Macro Risk

Tobias Adrian, Nina Boyarchenko, Francesco Furno
Domenico Giannone, Leonardo Iania, Michele Lenza
Francesca Loria, Cecilia Melo Fernandes, Sergio Sola

Preliminary and incomplete draft

Do not circulate or quote

March 1, 2025

Abstract

We guide the reader through key statistical techniques for monitoring and forecasting macroeconomic risk. Moving beyond standard linear point forecasts, we demonstrate how to construct flexible conditional distributions of future GDP growth. We show that several methods can be leveraged to achieve this goal: quantile regression, Markov switching models, and large-scale techniques. Since it captures the likelihood of all possible future outcomes, the conditional distribution of future GDP growth serves as an ideal tool for assessing macroeconomic vulnerabilities. The insights presented in this paper have implications for policymakers, practitioners, and academics.

Keywords: Risk measures, quantile regressions, Markov-switching models, machine learning.

JEL Codes: C53, E23, E27, E32, E44.

1 Introduction

Economic uncertainty is pervasive and influences decision-makers at all levels from policymakers striving to maintain financial stability, see for example International Monetary Fund (2024), to businesses and households making long-term plans, see for example Bernanke (1983). For this reason, we need tools that help to assess the likelihood of different potential paths. Specifically, a key challenge in macroeconomics is not only to forecast aggregate economic growth but also to assess the risks surrounding such projections, balancing risk in a timely and rigorous manner. Our paper provides a comprehensive statistical framework for this.

Traditional forecasting approaches often focus on point estimates or central tendency measures. Standard linear models, for instance, predict average economic conditions based on a set of (observed) conditional variables. And uncertainty around these predictions are often summarized by symmetric bands, governed by a fixed variance. These types of framework overlook the dynamic nature of economic uncertainty, failing to capture periods of heightened volatility or asymmetric risks that characterize financial crises and economic downturns, see Scotti (2023) and Boyarchenko(forthcoming).

As the GDP is a summary variable to assess macroeconomic vulnerabities, we focus our analysis on one-quarter-ahead GDP growth predictions. However, it is important to note that the principles outlined in this work can be applied to any forecasting problem where interest lies in characterizing its entire distribution—mean, median, quantiles, and tails. Section 10 provides additional applications of this methodology to macroeconomic variables beyond GDP growth.

Our approach throughout the paper is straightforward: we analyze financial sector conditions today to forecast GDP growth over the next quarter(s). Our analysis starts by assuming a linear relationship between these variables, see section 3.1. This allows us to clarify the intuition behind our logic of moving away from point forecast. In later sections, we relax this assumption and consider alternative, non-linear models.

To motivate our framework, consider the relationship between GDP growth (in blue) and financial conditions (in red), as depicted in Figure 1.

Figure 1: Financial conditions and GDP growth

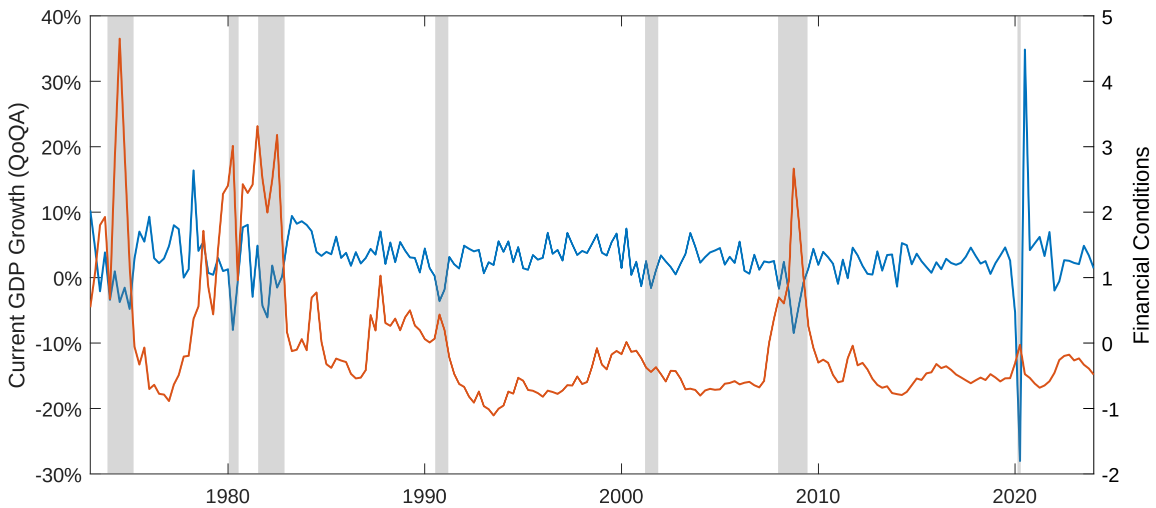

Note: The figure shows the time series evolution of GDP growth and the financial conditions (NFCI).

Financial conditions are represented by the National Financial Conditions Index (NFCI) compiled by the Federal Reserve Bank of Chicago. The NFCI aggregates various financial indicators, providing a comprehensive measure of the financial sector's overall health. An increase in this index signals tighter financial conditions. A visual inspection of the plots reveals a key insight: during "normal times"—broadly defined as periods outside the shaded recession zones—the correlation between the two variables is weak. However, as recessions approach, this correlation often strengthens, with financial conditions serving as an early indicator of negative and more volatile growth. Since macroeconomic risk is primarily related to downturns and heightened volatility, ignoring these dynamics might lead to a systematic understatement of risk. This is the central premise of our work. Models ignoring these non-linearities can provide a misleading representation of risk. We advocate for moving beyond these methods toward a framework that captures the full distribution of GDP growth, allowing for a more comprehensive and rigorous assessment of economic risks.

Figure 2: Scatter plots, OLS and quantile regression

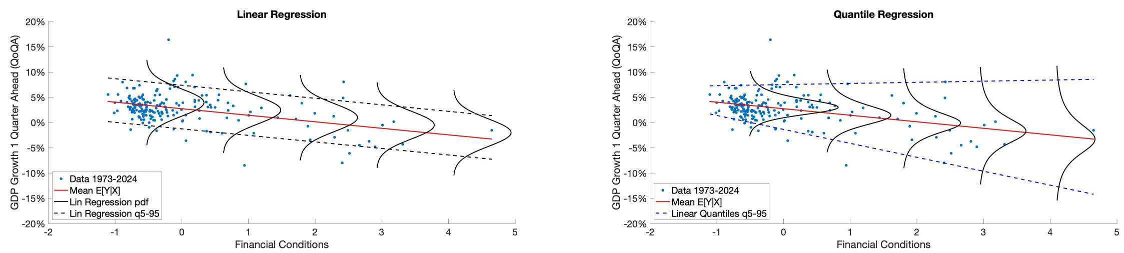

Note: The two graphs report the scatter plots of GPD growth and financial conditions with the OLS slopes (left plot) and the quantile slopes (right plot). Skew-t distributions are superimposed to the data.

To illustrate this, we turn to scatterplots of GDP growth and financial conditions, see Figure (2). In these plots, future GDP growth is positioned on the y-axis, while financial conditions serve as the predictor on the x-axis. In applications where we want to use current financial condition to forecast the average, expected GDP growth, White (1980) show us that the Ordinary Least Squares (OLS, left chart) downward-sloping relationship (red line) is valid even in the presence of heteroskedasticity. Multiplying this slope factor by the current financial conditions and adding a constant results in the time series of expected GDP growth, depicted by the black solid line of the left panel of Figure 3.

When forecasting economic downturns, we must move beyond point forecasts and consider the full predictive distribution of our model. A probabilistic approach is crucial to assess, for instance, the likelihood that, given current financial conditions, GDP growth will turn negative. This requires analyzing the entire distribution of possible outcomes rather than just expected values. A key shortcoming of standard OLS models emerges in this context.

The bell-shaped curves in the left chart of Figure 2 represent the conditional probability densities of OLS forecasts, obtained by fitting to the model's residual a normal distribution. As financial conditions change, only the mean of the distribution shifts, while its shape remains unchanged. This limitation prevents OLS models from capturing the non-linearity in the relationship between financial conditions and GDP growth, particularly in downturns.

To illustrate this issue, we report the 5th percentile of the OLS predictive distribution (lower dashed lines in the left panel of Figure 2), which separates the worst 5% of predicted

GDP growth outcomes from the rest. The 5th percentile—similar to Value at Risk in finance—serves as a measure of downside risk, commonly referred to in macroeconomic forecasting as Growth at Risk (GaR) Adrian et al. (2019). Under OLS assumptions, this percentile remains parallel to the central forecast, implying that the deviation of GDP growth from its expected value is independent of financial conditions. However, Figure 1 contradicts this, showing that extreme financial tightening is associated with disproportionately large GDP declines.

This underlies the inadequacy of linear (regression) models in capturing downside risks during financial distress. Quantile regression models, by contrast, overcome these limitations. Unlike OLS, they allow for differential slopes across quantiles, sections of the distribution of future GDP growth to respond differently to financial conditions. The principle of (linear) quantile regression is straightforward and it is easily explained in a univariate setting. Like any linear model, it optimally summarizes a scatter of observed data with a straight line. However, in quantile regression, "optimal" means minimizing a weighted average of absolute errors, where the weights depend on the chosen quantile level (see Section 2.1.2). Hence, this setting provides a more nuanced view of the relationship between financial conditions and GDP growth, particularly in capturing asymmetric risks and tail events.

To illustrate this point, the right panel of Figure 2 reports the quantile regression slopes for our dataset. Unlike in the OLS case, the 5th and 95th quantile regression lines vary significantly across quantiles (dashed lines). While the median relationship between financial conditions and GDP growth is relatively weak, the lower quantiles exhibit much steeper negative slopes. In other words, as financial conditions tighten, the upper quantile remains stable, the median response is muted, but the lower quantile declines sharply, indicating a higher probability of severe downturns. This indicates that financial conditions are particularly relevant for downside risks. To further assess the implications for risk evaluation, we overlay the predictive distributions of GDP growth under the quantile regression framework. This reveals a skewed distribution under tighter financial conditions, with a greater probability mass assigned to negative outcomes, underscoring the increased downside risk.

These findings are further corroborated by plotting the estimated quantiles over time (Figure 3).

Figure 3: Time series evolution of OLS and regression-based quantiles.

Note: The two graphs report the time series of OLS-based (left plot) and quantile regression-based (right plot) quantiles.

The figure displays the estimated median of predicted GDP growth (black solid line), along with the upper and lower quartiles (darker shaded areas), the 5^{th} and 95^{th} percentiles (lighter areas), and realized GDP growth (blue line). Under OLS, the shaded bands shift in parallel with the median. However, under quantile regression, a markedly different pattern emerges: the upper quantile remains relatively stable, while lower quantiles exhibit substantial variation, reinforcing the asymmetric nature of economic risk. Traditional least squares models, by focusing solely on mean outcomes, systematically underestimate recession probabilities during periods of financial stress and overestimate them when financial conditions are loose.

If we plot these two figures together with the entire probability distribution, these key takeaways become even more evident (Figure 4).

Figure 4: 3-D predictive distributions

Note: The two graphs report the 3-D plots of the evolution of Gaussian-based (left plot) and Skew-t (right plot) predictive distributions.

Under OLS (Figure 4, left chart), changes in financial conditions shift the location of the distribution, but its variance remains unchanged, preserving its Gaussian shape. As such, the probability mass around the mean will be always the same over time independently of financial conditions, as is the estimated probability of growth falling—say—one or two standard deviations below the mean. Contrast this with the quantile regression model (Figure 4, right chart). Over time, both the location and the shape of GDP growth change depending on financial conditions. When growth is high, and financial conditions are loose the distribution is concentrated around its mean. Yet, when growth is lower and financial conditions tighten, the distribution moves down, and variance increases but with a skew affecting the left tail more.

This is clearly visible if we compare the predicted distribution of GDP growth during the Great Recession (2008) and a more stable period (2014) as in Figure 5.

Figure 5: Predictive distributions in selected time periods

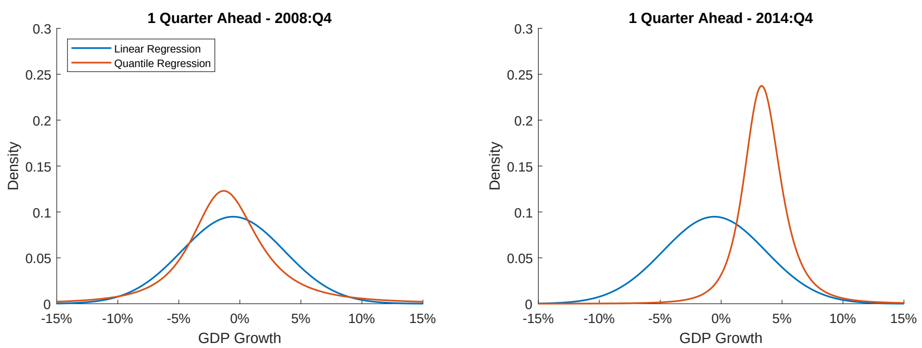

Notes: Normal and Skew-t based distributions in the 2008Q4 (left plot) and in 2016Q4.

The blue curves represent OLS-based predictions, while the red curves depict predictions based on quantile regression. During the recession, quantile regression assigns greater probability mass to negative outcomes than OLS, indicating higher downside risk. Conversely, in 2014, OLS predicts a symmetric shift in GDP growth, while quantile regression yields a more concentrated distribution around positive values, reflecting reduced downside risk.

These findings highlight the limitations of standard linear regression in assessing macroeconomic risk and emphasize the necessity of modeling the full predictive distribution. By adopting a distributional forecasting approach, we provide a more comprehensive framework for macroeconomic risk assessment, with significant implications for policymakers. Central banks and policymakers must consider not only expected growth but also tail risks when making decisions, ensuring a more robust response to financial instability.

In Section 2, we provide the general definition. Section 3 focuses on quantiles, where we demonstrate how to estimate the lines, perform inference on the slopes, construct Growth-at-Risk, and assess its accuracy. In Section 4, we extend the analysis to density, showing how to derive the entire density from estimated quantiles, conduct inference, refine risk measurement, and evaluate the results. Finally, in Section 5, we explore alternative methods.

2 Defining and Motivating Macroeconomic Risk

In our main application, we focus on forecasting U.S. real GDP growth one quarter ahead, following Adrian et al. (2019). With quarterly data, the annualized GDP growth rate is defined as:

$$Y_{t+1} = \left(\frac{W_{t+1}}{W_t} \right)^4 - 1 \tag{1}$$

where W_t represents the level of real GDP, and Y_t denotes the current GDP growth rate. Our forecasts will be conditional on X_t, a measure of current financial conditions, such as the National Financial Conditions Index (NFCI) compiled by the Federal Reserve Bank of Chicago.

In this section we will discuss the limitations of relying on point forecasts and standard OLS models when assessing economic risk, while in Subsection 2.1 we will highlight the importance of considering the entire predictive distribution of GDP growth rather than just its expected value, and review key concepts related to conditional distributions and quantiles, which are central to our analysis.

To motivate our framework, consider the relationship between GDP growth (in blue) and financial conditions (in red), as depicted in Figure 1.

A key observation from the graph is that outside the shaded recession periods ("normal times") the correlation between GDP growth and financial conditions appears weak. However, as recessions approach, this correlation strengthens. Financial conditions often serve as an early indicator of worsening (and more volatile) economic conditions. Precisely, when data display such-non linearities, they will influence not only the mean, but the *entire distribution* of our target variable: the effect will be not only on first moments (i.e. the mean, or location of the distribution) but possibly on all the higher moments (i.e. variance and behavior of the tails). But precisely because most linear predictive models assume that the conditioning variable impact only the central location of the forecast, they tend to understate risk.

Linear predictive models, such as those based on Ordinary Least Squares (OLS), are well-suited for point estimation—in our case, providing an estimate of expected GDP growth conditional on financial conditions. Formally, the forecast in an OLS model can be summarized as:

$$\mathbb{E}[Y_{t+1}|X_t] = \beta_0 + \beta_1 X_t. \tag{2}$$

Let's look at the elements of Equation 2. The term $\mathbb{E}[Y_{t+1}|X_t]$ is the prediction based on the OLS model. It corresponds to the mean of the forecasting distribution. It is often referred to as the "point forecast". The coefficient β_0 is the intercept. In the absence of an explanatory variable, it represents the average value of the target variable, Y_{t+1}. The coefficient β_1 captures how expected future GDP growth is related (*on average*) to the current level of financial conditions. Equation 2 implies that changes in financial conditions lead only to shifts in the mean prediction, without affecting the overall shape of the forecasting distribution. The strength of the model is its simplicity. But also its weakness.

To assess (macro) risk, it is essential to make assumptions on the distribution of the forecasted variable. In the OLS framework, a common approach is to impose distributional assumptions on the error term in the underlying linear model that generates the forecast, as specified in Equation 2. This relationship is given by:

$$Y_t = \beta_0 + \beta_1 X_{t-1} + \varepsilon_t. \tag{3}$$

In our setting, Equation 3 implies that using financial conditions data available up to time $t-1$, we obtain sample estimates of the parameters β_0 and β_1. The most recent observed data point, X_t, then determines our forecast of real GDP growth for the next quarter, as described in Equation 2. The assumed distribution of ε_t plays a crucial role in determining the uncertainty surrounding the point forecast.

Standard OLS inference typically assumes that the error term ε_t is homoskedastic (i.e., has a constant variance across observations) and/or follows a Gaussian distribution. However, in economic settings, these assumptions are frequently violated—especially during downturns, when volatility tends to rise. While such violations do not undermine the validity of inference—White (1980) showed that OLS coefficient estimates remain unbiased and consistent even under heteroskedasticity and autocorrelation—they will yield systematically wrong forecasts of risks.

The core issue is that OLS models the relationship between predictors and the conditional mean of the forecasted dependent variable but treats the rest of the distribution as unconditional, i.e. independent of the predictors. In our example, OLS links financial conditions to the expected value of GDP growth - through the coefficient β_1- but it is mute about the change in likelihood of extreme outcomes. This is in contrast with the evidence shown in Figure 1, where tightening financial conditions signal a vulnerable economic outlook, with an increase in the probability of severe downturns.

Standard OLS models fail to capture this dynamic because they do not allow the shape of the (forecasting) distribution of the predicted variable, Y_{t+1}, to vary with the level of the predictor, X_t. In our workhorse example Y_{t+1} is GDP growth and X_t are financial conditions, but the principle applies to any forecasting problem summarized by Equation 2. Proper risk assessment, therefore, requires a model that not only estimates the expected value of GDP growth but also links the entire distribution—particularly the likelihood of extreme negative realizations—to current financial conditions. Figure 2, illustrates this point clearly. In the two plots of the figure, future GDP growth is positioned on the y-axis, while financial conditions serve as the predictor on the x-axis. The downward-sloping OLS regression line (in red in the left chart) is the graphical representation of Equation 2. β_1 is its slope of the red line and β_0 is the intercept. The bell-shaped curves represent the conditional probability densities of the OLS forecasts. The conditionality implies that their location depends on the value of X_t, which also determines where their highest peak is, i.e. at the interception with the red line (corresponding to the point forecast and mean of the distribution). As financial conditions deteriorate, we move down the red line and the mean of the predictive distribution shifts, while the shape remains unchanged.

Risk measures are based on the shape of the predictive distribution. For instance, the 5% Growth at Risk (GaR) is a quantile of the distribution, separating the worst 5% predicted GDP growth outcomes from the rest. This measure provides insights into downside economic risks. In the OLS framework, the 5% GaR is simply a constant downward shift of the mean prediction, as reported by the dashed lower lines below on the left plot of Figure 2. In other words, it assumes that extreme declines in GDP growth are linked to financial conditions through the same coefficient as the mean forecast, β_1 in Equation 2. Only a fixed adjustment differentiates the two curves. However, this assumption is inconsistent with what illustrated in Figure 1. The relationship between real GDP growth and financial conditions differs depending on whether the economy is in a stable period ("normal times") or is approaching a recession. This highlights a fundamental shortcoming of OLS: it fails to link the shape of the distribution to the explanatory variable. In our example, it results in the failure to capture nonlinear dynamics of risk, particularly how downside risks intensify under worsening financial conditions.

Quantile regression models, by contrast, overcome these limitations by allowing the entire distribution of GDP growth to respond differently to financial conditions. Intuitively, in a linear quantile regression setting, the coefficients in Equation 2 change depending of which part of the distribution we are focusing on. Consider the right panel of Figure 2 to build this intuition further using our workhorse model (Section 3.1 will explore this in full details). The dashed lines represent the (quantile) regression lines at the 5th and 95th percentiles. Unlike in the OLS setting, where the percentiles are parallel, here we observe that the slopes of the quantile regression lines vary significantly across quantiles. While the median relationship between financial conditions and GDP growth is relatively weak, the lower quantiles exhibit much steeper negative slopes. As financial conditions tighten, the upper quantile remains stable, the median response is muted, but the lower quantile declines sharply. This is in line with the intuition suggested by Figure1, which indicates that financial conditions are especially relevant for downside risks, making quantile regression a more appropriate tool for forecasting extreme events.

To further assess the implications for risk evaluation, we overlay the predictive distributions of GDP growth under the quantile regression framework. The distribution is obtained by fitting a Skew-*t* distribution to the several the conditional quantiles, obtained for several (equidistant) values of financial conditions. Section 4 will discuss this procedure in detail. The forecasting distribution is skewed under tighter financial conditions, with a greater probability mass assigned to negative outcomes, underscoring the increased downside risk.

The different implications for risk assessment of an OLS-based framework and a quantile-regression framework are even more evident if we look at Figure 3, which shows the evolution over time of the estimated upper and lower quantiles of future GDP growth.

The figure presents the estimated median of predicted GDP growth (black solid line), along with the upper and lower quartiles (darker shaded areas), the 5^{th} and 95^{th} percentiles (lighter shaded areas), and the realized GDP growth (blue line). The left panel displays the results obtained using OLS, while the right panel those obtained with quantile regression.

These plots follow naturally from the coefficients of the estimated equations. At each point in time, the estimated predicted values are obtained by multiplying those coefficients (from either OLS or quantile regression) by the financial conditions—our conditioning variable—and adding the constant. In the case of OLS, however, estimated coefficient and constant are the same regardless of the value (i.e. quantiles) of GDP growth. The model estimates the mean of future GDP growth, and the predictive bands, including the 5^{th} and 95^{th} percentiles, are simply obtained by adding and subtracting a constant proportional to the estimated variance of the error term. Therefore, regardless of the value of financial conditions, the entire predictive distribution of GDP growth moves in parallel around the predicted mean, with no change in its shape. This limitation is evident in the left panel of Figure 3. When predictions are based on an OLS model, standard risk measures, such as the lowest 5^{th} percentile of the predicted distribution (i.e. the 5% Growth-at-Risk), move symmetrically with the mean and fail to reflect changing risk dynamics. The realized GDP growth (blue line) frequently falls outside the predictive bands, particularly in recessions, highlighting how OLS underestimates downside risks when financial conditions tighten.

Instead, the pattern that emerges when considering predictions from the quantile regression model is fundamentally different. Because the estimated coefficients (both constant and slope) vary depending on the quantile of the dependent variable, the entire predicted distribution of GDP growth moves with financial conditions; it shrinks, expands, becomes more or less asymmetric with thinner or fatter tails. This flexibility captures the asymmetric relationship between financial conditions and GDP growth. The right panel of Figure 3 shows that the upper quantile remains relatively stable, while the lower quantiles exhibit substantial variation, expanding notably during periods of economic downturn. This behavior reflects the empirical reality that downside risks intensify in times of financial stress, while upside risks remain largely unchanged.

The implications for macroeconomic risk assessment are profound. Standard OLS-based methods, which focus exclusively on the conditional mean, systematically underestimate recession probabilities when financial conditions deteriorate and overestimate them when financial conditions improve. In contrast, quantile regression, by explicitly modeling the entire distribution of GDP growth, adjusts predictions to changing financial conditions in a way that better reflects the underlying economic dynamics. This results in a more rigorous and informative framework for forecasting macroeconomic risks. This is very much evident if we look at the predicted 5^{th} percentile of GDP growth: it drops significantly in times of financial stress assigning greater probability mass to negative outcomes. This feature is entirely absent from the OLS model, where the predictive bands shift uniformly and fail to capture the true nature of risk.

In our quest to construct macro risk measures and navigate the turbulent waters of uncertainty and risk, estimated quantiles, and—more broadly—the entire Cumulative Distribution Function (CDF) will serve as our compass. These fundamental concepts not only underpin our understanding of distributional dynamics but also enable us to move beyond point forecasts and develop a fully probabilistic perspective on risk. In the next section, we revisit these key notions, laying the groundwork for constructing predictive PDFs. This, in turn, will allow us to derive virtually any risk measure needed for a comprehensive assessment of macroeconomic vulnerabilities.

2.1 The (conditional) distribution and quantiles

2.1.1 Definitions

Our objective is to forecast risk. As articulated by Knight (1921), risk refers to events that can be described by a known or knowable probability distribution. Hence, regardless of how we conceptualize risk, measuring it requires estimating the predictive distribution of the variable of interest, which in our case is the h-period ahead GDP growth, Y_{t+h}. In statistical terms, Y_{t+h} is a random variable that takes values across the real line.

The concept of a probability distribution is fundamental to any risk assessment, making it essential to revisit this notion. A straightforward way to summarize a probability distribution is through the Cumulative Distribution Function (CDF). The conditional CDF, $F_{Y_{t+h}|X_t}(y)$, gives the probability that Y_{t+h} takes a value below y, given the information $X_t = x$:

$$F_{Y_{t+h}|X_t}(y) = P(Y_{t+h} \leq y \mid X_t) := h(y, X_t). \tag{4}$$

The CDF fully describes the probability of all possible realizations of Y_{t+h}, conditional on X_t. The prediction is *"conditional"* because it depends on a set of variables X_t that are observed at the time of prediction and incorporated into our forecasting model. In our case, X_t represents financial conditions, meaning the CDF of future GDP growth depends on the financial environment at the time the forecast is made. If X_t consists only of a constant, the CDF is termed *"unconditional"* as it does not rely on additional variables.

Unlike standard linear regression, which links the mean of the dependent variable to explanatory variables, our approach evaluates how the entire distribution of future outcomes depends on X_t. This allows us to move beyond point forecasts and analyze the full predictive distribution, making it dependent on one or more observed variables.

The term $F_{Y_{t+h}|X_t}(y)$ is often called the predictive distribution since Y_{t+h} represents a future outcome. Various risk measures can be derived from it. These include direct computations from the CDF, such as the probability of a recession, $P(Y_{t+h} \leq 0 \mid X_t)$, or functionals of the CDF, such as the expected shortfall, which averages the worst 5% of predicted GDP values. For example, the risk of a recession over the next year—defined as a period of negative GDP growth—can be determined by forecasting the probability distribution of GDP growth over the next four quarters, $F_{Y_{t+4}|X_t}(y)$, evaluated at zero. If the CDF is absolutely continuos, we can write

$$F_{Y_{t+4}|X_t}(0) = \int_{-\infty}^{0} f_{Y_{t+4}|X_t}(y)\, dy.$$

where $f_{Y_{t+4}|X_t}(y)$ is the conditional probability density function (PDF).

A particularly useful tool derived from the CDF is the *quantile function*, which links the probability of an outcome with forecasted GDP values. The quantile function is defined as the inverse of the CDF, meaning that the CDF evaluated at the τ-quantile, $Q_{Y_{t+h}|X_t}(\tau)$, equals τ:

$$F_{Y_{t+h}|X_t}(Q_{Y_{t+h}|X_t}(\tau)) = \tau.$$

In other words, the quantile function provides the value of y such that the cumulative probability equals τ. Since τ represents a probability, it lies in the range $[0, 1]$.

The quantile function is particularly useful in risk analysis, as it helps identify thresholds for extreme events. A well-known example is the Value at Risk (VaR), which is defined as the τ-quantile for a given probability level. For instance, the 5% Value at Risk (VaR), denoted as $Q_{Y_{t+h}|X_t}(0.05)$, represents the level of GDP growth that separates the worst 5% of outcomes from the remaining 95%. In other words, if $\tau = 0.05$, there is a 5% probability that GDP growth will be lower than this threshold, while 95% of the time, GDP growth will exceed it. In a macroeconomic context, this measure is labeled this measure Grow At Risk, GaR (Adrian et al. (2019)).

However, our use of quantiles extends far beyond the simple computation of risk measures like VaR or GaR. Instead of focusing solely on specific probability thresholds, we will leverage quantiles to estimate the entire conditional distribution of GDP growth. By analyzing multiple quantiles across different probability levels, we can reconstruct the full predictive distribution of Y_{t+h} given X_t. This provides a more comprehensive understanding of risk, as it allows us to assess not only the probability of extreme events but also the entire range of possible outcomes and their respective likelihoods.

The quantile function described above is *"conditional"*, meaning its value depends on both X_t and τ. To build intuition, we will first discuss the estimation of unconditional quantiles before extending the concept to conditional quantiles.

2.1.2 Estimation of unconditional quantiles

As clarified in the previous section, quantiles and the CDF are closely related. In this section, we introduce the step-by-step estimation of the CDF and two methods for estimating quantiles. For clarity, we first focus on the unconditional CDF. Towards the end of the section, we introduce conditional quantiles and demonstrate how assuming linearity simplifies the expressions. The next section will explore the linear case in detail.

To build intuition, consider our variable of interest and assume we observe a time series of quarterly GDP growth rates, denoted as $y_1, \ldots, y_T$. We analyze a sample—a set of realizations—of the random variable Y_{t+h} and illustrate how to construct the CDF and quantiles. To distinguish empirical estimates from true values (i.e., those obtained in an infinitely large sample), empirical values are denoted with a hat.

Suppose we want to assess the (unconditional) probability that GDP growth is negative, which corresponds to evaluating the CDF at zero. A straightforward empirical approach is to compute the fraction of quarters in which GDP growth was negative. This fraction provides an estimate of the CDF at zero, representing the unconditional probability of negative GDP growth. Extending this method to all observed GDP growth values in the sample produces an empirical CDF, which is a step function increasing at each observed data point.

Figure 6 illustrates this approach. The right panel shows the time series of GDP growth, with red horizontal dashed lines identifying the fractions of observations corresponding to the 5^{th}, 25^{th}, 50^{th}, 75^{th}, and 95^{th} percentiles. The left panel represents the same data, now sorted from smallest to largest. The x-axis shows these ordered values, while the y-axis represents the fraction of observations that fall below each value. By construction, this step function is the empirical CDF of the GDP growth series, ranging between 0 and 1.

Formally, given the observed realizations y_{t+h} for $t = 1, \ldots, T - h$, the empirical unconditional predictive CDF, i.e., without incorporating independent variables X_t, is:

$$\hat{P}(Y_{t+h} \leq y) = \hat{F}_{Y_{t+h}}(y) = \frac{1}{T-h} \sum_{t=1}^{T-h} d_{t+h}(y),$$

where $d_{t+h}(y) = \mathbb{1}(y_{t+h} \leq y)$, with $\mathbb{1}(\cdot)$ denoting the indicator function that equals 1 if $y_{t+h} \leq y$ and 0 otherwise.

There are two ways to obtain empirical quantiles. First, they can be derived by sorting the data. The empirical quantile (left side of Figure 6) is a point on the horizontal axis that leaves a fraction τ of ranked observations to its left. The corresponding real number is denoted by $\hat{Q}_{Y_{t+h}}(\tau)$. By construction, exactly a fraction τ of observations falls below $\hat{Q}_{Y_{t+h}}(\tau)$. Sorting the observed values, the empirical quantile $\hat{Q}_{Y_{t+h}}(\tau)$ is given by:

$$\hat{Q}_{Y_{t+h}}(\tau) = \text{sorted}(y_{h+1}, y_{h+2}, \ldots, y_T) \text{ at position } \lfloor (T-h) \cdot \tau \rfloor.$$

In Figure 6, the vertical axis in the right panel, which represents the time series of data, becomes the horizontal axis in the left panel, where the empirical CDF is displayed. This transformation reflects the process of sorting the data before computing cumulative probabilities.

Figure 6: Empirical quantiles and time series of GDP growth

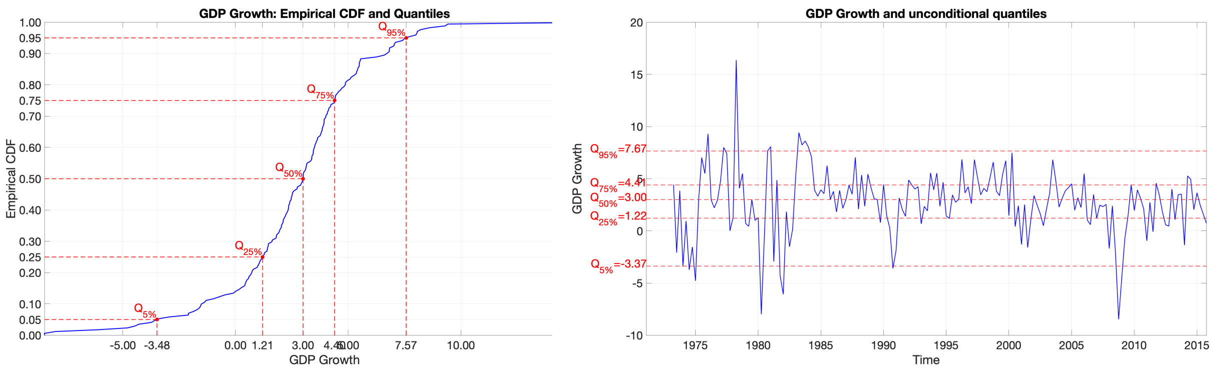

The sorting method becomes impractical for estimating conditional quantiles. In such cases, a more efficient approach involves solving a minimization problem. Before introducing the quantile setting, consider the mean of a distribution, which is defined as the center of a distribution via the following optimization problem:

$$\hat{\mu} = \arg\min_{c} \frac{1}{T-h} \sum_{t=1}^{T-h} (y_{t+h} - c)^2. \tag{5}$$

Thus, the arithmetic mean can be seen as the value that minimizes the average sum of squared deviations from its value. This example highlights three key points. First, the mean minimizes the sum of squared deviations. Second, each term in the summation is weighted equally by $\frac{1}{T-h}$. Third, introducing a conditioning variable X_t leads to the well-known ordinary least squares (OLS) estimates. It involve adding in the loss function of Equation 5 an additional term related to the conditioning variable(s).

The quantile optimization problem, introduced by Koenker and Bassett (1978), follows a similar approach but with two key differences. First, absolute deviations (or residuals) are used instead of squared ones. Second, deviations are weighted based on the quantile level. Thus, the minimization problem takes the following form:

$$\hat{Q}_{Y_{t+h}}(\tau) = \arg\min_{q} \sum_{t=1}^{T-h} \rho_\tau(y_{t+h} - q), \tag{6}$$

where the objective function $\rho_\tau(y_{t+h} - q)$ is labeled the check function or pinball losses. Calling the residuals as $u = y_{t+h} - q$, $\rho_\tau(u)$ is defined as:

$$\rho_\tau(u) = \begin{cases} \tau \cdot |u|, & \text{if } u \geq 0, \\ (1-\tau) \cdot |u|, & \text{if } u < 0. \end{cases}$$

This approach enables estimation of quantiles for any given τ, assigning weights equal to τ or $(1-\tau)$ to positive and negative errors, respectively. To illustrate the intuition behind this optimization framework, consider the special case where $\tau = 0.50$. In this setting, the objective function simplifies to:

$$\rho_\tau(u) = \begin{cases} 0.5 \cdot |y_{t+h} - q|, & \text{if } y_{t+h} - q \geq 0, \\ 0.5 \cdot |y_{t+h} - q|, & \text{if } y_{t+h} - q < 0. \end{cases}$$

The solution to this minimization problem is the value of q that divides the sample into two equal parts, corresponding to the median. The middle plot in the bottom row of Figure 7 illustrates the behavior of the pinball loss function for error values (u) ranging from -1 to +1.

The pinball losses of negative (positive) errors are plotted in red (green), with increasing intensity as the value of the function grows. Negative and positive errors are weighted equally, and their importance increases linearly with the absolute value of the error. The V-shaped plot represents the pinball loss function for the median, while the parabolic curve corresponds to the loss function for the mean. The median pinball loss is computed by simply multiplying the errors by their respective weight, 0.5.

For continuous variables, minimizing the sum of the pinball losses in Equation 6 is equivalent to shifting the V-shaped function along the x-axis to minimize the combined area of the two resulting triangles. The optimal position is at the center of the axis, thereby dividing the data into two equal parts.

This intuition can be further clarified through a spatial analogy. Consider the problem of determining the optimal location for a shop along a street, where houses are evenly distributed, and each has identical demand. To minimize total transportation costs, the shop should be positioned such that half the houses lie to its left and half to its right. This location minimizes the sum of absolute transportation distances, paralleling the minimization of absolute deviations in the estimation of the median. If the shop is placed too far left, more houses will be on the right, increasing total travel costs. If the shop is placed too far right, the opposite occurs. The optimal location is where half of the houses are on the left and half on the right—this is precisely the median.

Comparing the median pinball function and the loss function of the mean highlights why the median is less influenced by outliers. While the loss function for the mean increases quadratically with the error, the median loss function increases linearly. Consequently, extreme values have relatively less impact on the minimization process. This property extends to the regression setting.

Figure 7: Quantile estimation and check function

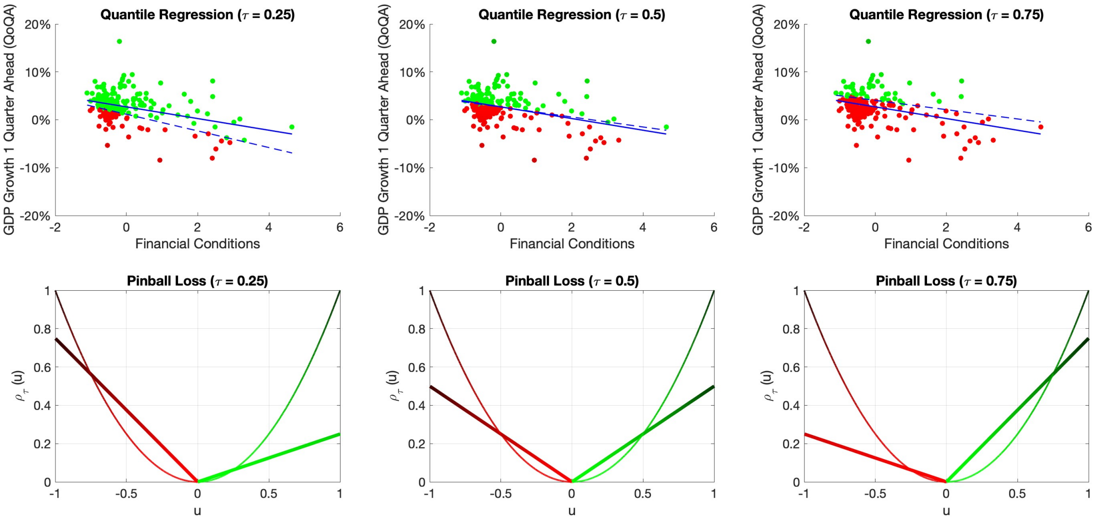

Notes: Effects of different weights in the check function.

The formulation of the pinball function shows that for quantiles below (above) the median, the minimization adjusts weights, placing greater emphasis on negative (positive) errors. Specifically, the weights are determined by the chosen quantile, such that:

- For $y_{t+h} - q < 0$ (or $u < 0$), higher weights are assigned for $\tau < 0.5$.
- For $y_{t+h} - q \geq 0$ (or $u \geq 0$), higher weights are assigned for $\tau > 0.5$.

The results of this asymmetric weighting process are depicted in the bottom row of Figure 7.

The left and right plots illustrate the pinball loss functions for the 25$^{\text{th}}$ and 75$^{\text{th}}$ quantiles, respectively, while the middle plot represents the median case. Again, in each plot negative errors are shown in red and positive errors in green, with color intensity increasing proportionally to the magnitude of the error.

For the estimation of the 25$^{\text{th}}$ quantile, the pinball loss function assigns greater weight to negative errors, whereas for the 75$^{\text{th}}$ quantile, positive errors receive more weight. This weighting structure reflects the fact that estimating the 25$^{\text{th}}$ quantile prioritizes the lower 25% of the ranked data, while estimating the 75$^{\text{th}}$ quantile emphasizes the upper 25%. This property is particularly important in a conditional setting, as it enables the estimation of the relationship between the dependent and explanatory variables while focusing on a specific segment of the CDF of the dependent variable.

So far, we have discussed unconditional quantiles. However, once the problem of estimating quantiles is defined within a minimization framework, extending it to a conditional setting follows naturally.

Recall from Figure 6 that quantiles represent the inverse of the cumulative distribution function. In the conditional setting, this inverse function depends on the control variables X_t for each quantile τ. However, without imposing additional assumptions, the precise functional relationship between the quantiles and X_t remains unknown:

$$Q_{Y_{t+h}|X_t}(\tau) = F^{-1}_{Y_{t+h}|X_t}(\tau) := h^{-1}(\tau, X_t).$$

A common approach, explored in the next section, is to adopt a linear approximation: $h^{-1}(\tau, X_t) = X_t'\beta_\tau$, where β_τ are the quantile regression coefficients.

By substituting this linear function into the original minimization problem, we obtain:

$$\hat{Q}_{Y_{t+h}}(\tau) = \arg\min_{\beta_\tau} \sum_{t=1}^{T-h} \rho_\tau(y_{t+h} - X_t'\beta_\tau).$$

This optimization is performed for each quantile level τ of interest, yielding a corresponding set of quantile regression coefficients β_τ. It is important to notice that the interpretation of the pinball loss function is the same as in the unconditional case, see bottom row of Figure 7.

In the corresponding scatter plots, depicted in the top rows of Figure 7, the estimated quantile line runs through the data, with deviations colored according to the pinball loss plots. In the left panel, representing the 25^{th} quantile, the check function disproportionately penalizes negative deviations, pulling the estimated line downward. This is because the lowest 25% of the data are more correlated with financial conditions than the upper 75%. Conversely, in the right panel, representing the 75^{th} quantile, positive deviations are assigned greater weight, causing the estimated line to tilt upward. This asymmetrical relationship between GDP growth and financial conditions is not captured by the OLS line which relates to the same minimization problem than the mean and hence only relates to the center of the distribution.

It is important to emphasize that this minimization framework is flexible and applies to any approximation of the quantile function $h^{-1}(\tau, X_t)$. As a result, one can extend beyond the linear case and incorporate more advanced methodologies, including supervised and unsupervised machine learning techniques. This will be discussed in Section 7.

3 Growth-at-Risk and Quantile Regressions

3.1 Quantile Regressions

Accurate predictions of risks would ideally entail making inference about the whole distribution of Y_{t+1}. If we were able to observe this directly, we would be able to compute all relevant measures of risk. Unfortunately, this is not directly observable from the data. Therefore, we need to use the limited information available and leverage on a set of predictors X_t observed at time t. This however allows us to estimate some portions or relevant characteristics, like moments and quantiles, of the conditional distribution $Y_{t+1}|X_t$. The quantile regression framework extends the traditional linear regression framework, where the objective is to estimate the conditional mean of Y_{t+1} given X_t:

$$\mathbb{E}[Y_{t+1}|X_t] = \beta_0 + \beta_1 X_t.$$

by estimating a whole set of conditional quantiles of the distribution of $Y_{t+1}|X_t$:

$$\mathbb{Q}_{Y_{t+1}}[\tau|X_t] = \beta_0^{(\tau)} + \beta_1^{(\tau)} X_t.$$

Similarly, distribution regression methods allow us to estimate conditional cumulative distribution functions:

$$F_{Y_{t+1}|X_t}(y) = \mathbb{E}[\mathbb{1}(Y_{t+1} \leq y)|X_t] = g(\beta_0 + \beta_1 X_t),$$

where $g(\cdot)$ is a link function, such as probit or logit.

The estimation method follows Koenker (2005), which imposes the assumption of linearity between the conditional quantile of interest and the predictors:

$$Q_{Y_{t+1}|X_t}(\tau) = \beta_0^{(\tau)} + \beta_1^{(\tau)} X_t.$$

The choice of the quantiles to estimate can depend on the scope of the exercise, but most applications focus on five quantiles: $\tau = 0.05, 0.25, 0.50, 0.75, 0.95$. The regression coefficients are obtained by minimizing the following tick loss function (as defined in Section 2.1.2).

$$\hat{\beta}^{(\tau)} = \arg\min_{\beta} \sum_{t=1}^{T} \rho_\tau (Y_{t+1} - \beta_0^{(\tau)} - \beta_1^{(\tau)} X_t),$$

where $\rho_\tau(u) = u(\tau - \mathbb{1}(u < 0))$ is the check function that assigns asymmetric weights to residuals, depending on whether they fall above or below the estimated quantile. This problem is solved using standard numerical optimization routines, typically initialized with OLS estimates.

Estimated quantiles of the conditional distribution of GDP growth are then computed based on the estimated coefficients and values of financial conditions—our only predictor:

$$\hat{Q}_{Y_{t+1}|X_t}(\tau) = \hat{\beta}_0^{(\tau)} + \hat{\beta}_1^{(\tau)} X_t.$$

Repeating this procedure for all the values of financial conditions yields five time series of conditional quantiles for GDP growth

$$\{\hat{Q}_{Y_{t+1}|X_t}(0.05), \hat{Q}_{Y_{t+1}|X_t}(0.25), \hat{Q}_{Y_{t+1}|X_t}(0.50), \hat{Q}_{Y_{t+1}|X_t}(0.75), \hat{Q}_{Y_{t+1}|X_t}(0.95)\}_{t=1}^{T-1}.$$

These estimates are represented in Figure 8. The financial conditions index is on the horizontal axis, GDP growth on the vertical axis; the blue line is the estimated conditional mean function (i.e. the OLS estimate). By definition of the OLS, the conditional mean function is chosen to minimize the sum of squared residuals: under the assumption of linearity in the coefficients, it represents the best linear approximation of the mean of the conditional distribution. The other lines represent the estimates of three selected quantile functions: the 5% quantile (black line), the 50% quantile (i.e. the median, red line), and the 95% quantile (green line). As discussed in Section 2.1.2, the median minimizes the sum of absolute deviations and ends up splitting the sample in two subsample of equal size, while the black and green lines, representing the lower and upper quantiles, are estimated using the check function, which assigns different weights to residuals above and below the estimated quantile.

Figure 8: Quantile Regression vs. OLS

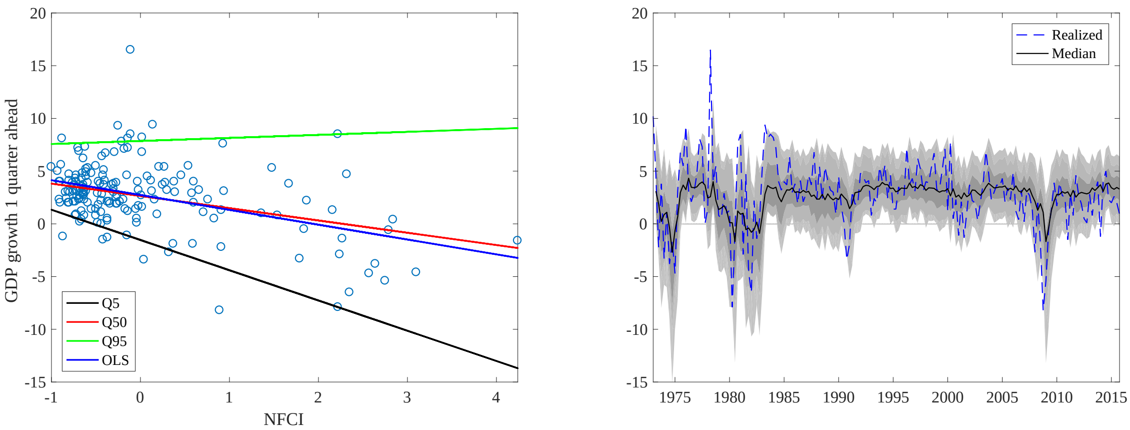

What emerges clearly from Figure 8–and it is key to our analysis–is that the OLS regression suggests only a mild negative relationship between financial conditions and GDP growth. Instead, estimated quantile functions display significantly different slopes. The 95% quantile regression line is nearly flat, suggesting that financial conditions do not significantly impact the likelihood of extremely high GDP growth realizations, while the 5% quantile regression line has a large negative slope, indicating that worsening financial conditions significantly increase the probability of experiencing extreme negative GDP growth outcomes.

This asymmetric response highlights the primary advantage of quantile regression over OLS in risk modeling. OLS, by construction, focuses on the conditional mean and assumes symmetric effects, failing to capture changes in the shape of the predictive distribution. Quantile regression, on the other hand, allows for a more flexible representation of risk, modeling different elements of the conditional distribution of GDP growth–which is particularly important in times of financial stress.[1]

Standard measures of risk, when talking about forecasting models for GDP growth ougth to answer the question: ”how bad can things get?”, or–maybe better–”what is the probability that GDP growth will fall below a certain value?”. The **Growth at Risk** metric answers precisely these questions. It equals the p-quantile of Y_{t+1} and mimics the notion of Value-at-Risk (VaR) from financial theory. More formally:

$$GaR_p(Y_{t+1}) = Q_{Y_{t+1}}(p) = F^{-1}_{Y_{t+1}}(p)$$

where p is usually chosen to be equal to 0.05 or 0.25.

Let’s consider an example. Suppose we are interested in the 5th percentile of the distribution, or the 5% worst predicted outcome. If GaR is -1.0% it means that our model predicts that GDP growth in the next quarter will be -1.0% or worse with a 5% probability. GDP@Risk is the upper bound of these 5% worst predicted outcome. In a sense, it’s the best realization we expect to get once the realized GDP growth falls in 5% worst predicted outcomes. Also, notice that GaR does not tell us whether GDP growth will be more likely to be -1.0% or -5.0%, it just tells us that growth will be -1.0% or worse with a given probability. For this reason, GaR is usually complemented

[1]By adding more assumptions we will be able to move from modeling single quantiles to the whole conditional distribution of GDP growth as discussed in Section 4

by additional measures. In Section 4.2 we will consider another measure of risk: the expected shortfall, which answers the quetsion: "if things get bad, how bad they will get?".

However, to answer this question, our set of estimated conditional quantile functions will no longer be sufficient. We will need an estimate of the full conditional distribution of GDP growth, which we will derive in Section 4.

Report some summary statistics, like the scatter with median and interquartile range, p50 and P75-p15, P5 and P50-P25. See Adrian et al. (2019)

3.2 Inference

We now turn to the problem of assessing the significance of the estimated quantile regression coefficients. Inference on quantiles in time series regressions is a relatively new area of research. Most of the existing literature has focused on the cross-sectional case, often overlooking key challenges specific to time series, such as serial correlation. A key reference in this field is Koenker (2005). His book dedicates some sections to the asymptotics of non-i.i.d. cases but largely ignores the issue of serially correlated errors, which is crucial in time series analysis.

Recently, there has been a growing interest in inference methods for macroeconomic applications, driven in part by the increasing use of distributional measures of risk, as introduced in the seminal work on Vulnerable Growth of Adrian et al. (2019). Adrian et al. (2019) examine whether empirical quantiles differ from those obtained under the assumption of a linear data-generating process. Gregory et al. (2018) and references therein tackle the issue of autocorrelation by developing an ad-hoc double-smoothing bootstrap procedure. Galvao and Yoon (2024) introduces a quantile-equivalent version of the HAC covariance matrix proposed by Newey and West (1987). Finally, Hoga and Schulz (2025) propose a two-step procedure based on a self-normalization approach, which also extends to expectiles.

In this section, we cover all these approaches. However, before doing so, we first introduce the problem at hand to establish a clear foundation. To highlight the key challenges involved in inference for a time series quantile setting, we draw a parallel with OLS regression. Recall Equation (3), which represents the in-sample regression used to estimate the model parameters β_0 and β_1:

$$Y_t = \beta_0 + \beta_1 X_{t-1} + \varepsilon_t.$$

Estimating this equation with our data yields the sample estimates, denoted by $\hat{\beta}_0$ and $\hat{\beta}_1$. If we were to repeat the estimation with another sample—say, a larger one—we would obtain different values of $\hat{\beta}_0$ and $\hat{\beta}_1$. In the limit, as the sample size increases indefinitely, these estimates converge to the true parameters β_0 and β_1.

In hypothesis testing, we evaluate the likelihood that an observed effect is due to chance. For example, to test whether the estimated coefficient $\hat{\beta}_1$ is statistically different from zero, we assess the probability of obtaining such an extreme estimate under the null hypothesis that $\beta_1 = 0$. If this probability (the p-value) is very low—say, less than 5%—then it is unlikely that $\hat{\beta}_1$ would have arisen by chance, leading us to conclude that β_1 is likely different from zero.

These conclusions depend crucially on the distribution of $\hat{\beta}_1$, which in turn hinges on its variance. Different assumptions about the error terms ε_t in Equation (3) lead to different variance structures. Let us examine this in more detail.

The variance of $\hat{\beta}_1$ can be derived from its expression. (For notational convenience, we assume that the data are demeaned.) In this case, the OLS estimator for the slope is given by

$$\hat{\beta}_1 = \frac{\sum_{t=1}^n X_{t-1}Y_t}{\sum_{t=1}^n X_{t-1}^2} = \beta_1 + \frac{\sum_{t=1}^n X_{t-1}\varepsilon_t}{\sum_{t=1}^n X_{t-1}^2},$$

where the second equality follows from substituting Y_t from Equation (3).

Since β_1 is a constant and $\sum_{t=1}^n X_{t-1}^2$ is observed (and hence non-random), the variance of $\hat{\beta}_1$ is given by

$$\text{Var}(\hat{\beta}_1) = \text{Var}\left(\frac{\sum_{t=1}^n X_{t-1}\varepsilon_t}{\sum_{t=1}^n X_{t-1}^2}\right) = \frac{1}{\left(\sum_{t=1}^n X_{t-1}^2\right)^2}\text{Var}\left(\sum_{t=1}^n X_{t-1}\varepsilon_t\right). \tag{7}$$

This expression can take various forms depending on the assumptions we make about the errors ε_t. For instance, if the errors have constant variance and are uncorrelated over time, then

$$\text{Var}\left(\sum_{t=1}^n X_{t-1}\varepsilon_t\right) = \sigma^2\sum_{t=1}^n X_{t-1}^2,$$

and Equation (7) simplifies to

$$\text{Var}(\hat{\beta}_1) = \frac{\sigma^2\sum_{t=1}^n X_{t-1}^2}{\left(\sum_{t=1}^n X_{t-1}^2\right)^2} = \frac{\sigma^2}{\sum_{t=1}^n X_{t-1}^2}.$$

In matrix notation—grouping all observations of the predictor in the vector X—this result is equivalent to writing

$$\text{Var}(\hat{\beta}_1) = \sigma^2(X^TX)^{-1}. \tag{8}$$

When the error term exhibits autocorrelation and/or heteroskedasticity, the variance of the estimator in Equation (3) must account for the nonzero covariances between error terms at different time points. In this case, the variance expression becomes

$$\text{Var}(\hat{\beta}_1) = (X^TX)^{-1}X^T\Omega X(X^TX)^{-1}, \tag{9}$$

where Ω represents the true covariance matrix of the error vector ε, and $X^T\Omega X$ reflects the weighting of these covariances by the predictor values. The expression in Equation (9) is sometimes called the "sandwich formula," with the outer parts $(X^TX)^{-1}$ depending solely on the observed data, and the middle part containing the unknown covariance structure of the errors.

To make inference on the coefficient's value, we need to know its distribution, which can be obtained by letting the sample size increase and applying the Law of Large Numbers and the Central Limit Theorem. In the limit, we have

$$\sqrt{n}(\hat{\beta}_1 - \beta_1) \xrightarrow{d} N(0, \Sigma),$$

where Σ denotes the asymptotic variance. In practice, because Ω is typically unknown, researchers often estimate it or employ robust methods—such as heteroskedasticity and autocorrelation consistent (HAC) standard errors—to ensure that inference is valid. This issue carries out also in a quantile setting.

In a quantile setting, we use similar principles to derive the asymptotic distribution of the estimated parameters. This distribution converges to a normal distribution as the sample size increases, but crucially, the asymptotic variance depends on the quantile being considered. To illustrate, consider the quantile regression model

$$Q_{Y_t|X_{t-1}}(\tau) = \beta_0^{(\tau)} + \beta_1^{(\tau)} X_{t-1} + \varepsilon_t^{(\tau)}.$$

When the error terms are i.i.d., the asymptotic variance of the quantile regression estimator is analogous to Equation (8) for OLS. In the quantile setting, it becomes:

$$\mathrm{Var}(\hat{\beta}_1^{(\tau)}) = \frac{\tau(1-\tau)}{f_{\varepsilon_t^{(\tau)}}(0)^2} (X^T X)^{-1}, \tag{10}$$

where $f_{\varepsilon_t^{(\tau)}}(0)$ is the value of the probability density function of the error term evaluated at zero (i.e., at the τ quantile). This implies that $f_{\varepsilon_t^{(\tau)}}(0)$ generally decreases as we move away from the median toward the tails of the distribution, thereby increasing the asymptotic variance of the estimated quantile regression coefficients. **ADD TEXT ON AUTOCORRELATION + Make link with what comes next**

To address the question of significance, we first examine in Section 3.2.1 how the estimated quantiles differ from those generated by a linear regression model. We then move on, in Sections 3.2.2 and 3.2.3, to constructing the distribution of the estimated parameters. This allows us to evaluate the probability that the true parameter has a sign opposite to our estimate, which would undermine our conclusions. Section 3.2.2 assumes that the parameter's distribution follows a normal distribution, while Section 3.2.3 takes a different approach, constructing the parameter's distribution using the bootstrap method by generating artificial data from a given distribution. Komunjer (2005)

Figure 9: Add description

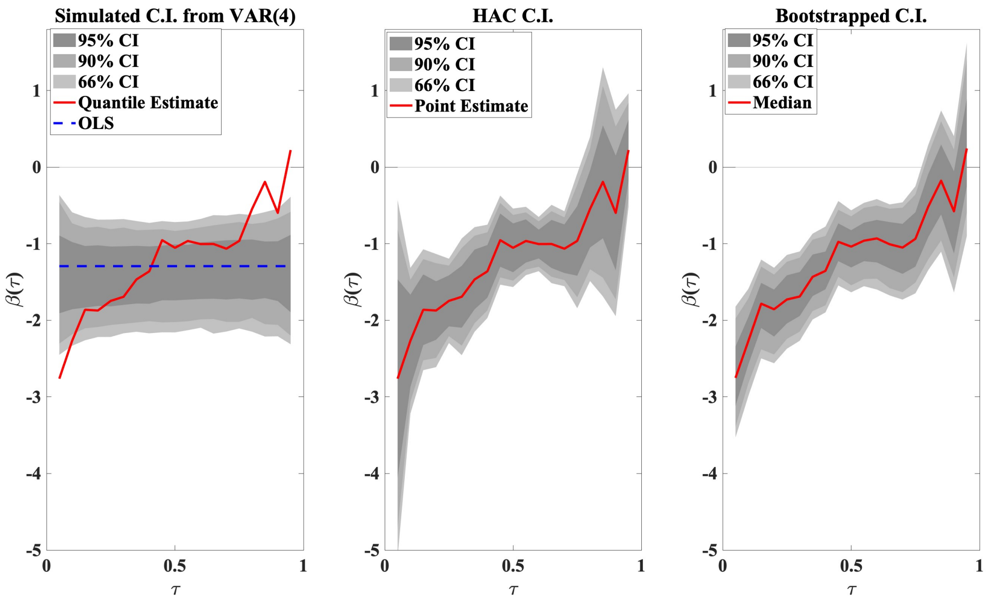

3.2.1 VAR-based simulation

The left plot of Figure 9 is generated as follows. We start from the assumption that the data is generated from a linear model, a VAR(4), and we estimate the parameters of this model. Then, based on the estimated parameters, we generate 1000 simulated paths for our variables. Each simulation is done as follows. Starting from an initial condition, we generate a path of i.i.d. shocks which are then fed into the model to generate the path. Given the simulated path, we run quantile regressions and store the estimated coefficients. Repeating this process 1000 times gives us, for every quantile level, a set of 1000 coefficients. The plot in the graph reports the confidence bound obtained from these simulation-based coefficients. For example, for $\tau = 0.50$ the plot reports the median (blue, dashed lines), the 67% (darkest shades), the 95% (lighter shades), and the 99% (lightest shades) confidence band of the 1000 simulation-based medians estimated by running a quantile regression for each simulation. The red lines are the quantile regression estimates obtained from the original data. The quantile estimates are linear and close to the OLS coefficients (dashed, black lines). This is not surprising as they are generated from a linear model. The simulations allow us to assess the uncertainty around these estimates (shaded areas). The shaded areas become larger for extreme quantiles, because in the estimation large weight is put on only a few data points. However, it is important to stress that even in the estimation of extreme quantiles, we are using all the data available. Only we weigh them differently. This is why uncertainty increases for extreme quantiles but does not become extremely large or go infinite.

Comparing the sample quantiles with the simulated ones, we notice that the sample quantiles of the financial conditions growth fall mostly outside the (darkest) shades. Hence, they cannot be reconciled with a model and are generated by some linearity that is present in the data.

3.2.2 HAC-based inference

TBD

3.2.3 Bootstrap

TBD Gregory et al. (2018) Silverman and Young (1987)

3.3 Out of sample analysis

We now move to forecast evaluation. Since our goal is forecasting risk, this is at the core of the analysis. This topic focuses on the importance of out-of-sample prediction (external validity) compared to in-sample fitting (internal validity). In-sample fitting evaluates how well a model explains the data it was trained on, while out-of-sample prediction tests the model on data outside the training set, ensuring it generalizes and avoids overfitting. Hence, with internal validity, you have your observation from 1 to t, and then you check how you fit that observation from 1 to t. With external validity, or out-of-the-sample, you estimate your parameters up in the sample, and then you see how you do outside of the sample, hence from observation form $t + 1$ onward.

The analogy of maps helps explain this: a detailed map (overfitting) may lead to errors in navigation (forecasting), whereas a simpler map (appropriate model complexity) might reduce errors in new scenarios. Forecasting accuracy hinges on external validity, especially in models with many parameters where uncertainty in estimation becomes critical.

To evaluate the predictive performance of the estimated quantile regression models, we perform out-of-sample (OOS) evaluation — based on quantile loss functions — and assess the calibration of the predicted quantiles. It is important to stress the difference between the two procedures. Out-of-sample evaluation is based on a loss function which penalizes a prediction the further away is from the (out-of-sample) realization. Lower losses indicate better forecasts. Out-of-sample calibration ensures that estimated quantiles align with observed (out-of-sample) frequencies. Hence, it hinges upon the empirical CDF of the out-of-sample realizations, which is compared to the forecasted quantile. These complementary assessments provide a comprehensive validation of the quantile regression model's reliability and generalization performance. In what follows we introduce the general setup and then we move to the two OOS assessments methods.

3.3.1 Setup

These out of sample predictions are made via the following recursively procedure. The first predictions, corresponding to 1993Q1 (one quarter ahead, left graph) and 1993Q4 (one year ahead), are performed with the estimation sample that ranging from 1973Q1 to 1992Q4:[2]

1. using the sample 1973Q1 to 1992Q4, we estimate quantile coefficients $\hat{\beta}_{\tau,T_1}$, where $\tau = [0.05, 0.25, 0.5, 0.75, 0.95]]$ and $T_1 = [1, T] = [1973Q1, 1992Q4]$;
2. we use $\hat{\beta}_{\tau,T_1}$ and the observations in $T = 1992Q4$ to forecast the $\hat{Q}_{Y_{T+h}|X_T}(\underline{\tau}$ for $h = [1, 4]$.

Then the procedure is repeated for $T_1 = [2, T+1] = [1973Q2, 1993Q1]$ to get $\hat{Q}_{Y_{T+1+h}|X_{T+1}}(\underline{\tau})$ and so on until the end of the sample.

Forecasting accuracy hinges on external validity, especially in models with many parameters where uncertainty in estimation becomes critical. The charts compare in-sample (black lines) and out-of-sample (blue lines) predictions.

[2]Here we are not accounting for data revisions, neither in GDP growth nor in the NFCI. In this sense the exercise in pseudo out of sample. This in principle could be a major problem since the NFCI is obtained from a Kalman smoother based on dynamic factor model estimated using the Quasi-Maximum Likelihood approach of Doz et al. (2012).

A full real time exercise has been conducted by Amburgey and McCracken (2023b). The general results are confirmed.

Figure 10: Add description

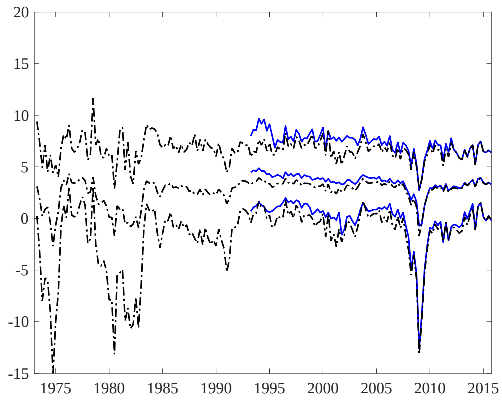

The outcome of this procedure is a time-series of 20 years of quantile forecasts for each of the two forecast horizons. The similarity between the in-sample quantiles and out-of-sample predictions indicates no overfitting and that the model performs well out-of-sample, even at the extremes (tails). The intuition why it works, you can think is that when you estimate the lowest quantile, it's not just using 5% of the observations. You are using all the observations. Of course, you are weighting more these 5% observations, but you also use the other. That's why you get estimated that are stable and reliable.

3.3.2 Out-of-Sample Calibration

Calibration evaluates whether the estimated quantiles $\hat{Q}_{Y_{t+h}|X_t}(\tau)$ correctly capture the empirical distribution of the observed outcomes Y_{t+h}. A well-calibrated model ensures that, for a given quantile τ, approximately $100\tau\%$ of the actual observations fall below the predicted quantile.

Formally, the **empirical coverage probability** is defined as:

$$\hat{C}(\tau) = \frac{1}{T}\sum_{t=1}^{T} \mathbb{1}\{Y_{t+h} \leq \hat{Q}_{Y_{t+h}|X_t}(\tau)\}, \tag{11}$$

where $\mathbb{1}\{\cdot\}$ is an indicator function that equals 1 if $Y_{t+h} \leq \hat{Q}_{Y_{t+h}|X_t}(\tau)$ and 0 otherwise. Ideally, we expect $\hat{C}(\tau) \approx \tau$ for all estimated quantiles τ.

3.3.3 Out-of-Sample Evaluation

Our aim here is to assesses the accuracy of the estimated conditional quantiles $\hat{Q}_{Y_{t+h}|X_t}(\tau)$.

A key evaluation metric is the **quantile loss function**, also known as the **pinball loss** or **lin-lin loss** , defined as:

$$L_\tau(Y_{t+h}, \hat{Q}_{Y_{t+h}|X_t}(\tau)) = \begin{cases} \tau(Y_{t+h} - \hat{Q}_{Y_{t+h}|X_t}(\tau)) & \text{if } Y_{t+h} \geq \hat{Q}_{Y_{t+h}|X_t}(\tau), \\ (1-\tau)(\hat{Q}_{Y_{t+h}|X_t}(\tau) - Y_{t+h}) & \text{if } Y_{t+h} < \hat{Q}_{Y_{t+h}|X_t}(\tau). \end{cases} \tag{12}$$

For each quantile τ, we compute the quantile loss over the test set and take the average to obtain an overall measure of predictive accuracy. Lower quantile loss values indicate better predictive performance.

3.4 Robustness

We should talk about the robustness of the quantile regression with outliers. Report the comparison of the quantile loss with the quadratic and discuss the difference in the influence function, with quadratic loss overemphasizing extreme deviations. Then, we should extend the sample to include COVID-19. We should report a chart showing how quantile predictions in real-time remain robust while OLS becomes crazy. References here should include work on robustness (Croux). In Growth-at-Risk during covid, we have the paper Chernis et al. (2023).

In this section we should also talk about applications to other countries, different sample, and argue that the finding of growth vulnerability is very robust.

- Coe and Vahey (2020) for secular data, from late 19th century.
- (Amburgey and McCracken, 2023b) on data revisions
- Work on other countries, panel, ... Figueres and Jarociński (2020) for the euro area. Adrian et al. (2022); Gächter et al. (2023) for international, add

4 Beyond Quantiles: Density-Based Risk Measures

4.1 Recovering full density: 2-step with skewed t-distribution

In Section 3.1 we showed how to derive the estimation of quantiles of the conditional distribution of GDP growth, approximated using linear functions, and discussed their relationship with uncertainty, and some measures of risk directly measurable from the estimated quantiles. We now take our analysis one step further and move beyond quantiles to construct a full predictive distribution of GDP growth using the estimated quantiles of Section 3.1, plus some assumptions on the shape of the distribution that we want to find. This will yield a complete predictive distribution, offering a richer representation of uncertainty and variability in the data.

First, recall from Section 2.1 that the quantile function is the inverse of the cumulative distribution function (CDF),

$$Q_{Y_{t+h}|X_t}(\tau) = F^{-1}_{Y_{t+h}|X_t}(\tau) := h^{-1}(\tau, X_t) \tag{13}$$

which we have approximated it using a linear function (section 3.1),

$$Q_{Y_{t+h}|X_t}(\tau) \approx X_t'\beta_\tau \tag{14}$$

We want to derive the entire conditional distribution of GDP growth, or the conditional PDF, $f_{Y_{t+h}|X_t}$. Mathematically, this function is the first derivative of the CDF. Do do so, we proceed in three steps: (i) we estimate the conditional quantiles based on observed data (financial conditions in our case), (ii) we select–by assumption–a class of potential candidates for our conditional distribution, and (iii) from the class of potential candidates, we find the distribution that best matches our estimated quantiles.

Step (i) has been covered in Section 3.1, and we focus here on steps (ii) and (iii). Among all the families of parametric probability density functions, we choose the Skew-t distribution introduced by Azzalini and Capitanio (2003), as proposed by Adrian et al. (2019). This distribution provides a flexible parametric form, which is particularly suited to analyze phenomena that occasionally have fat tails, and asymmetries which–as discussed in sections 1 and 2–is very much our case.

The family of Skew-t distribution that we use, have four parameters: μ, the location parameter, regulating the central tendency; σ, the scale parameter, controlling dispersion (variance); α, the shape parameter, capturing asymmetry; and ν, the degrees of freedom parameter, accounting for fat tails. This distribution is very versatile and encompasses many other distributions as special cases. For example, when the degrees of freedom ν approach infinity and α equals zero, the distribution simplifies to a normal distribution. When the degrees of freedom approach infinity but α is nonzero, the distribution becomes an asymmetric normal distribution.

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

To find the Skew-t distribution that best describes the conditional distribution of GDP growth, we estimate these four parameters so that the the estimated distribution will have quantiles that are as close as possible to the conditional quantiles that we estimated in Section 3.1. Because we have four unknown parameters, we will need (at least) four conditional quantiles. We use:

$$\mathcal{T}=\{0.05,0.25,0.75,0.95\}.$$

Since the estimated quantiles depend on X_t, the outcome of this procedure will be a time series of Skew-t distributions. In other words at every point in time, we have a cross-section of four quantiles that we use to estimate the the time series evolution of the Skew-t parameters:

$$\{\hat{\mu}_t,\hat{\sigma}_t,\hat{\alpha}_t,\hat{\nu}_t\}=\arg\min_{\mu_t,\sigma_t,\alpha_t,\nu_t}\sum_{\tau\in\mathcal{T}}\left(\hat{Q}_{Y_{t+h}|X_t=x_t}(\tau)-Q(\tau;\mu_t,\sigma_t,\alpha_t,\nu_t)\right)^2$$

For each time t, this minimization procedure selects the parameters $\{\hat{\mu}_t,\hat{\sigma}_t,\hat{\alpha}_t,\hat{\nu}_t\}$ such that the corresponding Skew-t distribution has quantiles as close as possible to those obtained from the linear model. It is important to emphasize, however, that this yields a *conditional* PDF, whose shape depends on X_t through the quantile regression coefficients.

Let's analyze the results from this procedure. Figure 11 shows the derived quantiles for two difference periods 2006 Q2 (left) and 2008 Q4 (right). The yellow line represents the estimated quantiles (i.e. $X_t'\beta_\tau$.) following the procedure of Section 3.1, where conditional quantiles are derived through a regression that uses the financial conditions index as the only control variable. The dashed red line and the blue line, instead, represent the results from two fitted Skew-t distributions: one that includes only lagged GDP as the control variable, and one that includes both lagged GDP and financial conditions as control variables.

Figure 11: NB: Figure to be adjusted

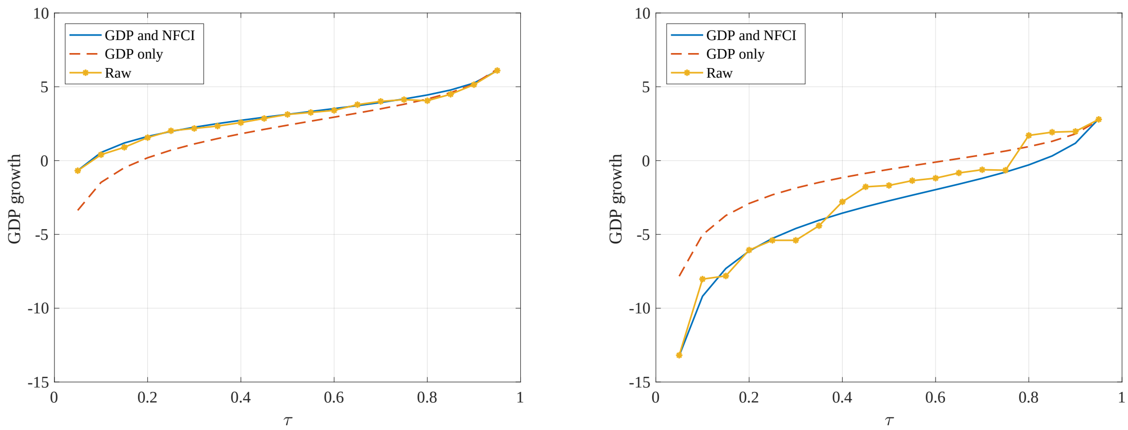

The figure shows that our parametric fit provides a very good approximation to the estimated quantiles: the blue line tracks very closely the yellow line, especially in the left panel. But comparing the quantiles obtained using different models leads to a much more central finding, which speaks to the importance of financial conditions in predicting downside risks: the model that includes financial conditions does a better job at capturing downside risks during periods of economic stress. In 2006 Q2, a period of relative stability without significant financial stress, the blue and red quantiles closely align. This indicates that financial conditions add little incremental predictive power under stable economic conditions. The quantiles follow a smooth upward trajectory, and the differences between models are subtle, reflecting a low-risk environment. In 2008 Q4, a period of heightened financial stress preceding the Lehman collapse, the differences between the blue and red quantiles become more pronounced. Including financial conditions (blue line) shifts the quantiles, particularly in the lower tails (e.g., the 5^{th} and 25^{th} percentiles), highlighting increased downside risk when financial stress is accounted for.

Finally, once the have computed the estimated quantiles $\widehat{Q}_{Y_{t+h}|X_t}$, as in Figure 11), we can simply invert them to compute the estimated conditional CDF, $\widehat{F}_{Y_{t+h}|X_t}(\hat{\mu}_t, \hat{\sigma}_t, \hat{\alpha}_t, \hat{\nu}_t)$:

Figure 12: Predictive distributions in selected time periods

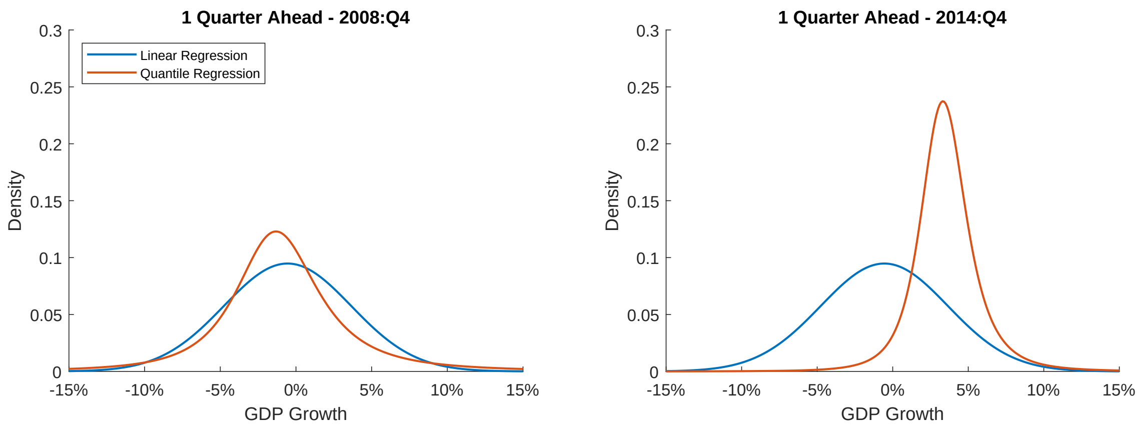

Notes: Normal and Skew-t based distributions in the 2008Q4 (left plot) and in 2016Q4.

Figure 12 illustrates the estimated conditional CDFs for the regression model that includes both lagged GDP growth and financial conditions for the two same periods as in Figure 11, but compares the CDF obtained with quantile regression (red line) with the standard OLS case (blue line). In the 2006 Q2 panel, a period marked by financial stability, the predictive model that includes financial conditions (solid blue line) shows a more concentrated distribution with less variance compared to the simpler autoregressive model using only lagged GDP (dashed red line). This tighter distribution suggests a lower risk of GDP falling below zero, indicating that including financial conditions leads to more accurate forecasts and a decreased likelihood of negative GDP growth.

Conversely, the 2008 Q4 panel, which precedes the Lehman collapse and represents a period of intense financial stress, displays significant differences. The autoregressive model (dashed red line) shifts modestly leftward, lowering the expected GDP slightly while maintaining a similar variance. However, incorporating financial conditions into the predictive model (solid blue line) results in a distribution that not only shifts leftward but also broadens considerably, signaling an increased probability of negative GDP outcomes and heightened uncertainty during stressful periods.

This expansion of the predictive distribution under stress underscores a crucial point: downside risks escalate significantly, whereas the potential for positive outcomes remains relatively unchanged. For instance, the probability of recession (GDP below zero) might typically be around 50

Financial stress periods, like 2008 Q4, often coincide with surges in financial condition indices and metrics such as the VIX, which measures stock market volatility. When the VIX increases, it indicates tightened financial conditions correlating with heightened macroeconomic risks. A similar pattern emerged during the COVID-19 pandemic, although aggressive policy responses helped alleviate some of the financial pressures unlike during the 2008 crisis.

The primary insight from comparing these two periods is that models integrating financial conditions deliver a more precise and robust assessment of risk, especially under stress. They demonstrate how financial stress can increase variance and skew GDP growth forecasts toward more negative outcomes, highlighting the importance of financial markets as predictors of economic vulnerability.

When we apply this procedure for all periods in our sample, we can track the dynamics of the fitted distribution over time (Figure 13):

Figure 13: Fitted Skew-t distribution

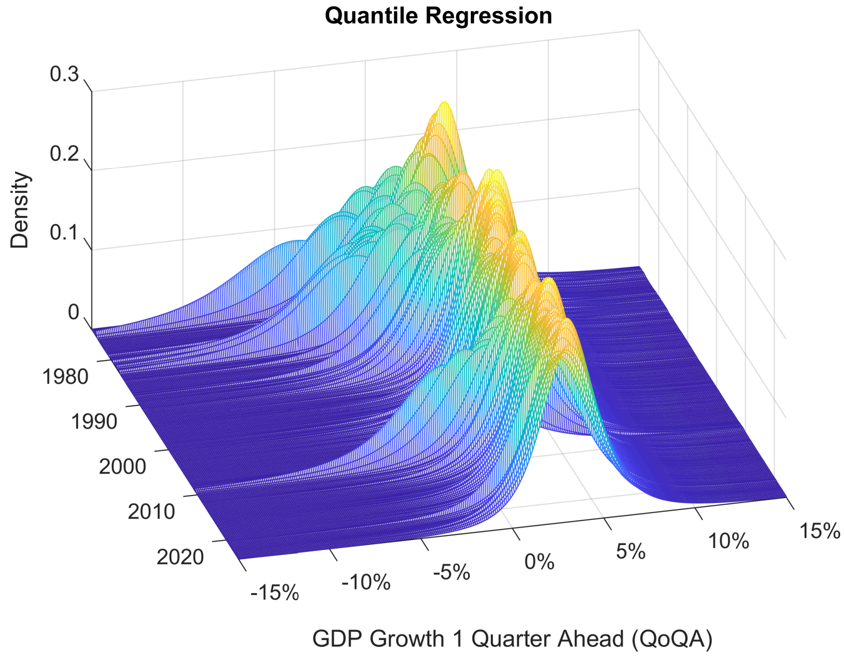

Note:

4.2 Additional measures of risk: Beyond G@R

Now that we have established a predictive distribution, we can focus on assessing risk — which in our benchmark example can be summarized as the vulnerability of GDP growth. To do so, we define a set of probability-based risk measures that fully leverage our predictive distribution, extending beyond Growth-at-Risk (GaR), the quantile-based measure of risk defined in Section 3.1. We present metrics providing a deeper understanding of how severe, likely, or adverse potential outcomes may be.

Recall that Growth-at-Risk (GaR) identifies the threshold below which GDP is expected to fall with a given probability. However, it provides no information about the distribution of losses beyond this level. The Expected Shortfall (ES), address this limitation by looking at the average losses, conditional on the event that the GaR is surpassed. Is is also known as Conditional Value at Risk (CVaR) and it answers the question: "If things do get bad, on average, how bad will the outcomes be?"

Expected Shortfall calculates the average GDP growth conditional on it being below a specific threshold (e.g., GaR). It quantifies the magnitude of (average) expected extreme negative outcomes.

Formally, the Expected Shortfall (SF) is the expected value of Y_{t+h}, conditional on the fact that Y_{t+h} is below a given threshold (or quantile) $\underline{\tau}$. This threshold is typically the Value at Risk (VaR). For instance, if $\underline{\tau}$ is set at 5%, the expected value conditional on being below this quantile is calculated as:

$$SF_t = \mathbb{E}\left[Y_{t+h} \mid Y_{t+h} < Q_{Y_{t+h}|X_t}(\underline{\tau})\right] = \frac{1}{\underline{\tau}} \int_0^{Q_{Y_{t+h}|X_t}(\underline{\tau})} y f_{y_{t+h}|x_t}(y|x_t) dy$$

Here, $f_{y_{t+h}|x_t}(y|x_t)$ represents the Probability Density Function (PDF). In simple terms, the Expected Shortfall answers the question: *"If GDP growth falls below a risky* *threshold (the growth-at-risk level), what is the average GDP drop we can expect?"* Essentially, it measures how bad things are likely to get in adverse scenarios.

The Expected Longrise (LR) mirrors the Expected Shortfall but focuses on the upper tail of the distribution. It evaluates positive events and tells us the expected value of Y_{t+h}, conditional on it being above a certain quantile threshold $\bar{\tau}$. In other words, it captures how good things could be in favorable scenarios.

$$LR_t = \mathbb{E}\left[Y_{t+h}|Y_{t+h} > Q_{Y_{t+h}|X_t}(\bar{\tau})\right] = \frac{1}{1-\bar{\tau}} \int_{\bar{\tau}}^{1} Q_{Y_{t+h|X_t}(\tau)}(\tau) d\tau$$

These two metrics together provide a comprehensive view of the extremes of the distribution, with the ES (EL) measuring the average downside risk (upside potential) below (above) a given threshold, while EL measures the average upside potential above a given threshold.

Downside entropy measures the relative likelihood of negative outcomes compared to the average scenario. It answers: "How likely are negative outcomes to occur, relative to the average?"

Downside (Upside) Relative Entropy is denoted as $\mathcal{L}_t^D(f_t; g)$ ($\mathcal{L}_t^U(f_t; g)$) and quantifies the difference between two probability distributions.

$$\mathcal{L}_t^D(f_t; g) = -\int_{-\infty}^{F_t^{-1}(50)} (\log g(y) - \log f_t(y)) f_t(y) \, dy$$
$$\mathcal{L}_t^U(f_t; g) = -\int_{F_t^{-1}(50)}^{\infty} (\log g(y) - \log f_t(y)) f_t(y) \, dy$$

Hence we compare the conditional PDF of GDP growth, denoted as $f_{y_{t+h}|x_t}(y \mid x_t)$ (also labeled as $f_t(y)$), and the unconditional PDF, denoted as $g(y)$. The Relative Entropy is calculated as the weighted average of the logarithm of the ratio between these two distributions, integrated over the entire support of the distributions. The intuition is the following:

- The ratio of the two distribution functions over a specific interval reflects how different the two distributions are within that interval.
- Taking the logarithm of the ratio scales the differences and prevents extreme values from disproportionately influencing the weighted average.
- The weights are the values of the PDF over the intervals used to compute the ratio. This ensures that differences in more likely events (those with higher PDF values) are weighted more heavily.

In summary, Relative Entropy provides a measure of how much the conditional PDF (capturing current or state-dependent dynamics) diverges from the unconditional PDF (representing the baseline or overall distribution), while accounting for the likelihood of different outcomes. In a sense is the expected difference between the two distribution. When $\mathcal{L}_t^U(f_t; g)$ is high the model indicates that the on average, negative growth outcomes are more likely than compared to the prediction of the unconditional distribution.

Another key measure of risk to GDP growth is by **Recession Risk** (RR). Recession periods tend to correspond to periods of negative GDP growth. Therefore, we can approximately compute recession risk from the distribution of GDP growth as follows:

$$RR_{t+1} \approx Prob(Y_{t+1} \leq 0) = F_{Y_{t+1}}(0)$$

which tells us that we can measure recession risk using the statistical notion of a cumulative density function (CDF) of GDP growth and evaluate it at zero. This also implies that we can estimate recession risk using statistical tools such as **distribution regressions**.

In Figure 14a, the graph on the left represents the PDF of GDP growth for 2006 Q2, a normal economic period. It illustrates the likelihood of different GDP growth outcomes, with positive growth (the upside) shown above zero and negative growth (the downside) below zero. The Growth-at-Risk (GaR) is indicated in red, marking a threshold below which there is a 5% probability of outcomes. This corresponds to the gray-shaded area on the downside tail. Conversely, the upper tail (highlighted in green) captures the potential for favorable growth outcomes above a specific quantile.

The right-hand graph shows the CDF and its inverse, the quantile function. Transitioning from the PDF to the CDF involves accumulating probabilities. The CDF at a given point represents the total probability of outcomes less than or equal to that point. The quantile function, obtained by inverting the CDF, maps probability levels (e.g., 5%) to the corresponding GDP growth values.

The ES, represented by the red area in the downside tail, is the expected value of GDP growth conditional on being below the Growth-at-Risk threshold. It measures the average severity of adverse outcomes. It can be computed using the quantile function. Specifically, the ES is equivalent to the integral of the quantile function over the interval $[0, \tau]$, where τ denotes the risk level, such as 5%.

These relationships between the PDF, CDF, and quantile function illustrate a unified framework for understanding risks and opportunities in GDP growth. By linking probabilities to potential outcomes, the Expected Shortfall and Longrise provide valuable insights into the distribution's extremes, quantifying both downside risks and upside potential.

Figure 14: NB: Figure to be adjusted

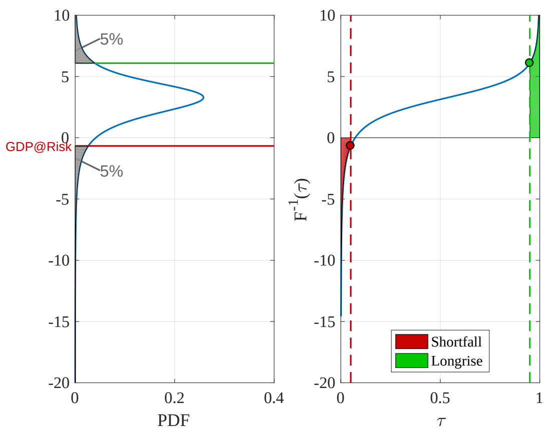

(a) A subfigure

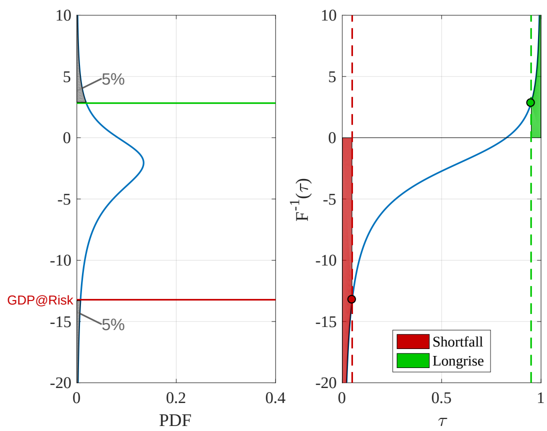

(b) A subfigure

Figure 14b allows to illustrate the differences between a normal economic period (2006Q2, Figure 14a) and a period of financial stress in terms of GDP growth risk measures.

In 2006Q2, the PDF is symmetric, with a well-centered peak near zero, reflecting stable economic conditions. The downside risk, represented by the shaded gray area (5%), is minimal, with a Growth-at-Risk (GaR) threshold slightly below zero. The Expected Shortfall and the Expected Longrise are both relatively small, consistent with a stable period where extreme events are unlikely.

In contrast, the 2008Q4 graphs reflect the heightened financial stress of the period. The PDF becomes asymmetric, with a significant shift to the left, reflecting economic contraction during the financial crisis. The downside tail is much longer and steeper, indicating increased variability and a higher likelihood of adverse GDP growth outcomes. The Growth-at-Risk threshold is now significantly lower, around -15, illustrating heightened downside risks. The upside potential is reduced with respect to 2006Q2, but the this decrease is minimal with respect to the strong increase in downside risk. This is due to the asymmetric behavior of of GDP growth in stress periods: The CDF shifts leftward due to the negative skewness in the distribution. The Expected Shortfall, represented by the red area, becomes much larger, indicating severe average losses conditional on being below the GaR threshold. The Expected Longrise, represented by the green area, is slightly larger than in 2006Q2 but remains relatively small compared to the Expected Shortfall, showing that positive opportunities are limited in periods of financial stress.

The key differences between the two periods are evident in the downside risks, upside potential, and distribution symmetry. In 2006Q2, the downside risk is minor, with a shallow tail and a small Expected Shortfall. In 2008Q4, the downside risk increases significantly, with a steep and extended tail leading to a much larger Expected Shortfall. Both periods have limited upside potential; however, the Expected Longrise in 2008Q4 is slightly higher due to increased variability, but it remains much smaller than the downside risks. The 2006Q2 distribution is symmetric, typical of stable economic periods, while the 2008Q4 distribution exhibits a strongly negative skew, highlighting the economic contraction and heightened downside risks during the financial crisis.

The financial stress in 2008Q4 shifts the entire distribution leftward, emphasizing the increased likelihood of negative GDP growth outcomes. This comparison highlights how economic stress alters the GDP growth distribution, increasing downside risks while leaving upside potential relatively unchanged. The Expected Shortfall becomes a critical measure in such stressed periods, as it captures the average losses under adverse scenarios. This contrast underscores the importance of accounting for financial conditions when assessing macroeconomic risks.

4.3 Out-of-Sample Analysis

Here we should add the additional evaluations we can do with the full density: log scores evaluation, ES in and out of sample. Let's remember to add for the log score, the test of equal predictive ability relative to a naive model that uses the empirical quantiles (basically dropping financial conditions and regressing only on a constant)

- With a density, we can also evaluate calibration by looking at the moments of the PIT D'Agostino et al. (2013).
- for calibration we use Rossi and Sekhposyan (2015) (see replication package)
- we can show the comparing of expected shortfall in sample and OOS
- Stress importance of: in-sample vs pseudo out-of-sample vs real-time

- For real-time we can use: real GDP and NFCI vintages for US, Amburgey and McCracken (2023b)
- we need to discuss inference on accuracy and calibration. Diebold-Mariano...

To cite/look: Corradi et al. (2023) Corradi and Llorens-Terrazas (2024) (this one in extentions).

4.3.1 Out-of-Sample Evaluation

Predictive scores assess the accuracy of probabilistic forecasts by measuring the probability assigned to observed outcomes—the higher, the better. Suppose we are in **2006 Q1** and forecast GDP for **2006 Q2**. Once realized, we evaluate the forecast by checking the probability density assigned to the actual GDP value before observation, known as the **predictive score**. A higher density at the realized value indicates a more **accurate** forecast.

Figure 15: Add description

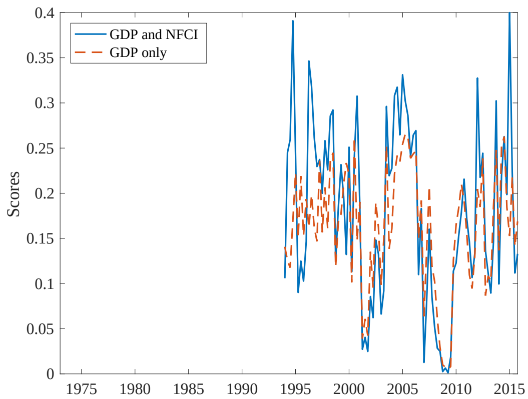

This concept applies across different economic conditions. For example, if GDP growth in **2008 Q4** was positive, incorporating **financial conditions** might have worsened the forecast, while for growth below **3%**, it could have improved it. Evaluating forecasts **in real-time**, using only available information at the time, helps assess their actual performance.

Predictive scores compare models: a **model with financial conditions** (blue line) may produce **sharper predictions**, assigning more probability mass to observed outcomes, while a **model without financial conditions** (red line) may be less precise. By analyzing **one-quarter-ahead** and **four-quarters-ahead forecasts**, we can determine which model performs better.

Interestingly, the model incorporating financial conditions tends to perform well in periods of **low financial stress**, where the blue line (model with financial conditions) is higher than the red line (model without). However, during recessions, the advantage of including financial conditions diminishes. In extreme downturns, such as the **Great Recession**, the difference between the two models narrows, particularly in the lower tail of the distribution. This occurs because, in crisis periods, while financial conditions help improve forecasts, both models tend to predict similar downside risks, reducing their relative difference.

In normal times, financial conditions provide sharper predictions by increasing confidence in stable outcomes. Since there is no financial stress, the model assigns higher probability to the realized GDP value, resulting in a **more concentrated probability density function (PDF)**. This increased certainty explains why the model performs best when financial markets are stable.

A common argument is that forecasting risk is difficult because recessions are hard to predict. This is true—recessions are tail events. However, a good risk forecast, including downside risk, is not necessarily one that perfectly predicts recessions. Instead, it is a model that distinguishes between different risk scenarios.

This becomes clearer when considering recession forecasting as a binary event, similar to predicting **rain or no rain**. Financial conditions provide valuable information: when they are stable, we can be fairly confident that a recession is not occurring. This insight allows us to tilt the predictive density toward the tail in times of financial stress.

There is also research on weighting density forecasts toward specific regions of interest. For example, Diks et al. (2011) explores methods for evaluating density forecasts by assigning greater weight to particular areas, which can also be useful for combining multiple predictions.

For a probabilistic forecast of Y_{t+h}, the **log predictive score** (or **log scoring rule**) is:

$$S_{\log}(Y_{t+h}) = -\log f_{Y_{t+h}|X_t}(Y_{t+h}) \tag{15}$$

where $f_{Y_{t+h}|X_t}(y)$ is the predicted probability density function (PDF). A **lower score** indicates **better predictive performance**, as it means the model assigned higher probability to the observed outcome.

Provide the intuition about the scoring rule using the example of Crump et al. (2024) with a normal predictive density

For a full probabilistic evaluation, we also consider the **continuous ranked probability score (CRPS)**:

$$\text{CRPS}(F, Y_{t+h}) = \int_{-\infty}^{\infty} (F_{Y_{t+h}|X_t}(y) - \mathbb{1}\{Y_{t+h} \leq y\})^2 dy \tag{16}$$

where $F_{Y_{t+h}|X_t}(y)$ is the forecasted cumulative distribution function (CDF), and $\mathbb{1}\{Y_{t+h} \leq y\}$ is the empirical step function indicating whether Y_{t+h} is less than or equal to y.

CRPS generalizes quantile loss by integrating over all quantiles to evaluate the full probabilistic forecast. **Lower CRPS values** indicate better forecast performance.

4.3.2 Out-of-Sample Calibration

Calibration of probabilistic forecasts evaluates whether predicted probabilities align with observed frequencies. A common method for assessing calibration is the Probability Integral Transform (PIT), defined as:

$$F_t(y_{t+h}) \tag{17}$$

where $F_t(y_{t+h})$ represents the cumulative predictive density function evaluated at the realized outcome y_{t+h}. If the predictive densities are well calibrated, the PIT values should be uniformly distributed.

Figure 16: Add description

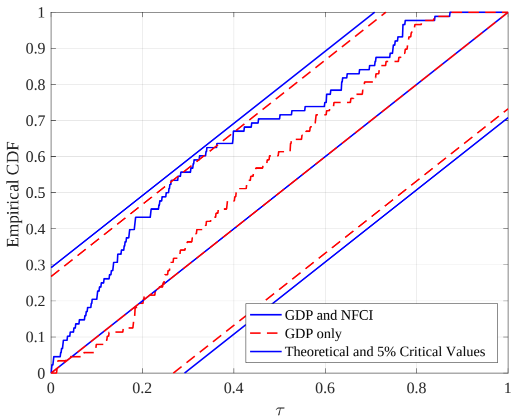

To empirically assess calibration, we compute the empirical cumulative distribution of the PITs:

$$\frac{1}{T_1 - T_0 + 1} \sum_{T_0 - h}^{T_1 - h} \mathbb{1}\left\{F_t(y_{t+h}) < r\right\} = \widehat{Prob}[y_{t+h} < r] \tag{18}$$

A well-calibrated forecast should produce an empirical CDF that aligns with the 45-degree line, ensuring that the proportion of realizations falling below a given quantile matches the predicted probability.

To build intuition, consider calibration in a binary prediction setting. Suppose a model forecasts recessions over the last 50 years. If the model predicts recessions 50% of the time, but historical data shows recessions occur in only 10% of periods, the model is poorly calibrated. Calibration does not measure the accuracy of individual forecasts but instead checks whether the average predicted frequency matches reality.

For continuous forecasts, the same principle applies to quantiles of a predictive distribution. Suppose a model estimates that GDP growth will be below a threshold (e.g., 5% value-at-risk) with 5% probability. If well-calibrated, the realized GDP should fall below this threshold exactly 5% of the time. More generally, for each quantile τ, we compute:

$$\hat{\Phi}(r) = \frac{1}{T_1 - T_0 + 1} \sum_{T_1 - h}^{T_0 - h} \left[\mathbb{1}\left\{F_t(y_{t+h}) < r\right\} - r\right] = \widehat{Prob}[y_{t+h} < r] - r \tag{19}$$

When the density forecasts are estimated using a rolling window and regularity conditions hold, we obtain the asymptotic convergence:

$$\hat{\Psi}(r) \Rightarrow \Psi(r) \quad \text{as} \quad P \to \infty \tag{20}$$

where $\Psi(r)$ follows a Gaussian process:

$$\Psi(r) := \sqrt{P}\hat{\Phi}(r), \quad P = T_1 - T_0 + 1. \tag{21}$$

If the forecast errors are uncorrelated (valid for $h = 1$), the autocovariance of $\Psi(r)$ simplifies to:

$$\mathrm{E}\left[\Psi(r_1)\Psi(r_2)\right] = \min\{r_1, r_2\} - r_1 r_2. \tag{22}$$

For autocorrelated residuals, long-run variance adjustments are necessary. The confidence bands in empirical applications account for this, though they remain indicative due to the use of real-time evaluation with an expanding sample.

Calibration ensures that forecasted probabilities match realized outcomes over time but does not assess forecast precision. This distinction is related to sharpness, which reflects the concentration of the predictive density. A model can be perfectly calibrated while providing little information about specific outcomes.

Consider an example where a naive model predicts recessions using a random card draw, where a recession is declared if a specific card is drawn (e.g., one out of 10 possible cards). Since recessions occur roughly 10% of the time, this model is perfectly calibrated but completely uninformative. In contrast, a sharper forecast assigns higher probability mass to the most likely outcomes.

A similar analogy applies to weather forecasting. Suppose one model predicts rain in Belgium with 70% probability every day. While well-calibrated, this forecast lacks precision. A more informative model conditions predictions on the season, forecasting 90% rain probability in winter and 30% in summer, leading to sharper predictions.

Applying this to macroeconomic forecasting, models incorporating financial conditions tend to be sharper because they differentiate stable and stressed environments. While they may slightly reduce calibration, they increase forecast accuracy by assigning higher probability to likely events.

Calibration is inherently an in-sample property. If assessed within the training sample, quantile regression will always produce a 45-degree line due to the properties of quantile estimation—5% of observations will always fall below the 5th percentile, 50% below the median, and so forth. However, meaningful forecast evaluation requires out-of-sample calibration. If a model remains well-calibrated when applied to new data, it indicates strong generalization.

Both calibration and sharpness are necessary for effective forecasting. Calibration ensures probability estimates are unbiased over time, while sharpness improves risk differentiation, making predictions more actionable. Financial conditions enhance sharpness, leading to more reliable forecasts, even if minor deviations from perfect calibration occur.

In the calibration report also the test based on the moments when you apply the normal CDF to the PIT, as in D'Agostino et al. (2013)

5 Alternative Models of Risk

The idea here is that we show how to use each method to estimate $f(Y|X)$ and compare it to the estimation with quantile regressions. For each approach, we compute the "moments" of interest from the density.

We need to introduce the section by emphasizing that there are several alternative methods to predict growth. Here we show why and how they are all connected. Each method provides a different perspective on forecasting at risk. Most importantly, all methods deliver the same broad conclusion that growth is vulnerable and provide different insights on why this is the case. Also, we will focus on univariate predictions for expositional simplicity. Several models can be generalized to multivariate, with all variables being endogenously determined, hence generating risk endogenously. We will not cover the details of these multivariate models. However, the single equation will provide the general intuition.

Imagine a line that goes from "more restrictions" to "less restrictions" on which we rank the various approaches:

Mean-Var Model ←——— MS Model ——— QR & Skewed-t ——— Kernel-Based →

more restricted **less restricted**

5.1 Mean-variance model

- We estimate it both with MLE and as a special case of the skewed-t (by allowing only location-scale)
- This will show that the mean-variance model is a more "restricted" approach than our baseline
- Key references:
  - Adrian et al. (2019) showed the equivalence with quantile regression
  - Adrian et al. (2023) use the location scale model with consumption growth and financial conditions to estimate the price of risk
  - Furceri et al. (2024) uses it in a panel setting
  - Carriero et al. (2024) extended it to a multivariate model (VAR)
  - Caldara et al. (2024) extend it to a multivariate model (factor model)

5.2 Distributional regressions and kernel-based methods

Here we focus on non-parametric conditional CDF estimation, multi-modality. We also talk of logit/probit.

5.2.1 Non parametric

References: Adrian et al. (2021) (multi-modality, multivariate extension with endogenous determination of financial and real risk), Li et al. (2013),Adrian et al. (2019) (non-parametric delivers similar results), Fernández-Villaverde et al. (2019) (uses Adrian et al. (2019) to validate structural models) Boyarchenko et al. (2024), Azzalini (1981)

5.2.2 Distributional regressions

References: Foresi and Peracchi (1995), Furno and Giannone (2024),Boyarchenko et al. (2020)

5.3 Markov-switching model

- We need a simplified version that works well enough to produce the results of interest

6 Panel

Quantile regression: Adrian et al. (2022), Gächter et al. (2023)
Location scale: Furceri et al. (2024)

7 Dynamic Machine Learning and Big Data (TBD)

The idea here is to estimate models of risk using ML methods and a large set of predictors (Big Data). There is almost no literature here, and we will need to do some work. Instead of simply summarizing these papers, we should show how to use ML approaches to estimate models of risk. A simple way to do so, is to set up the problem and then use ML/DNN/DL algorithms that are widely available on the web.

7.1 Compressing predictors

Most of the literature is already using a big data approach. In fact, the financial condition indexes that most of the literature uses are the result of compression of many financial indicators. The NFCI discussed used throughout the paper is a common factor extracted by several financial indicators using the dynamic factor model of Doz et al. (2012). For a detailed analysis, see Amburgey and McCracken (2023b).

7.2 Quantile Factor Models

Chen et al. (2021)

7.3 Pooling predictions

One approach to handle big data is to predict risk with one regressor at the time and then pool the predictions. Since the weights are assumed to be positive and sum to one, the resulting predictions are sparse and stable and hence can hangle big data (Brodie et al., 2009; Conflitti et al., 2015). For applications to forecasting risk, see Crump et al. (2024); Hengge (2019); Furceri et al. (2024)

7.4 Deep Neural Networks

- Fan et al. (2024): Growth-at-risk with Deep Neural Networks for the US.
- International Monetary Fund (2024): Growth-at-risk with Deep Neural Networks for the World.

7.5 Random Forest

- Lenza et al. (2023): Inflation-at-risk using Random Forests in the Euro Area
- Breiman (2001),Meinshausen (2006): Regression Forest, for point and density forecast
- Furceri et al. (2024): Regression Forest, for Debt at Risk
- Medeiros et al. (2021): application of random forest for US inflation point forecast

7.6 Combinations

Crump et al. (2024), Furceri et al. (2024): combination of density forecast in handle high dimension.

7.7 Non parametric Machine Learning

Boyarchenko et al. (2024)

8 Causal analysis

Furceri et al. (2025)

Loria et al. (2025)

9 Term structure of macro risk

Adrian et al. (2022)

10 Additional Applications (TBD)

The additional application section is articulates along two main directions: (i) we replicate the results of the previous section for other geographical areas, in particular for the Euro Area, and (ii) we explore additional applications related to macro risk. With respect to the latter point, below we provide a list of candidate extension, with key references. We replicate all the figures we produced for GDP at risk in the US, both I the intro and in the main section. However, we will report only selected charts in the text while the rest will be put in appendix.

References: Carriero et al. (2016), Chavleishvili and Manganelli (2024), Adams et al. (2021), Goulet Coulombe et al. (2022), Kiley (2022), Amburgey and McCracken (2023a), Botelho et al. (2023), Boyarchenko et al. (2023), Chavleishvili et al. (2023) and Chavleishvili and Kremer (2023)

10.1 Capital flows at risk

Gelos et al. (2022) Eguren-Martin et al. (2021)

10.2 Inflation at risk

López-Salido and Loria (2024), Korobilis et al. (2021). In what follow, there is a list of applications.

10.3 Debt at risk

Furceri et al. (2024)

10.4 Valuations at risk

Gneiting and Raftery (2007), Rossi and Sekhposyan (2015), D'Agostino et al. (2013),Diebold and Mariano (2002), Giacomini and White (2006), Amisano and Giacomini (2007)

10.5 House Price at Risk

Álvaro Menéndez Pujadas et al. (2022); Hafemann (2023)

10.6 Recession risk

Furno and Giannone (2024)

10.7 Alternative predictors

- Components of financial condition indexes Adrian et al. (2019), Gächter et al. (2023)
- Uncertainty and risk:Hengge (2019), Caldara and Iacoviello (2022), Keijsers and van Dijk (2024).
- macroeconomic predictors Plagborg-Møller et al. (2020)
- Text as data:

  Sharpe et al. (2023)

11 Conclusions

References

Adams, P. A., T. Adrian, N. Boyarchenko, and D. Giannone (2021). Forecasting macroeconomic risks. *International Journal of Forecasting 37*(3), 1173–1191.

Adrian, T., N. Boyarchenko, and D. Giannone (2019). Vulnerable growth. *The American Economic Review 109*(4), 1263–1289.

Adrian, T., N. Boyarchenko, and D. Giannone (2021). Multimodality in macrofinancial dynamics. *International Economic Review 62*(2), 861–886.

Adrian, T., F. Duarte, and T. Iyer (2023). The market price of risk and macro-financial dynamics. *IMF Working Paper* (WP/23/199).

Adrian, T., F. Grinberg, N. Liang, S. Malik, and J. Yu (2022). The term structure of growth-at-risk. *American Economic Journal: Macroeconomics 14*(3), 283–323.

Amburgey, A. and M. W. McCracken (2023a, August). Growth-at-Risk is Investment-at-Risk. Working Papers 2023-020, Federal Reserve Bank of St. Louis.

Amburgey, A. J. and M. W. McCracken (2023b). On the real-time predictive content of financial condition indices for growth. *Journal of Applied Econometrics 38*(2), 137–163.

Amisano, G. and R. Giacomini (2007). Comparing density forecasts via weighted likelihood ratio tests. *Journal of Business & Economic Statistics 25*(2), 177–190.

Azzalini, A. (1981). A note on the estimation of a distribution function and quantiles by a kernel method. *Biometrika 68*(1), 326–328.

Azzalini, A. and A. Capitanio (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*(2), 367–389.

Bernanke, B. S. (1983). Irreversibility, uncertainty, and cyclical investment. *The Quarterly Journal of Economics 98*(1), 85–106.

Botelho, V., C. Foroni, and A. Renzetti (2023, August). Labour at risk. Working Paper Series 2840, European Central Bank.

Boyarchenko, N. (forthcoming). Outlook-at-risk. In D. Aikman and P. Gai (Eds.), *The Research Handbook of Macroprudential Policy*. Edward Elgar Publishing. Forthcoming.

Boyarchenko, N., R. K. Crump, K. Dogra, L. Elias, and I. Lopez Gaffney (2024, May). The nonlinear case against leaning against the wind. Staff Reports 1100, Federal Reserve Bank of New York.

Boyarchenko, N., R. K. Crump, L. Elias, and I. L. Gaffney (2023, May). Look Out for Outlook-at-Risk. Liberty Street Economics 20230517, Federal Reserve Bank of New York.

Boyarchenko, N., D. Giannone, and A. Kovner (2020, November). Bank Capital and Real GDP Growth. Staff Reports 950, Federal Reserve Bank of New York.

Breiman, L. (2001). Random forests. *Machine learning 45*(1), 5–32.

Brodie, J., I. Daubechies, C. De Mol, D. Giannone, and I. Loris (2009). Sparse and stable Markowitz portfolios. *Proceedings of the National Academy of Sciences 106*(30), 12267–12272.

Caldara, D. and M. Iacoviello (2022). Measuring geopolitical risk. *American Economic Review 112*(4), 1194–1225.

Caldara, D., H. Mumtaz, and M. Zhong (2024, September). Risk in a data-rich model. *Mimeo*. Preliminary and Incomplete.

Carriero, A., T. E. Clark, and M. Marcellino (2016, July). Common Drifting Volatility in Large Bayesian VARs. *Journal of Business & Economic Statistics 34*(3), 375–390.

Carriero, A., T. E. Clark, and M. Marcellino (2024, August). Capturing Macro-Economic Tail Risks with Bayesian Vector Autoregressions. *Journal of Money, Credit and Banking 56*(5), 1099–1127.

Chavleishvili, S. and M. Kremer (2023, August). Measuring systemic financial stress and its risks for growth. Working Paper Series 2842, European Central Bank.

Chavleishvili, S., M. Kremer, and F. Lund-Thomsen (2023, July). Quantifying financial stability trade-offs for monetary policy: a quantile VAR approach. Working Paper Series 2833, European Central Bank.

Chavleishvili, S. and S. Manganelli (2024). Forecasting and stress testing with quantile vector autoregression. *Journal of Applied Econometrics 39*(1), 66–85.

Chen, L., J. J. Dolado, and J. Gonzalo (2021). Quantile factor models. *Econometrica 89*(2), 875–910.

Chernis, T., P. J. Coe, and S. P. Vahey (2023). Reassessing the dependence between economic growth and financial conditions since 1973. *Journal of Applied Econometrics 38*(2), 260–267.

Coe, P. J. and S. P. Vahey (2020, April). Financial conditions and the risks to economic growth in the United States since 1875. CAMA Working Papers 2020-36, Centre for Applied Macroeconomic Analysis, Crawford School of Public Policy, The Australian National University.

Conflitti, C., C. De Mol, and D. Giannone (2015). Optimal combination of survey forecasts. *International Journal of Forecasting 31*(4), 1096–1103.

Corradi, V., J. Fosten, and D. Gutknecht (2023). Out-of-sample tests for conditional quantile coverage: An application to growth-at-risk. *Journal of Econometrics 236*(2), 758–777.

Corradi, V. and J. Llorens-Terrazas (2024). Monitoring joint tail risks: An application to growth and inflation. *SSRN Electronic Journal*.

Crump, R. K., M. Everaert, D. Giannone, and C. S. Hundtofte (2024). *Changing Risk-Return Profiles*, pp. 283–302. Cham: Springer Nature Switzerland.

D'Agostino, A., L. Gambetti, and D. Giannone (2013). Macroeconomic forecasting and structural change. *Journal of Applied Econometrics 28*(1), 82–101.

Diebold, F. X. and R. S. Mariano (2002). Comparing predictive accuracy. *Journal of Business & Economic Statistics 20*(1), 134–144.

Diks, C., V. Panchenko, and D. van Dijk (2011). Likelihood-based scoring rules for comparing density forecasts in tails. *Journal of Econometrics 163*(2), 215–230.

Doz, C., D. Giannone, and L. Reichlin (2012). A Quasi–Maximum Likelihood Approach for Large, Approximate Dynamic Factor Models. *The Review of Economics and Statistics 94*(4), 1014–1024.

Eguren-Martin, F., C. O'Neill, A. Sokol, and L. von dem Berge (2021). Capital flows-at-risk: push, pull and the role of policy. *ECB Working Paper Series* (2538).

Fan, T.-H., Z.-X. K. Wang, R. Ghigliazza, D. Giannone, and T. C. Chi (2024, October). Forecasting macroeconomic risk: A machine learning approach. *Mimeo*. Presented at the 18th International Joint Conference on Computational and Financial Econometrics.

Fernández-Villaverde, J., F. Mandelman, Y. Yu, and F. Zanetti (2019, August). Search complementarities, aggregate fluctuations, and fiscal policy. Working Paper 26210, National Bureau of Economic Research. Revised April 2022.

Figueres, J. M. and M. Jarociński (2020). Vulnerable growth in the euro area: Measuring the financial conditions. *Economics Letters 191*(C).

Foresi, S. and F. Peracchi (1995). The conditional distribution of excess returns: An empirical analysis. *Journal of the American Statistical Association 90*(430), 451–466.

Furceri, D., D. Giannone, F. Kisat, and W. R. Lam (2024, November). Global debt-at-risk. *IMF, Mimeo*. Preliminary: Not for Circulation.

Furceri, D., D. Giannone, P. Pizzuto, and F. Frangiamore (2025, February). Government spending multipliers in good and bad times revisited: An iv quantile regression approach. *Mimeo*. Preliminary draft.

Furno, F. and D. Giannone (2024). Nowcasting recession risk. *Research Methods and Applications on Macroeconomic Forecasting*.

Galvao, A. F. and J. Yoon (2024). Hac covariance matrix estimation in quantile regression. *Journal of the American Statistical Association 119*(547), 2305–2316.

Gelos, G., L. Gornicka, R. Koepke, R. Sahay, and S. Sgherri (2022). Capital flows at risk: Taming the ebbs and flows. *Journal of International Economics 134*, 103555.

Giacomini, R. and H. White (2006). Tests of conditional predictive ability. *Econometrica 74*(6), 1545–1578.

Gneiting, T. and A. E. Raftery (2007). Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association 102*(477), 359–378.

Goulet Coulombe, P., M. Leroux, D. Stevanovic, and S. Surprenant (2022). How is machine learning useful for macroeconomic forecasting? *Journal of Applied Econometrics 37*(5), 920–964.

Gregory, K. B., S. N. Lahiri, and D. J. Nordman (2018). A smooth block bootstrap for quantile regression with time series. *The Annals of Statistics 46*(3), 1138–1166.

Gächter, M., M. Geiger, and E. Hasler (2023). On the Structural Determinants of Growth-at-Risk. *International Journal of Central Banking 19*(2), 251–293.

Hafemann, L. (2023). House prices at risk: A framework for assessing vulnerabilities in the german housing market. Technical paper, Deutsche Bundesbank.

Hengge, M. (2019). Uncertainty as a Predictor of Economic Activity. IHEID Working Papers 19-2019, Economics Section, The Graduate Institute of International Studies.

Hoga, Y. and C. Schulz (2025). Self-normalized inference in (quantile, expected shortfall) regressions for time series. *arXiv preprint arXiv:2502.10065*.

International Monetary Fund (2024, October). *Global Financial Stability Report: Steadying the Course: Uncertainty, Artificial Intelligence, and Financial Stability.* Washington, DC: International Monetary Fund.

Keijsers, B. and D. van Dijk (2024). Does economic uncertainty predict real activity in real time? *International Journal of Forecasting*.

Kiley, M. T. (2022, August). Unemployment Risk. *Journal of Money, Credit and Banking 54*(5), 1407–1424.

Knight, F. H. (1921). *Risk, Uncertainty and Profit.* Boston, MA: Houghton Mifflin Co.

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

Koenker, R. and G. Bassett (1978). Regression quantiles. *Econometrica 46*(1), 33–50.

Komunjer, I. (2005). Quasi-maximum likelihood estimation for conditional quantiles. *Journal of Econometrics 128*(1), 137–164.

Korobilis, D., B. Landau, A. Musso, and A. Phella (2021, October). The time-varying evolution of inflation risks. Working Paper Series 2600, European Central Bank.

Lenza, M., I. Moutachaker, and J. Paredes (2023). Density forecasts of inflation: A quantile regression forest approach. Working Paper Series 2830, European Central Bank.

Li, Q., J. Lin, and J. S. Racine (2013). Optimal Bandwidth Selection for Nonparametric Conditional Distribution and Quantile Functions. *Journal of Business and Economic Statistics 31*(1), 57–65.

Loria, F., C. Matthes, and D. Zhang (2025). Assessing Macroeconomic Tail Risk. *The Economic Journal 135*(665), 264–284.

López-Salido, D. and F. Loria (2024). Inflation at risk. *Journal of Monetary Economics 145*, 103570. Inflation: Expectations & Dynamics October 14-15, 2022.

Medeiros, M. C., G. F. R. Vasconcelos, Álvaro Veiga, and E. Zilberman (2021). Forecasting inflation in a data-rich environment: The benefits of machine learning methods. *Journal of Business & Economic Statistics 39*(1), 98–119.

Meinshausen, N. (2006). Quantile regression forests. *Journal of Machine Learning Research 7*, 983–999.

Newey, W. K. and K. D. West (1987). A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica 55*(3), 703–708.

Plagborg-Møller, M., L. Reichlin, G. Ricco, and T. Hasenzagl (2020). When is growth at risk? *Brookings Papers on Economic Activity 2020*(1), 167–229.

Rossi, B. and T. Sekhposyan (2015, May). Macroeconomic uncertainty indices based on nowcast and forecast error distributions. *American Economic Review 105*(5), 650–55.

Scotti, C. (2023). Financial shocks in an uncertain economy. Presented at the ECB Forum on Central Banking 2023.

Sharpe, S. A., N. R. Sinha, and C. A. Hollrah (2023). The power of narrative sentiment in economic forecasts. *International Journal of Forecasting 39*(3), 1097–1121.

Silverman, B. W. and G. A. Young (1987). The bootstrap: To smooth or not to smooth? *Biometrika 74*(3), 469–479.

White, H. (1980). A heteroskedasticity-consistent covariance matrix estimator and a direct test for heteroskedasticity. *Econometrica 48*(4), 817–838.

Álvaro Menéndez Pujadas, L. Hospido, and J. M. Montero (2022). House prices at risk: An empirical approach to downside risks in the spanish housing market. Working Paper 2244, Banco de España.