Federal Reserve Bank of New York
Staff Reports

Vulnerable Growth

Tobias Adrian
Nina Boyarchenko
Domenico Giannone

FRBNY
Staff
REPORTS

Vulnerable Growth
Tobias Adrian, Nina Boyarchenko, and Domenico Giannone
Federal Reserve Bank of New York Staff Reports, no. 794
September 2016; revised November 2017
JEL classification: C12, 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 the conditional volatility and a decline in the conditional mean of GDP growth, leading the lower quantiles of GDP growth to vary with financial conditions and the upper quantiles to be stable over time: Upside risks to GDP growth are low in most periods while downside risks increase as financial conditions become tighter. We argue that amplification mechanisms in the financial sector generate the observed growth vulnerability dynamics.

Key words: downside risk, entropy, quantile regressions

1 Introduction

Economic forecasts usually provide point estimates for the conditional mean of GDP growth and other economic variables. However, such point forecasts 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 the business cycle. Globally, a number of inflation targeting central banks publish GDP growth and inflation distributions. At the same time, surveys of economists (the Blue Chip Economic Survey), market participants (the Federal Reserve Bank of New York's Primary Dealer Survey) and professional forecasters (the Federal Reserve Bank of Philadelphia's Survey of Professional Forecasters) all collect respondents' beliefs regarding 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 current financial and economic conditions. We estimate the distribution semi-parametrically using quantile regressions. Our main finding is that the estimated lower quantiles of the distribution of future GDP growth exhibit strong variation as a function of current financial conditions, while the upper quantiles are stable over time. Moreover, we show that current 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, a flexible distribution function with four parameters. This allows us to transform the empirical quantile distribution into an estimated conditional distribution of GDP growth, plotted in Figure 1. Two features are striking about the estimated distribution. First, the entire distribution, and not just the central tendency, evolves over time. For example, recessions are associated with left-skewed

1

Figure 1. Distribution of GDP growth over time. One-year-ahead predictive distribution of real GDP growth, based on quantile regressions with current real GDP growth and NFCI as conditioning variables.
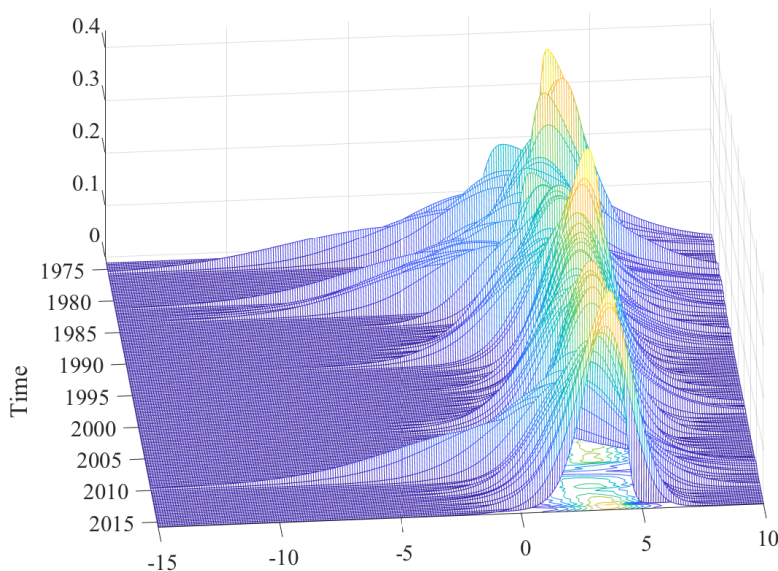
distributions while, during expansions, the conditional distribution is closer to being symmetric. Second, the probability distributions inherit the stability of the right tail from the estimated quantile distribution, while the median and the left tail of the distribution exhibit strong time series variation. This asymmetry in the evolution of the conditional distribution of future GDP growth indicates that downside risk to growth varies much more strongly over time than upside risk.

We summarize the downside and upside risks to the median GDP growth forecast using two metrics: (1) the upside and downside entropy of the unconditional distribution of GDP growth relative to the empirical conditional distribution; (2) the expected shortfall and its upper tail counterpart ("expected longrise"). While downside relative entropy captures the conditional risks to the downside in excess of the downside risks predicted by the unconditional distribution, expected shortfall measures the total probability mass that the conditional distribution assigns to the left tail of the distribution. Similarly, upside relative entropy captures the conditional risks to the upside in excess of the upside risks predicted

by the unconditional distribution, and the expected longrise measures the total probability mass that the conditional distribution assigns to the right tail of the distribution. We find that both measures of downside risk move with financial conditions, whereas both measures of upside risk are significantly more stable over time. This asymmetry echoes our finding that the dependence of future GDP growth on current financial conditions is significantly stronger for the lower quantiles of the distribution than for the upper quantiles.

We find that the conditional mean of GDP growth is negatively correlated with conditional volatility and measures of downside risk, with changes in conditional volatility leading changes in conditional mean by at least one quarter. These findings are consistent with the "volatility paradox" postulated in the recent literature on intermediary asset pricing (see e.g Brunnermeier and Sannikov, 2014; Adrian and Boyarchenko, 2012): Periods of low volatility precede negative growth outcomes.

We perform many robustness tests to our findings. First, we show that out-of-sample estimates of the conditional distribution of future growth 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. Furthermore, we analize predictive scores and probability integral transforms to document our strong out-of-sample performance. Second, we demonstrate that the strong time variation of lower quantiles of future GDP growth is not an artifact of our two-step linear quantile regression estimation procedure, but also arises both when we estimate the conditional distribution either fully parametrically or fully non-parametrically. Finally, in the appendix, we present alternative measures of financial conditions, focusing on specific variables whose predictive power for growth that has been emphasized in the recent macro finance literature, such as credit spreads, the term spread, and equity volatility. We find that the conditional quantile function is most sensitive to the overall financial conditions index, followed by equity volatility, term spread, and credit spread.

Our findings have strong implications for the recent macro-finance literature that emphasizes the link between financial stability and macroeconomic performance. We document a

non-linear relationship between financial conditions and the conditional distribution of GDP growth, suggesting that DSGE models with frictions in either the supply of or the demand for credit should allow for non-linear equilibrium relationships. In such models, financial conditions create downside risk to the economy. For example, the buildup of leverage and maturity transformation in the financial sector can give rise to financial vulnerability. The arrival of unexpected productivity or credit demand shocks then causes the financial sector to disinvest from the real economy, which depresses real economic growth.

The relationship between financial conditions and downside risk to GDP growth could instead be non-causal, with financial conditions merely providing a better signal of negative shocks to the economy. In this case, models that focus on the impact of Bayesian learning on aggregate outcomes (see e.g. Orlik and Veldkamp, 2014; Johannes, Lochstoer, and Mou, 2016) should allow agents to use financial conditions as a signal in forming their beliefs.

Finally, our evidence suggests that GDP vulnerability changes at relatively high frequencies. Such frequent changes in downside risk are in contrast to what would be predicted by the recent literature on disaster risk and economic growth (see e.g. Barro, 2009; Gabaix, 2012; Wachter, 2013; Barro and Ursúa, 2012; Gourio, 2012), but are consistent with the evidence from the term structure of asset prices across multiple markets.

A large literature has documented the decline of GDP volatility before the financial crisis of 2008 (see e.g. Kim and Nelson, 1999; McConnell and Perez-Quiros, 2000; Blanchard and Simon, 2001; Bernanke, 2004; Giannone, Lenza, and Reichlin, 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. Our striking finding is that GDP growth volatility is nearly entirely driven by the left side of the conditional distribution. In fact, we can attribute the decline in volatility during the Great Moderation period prior to the financial crisis 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 estimating and 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 the quantile regressions of Koenker and Bassett (1978) to estimate the conditional quantile function of future GDP growth as a function of current financial and economic conditioning variables. Ghysels (2014) and Schmidt and Zhu (2016) provide recent applications of quantile regressions to the estimation of the distribution of stock returns. 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. The procedure is computationally straightforward, and allows us to transform the inverse cumulative distribution function from the quantile regression into a density function. Alternative ways to estimate conditional predictive distributions for GDP growth proposed in the literature include the two state Markov chain (Hamilton, 1989), the Bayesian vector autoregression with stochastic volatility (Cogley, Morozov, and Sargent, 2005; Primiceri, 2005; Clark, 2012; D'Agostino, Gambetti, and Giannone, 2013), and copula estimates (Smith and Vahey, 2016). Our approach makes fewer parametric assumptions and is computationally much less burdensome. Importantly, the two state Markov chain does not feature financial conditions as state variables, and the Bayesian vector autoregression features exogenous time variation of risk. Instead, we find that the conditional mean and the conditional volatility are negatively correlated as a function of current financial conditions.

Our approach differs from the recent literature that has analyzed GDP uncertainty in its finding of the preeminent role for downside risk, rather than symmetric measures of risk. Baker, Bloom, and Davis (2016) propose a measure of political uncertainty based on news announcements. Jurado, Ludvigson, and Ng (2015) and Clark, Carriero, and Massimiliano (2016) compute conditional volatility from a large number of macroeconomic variables. That literature also finds the conditional mean and volatility of GDP growth to be negatively correlated. The main difference to our work is our emphasis on financial conditions as determinants of the conditional GDP distribution. Intriguingly, our measure of downside vulnerability correlates with the macroeconomic uncertainty index of Jurado et al. (2015)

but is more stable during the Great Moderation period.

More closely related to our paper, Giglio, Kelly, and Pruitt (2016) also use a quantile regression approach to evaluate the ability of various measures of systemic risk proposed in the literature to predict real activity outcomes. They find that, although some measures of systemic risk are statistically significant predictors of the left tail of real activity outcomes, systemic risk measures are unable to predict the right tail of real activity outcomes. Our approach differs as we focus on the entire GDP distribution.

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 discusses out-of-sample results and alternative fully parametric and nonparametric approaches. Section 5 discusses implications of our findings for macroeconomic theory. Section 6 concludes.

2 Economic and Financial Conditions and GDP Growth

To gauge economic and financial conditions, we use real GDP growth and the National Financial Conditions Index (NFCI).[1] 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.[2] When the NFCI is positive, financial conditions are tighter than average. The methodology for the NFCI is described in Brave and Butters (2012) and is based on the quasi maximum likelihood estimators for large dynamic factor models developed by Doz, Giannone, and Reichlin (2012). The data for the NFCI starts in January 1973, which we use as starting point for our empirical investigation. We use real GDP data from the Bureau of Economic

---

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

[2]The list of indicators is provided here.

Analysis (BEA) to compute real GDP growth.[3]

Figure 2 shows the times series of real GDP growth and the NFCI.[4] The time series plot is our first indication of a non-linear relationship between future GDP growth and financial conditions: While GDP growth is, on average, much more volatile than the NFCI, extreme negative outcomes in GDP growth tend to coincide with extreme positive outcomes of the NFCI.

To characterize formally the conditional relationship between future GDP growth and current financial and economic conditions, we rely on quantile regressions. Let us denote by y_{t+h} the annualized average growth rate of GDP between t and $t+h$ and by x_t a vector containing the conditioning variables, including a constant. In a quantile regression of y_{t+h} on x_t the regression slope β_τ is chosen to minimize the quantile weighted absolute value of errors:

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

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

$$\hat{Q}_{y_{t+h}|x_t} (\tau|x_t) = x_t \hat{\beta}_\tau. \quad (2)$$

Koenker and Bassett (1978) show that $\hat{Q}_{y_{t+h}|x_t} (\tau|x_t)$ is a consistent linear estimator of the quantile function of y_{t+h} 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

---

[3]Downloaded from FRED.

[4]The NFCI is converted into quarterly frequency by averaging the weekly observations within each quarter. For the attribution of weeks to overlapping quarters we follow the convention of Federal Reserve Economic Data (FRED) Economic Data, which is the source of our data. Weeks that start in one quarter and end in the next one are fully assigned to the latter quarter. For example, a weekly period that starts on Monday September 28 2015 and ends on Friday October 2 2015 is included in the aggregated value for the fourth quarter.

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 current realization of real GDP growth, as well as the univariate 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. Indeed, Figure 4 shows that, at both the lower and the upper quantiles, the estimated slopes are significantly different, at the 10 percent level, from the OLS slope.[5] The regression slopes change dramatically for the NFCI across the quantiles, but are stable for current GDP growth. Importantly, the regression slopes for the NFCI do not change significantly when current GDP growth is also included in the regression, indicating that most of the explanatory power of future GDP vulnerability arises from the information content of financial conditions. On the other hand, for economic conditions, the quantile regression slopes are not statistically significantly different from each other nor from the linear regression slopes, suggesting that economic conditions are uninformative for predicting tail outcomes.

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. Deteriorations of financial conditions coincide with increases in the interquartile range and decreases in the median (Figure 6a and 6b). Thus, the left tail of the distribution shifts to the left, as illustrated in Figure 6c and 6d: The 5th quantile has a negative relationship with the interquartile range. On the other hand, for the upper quantiles, the movements in the

---

[5]The confidence bounds plotted in Figure 4 are the 95 percent confidence bounds for the null hypothesis that the true data-generating process is a flexible and general linear model for growth and financial conditions. In particular, we estimate a vector autoregression (VAR) with four lags, Gaussian innovations and a constant using the full-sample evolution of the NFCI and real GDP growth, and bootstrap 1000 samples to compute the bounds at different confidence level for the OLS relationship. Quantile coefficient estimates that fall outside this confidence bound thus indicate that the relation between GDP growth and the predictive variable is non-linear.

median and the interquartile range are offsetting. Thus, changes in financial conditions have relatively little predictive information for the upper quantiles of future GDP growth, as is visible in Figure 5. This strong asymmetry of the quantiles of GDP growth is striking. We study it more extensively in the next section.

3 Quantifying GDP Vulnerability

3.1 The Conditional GDP Distribution

The quantile regression (2) provides us with approximate estimates of the quantile function, an inverse cumulative distribution function. In practice, these estimates are difficult to map into a probability distribution function because of approximation error and estimation noise. 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:[6]

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

where $t(\cdot)$ and $T(\cdot)$ respectively denote the PDF and CDF of the Student t-distribution. The four parameters of the distribution pin down the location μ, scale σ, fatness ν, and shape α. Relative to the t-distribution, the skewed t-distribution adds the shape parameter which regulates the skewing effect of the CDF on the PDF. 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 a shape parameter α. The notable special case is the traditional t-distribution when $\alpha = 0$. In the case of both $\alpha = 0$ and $\nu = \infty$, the distribution reduces to a Gaussian

---

[6]An alternative approach to smoothing the quantile densities is to interpolate the quantile function using splines. Imposing monotonicity and smoothness requires additional modeling choices, as in for example Schmidt and Zhu (2016).

with mean μ and standard deviation σ. When $\nu = \infty$ and $\alpha \neq 0$, the distribution is a skewed normal.

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 our estimated quantile function $Q_{y_{t+h}|x_t}(\tau)$ from (2) and the quantile function of the skewed t-distribution $F^{-1}(\tau; \mu_t, \sigma_t, \alpha_t, \nu_t)$ from (3) to match the 5, 25, 75, and 95 percent quantiles

$$\{\hat{\mu}_{t+h}, \hat{\sigma}_{t+h}, \hat{\alpha}_{t+h}, \hat{\nu}_{t+h}\} = \underset{\mu,\sigma,\alpha,\nu}{\text{argmin}} \sum_{\tau} \left(\hat{Q}_{y_{t+h}|x_t}(\tau|x_t) - F^{-1}(\tau; \mu, \sigma, \alpha, \nu) \right)^2 \qquad (4)$$

where $\hat{\mu}_{t+h} \in \mathbb{R}$, $\hat{\sigma}_{t+h} \in \mathbb{R}^+$, $\hat{\alpha}_{t+h} \in \mathbb{R}$, and $\hat{\nu}_{t+h} \in \mathbb{Z}^+$.[7] This can be viewed as an exactly identified nonlinear cross sectional regression of the predicted quantiles on the quantiles of the skewed t-distribution.[8]

Figure 7 plots the estimated conditional quantile distribution $\hat{Q}_{y_{t+h}|x_t}(\tau|x_t)$ and two versions of the fitted inverse cumulative skewed t-distribution $F^{-1}(\tau; \hat{\mu}_{t+h}, \hat{\sigma}_{t+h}, \hat{\alpha}_{t+h}, \hat{\nu}_{t+h})$ – one conditional on both GDP growth and NFCI and one conditional on GDP growth alone – for three sample dates at different points of the business cycle: 2006Q2, which represented the end of the Federal Reserve's tightening cycle before 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. Figure 7 also shows that the distribution conditional on both economic and financial conditions can deviate substantially from the distribution conditional on economic conditions only. While the full conditional distribution is above the distribution conditional only on economic conditions during expansions, it is significantly below the distribution the conditional only on economic

---

[7]Notice that these parameters are functions of the conditioning variables x_t. We drop the explicit dependence for notational convenience.

[8]An alternative approach would be to use the entire 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.

conditions during recessions, especially in the left tail.

Figure 8 then plots the two versions of the fitted conditional probability density functions of GDP growth for the same three quarters. Comparing the conditional density across the three quarters, we see significant time variation in the density and that this time variation is primarily due to changes in the lower tail of the distribution. During business cycle upswings, such as in 2006Q2 and 2014Q4, the distribution conditional on both GDP growth and financial conditions has lower variance and greater positive skewness than the distribution conditional on GDP growth only. During downturns, such as 2008Q4, instead, the distribution conditional on both GDP growth and financial conditions has higher variance, greater negative skewness, and a lower mean than the distribution conditional on GDP growth only. These results might seem to contradict the recent evidence on the unpredictability of GDP growth during the Great Moderation (see e.g. D'Agostino, Giannone, and Surico, 2006; Rossi and Sekhposyan, 2010) and the weak and unstable predictive power of financial indicators (Stock and Watson, 2003). However, those studies focused on point forecasts, while we investigate the ability of current economic and financial conditions to predict the entire distribution of GDP growth.

In the online appendix, we also show the time series evolution of the fitted parameters in Figures A.1 and A.2. The time series pattern of the scale parameter σ_t resembles the financial conditions index most closely. The parameters $\hat{\alpha}_{t+h}, \nu_{t+h}$ that govern the shape and fatness of the distribution exhibit time series patterns that have low correlations with either the economic or the financial conditions index. This is because our fitting procedure is very nonlinear, thus giving rise to a time series pattern that would be difficult to detect using linear regressions or correlations. Furthermore, the α_{t+h}, ν_{t+h} parameters are relatively stable.

The picture that emerges is one where changes in location and scale are the most important determinants of the shifts in the conditional distribution. The strong time-variation in the lower quantiles of the distribution reflects the strong negative correlation between

11

the location and scale. Deteriorations in financial conditions coincide with declines in the location and increases in the scale of the distribution, and thus with a leftward shift of the distribution.

3.2 Measuring Vulnerability

The median of the predicted density provides the modal forecast for GDP growth next quarter and four quarters ahead. However, policy makers are often concerned with the downside and upside risks to the forecast or, in other words, how vulnerable the predicted path of GDP growth is to unexpected shocks. In this paper, we quantify upside and downside vulnerability of future GDP growth as the "extra" probability mass that the conditional density assigns to extreme right and left tail outcomes, relative to the probability of these outcomes under the unconditional density. By comparing the probability assigned to extreme outcomes by the conditional density to the probability assigned to the same outcomes by the unconditional density, we evaluate whether the predicted GDP distribution in a given quarter implies greater vulnerability around the modal forecast than the unconditional distribution.

More formally, we denote by $\hat{g}_{y_{t+h}}$ the unconditional density computed by matching the unconditional empirical distribution of GDP growth[9] and by $\hat{f}_{y_{t+h}|x_t}(y|x_t) = f(y; \hat{\mu}_{t+h}, \hat{\sigma}_{t+h}, \hat{\alpha}_{t+h}, \hat{\nu}_{t+h})$ the estimated skewed t-distribution. We define the upside, $\mathcal{L}_t^U$, and downside, $\mathcal{L}_t^D$, entropy of $\hat{g}_{y_{t+h}}(y)$ relative to $\hat{f}_{y_{t+h}|x_t}(y|x_t)$ as

---

[9]The unconditional density is time invariant and can be computed by to performing the two-step procedure where only the constant term is included in the quantile regression of the first step.

$$\mathcal{L}_t^D \left(\hat{f}_{y_{t+h}|x_t}; \hat{g}_{y_{t+h}} \right) = - \int_{-\infty}^{\hat{F}_{y_{t+h}|x_t}^{-1}(0.5|x_t)} \left(\log \hat{g}_{y_{t+h}}(y) - \log \hat{f}_{y_{t+h}|x_t}(y|x_t) \right) \hat{f}_{y_{t+h}|x_t}(y|x_t) \, dy,$$

(5)

$$\mathcal{L}_t^U \left(\hat{f}_{y_{t+h}|x_t}; \hat{g}_{y_{t+h}} \right) = - \int_{\hat{F}_{y_{t+h}|x_t}^{-1}(0.5|x_t)}^{\infty} \left(\log \hat{g}_{y_{t+h}}(y) - \log \hat{f}_{y_{t+h}|x_t}(y|x_t) \right) \hat{f}_{y_{t+h}|x_t}(y|x_t) \, dy,$$

(6)

where $\hat{F}_{y_{t+h}|x_t}(y|x_t)$ is the cumulative distribution associated with $\hat{f}_{y_{t+h}|x_t}(y|x_t)$ and $\hat{F}_{y_{t+h}|x_t}^{-1}(0.5|x_t)$ is the conditional median. Intuitively, downside entropy measures the divergence between the unconditional density and the conditional density that occurs below the median of the conditional density. When downside entropy is high, the conditional density assigns positive probability to more extreme left tail growth outcomes than the unconditional density. Similarly, upside entropy measures the divergence between the unconditional density and the conditional density that occurs above the median of the conditional density. When upside entropy is high, the conditional density assigns positive probability to more extreme right tail growth outcomes than the unconditional density. It is important to note that, unlike the full relative entropy between two distribution, upside and downside entropy can be negative, though not at the same time. That is, while the overall divergence between two distributions is positive, one density can put more mass in one tail of the distribution while the other puts more mass in another tail.

In Appendix A.1, we illustrate the properties of downside and upside entropies using two examples: one where GDP growth evolves according to a first-order autoregressive (AR(1)) process with normal innovations and the second in which GDP growth evolves according to an AR(1) process with innovations drawn from a mixture of normal distributions. When both the conditional and unconditional distributions of GDP growth are Gaussian, downside entropy is high when the median of the conditional distribution is lower than the median of the unconditional distribution. If the conditional and unconditional medians coincide, down-

side entropy equals upside entropy. Instead, when innovations to the conditional distribution are drawn from a mixture of truncated normals, the volatility conditional on GDP growth falling below the median is higher than the volatility conditional on GDP growth falling above the median, and upside entropy exceeds downside entropy even when the conditional and unconditional median of the GDP growth distribution coincide.

An alternative way of characterizing downside and upside risks to GDP growth is in terms of expected shortfall and its upper tail counterpart, which we term the expected longrise. For a chosen target probability π, the shortfall and longrise are defined, respectively, as

$$
SF_{t+h} = \frac{1}{\pi} \int_0^\pi \hat{F}_{y_{t+h}|x_t}^{-1}(\tau|x_t)d\tau; \qquad LR_{t+h} = \frac{1}{\pi} \int_{1-\pi}^1 \hat{F}_{y_{t+h}|x_t}^{-1}(\tau|x_t)d\tau.
$$

The information content of the relative entropy and the expected shortfall measures is distinct. While shortfall and longrise summarize the tail behavior of the conditional distribution in absolute terms, downside and upside entropy measure the tail behavior of the conditional distribution in excess of the tail behavior exhibited by the unconditional distribution. Thus, if both the unconditional and conditional distributions are negatively skewed, downside entropy will be low while expected shortfall will be high.

Figures 9a and 9b show the evolution of GDP upside and downside entropy one and four quarters ahead. Figures 9c and 9d plot the 5% expected shorfall and the 95% expected longrise. Despite differences in the information content of the two measures, they exhibit a surprising degree of similarity, indicating that the non-Gaussian features of the conditional distribution are largely absent from the unconditional 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 shortfall is significantly more volatile.

4 Out-of-sample Evidence and Alternative Approaches

4.1 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 done by using the proposed methodology in real time, with the caveat that we use final revised data only.[10]

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, expanding 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, estimating quantile regressions, matching the skewed t-distribution, and computing downside and upside entropy. The outcome of this procedure is a 20 year time-series of out-of-sample 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 quantiles 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

---

[10]Real-time data for the 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.

To assess the reliability of the predictive distribution, we measure the accuracy of a density forecast using the predictive score, computed as the predictive distribution generated by a model and evaluated at the realized value of the time series. Higher predictive scores indicate more accurate predictions because 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 predictive distribution conditional on both financial and economic conditions together with the scores of the predictive distribution conditional on economic conditions alone. The predictive score for the distribution conditional on both financial and economic conditions is frequently above that of the distribution conditional economic conditions only. Thus, the full conditional distribution is often more accurate – and rarely less accurate – than the one that conditions on economic conditions only, and the information contained in the conditioning variables 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 the 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 are presented in Figures 11c and 11d for both the conditional and unconditional distribution. Following Rossi and Sekhposyan (2017), we report confidence bands around the 45-degree line to account for sample uncertainty.[11] For both the full conditional predictive distribution and the predictive distribution that

---

[11]The confidence bands should be taken as 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 an expanding estimation window. For the one-quarter-ahead, the bands are based on the critical values derived under the null of uniformity and independence of the PIT. For the PITs of the one-year-ahead predictive distributions, bands are computed by bootstrapping under the assumption of uniformity only.

conditions on economic conditions only, the empirical distribution of the PITs is well within the confidence bands for the lower quantiles, though the empirical distribution falls outside the confidence band 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.

4.2   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. Our two step procedure is straightforward to estimate both in- and out-of-sample and is very flexible. In the online appendix, we further show that the procedure is robust to including alternative measures of financial conditions. In this subsection, we investigate two alternative econometric approaches: a fully parametric approach that estimates the conditional distribution of GDP growth via maximum likelihood and a fully non-parametric approach.

We begin with the fully parametric approach. Consider the following model

$$y_{t+1} = \gamma_0 + \gamma_1 x_t + \sigma_t \varepsilon_{t+1}; \qquad \qquad \ln\left(\sigma_t^2\right) = \delta_0 + \delta_1 x_t, \qquad (7)$$

where $\varepsilon_{t+1} \sim N(0,1)$ and x_t are conditioning variables. The model is estimated via maximum likelihood.[12] Figure 12a plots the conditional mean and the conditional lower and upper 5th quantiles for one-quarter-ahead GDP growth implied by the model in equation (7). The simple conditionally heteroskedastic 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.

Both the quantile regression and the conditionally heteroskedastic model are linear in

---

[12]We have also estimated the model with ARCH and GARCH components of Engle (1982), but it turns out that in quarterly data, the conditioning variables drive out those effects.

the conditioning variables, a condition which might seem restrictive. We obtain the same qualitative results more directly—less parametrically—allowing for general dependence of the quantiles on the conditioning variables. In particular, following Li, Lin, and Racine (2013) we estimate the conditional cumulative density function of one-quarter-ahead GDP growth y_{t+1} as a function of current quarter conditioning variables $x_t = (x_{1,t}, ..., x_{n,t})'$ as

$$\hat{F}_{y_{t+1}|x_t}\left(y|\,x\right) = \frac{\frac{1}{T-1}\sum_{t=1}^{T-1}\Phi\left(\frac{y-y_{t+1}}{\omega_0}\right)K_\omega\left(x_t,x\right)}{\frac{1}{T-1}\sum_{t=1}^{T-1}K_\omega\left(x_t,x\right)}, \tag{8}$$

where

$$K_\omega\left(x_t,x\right) = \prod_{i=1}^{n}\frac{1}{\omega_i}\phi\left(\frac{x_i - x_{i,t}}{\omega_i}\right),$$

and $\phi\left(\cdot\right)$ and $\Phi\left(\cdot\right)$ are the standard normal pdf and cdf, respectively; T is the number of observations; and $\omega_0, \omega_1, \ldots, \omega_n$ are bandwidths chosen by cross-validation, as described in Li et al. (2013).[13]

Figure 12b plots the conditional mean and the conditional lower and upper 5th quantiles for one-quarter-ahead GDP growth implied by the model in equation (8). As with the ML estimates, the conditional quantiles obtained using the less parametric estimation procedure produce the general feature that the right tail of the distribution is stable while the left tail moves with financial conditions. However, the fluctuations of the left tail are smoother, due to a conservative bandwidth selection that is necessary to counteract the tendency of the very flexible model to overfit the data.[14]

Overall, the results for both the fully parametric and the fully non-parametric models

---

[13]The authors establish optimality of this data-driven selection method under the assumption of independence. Although serial dependence might be a concern in our application, recent work by Li, Ouyang, and Racine (2009) on nonparametric regression with weakly dependent data suggests that these optimality results may still be valid.

[14]In Figure A.8 in the Online Appendix, we further show that both the ML approach and the non-parametric approach perform well out-of-sample, though slightly less so than the semi-parametric conditional quantile approach. Indeed, the scores of the predictive distributions are somewhat smaller.

confirm that our main findings are not an artifact of the two-step quantile regression procedure but reflect a robust property of the distribution of GDP growth conditional on financial conditions.

Alternative ways to estimate conditional GDP distributions have been proposed in the literature, such as the two state Markov chain of Hamilton (1989). Hamilton's model could be augmented to include financial conditioning variables that shift the transition probabilities between the two states using the estimation method of Filardo (1994). More generally, Chen, Fan, and Tsyrennikov (2006) present a semi-parametric estimator of multivariate copula models that would allow the mixing between distributions as a function of financial condition variables (see Smith and Vahey, 2016 for an application to modeling GDP growth). Another strand of the macroeconomic literature introduces time varying volatility into macroeconomic dynamics following the Bayesian vector autoregression of Primiceri (2005). For example, Clark (2012) features an exogenously time varying stochastic volatility process in a Bayesian vector autoregression. The incorporation of financial conditions in the volatility process in such Bayesian vector autoregressions might give rise to conditional density forecasts similar to the ones we obtain. Finally, alternative semi-parametric and non-parametric methods could be used, including the sieve regressions of Chen, Liao, and Sun (2014) or the efficient non-parametric methods of Li and Racine (2007) and Norets and Pati (2017).

5 Implications for Theories of Macroeconomic Dynamics

Our main finding is that, while economic conditions forecast median GDP growth, downside risks to GDP growth are predicted by financial conditions and upside risks are stable over time. This finding presents challenges to traditional macroeconomic modeling because it points to the importance of financial conditions above economic conditions in predicting the future evolution of GDP growth. Measures of financial conditions may help forecast downside risks to GDP growth because they capture frictions in either the supply of or the demand

for credit in the economy, or because financial conditions represent a more informative signal about potential future risks. In this section, we discuss these theoretical alternatives.

We begin with the literature on 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), which emphasizes the role that frictions play in the *demand* for credit. In these models, which rely on asymmetric information between lenders and borrowers and moral hazard on the part of borrowers, a tightening in the stance of monetary policy leads to a decrease in the net worth of productive firms, causing adverse selection and moral hazard problems to worsen. This generates an asymmetry in the business cycle, with tighter financial conditions preceding economic downturns. However, this literature has the drawback that, 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, we find that financial conditions are particularly correlated with the higher moments of GDP growth, suggesting that non-linear approximations to the solutions of such models may be more appropriate for capturing the empirical distribution of GDP growth.

The Great Recession prompted the literature on credit frictions to also consider the role that financial conditions play in the *supply* of credit in the economy. Brunnermeier and Sannikov (2014), He and Krishnamurthy (2012) and Adrian and Boyarchenko (2012) consider production economies where the productive sector relies on a financial sector for the supply of credit. In these models, changes in the vulnerability of the financial sector lead to decreases in the supply of credit to the economy, potentially causing a downturn in GDP growth. The simplicity of the production sector considered allows these models to be solved analytically without relying on (log-)linearization and to instead emphasize the non-linear amplification channel created by constraints on the intermediary sector. This literature also shows that, in the presence of leverage constraints on intermediary capital,

a volatility paradox arises, in which times of low volatility of output growth coincide with high average output growth and foreshadow times of low average output growth.

The conditional distribution of GDP growth we construct is consistent with these predictions. 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. These two effects combined generate a negatively skewed unconditional distribution, as downward shifts in the growth outlook are associated with an increase in risk.

Figure 14 further 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 negative correlation between the conditional mean and volatility is statistically significant for two years leads and two years lags of volatility. Thus, the negative association between risk and return is a robust feature at multiple horizons. At the one quarter horizon, a one-standard-deviation increase in the conditional volatility is associated with a 0.8-standard-deviation decrease in the next quarter's predicted conditional mean of GDP growth.

Finally, financial conditions may predict downside risk to GDP growth because financial conditions may provide a better signal of non-normal shocks to the economy. Models of the impact of Bayesian learning on aggregate outcomes (see e.g. Orlik and Veldkamp, 2014 and Johannes et al., 2016) do have the potential to generate a negative correlation between the conditional mean and the conditional volatility of GDP growth. However, our results emphasize the importance of allowing agents to use financial conditions as a signal in forming their beliefs, as we find conditional volatility to be predicted by financial, not economic, conditions.

One possible source of non-normal shocks that could be captured by financial rather than economic conditions is disaster risk. Catastrophic tail events, such as the Great De-

pression or the 2008 financial crisis, can be used to generate a large equity risk premium (see e.g. Barro, 2009; Gabaix, 2012; Wachter, 2013 and Barro and Ursúa, 2012). In turn, risk and uncertainty aversion to extreme tail events impact macroeconomic dynamics in equilibrium (see e.g. Gourio, 2012). In contrast, our evidence suggests that GDP vulnerability changes at relatively high frequencies and is thus unlikely to be driven by changes in the exposure to disaster risk. However, higher frequency changes in downside GDP vulnerability are consistent with the more recent evidence from the term structure of asset prices: The slopes of term structures of risk premia across multiple asset classes suggest 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, Forthcoming).

6   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, with the left tail of the distribution positively correlated with slack in financial conditions. We measure the vulnerability of GDP growth to downside risks as relative entropy of the unconditional relative to conditional predictive distribution, and show that growth vulnerability is correlated with financial conditions.

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 FOMC statements, and show that financial conditions are a significant explanatory variable in augmented Taylor rules. Caldara and Herbst (2016) argue that monetary policy can be characterized by a direct and economically significant reaction to changes in credit spreads. 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. Measuring downside growth vulnerability helps quantify the cost side of that tradeoff.

References

ADRIAN, T. AND N. BOYARCHENKO (2012): "Intermediary Leverage Cycles and Financial Stability," Staff Report N. 567, Federal Reserve Bank of New York.

ADRIAN, T. AND N. LIANG (2016): "Monetary Policy, Financial Conditions, and Financial Stability," Staff Report N. 690, Federal Reserve Bank of New York.

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 (Forthcoming): "Term structures of asset prices and returns," *Journal of Financial Economics*.

BAKER, S. R., N. BLOOM, AND S. J. DAVIS (2016): "Measuring economic policy uncertainty," *The Quarterly Journal of Economics*, 131, 1593–1636.

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.

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.

CALDARA, D. AND E. HERBST (2016): "Monetary Policy, Real Activity, and Credit Spreads: Evidence from Bayesian Proxy SVARs," Finance and Economics Discussion Series 2016-049, Board of Governors of the Federal Reserve System.

CHEN, X., Y. FAN, AND V. TSYRENNIKOV (2006): "Efficient Estimation of Semiparametric Multivariate Copula Models," *Journal of the American Statistical Association*, 101, 1228–1240.

CHEN, X., Z. LIAO, AND Y. SUN (2014): "Sieve Inference on Possibly Misspecified Seminonparametric Time Series Models," *Journal of Econometrics*, 178, 639–658.

CLARK, T. (2012): "Real-time Density Forecasts from Bayesian Vector Autoregressions with Stochastic Volatility," *Journal of Business & Economic Statistics*, 29, 327–341.

CLARK, T., A. CARRIERO, AND M. MASSIMILIANO (2016): "Measuring Uncertainty and Its Impact on the Economy," Working Paper 1622, Federal Reserve Bank of Cleveland.

COGLEY, T., S. MOROZOV, AND T. SARGENT (2005): "Bayesian fan charts for U.K. inflation: Forecasting and sources of uncertainty in an evolving monetary system," *Journal of Economic Dynamics and Control*, 29, 1893–1925.

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

D'AGOSTINO, A., L. GAMBETTI, AND D. GIANNONE (2013): "Macroeconomic forecasting and structural change," *Journal of Applied Econometrics*, 28, 82–101.

D'AGOSTINO, A., D. GIANNONE, AND P. SURICO (2006): "(Un)Predictability and macroeconomic stability," Working Paper Series 605, European Central Bank.

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," Staff Report N. 647, Federal Reserve Bank of New York.

DOZ, C., D. GIANNONE, AND L. REICHLIN (2012): "A Quasi–Maximum Likelihood Approach for Large, Approximate Dynamic Factor Models," *Review of Economics and Statistics*, 94, 1014–1024.

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

FILARDO, A. J. (1994): "Business-cycle phases and their transitional dynamics," *Journal of Business & Economic Statistics*, 12, 299–308.

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.

GIGLIO, S., B. KELLY, AND S. PRUITT (2016): "Systemic risk and the macroeconomy: An empirical evaluation," *Journal of Financial Economics*, 119, 457–471.

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

HAMILTON, J. (1989): "A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle," *Econometrica*, 357–384.

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

JOHANNES, M., L. A. LOCHSTOER, AND Y. MOU (2016): "Learning about consumption dynamics," *The Journal of Finance*, 71, 551–600.

JURADO, K., S. LUDVIGSON, AND S. NG (2015): "Measuring Uncertainty," *American Economic Review*, 105, 1177–1216.

KIM, C.-J. AND C. R. NELSON (1999): "Has the US economy become more stable? A Bayesian approach based on a Markov-switching model of the business cycle," *The Review of Economics and Statistics*, 81, 608–616.

KOENKER, R. AND G. BASSETT (1978): "Regression Quantiles," *Econometrica*, 46, 33–50.

LI, C., D. OUYANG, AND J. S. RACINE (2009): "Nonparametric regression with weakly dependent data: the discrete and continuous regressor case," *Journal of Nonparametric Statistics*, 21, 697–711.

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, 57–65.

LI, Q. AND J. RACINE (2007): *Nonparametric econometrics: Theory and practice*, Princeton University Press.

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.

NORETS, A. AND D. PATI (2017): "Adaptive Bayesian estimation of conditional densities," *Econometric Theory*, 33, 980–1012.

ORLIK, A. AND L. VELDKAMP (2014): "Understanding Uncertainty Shocks and the Role of Black Swans," Working Paper N. 20445, National Bureau of Economic Research.

PEEK, J., E. ROSENGREN, AND G. TOOTELL (2015): "Should U.S. Monetary Policy Have a Ternary Mandate?" in *Federal Reserve Bank of Boston 59th Economic Conference: Macroprudential Monetary Policy*.

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

ROSSI, B. AND T. SEKHPOSYAN (2010): "Have economic models' forecasting performance for US output growth and inflation changed over time, and when?" *International Journal of Forecasting*, 26, 808–835.

——— (2017): "Alternative tests for correct specification of conditional predictive densities," Working Paper, Texas A&M University.

SCHMIDT, L. AND Y. ZHU (2016): "Quantile Spacings: A Simple Method for the Joint Estimation of Multiple Quantiles Without Crossing," Working Paper, University of Chicago.

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

STOCK, J. AND M. WATSON (2003): "Forecasting Output and Inflation: The Role of Asset Prices," *Journal of Economic Literature*, 41, 788–829.

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 figure shows the time series of the quarterly growth rate of real GDP growth and the NFCI.
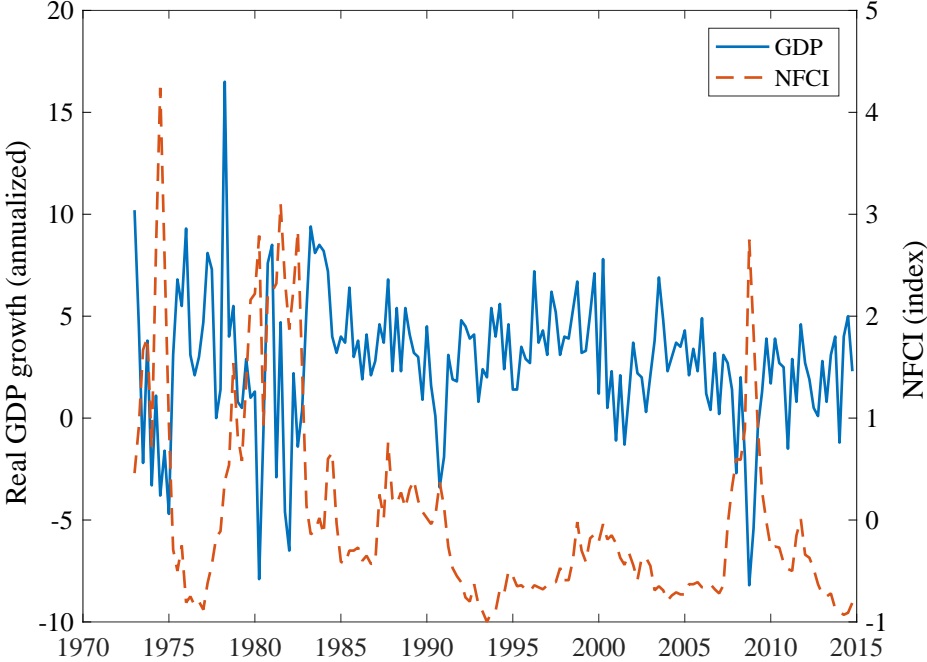
Figure 3. Quantile Regressions. The figure shows the univariate quantile regressions of one quarter ahead (left column) and four quarter ahead (right column) real GDP growth on current real GDP growth and the NFCI.

(a) One quarter ahead: GDP

(b) One year ahead: GDP

(c) One quarter ahead: NFCI

(d) One year ahead: NFCI
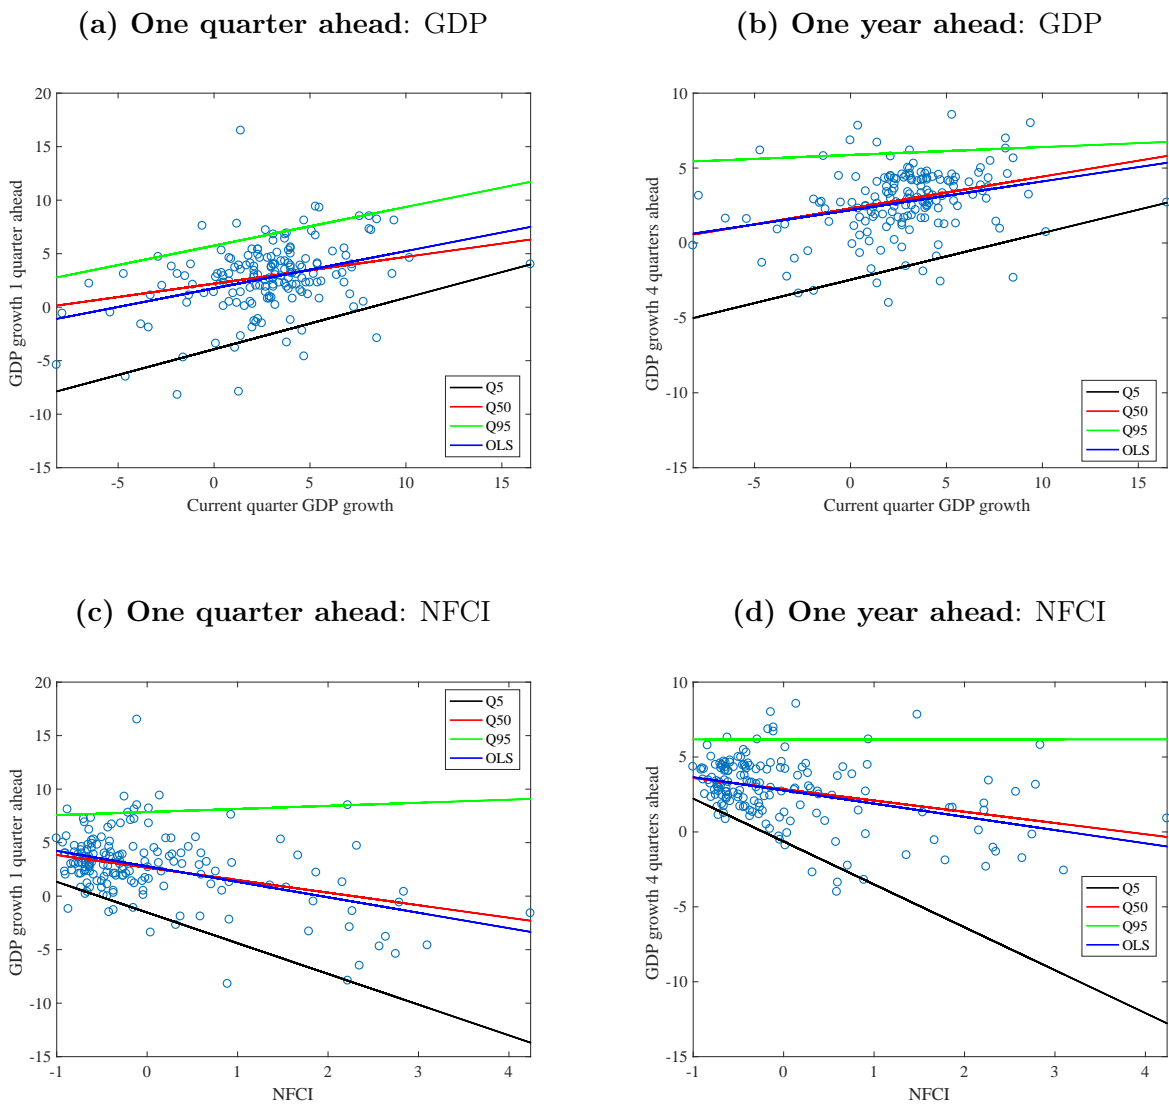
Figure 4. Estimated Quantile Regression Coefficients. The figure shows the estimated coefficients in quantile regressions of one quarter ahead (left column) and one year ahead (right column) real GDP growth on current real GDP growth and NFCI. We report confidence bounds for the null hypothesis that the true data-generating process is a general, flexible linear model for growth and financial conditions (VAR with 4 lags); bounds are computed using 1000 bootstrapped samples.

(a) **One quarter ahead**: GDP

(b) **One year ahead**: GDP

(c) **One quarter ahead**: NFCI

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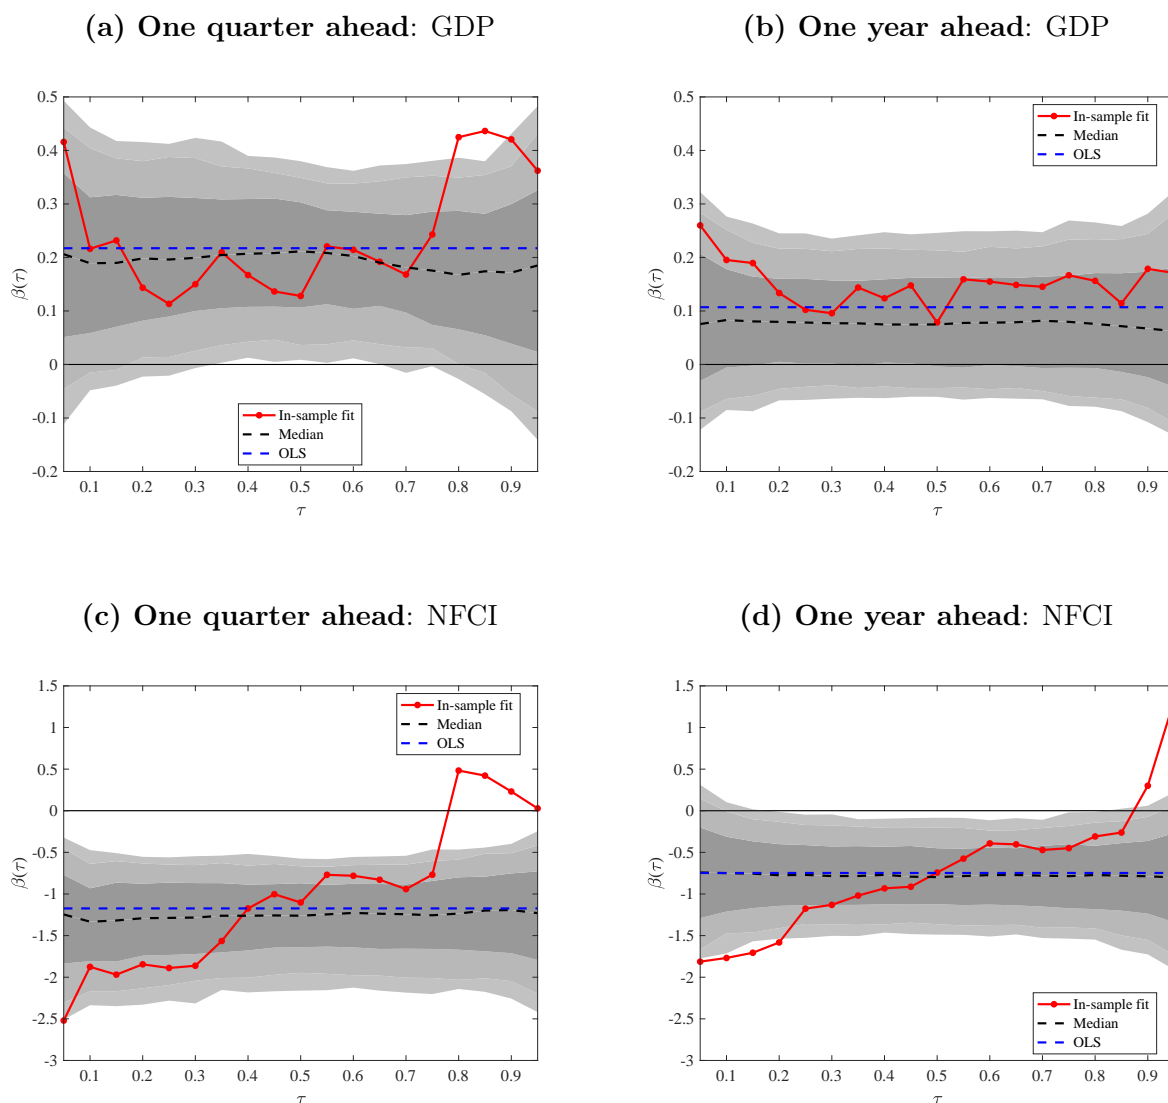
Figure 5. Predicted Distributions. The figure shows the time series evolution of the predicted distribution of one quarter ahead (panel a) and four quarter ahead (panel b) real GDP growth.

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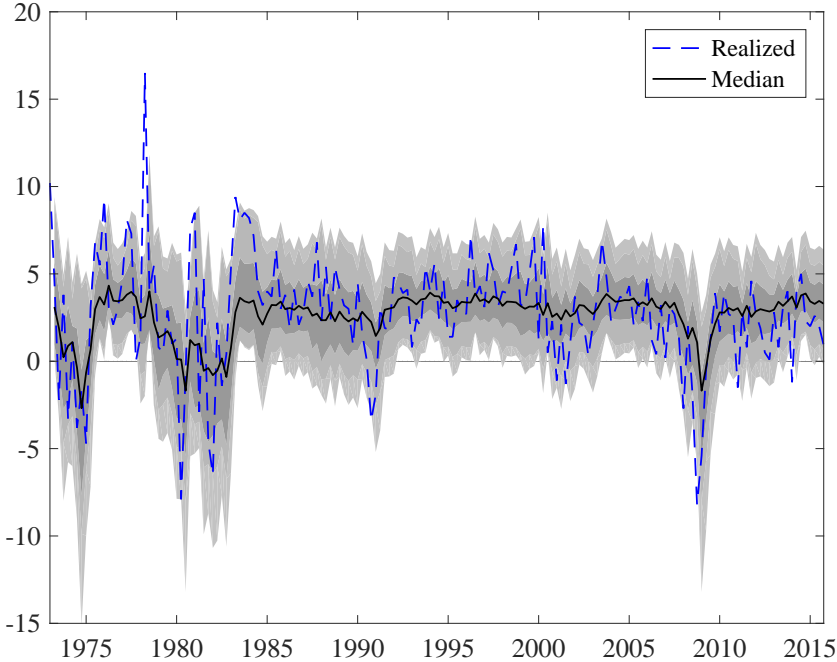
(b) One year ahead
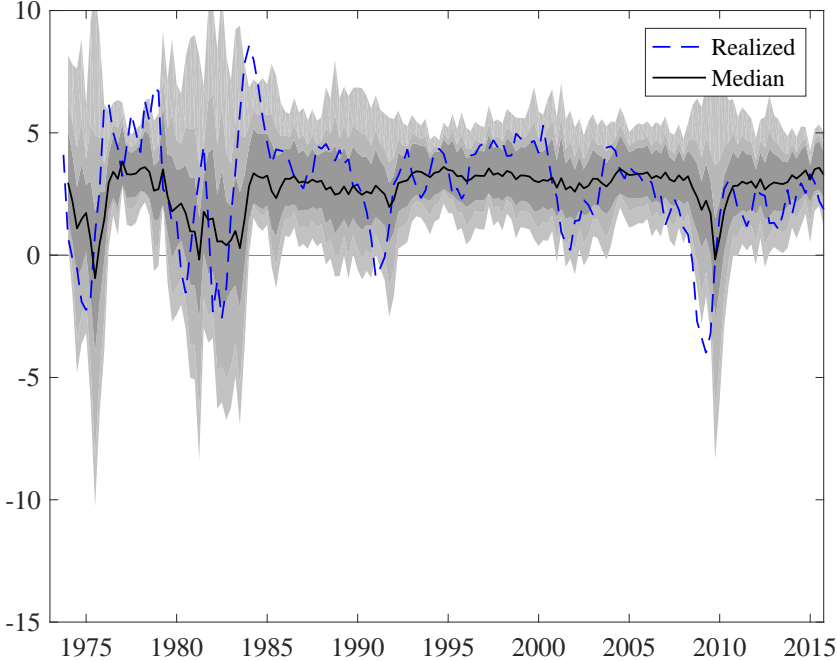
Figure 6. Median, Interquartile Range, and 5 Percent Quantile of the Predicted Distribution. The figure shows scatter plots of the median versus the interquartile range (6a and 6b) and the 5 percent quantile versus the interquartile range (6c and 6d).

(a) Median and Interquartile Range
One quarter ahead
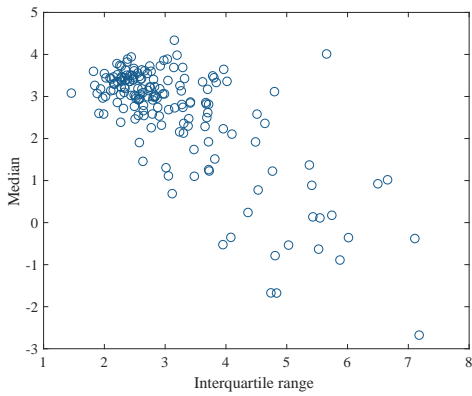
(b) Median and Interquartile Range
One year ahead
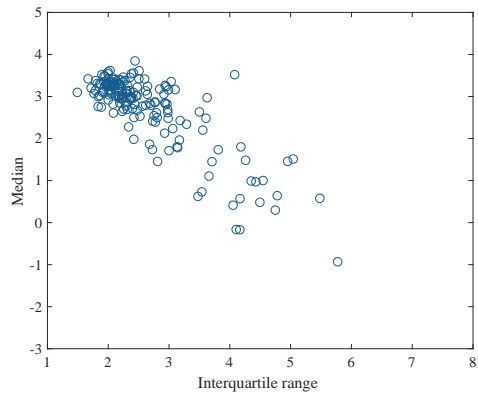
(c) Q5 and Interquartile Range
One quarter ahead
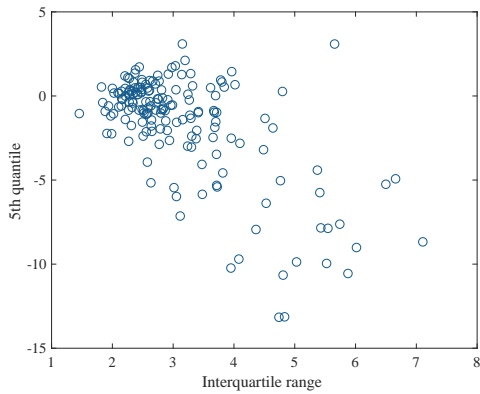
(d) Q5 and Interquartile Range
One year ahead
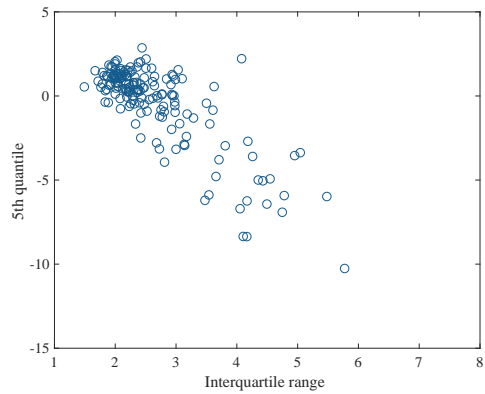
Figure 7. The Conditional Quantiles and the Skewed t-Distribution. The panels in this figure show the conditional quantiles together with the estimated skewed t-inverse cumulative distribution functions for one quarter and one year ahead GDP growth. For comparison, we also report the skewed t-inverse cumulative distribution functions obtained by fitting the quantiles obtained by conditioning only on current real GDP growth.

(a) One quarter ahead: Q2 2006

(b) One year ahead: Q2 2006

(c) One quarter ahead: Q4 2008

(d) One year ahead: Q4 2008

(e) One quarter ahead: Q4 2014

(f) One year ahead: Q4 2014
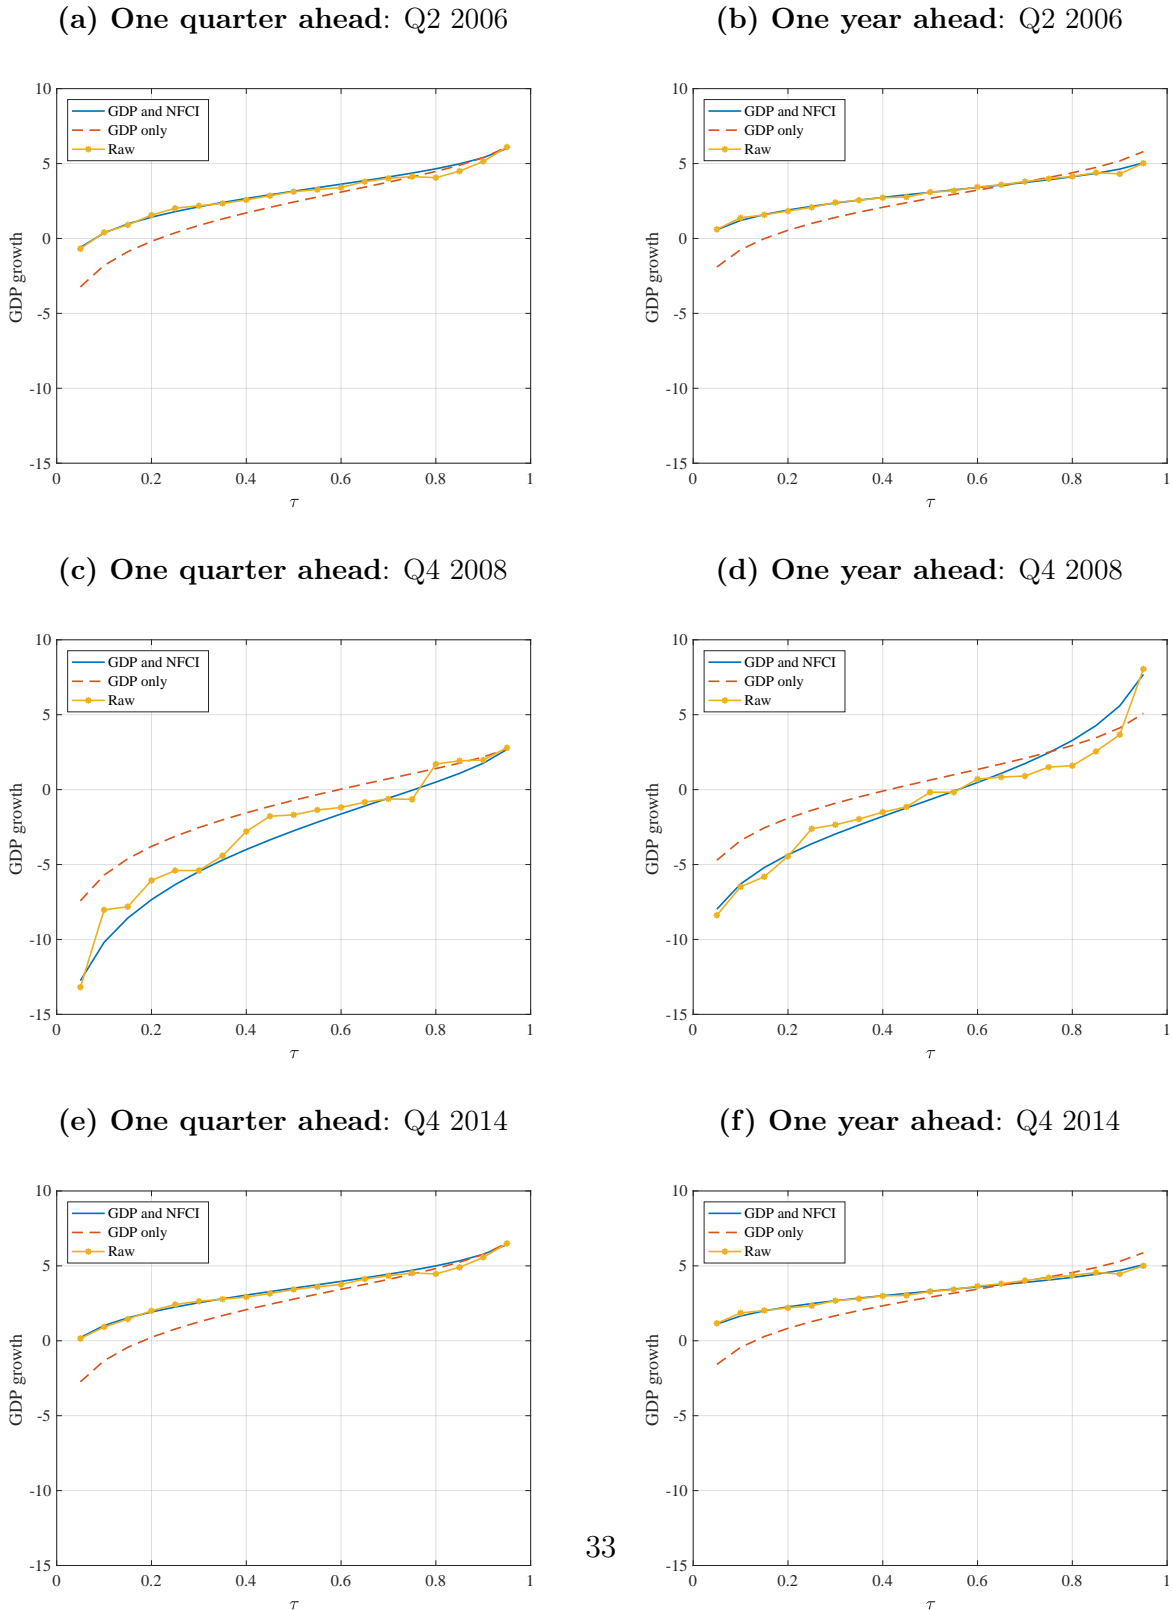
Figure 8. Probability Densities. The panels in this figure show the estimated skewed *t*-density functions for one quarter and one year ahead real GDP growth, with the density estimated conditional on current real GDP growth and NFCI. For comparison, we also report the skewed *t*-density functions obtained by conditioning on current real GDP growth only.

(a) One quarter ahead: Q2 2006
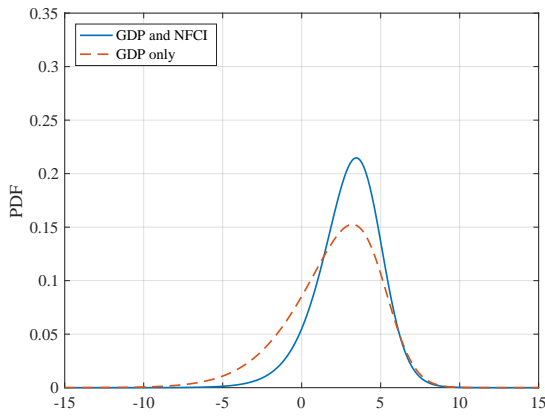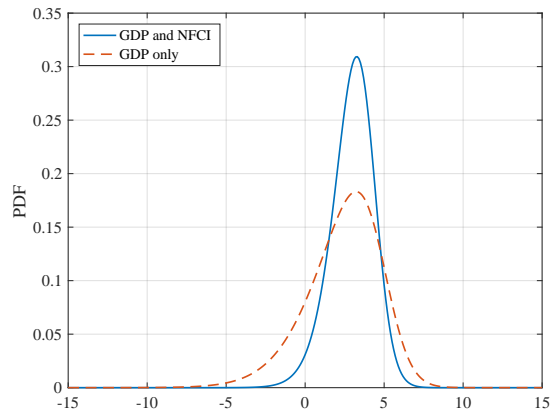
(b) One year ahead: Q2 2006

(c) One quarter ahead: Q4 2008
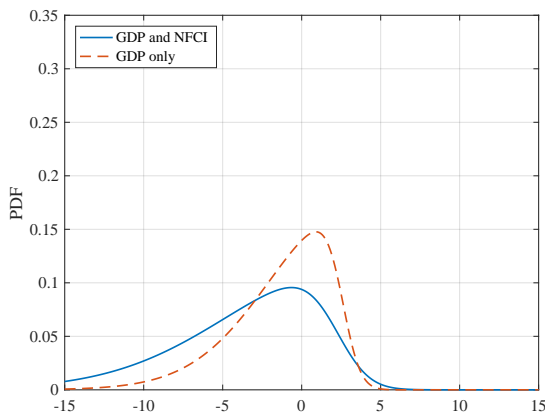
(d) One year ahead: Q4 2008
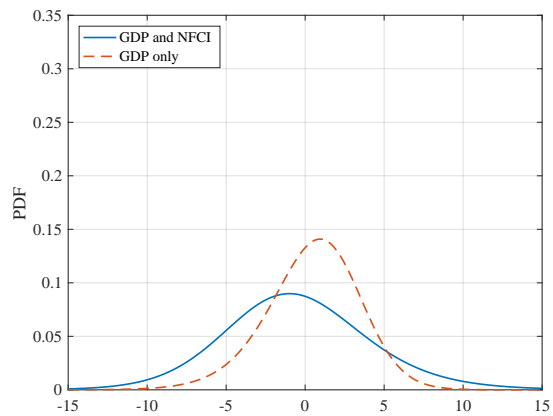
(e) One quarter ahead: Q4 2014
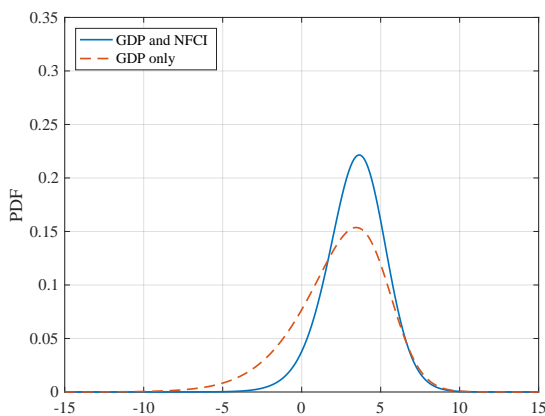
(f) One year ahead: Q4 2014
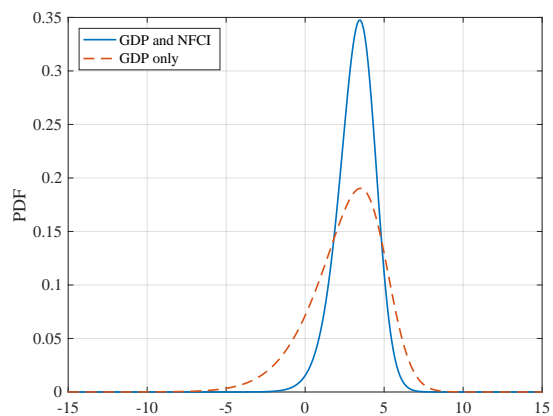
Figure 9. Growth Entropy and Expected Shortfall over Time. The figure shows the time series evolution of relative downside and upside entropy $\mathcal{L}_t^D$ and $\mathcal{L}_t^U$ together with the 5% expected shortfall ES_t.

(a) Entropy $\mathcal{L}_t$: One quarter ahead
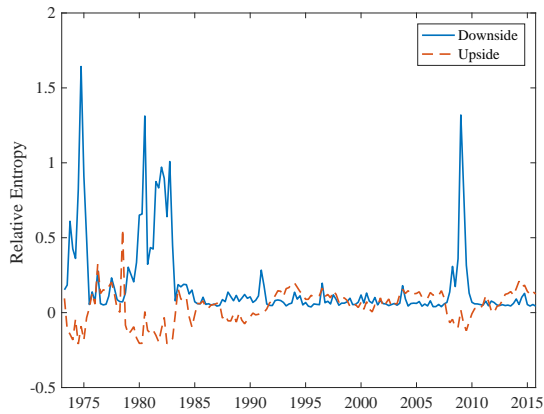
(b) Entropy $\mathcal{L}_t$: One year ahead
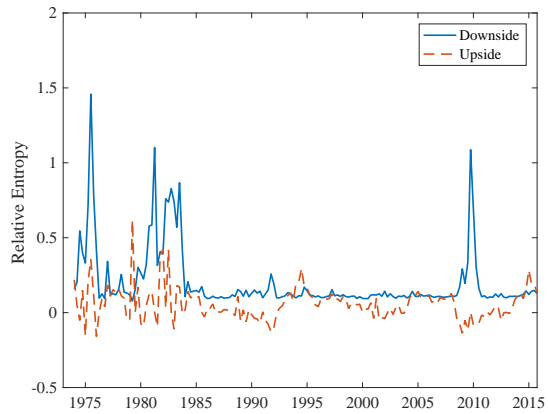
(c) 5% Expected Shortfall, ES_t, and Longrise, EL_t: One quarter ahead
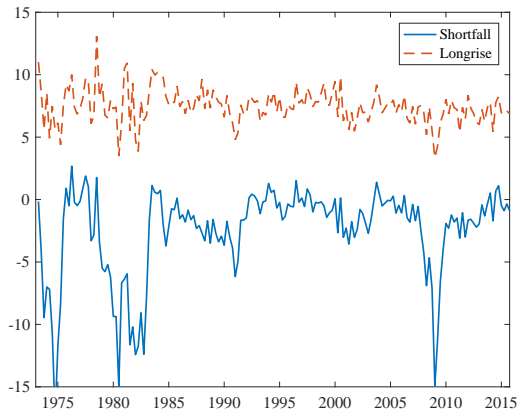
(d) 5% Expected Shortfall, ES_t, and Longrise, EL_t: One year ahead
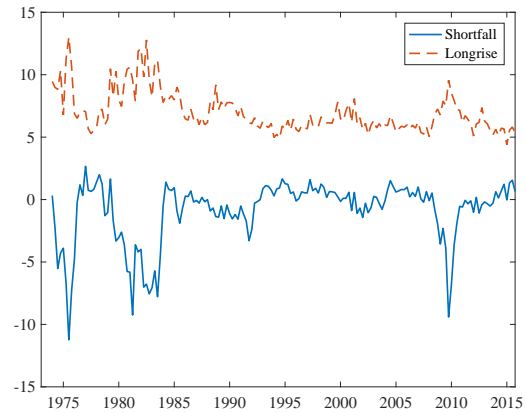
Figure 10. Out-of-Sample Predictions. The figure compares out-of-sample and in-sample predictive densities for GDP growth one and four quarter ahead. Figures 10a and 10b show the 5, 50, and 95 percent quantiles. Figures 10c and 10d show downside entropy.

(a) Quantiles: One quarter ahead
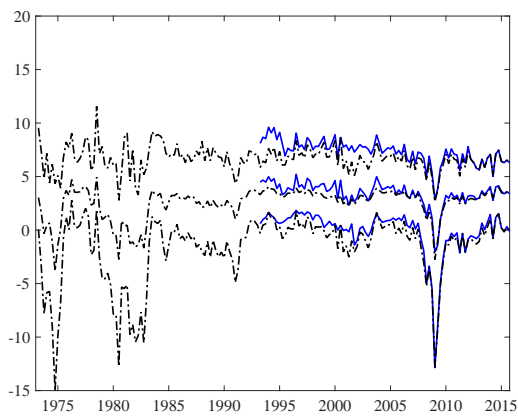
(b) Quantiles: One year ahead
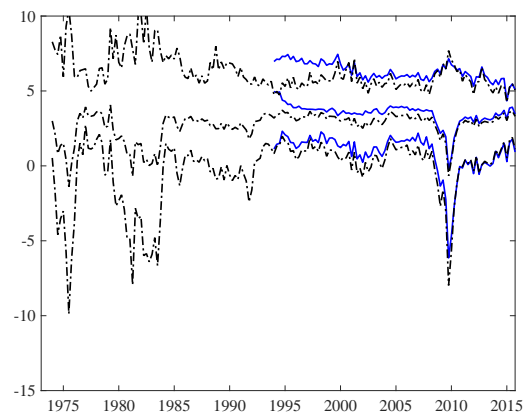
(c) Downside Entropy: One quarter ahead
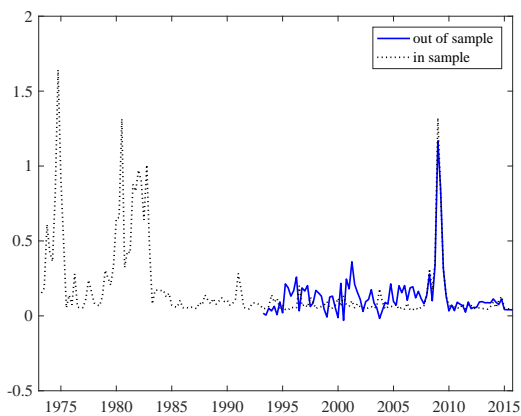
(d) Downside Entropy: One year ahead
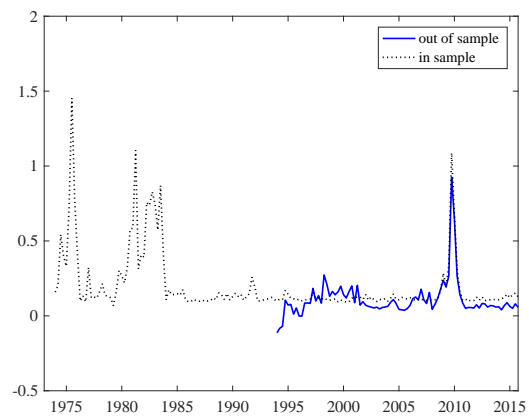
Figure 11. Out-of-Sample Accuracy. The figure reports the predictive scores and the cumulative distribution of the probability integral transform. Figures 11a and 11b compare the out- of-sample predictive scores of the predictive distribution conditional on both NFCI and real GDP growth, and the predictive distribution conditional on real GDP growth only. Figures 11c and 11d report the empirical cumulative distribution of the probability integral transform (PITs). Critical values are obtained as in Rossi and Sekhposyan (2017)
.

(a) Predictive scores: One quarter ahead   **(b) Predictive scores**: One year ahead
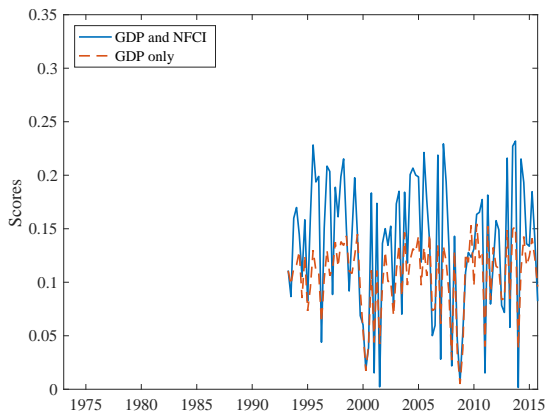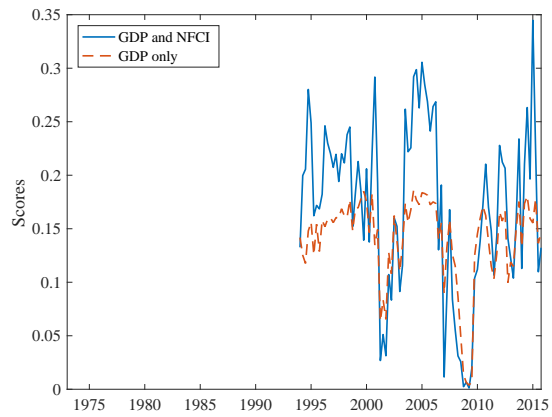
(c) PITs: One quarter ahead   **(d) PITs**: One year ahead
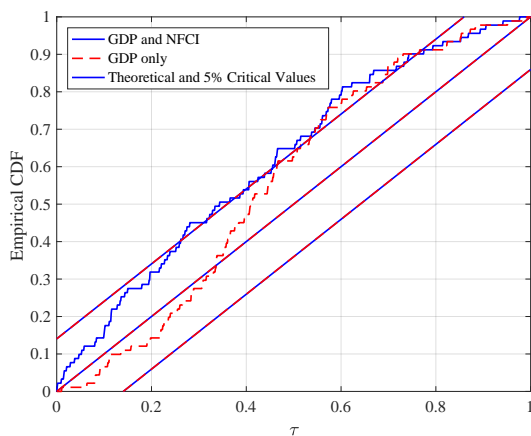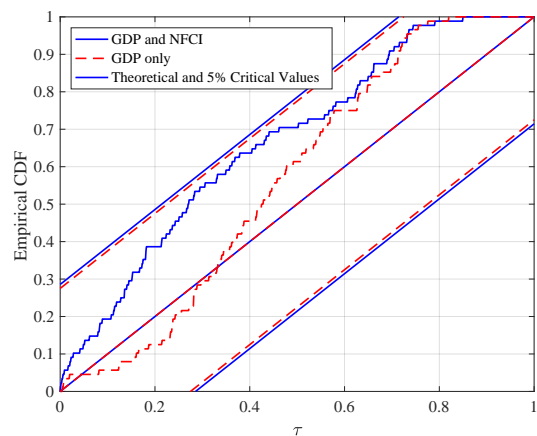
Figure 12. Alternative Econometric Approaches: Predicted Distributions. The figure shows the conditional mean, and the 95 percent and 5 percent standard error bands for GDP growth one-quarter-ahead. In the top panel, the distributions are estimated with a conditionally Gaussian model with conditioning variables in the mean equation and the volatility equations. In the bottom panel, the distributions are estimated non-parametrically.

(a) MLE
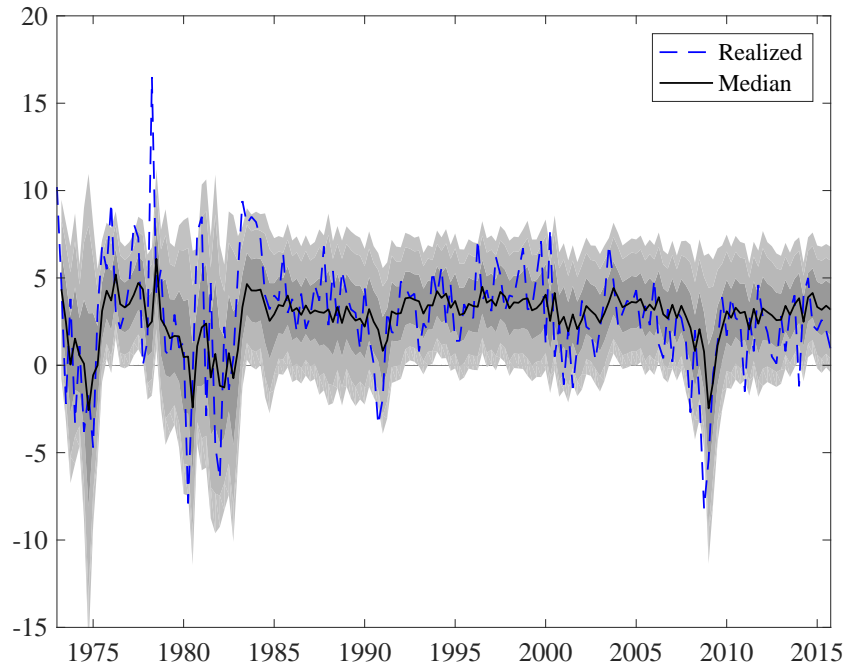
(b) Non-parametric
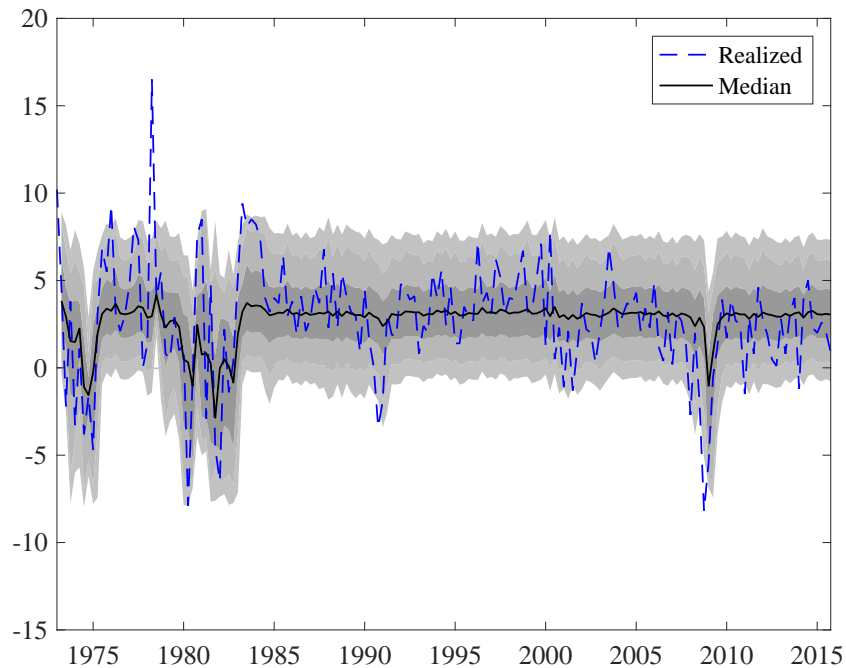
Figure 13. Mean, Volatility, and Entropy. The figure shows scatter plots of the mean versus the volatility (13a and 13b) and downside entropy versus the mean (13a and 13b) of the one-quarter-ahead and the one-year-ahead distributions of real GDP growth.

(a) Mean and Volatility: One quarter ahead
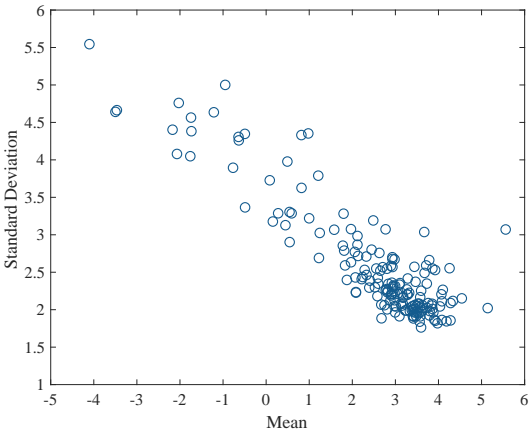
(b) Mean and Volatility: One year ahead
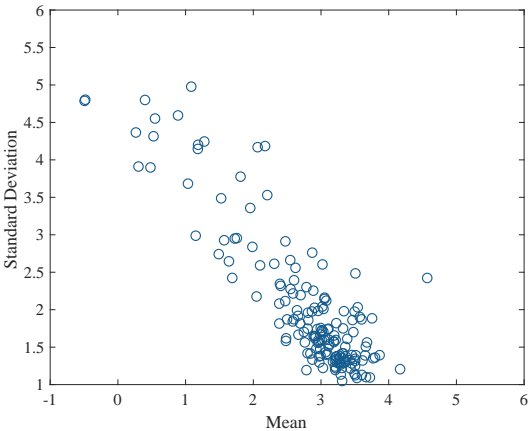
(c) Mean and Entropy: One quarter ahead
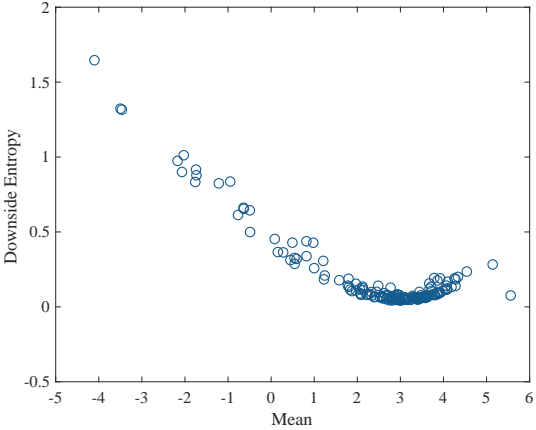
(d) Mean and Entropy: One year ahead
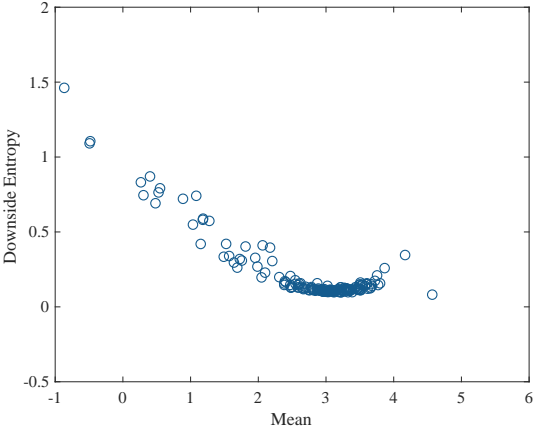
Figure 14. Median, Interquartile Range, and 5 Percent Quantile. The figure shows the cross correlograms between the mean and volatility, and mean and downside entropy.

(a) Cross Correlogram Mean and Volatility: One quarter ahead
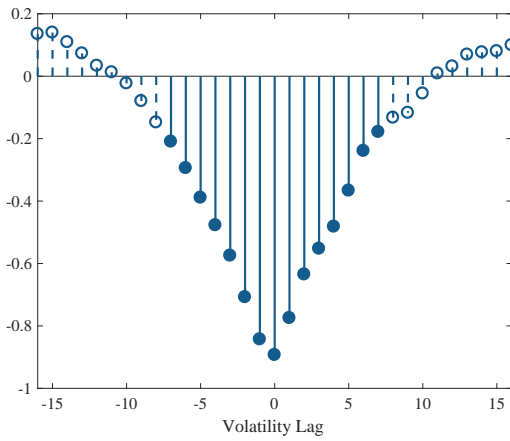
(b) Cross Correlogram Mean and Volatility: One year ahead
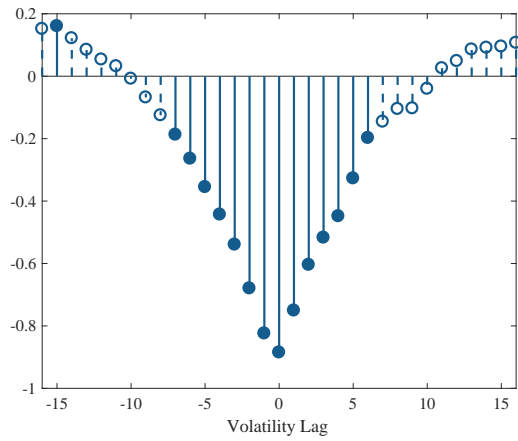
(c) Cross Correlogram Mean and Entropy: One quarter ahead
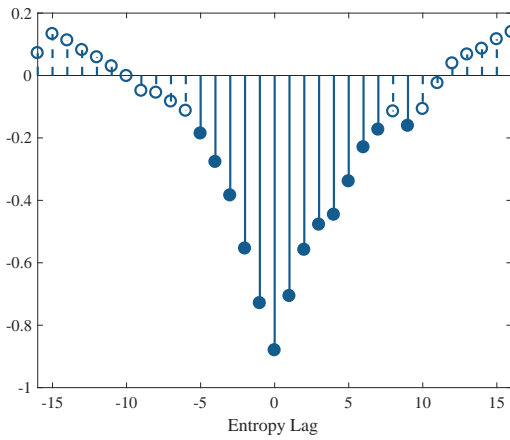
(d) Cross Correlogram Mean and Entropy: One year ahead
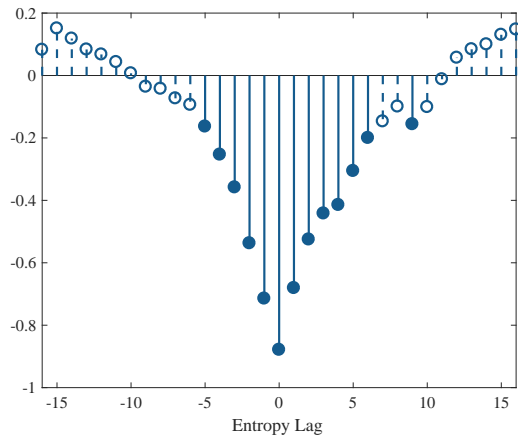
A Internet Appendix

A.1 Downside and Upside Entropy

In this section, we illustrate the properties of downside and upside entropy using two examples: one where GDP growth evolves according to a first-order autoregressive (AR(1)) process with normal innovations and one where GDP growth evolves according to an AR(1) process with non-normal innovations. These examples illustrate that downside entropy depends on the moments of the distribution of GDP growth conditional on the realization of GDP growth being below the median, and upside entropy depends on the moments of the distribution of GDP growth conditional on the realization of GDP growth being above the median.

Example 1. Assume that GDP growth evolves as an AR(1) process with normal innovations, so that

$$y_{t+1} = \mu_g + \varphi_g y_t + \epsilon_{t+1}, \qquad\qquad \epsilon \sim \mathcal{N}\left(0, \sigma_\epsilon^2\right),$$

where y_{t+1} is the one quarter annualized average growth rate of GDP and $\varphi_g < 1$ is the speed of mean-reversion of GDP growth to the unconditional growth rate $(1 - \varphi_g)^{-1} \mu_g$. In this case, the (one-step-ahead) conditional distribution of GDP growth is

$$f_t(y) = \left(2\pi\sigma_\epsilon^2\right)^{-\frac{1}{2}} \exp\left(-\frac{1}{2\sigma_\epsilon^2}(y - \mu_g - \varphi_g y_t)^2\right),$$

and the unconditional distribution is

$$g(y) = \sqrt{1 - \varphi_g^2}\left(2\pi\sigma_\epsilon^2\right)^{-\frac{1}{2}} \exp\left(-\frac{1 - \varphi_g^2}{2\sigma_\epsilon^2}\left(y - \frac{\mu_g}{1 - \varphi_g}\right)^2\right).$$

The downside and upside entropy in this case is given by

$$
\begin{aligned}
\mathcal{L}_t^D(f_t; g) =\ & -\int_{-\infty}^{\mu_t}\left(\log\sqrt{1 - \varphi_g^2} - \frac{1 - \varphi_g^2}{2\sigma_\epsilon^2}\left(y - \frac{\mu_g}{1 - \varphi_g}\right)^2 + \frac{1}{2\sigma_\epsilon^2}(y - \mu_t)^2\right) f_t(y)dy \\
=\ & -\left(\log\sqrt{1 - \varphi_g^2} - \frac{1 - \varphi_g^2}{2\sigma_\epsilon^2}\left(m_{-,t} - \frac{\mu_g}{1 - \varphi_g}\right)^2 + \frac{(m_{-,t} - \mu_t)^2}{2\sigma_\epsilon^2}\right)\int_{-\infty}^{\mu_t} f_t(y)dy \\
& -\frac{1}{\sigma_\epsilon^2}\left(m_{-,t} - \mu_t - (1 - \varphi_g^2)\left(m_{-,t} - \frac{\mu_g}{1 - \varphi_g}\right)\right)\int_{-\infty}^{\mu_t}(y - m_{-,t})f_t(y)dy \\
& + \frac{\varphi_g^2}{2\sigma_\epsilon^2}\int_{-\infty}^{\mu_t}(y - m_{-,t})^2 f_t(y)dy \\
=\ & -\frac{1}{2}\left(\log\sqrt{1 - \varphi_g^2} - \frac{1 - \varphi_g^2}{2\sigma_\epsilon^2}\left(m_{-,t} - \frac{\mu_g}{1 - \varphi_g}\right)^2 + \frac{1}{\pi}\right) + \frac{\varphi_g^2}{4}\left(1 - \frac{2}{\pi}\right);
\end{aligned}
$$

1

$$\mathcal{L}_t^U \left(f_t; g \right) = -\int_{\mu_t}^{\infty} \left(\log \sqrt{1 - \varphi_g^2} - \frac{1 - \varphi_g^2}{2\sigma_\epsilon^2} \left(y - \frac{\mu_g}{1 - \varphi_g} \right)^2 + \frac{1}{2\sigma_\epsilon^2} \left(y - \mu_t \right)^2 \right) f_t(y) dy$$

$$= -\left(\log \sqrt{1 - \varphi_g^2} - \frac{1 - \varphi_g^2}{2\sigma_\epsilon^2} \left(m_{+,t} - \frac{\mu_g}{1 - \varphi_g} \right)^2 + \frac{(m_{+,t} - \mu_t)^2}{2\sigma_\epsilon^2} \right) \int_{\mu_t}^{+\infty} f_t(y) dy$$

$$- \frac{1}{\sigma_\epsilon^2} \left(m_{+,t} - \mu_t - \left(1 - \varphi_g^2 \right) \left(m_{+,t} - \frac{\mu_g}{1 - \varphi_g} \right) \right) \int_{\mu_t}^{+\infty} \left(y - m_{+,t} \right) f_t(y) dy$$

$$+ \frac{\varphi_g^2}{2\sigma_\epsilon^2} \int_{\mu_t}^{+\infty} \left(y - m_{+,t} \right)^2 f_t(y) dy$$

$$= -\frac{1}{2} \left(\log \sqrt{1 - \varphi_g^2} - \frac{1 - \varphi_g^2}{2\sigma_\epsilon^2} \left(m_{+,t} - \frac{\mu_g}{1 - \varphi_g} \right)^2 + \frac{1}{\pi} \right) + \frac{\varphi_g^2}{4} \left(1 - \frac{2}{\pi} \right),$$

where $\mu_t = \mu_g + \varphi_g y_t$ is the conditional median of the distribution, $m_{+,t} \equiv \mathbb{E}_t \left[y_{t+1} \mid y_{t+1} \geq \mu_t \right] = \mu_t + \sigma_\epsilon \sqrt{\frac{2}{\pi}}$ is the mean of GDP growth conditional on being above the conditional median, and $m_{-,t} \equiv \mathbb{E}_t \left[y_{t+1} \mid y_{t+1} < \mu_t \right] = \mu_t - \sigma_\epsilon \sqrt{\frac{2}{\pi}}$ is the mean of GDP growth conditional on being below the conditional median.

Thus, when both the unconditional and conditional distributions are Gaussian, downside (upside) entropy is expressed in terms of the square difference between the mean conditional on the realization of GDP growth falling below (above) the mean and the unconditional mean and the variance conditional on the realization of GDP growth falling below (above) the mean. When the conditional mean μ_t coincides with the unconditional mean $(1 - \varphi_g)^{-1} \mu_g$, so that the only difference between the conditional and the unconditional distribution is in the variance, downside entropy equals the upside entropy, and risks to the upside and downside are balanced. When the conditional mean is higher than the unconditional mean, downside entropy is higher than upside entropy, and downside risk exceeds upside risk. This property of upside and downside entropy is illustrated in Figure A.5c.

Example 2. Assume now that GDP growth evolves as an AR(1) process with a mixture of normal innovations, so that

$$y_{t+1} = \mu_g + \varphi_g y_t + \epsilon_{t+1}, \qquad \epsilon \sim p_+ \mathcal{N}_+ \left(0, \sigma_{\epsilon,+}^2 \right) + p_- \mathcal{N}_- \left(0, \sigma_{\epsilon,-}^2 \right),$$

where p_+ is the probability of a positive innovation, and $\mathcal{N}_+$ ($\mathcal{N}_-$) denotes normal distributions truncated from below (above) at 0. That is, we allow for negative shocks to GDP growth to have a different variance than positive shocks to GDP growth, so that $\sigma_{\epsilon,-} > \sigma_{\epsilon,+}$. In this case, the (one-step-ahead) conditional distribution of GDP growth is

$$f_t(y) = p_+ \sqrt{\frac{2}{\pi \sigma_{\epsilon,+}^2}} \exp \left(-\frac{1}{2\sigma_{\epsilon,+}^2} \left(y - \mu_g - \varphi_g y_t \right)^2 \right) \mathbf{1}_{y \geq \mu_g + \varphi_g y_t}$$

$$+ (1 - p_+) \sqrt{\frac{2}{\pi \sigma_{\epsilon,-}^2}} \exp \left(-\frac{1}{2\sigma_{\epsilon,-}^2} \left(y - \mu_g - \varphi_g y_t \right)^2 \right) \mathbf{1}_{y < \mu_g + \varphi_g y_t}.$$

Assume further that the unconditional distribution is

$$g\left(y\right) = \sqrt{1 - \varphi_g^2}\left(2\pi\sigma_\epsilon^2\right)^{-\frac{1}{2}}\exp\left(-\frac{1 - \varphi_g^2}{2\sigma_\epsilon^2}\left(y - \frac{\mu_g}{1 - \varphi_g}\right)^2\right).$$

For simplicity, we will focus on the case $p_+ = 1/2$, so that there is an equal probability of positive and negative innovations to GDP growth. In this case, the conditional median is the same as in the conditionally Gaussian case, so that $\mu_t = \mu_g + \varphi_g y_t$. The downside and upside entropy are then given by

$$\mathcal{L}_t^D\left(f_t; g\right) = -\int_{-\infty}^{\mu_t}\left(\log\sqrt{1 - \varphi_g^2} + \log\frac{\sigma_{\epsilon,-}}{\sigma_\epsilon} - \frac{1 - \varphi_g^2}{2\sigma_\epsilon^2}\left(y - \frac{\mu_g}{1 - \varphi_g}\right)^2 + \frac{1}{2\sigma_{\epsilon,-}^2}\left(y - \mu_t\right)^2\right)f_t(y)dy$$

$$= -\left(\log\sqrt{1 - \varphi_g^2} + \log\frac{\sigma_{\epsilon,-}}{\sigma_\epsilon} - \frac{1 - \varphi_g^2}{2\sigma_\epsilon^2}\left(m_{-,t} - \frac{\mu_g}{1 - \varphi_g}\right)^2 + \frac{\left(m_{-,t} - \mu_t\right)^2}{2\sigma_{\epsilon,-}^2}\right)\int_{-\infty}^{\mu_t}f_t(y)dy$$

$$- \left(\frac{m_{-,t} - \mu_t}{\sigma_{\epsilon,-}^2} - \frac{\left(1 - \varphi_g^2\right)}{\sigma_\epsilon^2}\left(m_{-,t} - \frac{\mu_g}{1 - \varphi_g}\right)\right)\int_{-\infty}^{\mu_t}\left(y - m_{-,t}\right)f_t(y)dy$$

$$- \left(\frac{1 - \varphi_g^2}{2\sigma_\epsilon^2} - \frac{1}{2\sigma_{\epsilon,-}^2}\right)\int_{-\infty}^{\mu_t}\left(y - m_{-,t}\right)^2 f_t\left(y\right)dy$$

$$= -\frac{1}{2}\left(\log\sqrt{1 - \varphi_g^2} + \log\frac{\sigma_{\epsilon,-}}{\sigma_\epsilon} - \frac{1 - \varphi_g^2}{2\sigma_\epsilon^2}\left(m_{-,t} - \frac{\mu_g}{1 - \varphi_g}\right)^2 + \frac{1}{\pi}\right)$$

$$+ \frac{1}{4}\left(1 - \left(1 - \varphi_g^2\right)\frac{\sigma_{\epsilon,-}^2}{\sigma_\epsilon^2}\right)\left(1 - \frac{2}{\pi}\right);$$

$$\mathcal{L}_t^U\left(f_t; g\right) = -\int_{\mu_t}^{\infty}\left(\log\sqrt{1 - \varphi_g^2} + \log\frac{\sigma_{\epsilon,+}}{\sigma_\epsilon} - \frac{1 - \varphi_g^2}{2\sigma_\epsilon^2}\left(y - \frac{\mu_g}{1 - \varphi_g}\right)^2 + \frac{1}{2\sigma_{\epsilon,+}^2}\left(y - \mu_t\right)^2\right)f_t(y)dy$$

$$= -\left(\log\sqrt{1 - \varphi_g^2} + \log\frac{\sigma_{\epsilon,+}}{\sigma_\epsilon} - \frac{1 - \varphi_g^2}{2\sigma_\epsilon^2}\left(m_{+,t} - \frac{\mu_g}{1 - \varphi_g}\right)^2 + \frac{\left(m_{+,t} - \mu_t\right)^2}{2\sigma_{\epsilon,+}^2}\right)\int_{\mu_t}^{+\infty}f_t(y)dy$$

$$- \left(\frac{m_{+,t} - \mu_t}{\sigma_{\epsilon,+}^2} - \frac{\left(1 - \varphi_g^2\right)}{\sigma_\epsilon^2}\left(m_{+,t} - \frac{\mu_g}{1 - \varphi_g}\right)\right)\int_{\mu_t}^{+\infty}\left(y - m_{+,t}\right)f_t(y)dy$$

$$- \left(\frac{1 - \varphi_g^2}{2\sigma_\epsilon^2} - \frac{1}{2\sigma_{\epsilon,+}^2}\right)\int_{\mu_t}^{+\infty}\left(y - m_{+,t}\right)^2 f_t\left(y\right)dy$$

$$= -\frac{1}{2}\left(\log\sqrt{1 - \varphi_g^2} + \log\frac{\sigma_{\epsilon,+}}{\sigma_\epsilon} - \frac{1 - \varphi_g^2}{2\sigma_\epsilon^2}\left(m_{+,t} - \frac{\mu_g}{1 - \varphi_g}\right)^2 + \frac{1}{\pi}\right)$$

$$+ \frac{1}{4}\left(1 - \left(1 - \varphi_g^2\right)\frac{\sigma_{\epsilon,+}^2}{\sigma_\epsilon^2}\right)\left(1 - \frac{2}{\pi}\right),$$

where $m_{+,t} \equiv \mathbb{E}_t\left[y_{t+1}\middle| y_{t+1} \geq \mu_t\right] = \mu_t + \sigma_{\epsilon,+}\sqrt{\frac{2}{\pi}}$ is the mean of GDP growth conditional on

3

being above the conditional median, and $m_{-,t} \equiv \mathbb{E}_t\left[y_{t+1}|y_{t+1} < \mu_t\right] = \mu_t - \sigma_{\epsilon,-}\sqrt{\frac{2}{\pi}}$ is the mean of GDP growth conditional on being below the conditional median.

Thus, in this case as well, downside (upside) entropy is expressed in terms of the square difference between the mean conditional on the realization of GDP growth falling below (above) the median and the unconditional mean and the variance conditional on the realization of GDP growth falling below (above) the median. Unlike the fully Gaussian case described below, however, the conditional variance used is different for the upside and downside entropy. Hence, when the conditional median μ_t coincides with the unconditional median $(1-\varphi_g)^{-1}\mu_g$, downside entropy is lower than the upside entropy. Instead, to make the risks to the upside and downside balance, the conditional median of GDP growth has to be below the unconditional median, as can be seen in the numerical example in Figure A.5d.

A.2   GDP Vulnerability and Other Financial Indicators

The results in the main body of the paper 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 used as 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 (2017) demonstrate 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.

Finally, 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 A.6 shows the quantile coefficients for equity market option-implied volatility, the BAA-AAA credit spread, and the 10-year/3-month term spread.[15] 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.

Figure A.7 reports the coefficients obtained by replacing the NFCI with its risk, credit and leverage subindexes. Results broadly confirm the stronger relationship of the predictive distribution of real GDP growth with financial conditions at the lower quantiles.

---

[15]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).

Appendix References

ADRIAN, T., R. CRUMP, AND E. VOGT (2015): "Nonlinearity and flight to safety in the risk-return trade-off for stocks and bonds," Staff Report N. 723, Federal Reserve Bank of New York.

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

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.

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

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

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 (2017): "Credit-market sentiment and the business cycle," *The Quarterly Journal of Economics*.

REY, H. (2015): "Dilemma not Trilemma: The Global Financial Cycle and Monetary Policy Independence," Working Paper N. 21162, National Bureau of Economic Research.

Figure A.1. Predictive Distribution of GDP Growth: Location and Scale Parameters over Time. The figure shows the time series evolution of location μ_t and scale σ_t.

(a) Location μ_t: One quarter ahead
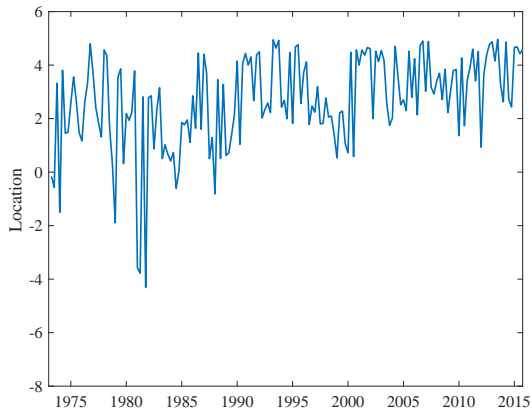
(b) Location μ_t: One year ahead
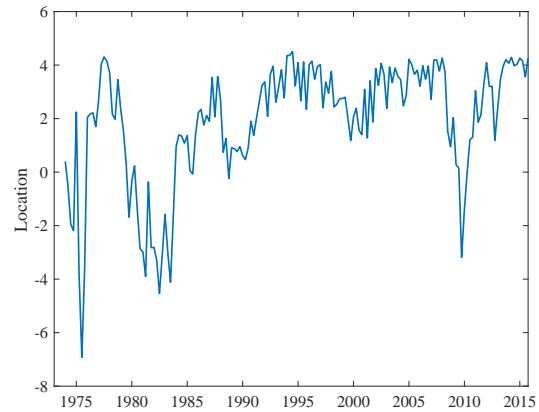
(c) Scale σ_t: One quarter ahead
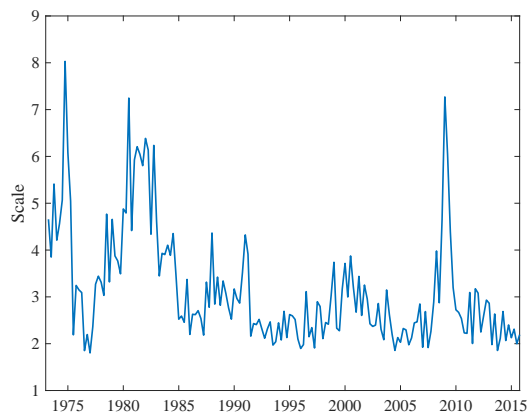
(d) Scale σ_t: One year ahead
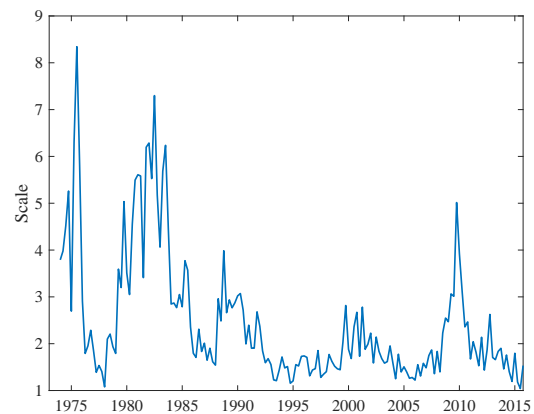
Figure A.2. Predictive Distribution of GDP Growth: Shape and Degrees of Freedom over Time. The figure shows the time series evolution of location shape α_t and degrees of freedom ν_t.

(a) Shape α_t One quarter ahead
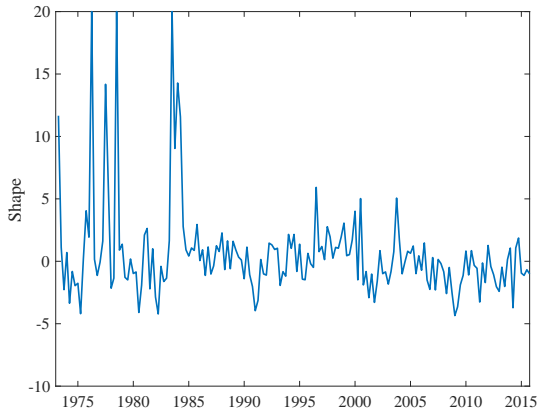
(b) Shape α_t One year ahead
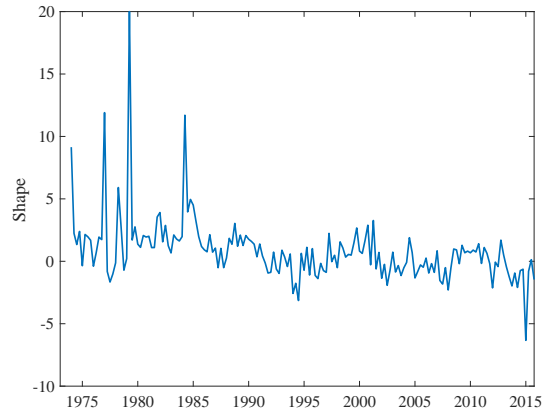
(c) Degrees of Freedom ν_t One quarter ahead
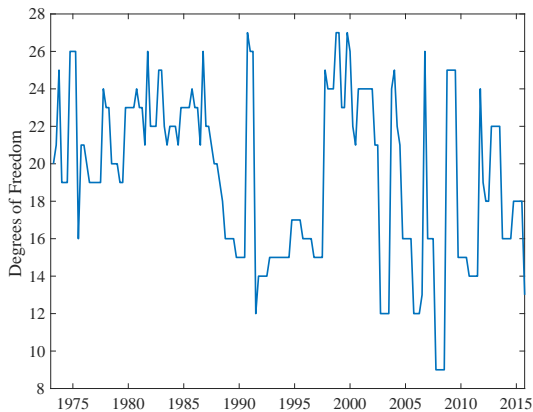
(d) Degrees of Freedom ν_t One year ahead
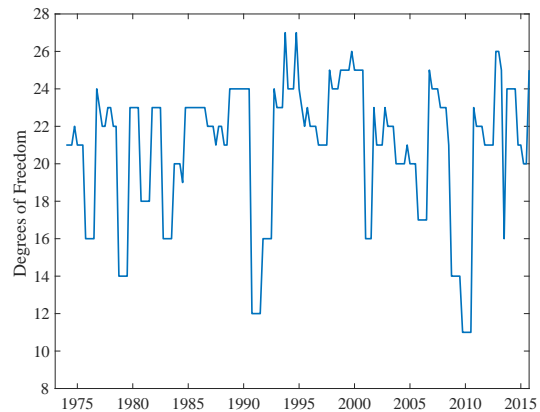
Figure A.3. Predictive Distribution of GDP Growth: Moments over Time. The figure shows the time series evolution of the conditional mean and variance.

(a) Mean: One quarter ahead
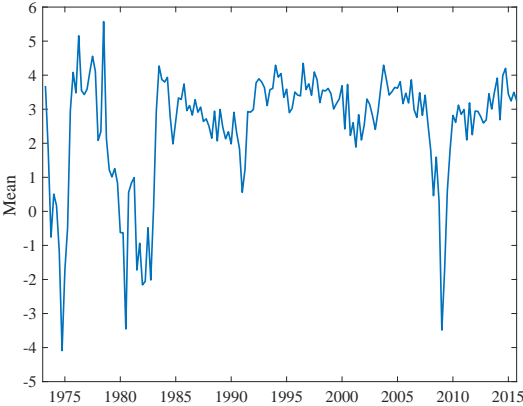
(b) Mean: One year ahead
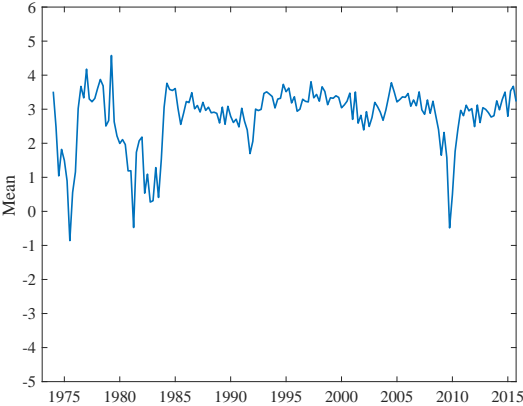
(c) Variance: One quarter ahead
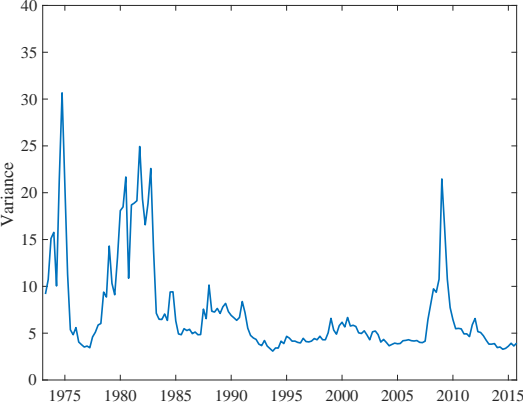
(d) Variance: One year ahead
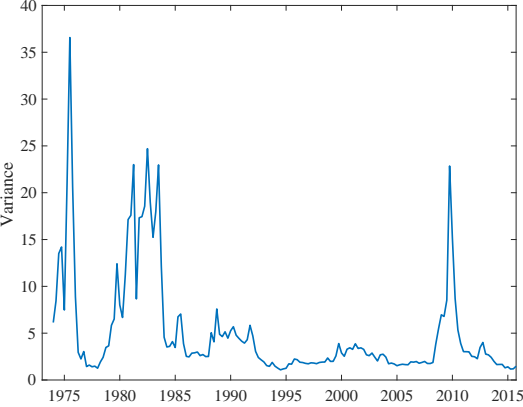
Figure A.4. Predictive Distribution of GDP Growth: Moments over Time. The figure shows the time series evolution of the conditional skewness, and kurtosis.

(a) Skewness: One quarter ahead
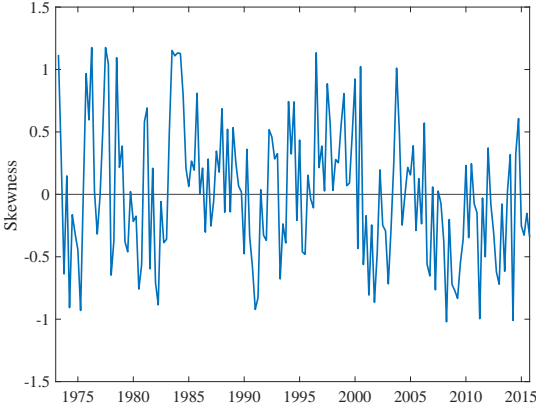
(b) Skewness: One year ahead
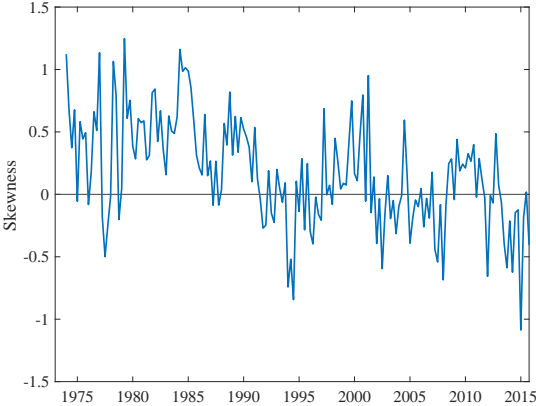
(c) Kurtosis: One quarter ahead
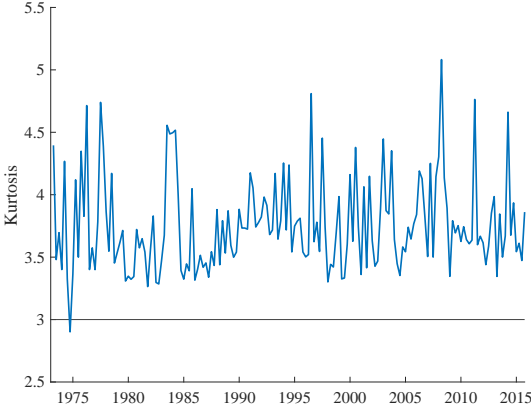
(d) Kurtosis: One year ahead
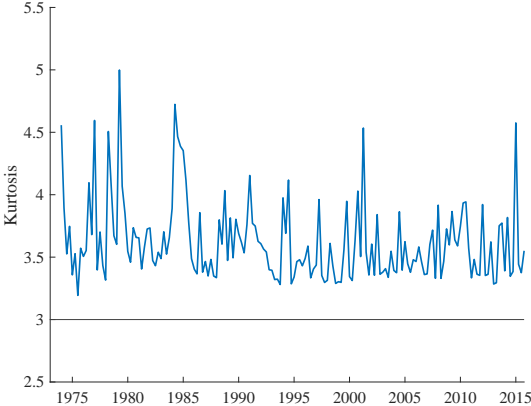
Figure A.5. Upside and Downside Entropy. The left column shows the conditional and unconditional distribution of GDP growth for the autoregressive GDP growth process with normal innovations in Example 1, and the corresponding downside and upside entropy as a function of the current realization of GDP growth. The right column shows the conditional and unconditional distribution of GDP growth for the autoregressive GDP growth process with truncated normal innovations in Example 2, and the corresponding downside and upside entropy as a function of the current realization of GDP growth. Parameters: $\mu_g = 0.2$, $\sigma_\epsilon = 3.5$, $\phi_g = 0.4$, $\sigma_{\epsilon,+} = 3.3$, $\sigma_{\epsilon,-} = 4.9$

(a) Distribution: Gaussian innovations   **(b) Distribution**: Truncated innovations
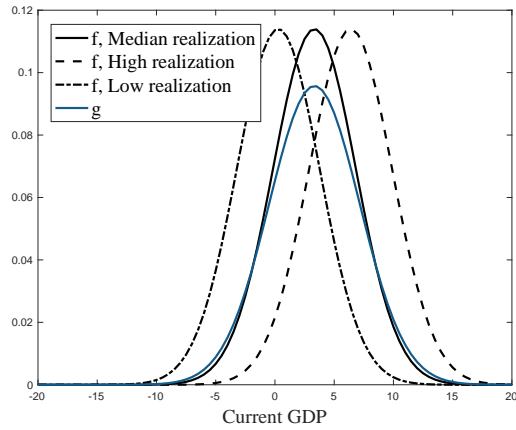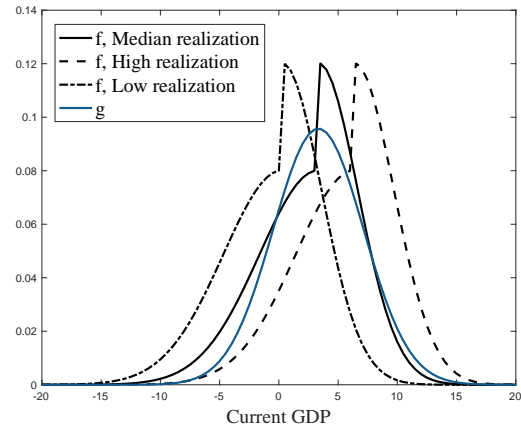
(c) Entropy: Gaussian innovations   **(d) Entropy**: Truncated innovations
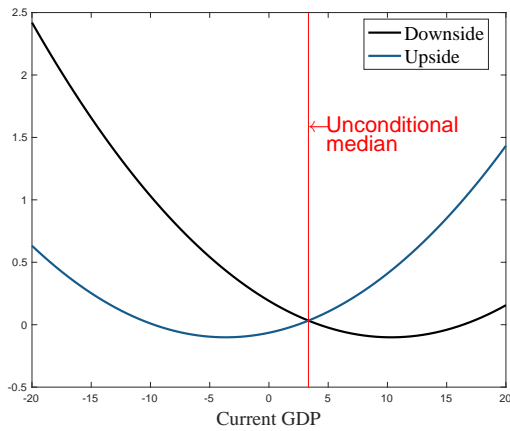
11
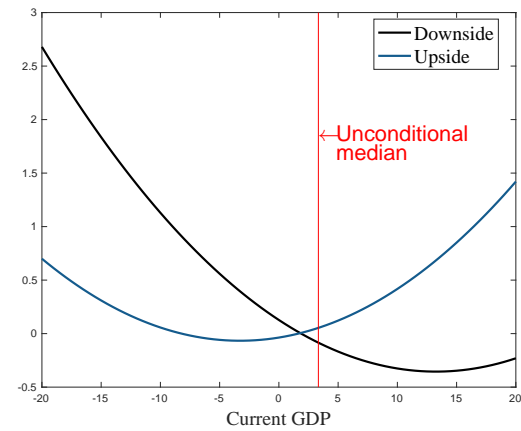
Figure A.6. GDP Growth and Other Predictors. The figure shows the estimated coefficients in quantile regressions of one quarter ahead (left column) and four quarter ahead (right column) real GDP growth on current real GDP growth and VXO; current real GDP growth and the Baa-Aaa spread; current real GDP growth and the term spread.
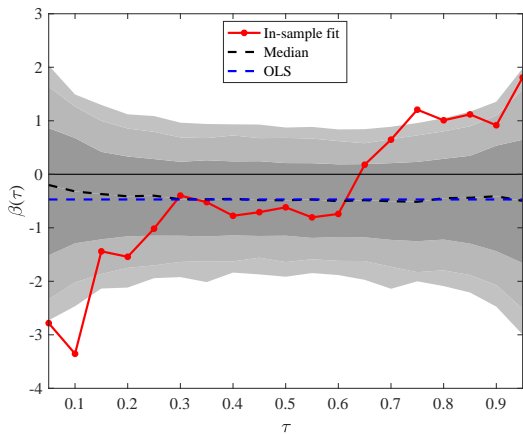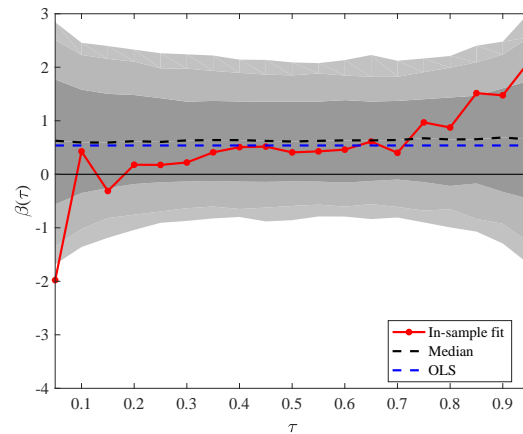
(a) Equity Volatility: One quarter ahead
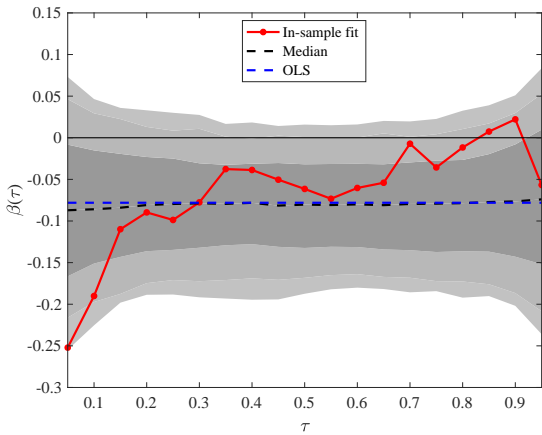
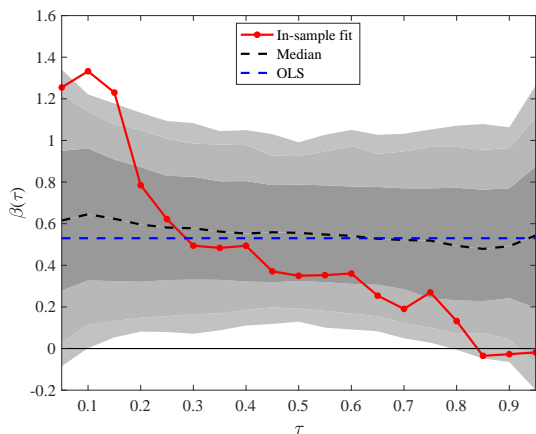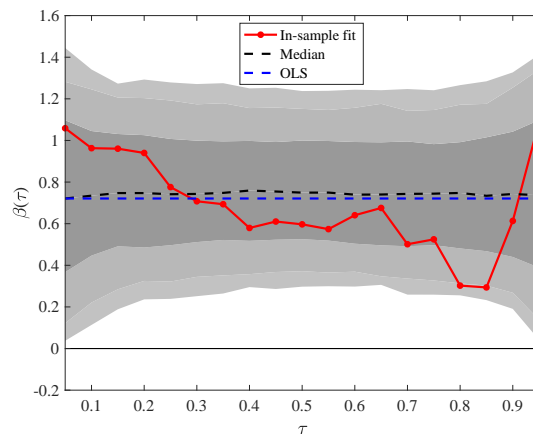**(b) Equity Volatility**: One year ahead
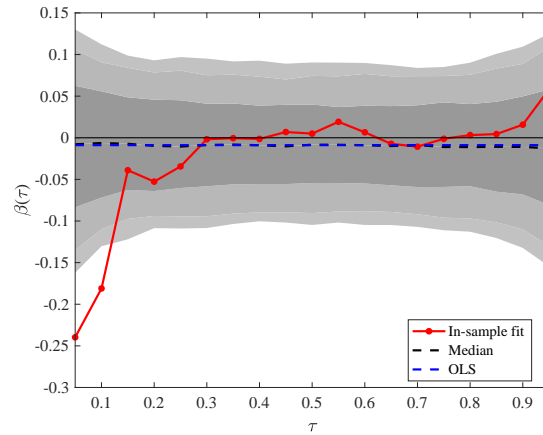
(c) Credit Spread: One quarter ahead

(d) Credit Spread: One year ahead

(e) Term Spread: One quarter ahead

(f) Term Spread: One year ahead

Figure A.7. GDP Growth and Subcomponents of NFCI. 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 current real GDP growth and the risk subindex of the NFCI; current real GDP growth and the credit subindex of the NFCI; current real GDP growth and the leverage subindex of the NFCI.
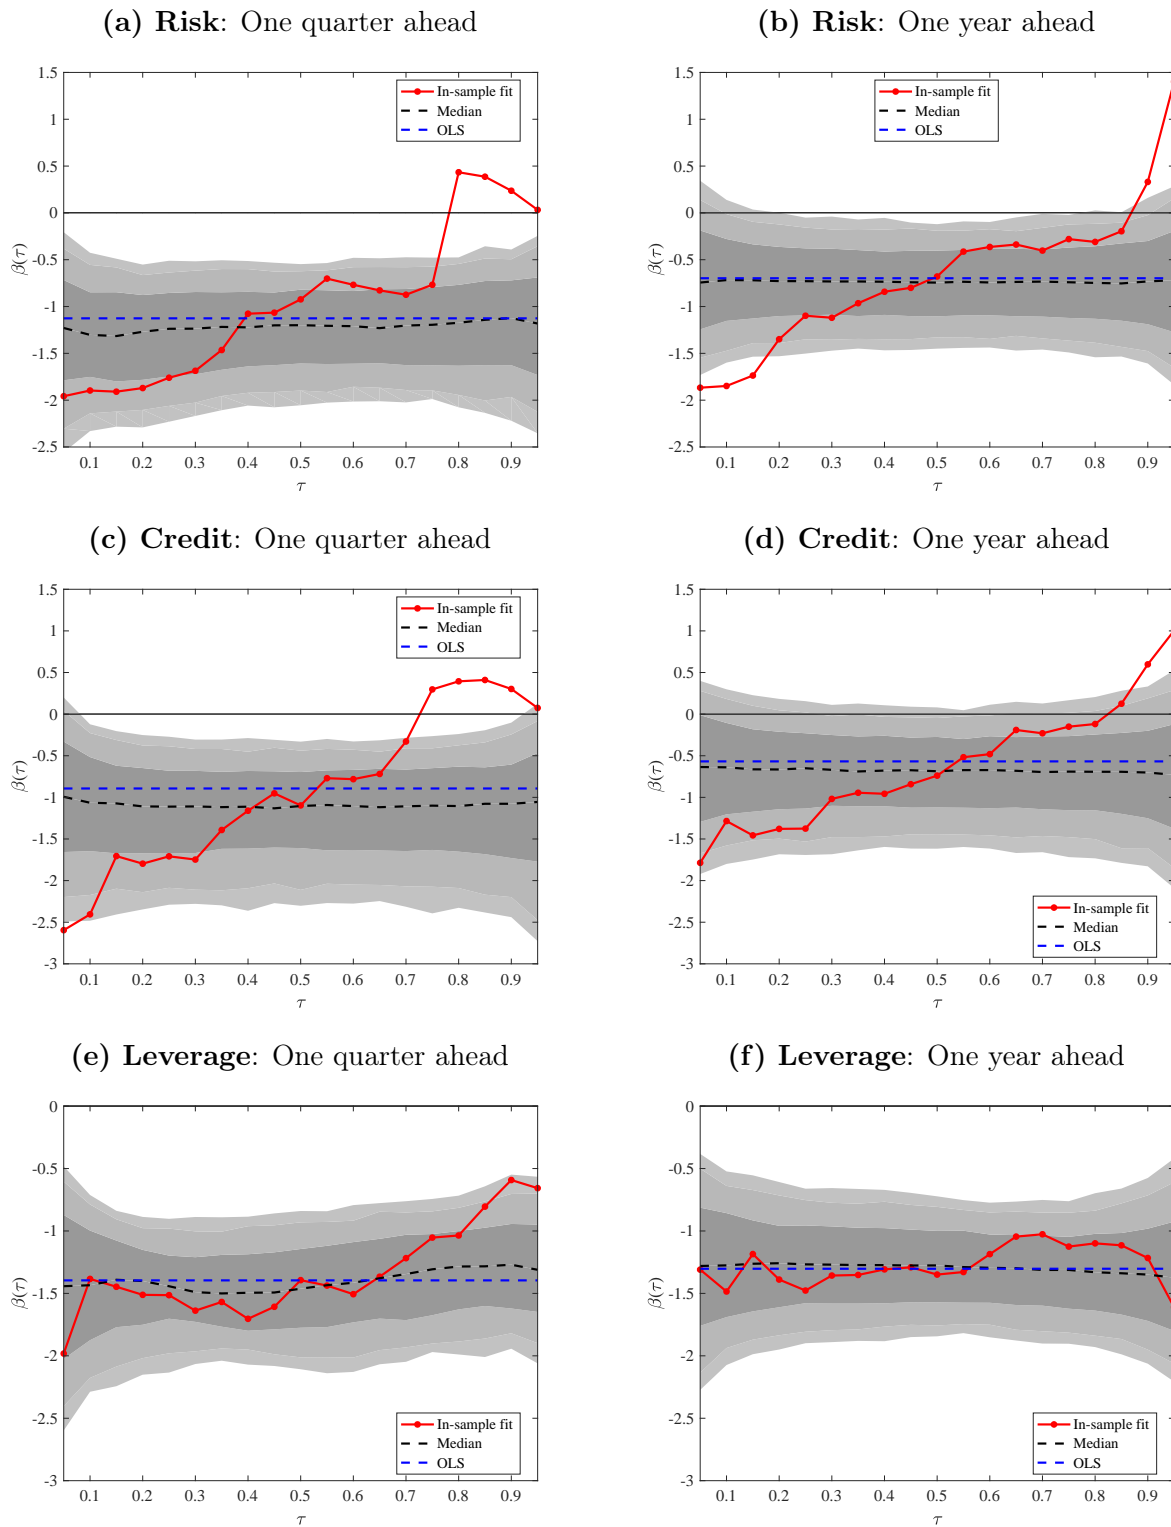

(a) **Risk**: One quarter ahead

(b) **Risk**: One year ahead

(c) **Credit**: One quarter ahead

(d) **Credit**: One year ahead

(e) **Leverage**: One quarter ahead

(f) **Leverage**: One year ahead

Figure A.8. Out-of-Sample Performance of Alternative Approaches. The figure reports the out-of-sample predictive density, the predictive scores and the cumulative distribution of the probability integral transform, generated with a conditionally Gaussian model (left column) and with a non-parametric model (right column). Figures A.8a and A.8b show the 5, 50, and 95 percent quantiles. Figures A.8c and A.8d compare the out- of-sample predictive scores of the predictive distribution conditional on both NFCI and real GDP growth, and the (semi-parametric) predictive distribution conditional on real GDP growth only. Figures A.8e and A.8f report the empirical cumulative distribution of the probability integral transform (PITs). Critical values are obtained as in Rossi and Sekhposyan (2017).
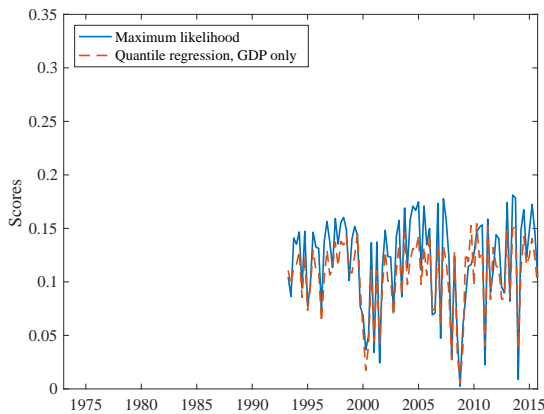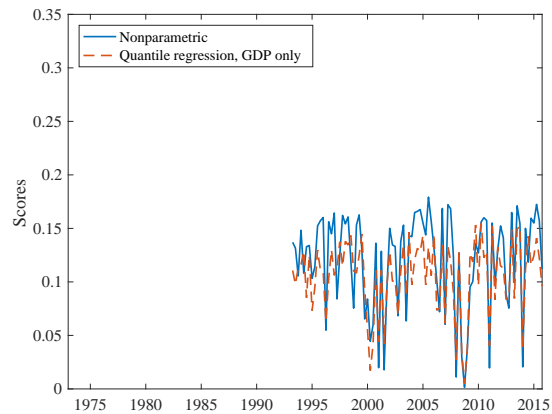
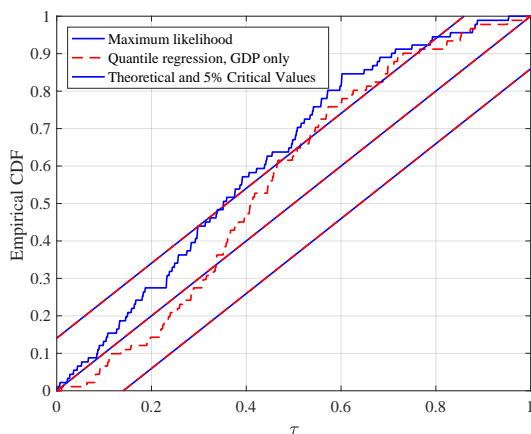**(a) Quantiles**: MLE
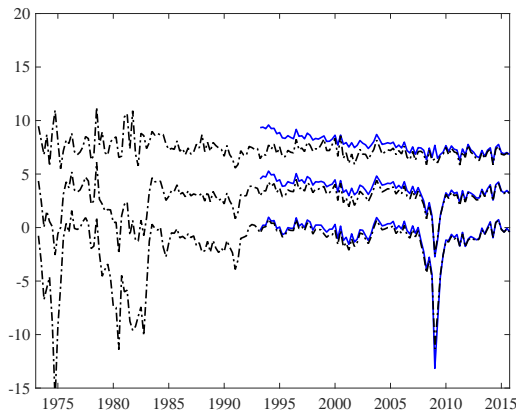
(b) Quantiles: Non-parametric
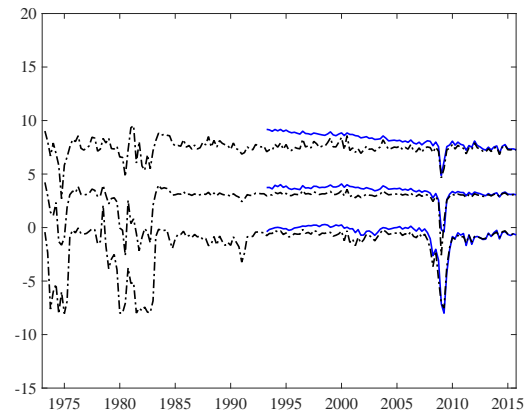
(c) Scores: MLE

(d) Scores: Non-parametric

(e) PITs: MLE

(f) PITs: Non-parametric

14