

Assessing Macroeconomic Tail Risk

Francesca Loria*

Christian Matthes[†]

Donghai Zhang[‡]

January 9, 2020

Abstract

What drives macroeconomic tail risk? To answer this question, we borrow a definition of macroeconomic risk from [Adrian *et al.* \(2019\)](#) by studying (left-tail) percentiles of the forecast distribution of GDP growth. We use local projections ([Jordà, 2005](#)) to assess how this measure of risk moves in response to economic shocks to the level of technology, monetary policy, and financial conditions. We find that contractionary shocks disproportionately increase downside risk, independently of what shock we study. Our results confirm that accounting for financial conditions is crucial for modeling the transmission of economic shocks to macroeconomic tail risk.

Keywords: Macroeconomic Risk, Shocks, Local Projections

JEL Classification: C21, C53, E17, E37

We thank Marvin Nöller for excellent research assistance. We also want to thank Danilo Cascaldi-Garcia, Dario Caldara, Fabio Canova, Giovanni Favara, Domenico Giannone, Ricardo Reis, Erick Sager, Rosen Valchev, and Alex Wolman as well as seminar/workshop participants at the Chinese University of Hong Kong, the University of Bonn, the International Finance Workshop on Quantile Regressions at the Board of Governors, the 2019 IAAE conference, and the "Modeling The Macroeconomy in Risky Times" Workshop in St. Louis for helpful comments.

Disclaimer: The views presented herein are those of the author and do not necessarily reflect those of the Federal Reserve Board, the Federal Reserve Bank of Richmond, the Federal Reserve System or their staff.

*Board of Governors of the Federal Reserve System, francesca.loria@frb.gov.

[†]Indiana University & Federal Reserve Bank of Richmond, matthesc@iu.edu.

[‡]Institute for Macroeconomics and Econometrics — University of Bonn, donghai.zhang@uni-bonn.de.

1 Introduction

Economic policymakers and market participants are generally not only worried about what changes to economic conditions will do to the economy *on average*, but also how these changes affect both the probability of *large* harmful events occurring as well as the magnitude of these events.¹ We borrow the concept of *value-at-risk* from the finance literature to operationalize the concept of macroeconomic tail risk. To be more precise, we follow [Adrian et al. \(2019\)](#) and study the distribution of macroeconomic risk by estimating a quantile forecast regression of real GDP growth over the next year for various quantiles. We focus on the 10th percentile and, as reference points, the median and 90th percentiles. We interpret this 10th percentile of the predictive distribution of future GDP growth as *macroeconomic tail risk*. Our contribution relative to [Adrian et al. \(2019\)](#) is that we ask how macroeconomic risk is shaped by various structural shocks. To do so, we use local projections ([Jordà, 2005](#)) that estimate the effects of these shocks on the risk measure introduced by [Adrian et al. \(2019\)](#). In particular, we consider shocks to the level of technology, monetary policy, and financial conditions. Doing so, we follow the large literature that directly uses measures of (or instruments for) structural shocks—see for example [Ramey \(2016\)](#), [Romer and Romer \(2004\)](#), or [Coibion and Gorodnichenko \(2012\)](#).

Standard impulse response functions in linear models such as Vector Autoregressions (VARs) are not built to answer these questions as they track average outcomes. While fully parametric VARs that specify the probability distribution of the one-step ahead forecast error could be used to derive such measures of risk, these VARs put many restrictions on the behavior of forecast densities.² Our goal is to provide a flexible, yet simple framework that can directly tackle these issues. A flexible approach is useful when studying GDP growth because GDP growth is far from Gaussian, as can be seen from papers that interpret the non-normality in GDP growth as coming from fat tails and/or stochastic volatility ([Fagiolo et al., 2008](#), [Cúrdia et al., 2014](#), [Primiceri, 2005](#), and [Justiniano and Primiceri, 2008](#)).

With the 10th percentile as well as the median and the 90th percentile in hand from our first-stage quantile regressions, we can further follow in the footsteps of [Adrian et al. \(2019\)](#) and fit a flexible (skewed-*t*) distribution to match various estimated quantiles. We can thus trace out how the entire predictive distribution of average future real GDP growth changes after

¹For research emphasizing that the Federal Reserve is concerned by downside risk, see [Kilian and Manganello \(2008\)](#). For direct evidence of a policymaker thinking about downside risk (which we will use synonymously with tail risk), see this March 2019 speech by Lael Brainard, member of the Board of Governors of the Federal Reserve System: <https://www.federalreserve.gov/newsevents/speech/brainard20190307a.htm>.

²Even non-linear time series models such as VARs with stochastic volatility and time-varying parameters along the lines of [Cogley and Sargent \(2005\)](#) and [Primiceri \(2005\)](#) impose substantial structure. [Carriero et al. \(2019\)](#) propose an extension to the standard stochastic volatility structure embedded in these models that helps to capture tail risks.

a shock hits the economy. We view these changes as summarizing the effect of shocks on the *economic outlook*. With this distribution at hand, we can ask how the occurrence of economic shocks changes the probability of seeing real GDP contract over the next year.

One key point to emphasize is that our approach is constructed to be as flexible as possible: In the initial quantile regression stage, we model each quantile separately instead of assuming a specific distribution for the forecast distribution of real GDP growth. In the second stage, we use local projections to impose as few restrictions on the data generating process as possible.³ Just as [Adrian et al. \(2019\)](#), we only use the skewed- t distribution *after* having estimated the quantiles separately.

A common pattern emerges when we study our shocks: Expansionary shocks compress the distribution of future GDP growth, thus making “bad” outcomes (those in the left tail) more tolerable. Unfortunately, as we show later, this result also implies that contractionary shocks make the 10th percentile fall more than the median—hence leading not only to poor average outcomes, but also to a further increase in downside risk. Complementing the analysis in [Adrian et al. \(2019\)](#), we find that the key channel through which shocks affect macroeconomic risk is via their effect on financial conditions.

The remainder of the paper is organized as follows. Section (2) presents our econometric methodology. Section (3) provides an intuition for how shocks might affect the shape of the distribution in different manners. Section (4) presents the main findings, Section (5) builds economic intuition and Section (6) concludes.

2 Econometric Methodology

Our approach consists of two steps, which we outline in this section. First, we compute a measure of macroeconomic tail risk following [Adrian et al. \(2019\)](#). Our second step then uses local projections ([Jordà, 2005](#)) to assess how structural shocks affect macroeconomic risk and the predictive distribution of average future real GDP growth more generally.

2.1 Conditional Quantiles

We compute conditional quantiles for real GDP growth following the method proposed by [Adrian et al. \(2019\)](#). In particular, we run a quantile regression ([Koenker and Bassett, 1978](#)) for the annualized average growth rate of real GDP between t and $t + 4$ (where t indicates quarters) by conditioning on a constant, the National Financial Conditions Index (NFCI), and real GDP growth at time t .⁴

³As shown by [Plagborg-Møller and Wolf \(2019\)](#), local projections and VARs asymptotically estimate the same impulse responses, but are on diametrically opposite ends of the bias-variance trade-off in finite samples.

⁴Our choice of conditioning variables (see [A](#)) is the same as in [Adrian et al. \(2019\)](#).

Formally, let y_{t+h} denote the value of real GDP growth between t and $t+h$ and let x_t denote a vector of conditioning variables, then the quantile regression solves the following optimization problem:

$$\hat{\gamma}_\tau = \underset{\gamma_\tau \in \mathbb{R}^k}{\operatorname{argmin}} \sum_{t=1}^{T-h} \left(\tau \cdot \mathbb{I}_{(y_{t+h} \geq x_t \gamma_\tau)} |y_{t+h} - x_t \gamma_\tau| + (1 - \tau) \cdot \mathbb{I}_{(y_{t+h} < x_t \gamma_\tau)} |y_{t+h} - x_t \gamma_\tau| \right), \quad (2.1)$$

where $\mathbb{I}(\cdot)$ denotes the indicator function and $\tau \in (0, 1)$ indicates the τ th quantile. The quantile of y_{t+h} conditional on x_t is then given by the predicted value from that regression⁵, defined as

$$\hat{Q}_{y_{t+h}|x_t}(\tau|x_t) = x_t \hat{\gamma}_\tau \equiv q_{\tau,t}. \quad (2.2)$$

Note that this approach does not require us to make any distributional assumptions about the error term in the quantile regression. Appendix A plots $q_{\tau,t}$ for various values of τ . In the following, we will analyze how different quantiles react to aggregate shocks.⁶

2.2 Impulse Responses

We estimate responses of different real GDP growth quantiles to a variety of aggregate shocks by applying the local projection method based on Jordà (2005). As a baseline, we run the following linear regression:

$$q_{\tau,t+s} = \alpha_{\tau,s} + \beta_{\tau,s} \text{shock}_t + \Psi(L)_{\tau,s} \text{controls}_t + \epsilon_{\tau,t+s}, \quad (2.3)$$

for $s = 0, \dots, S$ and where $q_{\tau,t+s}$ is the measure of risk at period $t+s$ for the τ th quantile (i.e., the quantile τ of the distribution of y_{t+s+h} conditional on information at time $t+s$) computed in the previous section, shock_t is the structural shock of interest, and $\Psi(L)_{\tau,s} \text{controls}_t$ is a lag polynomial of control variables which include the lagged quantiles and model-specific controls that we will discuss in detail later.⁷ Note that there are two distinct notions of “horizon” in our application. First, the horizon in the quantile regression h , which we keep fixed at 4 quarters. This first horizon captures how forward looking our measure of risk is. The second notion of horizon is s in the local projection, which we vary as we trace out how risk responds at different horizons to a shock at time t . The response of quantile q_τ at time $t+s$ to a *shock* at time t is then given by $\beta_{\tau,s}$.

⁵While Adrian *et al.* (2019) point out that the predicted value of y_{t+h} for any given τ is the conditional quantile at $t+h$ (conditional on time t information), we label the predicted value as today’s risk.

⁶It might be worthwhile to emphasize that x_t does *not* contain direct measures of structural shocks themselves, but only real GDP growth and the NFCI.

⁷To be specific, we have $\Psi(L)_{\tau,s} = \Psi_{1,\tau,s}L + \Psi_{2,\tau,s}L^2 + \dots + \Psi_{K,\tau,s}L^K$ so that $\Psi(L)_{\tau,s} \text{controls}_t$ only contains lags of the control variables.

We construct the impulse-response functions by the estimating the sequence of the $\beta_{\tau,s}$'s in a series of univariate regressions for each horizon. Confidence bands are based on Newey-West corrected standard errors that control for serial correlation in the error terms induced by the successive leading of the dependent variable.⁸

At this point it is useful to contrast our approach with another approach that aims to combine quantile regressions with local projections, an approach advocated for by [Linnemann and Winkler \(2016\)](#). We want to interpret the 10th percentile of 4-quarters ahead GDP growth as a measure of downside risk and we then ask how this measure of risk reacts to different shocks. We study a number of shocks and find it useful to use the same quantile (or measure of risk) for all shocks we study in our local projections. [Linnemann and Winkler \(2016\)](#), instead, are interested in one shock only and model the conditional quantiles *conditional on, among other things, a fiscal shock* and thus include the shock directly in the quantile regression. [Linnemann and Winkler \(2016\)](#) cannot distinguish between the two horizons h and s that we emphasized above (given that they ask a different question, they probably would not want to).⁹ An advantage of their approach is that they do not require a separate step to estimate the quantiles. We show later in section 5 that the asymmetry we emphasize is also present when directly estimating the effects of shocks via quantile regressions in a one-step approach, along the lines of [Linnemann and Winkler \(2016\)](#).

2.3 From Quantiles to Probabilities

With impulse responses to various quantiles at hand, we fit a flexible distribution to the estimated GDP growth quantiles after each shock. Why is it worthwhile to recover an estimate of the entire distribution? With such a distribution at hand, we can (i) get a better sense of how the entire distribution of real GDP growth shifts after a shock, and (ii) ask how a shock changes the probability of a contraction in real GDP (an exercise we carry out later).

Our flexible distribution is the skewed- t distribution developed by [Azzalini and Capitanio \(2003\)](#) and used by [Adrian et al. \(2019\)](#). It is given by the following density function

⁸On the right-hand side regressions we use both estimated lagged quantiles and shock series that generally also have been constructed via auxiliary regressions. Note that even in the presence of these generated regressors our parameter estimates are consistent (see [Murphy and Topel, 1985](#)). Previous literature has found ([Coibion and Gorodnichenko, 2012](#)) that it is usually not necessary to adjust standard errors because of the generated regressor nature of the shock series. For the right hand-side lagged quantiles, further evidence that our approach is not delivering misleading estimates is provided in section 5, where we directly estimate the same asymmetries using quantile regressions that do not use lagged quantiles as right-hand side variables.

⁹Another approach in empirical macroeconomics that uses quantile regressions is introduced in [Mumtaz and Surico \(2015\)](#), who use quantile autoregressive models to study state dependence in the consumption-interest rate relationship. Recent work that combines quantile regressions with VAR models to estimate impulse responses is presented in [Chavleishvili and Manganelli \(2017\)](#) and [Kim et al. \(2019\)](#).

for a data point y :

$$f(y|\mu, \sigma, \alpha, \nu) = \frac{2}{\sigma} t\left(\frac{y - \mu}{\sigma}; \nu\right) T\left(\alpha \frac{y - \mu}{\sigma} \sqrt{\frac{\nu + 1}{\nu + \left(\frac{y - \mu}{\sigma}\right)^2}}; \nu + 1\right). \quad (2.4)$$

As discussed in [Adrian *et al.* \(2019\)](#), t and T are the density and cumulative distribution function of the common t -distribution, μ is a location parameter, σ is a scale parameter, ν controls how fat the tails are (similar to the degrees of freedom in the common t -distribution), whereas α governs skewness because it controls how much the standard t -distribution is twisted (or skewed) according to T .

To translate the shocked quantiles to densities, we start with the average distribution of real GDP growth four quarters ahead using the sample of [Adrian *et al.* \(2019\)](#).¹⁰ We then calculate the impulse responses of the 5th, 25th, 75th, and 95th quantiles¹¹ to a given shock. Finally, at each impulse response horizon, the four parameters of the skewed- t distribution are set so as to exactly match those four quantiles.

3 Some Intuition for Impulse Responses of Quantiles

This section gives three examples where an initial distribution of a variable changes after a shock hits. We show these examples to illustrate how a change in quantiles is linked to changes in the distribution as a whole and how changes in specific moments translate into changes in quantiles. Our scenario is as follows: After an initial distribution of a scalar variable is hit by a shock, we trace out how this distribution changes on impact and in the period after impact. We consider three experiments:

1. The shock leads to an increase in the variance of our distribution, which is Gaussian.
2. The shock leads to an increase in the mean of our distribution, which is Gaussian.
3. The shock leads to an increase in the shape parameter of our distribution, which is distributed according to a Gamma distribution.

Figure 1 plots three panels for each experiment. The first panel in each row shows the initial distribution, the distribution when the shock hits, and the distribution in the period after the shock has materialized. The middle panel in each row shows the evolution of the 10th

¹⁰A plot of this average distribution can be found in the Appendix.

¹¹Following [Adrian *et al.* \(2019\)](#), we use impulse responses for the 5th, 25th, 75th, and 95th percentiles to match the percentiles that [Adrian *et al.* \(2019\)](#) used to compute the distributions in their paper. We show the impulse responses for those quantiles in Appendix E. They tell the same story as our choice of percentiles.

and 90th percentile for those three periods. The last panel in each row gives the impulse responses for the 10th and 90th percentiles *under the assumption that if the shock that moved the distributions had not taken on the specific value that led to the distribution moving, the distribution would have remained at its original position.*¹² As the impulse response plots the difference between the relevant percentiles and the original values, the impulse response figures only show values for two time periods (the period where the shock hits and the period after). Each row presents the figures for one experiment. Note that the levels of the percentiles are not directly interpretable as IRFs because we do not subtract the baseline value from the quantiles in those figures. As we can see, an increase in the variance of a symmetric distribution makes the quantiles drift apart in a mirror-image fashion, whereas a change in the mean of a symmetric distribution makes the quantiles move in parallel, which in turn makes the impulse responses lie on top of each other. With a non-symmetric distribution (or if a shock makes a distribution non-symmetric) the quantiles can drift apart, but not necessarily in a mirror-image fashion, as is the case in the last example.




Figure 1: Illustration of Changes in Percentiles.

¹²For our purposes, impulse responses are defined as the difference between two conditional expectations, where one of the expectations conditions on a specific value for one shock in one specific period.

Interpreting changes in multiple quantiles jointly can be challenging because we have to envision how the entire distribution changes. In Appendix D we plot changes in distributions to help the reader with interpretation. Nonetheless, it is useful to dig a bit deeper at this point. As an example, let us focus on the third experiment. As can be seen from the last panel on the bottom row of Figure 1, the 10th and 90th percentile drift apart because the 90th percentile *increases faster* than the 10th percentile. Thus the distribution *spreads out* as a result of the shock—this can also be seen by looking at the leftmost panel of the bottom row, where the yellow distribution is more spread out than the original blue distribution. Let us for a second imagine that this impulse response is the response to a “positive” shock and that quantiles react linearly to those shocks (as will be the case in our local projections), so that the response to a “negative” shock would just be the mirror image of the rightmost panel of the bottom row. What would happen to the distribution in that case? The 90th percentile would *decrease faster* than the 10th percentile. Hence the distribution would actually *compress* in that scenario.

Another scenario that could occur is that the impulse responses of various quantiles cross. It is important to emphasize that this does *not* mean that the quantiles themselves cross. In fact, a crossing of impulse responses of various quantiles can just mean that one quantile reverts back to its pre-shock value faster than another. In Appendix B we construct an example along the lines of those presented in Figure 1 where impulse responses cross without the actual quantiles crossing. One of our exercises below constructs the densities average real GDP growth over the next year (our variable of interest) implied by our impulse responses, partly to verify that the implied quantiles do not cross even though some of our estimated impulse responses do.

Our responses capture the difference between the expected path of the τ th quantile at horizon h after a given shock of specific (one standard deviation) size occurs and the expected path of the τ th quantile conditioning on such a specific shock value. That is, a response equal to 0 at horizon h means that the expected τ th quantile at horizon h is the same independently of whether we condition on a specific shock value in the initial period. In terms of magnitudes, the responses in our application capture the expected change (in percentage points) of the τ th quantile of average future GDP growth to a one-standard-deviation shock.

4 The Response of Tail Risk to Macroeconomic Shocks

We estimate responses in different quantiles of average one-year-ahead GDP growth to various aggregate shocks. In Appendix C, we show that our findings are robust to looking at real GDP growth at other horizons. All regressions are estimated at quarterly frequency and as a baseline we use four lags for all control variables. We give a brief overview of the various

specifications and data transformations. Most of the shocks considered here are reviewed in [Ramey \(2016\)](#). More details on our data sources are provided in Appendix [A](#).

Monetary Policy Shocks We explore two types of monetary policy shocks. First, we use the [Romer and Romer \(2004\)](#) (RR henceforth) narrative-based monetary shocks. They regress the federal funds target rate on Greenbook forecasts at each FOMC meeting date and use the residuals as the monetary policy shock. The sample period runs from 1973Q1 to 2007Q4. As a second measure, we use the monetary policy shocks identified by [Antolín-Díaz and Rubio-Ramírez \(2018\)](#) (AR henceforth) who add narrative sign restrictions to the VAR model in [Uhlig \(2005\)](#). Here the sample period is 1973Q1 to 2007Q3.¹³

For both types of shocks, we estimate the following regression:

$$q_{\tau,t+s} = \alpha_{\tau,s} + \beta_{\tau,s} shock_t + \Psi(L)_{\tau,s} \begin{bmatrix} q_{\tau,t} \\ shock_t \\ \pi_t^{cpi} \\ \pi_t^{com} \\ GDPgrowth_t \\ FFR_t \\ u_t \end{bmatrix} + \epsilon_{\tau,t+s}, \quad (4.1)$$

where $shock_t$ denotes RR or AR monetary shocks, π_t^{cpi} and π_t^{comm} respectively denote the consumer and the commodity price inflation (quarter over quarter), $GDPgrowth_t$ is the log of real GDP in first differences, FFR_t is the federal funds rate and u_t denotes the unemployment rate. Four lags of control variables are included.

Technology Shocks We consider two different technology shocks. First, a technology shock à la [Galí \(1999\)](#), constructed in a VAR¹⁴ by imposing that a technology shock is the only shock affecting labor productivity in the long-run. The sample period is 1975Q1 to 2007Q3. Second, we construct technology shocks by taking the growth rate of the [Fernald \(2012\)](#) utilization-adjusted TFP series for the aggregate economy. We refer to these shocks as “JF-TFP” shocks. They are available from 1974Q1 up to 2015Q3.

¹³Both the [Romer and Romer \(2004\)](#) and [Antolín-Díaz and Rubio-Ramírez \(2018\)](#) monetary policy shocks are aggregated to quarterly frequency by adding up the monthly values within each quarter.

¹⁴We estimate a VAR with four lags that includes three variables: changes in labor productivity, changes in hours, and changes in the GDP deflator. The technology shock is identified as the only shock affecting labor productivity in the long-run.

For each technology shock, we estimate the following regression:

$$q_{\tau,t+s} = \alpha_{\tau,s} + \beta_{\tau,s} shock_t + \Psi(L)_{\tau,s} \begin{bmatrix} q_{\tau,t} \\ shock_t \\ GDPgrowth_t \\ \Delta a_t \end{bmatrix} + \epsilon_{\tau,t+s}, \quad (4.2)$$

where $shock_t$ denotes JF or Galí technology shock and Δa_t is the log of productivity in first-differences. The latter is measured as real GDP divided by total hours.

Credit Spread Shocks We take the [Gilchrist and Zakrajšek \(2012\)](#) excess bond premium (EBP henceforth) updated by [Favara et al. \(2016\)](#) to construct an aggregate credit spread shock¹⁵. This shock can be identified by setting the controls in the local projection appropriately if one assumes that the bond premium affects interest rates contemporaneously but has no impact on prices and economic activity within a quarter. For a further discussion of how timing restrictions such as this can be incorporated in local projections, see [Barnichon and Brownlees \(2016\)](#) and [Plagborg-Møller and Wolf \(2019\)](#).

We estimate the following regression for the sample period from 1973Q1 to 2015Q4:

$$q_{\tau,t+s} = \alpha_{\tau,s} + \beta_{\tau,s} EBP_t + \lambda_{\tau,s} FFR_t + \Psi(L)_{\tau,s} \begin{bmatrix} q_{\tau,t} \\ EBP_t \\ \pi_t^{cpi} \\ GDPgrowth_t \\ FFR_t \end{bmatrix} + \epsilon_{\tau,t+s}. \quad (4.3)$$

4.1 Results

First, we present various impulse responses (i.e., $\beta_{\tau,s}$) based on equation 2.3. The $\beta_{\tau,s}$ coefficients can be interpreted as responses to one standard deviation shocks. We present results for three groups: monetary policy shocks, credit shocks, and technology shocks. Additional figures can be found in Appendix E. Throughout this paper, we plot impulse responses to *contractionary* shocks. We do this to help readers identify common patterns across shocks (which we discuss later). We plot the impulse responses of the 10th percentile, the median (along with its error bands), and the 90th percentile in Figure 2. The corresponding error bands for the other percentiles can be found in Appendix E. We then follow [Adrian et al. \(2019\)](#) in fitting a flexible (skewed- t) distribution on the estimated quantiles to show how

¹⁵We transform it to quarterly frequency by averaging the monthly values within each quarter.

the entire distribution of future real GDP growth shifts after shocks hit the economy. These results can be found in Appendix D.




Figure 2: Impulse Responses of Various Quantiles to Contractionary Shocks.

Note: Red (dashed) is response of the 10th quantile, black (solid) is the median response, blue (dotted) is response of the 90th quantile. Confidence bands correspond to median response, 90% significance level, based on Newey-West standard errors. The x-axis gives the horizon of the impulse response. The response on the y-axis is measured in percentage points.

Second, we use our fitted distribution for h -period-ahead real GDP growth to construct a time series of probabilities that real GDP growth h periods ahead is less than a threshold x .¹⁶ We then use these probabilities as left-hand side variables in local projection regressions to directly assess how important various structural shocks are for these probabilities. In Figure 3 we focus on a threshold of $x = 0$. That is, we are interested in whether a given structural shock increases the probability that real GDP will not grow over the next year ($h = 4$). We use the same specification for our local projections as in our benchmark analysis, just swapping

¹⁶With the estimated skewed t -distribution at hand, we can use numerical integration to compute this probability. Since this is a one-dimensional integration problem, standard numerical integration routines work very well in our context.

out the left-hand side variable.¹⁷




Figure 3: Impulse Responses of the Probability That Real GDP Will Not Grow Over The Next Year, Contractionary Shocks.

Note: Confidence bands correspond to a 90% significance level, based on Newey-West standard errors.

Notice that since we use linear local projections in the second stage of this exercise, our estimates are best interpreted as local linear approximations that are valid as long as the other conditioning variables in our local projections as well as the structural shocks do not take on values that are highly unusual.

4.2 Monetary Policy Shocks

The first panel of Figure 2 plots the responses to a contractionary RR monetary policy shock estimated via local projections. Those shocks affect the distribution of GDP growth disproportionately across quantiles. A contractionary (i.e., positive) monetary policy shock

¹⁷We have also carried out a similar exercise where, instead of using the fitted distribution to compute the relevant probabilities, we estimated 50 quantile regressions (for 50 evenly spaced quantiles) and constructed probabilities by interpolating between the quantiles. Results were very similar.

decreases the 10th percentile more than the median or the 90th percentile. This means that not only will a monetary policy shock lead to a decrease in *median* forecasted GDP growth four quarters ahead, but it will also make “bad” outcomes substantially worse by spreading out the left tail of the distribution.

The above result is robust to the use of an alternative monetary shock measure, namely the AR monetary shock (see the second panel in the top row of Figure 2). This shift is also evident from the top two panels of Figure 12, which plot the implied changes in the entire distribution of forecasted GDP growth.

Eventually (after between 10 to 15 quarters depending on the specific monetary shock measure), a contractionary monetary policy shock actually has a positive effect in that it shifts the distribution to the right (the impulse responses become positive). Given that we use linear local projections, the opposite mirror-image effect holds for expansionary shocks.

A contractionary monetary shock can have a substantial effect on the probability of real GDP not growing over the next year - the maximum point estimate over the horizons we consider is a 5 percent increase for a one standard deviation Romer and Romer (2004) shock.

4.3 Technology Shocks

The effects of a contractionary technology shock identified along the lines of Galí (1999) is shown in the third panel of Figure 2. The fourth panel of that figure shows the corresponding effects for a technology shock identified using Fernald (2012). We discuss the findings together since the results for both macroeconomic risk (i.e., the 10th percentile of the forecast distribution) and the entire shape of the economic outlook are similar across these specifications. An expansionary technology shock (the response to which is the mirror image of the responses plotted in the third and fourth panels of Figure 2) of any of the two types we consider here compresses the distribution of average future real GDP growth. This means that not only does a technology shock raise median GDP growth one year ahead, but it also makes low outcomes of future GDP growth more tolerable by shifting the distribution to the right—as can be seen in Figure 12. The only slight caveat to this interpretation is that at large horizons (more than three years out) the impulse response of the 10th percentile to a Fernald TFP shock becomes negative. This positive view of an expansionary technology shock comes with a downside: Contractionary technology shocks increases downside risk.

4.4 Credit Spread Shock

The fifth panel of Figure 2 plots the responses to a contractionary (i.e., positive) shock to the excess bond premium, which we interpret as an unexpected deterioration of financial conditions, just as Gilchrist and Zakrajšek (2012). The entire conditional distribution of GDP

growth is shifted, with the left tail being affected disproportionately more. On impact and up to one year, the interpretation of the effects of a contractionary credit shock is similar to the interpretation of the monetary policy shock given above. After one year, however, the responses of the 10th and 90th percentile cross, leading the distribution of future GDP growth to actually compress since the 10th percentile grows faster than the 90th percentile. One interpretation of these results is that policymakers counteract financial shocks, but that it takes around a year for these measures to take effect (potentially and partially due to lags in policy implementation). A one standard deviation contractionary credit shock increases the probability that real GDP will not grow over the next year by over *10 percent*, highlighting the influence that financial conditions can have on the future evolution of real GDP.

4.5 Additional Robustness Checks

A concern readers might have is that the presence of the recent financial crisis and the Great Recession in the sample might bias our findings. This concern is not warranted as we run the first stage quantile regression always on the same sample that the shock series in the second stage are available for. This means that the only results that could be influenced by these events are the JF-TFP and credit spread shock results (as the end of their sample extends beyond 2007:Q3). In Appendix E (Figure 14) we show that truncating the sample to before the financial crisis for those shocks still gives estimates that point to the same asymmetry as our benchmark results.

Another possible concern is that by focusing on extreme quantiles in a quarterly specification we are actually basing our local projections on estimates of quantiles that are only imprecisely estimated. To address this issue, we show in Appendix E (Figure 15 and Figure 16) that (i) our findings of asymmetric impulse responses are robust to using less extreme quantiles (the 25th and 75th quantiles, respectively), and (ii) that for those shock series that are available at a monthly frequency we find the same type of asymmetry (replacing real GDP growth with the growth rate of real industrial production).

5 Inspecting the Economic Mechanism

Through which channel do macroeconomic shocks affect the conditional distribution of GDP growth? To answer this question we carry out two exercises. First, we re-estimate our impulse responses of various quantiles of 4-quarter-ahead real GDP growth, but drop the financial conditions index from the right hand side of our quantile regressions. Thus, only past real GDP growth is used to predict future GDP growth in the first-step quantile regressions. Strikingly, as depicted in Figure 4, removing information about financial conditions

from the quantile regressions removes *all* asymmetric behavior of the various quantiles in response to the various shocks we consider. This result suggests that of our two conditioning variables, namely financial conditions and current GDP growth, it is through the former channel that shocks affect macroeconomic tail risk. Our finding is in line with [Adrian *et al.* \(2019\)](#), who point out that including the NFCI as a conditioning variable is key to capture downside risk. Furthermore, [Adrian *et al.* \(2019\)](#) point out that incorporating the NFCI series helps to forecast future real GDP growth. Thus the specification that we borrow from [Adrian *et al.* \(2019\)](#) is one that is useful in understanding the evolution of real GDP growth and, in particular, more useful than a specification that only conditions on past real GDP growth (where, as we show in this section, asymmetry is not present). [Adrian *et al.* \(2019\)](#) also discuss various equilibrium models in the literature that help explain the central role of financial conditions in shaping future real GDP growth.




Figure 4: Impulse Responses of Various Quantiles to *Contractionary* Shocks, Conditioning Only on Past GDP Growth.

Note: Red (dashed) is response of the 10th quantile, black (solid) is the median response, blue (dotted) is response of the 90th quantile. Confidence bands correspond to median response, 90% significance level, based on Newey-West standard errors.

Second, we can study how the the National Financial Conditions Index (NFCI) used to construct the quantiles of GDP growth in expression (2.1) responds to the shocks studied in this paper. In particular, in Appendix E (Figure 18) we report the impulse responses of that variable itself to the shocks we consider. Positive values indicate that financial conditions are tighter, while negative values indicate financial conditions that are looser. As expected, contractionary shocks make financial conditions tighter. The impulse responses of the 10th quantile of the conditional GDP growth distribution in Figure 2 inherit the (inverse) pattern of the response of financial conditions.¹⁸ The impulse responses of the probability that real GDP will not grow over the next year (presented in Figure 3) are similar in shape to the responses depicted in in the Appendix E in Figure 18, which we see as further evidence for the big role of financial conditions in shaping output.

We can thus conclude that contractionary shocks temporarily increase macroeconomic tail risk by tightening financial conditions.

To dig deeper into the role of financial conditions in shaping real GDP growth, we next turn to a different approach. We ask how the quantiles of real GDP growth are influenced by our structural shocks of interest *only via financial conditions*. To do so, we first instrument for current financial conditions (NFCI, to be exact) with our structural shock of interest and then estimate a quantile regression that includes this instrumented financial condition index on the right-hand side. More formally, we consider the first-stage OLS regression

$$NFCI_t = \sum_{j=0}^{20} \theta^j shock_{t-j} + \sum_{j=1}^4 NFCI_{t-j} + \sum_{j=1}^4 GDPgrowth_{t-j} + u_t. \quad (5.1)$$

We then compute the fitted or instrumented value

$$NFCI_t^{inst} = \sum_{j=0}^{20} \hat{\theta}^j shock_{t-j}. \quad (5.2)$$

In our second stage, we include $NFCI_t^{inst}$ in a quantile regression that also includes date t real GDP growth on the right-hand side:

$$Q_{y_{t+s}|x_t, \epsilon_t^m}(\tau) = \beta^s(\tau) NFCI_t^{inst} + \psi GDPgrowth_t + e_t. \quad (5.3)$$

This ensures that the shock can only operate through financial conditions directly.

¹⁸In Appendix E, we also show the corresponding figure for the other conditioning variable in the quantile regressions, GDP growth. There are substantially more pronounced differences in the responses of that variable to the shock relative to how the 10th percentile of the GDP growth forecast distribution reacts to shocks.

Our object of interest are the coefficients $\beta^s(\tau) * \theta^0$.¹⁹ To compute error bands, we follow [Linnemann and Winkler \(2016\)](#) and use a "blocks-of-blocks" bootstrap with 10,000 replications.²⁰ In Figure 5, we plot the responses to contractionary shocks of the 90th percentile and the 10th percentile as well as the associated 68 percent error bands. Our main message is present here as well: We see a substantially larger (in magnitude) response of the 10th percentile relative to the 90th percentile for all contractionary shocks, which is driven by the response of financial conditions to economic shocks.




Figure 5: **Impulse Responses to Contractionary Shocks, Financial Conditions Channel.**

Note: 68% significance level, standard errors computed with 10,000 "blocks-of-blocks" bootstrap replications.

Another feature of our results that stands out is that upside risk reacts substantially less to economic shocks than downside risk, as is evident from the different behavior of the 90th percentile and the 10th percentile in Figure 2. This is consistent with the finding in [Adrian et al. \(2019\)](#) that upside risk moves substantially less over time relative to downside risk. We also confirm this finding by directly computing the forecast error variance decomposition of

¹⁹Note how in this specification we do not condition on a specific estimate of a quantile to get estimates of asymmetry. As we use a bootstrap, this specification takes into account sampling uncertainty of the quantiles.

²⁰For a discussion of this type of bootstrap, see Chapter 12 of [Kilian and Lütkepohl \(2018\)](#).

[Gorodnichenko and Lee \(2019\)](#) for our specification, which we show in Appendix [F](#).

6 Conclusion

The impact of macroeconomic shocks on economic activity has been studied extensively. Their effect on lower quantiles—commonly referred to as “tail risk”—has been studied substantially less, even though it is of utmost importance to policymakers. This paper fills this gap by focusing on how macroeconomic shocks affect both tail risk and the entire distribution of future GDP growth. All shocks considered (monetary policy, credit spread, and productivity shocks) affect tail risk disproportionately more than other quantiles. Contractionary shocks thus deserve even more attention than what their effect on average outcomes suggests to the extent that they make poor economic conditions much more likely. As this is also true of monetary policy shocks, there is reason to be especially wary of contractionary policy shocks.

We find that financial conditions are the key channel through which shocks affect macroeconomic risk, complementing the findings in [Adrian *et al.* \(2019\)](#). This suggests that research on the financial transmission of shocks is key to understand economic growth and its vulnerability.

Finally, our findings can be important for the pricing of assets because asset pricing theory tells us that tail behavior in outcomes (which in turn influence payoffs of assets) has a substantial influence on the price of assets. Our results shed light on how economic shocks affect this tail behavior.

References

- ADRIAN, T., BOYARCHENKO, N. and GIANNONE, D. (2019). Vulnerable Growth. *American Economic Review*, **forthcoming**.
- ANTOLÍN-DÍAZ, J. and RUBIO-RAMÍREZ, J. F. (2018). Narrative Sign Restrictions for SVARs. *American Economic Review*, **108** (10), 2802–29.
- AZZALINI, A. and CAPITANIO, A. (2003). Distributions Generated by Perturbation of Symmetry with Emphasis on a Multivariate Skew t -distribution. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, **65**, 367–389.
- BARNICHON, R. and BROWNLEES, C. (2016). Impulse Response Estimation By Smooth Local Projections. (11726).
- CARRIERO, A., CLARK, T. E. and MARCELLINO, M. (2019). *Capturing Macroeconomic Tail Risks with Bayesian Vector Autoregressions*. Tech. rep.
- CHAVLEISHVILI, S. and MANGANELLI, S. (2017). Quantile Impulse Response Functions.
- COGLEY, T. and SARGENT, T. J. (2005). Drift and Volatilities: Monetary Policies and Outcomes in the Post WWII U.S. *Review of Economic Dynamics*, **8** (2), 262–302.
- COIBION, O. and GORODNICHENKO, Y. (2012). What Can Survey Forecasts Tell Us About Information Rigidities? *Journal of Political Economy*, **120** (1), 116–159.
- CÚRDIA, V., DEL NEGRO, M. and GREENWALD, D. L. (2014). Rare Shocks, Great Recessions. *Journal of Applied Econometrics*, **29** (7), 1031–1052.
- FAGIOLO, G., NAPOLETANO, M. and ROVENTINI, A. (2008). Are Output Growth-Rate Distributions Fat-Tailed? Some Evidence from OECD Countries. *Journal of Applied Econometrics*, **23** (5), 639–669.
- FAVARA, G., GILCHRIST, S., LEWIS, K. F. and ZAKRAJŠEK, E. (2016). Updating the Recession Risk and the Excess Bond Premium. (2016-10-06).
- FERNALD, J. G. (2012). A Quarterly, Utilization-Adjusted Series on Total Factor Productivity.
- GALÍ, J. (1999). Technology, Employment, and the Business Cycle: Do Technology Shocks Explain Aggregate Fluctuations? *American Economic Review*, **89** (1), 249–271.
- GILCHRIST, S. and ZAKRAJŠEK, E. (2012). Credit Spreads and Business Cycle Fluctuations. *American Economic Review*, **102** (4), 1692–1720.

- GORODNICHENKO, Y. and LEE, B. (2019). Forecast Error Variance Decompositions with Local Projections. *Journal of Business & Economic Statistics*, **forthcoming**.
- JORDÀ, O. (2005). Estimation and Inference of Impulse Responses by Local Projections. *American Economic Review*, **95** (1), 161–182.
- JUSTINIANO, A. and PRIMICERI, G. E. (2008). The Time-Varying Volatility of Macroeconomic Fluctuations. *American Economic Review*, **98** (3), 604–641.
- KILIAN, L. and LÜTKEPOHL, H. (2018). *Structural Vector Autoregressive Analysis*. No. 9781107196575 in Cambridge Books, Cambridge University Press.
- and MANGANELLI, S. (2008). The Central Banker as a Risk Manager: Estimating the Federal Reserve’s Preferences under Greenspan. *Journal of Money, Credit and Banking*, **40** (6), 1103–1129.
- KIM, T.-H., LEE, D. J. and MIZEN, P. (2019). Impulse Response Analysis in Conditional Quantile Models and an Application to Monetary Policy.
- KOENKER, R. and BASSETT, G. (1978). Regression Quantiles. *Econometrica*, **46** (1), 33–50.
- LINNEMANN, L. and WINKLER, R. (2016). Estimating Nonlinear Effects of Fiscal Policy Using Quantile Regression Methods. *Oxford Economic Papers*, **68** (4), 1120–1145.
- MUMTAZ, H. and SURICO, P. (2015). The Transmission Mechanism In Good And Bad Times. *International Economic Review*, **56**, 1237–1260.
- MURPHY, K. M. and TOPEL, R. H. (1985). Estimation and Inference in Two-Step Econometric Models. *Journal of Business & Economic Statistics*, **3** (4), 370–379.
- PLAGBORG-MØLLER, M. and WOLF, C. (2019). Local Projections and VARs Estimate the Same Impulse Responses.
- PRIMICERI, G. E. (2005). Time Varying Structural Vector Autoregressions and Monetary Policy. *Review of Economic Studies*, **72** (3), 821–852.
- RAMEY, V. (2016). Macroeconomic Shocks and Their Propagation. *Handbook of Macroeconomics*, **2**, 71 – 162.
- ROMER, C. D. and ROMER, D. H. (2004). A New Measure of Monetary Shocks: Derivation and Implications. *American Economic Review*, **94** (4), 1055–1084.
- UHLIG, H. (2005). What Are the Effects of Monetary Policy on Output? Results from an Agnostic Identification Procedure. *Journal of Monetary Economics*, **52** (2), 381–419.

A Data

This section gives a brief overview of the data we use throughout this paper, which are mostly available at FRED.

To estimate the quantile regression we use seasonally adjusted and annualized real GDP growth as well as the Chicago FED National Financial Conditions Index (NFCI). This index is not seasonally adjusted and downloaded at quarterly frequency by relying on the average of weekly values within a quarter.

The control variables in the local projection stage are given as follows. At quarterly frequency we take seasonal adjusted real GDP, the seasonal adjusted civilian unemployment rate, and total hours worked given by the hours of wage and salary workers on non-farm payrolls. The latter is used to compute productivity (real GDP divided by hours), respectively. Both the commodity price index and the consumer price index are available at monthly frequency. We take the CRB commodity index provided by [Ramey \(2016\)](#) and headline CPI (defined in FRED as “Consumer Price Index for all Urban Consumers: All Items”). Additionally, we take the monthly federal funds rate. All monthly series are aggregated to quarterly frequency by taking the quarterly average.

Finally, we utilize the following aggregate shocks. The [Romer and Romer \(2004\)](#) monetary shock is provided by [Ramey \(2016\)](#). We aggregate the monthly shock series to quarterly frequency by taking the quarterly sum. We take the narrative monetary policy shock provided by [Antolín-Díaz and Rubio-Ramírez \(2018\)](#), again aggregated to quarterly frequency by calculating the quarterly sum. To identify the credit shock we use the [Gilchrist and Zakrajšek \(2012\)](#) excess bond premium, frequently updated by [Favara *et al.* \(2016\)](#)²¹. The two technology shocks are identified by running the VARs described in Section (2).

²¹The series can be downloaded at <https://www.federalreserve.gov/econresdata/notes/feds-notes/2016/updating-the-recession-risk-and-the-excess-bond-premium-20161006.html>.




Figure 6: **Quantiles of One-Year-Ahead Real GDP Growth Distribution used in Local Projections.**




Figure 7: **Conditioning Variables in the Quantile Regressions.**

B Some More Intuition

In this section we show an example where the impulse responses of the quantiles cross. By construction, however, the quantiles themselves cannot cross because we directly model changes in the entire distribution. The example is similar to the first example from the main text, but we change the sequence of variances for the Gaussian distribution to achieve the crossing of the impulse responses of the 10th and 90th percentiles.




Figure 8: Example Where Impulse Responses Cross.

C Robustness at Other Horizons

In our benchmark analysis, we study the quantiles of the distribution of GDP growth over the next year. Instead, we can look at the distribution of quarterly GDP growth various quarters ahead. In this appendix, we show results for quarterly annualized real GDP growth 2 quarters ahead. The patterns are qualitatively very similar to our benchmark analysis. Furthermore, results for 1, 3, or 4 quarters ahead real GDP growth also show the same pattern.




Figure 9: Impulse Responses of Various Quantiles to Contractionary Shocks, Real GDP Growth 2 Quarters Ahead.

Note: Red (dashed) is response of the 10th quantile, black (solid) is the median response, blue (dotted) is response of the 90th quantile. Confidence bands correspond to median response, 90% significance level, based on Newey-West standard errors.




Figure 10: Impulse Responses of Various Quantiles to Contractionary Shocks, Real GDP Growth 2 Quarters Ahead.

Note: Red (dashed) is response of the 10th quantile, blue (dotted) is response of the 90th quantile. Confidence bands correspond to 10th quantile response, 90% significance level, based on Newey-West standard errors.

D The Average Distribution of 4-Quarter Ahead Average Real GDP Growth and Estimated Changes in the Distribution of Real GDP Growth

In order to get a sense how structural shocks affect the entire distribution of future GDP growth, we first compute the average distribution of four-quarters ahead real GDP growth in our total sample and then plot the difference between this initial distribution and the distribution affected by a specific shock at various horizons. In order to facilitate interpretation, each panel of Figure 12 plots three lines: the 10th percentile (in red), the median (in black), and the 90th percentile (in blue) of the original (average) distribution. This helps check in what direction a shock shifts the distribution. In particular, whenever a line is visible it means that posterior mass at that quantile of the original distribution has decreased.




Figure 11: The Average Distribution of 4-Quarter Ahead Average Real GDP Growth (Used to Initialize Distributions for Figure 12).

Note: At each horizon h , the impulse response captures the difference between the average distribution of future GDP growth and the distribution of the shocked quantiles at horizon h .




Figure 12: Difference in Fitted t -Distributions, Contractionary Shocks.

Note: Straight lines are 10th percentile (red), median (black), and 90th percentile (blue) of the average distribution of 4-quarters ahead real GDP growth in our sample.

E Additional Impulse Responses

In this section we show various additional impulse response figures, including one that displays the error bands for the 10th and 90th percentile responses that were not presented in the main text.




Figure 13: Impulse Responses of Various Quantiles to Contractionary Shocks.

Note: Red (dashed) is response of the 10th quantile, blue (dotted) is response of the 90th quantile. Confidence bands correspond to 10th quantile response, 90% significance level, based on Newey-West standard errors.




Figure 14: Impulse Responses (to Contractionary Shocks), Samples Ending in 2007:Q3.




Figure 15: Impulse Responses of Quantiles Used to Fit t -Distributions, Contractionary Shocks.




Figure 16: Impulse Responses (to Contractionary Shocks), Monthly Frequency.

Note: We use six lags of all controls. JF-TFP shocks are only available at a quarterly frequency - we reproduce our benchmark results for that shock here for the sake of comparison.




Figure 17: Impulse Responses (to Contractionary Shocks) of average real GDP growth over the next year y_{t+h} .




Figure 18: Impulse Responses (to Contractionary Shocks) of the Chicago FED National Financial Conditions Index.

F Forecast Error Variance Decomposition

In this section we briefly describe the forecast error variance decomposition (FEVD) based on local projections following [Gorodnichenko and Lee \(2019\)](#). For ease of notation let us define $shock_t \equiv z_t$ in equation 2.3. While we use the same set of control variables for the FEVD which we also use to construct the IRFs, we omit the vector of controls in this description for brevity.

Let $q_{\tau,t}$ be a given quantile of interest and z_t a series of shocks with mean zero and variance σ_z^2 . The variation in $q_{\tau,t}$ due to the series of shocks z_t can be described by the lag polynomial $\psi_z(L)z_t = \sum_{i=0}^{\infty} \psi_{z,i}z_{t-i}$, where the impulse response function is given by the series of coefficients $\{\psi_{z,i}\}$. The s -step ahead forecast error for $q_{\tau,t}$ can then be defined as

$$f_{t+s|t-1} \equiv q_{\tau,t+s} - P[q_{\tau,t+s} | \Omega_{t-1}], \quad (F.1)$$

where $P[\dots]$ is the projection of $q_{\tau,t}$ on the information set available at time $t-1$, defined as $\Omega_{t-1} \equiv \{q_{\tau,t-1}, z_{t-1}, q_{\tau,t-2}, z_{t-2}, \dots\}$. Now the forecast error can be decomposed into variations due to z_t and other shocks v_t , thus

$$f_{t+s|t-1} = \psi_{z,0}z_{t+s} + \dots + \psi_{z,s}z_t + v_{t+s|t-1}. \quad (F.2)$$

Finally, the share of the total variance of $f_{t+s|t-1}$ due to the series of shocks z_t to the total variance in the forecast error can be described as

$$s_s = \frac{\text{Var}(\psi_{z,0}z_{t+s} + \dots + \psi_{z,s}z_t)}{\text{Var}(f_{t+s|t-1})}. \quad (F.3)$$

[Gorodnichenko and Lee \(2019\)](#) show that with $Z_t^s = (z_{t+s}, \dots, z_t)'$ the expression for s_s can be written as

$$s_s = \frac{\mathbf{Cov}(f_{t+s|t-1}, Z_t^s) [\mathbf{Var}(Z_t^s)]^{-1} \mathbf{Cov}(Z_t^s, f_{t+s|t-1})}{\text{Var}(f_{t+s|t-1})}, \quad (F.4)$$

which can be understood as an R^2 of the population projection of $f_{t+s|t-1}$ on Z_t^s given in equation F.2. Therefore, [Gorodnichenko and Lee \(2019\)](#) suggest estimating the FEVD in two steps. First, estimate the forecast errors for the horizon s from local projections and second, regress these forecast errors on shocks that happen between t and $t+s$. The R^2 from the

latter regression then gives an estimate of s_s .²²

We follow these steps and estimate FEVDs for both the 10th percentile and the 90th percentile for all shocks under study. Note that we estimate one FEVD per shock, i.e., the shocks do not enter the estimation jointly. Hence, the FEVDs are not forced to add up to one, which is important as we use different shocks of the the same type, e.g., two different monetary policy shocks. Figure 19 shows the results. The first column shows the FEVDs for the 10th percentile, the second column plots FEVDs for the 90th percentile. In the first row we plot the contributions of the monetary policy shocks, the second row plots results for the technology shocks, and the third row depicts results for the credit shock. Up to a forecast horizon of $s = 12$ the monetary policy shocks explain a roughly constant share of the variance in the 10th percentile of about 15 percent. While the contribution of the AR monetary shock stays constant for longer horizons, the RR shock becomes more important and explains more than 40 percent of the variance in the 10th percentile at $s > 15$. For the 90th percentile the contribution of the two monetary shocks is similar. It is increasing in the horizon s , starting from about five percent to roughly 30 percent at the end of the forecast horizon. Turning to the technology shocks reveals a similar result. Especially the Gali shock explains much more of the variation in the left tail (up to 40 percent) than to the variation in the right tail (25 percent at its peak). For the other shock differences are not clear-cut. The credit shock contributes to between 20 and 30 percent of the variance in the 10th percentile and between 10 and 20 percent of the variance in the 90th percentile. Thus, the credit shock clearly contributes more to variation in the left tail than in the right tail of our distribution.

Overall, the FEVD underpins that most of the shocks under study are more important for movements in the left tail of the distribution than for movements in the right tail.

²²Two things are worth mentioning. First, [Gorodnichenko and Lee \(2019\)](#) show that if there is measurement error in z , the estimated s_s is downward biased and thus a conservative estimate of the contribution of shock z to the total variation in q_τ . Moreover, while the estimator for s_s is asymptotically unbiased without measurement error, there might be small sample biases. Therefore, [Gorodnichenko and Lee \(2019\)](#) recommend to use a small sample correction based on either bootstrap methods or by simulating asymptotic distributions. However, these corrections can lead to negative estimates for s_s , which are hard to interpret. Thus, we decide to report the potentially biased estimates of s_s , but results are qualitatively similar when using the bootstrap correction.




Figure 19: Forecast Error Variance Decomposition

Note: Plots show Forecast Error Variance Decompositions based on Local Projections following [Gorodnichenko and Lee \(2019\)](#).