Federal Reserve Bank of New York
Staff Reports

Vulnerable Growth

Tobias Adrian
Nina Boyarchenko
Domenico Giannone

Staff Report No. 794
September 2016

Vulnerable Growth

Tobias Adrian, Nina Boyarchenko, and Domenico Giannone

Federal Reserve Bank of New York Staff Reports, no. 794
September 2016
JEL classification: C22, E17, E37

Abstract

We study the conditional distribution of GDP growth as a function of economic and financial conditions. Deteriorating financial conditions are associated with an increase in conditional volatility and a decline in the conditional mean of GDP growth, leading to a highly skewed distribution. The lower quantiles of GDP growth exhibit strong variation as a function of financial conditions, while the upper quantiles are stable over time. Although measures of financial conditions have significant influence in forecasts of downside vulnerability, measures of economic conditions have significant predictive power only for the median of the distribution. These findings are robust both in and out of sample and to the inclusion of different measures of financial conditions. We quantify GDP vulnerability as relative entropy between the empirical conditional and unconditional distribution. We show that this measure of vulnerability is highly asymmetric: The contribution to the total relative entropy of the probability mass below the median of the conditional distribution is larger and more volatile than the contribution of the probability mass above the median. The asymmetric response of the distribution of GDP growth to financial and economic conditions—with adverse financial conditions increasing downside vulnerability of growth but not the median forecast—is challenging for standard models of the macroeconomy. We argue that the inclusion of a financial sector is crucial for generating the observed dynamics of growth vulnerability.

Keywords: Downside risk, entropy, quantile regressions

Adrian, Boyarchenko, Giannone: Federal Reserve Bank of New York (e-mails: tobias.adrian@ny.frb.org, nina.boyarchenko@ny.frb.org, domenico.giannone@ny.frb.org). The authors thank Richard Crump, Robert Engle, Eric Ghysels, James Hamilton, Lawrence Schmidt, Erik Vogt, and Jonathan Wright for helpful comments. The views expressed in this paper are those of the authors and do not necessarily reflect the position of the Federal Reserve Bank of New York or the Federal Reserve System.

1 Introduction

Economic forecasts usually provide point estimates for the conditional mean of GDP growth and other economic variables. Such point forecasts, however, ignore risks around the central forecast and, as such, may paint an overly optimistic picture of the state of the economy. In fact, policy makers' focus on downside risk has increased in recent years. In the U. S., the Federal Open Market Committee (FOMC) commonly discusses downside risks to growth in FOMC statements, with the relative prominence of this discussion fluctuating with economic conditions. At the same time, surveys of market participants (the Federal Reserve Bank of New York's Primary Dealer Survey), economists (the Blue Chip Economic Survey) and professional forecasters (the Federal Reserve Bank of Philadelphia's Survey of Professional Forecasters) all collect the respondents' beliefs of the probability distribution around the point forecast.

In this paper, we model empirically the full distribution of future real GDP growth as a function of financial and economic conditions. We estimate the distribution non-parametrically using quantile regressions. The estimated lower quantiles of the distribution of GDP growth exhibit strong variation as a function of financial conditions, while the upper quantiles are stable over time. Moreover, we show that economic conditions forecast the median of the distribution but do not contain information about the other quantiles of the distribution.

Next, we smooth the estimated quantile distribution every quarter by interpolating between the estimated quantiles using the skewed t-distribution. This allows us to transform the empirical quantile distribution into an estimated distribution of GDP growth, plotted in Figure 1. Two features are striking about the estimated distribution. First, the entire distribution – and not just the first two moments – evolves over time. For example, business cycle peaks are associated with right-skewed predicted distributions while troughs are associated with left-skewed distributions. Second, the probability distributions inherit the stability of

Figure 1. Distribution of GDP growth over time. One-quarter-ahead predicted distribution of real GDP growth, based on quantile regressions with CFNAI and NFCI as conditioning variables.

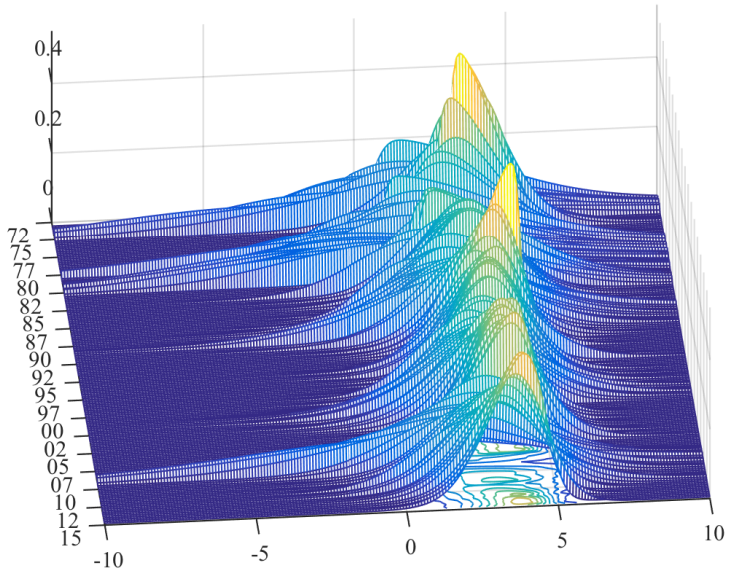

the right tail of the distribution from the estimated quantile distribution, while the median and the left tail of the distribution exhibit strong time series variation. This asymmetry in the evolution of the conditional GDP distribution suggests that downside risk to GDP varies much more strongly over time than upside risk.

We summarize the properties of the upside and downside risks to real GDP growth using two metrics: the upside and downside entropy of the unconditional distribution of GDP growth relative to the empirical conditional distribution, and by the expected shortfall and its upper tail counterpart. While upside entropy comoves with its downside counterpart, downside entropy is much more volatile. This asymmetry echoes our finding that the elasticity of GDP growth to financial conditions is significantly higher for the lower quantiles of the distribution than for the upper quantiles.

We perform many robustness tests to our findings. First, we present alternative measures of financial conditions, focusing on variables that have been emphasized in the recent macro finance literature such as credit spreads, the term spread, and equity volatility. Second, we

show that out-of-sample estimates of the conditional GDP distribution are very similar to the in-sample distribution. This leads the out-of-sample estimates of growth vulnerability to likewise be similar to the in-sample estimates. Third, we demonstrate that the strongly skewed conditional distribution of GDP growth is not an artifact of our two-step quantile regression estimation procedure but also arise when we estimate a simple conditional heteroskedasticity model for GDP growth and volatility using maximum likelihood.

A large literature has documented the decline of GDP volatility prior to the financial crisis of 2008 (see e.g. McConnell and Perez-Quiros, 2000; Blanchard and Simon, 2001; Bernanke, 2004; Giannone, Lenza, and Reichlin, 2008). While McConnell and Perez-Quiros (2000) argue that a structural break of GDP volatility occurred in 1984, Blanchard and Simon (2001) show evidence consistent with a slow decline in volatility over the post war period. In contrast to that influential literature, we focus not just on the second moment of GDP growth, but rather on the whole conditional distribution of GDP growth. Our striking finding is that GDP growth volatility is nearly entirely driven by the left side of the conditional distribution, thus providing testable implications for theoretical research. In fact, we can attribute the decline of volatility during the Great Moderation period to a decline in the downside risk to GDP growth.

From an econometric point of view, our paper is related to the statistical literature on evaluating conditional distributions. We develop a straightforward two-step procedure for the estimation of the conditional probability distribution function. In the first step, we employ quantile regressions of Koenker and Bassett (1978) to estimate the conditional quantile function of GDP growth as a function of lagged financial and economic conditioning variables (see Ghysels (2014) for a recent application of quantile regressions to the estimation of skewness). In the second step, we fit a parametric inverse cumulative distribution function with a known density function to the empirical conditional quantile function, for each quarter in the sample. That procedure is computationally straightforward, and allows us to transform the inverse cumulative distribution function from the quantile regression into

a density function. We then use an entropy metric to measure deviation of the conditional GDP distribution from the conditional gaussian distribution. We label the downside entropy – the contribution to relative entropy of the density below the median of the distribution – "growth vulnerability".

Our approach differs from the recent literature that has analyzed GDP uncertainty in its focus on the preeminent role for downside risk, rather than symmetric measures of risk. Baker, Bloom, and Davis (2013) propose a measure of political uncertainty based on news announcements and Jurado, Ludvigson, and Ng (2015) compute conditional volatility from a large number of macroeconomic variables. Bloom (2009) models uncertainty shocks in a macroeconomic context. The main difference of that strand of literature to our work is the emphasis on shocks that move the conditional variance. In contrast, our empirical distributions allow for differences in upside and downside volatility.

Our findings are closely related to the recent macro-finance literature that emphasizes the link between financial stability and macroeconomic performance. For example, the buildup of leverage and maturity transformation can give rise to financial vulnerability that increases the downside risk to GDP growth. Similarly, external imbalances can make economies more vulnerable to sudden stops, with potentially adverse consequences to real GDP growth. Brunnermeier and Sannikov (2014), He and Krishnamurthy (2012) and Adrian and Boyarchenko (2012) present macroeconomic production economies with financial intermediaries that give rise to time variation GDP vulnerability as a function of financial conditions. Though our baseline results do not distinguish between sources of financial instability, distributions predicted using equity implied volatility most resemble the baseline, suggesting that macrofinancial models that feature volatility-linked constraints may be most promising in generating the empirical link between financial conditions and growth vulnerability.

The rest of the paper is organized as follows. Section 2 presents the measures of economic and financial conditions, and relates them to GDP growth in a descriptive fashion. Section 3 presents our estimates of the conditional GDP distribution, and introduces the concept of

GDP vulnerability. Section 4 relates our estimates of vulnerability to alternative financial and economic indicators. Section 5 discusses out of sample results and alternative econometric approaches. Section 6 discusses implications of our findings for macroeconomic theories. Section 7 concludes.

2 Economic and Financial Conditions and GDP Growth

To gauge economic and financial conditions, we use the Chicago Fed National Activity Index (CFNAI) and the National Financial Conditions Index (NFCI).[1] The CFNAI is a monthly index designed to measure overall economic activity. It is a weighted average of 85 existing monthly indicators of national economic activity, normalized to have an average value of zero and a standard deviation of one, with positive realizations corresponding to growth above trend.[2] The economic indicators that are included in the CFNAI are drawn from four broad categories of data including production and income; employment, unemployment, and hours; personal consumption and housing; and sales, orders, and inventories. The derived index provides a single, summary measure of a factor common to these national economic data. The methodology for the CFNAI was developed by Stock and Watson (1999).

Similarly, the NFCI provides a weekly estimate of U.S. financial conditions in money markets, debt and equity markets, and the traditional and shadow banking systems. The index is a weighted average of 105 measures of financial activity, each expressed relative to their sample averages and scaled by their sample standard deviations.[3] The methodology for the NFCI is described in Brave and Butters (2012). The data for the NFCI starts in January 1973, which we use as starting point for our empirical investigation.

Figure 2 shows the times series of CFNAI and NFCI, as well as the quantile-quantile (QQ) plots of CFNAI and NFCI relative to one-step-ahead GDP growth. The QQ-plots

[1] The CFNAI and the NFCI are computed by the Federal Reserve Bank of Chicago, and available here and here, respectively.

[2] The list of 85 indicators is provided here.

[3] The list of indicators is provided here.

show the empirical quantiles of GDP growth on the x-axis against the empirical quantiles of either CFNAI or NFCI on the y-axis. While the relationship of the CFNAI with GDP growth is largely linear, the relationship of NFCI with GDP growth exhibits very pronounced nonlinearity. The nonlinearity indicates differences in the conditional distribution functions and foreshadow our findings that financial conditions are a more important indicator of growth vulnerability than economic conditions.

Our estimates of the conditional GDP distribution rely on quantile regressions. In a quantile regression of y_t on x_t (which includes $k-1$ conditioning variables and a constant), the regression slope β_τ is chosen to minimize the quantile weighted absolute value of errors:

$$\beta_\tau = \underset{\beta \in \mathbb{R}^k}{\operatorname{argmin}} \sum_{t=1}^{T} \left(\tau \cdot \mathbb{1}_{(y_t \geq x_t\beta)} \left| y_t - x_t\beta_\tau \right| + (1-\tau) \cdot \mathbb{1}_{(y_t < x_t\beta)} \left| y_t - x_t\beta_\tau \right| \right) \tag{1}$$

where $\mathbb{1}_{(\cdot)}$ denotes the indicator function. The predicted value from that regression is the quantile of y_t conditional on x_t

$$Q_{y_t|x_t}(\tau) = x_t\beta_\tau. \tag{2}$$

Koenker and Bassett (1978) show that $Q_{y_t|x_t}(\tau)$ is a consistent linear estimator of the quantile function of y_t conditional on x_t. The quantile regression differs from an ordinary least squares regression in two respects. First, the quantile regression minimizes the sum of absolute errors, rather than the sum of squared errors. Second, it puts differential weights on the errors depending on whether an error term is above or below the quantile.

Figure 3 shows the scatter plot of one quarter ahead and four quarters ahead GDP growth against the NFCI and the CFNAI, together with the quantile regression lines for the 5th, 50th and 95th quantiles and the OLS regression line. For the NFCI, the slopes differ significantly across quantiles and from the OLS regression line, indicating that the sensitivity of financial conditions to economic conditions evolves with the business cycle. Indeed, Figure 4 shows that, at the lower quantiles, the slope from the quantile regression are significantly different,

at the 10 percent level, from the OLS slope.[4] The regression slopes change dramatically for NFCI across the quantiles, but are stable for CFNAI. The slopes are most different from the linear benchmark for the lower quantiles. Importantly, the regression slopes for NFCI do not change significantly when the CFNAI is included as well in the regression, indicating that most of the explanatory power of future GDP vulnerability arises from the information content of financial conditions. For economic conditions, on the other hand, the quantile regression slopes are barely different from the linear regression slopes, indicating that economic conditions are less informative for tail risks.

Figure 5 shows one and four quarter GDP growth together with its conditional median and its conditional 5, 25, 75 and 95 percent quantiles. This figure demonstrates the main result of the paper: the asymmetry between the upper and lower conditional quantiles. While the lower quantiles vary significantly over time, the upper quantiles are stable. Figure 6 shows that the median and interquartile range are strongly negatively correlated. When financial conditions deteriorate, the interquartile range increases, while the median falls (Figure 6a and 6b). That implies that the lower quantile falls, as illustrated in Figure 6c and 6d: when the interquartile range increases, the 5th quantile falls. For upper quantiles, on the other hand, the movement in the median and the interquartile range are offsetting, thus changes in financial conditions have relatively little impact on those upper quantiles, as is visible in Figure 5. This strong asymmetry of the quantiles of GDP growth is striking. We study the differences between the upper and lower parts of the predicted GDP growth distribution more formally in the next section.

[4] The significance is computed using the distribution of coefficients estimated in 1000 bootstrapped samples generated using a vector autoregression (VAR) with 4 lags and constant fitted to the full-sample evolution of CFNAI, NFCI and real GDP growth.

3 Quantifying GDP Vulnerability

3.1 The Conditional GDP Distribution

While the quantile regression (2) provides us with an estimate of the quantile function, an inverse cumulative distribution function, in practice the estimates from the quantile regressions are noisy. We fit the skewed t-distribution developed by Azzalini and Capitanio (2003) in order to smooth the quantile function and recover a probability density function:[5]

$$f\left(y;\mu,\sigma,\alpha,\nu\right)=\frac{2}{\sigma}t\left(\frac{y-\mu}{\sigma};\nu\right)T\left(\alpha\frac{y-\mu}{\sigma}\sqrt{\frac{\nu+1}{\nu+\frac{y-\mu}{\sigma}}};\nu+1\right) \tag{3}$$

where $t(\cdot)$ and $T(\cdot)$ respectively denote the PDF and CDF of the skewed t-distribution. The four parameters of the distribution pin down the location μ, scale σ, skewness ν, and shape α. Relative to the t-distribution, the skewed t-distribution adds the shape variable which regulates the impact of degree of magnification of the T distribution on the t distribution. The skewed t-distribution is part of a general class of mixed distributions proposed by Azzalini (1985) and further developed by Azzalini and Dalla Valle (1996). The intuition for the derivation is that a base probability distribution—in this case $t\left(\frac{y-\mu}{\sigma};\nu\right)$—gets shaped by its cumulative distribution function, and rescaled by shape parameter α. The notable special case of the skewed t-distribution is the t-distribution when $\alpha=0$. In addition to $\alpha=0$, when $\nu=\infty$, the distribution reduces to the a Gaussian with mean μ and standard deviation σ. Relative to the standard t-distribution with skewness parameter ν, the shape parameter changes the fatness of the distribution.

For each quarter, we choose the four parameters $\{\mu_t,\sigma_t,\alpha_t,\nu_t\}$ of the skewed t-distribution f to minimize the squared distance between the our estimated quantile function $Q_{y_t|x_t}\left(\tau\right)$ from (2) and the quantile function of the cumulative skewed t-distribution $Q_f\left(\tau;\mu_t,\sigma_t,\alpha_t,\nu_t\right)$

[5] We have also estimated alternative flexible distribution functions, but found that the skewed t-distribution gave the best fit.

from (3) to match the 5, 25, 75, and 95 quantiles

$$\{\mu_t, \sigma_t, \alpha_t, \nu_t\} = \underset{\mu_t,\sigma_t,\alpha_t,\nu_t}{\operatorname{argmin}} \sum_{\tau} \left(Q_{y_t|x_t}(\tau) - Q_f(\tau; \mu_t, \sigma_t, \alpha_t, \nu_t)\right)^2 \tag{4}$$

where $\mu_t \in \mathbb{R}$, $\sigma_t \in \mathbb{R}^+$, $\alpha_t \in \mathbb{R}$, and $\nu_t \in \mathbb{Z}^+$. This can be viewed as an exactly identified nonlinear cross sectional regression of the predicted quantile function $Q_{y_t|x_t}(\tau)$ on the skewed t-distribution $f(y_t; \mu_t, \sigma_t, \alpha_t, \nu_t)$ across quantiles.[6]

Figure 7 plots the unconditional quantile distribution, the conditional quantile distribution $Q_{y_t|x_t}(\tau)$ and the fitted inverse cumulative skewed t-distribution $Q_f(\tau; \mu_t, \sigma_t, \alpha_t, \nu_t)$ for three sample dates at different points of the business cycle: 2006Q2, which represented the end of the Federal Reserve's tightening cycle prior to the financial crisis; 2008Q4 when the zero lower bound was reached just after the failure of Lehman; and 2014Q4, which is the last in-sample date of our dataset. In all three cases, the skewed t-distribution is sufficiently flexible to smooth the estimated quantile function while passing through all four target quantiles for both the one-quarter-ahead and the four-quarters-ahead forecasts. Finally, Figure 7 shows that the conditional distribution can deviate substantially from the unconditional distribution. While the median of the unconditional and conditional distributions coincides in 2006Q2, the unconditional distribution exhibits both greater upside and downside risks. In 2008Q4, the entire conditional distribution is below the unconditional distribution; in 2014Q4, the upper tails of the conditional and unconditional distributions coincide but the unconditional distribution predicts a lower median growth and has greater downside risks.

3.2 Measuring Vulnerability

We now turn to studying the differences between the conditional and unconditional GDP growth distributions more formally. Denote by g the unconditional distribution of GDP

[6] An alternative approach would be to use the entires quantile function to pin down the parameters of f, and allow the parameters of the skewed t-distribution to be over-identified. We follow the more parsimonious exactly-identified approach here.

growth computed from all available data $\{y_1, y_2, ...y_T\}$. We define the upside, $\mathcal{L}_t^U$, and downside, $\mathcal{L}_t^D$, entropy of g relative to the estimated skewed t-distribution f_t as

$$\mathcal{L}_t^D\left(f_t; g\right) = \int_{-\infty}^{F^{-1}(50)} \left(\log g\left(y\right) - \log f_t\left(y\right)\right) f_t\left(y\right) dy, \tag{5}$$

$$\mathcal{L}_t^U\left(f_t; g\right) = \int_{F^{-1}(50)}^{\infty} \left(\log g\left(y\right) - \log f_t\left(y\right)\right) f_t\left(y\right) dy. \tag{6}$$

The upside and downside entropies sum to the total relative entropy, $\mathcal{L}_t = \mathcal{L}_t^U + \mathcal{L}_t^D$, which measures the Kullback-Leibler divergence between two distributions and is the only definition of difference between distribution that satisfies the canonical extension of the properties (continuity, symmetry, additivity and maximality) of the Shannon entropy of a distribution $\mathcal{E}_t\left(f_t\right) = \int_{\mathbb{R}} \log f_t\left(y\right) f_t\left(y\right) dy$. Of particular interest to us is the interpretation of entropy as the summary measure of higher moments of a distribution $\mathcal{E}_t\left(f_t\right) = \sum_{j=1}^{+\infty} \frac{\kappa_j(f_t)}{j!}$, where κ_j is the j cumulant of the distribution n_t. Relative entropy then measures the difference between higher order cumulants of the fitted and the benchmark distribution.

Figures 8a and 8b show the evolution of GDP upside and downside entropy, one and four quarters ahead, respectively. In addition, we plot the 5% expected shortfall and its upper tail counterpart, which we term the 95% expected longrise (Figures 8c and 8d), defined as

$$SF_t = \frac{1}{0.05}\int_0^{0.05} F_t^{-1}(\tau)d\tau; \qquad LR_t = \frac{1}{0.05}\int_{0.95}^{1} F_t^{-1}(\tau)d\tau.$$

The information content of the relative entropy and the expected shortfall measures are distinct. While expected shortfall integrates the probability density below the 5th quantile, relative downside entropy measures the divergence of the Gaussian density from the reference density for realization of the state falling below the median implied by the reference density. Despite these differences, the two measures exhibit a surprising degree of similarity, indicating that the main deviation from Gaussianity is due to the behavior in the very far tail of the distribution. It is also noteworthy that, while the upside and downside entropy measures do

comove, downside entropy is more volatile and has much more pronounced nonlinearities. Similarly, the expected shortfall and longrise measures are positively correlated but expected shorfall is significantly more volatile.

Comparison of the upside and downside entropy metrics in Figures 8 again shows that variation of the conditional distribution is primarily due to variation in the left half of the distribution: downside entropy varies widely, while upside entropy is more or less constant.

4 GDP Vulnerability and Other Financial Indicators

Our results so far have relied on the NFCI, a composite financial conditions indicator that relies on information of 105 measures from money markets, debt and equity markets and the traditional and shadow banking systems. In order to shed light on the importance of the contribution of individual series, we now investigate three financial indicators that are of particular interest: equity market volatility, the credit spread, and the term spread.

Equity market volatility has been shown to be a significant indicator for the price of risk. Rey (2015) shows that global capital flows, global credit growth, and global asset prices comove tightly with the VIX. Longstaff, Pan, Pedersen, and Singleton (2011) estimate that the price of sovereign risk is strongly correlated with the VIX. Furthermore, Adrian, Crump, and Vogt (2015) show that a nonlinear transformation of the VIX forecasts stock and bond returns, suggesting that the pricing of risk depends on the VIX. In general equilibrium, pricing of risk is associated with GDP growth, and risk to GDP growth. Hence we expect the VIX to be a significant forecasting variable for GDP vulnerability.

A recent literature has linked downside risk to GDP, particularly during financial crises, to credit conditions as measured by credit spreads. Gilchrist and Zakrajšek (2012) construct the excess bond premium, a residual credit spread orthogonal to firm specific information on defaults, and show that that premium has considerable predictive power for future real activity. Using U.S. data from 1929 to 2013, López-Salido, Stein, and Zakrajšek (2016) show

that elevated credit-market sentiment is associated with a decline in economic activity two and three years in the future, driven by mean reversion in credit spreads. Using a long time series across a panel of countries, Krishnamurthy and Muir (2016) document that the transition into a crisis occurs when credit spreads increase markedly, indicating that crises involve a dramatic shift in expectations and are a surprise.

An extensive literature has shown the forecasting power of the term spread for recession (see Estrella and Hardouvelis, 1991, and the subsequent literature). The term spread is shown to predict recessions 12-18 months in advance, both in sample and out of sample, and is generally a more powerful predictor of recessions than other variables. The term spread generally works best as individual predictor (see Estrella and Mishkin, 1998). Harvey (1988) shows that a consumption Euler equation naturally gives rise to forecasting of the term spread for real activity.

Figure 9 shows the quantile coefficients for equity market option-implied volatility, the BAA-AAA credit spread, and the 10-year/3-month term spread.[7] Comparing the loadings on NFCI to the loading on these three individual components, we see that the conditional quantile function is most sensitive to the NFCI, followed by option-implied volatility, term spread and the credit spread. The term spread has the curious property of having a non-monotonic relationship with respect to the upper quantiles for the four quarter ahead prediction. At intermediate quantiles, the conditional quantile function has almost constant loadings on volatility. At very low quantiles, however, the quantile function has a significant negative relationship with volatility: high option-implied volatility is associated with large downside risks to GDP growth at both one and four quarter horizons. The credit spread carries surprisingly little information, as indicated by a very flat quantile coefficient curve, which is close to zero. In sum, these findings suggest that the NFCI financial conditions index is a robust proxy for how financial conditions affect the predicted distribution for GDP growth.

[7] We use the VXO instead of the VIX as it has a slightly longer time series, and we backfill the data to 1973 using realized equity market volatility (see Bloom, 2009).

5 Out-of-sample Evidence and Alternative Econometrics

5.1 Out of Sample Evidence

In this section we evaluate the out-of-sample performances of the methods. We backtest the model by replicating the analysis that an economist would have obtained by using the proposed methodology in real time, with the caveat that we use final revised data only.[8]

We produce predictive distributions recursively for two horizons (1 and 4 quarters), starting with the estimation sample that ranges from 1973Q1 to 1992Q4. More precisely, using data from 1973Q1 to 1992Q4, we estimate the predictive distribution for 1993Q1 (one quarter ahead) and 1993Q4 (one year ahead). We then iterate the same procedure updating the estimation sample, one quarter at a time, until the end of the sample (2015Q4). At each iteration, we repeat the estimation steps of Sections 2 and 3, and estimate quantile regression, match the skewed t-distribution, and compute downside and upside entropy. The outcome of this procedure is a time-series of 20 years of density forecasts for each of the two forecast horizons.

We perform two types of out-of-sample analyses. First, we study the robustness of the results shown so far by comparing the in-sample measures of vulnerability with their realtime counterparts. Second, we evaluate the out-of-sample accuracy and calibration of the density forecasts by analyzing the predictive score and the probability integral transform (PIT), that is, the predictive density and cumulative distribution evaluated at the outturn, respectively.

Results for the first exercise are presented in Figure 10. We report selected quintiles and downside entropy computed using the full sample (in-sample) and recursively (out-of-sample). The figure illustrates that the in-sample and out-of-sample estimates of the quantiles are virtually indistinguishable. The similarities are more striking as the financial crisis of 2007-09 is a significant tail event that is not in the data when estimating the

[8] Real time data for CFNAI and NFCI are only available for the recent past.

out-of-sample quantiles. The stability of the recursive estimates thus shows that downside vulnerability can be detected in real time.

Turning to the evaluation of the reliability of the predictive distribution, we measure the accuracy of a density forecast using the predictive score. This is the predictive distribution generated by a model, evaluated at the realized value of the time series. Higher predictive scores indicate more accurate predictions as they show that outcomes that the model considers more likely are closer to the ex-post realization. Figures 11a and 11b plot the scores of the conditional predictive distribution together with the scores of the unconditional distribution.[9] The predictive score for the conditional distribution is frequently above that of the unconditional distribution. Thus, the conditional distribution is often more accurate – and rarely less accurate – than the unconditional one, and the information contained in the conditioning variables (metrics of financial and economic conditions) is a robust and genuine feature of the data.

We conclude the out-of-sample evaluation by analyzing the calibration of the predictive distribution. We compute the empirical cumulative distribution of the PITs, which measures the percentage of observations that are below any given quantile. The model is better calibrated the closer empirical cumulative distribution of the PITs is to the 45 degrees line. In a perfectly calibrated model the cumulative distribution of the PITs is a 45-degree line, so that the fraction of realizations below any given quantile $Q_{y_{t+h|x_t}}(\tau)$ of the predictive distribution is exactly equal to τ. Results a presented in Figures 11c and 11d for both the conditional and unconditional distribution. Following Rossi and Sekhposyan (2015), we report confidence bands around the 45-degree line to account for sample uncertainty.[10][11] For both the conditional and unconditional predictive distribution, the empirical distribution of

[9] Recall that the unconditional distribution only uses a constant as the explanatory variable.

[10] The confidence bands should be taken as a general guidance since they are derived for forecasts computed using a rolling scheme, i.e. with a constant length of the estimation sample, while we use a rolling scheme, with an expanding estimation window.

[11] The confidence bands for one year ahead are in principle different for the conditional and unconditional distribution since they are based on a HAC estimator to account for serial correlation. However, with our data they are undistinguishable.

the PITs is well within the confidence bands for the lower quantiles and the upper quantiles, though the empirical distribution falls outside the confidence bounds in the center of the distribution. Overall, Figure 11 illustrates that the quantile regression approach generates robust predictive distributions, and is able to capture downside vulnerabilities particularly well.

5.2 Alternative Econometric Approaches

We view our two-step estimation procedure of fitting quantile regressions in the time series, and then the distribution across quantiles, as a methodological contribution of the paper. However, the main results of the paper can also be estimated using more traditional maximum likelihood based approaches. Consider the following model

$$y_t = \gamma_0 + \gamma_1 x_{t-1} + \beta_1 y_{t-1} + \sigma_t \varepsilon_t; \qquad \sigma_t^2 = \exp\left(\delta_0 + \delta_1 x_{t-1}\right) + \delta_3 \varepsilon_{t-1}^2, \tag{7}$$

where $\varepsilon_t \sim N(0,1)$ and x_t is the realization of NFCI. The model is an ARCH(1) that features financial and economic conditions x_t as conditioning variables for both the time varying mean and time varying volatility (the ARCH model was introduced by Engle, 1982). The model is straightforward to estimate via maximum likelihood.

Figure 12 plots the conditional mean and the conditional lower and upper 5th quantiles for one-quarter-ahead GDP growth (Figure 12a) and one-quarter S&P 500 stock market index return (Figure 12b) implied by the model in equation (7). The simple conditional heteroskedasticity model is able to reproduce the strongly skewed conditional GDP distribution by simply shifting the mean and volatility of GDP as a function of economic and financial conditions. Thus, our main results are not an artifact of our two-step quantile regression estimation procedure but reflect rather a robust property of the conditional distribution of GDP growth.

In contrast, the one quarter ahead S&P 500 return distribution features very little skew-

ness and is surprisingly symmetric. The symmetry of the equity return conditional distribution varies, however, with horizon. Ghysels (2014) computes the four quarter ahead S&P500 return distribution and uncovers more negative skewness. In fact, four quarters ahead, the S&P500 return distribution looks strikingly similar to the one quarter ahead GDP distribution. Schmidt and Zhu (2016) develop an alternative methodology to smooth across quantiles and report a conditional S&P500 return distribution at a daily frequency that looks similar to our Figure 12.

6 The Economics of GDP Vulnerability

6.1 Mean, Volatility, and Vulnerability

We turn now to the relationship between the mean of the forecasted distribution of GDP growth and the volatility and vulnerability of GDP growth. Figure 13 shows that, contemporaneously, high realizations of the conditional mean of GDP growth are associated with both low volatility and low downside entropy. This negative relationship between the conditional mean and measures of growth vulnerability is present at both the one quarter and four quarter horizons. Economically, this finding implies that deteriorating financial conditions lead to lower average growth and an increase in volatility. Combined, these two effects generate a negatively skewed unconditional distribution, as downward shifts in the growth outlook are associated with an increase in risk.

Figure 14 shows that the negative relationship between the conditional mean and measures of vulnerability is present at multiple leads and lags of the vulnerability measures. In particular, the correlation between the conditional mean and volatility is statistically significantly negative for lags (-6, +4) of volatility and becomes statistically significantly positive only at lags beyond 12 quarters. Thus, while there is a positive risk-return tradeoff for GDP growth at low frequencies, the negative association between risk and return is a robust feature at higher frequencies. Figure 14 also shows that volatility leads the conditional mean by

one quarter, with a one standard deviation increase in the conditional volatility associated with a .8 standard deviations decrease in the next quarter's predicted conditional mean of GDP growth.

The result that the negative association between the conditional mean and the conditional variance generates the highly asymmetric shape of the conditional GDP distribution is also confirmed by our alternative econometric procedure using the ARCH model (see equation (7) and Figure 12). In that approach, shocks are Gaussian with a time varying mean and time varying variance, both of which are functions of economic and financial conditions. Despite the conditional Gaussian shocks, the model generates a highly skewed GDP growth distribution, as deteriorating financial conditions are associated with larger variance and lower growth. This finding therefore does not rely on our usage of quantile regressions and the fitting of the skewed t-distribution.

The negative association between the conditional mean of GDP growth and the conditional variance of GDP already appear in the VARs of Jurado et al. (2015) and Bloom (2009) who model macroeconomic uncertainty. Our contribution relative to that literature is our emphasis on the entire conditional distribution of GDP growth, which shows that the negative association of location and scale of the GDP distribution generates a highly skewed distribution. Furthermore, we emphasize the role of financial conditions in modeling tail risk, while Jurado et al. (2015) use metrics of macroeconomic uncertainty, Bloom (2009) uses equity market volatility, and Baker et al. (2013) construct political uncertainty.

6.2 Implications for Theories of Macroeconomic Dynamics

Our main finding is that downside risk to GDP growth is forecast by financial conditions, while upside risk remains more or less constant. These findings present challenges for macroeconomic modeling. On the one hand, they point to the importance of financial conditions over and above economic conditions: We do not uncover any evidence of economic conditions forecasting GDP vulnerability. This points towards an independent role of the fi-

nancial sector, over and above the impact of financial conditions on the dynamics of real variables. Brunnermeier and Sannikov (2014), He and Krishnamurthy (2012) and Adrian and Boyarchenko (2012) present macroeconomic production economies with financial intermediaries whose vulnerability gives rise to time variation in GDP vulnerability as a function of financial conditions. These approaches also emphasize the link between financial stability and macroeconomic performance.

An earlier literature emphasizes the credit channel of monetary policy (see the overviews by Bernanke and Gertler, 1995, Boivin, Kiley, Mishkin et al., 2010, and Brunnermeier, Eisenbach, and Sannikov, 2013). While the intermediary asset pricing approach features financial frictions in the supply of credit, the credit channel of monetary policy emphasizes primarily frictions in the demand for credit. The credit channel does generate asymmetries in the business cycle. However, in the most common DSGE implementations of the credit channel, models are only solved up to first order, so that financial vulnerability does not play a role in equilibrium (see, in particular, Bernanke, Gertler, and Gilchrist, 1999 and Del Negro, Eusepi, Giannoni, Sbordone, Tambalotti, Cocci, Hasegawa, and Linder, 2013). In contrast, the intermediary asset pricing literature emphasizes the nonlinearities in volatility dynamics that arise as financial intermediaries become more constrained. Since financial intermediaries are the providers of credit in these models, nonlinearities in volatility dynamics feed through to investment decisions of the productive and household sectors of the economy, generating an asymmetric propagation of shocks at different points of the business cycle even if the fundamental shocks in the economy are normally distributed.

In addition to the independent role of financial relative to economic conditions, we focus on the evolution of the shape of the conditional GDP distribution. This is in contrast to the recent literature that emphasizes economic risk or uncertainty. Baker et al. (2013) propose a measure of uncertainty based on political news announcements and Jurado et al. (2015) compute conditional volatility from a large number of macroeconomic variables. Bloom (2009) models uncertainty shocks in a macroeconomic context. However, that strand of

work primarily focuses on the second moment, not on the whole shape of the conditional GDP distribution. Our key finding is that moments beyond the first two significantly vary over time, giving rise to a pronounced skew in the conditional GDP growth. Hence our empirical results strongly point towards downside risk, rather than standard deviation, as the relevant metric.

A recent literature that does generate asymmetry in the conditional distribution is focused on the role of disaster risk for macroeconomic performance and asset prices (see the overview by Barro and Ursúa, 2012). Catastrophic tail events, such as the Great Depression or the 2008 financial crises can be used to generate a large equity risk premium (see e.g. Barro, 2009; Gabaix, 2012; Wachter, 2013). The pricing of disaster risk in turn impacts macroeconomic dynamics in equilibrium (Gourio, 2012). In contrast to that literature, our evidence suggests that GDP vulnerability changes at relatively high frequencies that are unlikely to be driven by changes in disaster risk, and are more closely associated with amplification mechanisms emphasized in macro-finance models. This finding is in-line with the more recent evidence from the term structure of asset prices: the slope of term structures of risk premia across multiple asset classes suggests that the equilibrium pricing kernel reflects nonlinear shocks occurring at intermediate frequencies (see e.g. Van Binsbergen, Hueskes, Koijen, and Vrugt, 2013; Backus, Boyarchenko, and Chernov, 2016).

Our findings are closely related to the literature on the Great Moderation, which documents a decline of GDP volatility prior to the financial crisis of 2008 (see McConnell and Perez-Quiros, 2000; Blanchard and Simon, 2001; Bernanke, 2004; Giannone et al., 2008). In contrast to that influential literature, we focus not just on the second moment of GDP growth, but rather on the whole conditional distribution of GDP growth. Indeed, we find that GDP volatility is nearly entirely driven by the left side of the conditional distribution, thus providing strong implications for research. In fact, we can attribute the decline of volatility during the Great Moderation period to a decline in the downside risk to GDP growth.

The strong relationship between GDP vulnerability and financial conditions rationalizes the FOMC's increased emphasis on the notion of financial conditions. Peek, Rosengren, and Tootell (2015) document the increased frequency of the mention of financial conditions in the FOMC statement, and show that financial conditions are a significant explanatory variable in augmented Taylor rules. Monetary policy models that feature frictions in the financial intermediary sector rationalize these findings, as optimal Taylor rules are augmented by financial variables that can be interpreted as measures of GDP vulnerability (see Curdia and Woodford, 2010; Gambacorta and Signoretti, 2014). Adrian and Liang (2016) point out that monetary policy impacts financial conditions as well as vulnerabilities, thus producing an intertemporal tradeoff for monetary policy between present macroeconomic objectives and risks to objectives in the future. Downside growth vulnerability presents a tool to quantify that tradeoff.

7 Conclusion

The financial crisis of 2007–2009 and the ensuing Great Recession reignited academic interest in the volatility of GDP growth. In this paper, we argue that the entire distribution of GDP growth evolves over time as a function of financial and economic conditions. We measure the vulnerability of GDP growth to downside risks as the downside entropy of a conditional mean, constant variance Gaussian distribution relative to the empirical conditional distribution, and show that growth vulnerability is linked to financial conditions, while economic conditions forecast the center of the GDP distribution.

We estimate our measure of GDP growth vulnerability in two steps. First, we use quantile regressions to estimate the quantile function of GDP growth, with each quantile modeled as a linear function of indicators of financial and economic conditions. In the second step, we smooth the estimated quantile function by using a skewed student t-distribution to interpolate between quantiles. This smoothing step allows us to compute analytically the

Kullback-Leibler divergence of the conditional mean, constant variance Gaussian distribution from the fitted skewed student t-distribution every quarter, and to split the total divergence between the downside and the upside components.

We document that the conditional GDP growth evolution is highly asymmetric, with downside entropy much more volatile than upside entropy. The significant relationship between financial conditions and downside entropy, together with the weak link between both financial and economic conditions and upside entropy, strongly suggests that financial frictions are necessary in modeling the dynamics of the conditional GDP distribution. Recent macro-finance models with a financial sector appear promising in that regard.

References

Adrian, T. and N. Boyarchenko (2012): "Intermediary Leverage Cycles and Financial Stability," *Federal Reserve Bank of New York Staff Reports*, 567.

Adrian, T., R. Crump, and E. Vogt (2015): "Nonlinearity and flight to safety in the risk-return trade-off for stocks and bonds," *Federal Reserve Bank of New York Staff Report*, 723.

Adrian, T. and N. Liang (2016): "Monetary Policy, Financial Conditions, and Financial Stability," *FRB of New York Staff Report*.

Azzalini, A. (1985): "A class of Distributions which Includes the Normal Ones," *Scandinavian Journal of Statistics*, 171–178.

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, 367–389.

Azzalini, A. and A. Dalla Valle (1996): "The Multivariate Skew-normal Distribution," *Biometrika*, 83, 715–726.

Backus, D., N. Boyarchenko, and M. Chernov (2016): "Term structures of asset prices and returns," Tech. rep., National Bureau of Economic Research, Working Paper No. w22162.

Baker, S., N. Bloom, and S. Davis (2013): "Measuring Economic Policy Uncertainty," *Stanford University Working Paper*.

Barro, R. J. (2009): "Rare Disasters, Asset Prices, and Welfare Costs," *American Economic Review*, 99, 243–264.

Barro, R. J. and J. F. Ursúa (2012): "Rare Macroeconomic Disasters," *Annual Review of Economics*, 4, 83–109.

Bernanke, B. (2004): "The Great Moderation," in *The Taylor Rule and the transformation of monetary policy*, Institutions Press Publication Hoover.

Bernanke, B. and M. Gertler (1995): "Inside the Black Box: The Credit Channel of Monetary Policy Transmission," *Journal of Economic Perspectives*, 9, 27–48.

Bernanke, B. S., M. Gertler, and S. Gilchrist (1999): "The Financial Accelerator in a Quantitative Business Cycle Framework," *Handbook of Macroeconomics*, 1, 1341–1393.

van Binsbergen, J., W. Hueskes, R. Koijen, and E. Vrugt (2013): "Equity yields," *Journal of Financial Economics*, 110, 503–519.

Blanchard, O. and J. Simon (2001): "The Long and Large Decline in US Output Volatility," *Brookings Papers on Economic Activity*, 2001, 135–164.

Bloom, N. (2009): "The Impact of Uncertainty Shocks," *Econometrica*, 77, 623–685.

Boivin, J., M. T. Kiley, F. S. Mishkin, et al. (2010): "How Has the Monetary Transmission Mechanism Evolved Over Time?" *Handbook of Monetary Economics*, 3, 369–422.

Brave, S. and A. Butters (2012): "Diagnosing the Financial System: Financial Conditions and Financial Stress," *International Journal of Central Banking*, 8, 191–239.

Brunnermeier, M. K., T. M. Eisenbach, and Y. Sannikov (2013): "A Survey of Financial Frictions in Macroeconomics," in *Advances in Economics and Econometrics: Volume 2, Applied Economics: Tenth World Congress*, Cambridge University Press, vol. 50, 3.

Brunnermeier, M. K. and Y. Sannikov (2014): "A Macroeconomic Model with a Financial Sector," *American Economic Review*, 104, 379–421.

Curdia, V. and M. Woodford (2010): "Credit Spreads and Monetary Policy," *Journal of Money, Credit and Banking*, 42, 3–35.

Del Negro, M., S. Eusepi, M. P. Giannoni, A. M. Sbordone, A. Tambalotti, M. Cocci, R. Hasegawa, and M. H. Linder (2013): "The FRBNY DSGE Model," *FRB of New York Staff Report*.

Engle, R. (1982): "Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of United Kingdom Inflation," *Econometrica: Journal of the Econometric Society*, 987–1007.

Estrella, A. and G. Hardouvelis (1991): "The Term Structure as a Predictor of Real Economic Activity," *Journal of Finance*, 46, 555–576.

Estrella, A. and F. Mishkin (1998): "Predicting US Recessions: Financial Variables as Leading Indicators," *Review of Economics and Statistics*, 80, 45–61.

Gabaix, X. (2012): "Variable rare disasters: An exactly solved framework for ten puzzles in macro-finance," *Quarterly Journal of Economics*, 127, 645–700.

Gambacorta, L. and F. M. Signoretti (2014): "Should monetary policy lean against the wind?: An analysis based on a DSGE model with banking," *Journal of Economic Dynamics and Control*, 43, 146–174.

Ghysels, E. (2014): "Conditional Skewness with Quantile Regression Models: SoFiE Presidential Address and a Tribute to Hal White," *Journal of Financial Econometrics*, 12, 620–644.

Giannone, D., M. Lenza, and L. Reichlin (2008): "Explaining the Great Moderation: It is Not the Shocks," *Journal of the European Economic Association*, 6, 621–633.

Gilchrist, S. and E. Zakrajšek (2012): "Credit Spreads and Business Cycle Fluctuations," *American Economic Review*, 102, 1692–1720.

Gourio, F. (2012): "Disaster Risk and Business Cycles," *American Economic Review*, 102, 2734–2766.

Harvey, C. (1988): "The Real Term Structure and Consumption Growth," *Journal of Financial Economics*, 22, 305–333.

He, Z. and A. Krishnamurthy (2012): "A Model of Capital and Crises," *Review of Economic Studies*, 79, 735–777.

Jurado, K., S. Ludvigson, and S. Ng (2015): "Measuring Uncertainty," *American Economic Review*, 105, 1177–1216.

Koenker, R. and G. Bassett (1978): "Regression Quantiles," *Econometrica*, 33–50.

Krishnamurthy, A. and T. Muir (2016): "How Credit Cycles across a Financial Crisis," *Stanford University Working Paper*.

Longstaff, F. A., J. Pan, L. H. Pedersen, and K. J. Singleton (2011): "How Sovereign Is Sovereign Credit Risk?" *American Economic Journal: Macroeconomics*, 3, 75–103.

López-Salido, D., J. C. Stein, and E. Zakrajšek (2016): "Credit-market Sentiment and the Business Cycle," *National Bureau of Economic Research*, Working Paper.

McConnell, M. M. and G. Perez-Quiros (2000): "Output Fluctuations in the United States: What Has Changed Since the Early 1980's?" *American Economic Review*, 90, 1464–1476.

Peek, J., E. Rosengren, and G. Tootell (2015): "Should U.S. Monetary Policy Have a Ternary Mandate?" Tech. rep., Federal Reserve Bank of Boston.

Rey, H. (2015): "Dilemma not Trilemma: The Global Financial Cycle and Monetary Policy Independence," Tech. rep., National Bureau of Economic Research.

Rossi, B. and T. Sekhposyan (2015): "Alternative tests for correct specification of conditional predictive densities," Manuscript, Texas A&M University.

Schmidt, L. and Y. Zhu (2016): "Quantile Spacings: A Simple Method for the Joint Estimation of Multiple Quantiles Without Crossing," *University of Chicago Working Paper*.

Stock, J. and M. Watson (1999): "Forecasting Inflation," *Journal of Monetary Economics*, 44, 293–335.

Wachter, J. A. (2013): "Can Time-Varying Risk of Rare Disasters Explain Aggregate Stock Market Volatility?" *The Journal of Finance*, 68, 987–1035.

Figure 2. Raw Data. The top two figures show the time series of the CFNAI and NFCI. The lower two panels show the sample quantile-quantile relationship between one quarter ahead real GDP growth and CFNAI and NFCI.

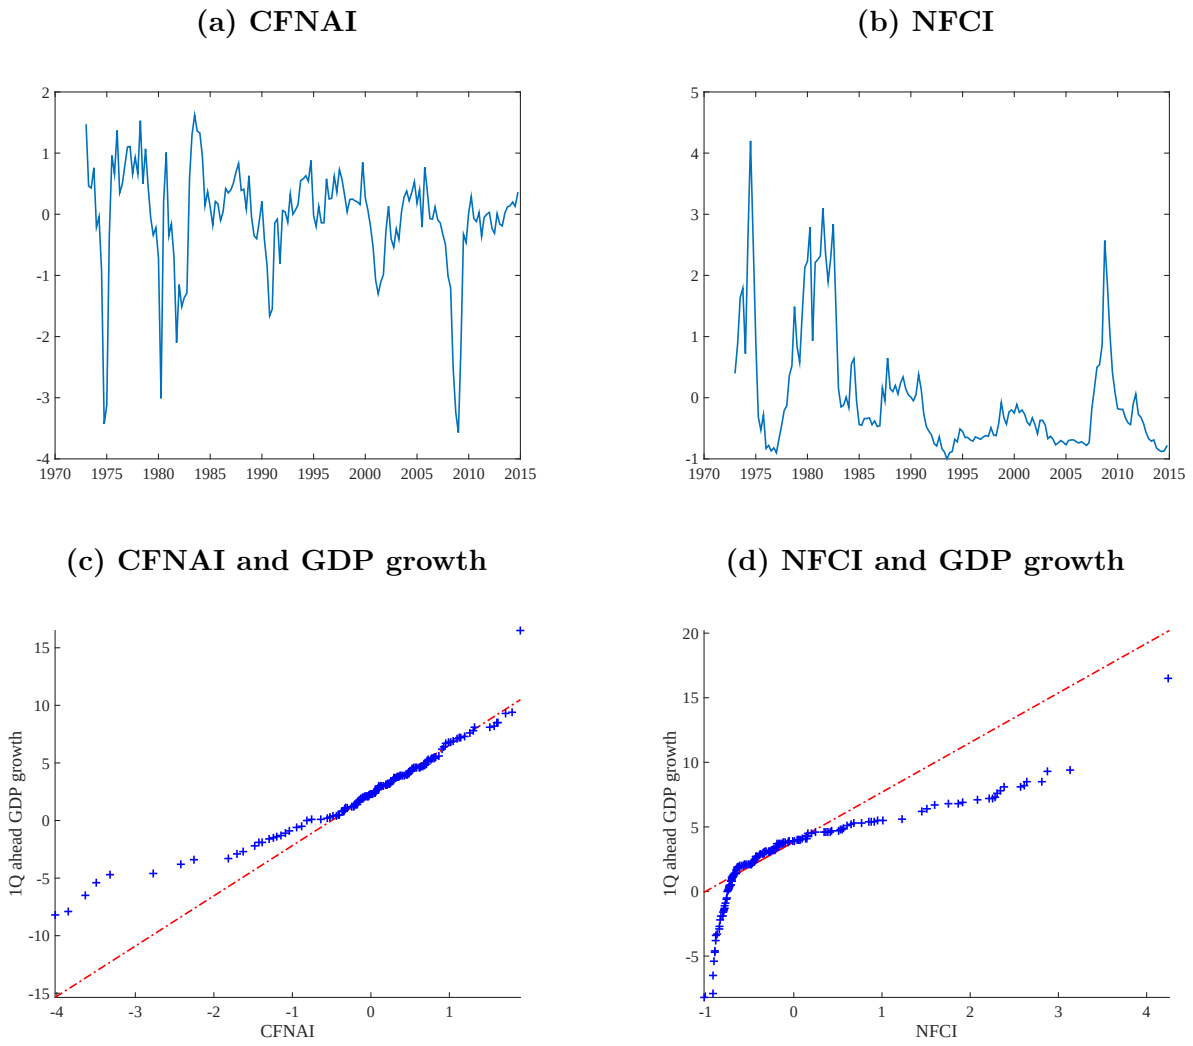

Figure 3. Quantile Regressions. The figure shows the quantile regressions of one quarter ahead (left column) and four quarter ahead (right column) real GDP growth on CFNAI and NFCI.

(a) One quarter ahead: CFNAI

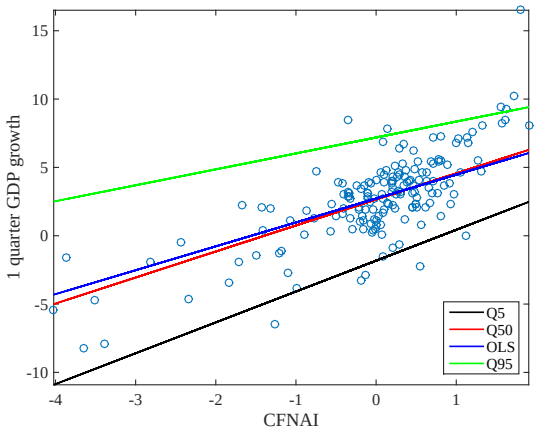

(b) Four quarters ahead: CFNAI

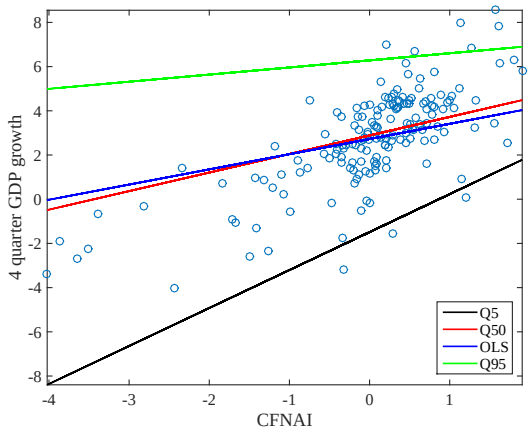

(c) One quarter ahead: NFCI

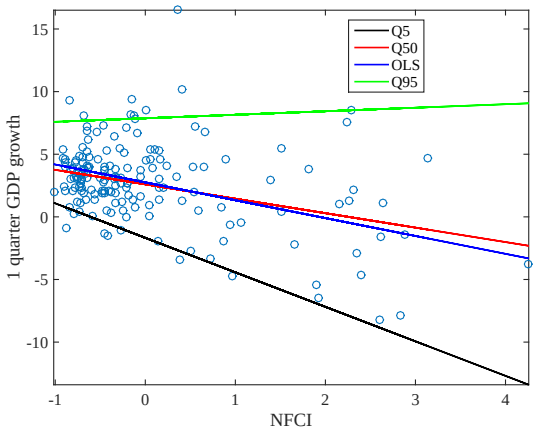

(d) Four quarters ahead: NFCI

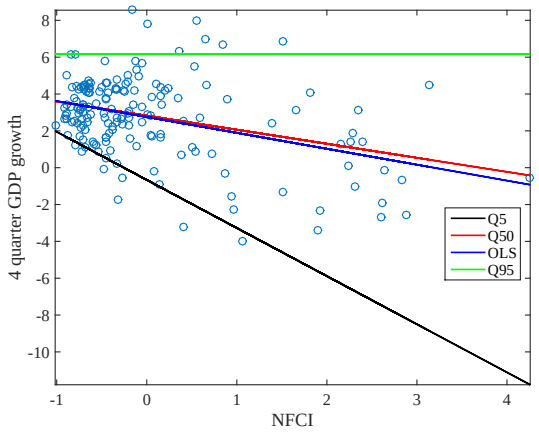

Figure 4. Estimated Quantile Regression Coefficients. The figure shows the estimated coefficients in quantile regressions of one quarter ahead (left column) and four quarter ahead (right column) real GDP growth on CFNAI and NFCI. Confidence bounds computed using VAR with 4 lags and 1000 simulation samples.

(a) One quarter ahead: CFNAI

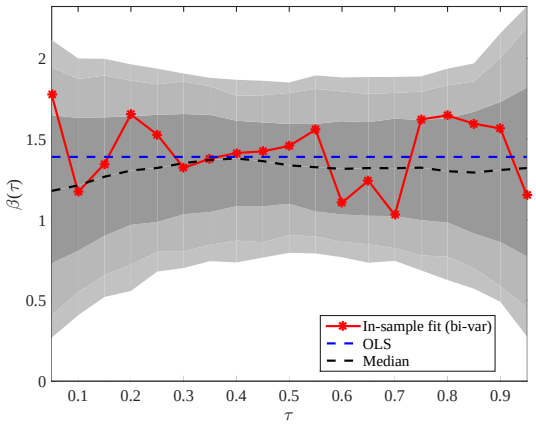

(b) Four quarters ahead: CFNAI

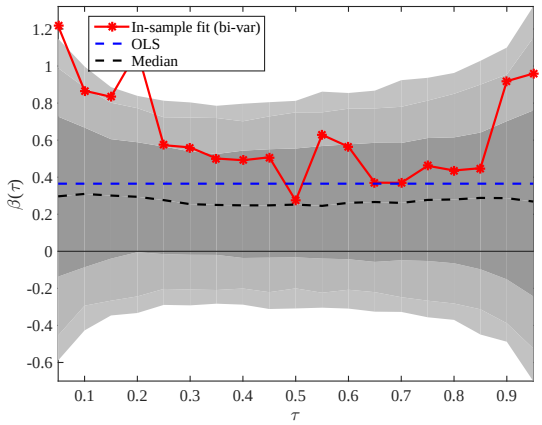

(c) One quarter ahead: NFCI

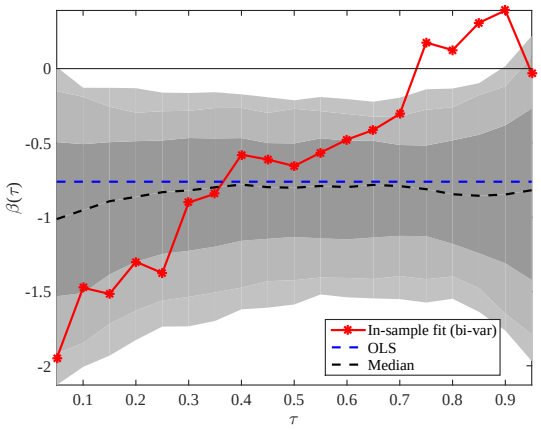

(d) Four quarters ahead: NFCI

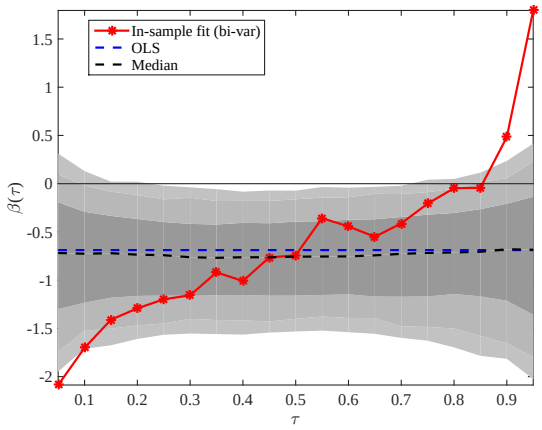

Figure 5. Predicted Distributions. The figure shows the time series evolution of the predicted distribution of one quarter ahead (panel a) and four quarter ahead (panel b) real GDP growth.

(a) One quarter ahead

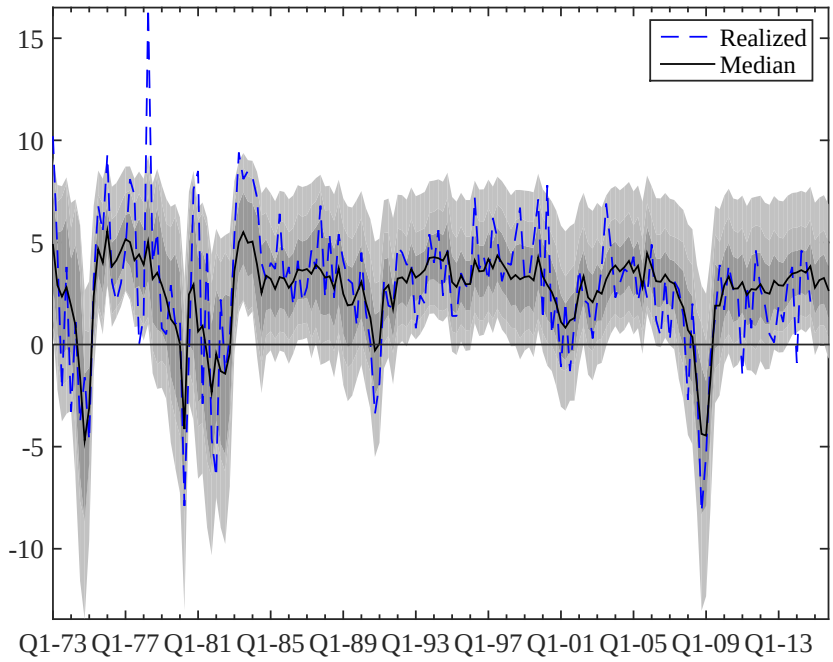

(b) Four quarters ahead

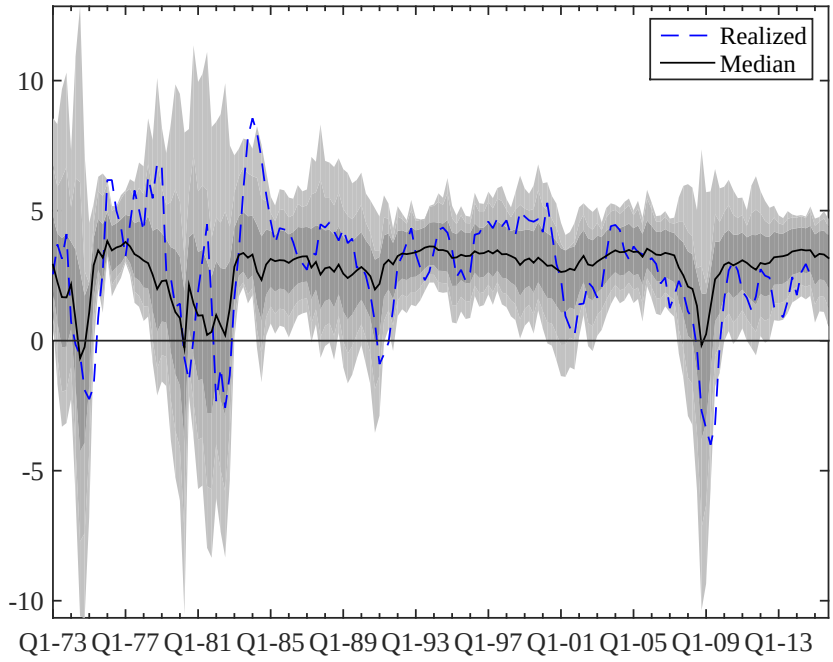

Figure 6. Median, Interquartile Range, and 5 Percent Quantile. The figure shows scatter plots of the median versus the interquartile range (6a and 6b) and the 5th percent quantile versus the interquartile range (6c and 6d).

(a) Median and Interquartile Range
One quarter ahead

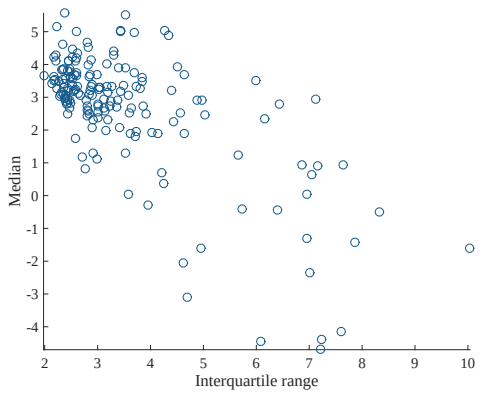

(b) Median and Interquartile Range
Four quarters ahead

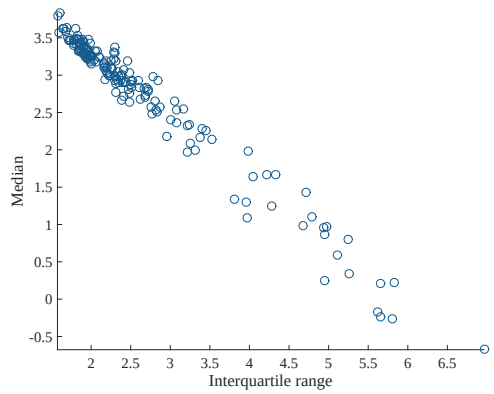

(c) Q5 and Interquartile Range
One quarter ahead

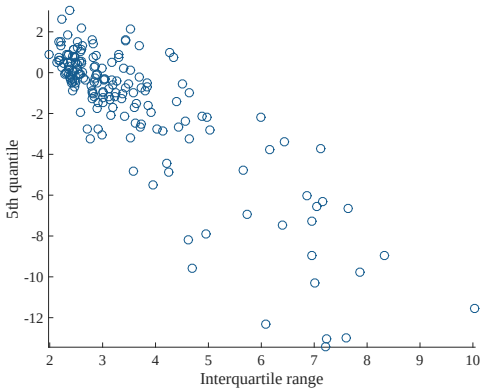

(d) Q5 and Interquartile Range
Four quarters ahead

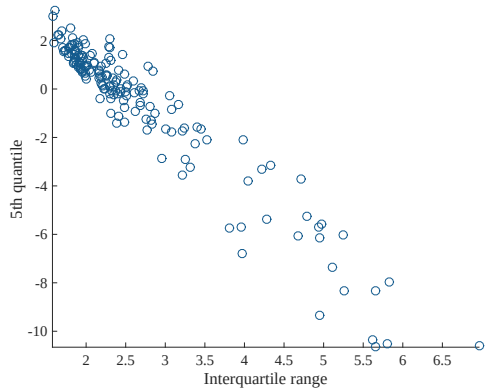

Figure 7. The Conditional Quantiles and the Skewed t-Distribution. The panels in this figure show the conditional GDP quantiles together with the estimated skewed t-inverse cumulative distribution functions for one quarter and four quarter ahead GDP. For comparison, we also report the skewed t-inverse cumulative distribution functions obtained by fitting the unconditional empirical distribution.

(a) One quarter ahead: Q2 2006

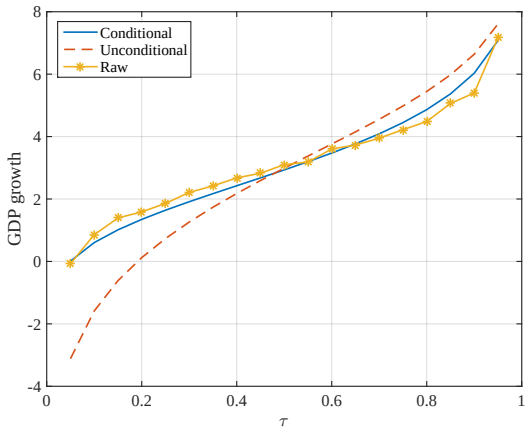

(b) Four quarters ahead: Q2 2006

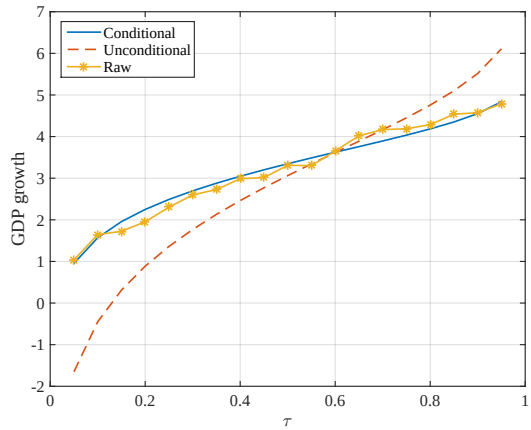

(c) One quarter ahead: Q4 2008

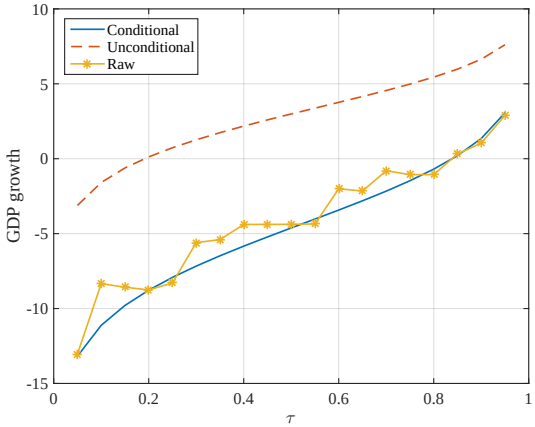

(d) Four quarters ahead: Q4 2008

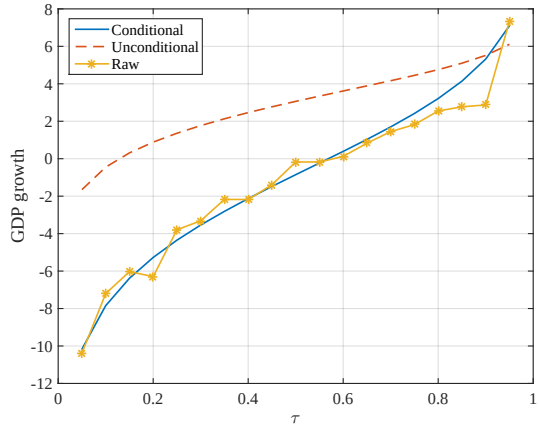

(e) One quarter ahead: Q4 2014

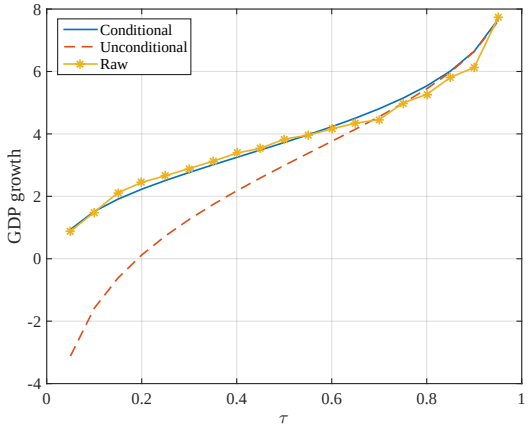

(f) Four quarters ahead: Q4 2014

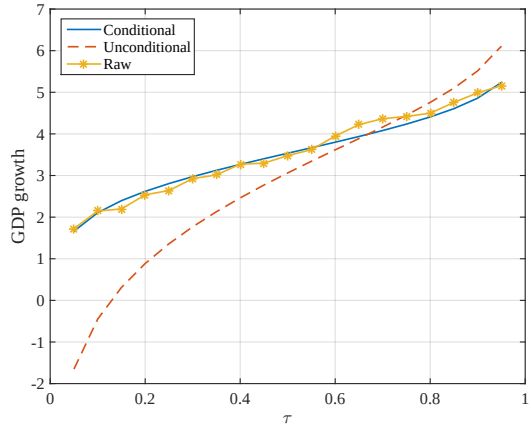

Figure 8. GDP Entropy and Expected Shortfall over Time. The figure shows the time series evolution of relative downside and upside entropy $\mathcal{L}_t^D$ and $\mathcal{L}_t^U$ together with the 5% expected shortfall ES_t. The one quarter ahead series are on the left, the four quarter ahead series are on the right.

(a) Entropy $\mathcal{L}_t$: One quarter ahead

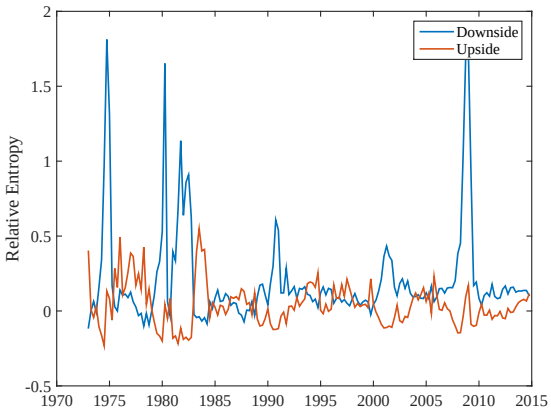

(b) Entropy $\mathcal{L}_t$: Four quarters ahead

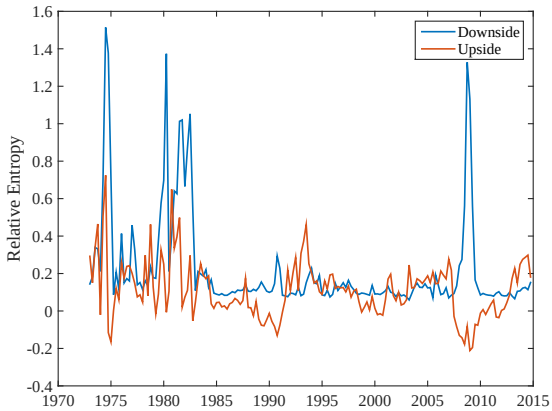

(c) 5% Expected Shortfall, ES_t**, and Longrise,** EL_t: One quarter ahead

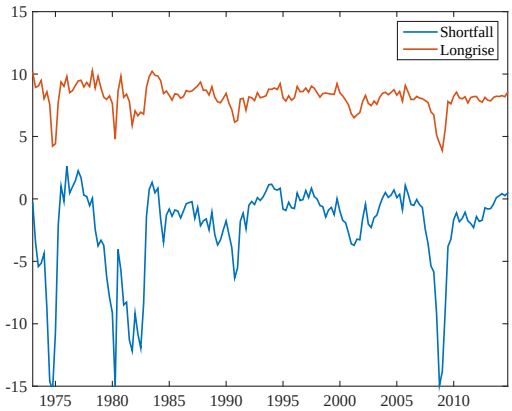

(d) 5% Expected Shortfall, ES_t**, and Longrise,** EL_t: Four quarters ahead

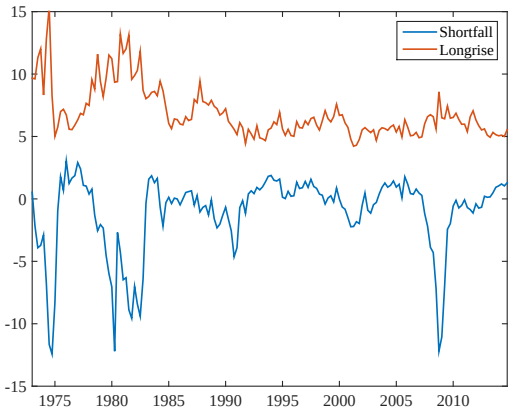

Figure 9. GDP Growth and Other Predictors. The figure shows the estimated coefficients in quantile regressions of one quarter ahead (left column) and four quarter ahead (right column) real GDP growth on CFNAI and VXO; CFNAI and the Baa-Aaa spread; CFNAI and the term spread.

(a) Equity Volatility One quarter ahead

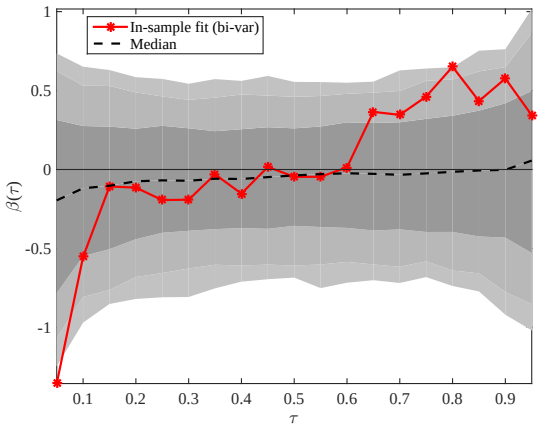

(b) Equity Volatility Four quarters ahead

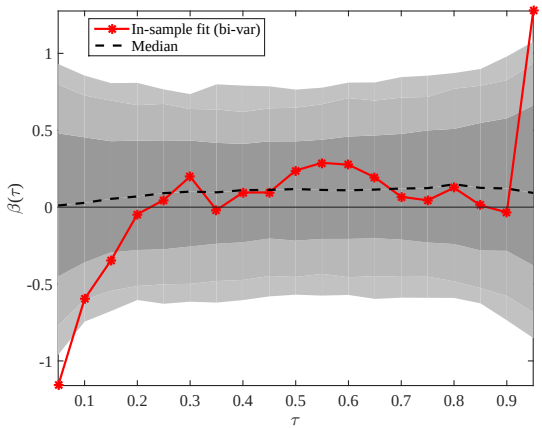

(c) Credit Spread One quarter ahead

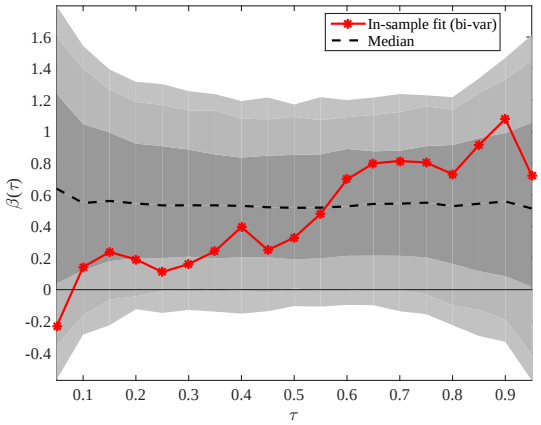

(d) Credit Spread Four quarters ahead

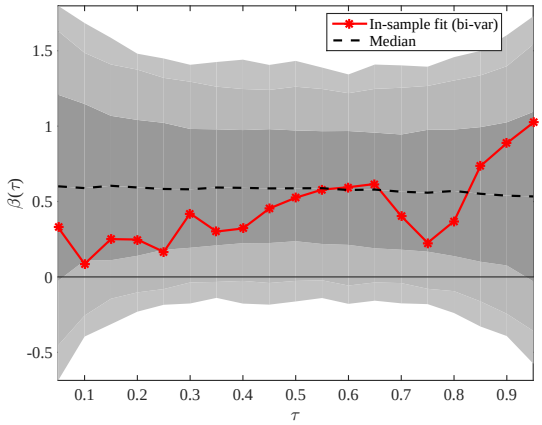

(e) Term Spread One quarter ahead

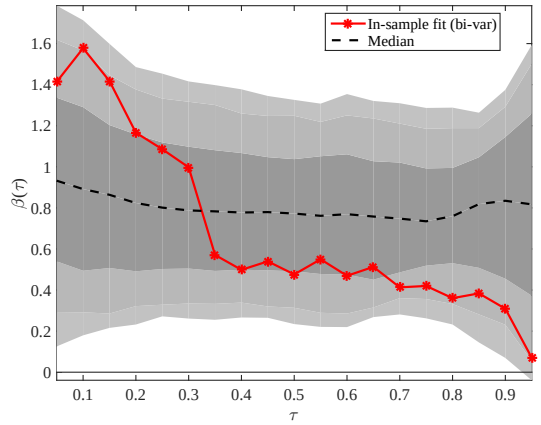

(f) Term Spread Four quarters ahead

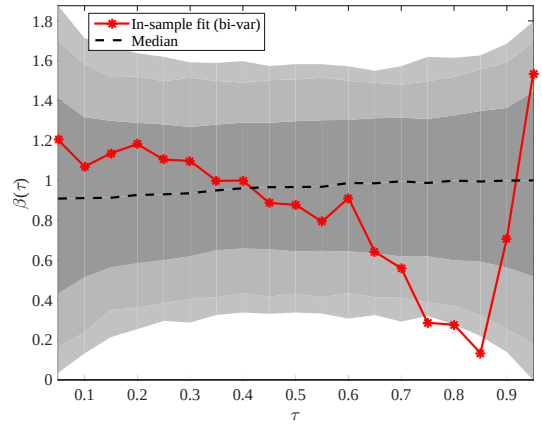

Figure 10. Out of Sample Predictions. The figure compares out-of-sample and in-sample predictive density one and four quarter ahead. Figures 10a and 10b show the 5, 50, and 95 percent quantiles. Figures 10c and 10d show downside entropy.

(a) Quantiles One quarter ahead

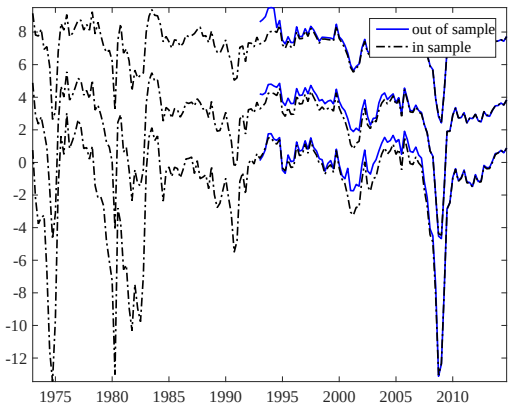

(b) Quantiles Four quarters ahead

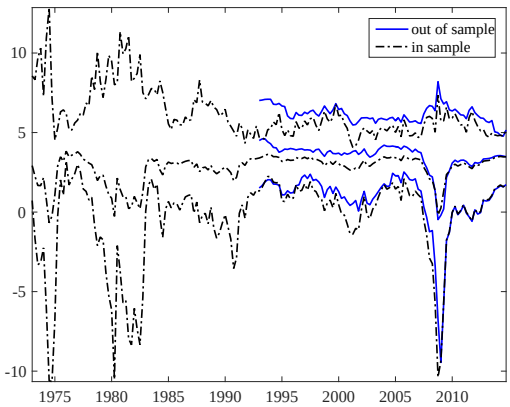

(c) Downside Entropy One quarter ahead

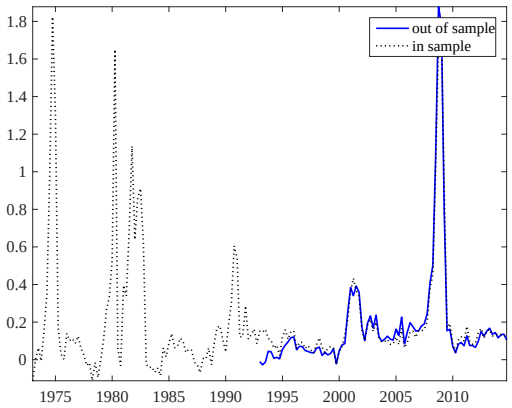

(d) Downside Entropy Four quarters ahead

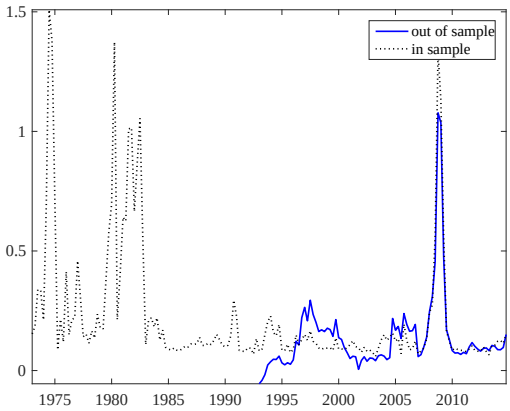

Figure 11. Out of Sample Accuracy. The figure reports the predictive scores and the cumulative distribution of the probability integral transform. Figures 11a and 11b compare the out- of-sample predictive scores of the conditional and unconditional predictive distribution. Figures 11a and 11b report the empirical cumulative distribution of the probability integral transform (PITs).

(a) Predictive scores One quarter ahead

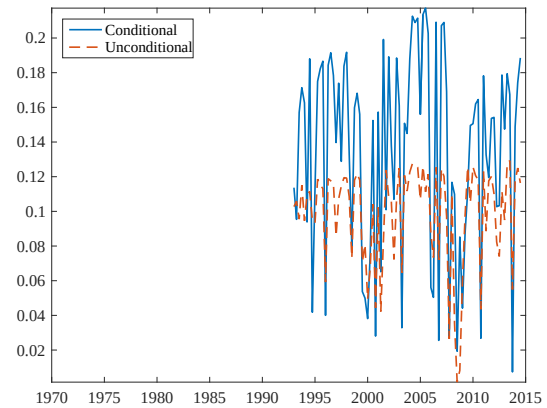

(b) Predictive scores Four quarters ahead

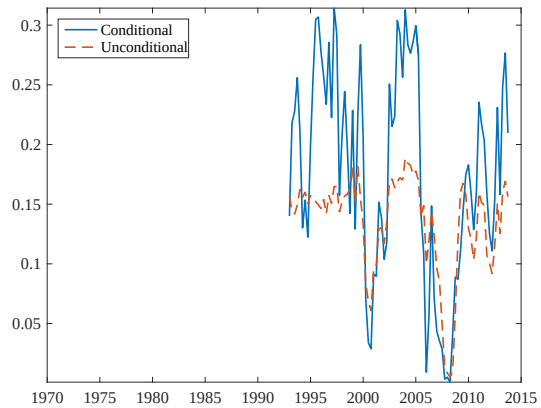

(c) PITs One quarter ahead

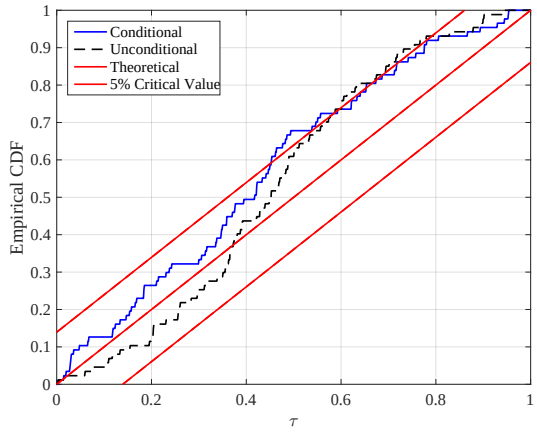

(d) PITs Four quarters ahead

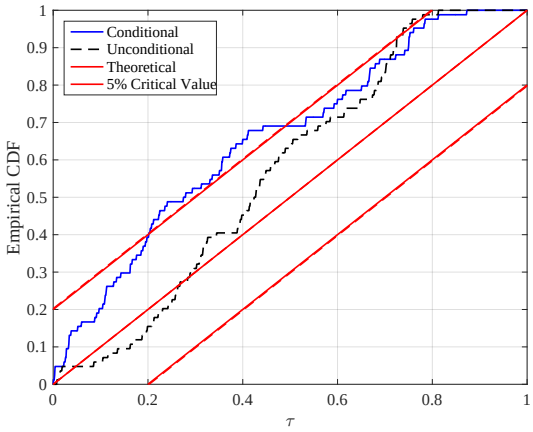

Figure 12. ARCH Estimation of GDP Growth and Stock Return Conditional Distributions. The figure shows the conditional mean, and the 95 percent and 5 percent standard error bands for one quarter GDP growth and one quarter S&P500 returns. The distributions are estimated with a ARCH(1) model with the NFCI as conditioning variable in the mean equation and the volatility equations.

(a) 1 Quarter Ahead GDP Growth

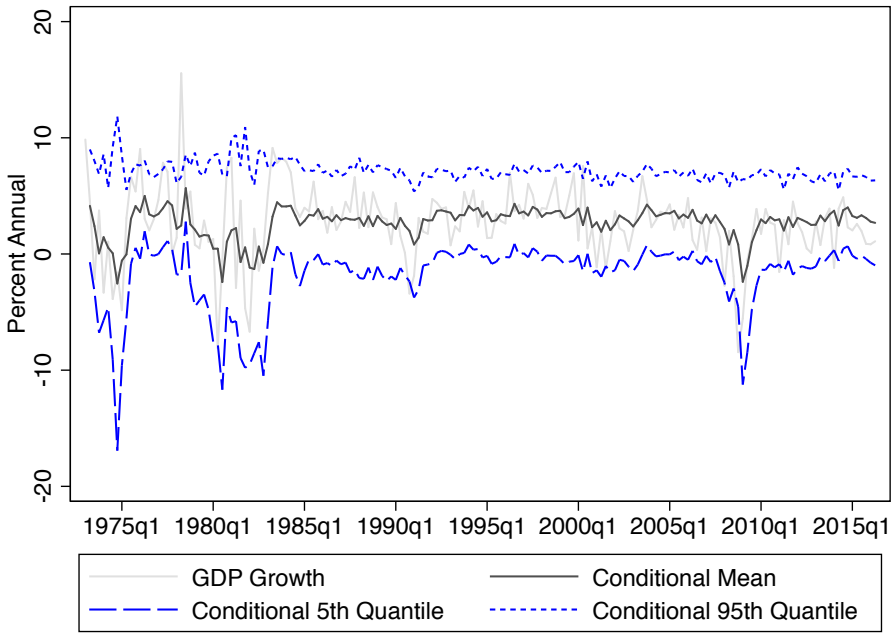

(b) 1 Quarter Ahead S&P500 Return

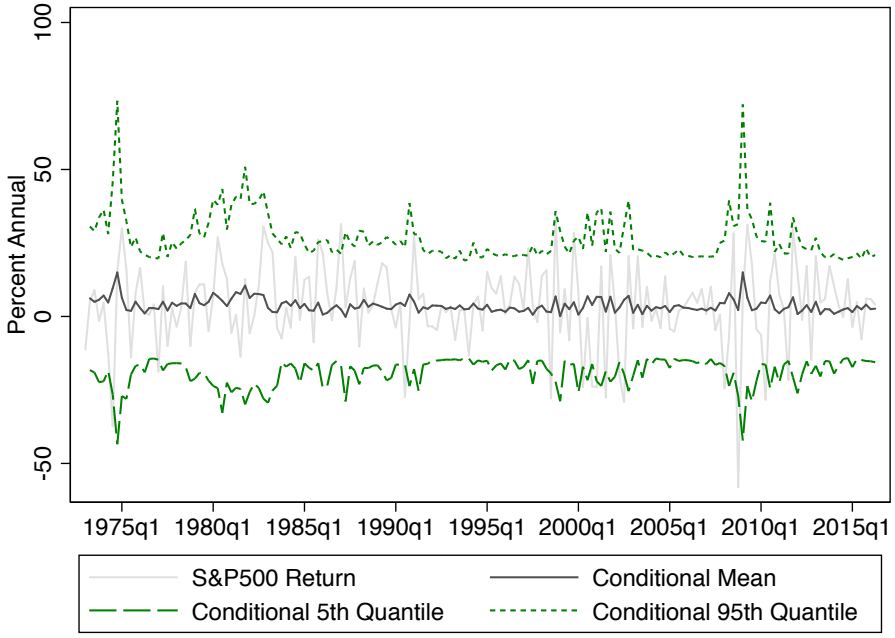

Figure 13. Mean, Volatility, and Entropy. The figure shows scatter plots of the mean versus the volatility (13a and 13b) and downside entropy versus the mean (13a and 13b) for one and four-quarters-ahead distribution of GDP growth.

(a) Mean and Volatility One quarter ahead

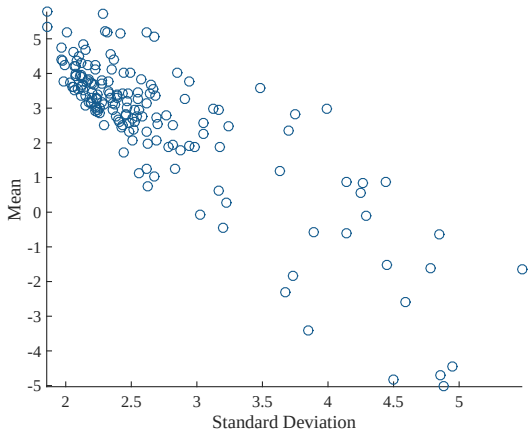

(b) Mean and Volatility Four quarters ahead

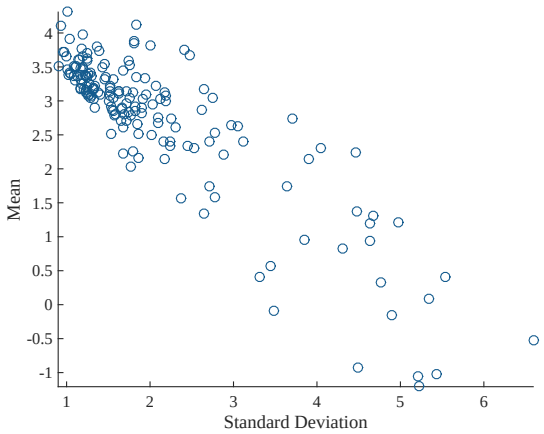

(c) Mean and Entropy One quarter ahead

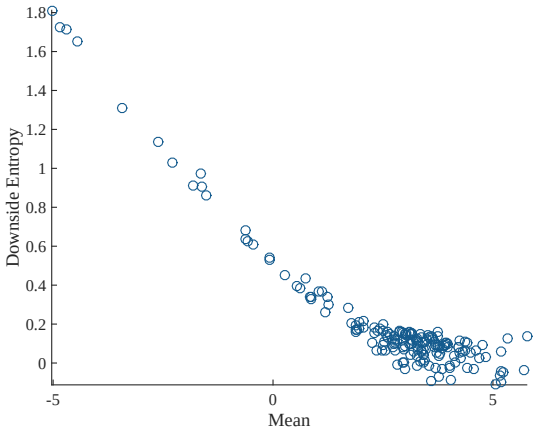

(d) Mean and Entropy Four quarters ahead

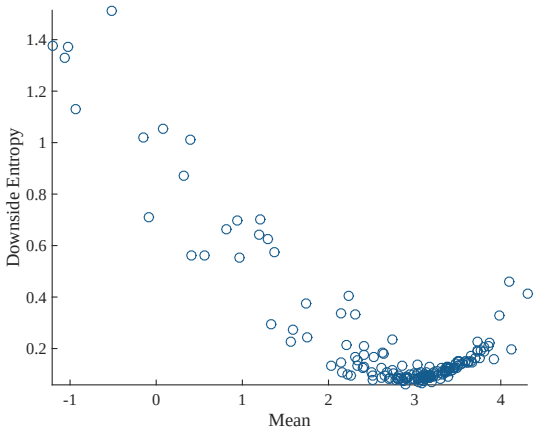

Figure 14. Median, Interquartile Range, and 5 Percent Quantile. The figure shows the cross correlograms between the mean and volatility, and mean and downside entropy.

(a) Cross Correlogram Mean and Volatility One quarter ahead

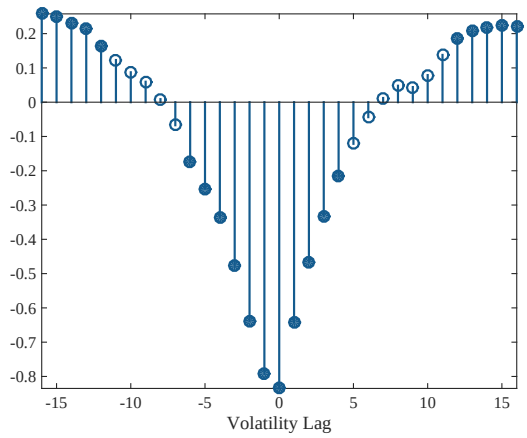

(b) Cross Correlogram Mean and Volatility Four quarters ahead

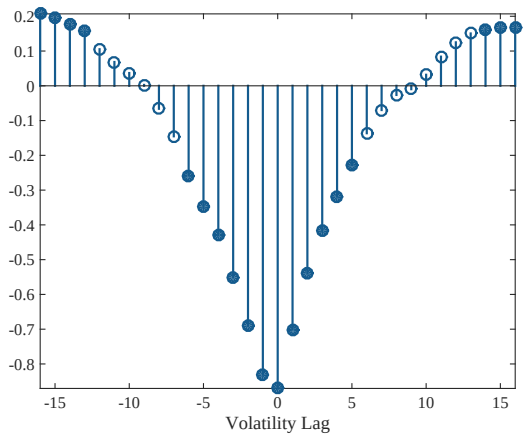

(c) Cross Correlogram Mean and Entropy One quarter ahead

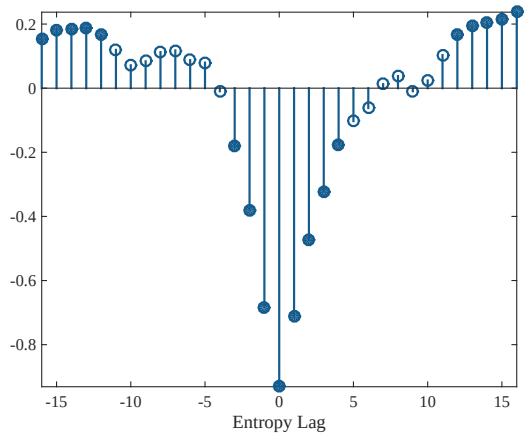

(d) Cross Correlogram Mean and Entropy Four quarters ahead

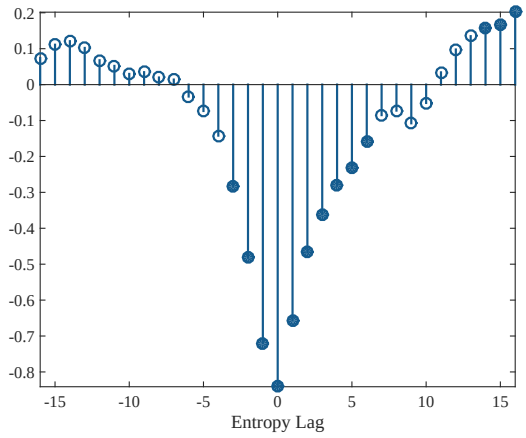

Internet Appendix

Figure A.1. GDP Location and Scale Parameters over Time. The figure shows the time series evolution of location μ_t and scale σ_t. The one quarter ahead parameters are in the left column, the four quarter ahead parameters are in the right column.

(a) Location μ_t

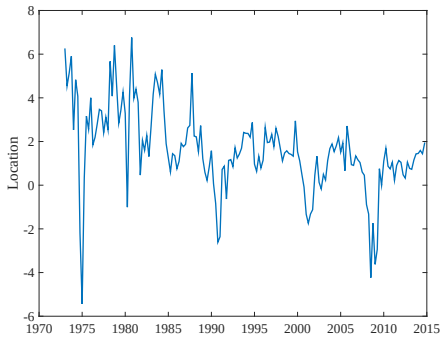

(b) Location μ_t

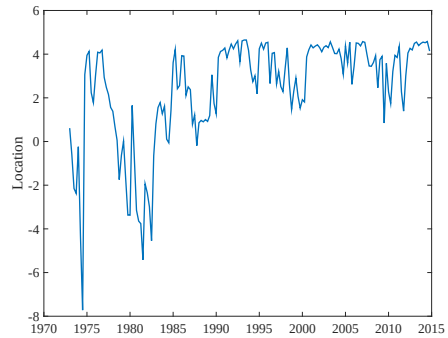

(c) Scale σ_t

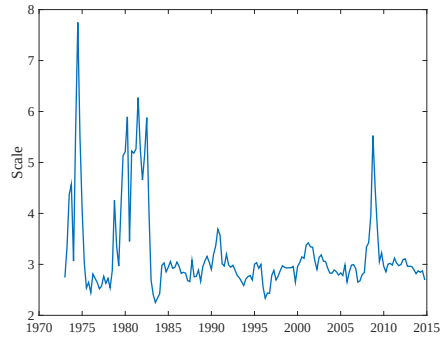

(d) Scale σ_t

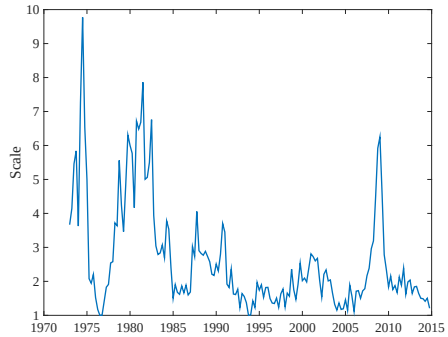

Figure A.2. GDP Shape and Degrees of Freedom over Time. The figure shows the time series evolution of location shape α_t and degrees of freedom ν_t. The one quarter ahead parameters are in the left column, the four quarter ahead parameters are in the right column.

(a) Shape α_t

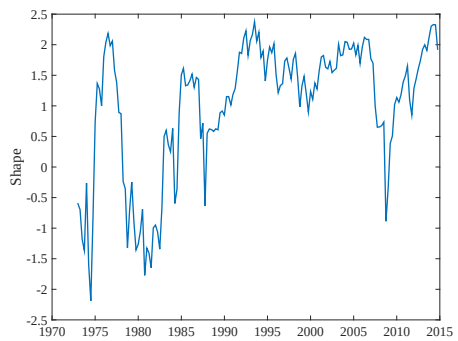

(b) Shape α_t

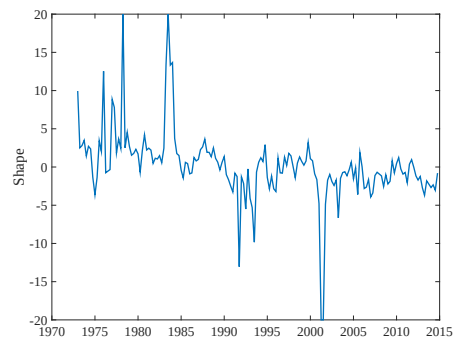

(c) Degrees of Freedom ν_t

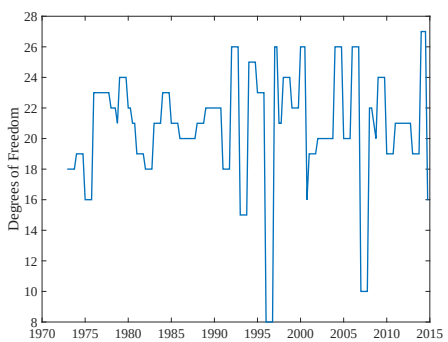

(d) Degrees of Freedom ν_t

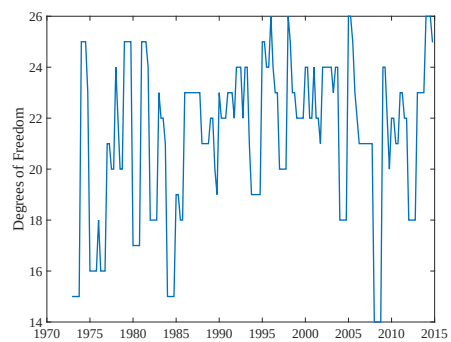

Figure A.3. GDP Moments over Time. The figure shows the time series evolution of the conditional mean and variance. The one quarter ahead parameters are in the left column, the four quarter ahead parameters are in the right column.

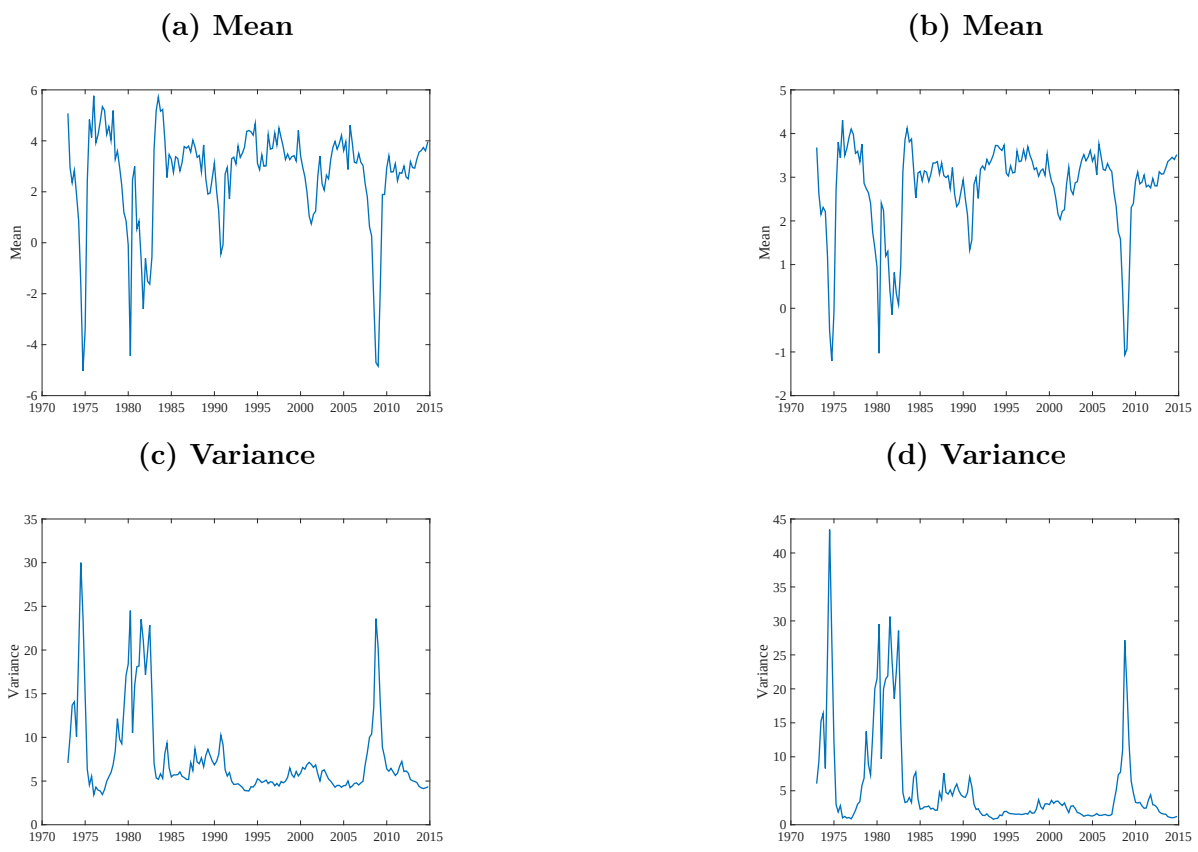

Figure A.4. GDP Moments over Time. The figure shows the time series evolution of the conditional skewness, and kurtosis. The one quarter ahead parameters are in the left column, the four quarter ahead parameters are in the right column.

(a) Skewness

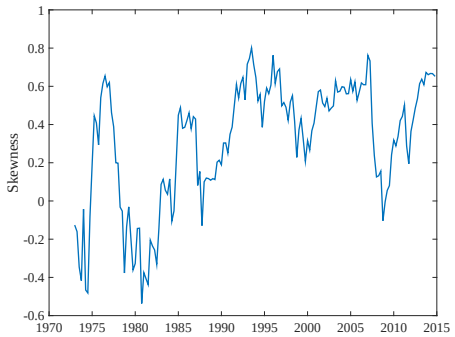

(b) Skewness

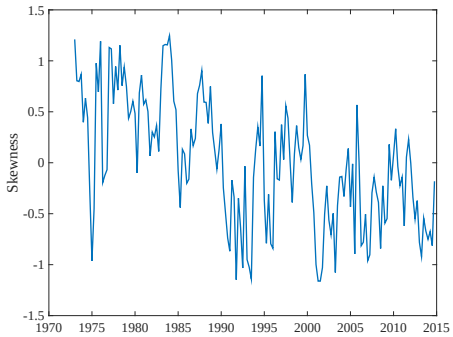

(c) Kurtosis

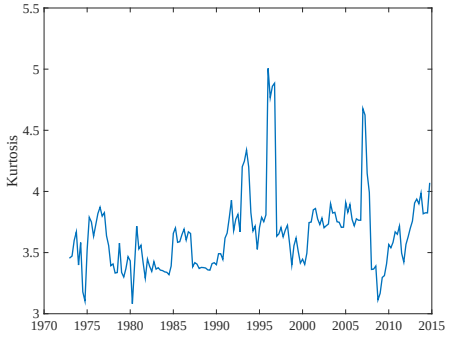

(d) Kurtosis

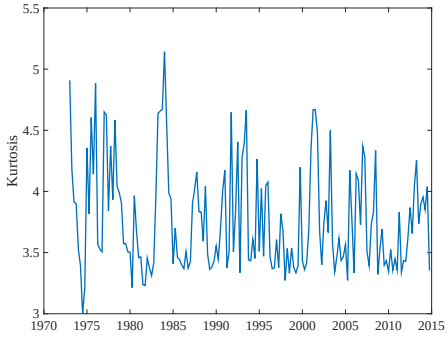

Figure A.5. Probability Densities. The panels in this figure show the estimated skewed t-density functions for one quarter and four quarter ahead GDP. For comparison, we also report the skewed t-density functions obtained by fitting the unconditional empirical distribution.

(a) One quarter ahead: Q2 2006

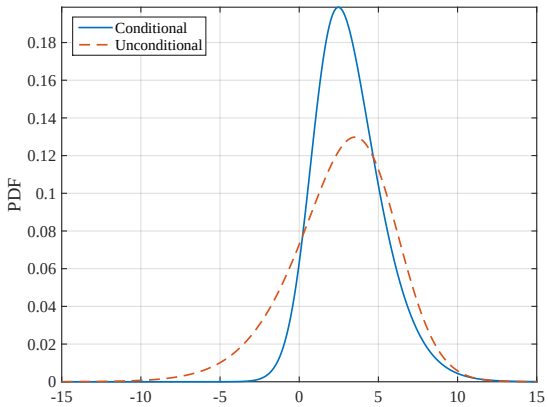

(b) Four quarters ahead: Q2 2006

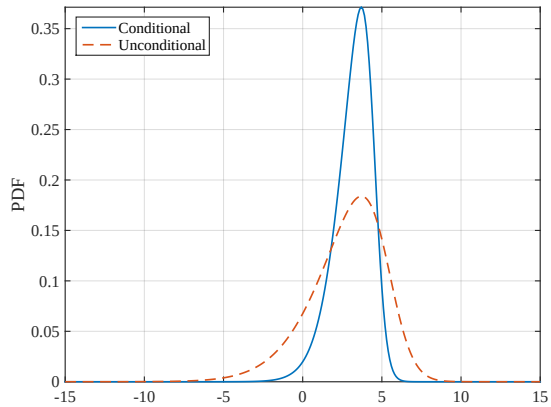

(c) One quarter ahead: Q4 2008

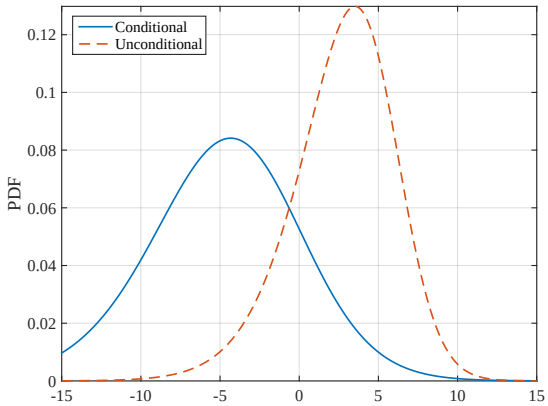

(d) Four quarters ahead: Q4 2008

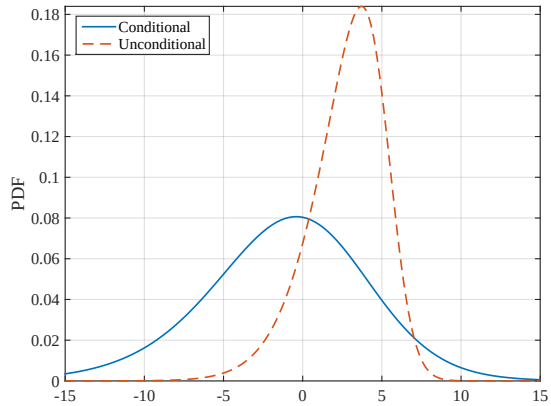

(e) One quarter ahead: Q4 2014

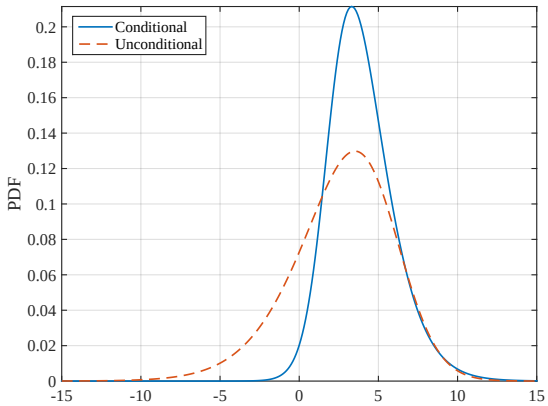

(f) Four quarters ahead: Q4 2014