BANK OF CANADA
BANQUE DU CANADA

**Staff Working Paper/Document de travail du personnel — 2021-24**

Last updated: May 25, 2021

# Shaping the Future: Policy Shocks and the GDP Growth Distribution

by François-Michel Boire,[1] Thibaut Duprey[2] and Alexander Ueberfeldt[2]

[1] Western University
London, Ontario, Canada N6A 3K7

[2] Financial Stability Department
Bank of Canada, Ottawa, Ontario, Canada K1A 0G9

fboire2@uwo.ca; tduprey@bank-banque-canada.ca; aueberfeldt@bank-banque-canada.ca



ISSN 1701-9397 

# Acknowledgements

The views expressed in this paper are those of the authors and do not necessarily reflect those of the Bank of Canada. We thank Paul Beaudry, Gabriel Bruneau, Heng Chen, Bruno Feunou, Sermin Gungor, Kerem Tuzcuoglu, and seminar participants at the Bank of Canada, Queens University, and the Winter 2019 Midwest Macro Meetings for their suggestions and comments, as well as Nicolas Whitman and Vatya Kishore for excellent research assistance.

# Abstract

We incorporate quantile regressions into a structural vector autoregression model to empirically assess how monetary and fiscal policy influence risks around future GDP growth. Using a panel of six developed countries, we find that both policy instruments affect the location of the distribution of future GDP growth, whereas fiscal shocks also impact the shape of the distribution. Fiscal stimulus generates upside risk, paving the path to a faster recovery, especially when the policy rate is constrained by the zero lower bound (ZLB). Unconventional monetary policy during ZLB episodes has a comparable effect on future GDP growth as conventional monetary policy.



*Topics: Central bank research, Econometric and statistical methods, Financial stability, Fiscal policy, Monetary policy*
*JEL codes: C32, C53, E52, E62*

# 1 Introduction

When central banks lower interest rates or central governments boost spending during downturns, they aim to improve the economic outlook and avert the worst possible outcomes. As such, central banks usually report a main scenario and provide a risk assessment around that outlook. Counterfactual analyses based on econometric models can then help policy-makers form expectations about the outcomes of policy interventions. However, traditional vector autoregression (VAR) models are mean-focused and rely on a normality assumption, and are thus not suited to the task of detecting the impact of policy interventions on complex dynamics of the distribution like asymmetries and fat tails.

We incorporate a conditional quantile representation into a vector autoregression to investigate the dynamics of the entire distribution of GDP growth following policy shocks. On one hand, the VAR structure allows us to characterize economic dynamics over time and give a structural meaning to the monetary and fiscal shocks. On the other hand, we model the distribution of GDP growth, our variable of interest, through quantile regressions (QR) to allow for more flexible distributions. Thus, we can capture heteroskedasticity, skewness and even multi-modality of GDP growth as the quantile estimators are distribution-free. The distribution-free quantiles of the distribution of GDP growth are recursively projected one step ahead and smoothed with a kernel to simulate the entire distribution of GDP growth. Thus, we can investigate how policy shocks shape the future GDP growth distribution. In Monte Carlo simulations, we find that whenever the variable of interest is governed by a skewed distribution, choosing our QR-VAR over a VAR dramatically reduces the bias and RMSE for the one-step-ahead conditional skewness. This only comes at the cost of a small positive bias in the one-step-ahead conditional variance.

In this paper, we apply our new quantile regression–vector autoregression hybrid model (QR-VAR) to assess how monetary and fiscal policies impact risks around future GDP growth. Our results thus provide an answer to the important question: How do fiscal and monetary policies change the odds of a good or bad realization of GDP growth?

We estimate our QR-VAR on a panel of six countries from 1964Q1 to 2019Q4 at the quarterly frequency: Australia, Canada, Finland, Japan, the United Kingdom and the United States. We use a panel estimation to ensure that a wide range of crises and a policy actions are represented in the data. Specifically, we focus on the impact of quantitative easing when policy rates are at the zero lower bound (ZLB). We find that by allowing for heterogeneous quantile effects with the QR-VAR we obtain a better root mean squared error (RMSE) than with a VAR model. This improvement is mostly due to crises episodes, while a regular VAR would perform better in normal times.

We are the first to track the impact of monetary and fiscal policy shocks on each quantile of the density forecast of GDP growth. Our setup allows for the identification of policy shocks as part of the same quantile framework, with traditional sign or zero restrictions. This is in contrast to Montes-Rojas (2019) and Duprey and Ueberfeldt (2020), who require structural shocks from a VAR computed outside their quantile model, and Chavleishvili and Manganelli (2019), who produce only reduced form shocks. We find that monetary policy shocks mostly shift the distribution of future

GDP growth up or down, without significantly changing the shape of the distribution. Conversely, fiscal policy shocks change the shape of the distribution of future GDP growth. A positive fiscal shock can improve the lower tail of the distribution of GDP growth, but the improvement is stronger for the upper tail. Our model therefore suggests that a fiscal expansion both mitigates the likelihood of severe downturns and increases the likelihood of a swift recovery. Because of the asymmetric effect of fiscal policy shocks, a VAR that does not account for the heterogeneity in the quantile response to fiscal shocks would overestimate the mean response.

We further explore the state dependence of quantile responses by accounting for ZLB events. In non-ZLB periods, the distribution of GDP growth is not significantly responsive to policy shocks. During ZLB events, however, we observe two effects. First, the lower tail of GDP growth is more responsive to fiscal shocks during ZLB than non-ZLB times. Second, the upper tail of GDP growth significantly increases in response to expansionary fiscal shocks, suggesting that fiscal policy is especially useful to speed up the recovery during ZLB events. In addition, we show that the type of government spending can influence the shape of the distribution of GDP growth. Central government spending creates more asymmetry than local government spending.

We contribute to two strands of the literature. First, our hybrid QR-VAR contributes to recent efforts to capture the asymmetry and skewness of certain macroeconomic variables (credit growth: Ranciere et al., 2008; GDP growth: Adrian et al., 2019 and International Monetary Fund, 2017) in autoregressive macroeconomic models. Some existing methods already rely on quantile regressions in a VAR framework. The conditional autoregressive value at risk (CAViaR) model of Engle and Manganelli (2004), extended to the multivariate case by White et al. (2015), makes the assumption that the underlying distribution is governed by scale effects. Chavleishvili and Manganelli (2019) further cast this approach in a VAR framework where the entire system is modelled in terms of quantiles, at the cost of using reduced-form quantile shocks that cannot be interpreted as policy shocks. Conversely, we approximate the distribution of one variable with QR, thereby preserving the interpretability of structural shocks for policy purposes. Schüler (2014) identifies reduced-form quantile shocks from the quantile covariance matrix. Linnemann and Winkler (2016) estimate a VAR with QR to track the quantiles but do not generate the whole conditional distribution of GDP growth to draw simulations from it. Duprey and Ueberfeldt (2020) combine VAR and QR in two separate steps, such that it does not allow for the projected quantiles to feed back into the VAR projection. Han et al. (2019) and Loria et al. (2019) use instead the local projection (Jordà, 2005) with QR to compute the quantiles of GDP growth at each point in time, but this does not allow one to compute the full density distribution of GDP growth and simulate the dynamics of the variables over time. We contribute to this growing literature by allowing structural shocks to modify the shape of the distribution of one variable by recovering its full conditional distribution, at the cost of retaining the standard VAR structure for all other variables. Alternative methods relying on the VAR framework move away from QR and relax other assumptions. Some approaches add stochastic volatility to the Bayesian VAR (BVAR) (Cogley et al., 2005; Primiceri, 2005) and highlight that negative skewness can be generated by a symmetric conditional model if there are simultaneous negative shifts in conditional means and positive shifts in conditional variances (Carriero et al., 2020). Other work deviates from the normality of the shocks, e.g. BVAR with non-Gaussian disturbances (Chiu et al., 2017) or Copula-VAR with skew-t marginals (Smith and Vahey, 2016). Eventually, asymmetries can be introduced with regime-specific dynamics that depend on latent (Markov-switching VAR; Hubrich and Tetlow, 2015) or observable regime switches (threshold VAR; Balke, 2000). We propose a flexible yet tractable approach to determine a representation of the possibly non-normal, asymmetric and multi-modal true data-generating process.

Second, our empirical application contributes to the literature on the state dependence of the impact of policy shocks. Only Linnemann and Winkler (2016), Duprey and Ueberfeldt (2020) and Loria et al. (2019) partly explore the quantile response to some policy shocks. Linnemann and Winkler (2016) find that fiscal policy improves lower tail GDP risk in the United States. Duprey and Ueberfeldt (2020) find that, for Canada, monetary policy does not change the lower tail of GDP growth more than the median, while macroprudential policy does. Loria et al. (2019) find that contractionary monetary policy shocks disproportionately increase downside risk in the United States. Although the literature on quantile responses to policy shocks is scarce, some of our results are consistent with existing work. In line with Sims and Wu (2020), our results suggest that conventional monetary policy and quantitative easing—identified with the shadow rate à la Wu and Xia (2016)—have a comparable impact. Our results on fiscal policy are consistent with the empirical work of Ramey and Zubairy (2018) and the modelling of Christiano et al. (2011), who show that fiscal spending has a larger impact on GDP growth during ZLB episodes. Our empirical results imply that fiscal multipliers are likely larger during ZLB episodes.[1] In the spirit of Auerbach and Gorodnichenko (2012a), we also find that fiscal policy leads to a strong reaction in the lower tail of the GDP growth distribution, which corresponds to recession states. Since ZLB events and the lower tail of GDP growth correspond to crises, our results are in line with Canova and Pappa (2011), who argue that fiscal policy has the largest impact on GDP growth when acting in a counter-cyclical manner.

The remainder of the paper is organized as follows. Section 2 describes a bivariate version of our hybrid QR-VAR to highlight the mechanics. Section 3 applies the QR-VAR to a panel of countries to investigate the quantile responses to monetary and fiscal shocks. Section 4 further investigates the state dependence of the quantile response during ZLB events. Section 5 concludes.

## 2 The model

For expositional purposes, we introduce our hybrid quantile VAR in the context of a simple bivariate process with one lag. In the empirical analysis of the paper, we estimate a multivariate version

[1] Most studies find that the fiscal multiplier is below one, with limited benefits from a fiscal expansion (Blanchard and Perotti, 2002; (Perotti, 2005); Galí et al., 2007; Ramey, 2011), with a multiplier closer to zero for open economies with flexible exchange rates (Ilzetzki et al., 2013). However, as pointed out by Ramey and Zubairy (2018), the precise computation of fiscal multipliers in a VAR framework relies on many ad hoc assumptions, especially for cross-country VAR setups estimated on time series in growth rate form. Thus, we refrain from emphasizing a specific fiscal multiplier value and instead make a relative statement about fiscal multipliers during ZLB episodes.

with several lags, possibly with interaction terms, but the same principles apply.

### 2.1 A hybrid of QR and structural VAR

Consider a bivariate vector stochastic process for variables $Y_1$ and $Y_2$ that evolves as follows, where upper cases represent random variables and lower cases the realizations:[2]

$$\begin{aligned} y_{1,t} &= \mathbf{\Gamma}_{1,1} y_{1,t-1} + \mathbf{\Gamma}_{1,2} y_{2,t-1} + \epsilon_{1,t} \\ y_{2,t} &= \mathbf{\Gamma}_{2,1} y_{1,t-1} + \mathbf{\Gamma}_{2,2} y_{2,t-1} + \epsilon_{2,t} \end{aligned} \tag{1}$$

In the standard VAR, the error term $\epsilon_t = (\epsilon_{1,t}, \epsilon_{2,t})'$ is such that $\mathbb{E}(\epsilon_t) = 0$ and $\mathbb{E}(\epsilon'_{t_i}\epsilon_{t_j}) = \mathbf{\Sigma}_\epsilon \mathbb{1}_{i=j}$, where $\mathbf{\Sigma}_\epsilon$ is a positive definite variance-covariance matrix.

Assuming $Y_1$ is governed by a conditional density distribution $f$, we replace the first equation of the VAR in Equation (1) to obtain our hybrid system of equations:

$$\begin{aligned} y_{1,t} &= \mathbb{E}(f(y_{1,t-1}, y_{2,t-1})) + \epsilon_{1,t}, \\ y_{2,t} &= \mathbf{\Gamma}_{2,1} y_{1,t-1} + \mathbf{\Gamma}_{2,2} y_{2,t-1} + \epsilon_{2,t}. \end{aligned} \tag{2}$$

The residual of the first equation is no longer a classical regression residual and is obtained as $\epsilon_{1,t} \sim \hat{f}(\boldsymbol{y}_{t-1}) - \mathbb{E}[\hat{f}(\boldsymbol{y}_{t-1})]$, where $\boldsymbol{y}_{t-1}$=$\{y_{1,t-1}, y_{2,t-1}\}$. By computing a function $f$ that moves away from a linear relationship, we expect to be able to analyze the potential for nonlinear amplification arising from a non-standard distribution of variable $Y_1$. One advantage of our modelling strategy is the flexibility with which we characterize distributions in a multivariate setup. We focus on the evolution of the distribution of one variable only, while other variables are modelled with linear regressions as in a standard VAR to reduce the complexity of the estimation.

We now turn to the estimation of the function $f$ by relying on quantile regressions and kernel density distributions.

### 2.2 Conditional density distribution

The conditional quantiles of the density distribution $f$ of variable $Y_1$ can be explicitly modelled with quantile regressions that map $\mathbf{Y}_{t-1}$=$\{Y_{1,t-1}, Y_{2,t-1}\}$ to the conditional quantiles of $Y_{1,t}$. Quantile processes capture heteroskedasticity and skewness, which go beyond parametric forms of asymmetry (e.g. asymmetric GARCH). However, quantile regressions do not directly produce the conditional distribution so another step is needed to recover the conditional density out of the fitted quantiles.

[2] For simplicity, we omit the constant term throughout the paper unless otherwise specified.

### 2.2.1 Conditional quantile estimation

We rely on Koenker and Bassett Jr (1978) and estimate linear conditional QR of $Y_1$ to estimate the quantiles of order $\tau \in \mathcal{T}$, where $\tau$ is a probability associated with a quantile:

$$y_{1,t} = \boldsymbol{\beta}_\tau \mathbf{y}_{t-1} + \nu_{\tau,t}. \tag{3}$$

The quantile-specific coefficients are denoted by a $\tau$ subscript. Now consider $Q_\tau(\cdot)$ the quantile operator of order $\tau$. The quantile-specific residuals $\nu_\tau$ are such that $Q_\tau(\nu_\tau) = 0$, which means that the only constraint of the QR model is that the quantile of order $\tau$ of the error term should be zero, and thus $\nu_\tau$ is otherwise distribution-free. We then obtain the conditional quantile model:

$$Q_\tau(Y_{1,t}|\mathbf{y}_{t-1}) = \boldsymbol{\beta}_\tau \mathbf{y}_{t-1}. \tag{4}$$

The quantile-specific coefficients $\boldsymbol{\beta}_\tau$ can be interpreted as the marginal rate of change of $Q_\tau(Y_{1,t}|\mathbf{Y}_{t-1})$ with respect to $\mathbf{Y}_{t-1}$. The quantile regression allows the response of the quantiles to differ in different regions of the distribution, which we will hereafter refer to as heterogeneous effects. To estimate the parameters $\boldsymbol{\beta}_\tau$ of Equation (4), we solve the following linear programming problem:

$$\hat{\boldsymbol{\beta}}_\tau = \underset{\boldsymbol{\beta}}{\operatorname{argmin}} \sum_{t=1}^{T} \rho_\tau(y_{1,t} - \boldsymbol{\beta}\mathbf{y}_{t-1}), \tag{5}$$

where $\rho_\tau(u) = u \cdot (\tau - \mathbb{1}_{u<0})$ is the objective function based on quantile-weighted absolute residuals. The solution to the first order conditions yields $\Pr(Y_{1,t} \leq \hat{\boldsymbol{\beta}}_\tau \mathbf{y}_{t-1}) = \tau$ and we indeed obtain the $\tau$th fitted quantile of $Y_1$.

To move closer to recovering the whole distribution of $Y_1$, we fit multiple quantile regressions of orders $\mathcal{T} = \{0.10, 0.11, ..., 0.90\}$ to obtain a set of conditional quantile predictions $Q_{\mathcal{T}} = \{Q_{0.10}, Q_{0.11}, \ldots, Q_{0.90}\}$. We truncate the top and bottom deciles of the probability grid because the estimation of extreme quantiles can be imprecise.

### 2.2.2 Recover the distribution from the quantiles

The resulting set of conditional quantile predictions $\hat{Q}_{\mathcal{T}}$ can be interpolated to obtain a smooth density distribution to be substituted back into Equation (2):

$$\hat{f}(y_{1,t}|\hat{Q}_{\mathcal{T}}(Y_{1,t}|\boldsymbol{y}_{t-1})).$$

A challenge of time-varying quantiles is the possibility that quantile predictions may cross, which directly contradicts the monotonicity of quantiles with respect to $\tau$.[3] Following Chernozhukov et al. (2010), we circumvent the issue of quantile crossings by sorting conditional quantile predic-

[3] This is a well-known consequence of the linear form of the conditional quantiles in Koenker's estimator. Quantiles are more likely to intersect in the periphery of the data's centre and, as a result, crossings are potentially more frequent for rare events like systemic crises.

tions. Although the model parameters remain unchanged, this method efficiently solves the crossing problem and preserves desirable asymptotic properties of the quantile estimator.[4]

Given the sorted conditional quantiles $\hat{Q}_{\mathcal{T}}(Y_{1,t}|\boldsymbol{y}_{t-1})$, we estimate the density $\hat{f}(y_{1,t}|\hat{Q}_{\mathcal{T}}(Y_{1,t}|\boldsymbol{y}_{t-1}))$ using kernel density estimation (KDE), following Gaglianone and Lima (2012), Davino et al. (2013), and Korobilis (2017). Our conditional density estimate takes the form:

$$\hat{f}(y_{1,t}|\hat{Q}_{\mathcal{T}}(Y_{1,t}|\boldsymbol{y}_{t-1})) = \frac{1}{n}\sum_{\tau\in\mathcal{T}}\frac{1}{\mathcal{H}\lambda_\tau}K\left(\frac{y_{1,t}-\hat{Q}_\tau(Y_{1,t}|\boldsymbol{y}_{t-1})}{\mathcal{H}\lambda_\tau}\right), \tag{6}$$

where the kernel $K(\cdot)$ is a well-defined distribution function, and $\mathcal{H} > 0$ is the bandwidth. We use the adaptive kernel of Koenker and Portnoy (1987) with $\lambda_\tau$ a local bandwidth factor defined as:

$$\lambda_\tau = \left(\frac{g(\hat{Q}_\tau(Y_{1,t}|\boldsymbol{y}_{t-1}))}{\tilde{f}(\hat{Q}_\tau(Y_{1,t}|\boldsymbol{y}_{t-1}))}\right)^\alpha,$$

where $\tilde{f}(\cdot)$ is a pilot density estimate *à la* Silverman (2018), $\alpha \in (0,1)$ is a sensitivity parameter that we set to 0.5 as suggested by Davino et al. (2013), and $g(\cdot) = f'(\cdot)$. Kernel density estimates are highly flexible and allow for the possibility of multi-modal distributions.[5] But the selection of kernel families and bandwidths requires empirically guided choices, such that practitioners often rely on rules of thumb.

For a quantile-based density estimation, Gaglianone and Lima (2012) use Epanechnikov kernels, and Korobilis (2017) use Gaussian kernels.

Since we apply our framework to macroeconomic data in growth rate space, a desirable property of the KDE is that the support of the response is unbounded. That is, we want a non-zero density estimate for all possible outcomes in $\mathbb{R}$, and thus we favour the Gaussian kernel.

Now we can compute the conditional mean of $Y_1$ via numerical integration of the KDE. Given a density estimate $\hat{f}(y_{1,t}|\hat{Q}_{\mathcal{T}}(Y_{1,t}|\boldsymbol{y}_{t-1}))$ with support $\mathbb{R}$, the point forecast is computed via numerical integration:

$$\hat{\mathbb{E}}_{t-1}(Y_{1,t}) \equiv \mathbb{E}[Y_{1,t}|\hat{f}(y_{1,t}|\hat{Q}_{\mathcal{T}}(Y_{1,t}|\boldsymbol{y}_{t-1}))] = \int_{\mathbb{R}} y_{1,t}\hat{f}(y_{1,t}|\hat{Q}_{\mathcal{T}}(Y_{1,t}|\boldsymbol{y}_{t-1}))dy_{1,t}. \tag{7}$$

Altogether, we obtain an algorithm to produce conditional distribution forecasts of variable $Y_1$.

The Monte Carlo exercise shows that, whenever the variable of interest is governed by a skewed distribution, the bias for the one-step-ahead conditional skewness is reduced by a factor of ten

[4] Several other solutions to the crossing problem have been proposed, usually more data-intensive, namely by He (1997) and Bondell et al. (2010), who impose a set of restrictions on the quantile process to preserve monotonicity, and Schmidt and Zhu (2016), who model the spacings between adjacent quantiles as positive functions.

[5] Another way to construct smooth density estimates from quantiles consists of interpolating the quantiles of a skew-t distribution, as seen in Adrian et al. (2019) and Ghysels et al. (2018). The skew-t distribution family encompasses a wide range of unimodal, asymmetric, and fat-tailed distributions. The present empirical application is robust to this alternative density estimation approach. QR-based density estimation has also been addressed by Koenker (2005), who suggests a piecewise linear interpolation of the density function, and Schmidt and Zhu (2016), who consider a simple interpolation of the quantiles to recover a smooth density.

for the QR-VAR compared to the VAR, and the RMSE is reduced by a factor of almost three (see Appendix A). This only comes at the cost of a small positive bias in the one-step-ahead conditional variance, and a one percent increase in the RMSE of the one-step-ahead conditional variance. The QR-VAR is therefore particularly useful when the variable of interest modelled by the quantile regressions is heteroskedastic and skewed. When the data-generating process does not exhibit skewness, then the VAR is unsurprisingly better, since it assumes zero conditional skewness already.

### 2.3 Shock identification

To achieve structural identification in the context of Equation (2), we assume that there is no contemporaneous correlation between the quantiles of $Y_{1,t}$ and the rest of the variables in the system. This is equivalent to saying that, on impact, a shock shifts the entire conditional distribution of $Y_1$ homogeneously.[6] After the initial impact, the realizations of variable $Y_1$ are drawn from the conditional density distribution with heterogeneous quantile effects.

With this ad-hoc approach, we can still perform the standard structural decomposition of the residuals of our hybrid, by combining the prediction error $\epsilon_{1,t}$ and the linear regression residual $\epsilon_{2,t}$. We impose a set of restrictions to identify a matrix $\mathbf{C}$ such that $\boldsymbol{\epsilon}_t = \mathbf{C}\mathbf{u}_t$, $\boldsymbol{\Sigma}_\epsilon = \mathbf{C}\boldsymbol{\Sigma}_u\mathbf{C}'$ and $\boldsymbol{\Sigma}_u$ is diagonal. This yields a structural VAR representation where the elements of innovations $u_t = \{u_{1,t}, u_{2,t}\}$ are orthogonal. Structural identification approaches like Cholesky decomposition, sign restrictions or zero restrictions can be used.

### 2.4 Bootstrapped confidence intervals

We bootstrap the residuals to obtain confidence for our hybrid approach. Since we use a non-parametric specification of the distribution of variable $Y_1$, confidence bands should be distribution-free as well. In addition, usual methods to compute confidence intervals differ for OLS and quantile regressions, which warrants a unified bootstrap approach to obtain confidence bands.

From Equation (2), we resample the same number of residuals $\tilde{\epsilon} = \{\tilde{\epsilon}_1, \tilde{\epsilon}_2\}$ as in the initial dataset by blocks of size $S$. The block resampling procedure aims to preserve any residual autocorrelation in the residuals, for instance, around a financial crisis, which is likely to affect the tail of the distribution relatively more so than the mean. If the model is estimated for a panel of countries, we cluster the resampling within each country. We then recursively construct the new variables $\tilde{\boldsymbol{Y}} = \{\tilde{Y}_1, \tilde{Y}_2\}$, given the initial estimates $\hat{\boldsymbol{\Gamma}}$ and $\hat{\boldsymbol{\beta}}_\tau$ with the recursion initialized by $\tilde{\boldsymbol{y}}_0 = \boldsymbol{y}_0$, where:

$$\begin{aligned} \tilde{y}_{1,t} &= \mathbb{E}\left(\hat{f}(y_{1,t}|\hat{Q}_{\mathcal{T}}(\tilde{Y}_{1,t}|\tilde{\boldsymbol{y}}_{t-1};\hat{\boldsymbol{\beta}}_{\mathcal{T}}))\right) + \tilde{\epsilon}_{1,t} \\ \tilde{y}_{2,t} &= \hat{\boldsymbol{\Gamma}}_{2,1}\tilde{y}_{1,t-1} + \hat{\boldsymbol{\Gamma}}_{2,2}\tilde{y}_{2,t-1} + \tilde{\epsilon}_{2,t}. \end{aligned} \tag{8}$$

[6] Strictly speaking, a shock is an unexpected surprise to a realization of a variable. It is unclear how one can generalize the concept of shock to be the unexpected surprise to the unobserved quantile of a variable, so we leave this issue to future research.

We re-estimate the model for $\tilde{\boldsymbol{Y}}$ and collect the new estimates $\tilde{\boldsymbol{\Gamma}}$, $\tilde{\boldsymbol{\beta}}_\tau$ and $\tilde{\boldsymbol{C}}$, allowing us to compute confidence bands around the expectation of $\boldsymbol{Y}$ and the quantiles of $Y_1$.

Note that this makes our work more challenging than one might expect. Confidence bands usually provide the uncertainty around the expected mean. But quantile estimates provide the uncertainty regarding the whole distribution of the variable $Y_1$, not only its mean. In addition, by bootstrapping the quantile estimates themselves, we can compute the uncertainty of the whole distribution of $Y_1$. Essentially, we can compute the uncertainty around the confidence bands of the variable $Y_1$, implying larger bands compared with a standard VAR analysis.

### 2.5 Impulse response

Our QR-VAR hybrid allows for the computation of quantile impulse responses of one variable of interest. But unlike in a regular structural VAR, the impulse response functions (IRFs) of the QR-VAR are not linear[7] and hence not analytically tractable. The one-period-ahead forecast $\hat{\mathbb{E}}_t(Y_{1,t+1})$ in Equation (7) requires an intermediary step that involves the computation of a conditional density. As a result, it is not feasible to compute the $n$-steps-ahead forecast with the recursive approach where we replace the conditioning variables by the one-period-ahead forecast as one would do in a linear structural VAR.

Hence, we resort to Monte Carlo simulations to compute the $n$-steps-ahead forecast. Although this approach can be computationally taxing for long-term forecasts, it is straightforward to implement controlled experiments by drawing from the estimated density distribution and conditioning on a user-chosen policy shock.

The steps to compute the generalized impulse response function (GIRF) and the quantile impulse response functions (QIRF) are described by the algorithm in Appendix E. For each bootstrapped dataset, we re-estimate the model and compute the impulse response function. We then take the median and upper 5th and lower 5th percentiles to compute the generalized impulse response and its confidence bands. When re-estimating the model, we also obtain the change in the percentiles of the distribution of GDP growth following a shock. For each percentile, the quantile impulse response is the mean change of the percentile of the GDP growth distribution across all bootstrapped samples.

## 3 Impact of monetary and fiscal shocks on the GDP growth outlook

We now apply our hybrid QR-VAR to analyze the nonlinear dynamic response of the density forecast of GDP growth following monetary and fiscal policy shocks. We assess if those policy

[7] For nonlinear VAR, as in our case, the impulse response function depends on initial conditions, is asymmetric with respect to the sign of the impulse, and does not scale to the magnitude of the impulse. See Koop et al. (1996), Gallant et al. (1993), and Potter (2000) for a comprehensive discussion on impulse response functions in nonlinear settings.

shocks can modify the shape of the future distribution of GDP growth via their differential impacts on quantiles. A particular question of interest is: Can policies change the odds of a good or bad realization of GDP growth?

If policy shocks impact all quantiles of the future GDP growth distribution homogeneously, then only the location of the density forecast changes and a mean model with a symmetry assumption is well suited.

### 3.1 Model estimation setup

We estimate a multivariate version of the hybrid QR-VAR model of Equation (2) with five variables and $L = 3$ lags.[8] We use data from 1964Q1 to 2019Q4, listed in Table 6. The quarterly annualized log-growth of real GDP per capita ($\Delta y$) is modelled by quantile regressions, so that we can focus on the nonlinear density forecast of GDP growth. Four other variables are modelled using linear relationships: quarterly annualized log-growth of public sector expenditures ($\Delta g$), the short-term nominal shadow rates are transformed to year-over-year changes to ensure stationarity ($\Delta r$),[9] quarterly annualized log-growth of the consumer price index ($\pi$) and the log of the CLIFS financial stress indices from Duprey et al. (2017) ($\varphi$). The stress index has important implications for the left tail of GDP growth, as pointed out by Adrian et al. (2019). All series are seasonally adjusted.

We use a panel dataset that includes six developed countries: Australia, Canada, Finland, Japan, the United Kingdom and the United States. The use of panel data allows for a better estimation of tail dynamics during crisis episodes and thus provide a better assessment of the potentially heterogeneous effects of monetary policy around 2008. Adding Finland supplements the data with a severe financial crisis period in the early 1990s. Similarly, adding Japan allows for a wider heterogeneity in monetary policy in the context of sustained low rates accompanied by public spending expansions. Adding Canada brings an episode of low rates without quantitative easing while the United States, the United Kingdom, Japan and Finland all experienced quantitative easing. Adding Australia also brings episodes of quantitative easing around 2008, though without reaching the zero lower bound. See Appendix D for further details about the data.

[8] See lag specification tests in Appendix B and location-scale specification tests in Appendix C. We also conducted a number of robustness checks, including a version of the model with the short-term rate in levels. See Appendix F for more details.

[9] When bootstrapping, we require that the new sample have a similar frequency of positive and negative policy rate changes of plus or minus 10 percent as the initial sample. We do this because the higher frequency of negative changes, due to a decreasing trend, is associated with the great moderation and a change in the frequency of extreme values in GDP growth, which would likely affect our ability to correctly bootstrap the tails of the GDP growth distribution.

$$
\begin{aligned}
\Delta y_t &= f(\Delta y_{t-1}, \Delta g_{t-1}, \Delta r_{t-1}, \pi_{t-1}, \varphi_{t-1}, \alpha_c; \boldsymbol{\beta}_\tau(L)) + \epsilon_{1,t} \qquad (9) \\
\Delta g_t &= \boldsymbol{\Gamma}_{2,0}\alpha_c + \sum_{l=1}^{L} \Big[\boldsymbol{\Gamma}_{2,1,l}\Delta y_{t-l} + \boldsymbol{\Gamma}_{2,2,l}\Delta g_{t-l} + \boldsymbol{\Gamma}_{2,3,l}\Delta r_{t-l} + \boldsymbol{\Gamma}_{2,4,l}\pi_{t-l} + \boldsymbol{\Gamma}_{2,5,l}\varphi_{t-l}\Big] + \epsilon_{2,t} \\
\Delta r_t &= \boldsymbol{\Gamma}_{3,0}\alpha_c + \sum_{l=1}^{L} \Big[\boldsymbol{\Gamma}_{3,1,l}\Delta y_{t-l} + \boldsymbol{\Gamma}_{3,2,l}\Delta g_{t-l} + \boldsymbol{\Gamma}_{3,3,l}\Delta r_{t-l} + \boldsymbol{\Gamma}_{3,4,l}\pi_{t-l} + \boldsymbol{\Gamma}_{3,5,l}\varphi_{t-l}\Big] + \epsilon_{3,t} \\
\pi_t &= \boldsymbol{\Gamma}_{4,0}\alpha_c + \sum_{l=1}^{L} \Big[\boldsymbol{\Gamma}_{4,1,l}\Delta y_{t-l} + \boldsymbol{\Gamma}_{4,2,l}\Delta g_{t-l} + \boldsymbol{\Gamma}_{4,3,l}\Delta r_{t-l} + \boldsymbol{\Gamma}_{4,4,l}\pi_{t-l} + \boldsymbol{\Gamma}_{4,5,l}\varphi_{t-l}\Big] + \epsilon_{4,t} \\
\varphi_t &= \boldsymbol{\Gamma}_{5,0}\alpha_c + \sum_{l=1}^{L} \Big[\boldsymbol{\Gamma}_{5,1,l}\Delta y_{t-l} + \boldsymbol{\Gamma}_{5,2,l}\Delta g_{t-l} + \boldsymbol{\Gamma}_{5,3,l}\Delta r_{t-l} + \boldsymbol{\Gamma}_{5,4,l}\pi_{t-l} + \boldsymbol{\Gamma}_{5,5,l}\varphi_{t-l}\Big] + \epsilon_{5,t}
\end{aligned}
$$

We introduce a dummy variable for each of our six countries, represented by $\alpha_c$, as in Kato et al. (2012) for the quantile regression.[10] The function $f(\Delta y_{t-1}, \Delta g_{t-1}, \Delta r_{t-1}, \pi_{t-1}, \varphi_{t-1}, \alpha_c; \boldsymbol{\beta}_\tau(L))$, where $\boldsymbol{\beta}_\tau(L)$ is a lagged polynomial of order $L$, is the kernel density estimate of GDP growth recovered from the conditional quantiles, where each quantile associated with probability $\tau \in \mathcal{T}$ is estimated as:

$$
Q_\tau(\Delta y_t) = \boldsymbol{\beta}_{\tau,0}\boldsymbol{\alpha}_{\tau,c} + \sum_{l=1}^{L} \Big[\boldsymbol{\beta}_{\tau,1,l}\Delta y_{t-l} + \boldsymbol{\beta}_{\tau,2,l}\Delta g_{t-l} + \boldsymbol{\beta}_{\tau,3,l}\Delta r_{t-l} + \boldsymbol{\beta}_{\tau,4,l}\pi_{t-l} + \boldsymbol{\beta}_{\tau,5,l}\varphi_{t-l}\Big]. \qquad (10)
$$

### 3.1.1 Fiscal and monetary policy shock identification

The five series of residuals $\boldsymbol{\epsilon}_t$ are decomposed into structural shocks by combining sign restrictions and zero restrictions in Table 1. We identify monetary policy shocks by assuming that they are negatively associated with GDP growth and inflation (see e.g. Uhlig, 2005).[11] We identify government spending shocks by using the Blanchard and Perotti (2002) identification scheme with zero restrictions, which assumes that government spending does not respond contemporaneously to shocks as it usually takes time for spending changes to materialize.[12] Demand and supply shocks are also identified to net out their contribution from other shocks. Note that if government spending is ordered first, we can simply take the Cholesky decomposition to obtain the zero restrictions and look for 1,000 random ortho-normal rotations of the subset to ensure other signs restrictions are met.

[10] Other methods than Koenker and Bassett Jr (1978) could be used for the panel quantile regression, like Canay (2011), that requires the assumption that the country fixed effect is the same across quantiles; Koenker (2004), with regularization methods especially useful for the case of a large number of fixed effects; or Machado and Santos Silva (2019). We have only a limited cross-section of six countries, for 224 quarters.

[11] This is a widely used type of identification scheme for monetary policy. Alternatively, one can use the narrative approach of monetary policy news shocks by Romer and Romer (2010).

[12] This is the most common type of identification scheme for government spending shocks. A popular alternative relies on the narrative identification of military spending news shocks by Ramey and Zubairy (2018), or sign restriction specifications (Mountford and Uhlig, 2009).

Table 1: Sign and zero restrictions for a structural identification scheme

| | | Spending shock | Supply shock | Monetary shock | Financial shock | Demand shock |
|---|---|---|---|---|---|---|
| Contemporaneous response of variable... | $\Delta g$ | + | 0 | 0 | 0 | 0 |
| | $\pi$ | | + | - | | + |
| | $\Delta r$ | | + | + | | + |
| | $\varphi$ | | | | + | |
| | $\Delta y$ | | - | - | | + |

### 3.2 Monetary policy shocks and the risk to GDP growth outlook

To understand the implications of monetary policy shocks for the anticipated GDP growth distribution, we analyze the implications of a one-period unanticipated decrease of 50 basis points in the short-term nominal rate (or the shadow rate, when applicable).

First, we focus on the impulse responses of different quantiles of the GDP growth distribution; see Figure 1, left column. Monetary policy easing increases all quantiles of GDP growth by a homogeneous amount and the difference across quantiles is not statistically significant. The response of the quantiles (plain line) is very similar to the response of the VAR model (star dotted line). This means that monetary policy incurs a homogeneous shift of the distribution of output growth forecasts but does not alter the shape in a meaningful way. The resulting effect can therefore be summarized as a change in the mean, with little to no change in variance or skewness. Recall that conditional quantiles do not respond to the impulse at the moment of impact (i.e. at period zero). Rather, the one-step-ahead forecasts of conditional quantiles will respond to the impulse. Indeed, a shock on output growth is defined as a term that is added to a random draw from the estimated density distribution, but the distribution from which we draw does not depend on current information. After impact, the density forecast of GDP can change as the initial shock starts propagating through the quantile-based kernel density approximation mechanism.

Next, we focus on the difference between the IRFs of our hybrid QR-VAR and that of a standard VAR. To obtain this comparison, we turn off the possibility of heterogeneous quantile effects following monetary policy shocks. As seen in the left column of Figure 2, the mean impulse response is similar in both cases, suggesting that a mean model reasonably approximates the impact of monetary policy on the distribution of GDP growth.

### 3.3 Fiscal policy shocks and the GDP growth outlook risk

To understand the implications of a fiscal shock, we consider a punctual 5 percent increase in the central government's current expenditures, unless stated otherwise. In the middle column of Figure 1, we see that fiscal stimulus significantly increases the top quantiles of GDP growth, whereas the rest of the distribution does not respond significantly. This asymmetric response suggests that fiscal policy changes the shape of the outlook distribution by making positive surprises more likely.

In other words, an expansionary fiscal policy increases the skewness coefficient of the distribution of GDP growth.

The low benefit of fiscal stimulus for the quantiles other than the top ones could result from the lack of coordination between monetary and fiscal policies. Empirically, if fiscal stimulus increases output, it leads to an endogenous monetary tightening, partly undoing the stimulus benefit of fiscal policy. To obtain the pure effect of fiscal spending, we keep the nominal rates constant at their initial position when performing fiscal shock simulations; see Figure 1, right column. Differences are small for our benchmark specification. So for the remainder of the paper, we continue to assume that fiscal and monetary policy are coordinated, i.e. the effect of a fiscal stimulus is not counteracted by monetary policy.

Next, we compare the hybrid QR-VAR with a standard VAR that abstracts from heterogeneous quantile effects. In the right column of Figure 2, the mean impulse response is lower when we allow for heterogeneous quantile effects (plain line) than the VAR (starred line). Indeed, the top quantiles respond more positively. A model that does not account for asymmetric effect is then likely to predict a mean response that is biased upward.

Note that in the United States, the ratio of central government current expenditures to GDP (excluding social-security-related spending and interest payments) is around 5 percent for the most recent periods.[13] Therefore, a 5 percent shock to central government spending is roughly equivalent to one-quarter of one percent of GDP. Such a shock would generate a peak increase in annualized GDP growth of 0.06 percentage points (with an upper bound of the confidence band of the mean response at 0.18 percentage points). This implies a mean fiscal multiplier of GDP of 0.25 (with the upper bound of the confidence band of the mean multiplier at 0.72). In a survey of 41 studies using both VAR and dynamic stochastic general equilibrium (DSGE), Mineshima et al. (2014) find that a plausible range for the fiscal multiplier is 0.5 to 0.9, with multipliers being possibly as low as 0. The mean fiscal multiplier from our QR-VAR thus belongs to the lower end of the multiplier estimates in the literature. With a standard VAR specification, the fiscal multiplier we obtain is 0.32. However, it is not too surprising that we find a lower mean fiscal multiplier, because we model upper and lower quantiles of GDP separately and our distribution of GDP growth may not be symmetric. Indeed, if we compute the fiscal multiplier for the 90th percentile of GDP growth, where the elasticity to fiscal shocks is largest, then the upper tail fiscal multiplier is now 0.5, with a confidence band of 0.12 to 0.88, much more in line with existing estimates.[14] With a fourth lag, the mean fiscal multiplier is 0.38 (QR-VAR) and 0.44 (VAR); see robustness in Appendix F.

Last, we conduct a series of robustness checks to account for the fact that government spending comes in different forms. Initially, we probe how important the type of government spending is in

[13] Data accessed from DLX Haver, in the category Major Economic and Financial Indicators.

[14] As pointed out by Ramey and Zubairy (2018), the precise estimation of fiscal multipliers hinges on details. Similar to our work, most of the literature estimates models in growth rates. But fiscal multipliers are obtained as the ratio between the change in the level of GDP and the change in the level of spending. Thus, models based on growth rates need to be scaled by some average ratio of GDP to government spending, potentially introducing a bias, especially when the ratio of government spending is heterogeneous across countries and steadily increases over time. With this in mind, we refrain from tackling the problem of computing fiscal multipliers.

our model. Figure 10 shows results for different definitions of government spending. For example, we can focus our attention on central government spending on its own, or we may also include local government spending. Moreover, we can focus on current expenses or on investment spending. If we consider current expenditures by either the central government (our benchmark; sub-figure a) or the broader government (federal and local; sub-figure b), we observe the same pattern: the 90th percentile (dashed line) responds more than the 10th percentile (dotted line). Furthermore, in both cases the peak impact for the 90th percentile is larger when looking at investment spending by the central government, although the responses of most other quantiles tend to be indistinguishable from zero. This corroborates the findings of Boehm (2020), according to which government investment multipliers are close to zero and government consumption multipliers are above zero.

Figure 1: Quantile impulse response to monetary and fiscal shocks

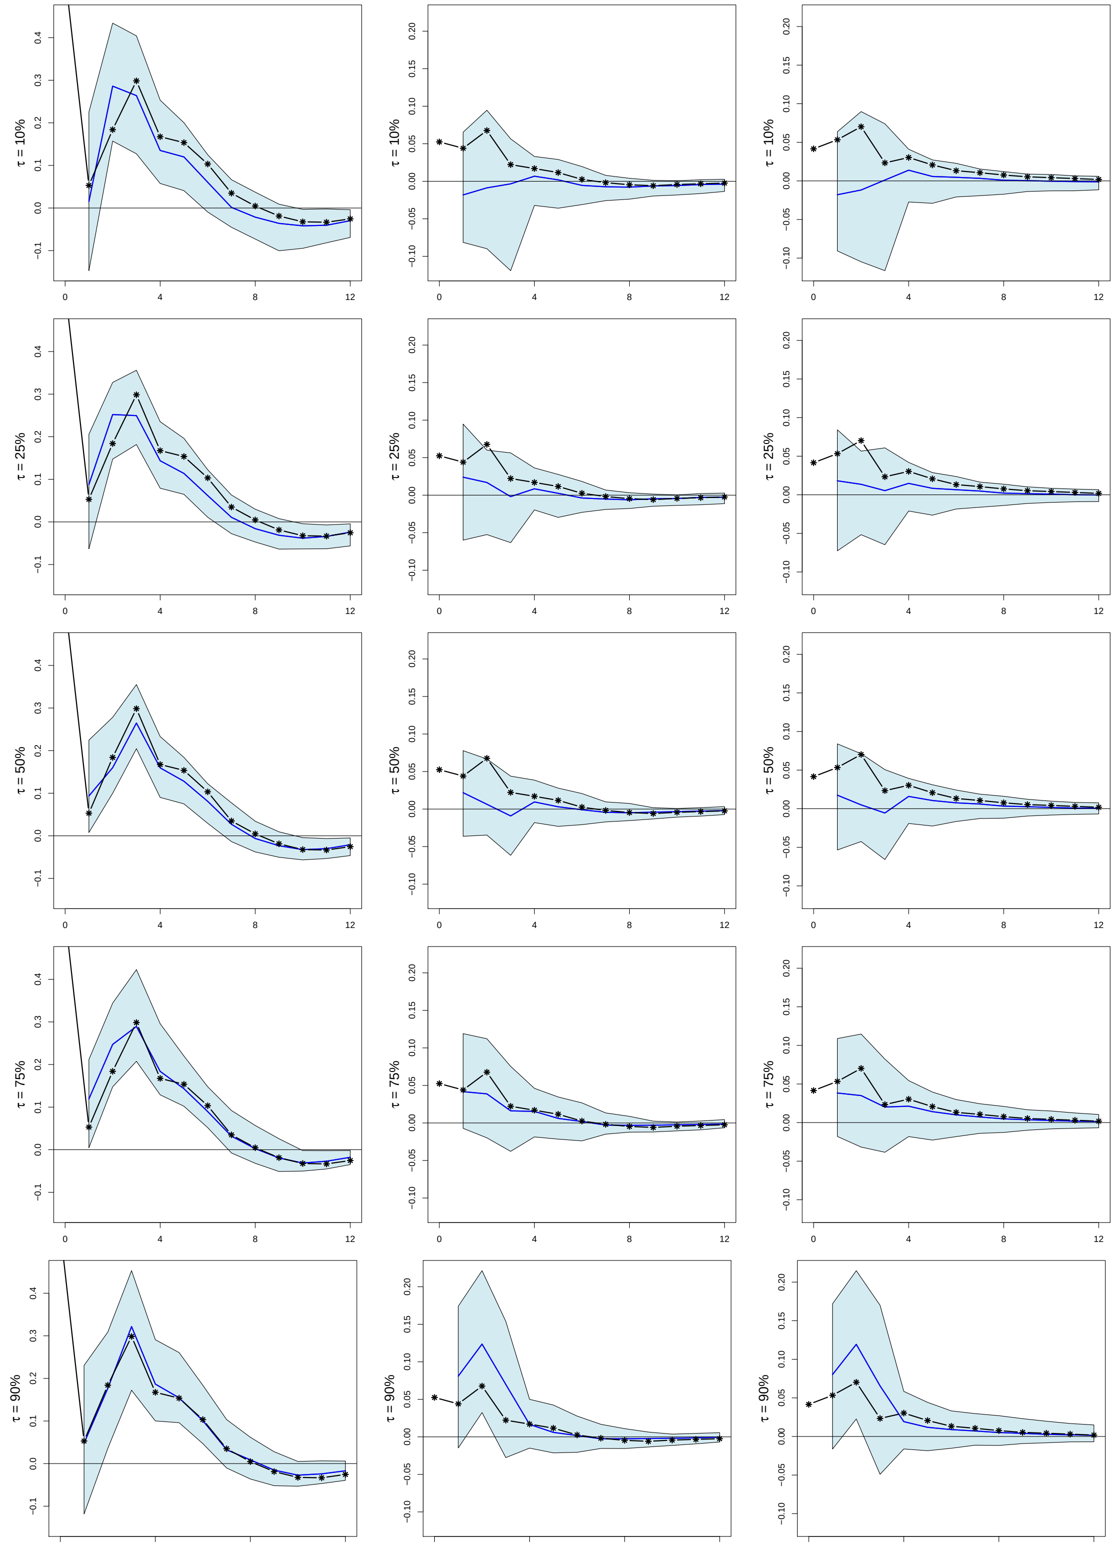


Notes: The response of different percentiles of the distribution of future GDP growth following, respectively, a 50-basis-point monetary policy easing (left), a 5 percent central government expenditure increase (middle), and a 5 percent central government expenditure increase while keeping the response of monetary policy fixed (right). The model is estimated with three lags; shaded areas are the 90 percent confidence intervals computed with block bootstrap. The starred black line is the mean response of the linear VAR.

Figure 2: Mean impulse response to monetary and fiscal shocks

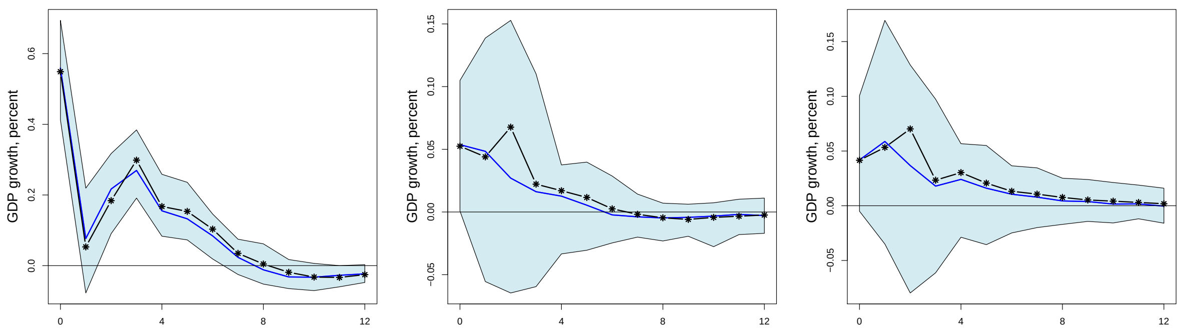


Notes: The response of the mean of the distribution of future GDP growth (plain line) following a 50-basis-point monetary policy easing (left), a 5 percent central government expenditure increase (middle), and a 5 percent central government expenditure increase while keeping the response of monetary policy fixed (right). The model is estimated with three lags; shaded areas represent the 90 percent confidence interval computed with block bootstrap. The starred black line is the mean impulse response from a comparable structural VAR that does not allow for heterogeneous effects of the response of GDP growth across quantiles. The difference between the plain line and the starred line is the mean bias introduced by shutting down the heterogeneous effect across quantiles.

Figure 3: Quantile response of GDP growth after different types of government spending shocks

(a) Expenditures, central (b) Expenditures, central and local (c) Investment, central

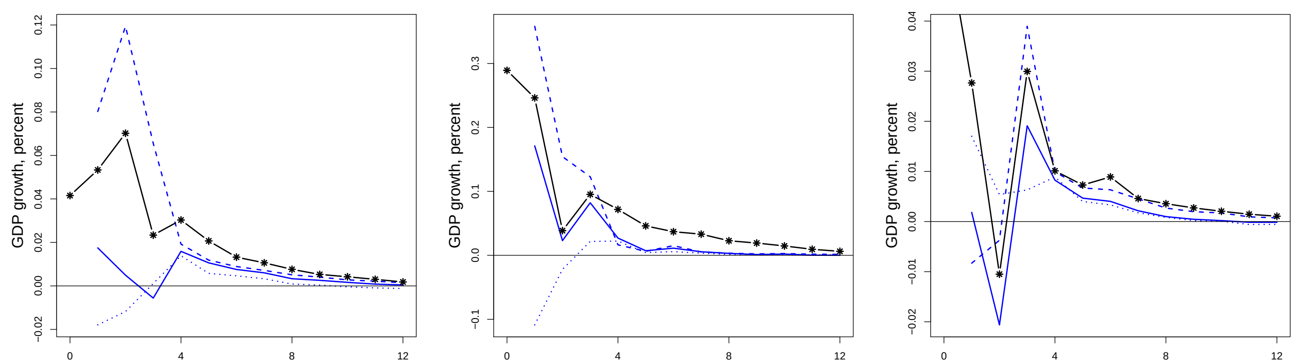


Notes: The response of the 90th (dashed lines), the 50th (plain line) and the 10th percentiles (dotted line) of future GDP growth (up to a 12-quarter horizon) to a 5 percent central government spending shock, a 5 percent central and local government spending shock and a 5 percent investment shock from the central government. Spending shocks are identified by zero restrictions assuming no contemporaneous response of government spending to other macro shocks. Model estimated with three lags. The starred black line is the mean response from of the linear VAR.

# 4 Policy shocks with the policy rate at the ZLB

We now focus our attention on the state dependence of the effects of fiscal and monetary policy tools on the distribution of GDP growth. In particular, we draw a contrast between the impacts of policy shocks when the ZLB is binding as opposed to normal times.

## 4.1 Model estimation setup

Because the ZLB is essentially a level concept, we now use a model with the level of the nominal policy rate $r$ instead of the differenced value, at the cost of not ensuring stationarity, and therefore not enforcing a stability condition for the VAR. We define the ZLB state $s_t(ZLB)$ as a period when the short-term rate is at or below 1 percentage point:[15]

$$s_t(ZLB) = \begin{cases} 1 & \text{if } r_t \leq 1\% \\ 0 & \text{otherwise.} \end{cases} \tag{11}$$

The rest of the model specification is similar to that in Equation (14).

Thinking more broadly about our modelling approach in the context of the ZLB episodes, we wish to emphasize a number of points. First, we assess the forecasting benefits of following a quantile approach compared with a simple OLS approach and report the RMSE of the one-step-ahead forecasts in Table 8 in Appendix F. Not surprisingly, the RMSE gains of the quantile regressions combined with a kernel estimation are largest for ZLB episodes, while OLS estimates generate a better forecasting performance in normal times and in-sample.

Next, we wish to clarify the need for a cross-country model and thus a larger set of data. Any single country experienced a unique ZLB episode only in the post-world-war period, with Japan experiencing the longest episode going back to the 1990s. The multi-country approach broadens our sample size of ZLB episodes, as well as their timing and circumstances.

Third, we use the shadow short-term rate. This allows us to capture the effect of quantitative easing, albeit in a simplified way. When the nominal rate hits the ZLB, conventional monetary policy becomes inactive and can be replaced by quantitative easing (QE), one of the unconventional monetary policy tools. This is where our cross-country strategy provides additional benefits, since our sample covers the main countries that implemented QE and for which we used the estimated shadow rate instead of the nominal short-term rate: the United States, Japan, the United Kingdom and Finland (as part of the EuroZone).[16]

[15] Given that the short-term rate is slightly above the policy rate, choosing a threshold of 1 percent is equivalent to a policy rate below 1 percent. This has an impact only for Canada around the oil crisis of 2015, where the policy rate was down to 0.5 percent, which was considered the effective lower bound at the time and the short-term rate fluctuated slightly above. Similar results are obtained with a cutoff of 0.5 percent or 0.25 percent, which would exclude only the Canadian episode.

[16] An alternative strategy to using the shadow rate could be to use a long-term rate that is unaffected by the ZLB, but would capture many QE benefits. We did not implement this approach, since the ZLB regime identification still relies on the short-term rate. Impulse responses are constructed from simulations, and we split those simulations between cases where the short-term rate remains above or falls below a ZLB threshold.

On a technical note, when we do the bootstrap, we ensure that the resampled dataset does not over- or under-represent ZLB events that could have an impact on the bootstrapped tail estimates of the GDP growth distribution and our ability to estimate the state-dependent parameters.

### 4.2 Monetary policy shock during ZLB episodes

As before, we consider the impact of stimulative monetary policy. However, it is important to note that the monetary shock in normal times corresponds to a conventional monetary policy shock, while the monetary shock during ZLB events corresponds to the implications from QE of an equivalent magnitude. The output growth responses for different quantiles are displayed in Figure 4, left column. The results in normal times are in plain blue and those for ZLB episodes in dashed red.

The response of the different quantiles of the GDP density forecast are similar between conventional and unconventional monetary policy, with only marginally lower point estimates for unconventional monetary policy. Note that the large confidence bands around the ZLB responses are pointing towards limited test power, given the relatively small number of ZLB observations, and do not necessarily indicate a weaker effect of QE. Our results suggest that both conventional and unconventional monetary policies shift the distribution of GDP growth homogeneously. However, the shape and scale of its distribution do not change.

Next, we focus on the mean impulse responses of GDP growth to monetary stimulus. In the left column of Figure 5, the mean impact of conventional and unconventional monetary policy cannot be distinguished across the two regimes, as we expected from the similar effects on the quantiles discussed above.

Our findings thus have two implications. First, unconventional monetary policy is effective at stimulating GDP even during ZLB episodes. Second, the approach of the standard New Keynesian DSGE literature focusing on the central moments of GDP is reasonable when conducting monetary policy experiments.

Figure 4: Quantile impulse response to monetary and fiscal shocks: normal times and ZLB episodes

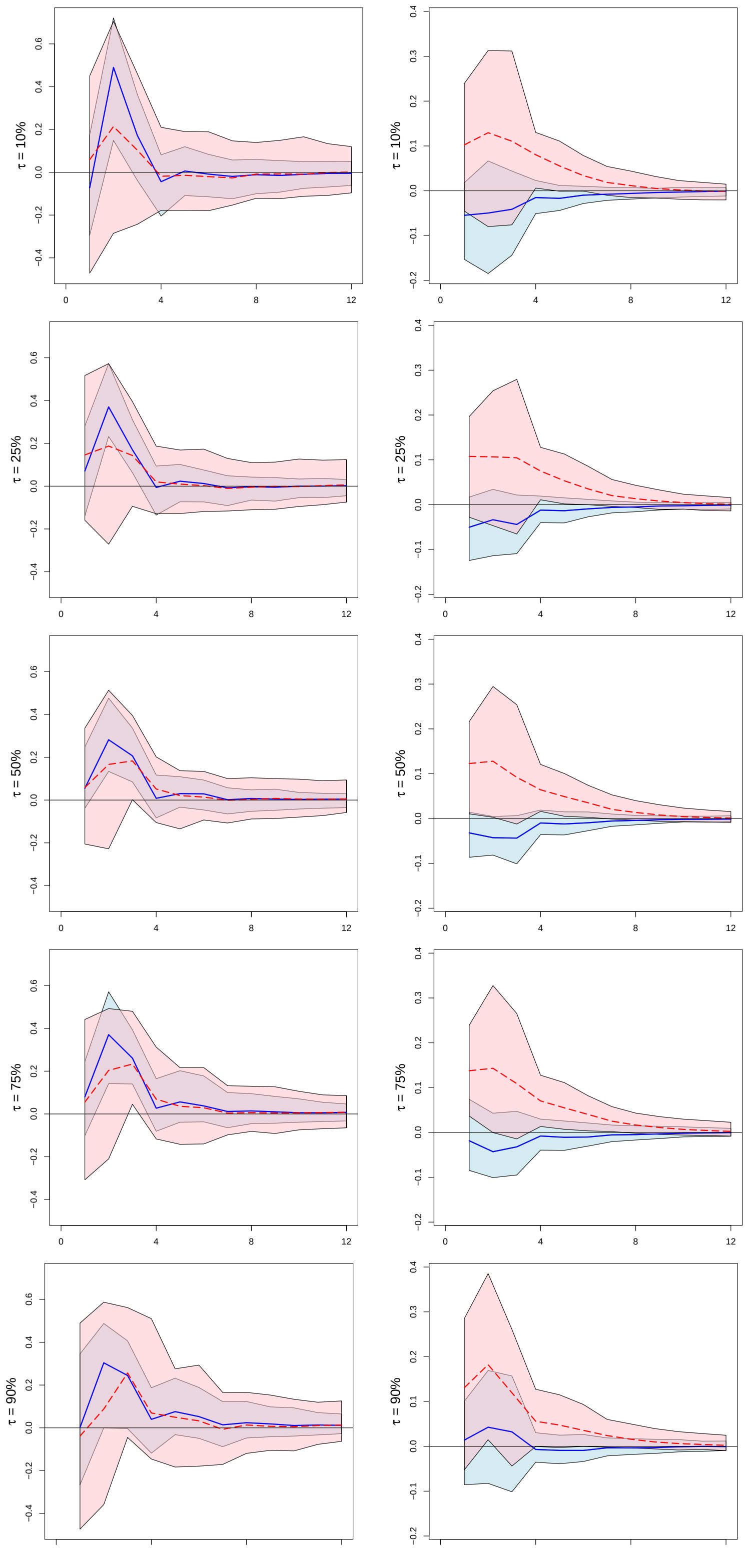


Notes: The figure shows the response of the 10th, 25th, 50th, 75th and 90th percentiles of the distribution of future GDP growth following a 50-basis-point monetary policy easing (left) and a 5 percent central government expenditure increase while keeping the response of monetary policy fixed (right). The plain blue lines correspond to the response in normal times, the dashed red lines the response during episodes of ZLB. The model is estimated with three lags using the shadow rate during ZLB episodes. The shaded areas represent the 90 percent confidence interval computed with block bootstrap.

Figure 5: Mean impulse response to monetary and fiscal shocks: normal times and ZLB episodes

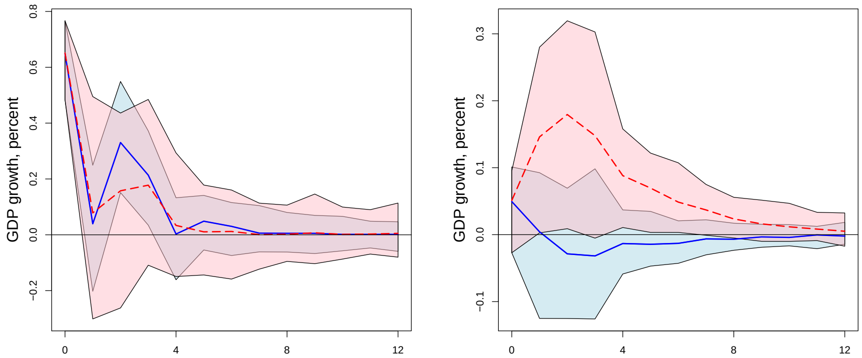


Notes: The figure shows the response of the mean of the distribution of future GDP growth in normal times (plain blue line) and during ZLB episodes (dashed red line) following a 50-basis-point (conventional or unconventional) monetary policy easing (left) and a 5 percent central government expenditure increase while keeping the response of monetary policy fixed (right). The model is estimated with three lags. The shaded areas represent the 90 percent confidence interval computed with block bootstrap.

### 4.3 Fiscal policy shocks during ZLB episodes

With the theoretical work by Christiano et al. (2011), we know that fiscal policy is more impactful during ZLB episodes. Now, we revisit the issue by asking whether it is purely a mean impact phenomenon or whether there is a distributional aspect to this insight. To do this, we look at the impulse responses of GDP growth for different quantiles following a central government expenditure increase; see Figure 4, right column. We differentiate between a shock during a ZLB episode (dashed red lines) and all other times (plain blue lines). Figure 4, right column, shows the mean response.

In the absence of a binding ZLB, the response to the government spending stimulus is negative for most quantiles and positive at the 90th quantile. However, none is significant. The point estimates are consistent with the view that fiscal policy can create a misallocation of resources and crowd out private consumption expenditures, thereby potentially creating downside risks to future GDP growth, while still increasing upside risks to future GDP. These differential responses in the tails are averaged out when looking at the mean response of GDP growth, such that fiscal policy does not effectively boost GDP growth in normal times.

When the policy rate is at the ZLB, the mean response to a fiscal shock is positive and significant. Since the ZLB is normally attained during severe economic downturns, this positive impulse arrives precisely when it is needed the most. This finding corroborates those of Ramey and Zubairy (2018) and Christiano et al. (2011), among others, who provide, respectively, empirical and theoretical evidence in favour of higher effectiveness of fiscal policy during ZLB episodes.

The quantile breakdown of the mean response reveals a significant positive response of the upper quantiles of the GDP growth outlook distribution. For the lower quantiles, the point estimate is now also positive, though this effect is initially not statistically significant. This suggests that an expansionary fiscal policy during ZLB episodes can successfully reduce the downside risk to GDP growth and increase the upside risk, thereby speeding up the recovery following the severe recession that triggered the ZLB episode.

## 5 Conclusion

To better understand the ability of fiscal and monetary policy to shape the future GDP growth distribution, we build a hybrid model that integrates the distributional insights from quantile regressions into a structural VAR framework that allows the identification of shocks.

Our hybrid model is a targeted approach to combine a structurally autoregressive vector with the explicit modelling of a nonparametric distribution via conditional quantiles. The method allows for only one variable to be modelled in a flexible distribution-free manner. Other variables that are less likely to exhibit non-normality are modelled using a standard linear model and estimated via OLS. In Monte Carlo simulations, we find that whenever the variable of interest is governed by a skewed distribution, choosing our QR-VAR over a VAR dramatically reduces the bias and RMSE for the one-step-ahead conditional skewness. This only comes at the cost of a small positive bias in the one-step-ahead conditional variance.

We are the first to analyze the quantile response of the probability density of future GDP growth to monetary and fiscal policy shocks where the shock identification is consistent with our quantile setup. We find that monetary policy shocks change the location of the distribution of GDP growth, while fiscal policy shocks also change the shape of the distribution. This impact of fiscal policy on the distribution of future GDP growth is more pronounced for central government spending and during ZLB episodes. Under these circumstances, fiscal stimulus generates upside risk, paving the path to a faster recovery.

Our results suggest that fiscal policy is well suited to hasten the recovery phase of the COVID-19 pandemic. As the economy gets past the initial shutdown, the recovery will increasingly rely on central government spending to speed up the recovery, while policy rates will likely remain low for a long time.

## References

Adrian, T., N. Boyarchenko, and D. Giannone (2019): "Vulnerable Growth," *American Economic Review*, 109, 1263–89. 1, 5, 3.1, 17, A

Adrian, T., F. Grinberg, N. Liang, and S. Malik (2021): "The Term Structure of Growth-at-Risk," *American Economic Journal: Macroeconomics*. A

Auerbach, A. J. and Y. Gorodnichenko (2012a): "Measuring the Output Responses to Fiscal Policy," *American Economic Journal: Economic Policy*, 4, 1–27. 1

Auerbach, A. J. and Y. Gorodnichenko (2012b): "Measuring the output responses to fiscal policy," *American Economic Journal: Economic Policy*, 4, 1–27. G.3

Balke, N. (2000): "Credit and economic activity: credit regimes and nonlinear propagation of shocks," *Review of Economics and Statistics*, 82, 344–349. 1

Blanchard, O. and R. Perotti (2002): "An Empirical Characterization of the Dynamic Effects of Changes in Government Spending and Taxes on Output," *The Quarterly Journal of Economics*, 117, 1329–1368. 1, 3.1.1

Boehm, C. (2020): "Government Consumption and Investment: Does the Composition of Purchases Affect the Multiplier?" *Journal of Monetary Economics*, 115, 80–93. 3.3

Bondell, H. D., B. J. Reich, and H. Wang (2010): "Noncrossing quantile regression curve estimation," *Biometrika*, 97, 825–838. 4

Canay, I. A. (2011): "A simple approach to quantile regression for panel data," *Econometrics Journal*, 14, 368–386. 10

Canova, F. and E. Pappa (2011): "Fiscal policy, pricing frictions and monetary accommodation," *Economic Policy*, 26, 555–598. 1

Carriero, A., T. E. Clark, and M. Marcellino (2020): "Capturing Macroeconomic Tail Risks with Bayesian Vector Autoregressions," Working Papers 202002R, Federal Reserve Bank of Cleveland. 1

Chavleishvili, S. and S. Manganelli (2019): "Forecasting and Stress Testing with Quantile Vector Autoregression," Tech. rep., ECB mimeo. 1

Chernozhukov, V., I. Fernández-Val, and A. Galichon (2010): "Quantile and probability curves without crossing," *Econometrica*, 78, 1093–1125. 2.2.2

Chiu, C.-W. J., H. Mumtaz, and G. Pinter (2017): "Forecasting with VAR models: Fat tails and stochastic volatility," *International Journal of Forecasting*, 33, 1124–1143. 1

Christiano, L., M. Eichenbaum, and S. Rebelo (2011): "When Is the Government Spending Multiplier Large?" *Journal of Political Economy*, 119, 78–121. 1, 4.3

Cogley, T., S. Morozov, and T. J. Sargent (2005): "Bayesian fan charts for U.K. inflation: Forecasting and sources of uncertainty in an evolving monetary system," *Journal of Economic Dynamics and Control*, 29, 1893–1925. 1

Davino, C., M. Furno, and D. Vistocco (2013): *Quantile regression: theory and applications*, vol. 988, John Wiley & Sons. 2.2.2, 2.2.2

Duprey, T., B. Klaus, and T. Peltonen (2017): “Dating systemic financial stress episodes in the EU countries,” *Journal of Financial Stability*, 32, 30 – 56. 3.1, 6

Duprey, T. and A. Ueberfeldt (2020): “Managing GDP Tail Risk,” Staff working paper, 2020-03, Bank of Canada. 1, A

Engle, R. F. and S. Manganelli (2004): “CAViaR: Conditional autoregressive value at risk by regression quantiles,” *Journal of Business & Economic Statistics*, 22, 367–381. 1, C

Gaglianone, W. P. and L. R. Lima (2012): “Constructing density forecasts from quantile regressions,” *Journal of Money, Credit and Banking*, 44, 1589–1607. 2.2.2, 2.2.2

Gallant, A. R., P. E. Rossi, and G. Tauchen (1993): “Nonlinear dynamic structures,” *Econometrica: Journal of the Econometric Society*, 871–907. 7

Galí, J., J. D. López-Salido, and J. Vallés (2007): “Understanding the Effects of Government Spending on Consumption,” *Journal of the European Economic Association*, 5, 227–270. 1

Ghysels, E., L. Iania, and J. Striaukas (2018): “Quantile-based Inflation Risk Models,” Tech. rep. 5

Han, H., W. Jung, and J. H. Lee (2019): “Estimation and Inference of Quantile Impulse Response Functions by Local Projections: With Applications to VaR Dynamics,” Tech. rep., working Paper. 1

He, X. (1997): “Quantile curves without crossing,” *The American Statistician*, 51, 186–192. 4

Hubrich, K. and R. J. Tetlow (2015): “Financial stress and economic dynamics: The transmission of crises,” *Journal of Monetary Economics*, 70, 100–115. 1

Ilzetzki, E., E. G. Mendoza, and C. A. Végh (2013): “How big (small?) are fiscal multipliers?” *Journal of Monetary Economics*, 60, 239–254. 1

International Monetary Fund (2017): “Financial Conditions and Growth at Risk,” Chapter 3 of the global financial stability report: Is growth at risk? 1

Jordà, Ò. (2005): “Estimation and inference of impulse responses by local projections,” *American economic review*, 95, 161–182. 1

Kato, K., A. F. Galvao, and G. V. Montes-Rojas (2012): “Asymptotics for panel quantile regression models with individual effects,” *Journal of Econometrics*, 170, 76–91. 3.1

Khmaladze, E. V. (1982): “Martingale approach in the theory of goodness-of-fit tests,” *Theory of Probability & Its Applications*, 26, 240–257. C

Koenker, R. (2004): “Quantile regression for longitudinal data,” *Journal of Multivariate Analysis*, 91, 74–89. 10

Koenker, R. (2005): *Quantile Regression*, Econometric Society Monographs, Cambridge University Press. 5

Koenker, R. and G. Bassett Jr (1978): "Regression quantiles," *Econometrica: Journal of the Econometric Society*, 33–50. 2.2.1, 10

Koenker, R. and G. Bassett Jr (1982): "Robust tests for heteroscedasticity based on regression quantiles," *Econometrica: Journal of the Econometric Society*, 43–61. C

Koenker, R. and S. Portnoy (1987): "L-estimation for linear models," *Journal of the American statistical Association*, 82, 851–857. 2.2.2

Koop, G., M. H. Pesaran, and S. M. Potter (1996): "Impulse response analysis in nonlinear multivariate models," *Journal of econometrics*, 74, 119–147. 7

Korobilis, D. (2017): "Quantile regression forecasts of inflation under model uncertainty," *International Journal of Forecasting*, 33, 11–20. 2.2.2, 2.2.2

Krippner, L. (2013): "Measuring the stance of monetary policy in zero lower bound environments," *Economics Letters*, 118, 135–138. 6

Linnemann, L. and R. Winkler (2016): "Estimating nonlinear effects of fiscal policy using quantile regression methods," *Oxford Economic Papers*, 68, 1120–1145. 1

Loria, F., C. Matthes, and D. Zhang (2019): "Assessing Macroeconomic Tail Risk," Tech. rep., finance and Economics Discussion Series 2019-026, Washington: Board of Governors of the Federal Reserve System. 1

Machado, J. A. and J. Santos Silva (2019): "Quantiles via moments," *Journal of Econometrics*, 213, 145–173. 10

Mineshima, A., M. Poplawski-Ribeiro, and A. Weber (2014): "Fiscal Multipliers," in *Post-crisis Fiscal Policy*, ed. by C. Cottarelli, P. Gerson, and A. Senhadji. 3.3

Montes-Rojas, G. (2019): "Multivariate Quantile Impulse Response Functions," *Journal of Time Series Analysis*. 1

Mountford, A. and H. Uhlig (2009): "What are the effects of fiscal policy shocks?" *Journal of Applied Econometrics*, 24, 960–992. 12

Perotti, R. (2005): "Estimating the effects of fiscal policy in OECD countries," *Proceedings*. 1

Potter, S. M. (2000): "Nonlinear impulse response functions," *Journal of Economic Dynamics and Control*, 24, 1425–1446. 7

Primiceri, G. E. (2005): "Time Varying Structural Vector Autoregressions and Monetary Policy," *Review of Economic Studies*, 72, 821–852. 1

Ramey, V. A. (2011): "Identifying Government Spending Shocks: It's all in the Timing," *The Quarterly Journal of Economics*, 126, 1–50. 1

Ramey, V. A. and S. Zubairy (2018): "Government spending multipliers in good times and in bad: evidence from US historical data," *Journal of Political Economy*, 126, 850–901. 1, 1, 12, 14, 4.3

Ranciere, R., A. Tornell, and F. Westermann (2008): "Systemic Crises and Growth*," *The Quarterly Journal of Economics*, 123, 359–406. 1

Romer, C. D. and D. H. Romer (2010): "The Macroeconomic Effects of Tax Changes: Estimates Based on a New Measure of Fiscal Shocks," *American Economic Review*, 100, 763–801. 11

Schmidt, L. and Y. Zhu (2016): "Quantile spacings: A simple method for the joint estimation of multiple quantiles without crossing," *Available at SSRN 2220901*. 4, 5

Schüler, Y. S. (2014): "Asymmetric effects of uncertainty over the business cycle: A quantile structural vector autoregressive approach," . 1

Silverman, B. W. (2018): *Density estimation for statistics and data analysis*, Routledge. 2.2.2

Sims, E. and C. Wu (2020): "Are QE and Conventional Monetary Policy Substitutable?" *International Journal of Central Banking*, 16, 195–230. 1

Smith, M. S. and S. P. Vahey (2016): "Asymmetric Forecast Densities for U.S. Macroeconomic Variables from a Gaussian Copula Model of Cross-Sectional and Serial Dependence," *Journal of Business & Economic Statistics*, 34, 416–434. 1

Uhlig, H. (2005): "What are the effects of monetary policy on output? Results from an agnostic identification procedure," *Journal of Monetary Economics*, 52, 381–419. 3.1.1

White, H., T.-H. Kim, and S. Manganelli (2015): "VAR for VaR: Measuring tail dependence using multivariate regression quantiles," *Journal of Econometrics*, 187, 169–188. 1

Wu, J. C. and F. D. Xia (2016): "Measuring the Macroeconomic Impact of Monetary Policy at the Zero Lower Bound," *Journal of Money, Credit and Banking*, 48, 253–291. 1

# A Monte Carlo study

In this appendix, we present the results from our comparative study of the traditional VAR and the QR-VAR hybrid models in the context of three autoregressive processes that can be represented as follows:

$$\begin{aligned} X_t &= \theta X_{t-1} + \varepsilon_t \\ \varepsilon_t &= \sigma_t Z_t \\ \sigma_t^2 &= \omega + \alpha(\varepsilon_{t-1})^2, \end{aligned} \tag{12}$$

where $X_t = [X_{1,t}, X_{2,t}]'$, $\varepsilon_t = [\varepsilon_{1,t}, \varepsilon_{2,t}]'$, $\sigma_t = I_{2\times 2}[\sigma_{1,t}, \sigma_{2,t}]'$, and $\theta$ and $\alpha$ are 2-by-2 coefficient matrices consistent with stationarity restrictions. The constant $\omega$ is fixed at $\omega = \begin{bmatrix} 0.5 & 0 \\ 0 & 0.5 \end{bmatrix}$. In all cases, we keep the autoregressive structure identical, i.e $\theta = \begin{bmatrix} 0.5 & 0 \\ 0.5 & 0.5 \end{bmatrix}$.

1. AR(1) with Gaussian innovations (Bivariate AR(1)):
   $\alpha = 0$ and $\varepsilon_t \overset{i.i.d}{\sim} N(0,1)$
2. ARCH(1,1) with Gaussian innovations (Bivariate ARCH(1,1)):
   $\alpha = \begin{bmatrix} 0.25 & 0.25 \\ 0.25 & 0.25 \end{bmatrix}$ and $\varepsilon_t \overset{i.i.d}{\sim} N(0,1)$
3. ARCH(1,1) with skewed Student-distributed innovations (Bivariate Skewed ARCH(1,1)):
   $\alpha = \begin{bmatrix} 0.25 & 0.25 \\ 0.25 & 0.25 \end{bmatrix}$ ; $\varepsilon_{1,t} \sim ST(\lambda{=}0.25)$[17] and $\varepsilon_{2,t} \sim N(0,1)$

Table 2 produces the RMSE of the one-step-ahead forecasts of the conditional mean, variance and skewness of variable $X_1$ for both the VAR and the QR-VAR hybrid. We use 1,000 Monte Carlo simulations of a sample with 200, 500 or 1,000 observations. In the absence of skewness, there is no significant difference in the RMSE of the conditional mean forecast between the QR-VAR and the VAR. There is a small increase in the RMSE of the conditional variance for the QR-VAR compared to the VAR, mostly in the absence of heteroskedasticity. This is not surprising, since the VAR model is perfectly specified for this case. This is partly because the VAR correctly assumes zero skewness. In the presence of skewness, there is no significant difference in the RMSE of the conditional mean or variance, but a very large reduction in the RMSE of the conditional skewness for the QR-VAR compared to the VAR, by a factor of about three.

Table 3 produces the associated bias of the one-step-ahead forecasts of the conditional mean, variance and skewness of variable $X_1$. In the absence of skewness, there is no bias of the conditional mean and a small positive bias of the conditional variance when choosing the QR-VAR over the VAR. Interestingly, using a QR-VAR in the absence of skewness does not create a significant bias in the estimated conditional skewness. In the presence of skewness, choosing a QR-VAR over a VAR dramatically reduces the bias of the conditional skewness by a factor of ten, without a significant bias of the conditional mean and only a marginal positive bias of the conditional variance.

Figures 6 and 7 focus on the Bivariate Skewed ARCH(1,1) and evaluate the robustness of the QR-VAR for varying degrees of skewness of the data-generating process. Figure 6 displays mean RMSEs of the conditional mean, variance and skewness one-step-ahead forecasts across a range of

[17] $\lambda$ is the slant parameter. See the R package `sgt` for more details. The use of the skewed-t is motivated by the work of Adrian et al. (2019) that highlight its relevance for modelling GDP growth.

slant parameter $\lambda \in (-1, 1)$ for 200 replications. The QR-VAR is not statistically different from the VAR with respect to the forecast errors of the mean (panel a) and variance (panel b) for the most common range of non-zero skewness. This is reassuring, as the VAR model corresponds to the best linear unbiased estimator of the conditional mean. However, the QR-VAR significantly outperforms the VAR when it comes to detecting asymmetries (panel c). If there is no skewness in the data-generating process with $\lambda$ in the neighborhood of zero, then the VAR that assumes a skewness of zero is unsurprisingly better since it is perfectly specified for this case.

Figure 7 displays the associated bias of the conditional mean, variance and skewness one-step-ahead forecasts across a range of slant parameter $\lambda \in (-1, 1)$. For the most common range of non-zero skewness, there is a tradeoff when using the QR-VAR over the VAR. The small increase in the bias for the conditional variance when using the QR-VAR rather than the VAR (panel b) is largely compensated by the large decrease in the bias for the conditional skewness (panel c). For the most extreme range of skewness, less common for economic variables, the tradeoff is between a small increase in the bias on the conditional mean (panel a) against a large reduction in the bias of the conditional skewness (panel c).

Table 2: RMSE of conditional mean, variance and skewness one-step-ahead forecasts

| | Bivariate AR(1) | | | Bivariate ARCH(1,1) | | | Bivariate Skewed ARCH(1,1) | | |
|---|---|---|---|---|---|---|---|---|---|
| Sample size | 200 | 500 | 1000 | 200 | 500 | 1000 | 200 | 500 | 1000 |
| *RMSE of the conditional mean* | | | | | | | | | |
| QR-VAR | .079 | .051 | .036 | .122 | .077 | .056 | .120 | .077 | .057 |
| | (.033) | (.022) | (.015) | (.052) | (.032) | (.022) | (.051) | (.031) | (.023) |
| VAR | .078 | .051 | .036 | .124 | .078 | .057 | .122 | .077 | .055 |
| | (.032) | (.022) | (.015) | (.053) | (.033) | (.023) | (.053) | (.032) | (.024) |
| *RMSE of the conditional variance* | | | | | | | | | |
| QR-VAR | .093 | .064 | .051 | .519 | .512 | .509 | .524 | .513 | .511 |
| | (.043) | (.028) | (.022) | (.059) | (.035) | (.026) | (.062) | (.038) | (.027) |
| VAR | .040 | .025 | .018 | .506 | .503 | .502 | .514 | .509 | .508 |
| | (.030) | (.019) | (.013) | (.050) | (.031) | (.024) | (.056) | (.035) | (.025) |
| *RMSE of the conditional skewness* | | | | | | | | | |
| QR-VAR | .289 | .190 | .134 | .304 | .205 | .147 | .277 | .189 | .141 |
| | (.118) | (.078) | (.056) | (.129) | (.084) | (.061) | (.119) | (.075) | (.057) |
| VAR | .000 | .000 | .000 | .000 | .000 | .000 | .374 | .373 | .383 |
| | (.000) | (.000) | (.000) | (.000) | (.000) | (.000) | (.169) | (.109) | (.076) |

Notes: The table shows the mean RMSE for the fitted conditional mean, variance and skewness on variable $X_1$ governed by the data-generating process (12) and modelled by the QR-VAR and the VAR. Standard error in parenthesis. 1,000 simulations of a sample with 200, 500 or 1,000 observations.

Table 3: Bias of conditional mean, variance and skewness one-step-ahead forecasts

| | Bivariate AR(1) | | | Bivariate ARCH(1,1) | | | Bivariate Skewed ARCH(1,1) | | |
|---|---|---|---|---|---|---|---|---|---|
| Sample size | 200 | 500 | 1000 | 200 | 500 | 1000 | 200 | 500 | 1000 |
| *Bias of the conditional mean* | | | | | | | | | |
| QR-VAR | .000 | .001 | .000 | .001 | -.004 | .000 | -.008 | -.009 | -.009 |
| | (.048) | (.033) | (.023) | (.067) | (.043) | (.031) | (.067) | (.044) | (.031) |
| VAR | .000 | .001 | .000 | .001 | -.004 | .000 | .000 | -.002 | -.002 |
| | (.048) | (.033) | (.023) | (.069) | (.044) | (.031) | (.069) | (.045) | (.032) |
| *Bias of the conditional variance* | | | | | | | | | |
| QR-VAR | .038 | .035 | .035 | .033 | .032 | .031 | .027 | .026 | .028 |
| | (.054) | (.035) | (.024) | (.104) | (.068) | (.048) | (.103) | (.066) | (.048) |
| VAR | -.004 | -.002 | -.001 | -.012 | -.004 | -.003 | -.014 | -.006 | -.002 |
| | (.049) | (.032) | (.022) | (.108) | (.073) | (.051) | (.114) | (.073) | (.053) |
| *Bias of the conditional skewness* | | | | | | | | | |
| QR-VAR | -.006 | .001 | .000 | -.007 | -.002 | -.002 | -.036 | -.038 | -.035 |
| | (.182) | (.122) | (.083) | (.206) | (.132) | (.096) | (.152) | (.093) | (.071) |
| VAR | .000 | .000 | .000 | .000 | .000 | .000 | -.373 | -.373 | -.383 |
| | (.000) | (.000) | (.000) | (.000) | (.000) | (.000) | (.171) | (.109) | (.076) |

Notes: The table shows the mean bias for the fitted conditional mean, variance and skewness on variable $X_1$ governed by the data-generating process (12) and modelled by the QR-VAR and the VAR. Standard error in parenthesis. 1,000 simulations of a sample with 200, 500 or 1,000 observations.

Figure 6: RMSE on the conditional mean, variance and skewness one-step-ahead forecasts for different slant parameter of the skew-t distribution

(a) Conditional mean (b) Conditional variance (c) Conditional skewness

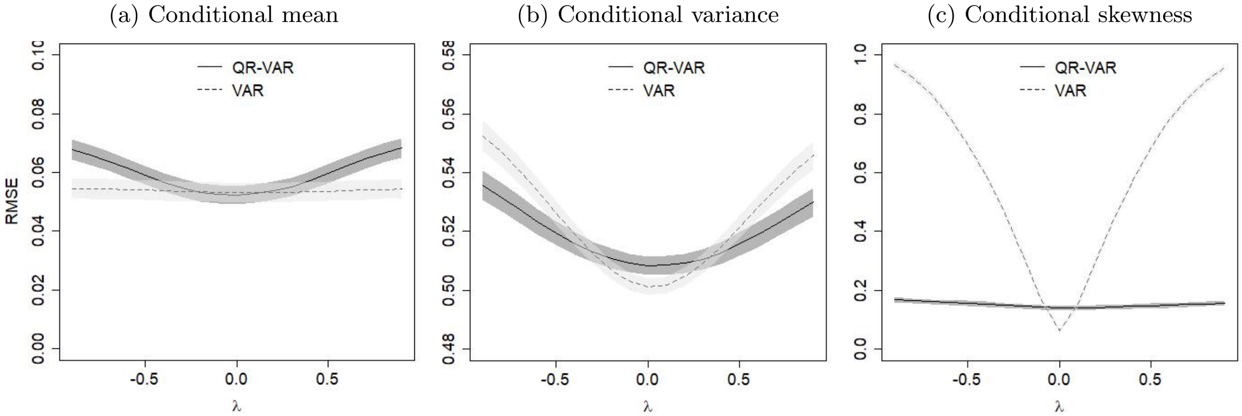


Notes: Mean RMSE on the conditional mean, variance and skewness forecast for different slant parameters in the bivariate skewed ARCH(1,1) data-generating process with 95% confidence intervals. Computed with 200 replications for each slant parameter $\lambda$ of the skew-t distribution, ranging from -.9 to .9 with a step of .1.

Figure 7: Bias on the conditional mean, variance and skewness one-step-ahead forecasts for different slant parameter of the skew-t distribution

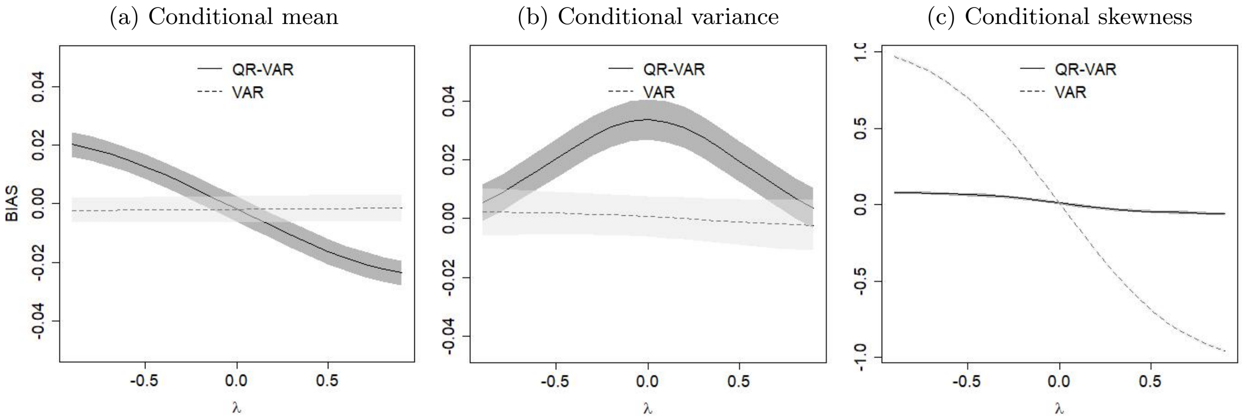


Notes: Mean bias on the conditional mean, variance and skewness forecast for different slant parameters in the bivariate skewed ARCH(1,1) data-generating process with 95% confidence intervals. Computed with 200 replications for each slant parameter $\lambda$ of the skew-t distribution, ranging from -.9 to .9 with a step of .1.

Our key takeaway is this: whenever the variable of interest is governed by a skewed distribution, then the QR-VAR generates a large reduction in both the RMSE and bias of the conditional skewness. This comes only at the cost of a small increase in the bias of the conditional variance. The QR-VAR is most valuable when the variable of interest modelled by the quantile regressions is skewed, like for instance GDP growth (Adrian et al., 2019; Duprey and Ueberfeldt, 2020; Adrian et al., 2021).

# B Lag order and model fit

We use three lags $L = 3$. This choice is informed by the out-of-sample mean squared error of the model; see Table 4 in the Appendix. In-sample, the RMSE is smallest for four lags when computed by OLS. If we instead assess the ability of the nonlinear models to fit the Great Recession out-of-sample, the RMSE is lowest when using the kernel smoother of the quantile estimates with two lags. This is in line with the idea that OLS performs well in normal times but nonlinear models that rely on quantile regressions are better at capturing tail events. We also perform cross-validation by removing one country at a time from the estimation sample and computing the RMSE for the country not used for the estimation. For the test datasets, the RMSE is smallest for three lags using a skew-t distribution or the kernel estimation.

When implementing block bootstrap, we pick $S = 8$ quarters in our estimation to balance the accuracy of the confidence bands of the tail quantiles and the overall randomization.[18]

[18]The choice of $S$ is not crucial for our empirical results. But a lower $S$ makes it more likely that some tail quantiles end up outside of their bootstrapped confidence bands, as the tail effect is weaker with smaller $S$.

Table 4: RMSE of one-step-ahead forecasts of GDP growth across model specifications

| Lag order | In-sample until 2007 | | | | Out-of-sample from 2007 onwards | | | |
|---|---|---|---|---|---|---|---|---|
| | **OLS** | Skew-t | P50 | Kernel | OLS | Skew-t | P50 | **Kernel** |
| Lag 1 | 4.006 | 4.023 | 4.068 | 4.060 | 3.710 | 3.629 | 3.555 | 3.454 |
| Lag 2 | 3.878 | 3.892 | 3.913 | 3.926 | 3.633 | 3.567 | 3.496 | **3.390** |
| Lag 3 | 3.766 | 3.778 | 3.823 | 3.831 | 3.622 | 3.558 | 3.471 | 3.408 |
| Lag 4 | **3.702** | 3.721 | 3.767 | 3.776 | 3.801 | 3.719 | 3.666 | 3.582 |
| | In-sample until 2010 | | | | Out-of-sample from 2010 onwards | | | |
| | **OLS** | Skew-t | P50 | Kernel | OLS | Skew-t | P50 | **Kernel** |
| Lag 1 | 3.975 | 3.993 | 4.020 | 4.019 | 3.635 | 3.600 | 3.495 | 3.418 |
| Lag 2 | 3.906 | 3.922 | 3.947 | 3.951 | 3.271 | 3.279 | 3.186 | 3.055 |
| Lag 3 | 3.929 | 3.937 | 3.972 | 3.979 | 2.568 | 2.523 | 2.512 | **2.425** |
| Lag 4 | **3.875** | 3.880 | 3.929 | 3.931 | 2.817 | 2.765 | 2.779 | 2.644 |
| | Cross-validation: training | | | | Cross-validation: test | | | |
| | **OLS** | Skew-t | P50 | Kernel | OLS | **Skew-t** | P50 | Kernel |
| Lag 1 | 3.805 | 3.816 | 3.835 | 3.822 | 4.214 | 4.279 | 4.184 | 4.247 |
| Lag 2 | 3.696 | 3.711 | 3.738 | 3.725 | 4.058 | 4.049 | 4.069 | 4.076 |
| Lag 3 | 3.623 | 3.631 | 3.666 | 3.661 | 4.043 | **4.009** | 4.045 | 4.025 |
| Lag 4 | **3.611** | 3.621 | 3.660 | 3.657 | 4.080 | 4.036 | 4.151 | 4.073 |

Notes: The table shows the root mean squared error (RMSE) for one-step-ahead GDP growth modelled as the mean fitted by OLS, the mean recovered from a skew-t distribution on a limited set of fitted quantiles, the median fitted by a quantile or the mean recovered from a kernel smoothing of a large number of fitted quantiles, with one to four lags. For the first two parts of the table, the RMSE is computed in-sample until 2007 or 2010 and on the corresponding out-of-sample. For the third part of the table, the RMSE is computed with cross-validation: the training dataset consists of all but one country, and the test dataset is the country not used for the estimation. We report the RMSE averaged over all possible folds of the dataset. In bold are the models that provide the smallest RMSE. For comparison, the mean of the annualized real GDP growth per capita is 1.89 and the unconditional standard deviation is 3.99.

# C Slope heterogeneity and location-scale effects

QRs are very useful in cases where the distribution displays heterogeneous effects across quantiles, but conventional methods will easily outperform QRs when that is not the case. Specifically, the Wald test statistic described in Koenker and Bassett Jr (1982) tests the null hypothesis that the slope coefficients of a collection of QR models are identical. We consider five QRs for GDP growth of probability orders $\mathcal{T} = \{0.10, 0.25, 0.50, 0.75, 0.90\}$ to conduct these tests. First, we test the null hypothesis that all slope coefficients are equal across all $\mathcal{T}$, which is strongly rejected.

Second, we perform pairwise tests in order to reveal whether a specific region of the distribution shows co-movements across quantiles. One can imagine a case where the centre of the distribution is changing homogeneously, but not the tails. In such instances, a pairwise test between two central quantiles would fail to reject the null hypothesis, whereas a pairwise test between a central quantile and a tail quantile would reject it. Test results in Table 5 provide strong evidence that no two quantiles are responding in same way.

Table 5: Pairwise Wald test p-values for slope homogeneity with Bonferroni correction

| $\tau$ | 0.10 | 0.25 | 0.50 | 0.75 | 0.90 |
|---|---|---|---|---|---|
| 0.10 | 1 | 0.00 | 0.00 | 0.00 | 0.00 |
| 0.25 | | 1 | 0.01 | 0.00 | 0.00 |
| 0.50 | | | 1 | 0.00 | 0.00 |
| 0.75 | | | | 1 | 0.00 |
| 0.90 | | | | | 1 |

Next, we use a Kolmogorov-type goodness-of-fit test proposed by Khmaladze (1982) to test location and location-scale model specifications. A location model is a statistical model where the distribution is solely characterized by a location parameter, whereas a location-scale model would also include a scale (variance) parameter. For instance, simple linear regression models assume a conditional distribution from the location family, and GARCH models from the location-scale family. The CAViaR model of Engle and Manganelli (2004) would then be well specified in the case of a location-scale family, for instance. Consistent with our findings in Table 5, the location family specification is rejected at the 1 percent level. Evidence against the location-scale family specification is weaker, but still significant at the 5 percent level.[19] These results support the use of a density estimation methodology that is not restricted to either specification.

[19] Wald and Khmaladze statistical hypothesis tests are computed in `R` with the `quantreg` package. Goodness-of-fit test significance levels are computed with the function `KhmaladzeFormat` from the `Qtools` package.

# D Panel data

Table 6: Dataset from 1964Q1 to 2019Q4

| **Data** | **Source** | **Description** |
|---|---|---|
| Real gross domestic product | OECD | Local currency |
| Consumer price index | OECD | |
| Real central government total outlays | Haver | Current Consumption, Transfer Payments, Subsidies, Capital Outlays. Local currency. Nominal data deflated using the GDP deflator from the OECD. Starts in 1970Q2 for Japan and 1973Q3 for Australia |
| Real general government final consumption expenditure | OECD | All levels of government (central, provincial, local). Expenditures for collective consumption (defence, justice, etc.) or for individual consumption (health care, housing, education, etc.). Does not include expenditures on assets nor transfer payments. Local currency |
| Real general government fixed capital formation | OECD; Statistics Canada; Cabinet Office of Japan | All levels of government (central, provincial, local). Expenditures on assets: buildings, machinery and equipment, roads, bridges or expenditures on research and development. Local currency. Starts in 1987Q1 for the United Kingdom, 1980Q1 for Japan and 1990Q1 for Finland |
| Shadow rate | Krippner (2013) | For the United States, Finland, Japan and the United Kingdom |
| Short-term interest rate | OECD | Starts in 1968Q1 for Australia |
| Country level index of financial stress | Duprey et al. (2017) | Adjusted for cross-country comparability; index covers the foreign exchange, the equity and the government bonds market, and starting in the 1970s/80s the banking and housing sectors |
| Population | Haver | |

Table 7: Summary statistics for the variables used in the benchmark model

| | Mean | Variance | Excess skewness | Excess kurtosis |
|---|---|---|---|---|
| Annualized growth of GDP per capita $\Delta y$ | | | | |
| Japan | 2.01 | 16.69 | -0.87 | 4.26 |
| United States | 1.89 | 9.37 | -0.27 | 2.68 |
| United Kingdom | 1.86 | 13.06 | 0.60 | 6.57 |
| Canada | 1.60 | 10.01 | -0.36 | 1.13 |
| Australia | 1.66 | 11.26 | 0.40 | 2.53 |
| Finland | 2.21 | 33.61 | -0.26 | 4.33 |
| Annualized growth of central government expenditure per capita $\Delta g$ | | | | |
| Japan | 2.77 | 825.59 | 0.20 | 0.87 |
| United States | 2.72 | 404.97 | 0.44 | 4.33 |
| United Kingdom | 2.25 | 65.31 | 1.11 | 7.31 |
| Canada | 1.01 | 124.81 | 0.77 | 9.40 |
| Australia | 2.03 | 421.74 | 0.41 | 3.58 |
| Finland | 2.39 | 256.68 | 0.04 | 12.10 |
| Annualized growth of central and local government expenditure per capita $\Delta g$ | | | | |
| Japan | 2.51 | 12.55 | -0.50 | 16.90 |
| United States | 0.74 | 11.65 | 0.63 | 2.76 |
| United Kingdom | 1.38 | 22.79 | 0.75 | 4.10 |
| Canada | 1.27 | 17.19 | 0.46 | 0.92 |
| Australia | 0.56 | 9.44 | -0.33 | 4.02 |
| Finland | 2.12 | 52.67 | -1.04 | 13.91 |
| Annualized growth of central government investment per capita $\Delta g$ | | | | |
| Japan | -0.23 | 168.67 | 0.28 | 1.00 |
| United States | 3.32 | 90.41 | 0.52 | 1.53 |
| United Kingdom | 2.48 | 1386.22 | 0.47 | 3.16 |
| Canada | 1.60 | 69.47 | 0.22 | 0.34 |
| Australia | 2.15 | 1458.71 | 0.13 | 3.17 |
| Finland | 0.91 | 379.82 | 0.02 | 3.24 |
| Annualized inflation rate $\pi$ | | | | |
| Japan | 2.33 | 20.30 | 3.16 | 16.24 |
| United States | 3.84 | 9.14 | 0.88 | 3.02 |
| United Kingdom | 5.12 | 25.37 | 2.18 | 5.92 |
| Canada | 3.82 | 10.41 | 0.86 | 0.44 |
| Australia | 4.74 | 17.90 | 1.32 | 1.98 |
| Finland | 4.38 | 19.13 | 1.23 | 1.04 |
| Change in the short-term rate $\Delta r$ | | | | |
| Japan | -0.30 | 2.16 | 1.05 | 2.11 |
| United States | -0.03 | 4.24 | -0.05 | 0.76 |
| United Kingdom | -0.07 | 5.67 | -0.05 | 0.92 |
| Canada | -0.04 | 4.43 | 0.01 | 2.42 |
| Australia | -0.21 | 5.43 | -0.36 | 1.85 |
| Finland | -0.29 | 2.96 | -0.52 | 2.31 |
| Log of the country level index of financial stress $\varphi$ | | | | |
| Japan | -2.08 | 0.63 | -0.25 | -0.70 |
| United States | -2.27 | 0.83 | -0.20 | -0.69 |
| United Kingdom | -2.43 | 0.68 | -0.05 | -0.27 |
| Canada | -2.47 | 0.66 | -0.07 | -0.57 |
| Australia | -1.98 | 0.54 | -0.21 | -0.14 |
| Finland | -2.24 | 0.55 | 0.14 | -0.43 |

# E Algorithm for impulse responses

Consider B the bootstrapped version of the initial dataset, such that we have B+1 dataset including the initial dataset. The algorithm to compute the generalized impulse response function (GIRF) and the quantile impulse response functions (QIRF) is as follows:

---

**for** each dataset b = 0 to B **do**

  Estimate $\hat{\mathbf{\Gamma}}(b)$, $\hat{\boldsymbol{\beta}}_\tau(b)$ and $\hat{\boldsymbol{C}}(b)$ on a dataset of size $N$

  **for** each initial condition n = 1 to N **do**

    Pick $\boldsymbol{y}(n)$ as initial condition, the $n$th observation of the dataset.

    $\boldsymbol{y}(b, n, 0; \delta) = \boldsymbol{y}(n) + \mathbf{C}(b)\delta$ for an initial impulse $\delta = \{0, 1\}$ or $\{1, 0\}$
    $\boldsymbol{y}(b, n, 0; 0) = \boldsymbol{y}(n)$ for the control with no initial impulse

    For $Y_1$ with a conditional density, draw $\{\omega(h)\}_{h=1}^H$ from the uniform distribution $\mathcal{U}(0, 1)$

    For $Y_2$ with the VAR structure, resample from the structural residuals $\{u_2(h)\}_{h=1}^H$

    **for** each horizon h = 1 to H **do**

      For $Y_1$ with a conditional density:
      $y_1(b, n, h; \delta)$ is the percentile $\omega(h)$ of $\hat{f}(y_1(b, n, h; \delta)|\hat{Q}_{\mathcal{T}}(Y_1|\boldsymbol{y}(b, n, h{-}1; \delta); \boldsymbol{\beta}_\tau(b)))$
      $y_1(b, n, h; 0)$ is the percentile $\omega(h)$ of $\hat{f}(y_1(b, n, h; 0)|\hat{Q}_{\mathcal{T}}(Y_1|\boldsymbol{y}(b, n, h{-}1; 0); \boldsymbol{\beta}_\tau(b)))$

      For $Y_2$ with the VAR structure:
      $y_2(b, n, h; \delta) = \hat{\mathbf{\Gamma}}(b)\mathbf{y}(b, n, h{-}1; \delta) + \hat{\boldsymbol{C}}(b)u_2(h)$
      $y_2(b, n, h; 0) = \hat{\mathbf{\Gamma}}(b)\mathbf{y}(b, n, h{-}1; 0) + \hat{\boldsymbol{C}}(b)u_2(h)$

    **end for**
  **end for**
**end for**

$\text{GIRF}_{\mathbf{Y}}(b{=}0, h; \delta) = \mathbb{E}_n(\mathbf{y}(b{=}0, n, h; \delta) - \mathbf{y}(b{=}0, n, h; 0))$
$\text{GIRF\_CI}_{\mathbf{Y}}(h; \delta) = \{\mathbb{Q}_{.05}(\text{GIRF}_{\mathbf{Y}}(b, h; \delta)); \mathbb{Q}_{.95}(\text{GIRF}_{\mathbf{Y}}(b, h; \delta))\}$

$\text{QIRF}_{Y_1,\mathcal{T}}(b{=}0, h; \delta) = \mathbb{E}_n(Q_{\mathcal{T}}(Y_1|\boldsymbol{y}(b{=}0, n, h{-}1; \delta)) - Q_{\mathcal{T}}(Y_1|\boldsymbol{y}(b{=}0, n, h{-}1; 0))))$
$\text{QIRF\_CI}_{Y_1,\mathcal{T}}(h; \delta) = \{\mathbb{Q}_{.05}(\text{QIRF}_{Y_1,\mathcal{T}}(b, h; \delta)), \mathbb{Q}_{.95}(\text{QIRF}_{Y_1,\mathcal{T}}(b, h; \delta))\}$

---

The generalized IRF (GIRF) for all variables $\mathbf{Y}$ and quantile IRF (QIRF) for variable $Y_1$ for percentiles $\mathcal{T} = \{0.10, 0.25, 0.50, 0.75, 0.90\}$ are computed on the initial dataset $b = 0$. $\mathbb{E}_n$ is the mean over the initial conditions N. $\mathbb{Q}_{.05}$ and $\mathbb{Q}_{.95}$ are the percentiles across all resampled dataset B. Note that QIRF is not defined at $h$=1: it would depend on the heterogeneity in the variables at period 0, but, by construction, there is no difference before a shock occurs.

# F Additional robustness

Figures 8 and 10 show robustness results for different lag structures, from two to four lags. Figure 9 shows similar percentiles of the GDP density forecast when the model is estimated with the short-term rate in level, even if the model is no longer stationary. Figure 11 shows the response to various government spending shocks with four lags instead of the benchmark of three lags.

Figure 8: Quantile response of GDP growth after a monetary policy shock - robustness for different lags

(a) Model with two lags (b) Model with three lags (c) Model with four lags

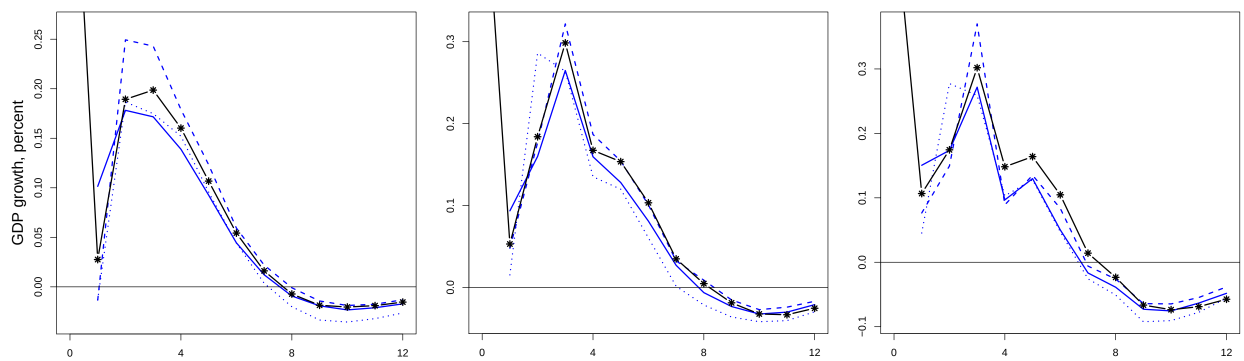


Notes: This figure displays the response of the 90th percentile (dashed lines), the 50th percentile (plain line) and the 10th percentile (dotted line) of the probability density of future GDP growth, up to a 12-quarter horizon, following a 50-basis-point easing of monetary policy for a model estimated with two, three or four lags. The starred black line is the mean response from of the linear VAR.

Figure 9: Quantile response of GDP growth after a policy shock - robustness for policy rate in level

(a) Monetary policy shock (b) Government spending shock

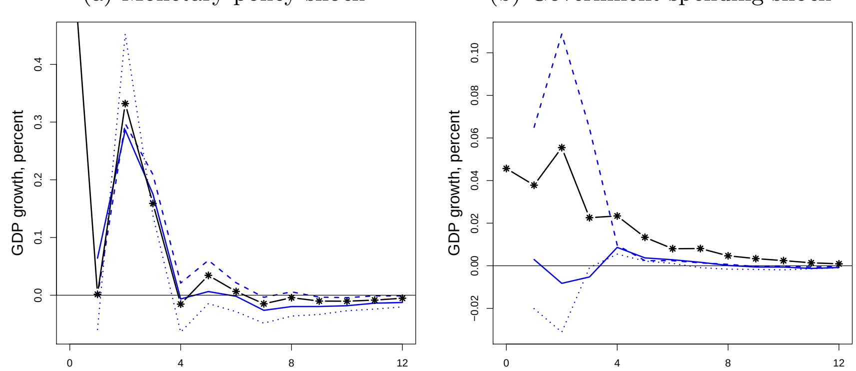


Notes: This figure displays the response of the 90th percentile (dashed lines), the 50th percentile (plain line) and the 10th percentile (dotted line) of the probability density of future GDP growth, up to a 12-quarter horizon, to a 50-basis-point easing of monetary policy (left) and a 5 percent spending shock to the expenditures of the central government (right), for a model estimated with the short-term rate in level with three lags. The starred black line is the mean response from of the linear VAR.

Figure 10: Quantile response of GDP growth after a government spending shock - robustness for different lags

(a) Model with two lags (b) Model with three lags (c) Model with four lags

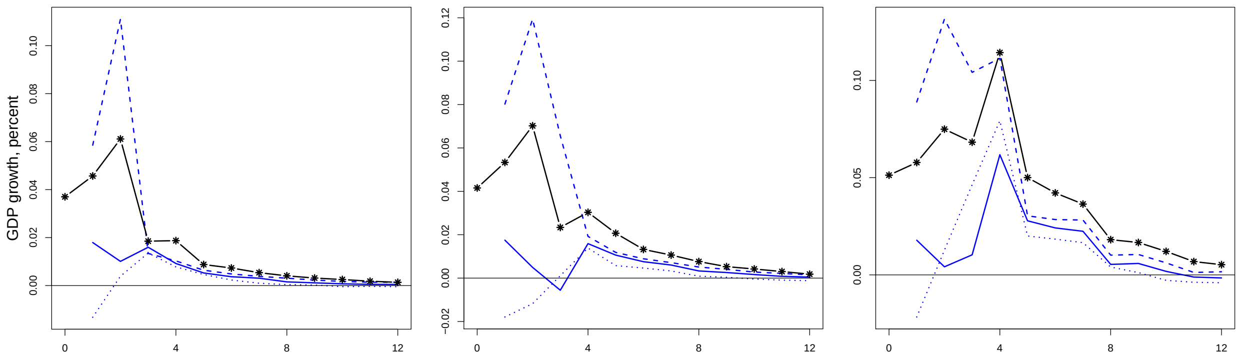


Notes: This figure displays the response of the 90th percentile (dashed lines), the 50th percentile (plain line) and the 10th percentile (dotted line) of the probability density of future GDP growth, up to a 12-quarter horizon, following a 5 percent spending shock to the expenditures of the central government for a model estimated with two, three or four lags. Spending shocks are identified by zero restrictions assuming no contemporaneous response of government spending to other macro shocks. Monetary policy is kept fixed. The starred black line is the mean response from of the linear VAR.

Figure 11: Quantile response of GDP growth after different types of government spending shocks - robustness different lags

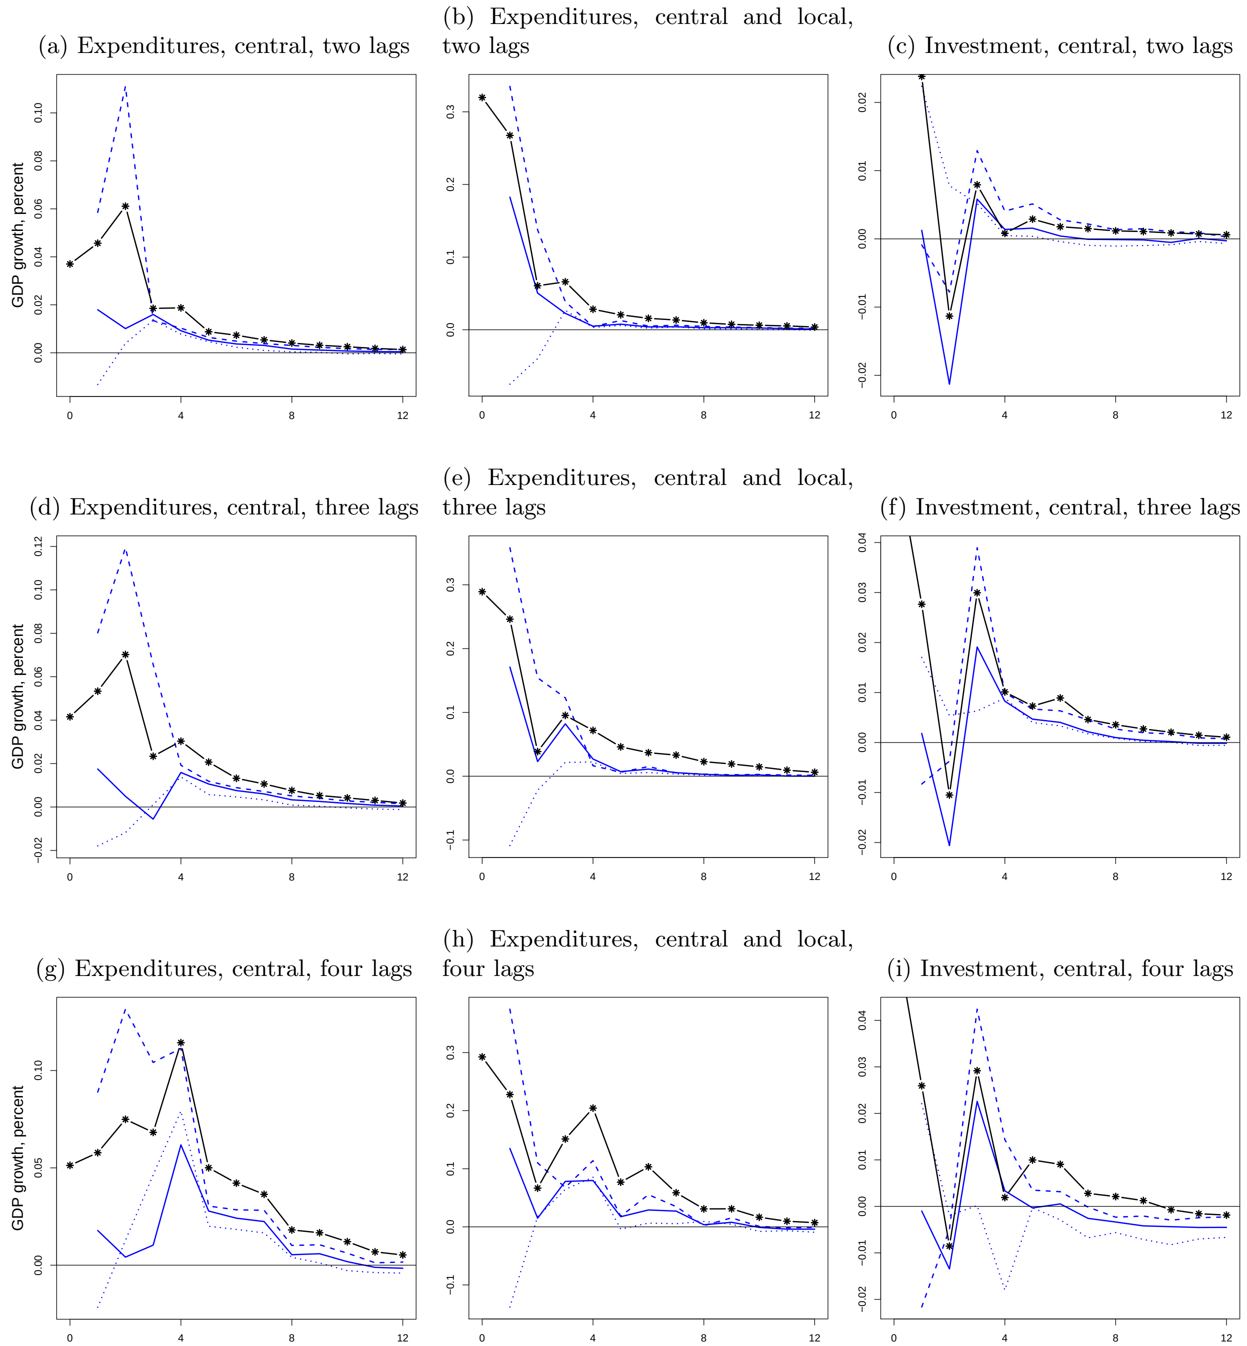


(a) Expenditures, central, two lags

(b) Expenditures, central and local, two lags

(c) Investment, central, two lags

(d) Expenditures, central, three lags

(e) Expenditures, central and local, three lags

(f) Investment, central, three lags

(g) Expenditures, central, four lags

(h) Expenditures, central and local, four lags

(i) Investment, central, four lags

Notes: This figure displays the response of the 90th percentile (dashed lines), the 50th percentile (plain line) and the 10th percentile (dotted line) of the probability density of future GDP growth, up to a 12-quarter horizon, following a 5 percent spending shock to the expenditures of the central government, the expenditures from both central and local governments, and the investments of the central government. Spending shocks are identified by zero restrictions assuming no contemporaneous response of government spending to other macro shocks. Model estimated with four lags. Monetary policy is kept fixed. The starred black line is the mean response from of the linear VAR.

Table 8: RMSE of one-step-ahead forecasts of GDP growth - ZLB episodes and other periods

| Lag order | In-sample until 2007 | | | | Out-of-sample from 2007 onwards | | | |
|---|---|---|---|---|---|---|---|---|
| No ZLB | **OLS** | Skew-t | P50 | Kernel | OLS | Skew-t | **P50** | Kernel |
| Lag 1 | 4.056 | 4.075 | 4.115 | 4.122 | 4.372 | 4.175 | 3.908 | 3.927 |
| Lag 2 | 3.983 | 3.999 | 4.032 | 4.044 | 4.373 | 4.166 | **3.885** | 3.938 |
| Lag 3 | 3.850 | 3.865 | 3.897 | 3.924 | 4.462 | 4.172 | 4.017 | 4.014 |
| Lag 4 | **3.768** | 3.792 | 3.830 | 3.848 | 4.774 | 4.512 | 4.364 | 4.344 |
| ZLB | OLS | Skew-t | P50 | **Kernel** | OLS | Skew-t | P50 | **Kernel** |
| Lag 1 | 2.236 | 2.215 | 2.251 | 2.200 | 3.173 | 3.151 | 3.153 | 3.104 |
| Lag 2 | 2.202 | 2.167 | 2.201 | 2.188 | 3.142 | 3.078 | 3.069 | 3.057 |
| Lag 3 | 2.260 | 2.281 | 2.318 | 2.209 | 3.209 | 3.067 | 3.117 | **3.040** |
| Lag 4 | 2.201 | 2.180 | 2.237 | **2.159** | 3.252 | 3.144 | 3.128 | 3.084 |

| | Cross-validation: training | | | | Cross-validation: test | | | |
|---|---|---|---|---|---|---|---|---|
| No ZLB | **OLS** | Skew-t | P50 | Kernel | OLS | Skew-t | P50 | **Kernel** |
| Lag 1 | 3.958 | 3.973 | 3.995 | 3.993 | 4.441 | 4.385 | 4.365 | 4.383 |
| Lag 2 | 3.891 | 3.903 | 3.918 | 3.928 | 4.339 | 4.353 | 4.346 | 4.294 |
| Lag 3 | 3.778 | 3.797 | 3.816 | 3.829 | 4.245 | 4.337 | 4.237 | **4.227** |
| Lag 4 | **3.756** | 3.779 | 3.809 | 3.809 | 4.257 | 4.355 | 4.264 | 4.256 |
| ZLB | OLS | Skew-t | P50 | **Kernel** | OLS | Skew-t | P50 | **Kernel** |
| Lag 1 | 2.758 | 2.745 | 2.784 | 2.735 | 2.766 | 2.571 | 2.598 | 2.562 |
| Lag 2 | 2.705 | 2.705 | 2.724 | 2.702 | 2.641 | 2.658 | 2.518 | 2.565 |
| Lag 3 | 2.758 | 2.730 | 2.762 | 2.718 | 2.598 | 2.551 | 2.507 | 2.514 |
| Lag 4 | 2.657 | 2.633 | 2.675 | **2.632** | 2.513 | 2.478 | 2.513 | **2.469** |

Notes: The table shows the root mean squared error (RMSE) for one-step-ahead GDP growth modelled as the mean fitted via ordinary least square (OLS), the mean recovered from a skew-t distribution on a limited set of fitted quantiles, the median fitted by a quantile or the mean recovered from a kernel smoothing of a large number of fitted quantiles, with one to four lags. The RMSE is computed during ZLB episodes and outside of these. For the first part of the table, the RMSE is computed in-sample until 2007 and for subsequent out-of-sample periods. For the second part of the table, the RMSE is computed with cross-validation: the training dataset consists of all but one country, and the test dataset is the country not used for the estimation. We report the RMSE averaged over all possible folds of the dataset that experienced ZLB episodes (United States, United Kingdom, Japan, Finland). We bold the models that provide the smallest RMSE in the respective class. For comparison, the mean of the annualized real GDP growth per capita is 1.89 and the unconditional standard deviation is 3.99.

# G Policy shocks during recessions

We now focus on the state dependence of the shock propagation to the distribution of GDP growth during recessions.

## G.1 Model estimation setup

The state $s_t(REC)$ is a recession when real GDP growth is negative for two consecutive quarters:

$$s_t(REC) = \begin{cases} 1 & \text{if } \Delta y_t < 0 \text{ and } \Delta y_{t-1} < 0 \\ 0 & \text{otherwise.} \end{cases} \tag{13}$$

We adjust the previous model of Equation (9) by adding the recession dummy $s_t$ and an interaction term with the monetary and fiscal policy variables (the parameters indexed by $s_t$). We do not interact all variables with the recession dummy to keep the model parsimonious.[20]

$$\begin{aligned}
\Delta y_t &= f(\Delta y_{t-1}, \Delta g_{t-1}, \Delta r_{t-1}, \pi_{t-1}, \varphi_{t-1}, \alpha_c, s_t; \boldsymbol{\beta}_\tau(L)) + \epsilon_{1,t} \qquad (14)\\
\Delta g_t &= \boldsymbol{\Gamma}_{2,0}\alpha_c + \sum_{l=1}^{L}\Big[\boldsymbol{\Gamma}_{2,1,l}\Delta y_{t-l} + \boldsymbol{\Gamma}^{s_t}_{2,2,l}\Delta g_{t-l} + \boldsymbol{\Gamma}^{s_t}_{2,3,l}\Delta r_{t-l} + \boldsymbol{\Gamma}_{2,4,l}\pi_{t-l} + \boldsymbol{\Gamma}_{2,5,l}\varphi_{t-l}\Big] + \boldsymbol{\Gamma}_{2,6}s_t + \epsilon_{2,t}\\
\Delta r_t &= \boldsymbol{\Gamma}_{3,0}\alpha_c + \sum_{l=1}^{L}\Big[\boldsymbol{\Gamma}_{3,1,l}\Delta y_{t-l} + \boldsymbol{\Gamma}^{s_t}_{3,2,l}\Delta g_{t-l} + \boldsymbol{\Gamma}^{s_t}_{3,3,l}\Delta r_{t-l} + \boldsymbol{\Gamma}_{3,4,l}\pi_{t-l} + \boldsymbol{\Gamma}_{3,5,l}\varphi_{t-l}\Big] + \boldsymbol{\Gamma}_{3,6}s_t + \epsilon_{3,t}\\
\pi_t &= \boldsymbol{\Gamma}_{4,0}\alpha_c + \sum_{l=1}^{L}\Big[\boldsymbol{\Gamma}_{4,1,l}\Delta y_{t-l} + \boldsymbol{\Gamma}^{s_t}_{4,2,l}\Delta g_{t-l} + \boldsymbol{\Gamma}^{s_t}_{4,3,l}\Delta r_{t-l} + \boldsymbol{\Gamma}_{4,4,l}\pi_{t-l} + \boldsymbol{\Gamma}_{4,5,l}\varphi_{t-l}\Big] + \boldsymbol{\Gamma}_{4,6}s_t + \epsilon_{4,t}\\
\varphi_t &= \boldsymbol{\Gamma}_{5,0}\alpha_c + \sum_{l=1}^{L}\Big[\boldsymbol{\Gamma}_{5,1,l}\Delta y_{t-l} + \boldsymbol{\Gamma}^{s_t}_{5,2,l}\Delta g_{t-l} + \boldsymbol{\Gamma}^{s_t}_{5,3,l}\Delta r_{t-l} + \boldsymbol{\Gamma}_{5,4,l}\pi_{t-l} + \boldsymbol{\Gamma}_{5,5,l}\varphi_{t-l}\Big] + \boldsymbol{\Gamma}_{5,6}s_t + \epsilon_{5,t}
\end{aligned}$$

The function $f(\Delta y_{t-1}, \Delta g_{t-1}, \Delta r_{t-1}, \pi_{t-1}, \varphi_{t-1}, \alpha_c, s_t; \boldsymbol{\beta}_\tau(L))$ is the kernel density of GDP growth recovered from the ordered quantiles $\mathcal{T}$ where each quantile $\tau \in \mathcal{T}$ is estimated as:

$$Q_\tau(\Delta y_t) = \boldsymbol{\beta}_{\tau,0}\boldsymbol{\alpha}_{\tau,c} + \sum_{l=1}^{L}\Big[\boldsymbol{\beta}_{\tau,1,l}\Delta y_{t-l} + \boldsymbol{\beta}^{s_t}_{\tau,2,l}\Delta g_{t-l} + \boldsymbol{\beta}^{s_t}_{\tau,3,l}\Delta r_{t-l} + \boldsymbol{\beta}_{\tau,4,l}\pi_{t-l} + \boldsymbol{\beta}_{\tau,5,l}\varphi_{t-l}\Big] + \beta_{\tau,6}s_t. \tag{15}$$

When we do the bootstrap, we require that the number of recession episodes in the resampled data be plus or minus 10 percent of the frequency of the actual dataset. Otherwise, an over-representation of the recession events could have an impact on the bootstrapped tail estimates of the GDP growth distribution. We further require that at least 5 percent of the resampled dataset fall in the recession regime, to make sure that we have enough observations in the randomly resampled dataset to estimate the state dependence parameters.

Table 9 shows the RMSE gain outside and during recession episodes. In the latter case, nonlinear models tend to outperform OLS in out-of-sample one-step-ahead forecasts.

[20] Interacting all variables would yield qualitatively similar results, albeit at the cost of doubling the number of parameters.

Table 9: RMSE of one-step-ahead forecasts of GDP growth - recessions and normal times

| Lag order | In-sample until 2007 | | | | Out-of-sample from 2007 onwards | | | |
|---|---|---|---|---|---|---|---|---|
| No recession | **OLS** | Skew-t | P50 | Kernel | OLS | Skew-t | P50 | **Kernel** |
| Lag 1 | 3.639 | 3.654 | 3.690 | 3.719 | 3.345 | 3.219 | 2.987 | 2.998 |
| Lag 2 | 3.562 | 3.582 | 3.616 | 3.636 | 3.320 | 3.195 | 2.964 | 2.966 |
| Lag 3 | 3.541 | 3.561 | 3.592 | 3.611 | 3.131 | 3.040 | 2.847 | **2.839** |
| Lag 4 | **3.485** | 3.509 | 3.548 | 3.557 | 3.290 | 3.216 | 3.099 | 3.054 |
| Recession | OLS | **Skew-t** | P50 | Kernel | OLS | Skew-t | **P50** | Kernel |
| Lag 1 | 5.220 | 5.209 | 5.294 | 5.174 | 6.757 | 6.683 | **6.400** | 6.464 |
| Lag 2 | 5.048 | 5.055 | 5.098 | 5.042 | 6.819 | 6.617 | 6.595 | 6.456 |
| Lag 3 | 4.582 | 4.594 | 4.772 | 4.736 | 6.854 | 6.709 | 6.660 | 6.571 |
| Lag 4 | 4.469 | **4.461** | 4.690 | 4.701 | 6.966 | 6.861 | 6.612 | 6.607 |

| | Cross-validation: training | | | | Cross-validation: test | | | |
|---|---|---|---|---|---|---|---|---|
| No recession | **OLS** | Skew-t | P50 | Kernel | OLS | Skew-t | P50 | **Kernel** |
| Lag 1 | 3.419 | 3.437 | 3.470 | 3.482 | 3.522 | 3.495 | 3.528 | 3.547 |
| Lag 2 | 3.356 | 3.378 | 3.411 | 3.419 | 3.451 | 3.431 | 3.499 | 3.447 |
| Lag 3 | 3.318 | 3.333 | 3.371 | 3.378 | 3.410 | 3.395 | 3.415 | **3.395** |
| Lag 4 | **3.315** | 3.333 | 3.377 | 3.379 | 3.487 | 3.458 | 3.558 | 3.438 |
| Recession | **OLS** | Skew-t | P50 | Kernel | OLS | Skew-t | P50 | **Kernel** |
| Lag 1 | 5.586 | 5.547 | 5.546 | 5.483 | 5.616 | 5.681 | 5.671 | 5.580 |
| Lag 2 | 5.454 | 5.489 | 5.452 | 5.399 | 5.682 | 5.773 | 5.783 | 5.550 |
| Lag 3 | 5.265 | 5.298 | 5.346 | 5.301 | 5.706 | 5.753 | 5.661 | **5.483** |
| Lag 4 | **5.194** | 5.222 | 5.327 | 5.319 | 5.936 | 6.030 | 6.033 | 5.809 |

Notes: The table shows the root mean squared error (RMSE) for one-step-ahead GDP growth modelled as the mean fitted by ordinary least square (OLS), the mean recovered from a skew-t distribution on a limited set of fitted quantiles, the median fitted by a quantile or the mean recovered from a kernel smoothing of a large number of fitted quantiles, with one to four lags. The RMSE is computed on recession events or not. For the first part of the table, the RMSE is computed in-sample until 2007 and on the corresponding out-of-sample. For the second part of the table, the RMSE is computed with cross-validation: the training dataset consists of all but one country, and the test dataset is the country not used for the estimation. We report the RMSE averaged over all possible folds of the dataset. In bold are the models that provide the smallest RMSE. For comparison, the mean of the annualized real GDP growth per capita is 1.89 and the unconditional standard deviation is 3.99.

## G.2 Monetary policy shock during recessions

We now consider the impulse response function of the GDP growth outlook following a monetary policy shock (Figure 12, left column). However, we distinguish between the response of an economy in recession (dashed red lines) and normal times (plain blue lines). Figure 13, left column, shows the mean impulse response.

Most quantile or mean responses to the monetary policy shock are very similar during recessions and normal times: monetary policy stimulus significantly shifts the future GDP growth distribution up in both normal and recession times. However, the upper quantiles during a recession initially have a higher point estimate. Although the confidence bands of the recession times are too large to yield a significant difference with normal times, the point estimate in recession times lies outside the normal times confidence bands. This suggests that monetary policy helps to increase the upside risk to GDP growth during recession events relatively more than in normal times. This marginally increases the mean response of GDP growth to a monetary policy shock during a recession.

## G.3 Fiscal policy shock during recessions

Turning to the consequences of a fiscal expansion during recession episodes, we discuss the impulse responses of different quantiles of the GDP growth outlook (Figure 12, right column). Here we consider a central government expenditure increase during a recession (dashed red lines) and in normal times (plain blue lines). Figure 12, right column, shows the mean response.

The average response in recessions is similar to that in normal times. But there is some noticeable difference across point estimates of quantiles; see Figure 14.

The response of the upper quantiles is larger and positive (significant for the 90th percentile for a one-sided test) while the response of the lower quantiles is smaller, non-significant and possibly negative. This suggests that government spending works mostly by increasing upside risks and making a recovery faster. Our results complement Auerbach and Gorodnichenko (2012b), who find that government spending was most effective during recessions.

The point forecast is lower during recessions for the lower tail. Because of large confidence bands during recessionary episodes, the difference is not significant. However, the point forecast of the lower tail of GDP growth is outside the confidence bands of the response in normal times, suggesting a possible misallocation of resources that can further amplify tail risks during a recession. The potential negative impact in the tail quantile responses is averaged out when looking at the mean response, which is slightly less positive during a recession event than in normal times.

Figure 12: Quantile impulse response to monetary and fiscal shocks: normal times and recessions

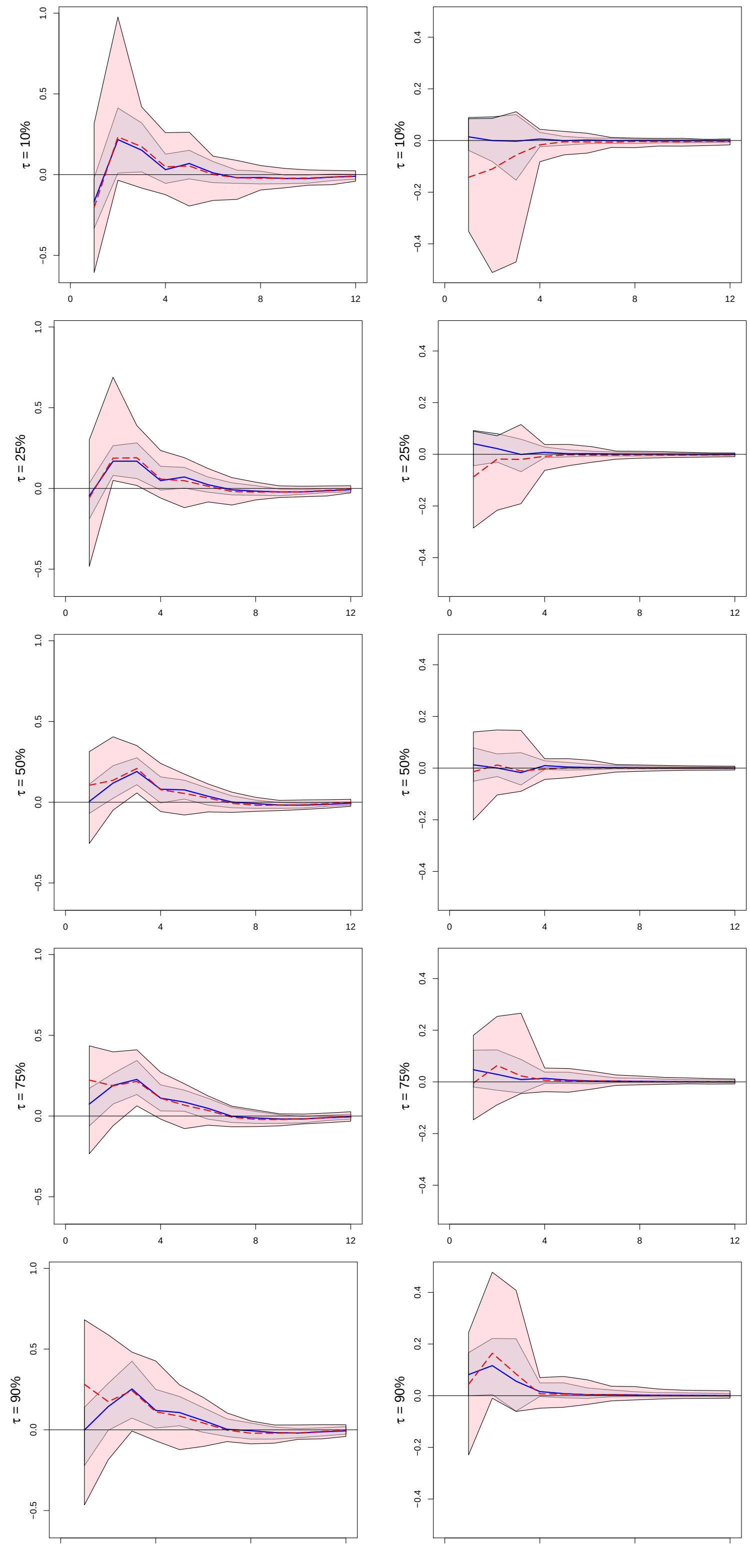


Notes: The figure shows the response of the 10th, 25th, 50th, 75th and 90th percentiles of the distribution of future GDP growth following a 50-basis-point monetary policy easing (left) and a 5 percent central government expenditure increase, while keeping the response of monetary policy fixed (right). The plain blue lines correspond to the response in normal times, the dashed red lines the response during recessions. The model is estimated with three lags. The shaded areas represent the 90 percent confidence interval computed with block bootstrap.

Figure 13: Mean impulse response to monetary and fiscal shocks: normal times and recessions

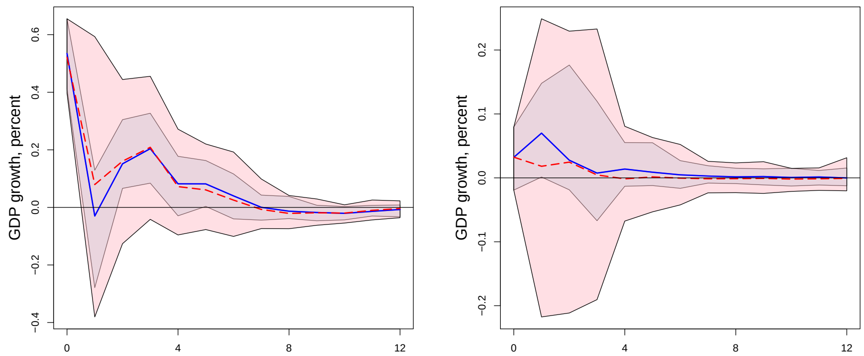


Notes: The figure shows the response of the mean of the distribution of future GDP growth in normal times (plain blue line) and during recessions (dashed red line) following a 50-basis-point monetary policy easing (left) and a 5 percent central government expenditure increase, while keeping the response of monetary policy fixed (right). The model is estimated with three lags. The shaded areas represent the 90 percent confidence interval computed with block bootstrap.

Figure 14: Summary of the quantile response to monetary and fiscal shocks: normal times and recessions

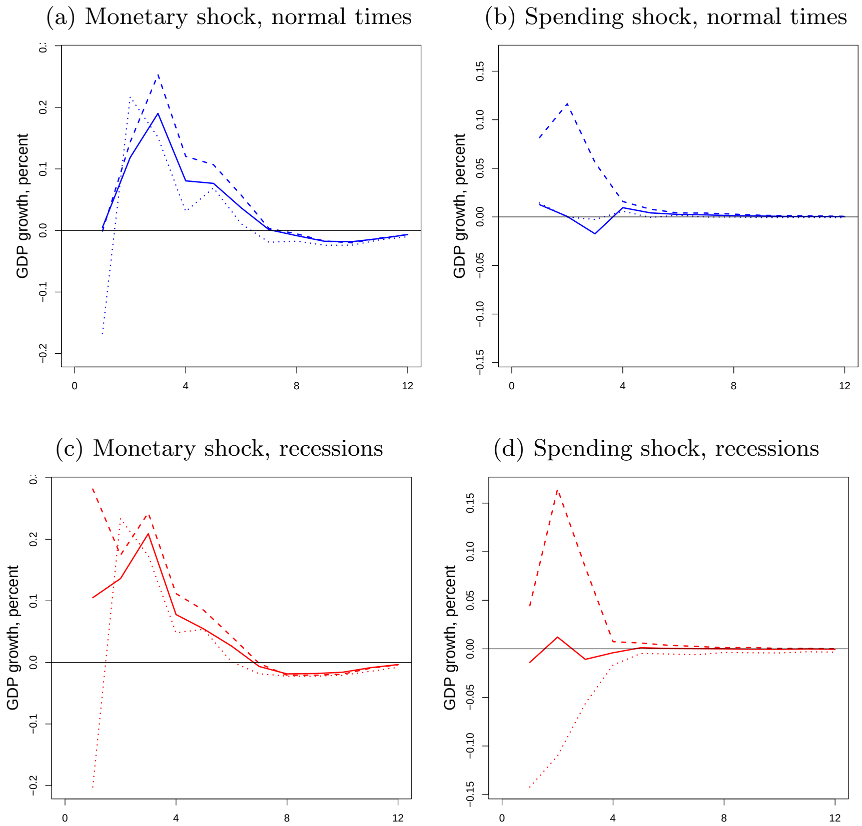


Notes: This figure displays the response of the 90th percentile (dashed lines), the 50th percentile (plain line) and the 10th percentile (dotted line) of the probability density of future GDP growth, up to a 12-quarter horizon, following a 50-basis-point monetary policy easing (left) and a 5 percent central government expenditure increase, while keeping the response of monetary policy fixed (right). Spending shocks are identified by zero restrictions assuming no contemporaneous response of government spending to other macro shocks. Model estimated with three lags.