and employs a high frequency identification approach. Using monetary policy shocks from Gertler and Karadi (2015) provides similar results to our benchmark shocks taken from Wieland and Yang (2020). In summary, the conditional mean decreases after a few quarters (although the effect is

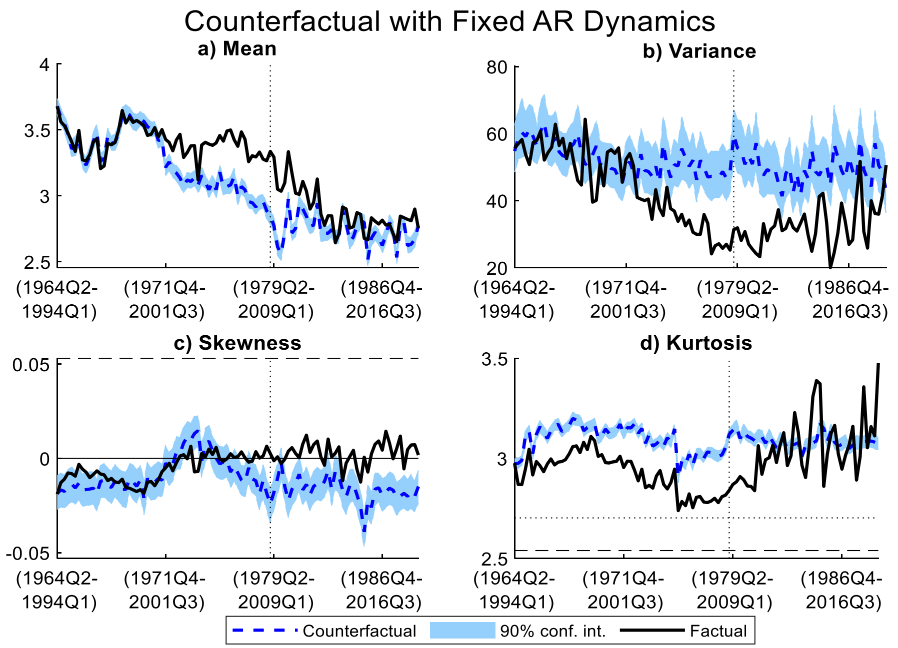

Fig. 20 Moments of the long-run output growth distribution of the counterfactual with fixed AR parameters. The x-axis indicates the range of quarters of the estimation window. The dotted vertical line indicates the estimation window when the GFC enters the estimation (1979Q1-2008Q4).

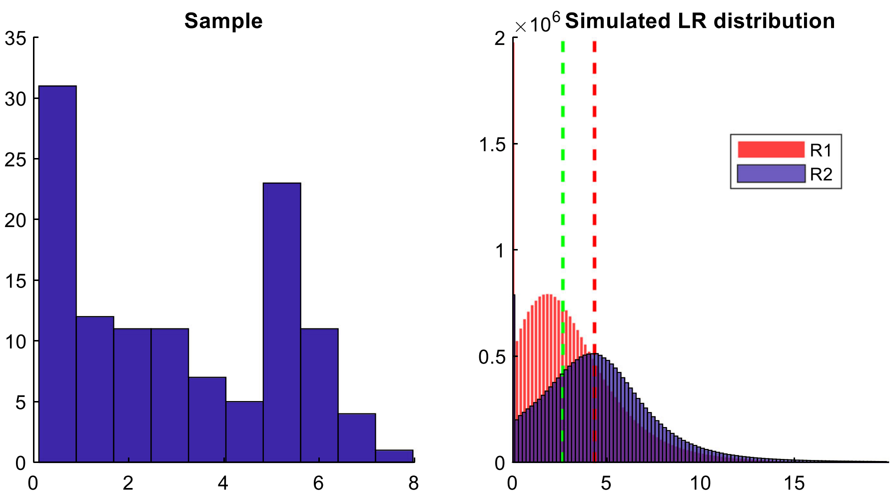

Fig. 21 Histogram of the observed interest rate in estimation window 1989Q2-2019Q1 and corresponding simulated long-run distribution

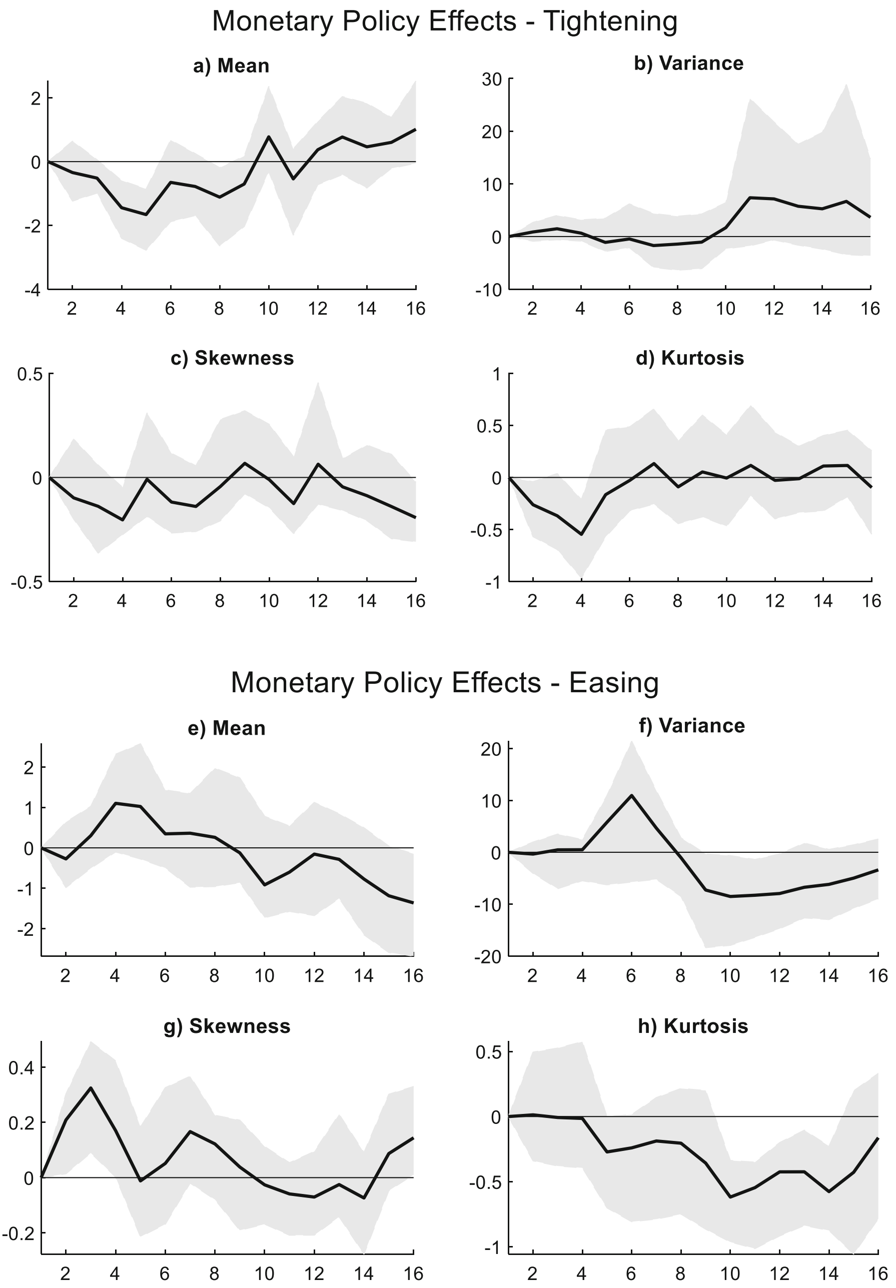

Fig. 22 Response of the conditional mean, variance, skewness and kurtosis of the four-quarters-ahead output growth distribution to an unexpected 100-basis-points monetary policy shock. The top panel shows policy tightening, the bottom panel shows policy easing. The shaded areas indicate the 90% confidence intervals. The separate policy shocks are employed in a way that if tightening shocks are considered only, unexpected easing is represented by 0 and vice versa

on the edge of statistical significance), skewness decreases temporarily and kurtosis is largely unaffected, at least in the medium run (Fig. 23).

Finally, a robustness check employing the monetary policy shocks from periods of both conventional and unconventional monetary policies is discussed in relation to Fig. 24. We use monetary policy shocks from Bu et al. (2021). The local projection regression (4) is estimated for the period 1994Q1-2015Q2. A less significant effect on skewness in the short-run and medium-run in comparison with the benchmark case (which covered conventional monetary policy only) is found together with a more profound effect on the variance. It follows that the combination of conventional and unconventional monetary policy shocks has similar effects to financial shocks; they exhibit a negligible effect on asymmetry for tightening and a more profound effect on variance.

Estimate of the 1st percentile of the US output growth distribution (C.2)

Figure 25 presents the 1st and 10th percentiles of both the conditional and unconditional output growth distributions. It demonstrates that focusing on the latter as a measure of downside risk is warranted; moving toward lower quantiles does not offer additional insights. The fact that the 1st percentile is more variable than the 10th (both short run and long run) is to be expected.

The effect of (not) imposing the ZLB in the simulation of predictive distributions (C.3)

During the simulation of predictive distributions in the main text, the zero lower bound was implemented in the most natural way by imposing a zero interest rate whenever its path tended to fall below zero. The issue of the ZLB is revisited in Fig. 26, which provide the simulation results without the ZLB being imposed. It is apparent that the main features of the conditional and long-run distributions remain intact, and so do our key findings in the main text.

In particular, even without the ZLB the post-2014 reversal of the trend in the 10th percentile of the unconditional output growth distribution can still be observed (Fig. 26, the right panel). Similarly, the change in unconditional skewness is still the most important development in the unconditional output growth distribution after the GFC enters the estimation windows (Fig. 27c).

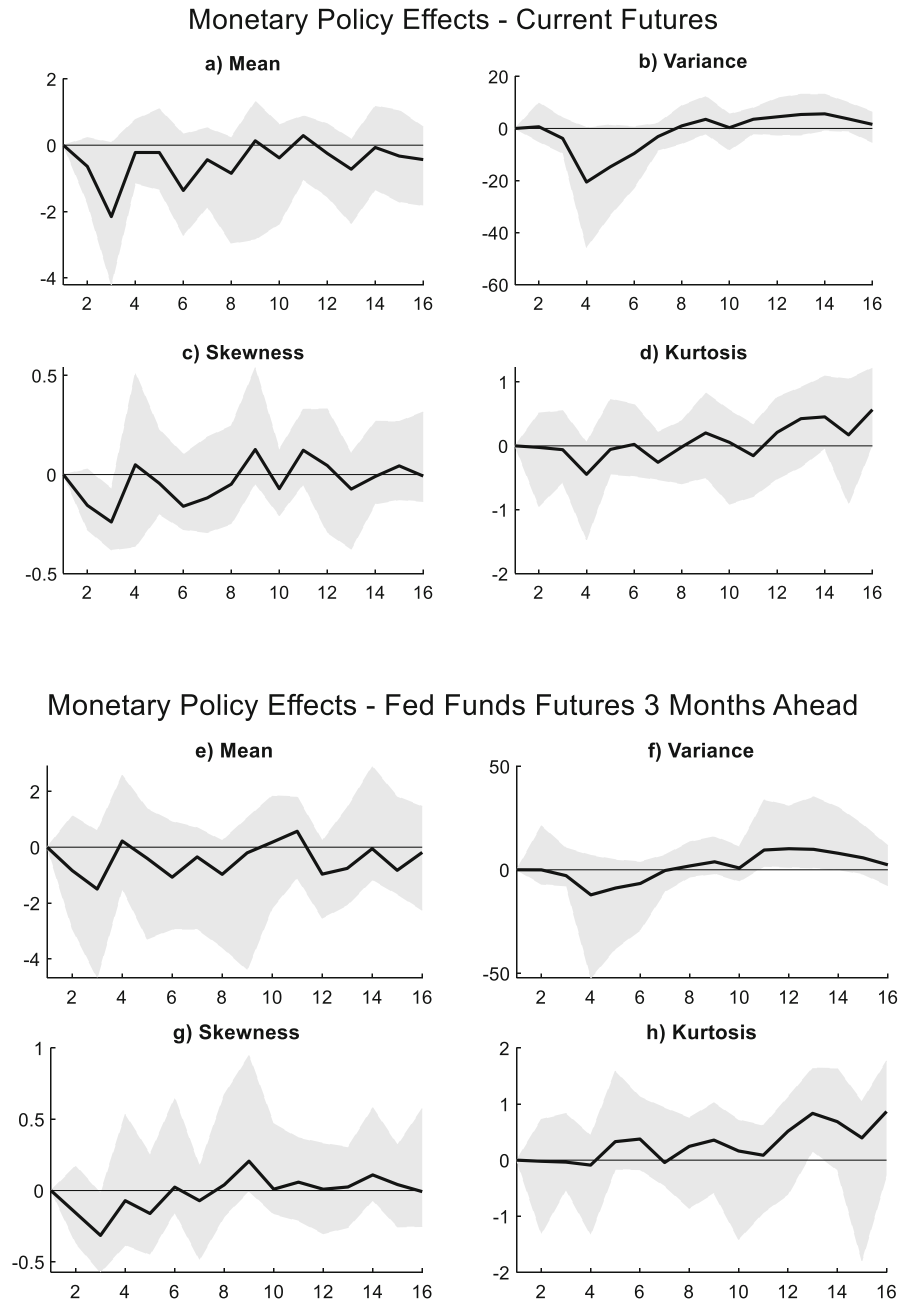

Fig. 23 Response of the conditional mean, variance, skewness and kurtosis of the four-quarters-ahead output growth distribution to an unexpected 100-basis-points monetary policy tightening with shocks from Gertler and Karadi (2015). The top panel is based on current futures, the bottom panel is based on the Federal Funds futures 3 months ahead. The shaded areas indicate the 90% confidence intervals

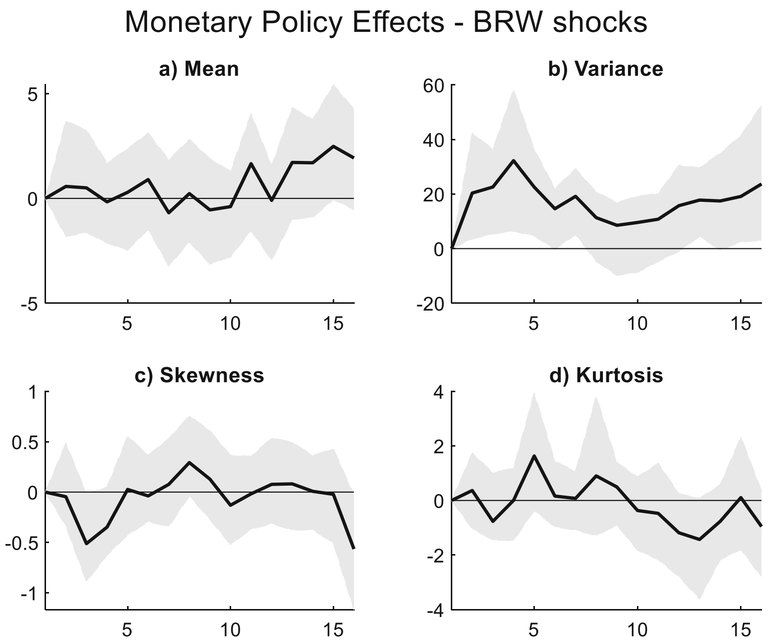

Fig. 24 Response of the conditional mean, variance, skewness and kurtosis of the four-quarters-ahead output growth distribution to an unexpected 100-basis-points monetary policy tightening with shocks from Bu et al. (2021). The top panel is based on current futures, the bottom panel is based on the Federal Funds futures 3 months ahead. The shaded areas indicate the 90% confidence intervals

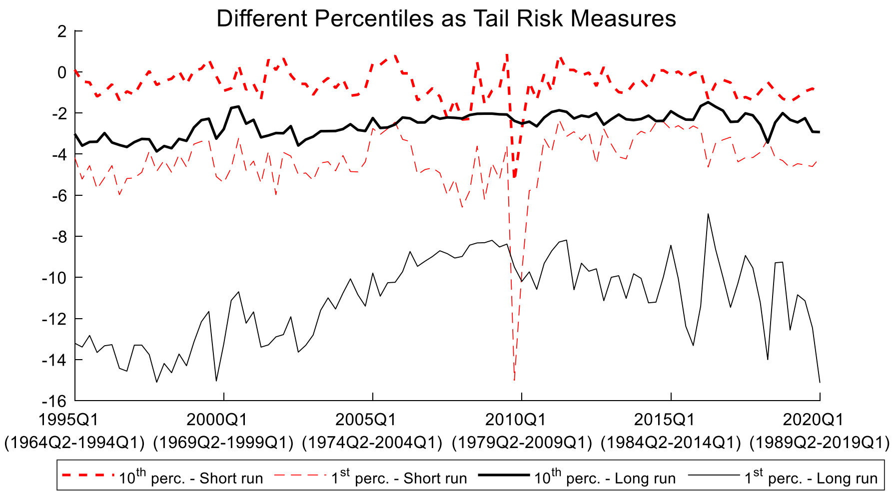

Fig. 25 The 1st and 10th percentiles of the four-quarters-ahead and long-run output growth distribution. The *x*-axis indicates the quarter for which the percentiles of the conditional output growth distribution are estimated. The periods in parenthesis indicate the windows used to estimate the long-run output growth percentiles

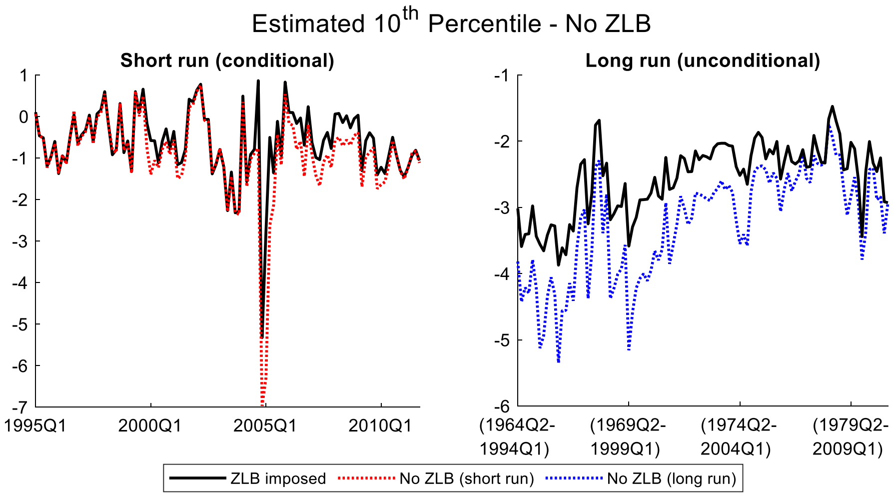

Fig. 26 The 10th percentile of the four-quarters-ahead (left panel) and long-run (right panel) output growth distribution from the model specification with and without imposing the zero lower bound. The *x*-axis indicates the quarter for which the percentiles of conditional output growth distribution are estimated. The periods in parenthesis indicate the windows used to estimate the long-run output growth percentiles. The dashed lines indicate 90% confidence intervals

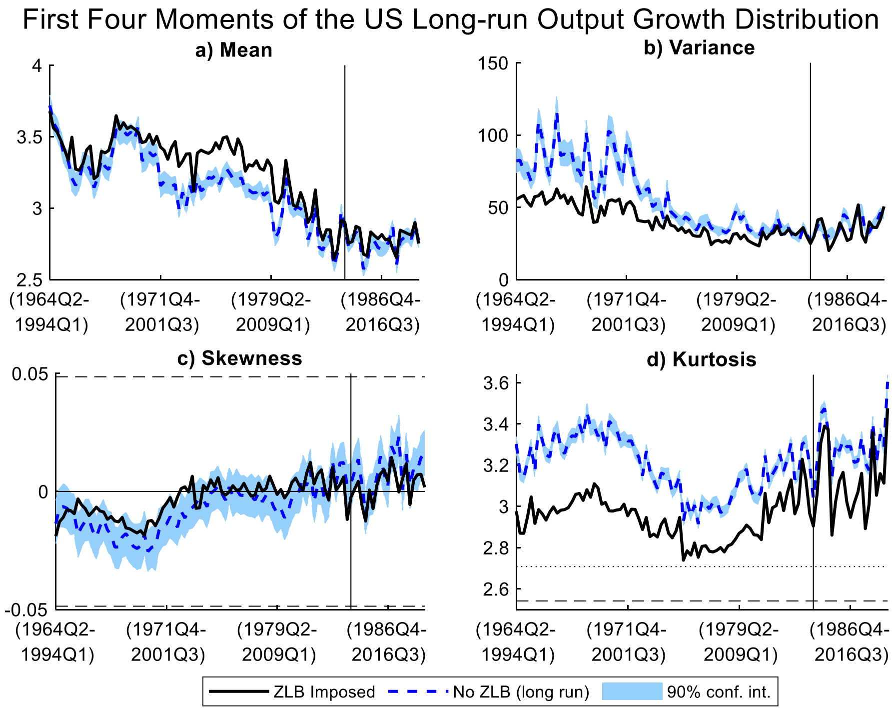

Fig. 27 Moments of the long-run output growth distribution of the counterfactual without imposing the ZLB. The x-axis indicates the range of quarters of the estimation window. The vertical line indicates the estimation window when the GFC enters the estimation (1979Q1-2008Q4)

References

Adrian T, Boyarchenko N, Giannone D (2019) Vulnerable growth. Am Econ Rev 109(4):1263–1289

Adrian T, Boyarchenko N, Giannone D (2021) Multimodality in macro-financial dynamics. Int Econ Rev. https://doi.org/10.1111/iere.12501

Adrian T, Grinberg F, Liang N, Malik S, Yu J (2022) The term structure of growth-at-risk. Am Econ J Macroecon 14(3):283–323

Aikman D, Bridges J, Hoke SH, O'Neil C, Raja A (2019) Credit, capital and crises: a GDP-at-risk approach. Staff working paper no. 824, Bank of England

Akinci O, Queralto A (2017) Credit spreads, financial crises, and macroprudential policy. Federal Reserve Bank of New York Staff Reports No. 802

Alessi L, Ghysels E, Onorante L, Peach R, Potter S (2014) Central bank macroeconomic forecasting during the global financial crisis: the European central bank and federal reserve bank of New York experiences. J Bus Econ Stat 32:483–500

Antolin-Diaz J, Drechsel T, Petrella I (2017) Tracking the slowdown in long-run GDP growth. Rev Econ Stat 99:345–356

Balke NS (2000) Credit and economic activity: credit regimes and nonlinear propagation of shocks. Rev Econ Stat 82(2):344–349

Banerjee R, Contreras J, Mehrotra A, Zampolli F (2020) Inflation at risk in advanced and emerging market economies. BIS working paper, no 883

Baumeister C, Benati L (2013) Unconventional monetary policy and the great recession: estimating the macroeconomic effects of spread compression at the zero lower bound. Int J Cent Bank 9(2):165–212

Bianchi F (2013) Regime switches, agents' beliefs, and post-world war II U.S. macroeconomic dynamics. Rev Econ Stud 80:463–490
Borio C, Drehmann M, Xia D (2020) Forecasting recession: the importance of the financial cycle. J Macroecon 66(585):103258
Brownlees C, Souza ABM (2021) Backtesting global growth-at-risk. J Monet Econ 118:312–330
Bu C, Rogers J, Wu W (2021) A unified measure of fed monetary policy shocks. J Monet Econ 118:331–349
Caldara D, Cascaldi-Garcia D, Cuba-Borda P, Loria F (2020) Understanding growth-at-risk: a Markov-switching approach, mimeo
Carriero A, Clark TE, Marcellino M (2020) Capturing macroeconomic tail risks with Bayesian vector autoregressions. Working paper 20-02, Federal Reserve Bank of Cleveland
Cecchetti S (2008) Measuring the macroeconomic risks posed by asset price booms. In: Campbell J (ed) Asset prices and monetary policy. University of Chicago Press for NBER
Chen CWS, Lee JC (1995) Bayesian inference of threshold autoregressive models. J Time Ser Anal 16:483–492
Chen CWS, So MKP (2006) On a threshold heteroscedastic model. Int J Forecast 22:73–89
Chung H, Laforte JP, Reifschneider D, Williams JC (2012) Have we underestimated the likelihood and severity of zero lower bound events? J Money, Credit, Bank 44(1):47–82
Cogley T, Morozov S, Sargent TJ (2005) Bayesian fan charts for U.K. inflation: forecasting and sources of uncertainty in an evolving monetary system. J Econ Dyn Control 29(11):1893–1925
Delle Monache D, De Polis A, Petrella I (2020) Modeling and forecasting macroeconomic downside risk. CEPR Discussion Papers 15109
Duprey T, Ueberfeldt A (2020) Managing GDP tail risk. Staff Working Paper 2020-03, Bank of Canada
Eguren-Martin F, O'Neill C, Sokol A, von dem Berge L (2020) Capital flows-at-risk: push, pull and the role of policy. Staff Working Paper No. 881, Bank of England
Estrella A, Trubin MR (2006) The yield curve as a leading indicator: some practical issues. Curr Issues Econ Finance 12(5):1–7
Fagiolo G, Napoletano M, Roventini A (2008) Are output growth-rate distributions fat-tailed? Some evidence from OECD countries. J Appl Economet 23:639–669
Favara G, Gilchrist S, Lewis KF, Zakrajšek E (2016) Updating the recession risk and the excess bond premium. FEDS Notes, Board of Governors of the Federal Reserve System, Washington. https://doi.org/10.17016/2380-7172.1836
Franta M, Gambacorta L (2020) On the effects of macroprudential policies on growth-at-risk. Econ Lett 196(C):109501
Gertler M, Karadi P (2015) Monetary policy surprises, credit costs, and economic activity. Am Econ J Macroecon 7(1):44–76
Geweke J, Keane M (2007) Smoothly mixing regressions. J Econometr 138:252–290
Gilchrist S, Zakrajšek E (2012) Credit spreads and business cycle fluctuations. Am Econ Rev 102(4):1692–1720
Groeneveld RA, Meeden G (1984) Measuring skewness and kurtosis. The Statistician 33:391–399
Hogg RV (1974) Adaptive robust procedures: a partial review and some suggestions for future applications and theory. J Am Stat Assoc 69:909–923
Holston K, Laubach T, Williams JC (2017) Measuring the natural rate of interest: international trends and determinants. J Int Econ 108(1):39–75
Hwang Y (2019) Forecasting recessions with time-varying models. J Macroecon 62(C):103153
Jordà Ò (2005) Estimation and inference of impulse responses by local projections. Am Econ Rev 95(1):161–182
Jung W, Lee JH (2019) Quantile impulse responses analysis with applications in macroeconomics and finance. Mimeo, University of Illinois
Kilian L, Manganelli S (2008) The central banker as a risk manager: estimating the federal reserve's preferences under Greenspan. J Money, Credit, Bank 40(6):1103–1129
Kim TH, Lee DJ, Mizen P (2019) Impulse response analysis in conditional quantile models and an application to monetary policy, mimeo. Yonsei University, Economic Research Institute
Kiss T, Nguyen H, Österholm P (2023) Modelling Okun's law: Does non-Gaussianity matter? Empir Econ 64(5):2183–2213
Liu P, Theodoridis K, Mumtaz H, Zanetti F (2019) Changing macroeconomic dynamics at the zero lower bound. J Bus Econ Stat 37(3):391–404

Loria F, Matthes C, Zhang D (2023) Assessing macroeconomic tails risk. Available at SSRN https://ssrn.com/abstract=4002665.
Perron P, Yamamoto Y (2022) The great moderation: updated evidence with joint tests for multiple structural changes in variance and persistence. Emp Econ 62(3):1193–1218
Plagborg-Møller M, Wolf CK (2021) Local projections and VARs estimate the same impulse responses. Econometrica 89(2):955–980
Plagborg-Møller M, Reichlin L, Ricco G, Hasenzagl T (2020) When is growth at risk? In: Brookings papers on economic activity. Spring, pp 167–229
Raftery AE, Lewis S (1992) How many iterations in the Gibbs sampler? In: Bernardo J, Berger J, Dawid AP, Smith AFM (eds) Bayesian statistics. Oxford University Press, pp 763–773
Romer CD, Romer DH (2004) A new measure of monetary shocks: derivation and implications. Am Econ Rev 94(4):1055–1084
Stock JH, Watson MW (2003) Has the business cycle changed and why? NBER chapters. In: NBER macroeconomics annual 2002, vol 17, pp 159–230
Villani M (2009) Steady-state priors for vector autoregressions. J Appl Economet 24:630–650
Wieland JF, Yang MJ (2020) Financial dampening. J Money, Credit, Bank 52(1):79–113