

When Growth-at-Risk Hits the Fan: Comparing Quantile-Regression Predictive Densities with Committee Fan Charts*

Very Preliminary and Incomplete — Please Do Not Circulate

Simon Lloyd[†]

Ed Manuel[‡]

May 30, 2022

Abstract

Both monetary and macroprudential policymakers use conditional density forecasts of GDP to inform their policy decisions. The former tend to use ‘fan charts’ that convey uncertainty around central projections from linearised macroeconomic models. The latter draw on estimates of ‘growth-at-risk’, typically estimated using quantile-regression techniques, reflecting their focus on the tails of the GDP distribution. Focusing on the UK, we study how the fan charts constructed by the Bank of England’s Monetary Policy Committee compare to predictive densities derived from statistical growth-at-risk models. We find that GDP-at-risk models provide improved estimates of the left-tail of the GDP-growth distribution, in particular at medium-term horizons compared to the GDP fan charts. However, GDP-at-risk models generally perform worse than the fan charts at the centre of the distribution. Combining forecast densities in a parsimonious manner provides the best forecasts, with limited losses compared to optimal density combination methods. Doing so offers central banks the opportunity to improve GDP density forecasts, and unify the narrative provided by monetary and macroprudential policymakers about the economic outlook and uncertainty around it.

JEL Codes: E44, E47, E58, G01.

Key Words: Fan Charts; Forecast Densities; Growth-at-Risk; Quantile Regression.

*We thank Manfred Kremer, James Mitchell and Tatevik Sekhposyan for helpful comments. We are also grateful to Marco Garofolo and Alex Rattan for sharing helpful insights on the Bank of England’s fan charts. The views expressed here are those of the authors, and do not necessarily reflect those of the Bank of England.

[†]Bank of England. Email Address: simon.lloyd@bankofengland.co.uk.

[‡]Bank of England. Email Address: edward.manuel@bankofengland.co.uk.

1 Introduction

As part of a drive to foster greater financial stability since the global financial crisis, a number of central banks have been tasked with mandates to establish macroprudential policy frameworks.¹ For many of these central banks, financial-stability policymaking overlaps with monetary policy setting. Some of these commonalities are practical. For example, at the Bank of England, four of the nine members of its Monetary Policy Committee (MPC) also sit on the thirteen-member Financial Policy Committee (FPC).² But there are other synergies in communications and messaging. For instance, the Bank of England’s FPC regularly cites the MPC’s central projections for the UK economic outlook in its public outputs, and uses it as one input to inform its assessment of the overall risk environment when setting macroprudential policy.

In spite of these links, there are differences in the underlying focus of these policy frameworks. Within inflation-targeting regimes, monetary policy is typically focused on central forecasts for economic activity, unemployment and inflation. On the other hand, macroprudential policymakers’ focus tends to be on mitigating the economic impacts of ‘tail events’, i.e. ‘1-in- x ’ bad outcomes like financial crises, for GDP and welfare.

Despite the fact monetary policymakers now regularly provide ‘fan charts’ (i.e. conditional distributions for variables of interest across each horizon of the forecast) around their central forecasts to communicate uncertainty around their projections, these are seldom used by financial stability bodies to inform their views of potential tail outcomes. Indeed, many of the forecasting models used by central banks (e.g. COMPASS at the Bank of England and RAMSES at the Riksbank) are solved using linear approximation methods—often necessarily so to foster estimation—so are not well-suited to capture non-linearities and multi-modalities in economic projections relevant for financial-stability policymakers. While numerous frameworks allow for some asymmetries in forecast densities,³ these are often grounded in policy committee judgement and an inspection of historical forecast errors, rather than an assessment of how macro-financial vulnerabilities endogenously affect these higher moments.

Instead, with a view to financial-stability policy, a growing body of work has sought to apply quantile regression techniques to estimate the level and drivers of macroeconomic tail risks (see, e.g., [Adrian, Boyarchenko, and Giannone, 2019](#); [Adrian, Grinberg, Liang, Malik,](#)

¹In some jurisdictions, institutions other than the central bank have been handed financial stability mandates (e.g. Germany, Sweden). However, in these countries, it is still common for senior central bank staff to be involved in macroprudential policymaking (e.g. Sweden).

²These four members comprise the Governor, as well as three Deputy Governors—for Monetary Policy, Financial Stability, and Markets and Banking. They also concurrently sit on the Bank’s Prudential Regulation Committee (PRC), which focuses on microprudential regulation.

³For example, the Bank of England fan charts are constructed using a two-part normal distribution, which admits skewness, but not kurtosis.

and Yu, 2022; Aikman, Bridges, Hacıoglu Hoke, O'Neill, and Raja, 2019; Lloyd, Manuel, and Panchev, 2022). This literature typically focuses on estimating a low quantile of the distribution of (projected) GDP growth—termed ‘growth-at-risk’ or ‘GDP-at-risk’—precisely the tail of the distribution relevant for financial stability and macroprudential policy (see, e.g., Carney, 2020; Suarez, 2021). Quantile regression techniques are well-suited to doing this, by accounting for the potentially heterogeneous effects of variables across the entire distribution of variables of interest (e.g. GDP and inflation). As a consequence, the headline estimated drivers of growth-at-risk (e.g. credit growth) can differ from those captured in central-forecasting models.

These quantile-regression techniques can go beyond a focus on the tails of distributions too. By estimating the relationships between variables across the distribution, quantile regression methods can be used to construct forecasts of the entire conditional distribution of future GDP growth across horizons—in effect, providing an alternative to monetary policymakers’ fan charts. Moreover, relative to an approach that constructs fan charts using linearised forecasting models and historical forecast errors, the predictive densities from quantile regressions could yield more flexible—and potentially more accurate—outputs in light of the fact they can estimate moves over time in higher-order GDP moments conditional on the macro-financial environment.

This raises a natural set of questions, which we seek to answer in this paper. How do fan charts from monetary policymakers—constructed using a combination of judgement and historical forecast errors—compare to the predictive densities derived from statistical growth-at-risk models that rely on quantile regression techniques and are employed by financial stability policymakers? Can central bank fan charts be improved by using insights from quantile regression techniques? Do combinations of fan charts and growth-at-risk methods yield ‘better’ estimates? Can combinations of these estimates offer central banks the opportunity to unify the narrative provided by monetary and macroprudential policymakers?

These questions are of additional importance in light of work by Adrian and Duarte (2018) who model the effects of monetary policy on financial vulnerabilities in the tails of the GDP-growth distribution. The occasionally binding value-at-risk constraints that generate these effects can give rise to multi-modalities in the GDP-growth distribution (Adrian, Boyarchenko, and Giannone, 2021), which can be captured by quantile-regression techniques.

To answer these questions, we focus on GDP fan charts produced by the Bank of England’s MPR in their quarterly *Monetary Policy Reports* (MPR)—formerly known as the *Inflation Report* (IR)—and compare them to the predictive densities from GDP-at-risk models.⁴ We show that GDP-at-risk models can provide better estimates specifically of the left-tail of the GDP growth

⁴Formally, our growth-at-risk setting builds on models developed within the Bank of England (Aikman et al., 2019; Lloyd et al., 2022).

distribution than the Bank of England’s GDP fan charts. However, GDP-at-risk models generally perform worse than the fan charts at the centre of the distribution. Building on these findings, we propose a practical solution for central banks to improve the calibration of their fan charts by incorporating quantile regression methodologies into their construction, unifying the narrative provided by monetary and macroprudential policymakers.

In comparing GDP-at-risk model estimates to the Bank of England’s GDP fan charts, we emphasise three main findings. First, we present evidence around the optimal modelling choices for standard GDP-at-risk models. We run a horse-race between different quantile regression models for the conditional distribution of UK GDP growth. We find that exploiting panel data, and incorporating information on vulnerabilities abroad improves out-of-sample performance.

Second, we find GDP-at-risk models can provide improved estimates of the left-tail of the GDP growth distribution in particular at medium term horizons compared to historical Bank of England GDP fan charts. However, the GDP-at-risk models we consider generally perform worse than the Bank of England GDP fan charts in predicting the centre of the distribution.

Third, combining forecasts from Bank of England fan charts and GDP-at-risk models provides the best forecasts. While a range of combination methods may be employed, we propose a simple method which can improve fan chart calibration and be readily incorporated by central banks in their construction of their fan charts.

Related Literature Our paper is related to two main strands of literature.

First, our papers relates to a range of work that assesses the ability of quantile regression techniques to forecast the conditional GDP growth distribution (see, e.g., [Giglio et al., 2016](#); [Aikman et al., 2018, 2019](#); [Adrian et al., 2019, 2022](#); [Plagborg-Møller et al., 2020](#); [Brownlees and Souza, 2021](#); [Adams et al., 2021](#); [Lloyd et al., 2022](#)). We extend this literature by explicitly comparing the predictive densities from growth-at-risk models to Bank of England GDP fan charts—an exercise that is particularly informative for the question of whether growth-at-risk models can be of practical use for central banks for forecasting purposes. We also contribute to this literature by assessing the out-of-sample accuracy of a range of quantile regression models, providing novel evidence e.g. on the importance of using panel data to improve out-of-sample forecasts of the conditional GDP growth distribution in a quantile-regression setting.

Second, we contribute to literature evaluating central bank density forecasts, particularly those from the Bank of England (see, e.g., [Clements, 2004](#); [Wallis, 2004](#); [Mitchell and Hall, 2005](#); [Elder, Kapatnios, Taylor, and Yates, 2005](#); [Dowd, 2007](#)). We broadly confirm the findings of these papers, showing that the Bank of England’s GDP fan charts provide compet-

itive density forecasts relative to benchmark empirical models with no significant evidence of mis-specification. We also build on work that assesses whether the width and skew of central banks' fan charts are well-calibrated (see, e.g., [Knüppel and SchulteFrankenfeld, 2012, 2019](#); [Galvão, Garratt, and Mitchell, 2021](#)). Here we find evidence that judgements specifically around higher-order moments for the Bank of England's GDP fan charts can be improved through the use of statistical models, and quantile regression models in particular.

Our key contribution is to bring these two literatures together to ask: how do Bank of England GDP density forecasts compare to those from GDP-at-risk models? In doing so, we directly address a question of key interest for central banks and other international policy institutions, such as the International Monetary Fund. In particular, we present novel evidence around the effects of combining GDP-at-risk model estimates with central bank forecasts, and as part of this propose a practical solution for central banks to improve the calibration of their fan charts by incorporating quantile regression methodologies into their construction.

The remainder of this paper is structured as follows. Section 2 provides detail on the Bank of England's Fan Charts, published since 1996. Section 3 describes our quantile-regression GDP-at-risk framework and documents results from a horse-race between various model specifications. Section 4 compares and evaluates the predictive densities from our selected GDP-at-risk model with Bank of England GDP fan charts, and considers combination methods. Section 5 concludes, listing further work and outstanding questions for this project.

2 Bank of England Fan Charts

The Bank of England first released fan charts for inflation in 1996, and the first fan chart for UK GDP was presented in the November 1997 IR. Since then, these fan charts have been published on a near-quarterly basis in the Bank's MPR—formerly the IR. The methodology underlying the construction of these fan charts has been broadly unchanged since the fan chart's inception, and is described in [Britton, Fisher, and Whitley \(1998\)](#).

2.1 Choice of Distribution

When constructing fan charts for GDP, the Bank of England assume that the conditional distribution of GDP growth follows a two-piece normal distribution. The split normal distribution arises from merging two opposite halves of normal distributions in their common mode. The resultant distribution can display either positive or negative skewness, depending on the relative variances of the two normal distributions used to construct it.

The two-piece normal distribution is governed by three parameters: the mode μ , a measure of uncertainty σ , and a measure of the balance of risks γ . The probability density function for the two-piece normal is defined as:

$$f(x|\mu, \sigma, \gamma) = A \exp \left[-\frac{1}{2\pi\sigma^2} \left((x - \mu)^2 + \gamma \left(\frac{x - \mu}{|x - \mu|} \right) (x - \mu)^2 \right) \right] \quad (1)$$

where:

$$A = \frac{2}{\left[\frac{1}{\sqrt{1-\gamma}} + \frac{1}{\sqrt{1+\gamma}} \right]} \frac{1}{\sqrt{2\pi\sigma^2}} \quad (2)$$

This choice of distribution simplifies the calibration of the fan chart by reducing it to the estimation of just three parameters, each with an intuitive interpretation. Part of the motivation for choosing the two-piece normal distribution, rather than a bell-shaped distribution, is that it allows the fan chart to exhibit a degree of asymmetry in the form of variable skew (Britton et al., 1998).

2.2 Calibrating the Parameters

The Bank of England's fan charts are constructed by calibrating the three key parameters governing the distribution function in equations (1) and (2). This calibration is informed by a combination of statistical tools and judgements by the MPC.

Mode The modal path for GDP is informed by the Bank of England's suite of forecasting models which have been periodically updated to align with latest forecasting techniques (see Burgess, Fernandez-Corugedo, Groth, Harrison, Monti, Theodoridis, and Waldron, 2013, for a description of the latest approach, alongside a brief summary of previous models). Under the current approach, the central forecast is constructed using COMPASS, a small-open economy New Keynesian DSGE model estimated on UK data using Bayesian methods, alongside a range of judgements in part informed by other forecasting models.

Uncertainty and Skewness The uncertainty parameter σ —which governs the width of the fan chart—is partly informed by historical forecast errors. It is also informed by forward-looking judgements as to whether uncertainty over the forecast horizon is likely to be greater or less than past experience. Similar judgements regarding whether risks are likely to be skewed to the upside or downside underpin the calibration of the γ parameter. These judgements may be informed by, amongst other things, the calibration of a range of alternative scenarios under different assumptions to the central forecast (e.g. different assumptions for the path of government policy). Constructing the fan chart by simulating a range of scenarios and then

attaching probability weights to each approximates the task of simulating all possible forecast variants.

Figures 1 and 2 demonstrate how the parameters underlying the Bank of England’s GDP fan charts have evolved over time at the 1- and 8-quarter-ahead forecast horizons, respectively. At medium-term horizons in particular there are clear moves in the width of the GDP fan chart around key macroeconomic events, with the uncertainty parameter at $h = 8$ highest following the collapse of Lehman Brother’s in 2008, during the height of the euro area sovereign debt crisis in 2012 and following the Brexit referendum result in the second half of 2016.

3 Quantile-Regression Model for GDP-Growth Distribution

Having described the approach to constructing the Bank of England’s fan charts, we now present our statistical methodology for attaining alternative estimates of the predictive density of UK GDP growth. The model we propose is estimated using quantile-regression techniques, and draws on the GDP-at-risk literature. GDP-at-risk is defined as a low quantile of the conditional distribution of GDP growth—for our purposes we equate it with the 5th percentile of the distribution. While other literature has typically focused specifically on estimating GDP-at-risk (i.e. a point in the left tail of the GDP distribution), we use the quantile regression framework below to study the entire predictive distribution of GDP growth in order to compare to the distributions from the Bank of England fan charts.

3.1 Framework

We use quantile regressions to estimate the entire conditional distribution of GDP growth, following a wide literature on GDP-at-risk (see, e.g., [Aikman et al., 2019](#); [Adrian et al., 2019, 2022](#); [Lloyd et al., 2022](#)).

Denoting time by $t = 1, \dots, T$ and the countries for whom we estimate the conditional distribution of GDP with $i = 1, \dots, N$, we estimate the following local-projection model ([Jordà, 2005](#)) for the conditional quantile function Q of h -period-ahead 4-quarter GDP growth $\Delta^h y_{i,t+h}$:

$$Q_{\Delta^h y_{i,t+h}}(\tau | \mathbf{X}_{i,t}) = \alpha_i^h(\tau) + \beta^h(\tau) \mathbf{X}_{i,t} \quad (3)$$

where Q computes quantiles τ of the distribution of $\Delta^h y_{i,t+h}$ given a set of covariates $\mathbf{X}_{i,t}$ and a potentially country- and quantile-specific constant $\alpha_i^h(\tau)$. To align the quantile-regression output with the Bank of England’s fan charts, we set our dependent variable as h -quarter-ahead year-on-year GDP growth, i.e. $\Delta^h y_{i,t+h} \equiv (y_{i,t+h}/y_{i,t+h-4} - 1) \times 100$. Equation (3)

Figure 1: Bank of England GDP Fan Chart Parameters over time: 1-quarter horizon



Figure 2: Bank of England GDP Fan Chart Parameters over time: 8-quarter horizon



Note: Estimates of the mode, uncertainty and skewness parameters from the Bank of England's GDP fan charts at $h = 1$ and $h = 8$ over time.

captures our broad framework and encompasses estimation using both panel data (i.e. $N > 1$) and single-country data (i.e. $N = 1$), as well as a range of choices for covariates $\mathbf{X}_{i,t}$ and assumptions for the constant term $\alpha_i^h(\tau)$. We discuss in the next sections the data and exact specifications we use to estimate equation (3).

3.2 Data

The dataset we use to estimate equation (3) builds on that in Aikman et al. (2019) and Lloyd et al. (2022). It contains data on a range of macro-financial variables for 11 advanced economies: Australia, , Canada, Denmark, Finland, France, Germany, Italy, Norway, Spain, Sweden, Switzerland, the UK and the US. The dataset spans the period 1981Q1 to 2018Q4.⁵

For our set of covariates $\mathbf{X}_{i,t}$ we rely on evidence from the literature regarding the key predictors of the conditional GDP growth distribution. Given the Bank of England fan charts cover the forecast horizon $h = 1$ to $h = 12$, we include variables that have predictive power at both near-term and medium-term horizons. In particular, we rely on evidence from Aikman et al. (2019), Adrian et al. (2022) and Lloyd et al. (2022) who demonstrate the predictive power of a range of variables out to $h = 12$.⁶ Following these studies, we focus on three broad categories of covariates $\mathbf{X}_{i,t}$: (i) price-based financial indicators (e.g. equity price volatility, corporate bond spreads etc.); (ii) quantity-based leverage indicators (e.g. changes in the level of credit-to-GDP); and (iii) controls for prevailing macroeconomic conditions (e.g. GDP growth, inflation, interest rates). Following Lloyd et al. (2022) specifically, we consider both domestic and foreign-weighted indicators within our three broad categories of covariates, e.g. including a measure of global leverage alongside domestic leverage. One of the key findings from these studies specifically is the importance of separating out financial *price* variables and credit *quantity* variables due to their different cyclical behaviour - rather than conflating them into a single index of ‘financial conditions’. In particular, price-based measures have been found to have a significant effect on the conditional distribution of GDP growth (and on the left-tail in particular) in the near term, while quantity-based measures have significant effects in the medium term.

Given our focus on comparing GDP-at-Risk model estimates to real-time forecasts from the Bank of England’s MPR, we include lagged variables to reflect data lags for e.g. GDP and credit variables. For high-frequency financial market variables (e.g. equity volatility) we include contemporaneous variables.

⁵See Lloyd et al. (2022, Appendix A) for a full description of data sources.

⁶Our focus differs importantly from, e.g., Adrian et al. (2019), Plagborg-Møller et al. (2020), and Brownlees and Souza (2021), who focus only on predicting the conditional growth distribution in the *near* term, i.e. out to $h = 4$.

3.3 Model Specifications

Given a wide range of potential specifications to estimate equation (3), we first run a horse-race between a range of specifications, before comparing our preferred model to the MPR GDP fan charts in Section 4. We describe the details of each of the specifications we estimate below, before presenting both in- and out-of-sample results.

Baseline In our baseline specification, we estimate equation (3) using data on all 11 advanced economies in our dataset spanning 1981Q1 to 2018Q4. We include a fairly minimal number of indicators, using just three indicators in the variable set $\mathbf{X}_{i,t}$: the one-quarter realised volatility of equity prices; the (one-quarter lagged) three-year percentage point change in the aggregate private non-financial credit-to-GDP ratio; and the one-quarter lagged growth of real GDP. As in Aikman et al. (2019) and Lloyd et al. (2022), we use equity price volatility as a proxy for financial market conditions, while the change in credit-to-GDP captures leverage in the real economy. We use equity price volatility in our baseline given its availability across countries and over time relative to other price-based financial variables (e.g. corporate bond spreads). And we favour the three-year change in credit-to-GDP to capture *persistent* changes in credit, which are thought to pose risks to financial stability and are leading indicators of macroeconomic crises (Schularick and Taylor, 2012). We estimate country- and quantile-specific fixed effects $\alpha_i^h(\tau)$ following Kato, Galvao, and Montes-Rojas (2012) and Galvao, Gu, and Volgushev (2020). We use the one-quarter lag for GDP and credit variables to reflect data lags.

Alternate Financial Conditions Index Adrian et al. (2022) use a Financial Conditions Index (FCI) as proxy for financial conditions. The FCI, constructed per the method of Koop and Korobilis (2014), is a summary measure that extracts common variation across a range of asset prices. This captures wider variables and so may improve forecasting performance. So we also estimate specification where we replace equity price volatility with FCI. Relative to the baseline, we estimate slightly smaller panel given limited data availability for the FCI (dropping 3 countries relative to the baseline).

Include Global Variables Following Lloyd et al. (2022), we also estimate a model with global vulnerabilities. To do this, we include foreign-weighted measures of the VIX, (lagged) credit-to-GDP, and (lagged) GDP growth. We construct foreign-weighted variables using data on bilateral trade linkages. Using data from IMF Direction of Trade Statistics, we define the weights $\omega_{i,j,t}$ as the fraction of country i 's exports to country j at time t . This scheme will place higher weight on countries that country i exports more extensively to, reflecting the fact that a down-

turn in one country j may spill over to another i through reduced demand for country- i exports. Owing to constraints on data availability, we focus on the same set of *foreign* countries used in the *domestic* variable set, i.e. $N = N^* = 11$.

Additional Domestic Variables We also estimate a variant of the model with a wider set of domestic covariates. These include the variables used in our baseline specification, plus the domestic 3-year house price growth, the capital ratio (a measure of overall banking system resilience), the current account, the 1-year change in headline central bank policy rates and 1-year inflation. [Aikman et al. \(2019\)](#) use a similar set of covariates. As before, we only include one-quarter lag versions of each variable to reflect data lags (except for the central bank policy rate where we use contemporaneous variable).

We also consider this specification with and without the three global variables described above.

Alternate Fixed Effects We also experiment with a specification that deviates from the baseline by imposing an alternative country fixed-effects structure, in which the fixed effect is the same across quantiles for a given country, i.e. $\alpha_i^h(\tau)$ for all τ , alongside a quantile-specific intercept ([Canay, 2011](#)). This assumption is more restrictive than the baseline model (which allows for fixed effects varying across countries and quantiles), but may be useful in improving out-of-sample forecast accuracy by imposing additional structure to the quantile regression and effectively reducing dimensionality.

UK-Specific Model The models described so far have a panel dimension, which may improve forecasts with respect to a single-country model by increasing the number of observations, which is likely to be particularly important when forecasting the lower quantiles of the distribution where there are relatively few observations. However, using panel data may also reduce forecast accuracy, particularly if structural features of the UK economy differ importantly from other countries in our panel such that the panel-wide coefficient estimates provide a poor fit for the UK specifically. To test for this, we estimate a specification with same covariates as in our baseline, but using only data for the UK (i.e. moving away from panel model by setting $N = 1$). We focus on this specification as our ‘UK-specific’ model in following sections, although in the Appendix we also show results for some of the other specifications described above estimated with only UK data (e.g. specifications with FCI instead of equity price volatility).

3.4 In-Sample Results

We begin by setting out in-sample coefficient estimates in each specification across horizons, estimated at both the 5th and 50th percentiles. As Table 1 summarises, the headline results are similar across model specifications and in line with the broader GDP-at-risk literature. In particular a sharp tightening in (domestic and global) financial conditions leads to a significant worsening in the growth outlook (and particularly on downside growth risks) in the near-term, while rapid credit-to-GDP growth leads to a worsening in expected GDP and especially the 5th percentile in the near-to-medium term.

3.5 Out-of-Sample Results

In this section, we present out-of-sample results from the model. In effect, we back-test the model by replicating the results that an analyst would have produced in real-time, given access to data available at that time.⁷ This is out-of-sample exercise is crucial for comparing the predictive densities from the model to the Bank of England fan charts in Section 4, given the fan chart are produced in real-time.

To evaluate forecast accuracy, we calculate quantile scores for out-of-sample forecasts from our various model specifications. One advantage of this test is that we do not need to rely on additional assumptions to recover estimates of the entire conditional GDP density and can focus instead on individual quantile estimates directly estimated within the quantile regression. We can also locate exactly where the model performs well, or poorly, in terms of its ability to predict different quantiles of GDP growth.

The average quantile score at quantile τ and forecasting horizon h is defined as:

$$QS^h(\tau) = \frac{1}{N_\nu} \sum_{\nu} \rho_{\tau} \left(y_{t_{\nu}+h} - \hat{P}_{\nu,h}^{-1}(\tau) \right) \quad (4)$$

where N_ν denotes the number of forecast vintages, $\rho_{\tau} \equiv \rho_{\tau}(u) = u [\tau - \mathbf{1}(u < 0)]$ is the check-function and $\hat{P}_{\nu,h}$ is the cumulative distribution function.

Table 2 presents quantile scores across quantiles for different model variants, where lower numbers indicate greater forecast accuracy.

Discussion Table 2 warrants further discussion. While our in-sample results align with an extensive literature on GDP-at-risk, Table 2 presents novel evidence on the out-of-sample accuracy of various GDP-at-risk models. Interestingly we find the model with alternate fixed

⁷One small caveat to this is that, with the exception of UK quarterly GDP growth data, we rely on the latest vintage of data, and therefore do not account for back-data revisions over time

Table 1: Coefficient estimates for benchmark model and robustness exercises

	(A) Baseline					(B) Alternative Financial Conditions				
	$h = 1$	$h = 4$	$h = 8$	$h = 12$	$h = 20$	$h = 1$	$h = 4$	$h = 8$	$h = 12$	$h = 20$
Foreign variables										
For. Credit-to-GDP	-0.243 [^] [-0.066]	-0.597 [^] [-0.188 [^]]	-0.655 ^{**} [-0.357 [^]]	-0.429 [*] [-0.371 [^]]	-0.341 [^] [-0.363 [^]]	-0.013 [0.104 [^]]	-0.451 [^] [-0.018]	-0.699 [*] [-0.182 [^]]	-0.414 [^] [-0.349 [^]]	-0.679 ^{***} [-0.510 ^{***}]
For. Fin. Mkt. Vol.	-0.722 [^] [-0.440 [^]]	-0.251 [^] [-0.138]	0.135 [0.029]	0.160 [^] [0.074]	0.017 [0.023]					
For. FCI						-1.140 ^{***} [-0.056]	-0.781 [^] [0.078]	0.096 [0.157 [^]]	0.303 [^] [0.215 [^]]	0.094 [0.274 [^]]
For. GDP gr.	1.171 ^{***} [0.958 ^{***}]	0.720 ^{**} [0.635 ^{***}]	0.201 [^] [0.304 [*]]	0.026 [0.111 [^]]	0.073 [^] [0.083 [^]]	1.171 ^{***} [0.958 ^{***}]	0.720 ^{**} [0.635 ^{***}]	0.201 [^] [0.304 [*]]	0.026 [0.111 [^]]	0.073 [^] [0.083 [^]]
Domestic variables										
Credit-to-GDP	-0.166 [^] [-0.252 [^]]	-0.499 ^{***} [-0.454 ^{***}]	-0.490 ^{***} [-0.518 ^{***}]	-0.389 ^{***} [-0.469 ^{***}]	-0.336 ^{***} [-0.404 ^{***}]	-0.118 [-0.159 [^]]	-0.407 ^{**} [-0.189 [^]]	-0.355 [*] [-0.312 ^{**}]	-0.360 ^{**} [-0.384 ^{***}]	-0.407 ^{***} [-0.416 ^{***}]
Fin. Mkt. Vol.	-0.096 [-0.124]	-0.103 [-0.178 [^]]	-0.105 [-0.008]	-0.066 [-0.003]	0.019 [0.041]					
FCI						-0.382 [^] [-0.313 ^{**}]	-0.312 [^] [-0.191 [^]]	0.049 [-0.144 [^]]	-0.037 [-0.140 [^]]	0.085 [^] [0.000]
GDP gr.	0.188 [^] [0.274 [^]]	0.258 [*] [0.266 ^{***}]	0.123 [^] [0.131 [*]]	0.088 [^] [0.072 [^]]	-0.033 [^] [-0.013]	0.258 [^] [0.514 ^{***}]	0.318 [*] [0.316 ^{***}]	0.211 [^] [0.155 [*]]	0.107 [^] [0.044]	-0.011 [-0.001]
N (N^*)			13 (13)					10 (10)		
Weights (Sample)		Trade (1981Q1-2018Q4)					Trade (1981Q1-2018Q4)			
	$h = 1$	$h = 4$	$h = 8$	$h = 12$	$h = 20$	$h = 1$	$h = 4$	$h = 8$	$h = 12$	$h = 20$
Foreign variables										
For. Credit-to-GDP	-0.334 [^] [0.240 ^{**}]	-0.428 [^] [0.100 [^]]	-0.405 [^] [-0.016]	-0.208 [^] [-0.106]	-0.224 [^] [-0.195 [^]]	-0.278 [^] [0.166 [^]]	-0.536 [^] [0.108 [^]]	-0.628 ^{**} [0.027]	-0.417 [*] [-0.056]	-0.305 [^] [-0.261 [^]]
For. Fin. Mkt. Vol.	-0.737 [^] [-0.258 [^]]	-0.233 [-0.230 [^]]	0.104 [-0.067]	0.140 [^] [-0.047]	-0.025 [0.027]	-0.543 [^] [-0.042]	-0.176 [-0.021]	0.117 [0.047]	0.176 [^] [-0.008]	-0.073 [-0.042]
For. GDP gr.	1.076 ^{***} [0.792 ^{***}]	0.731 ^{**} [0.450 ^{***}]	0.207 [^] [0.229 ^{**}]	0.091 [^] [0.100 [^]]	0.056 [^] [0.075 [^]]	1.283 ^{***} [0.711 ^{***}]	0.779 ^{**} [0.491 ^{***}]	0.340 [^] [0.318 ^{***}]	0.131 [^] [0.197 ^{**}]	0.099 [^] [0.124 [^]]
Domestic variables										
Credit-to-GDP	-0.241 [^] [-0.205 ^{**}]	-0.418 ^{***} [-0.286 ^{***}]	-0.347 ^{***} [-0.384 ^{***}]	-0.316 ^{***} [-0.407 ^{***}]	-0.253 ^{***} [-0.328 ^{***}]	-0.028 [-0.111 [^]]	-0.402 ^{***} [-0.192 [*]]	-0.456 ^{***} [-0.323 ^{**}]	-0.364 ^{**} [-0.400 ^{***}]	-0.316 ^{***} [-0.408 ^{***}]
Fin. Mkt. Vol.	-0.120 [-0.198 [^]]	-0.192 [^] [0.011]	-0.135 [^] [-0.021]	-0.102 [^] [-0.021]	0.020 [-0.013]	-0.220 [-0.291 [^]]	-0.095 [-0.094 [^]]	-0.053 [-0.069]	-0.046 [-0.025]	0.111 [^] [0.079 [^]]
GDP gr.	0.243 [^] [0.461 ^{***}]	0.293 ^{**} [0.224 ^{**}]	0.155 [^] [0.129 [*]]	0.061 [^] [0.080 [^]]	-0.016 [0.044 [^]]	0.326 [^] [0.477 ^{***}]	0.478 ^{***} [0.279 ^{***}]	0.176 [^] [0.103 [^]]	0.092 [^] [0.031]	-0.006 [-0.037 [^]]
House price gr.	0.543 ^{***} [0.063]	0.187 [^] [0.055]	-0.047 [-0.012]	-0.199 [^] [-0.116 [^]]	-0.175 [*] [-0.280 ^{***}]					
Capital ratio	-0.393 [*] [0.003]	-0.042 [-0.082]	0.099 [-0.102 [^]]	0.11 [-0.074]	0.143 [^] [-0.008]					
Inflation	-1.054 ^{***} [-0.327 [^]]	-0.715 [*] [-0.278 [^]]	-0.206 [-0.149 [^]]	-0.108 [-0.070]	0.044 [0.051]					
Policy Rate	-0.122 [-0.220 [*]]	-0.450 ^{**} [-0.375 ^{***}]	-0.541 ^{***} [-0.368 ^{***}]	-0.429 ^{**} [-0.325 ^{***}]	-0.238 ^{**} [-0.171 ^{**}]					
N (N^*)			13 (13)					11 (11)		
Weights (Sample)		Trade (1981Q1-2018Q4)					Financial (1981Q1-2018Q4)			
	$h = 1$	$h = 4$	$h = 8$	$h = 12$	$h = 20$	$h = 1$	$h = 4$	$h = 8$	$h = 12$	$h = 20$
Foreign variables										
For. Credit-to-GDP	0.004 [0.171 [^]]	-0.745 [^] [-0.039]	-0.869 ^{**} [-0.196 [^]]	-0.539 ^{**} [-0.314 [^]]	-0.435 ^{**} [-0.369 [*]]	-0.012 [0.180 [^]]	-0.169 [-0.119 [^]]	-0.441 [*] [-0.263 [^]]	-0.515 ^{***} [-0.325 ^{**}]	-0.437 ^{***} [-0.302 [*]]
For. Fin. Mkt. Vol.	-1.100 [^] [0.028]	-0.548 [^] [0.005]	0.229 [^] [0.123]	0.381 [*] [0.157 [^]]	0.159 [^] [0.200 [^]]	-0.119 [-0.231 [^]]	0.215 [^] [-0.042]	0.164 [^] [0.021]	0.089 [^] [0.094 [^]]	-0.095 [^] [0.204 ^{***}]
For. GDP gr.	0.623 [*] [0.384 ^{**}]	0.453 [^] [0.157 [*]]	-0.009 [0.047 [^]]	-0.042 [0.005]	0.036 [0.021]	0.694 ^{**} [0.677 ^{***}]	0.364 ^{**} [0.360 ^{***}]	0.149 [^] [0.177 [*]]	0.100 [^] [0.117 [*]]	0.057 [^] [0.069 [^]]
Domestic variables										
Credit-to-GDP	0.126 [^] [-0.098 [^]]	-0.323 [*] [-0.100 [^]]	-0.328 ^{**} [-0.194 [*]]	-0.310 ^{**} [-0.278 ^{**}]	-0.300 ^{***} [-0.333 ^{***}]	-0.294 [^] [-0.351 ^{**}]	-0.600 ^{***} [-0.416 ^{**}]	-0.629 ^{***} [-0.442 ^{**}]	-0.556 ^{***} [-0.456 ^{***}]	-0.341 ^{***} [-0.404 ^{***}]
Fin. Mkt. Vol.	-0.409 [^] [-0.490 ^{**}]	-0.008 [-0.299 ^{**}]	-0.205 [^] [-0.253 [*]]	-0.256 [^] [-0.158 [^]]	-0.095 [^] [-0.072 [^]]	-0.034 [0.056]	-0.027 [0.092]	0.056 [0.152 [^]]	0.048 [0.160 [^]]	0.078 [^] [0.141 [*]]
GDP gr.	0.500 [*] [0.429 ^{***}]	0.280 [^] [0.240 ^{**}]	0.059 [0.064]	0.044 [0.062 [^]]	-0.007 [0.027]	0.084 [0.132 [^]]	0.115 [^] [0.078 [^]]	-0.051 [^] [0.023]	-0.058 [^] [-0.019]	-0.075 ^{**} [-0.059 [*]]
N (N^*)			13 (19)					13 (13)		
Weights (Sample)		Trade (1981Q1-2018Q4)					Trade (1991Q3-2005Q4)			

Notes: Coefficient estimates for 5th pctl. [and median]. Significance, from block bootstrap, at 32%, 10%, 5% and 1% levels denoted by [^], ^{*}, ^{**} and ^{***}.

Table 2: Absolute Quantile Scores for GDP-at-risk Model Variants

	$\tau = 0.05$	$\tau = 0.25$	$\tau = 0.5$	$\tau = 0.75$	$\tau = 0.95$
A: Baseline					
$h = 1$	0.138	0.306	0.358	0.300	0.128
$h = 4$	0.362	0.547	0.629	0.506	0.192
$h = 8$	0.420	0.725	0.839	0.631	0.216
$h = 12$	0.573	1.055	1.040	0.772	0.239
B: Alternate F.E.s					
$h = 1$	0.157	0.436	0.542	0.431	0.157
$h = 4$	0.336	0.574	0.674	0.548	0.196
$h = 8$	0.328	0.619	0.797	0.609	0.198
$h = 12$	0.454	0.842	0.975	0.739	0.225
C: UK-specific					
$h = 1$	0.169	0.462	0.554	0.426	0.115
$h = 4$	0.407	0.697	0.724	0.524	0.155
$h = 8$	0.715	0.878	0.776	0.597	0.218
$h = 12$	1.077	1.058	0.979	0.636	0.199

Note: Absolute quantile score, defined in equation (4), for various GDP-at-risk model variants.

effects, following [Canay \(2011\)](#), provide the best out-of-sample accuracy across quantiles and horizons. While there has been debate in the quantile regression on the appropriate fixed effects estimator in panel settings (see [Lamarche, 2021](#), for a review), to the best of our knowledge we are the first to consider the importance of this assumption specifically in respect to improving the accuracy of forecasts of the GDP growth distribution. For example, [Brownlees and Souza \(2021\)](#) demonstrate that quantile regression GDP-at-risk models perform poorly relative to GARCH models in-terms of forecast accuracy, but only consider country-specific quantile regressions and panel regressions based on fixed effects from [Kato et al. \(2012\)](#).

In general, our results highlight the importance of imposing structure in the quantile regression in order to improve out-of-sample forecast accuracy. In particular, limiting the number of right-hand side variables (i.e. imposing $\beta = 0$ for some x_i in the full dataset \mathbf{X}_i), estimating coefficients using panel data (i.e. imposing $\beta_i = \beta$ for all i), and restricting the fixed effects assumption (e.g. by imposing $\alpha_i^h(\tau)$ for all τ) seems to improve forecast accuracy.

Having run forecast accuracy horse-race between various quantile-regression models, we now use specification using [Canay \(2011\)](#) fixed effects to compare its estimates of the GDP growth distribution to the MPR.

Table 3: Relative Quantile Scores for Preferred GDP-at-risk Model Relative to MPR Forecasts

	$\tau = 0.05$	$\tau = 0.25$	$\tau = 0.5$	$\tau = 0.75$	$\tau = 0.95$
$h = 1$	0.740	1.090	1.116	0.994	0.699
$h = 4$	0.958	1.064	1.136	1.167	1.084
$h = 8$	0.759	1.081	1.281	1.246	1.313
$h = 12$	0.735	1.147	1.242	1.164	1.128

Note: Relative quantile score for preferred GDP-at-risk model relative to MPR. Statistics below 1 denote greater forecast accuracy of GDP-at-risk model relative to MPR.

4 Comparing GDP-at-Risk Model Estimates with the Fan Charts

In this section, we first compare the quantile forecasts from our GDP-at-risk model presented in previous section with the Bank of England GDP fan charts, focusing on a range of dimensions. We then turn to combination methods, to investigate whether we can improve density forecasts by combining GDP-at-risk model estimates with those from the Bank of England fan charts.

4.1 Comparison: Results

Quantile Scores Table 3 compares quantile scores of GDP forecasts from the Bank of England GDP fan charts with estimates from our GDP-at-risk model, where a number below 1 implies a lower quantile score (i.e. improved forecast accuracy) for the GDP-at-Risk model. A number of key results stand out. First the GDP-at-risk model consistently provides better forecasts of the 5th percentile of the GDP growth distribution relative to the Bank of England GDP fan charts. The improvement in left-tail forecast accuracy is most pronounced at medium-term horizons, at $h = 8$ and $h = 12$. However, the Bank of England GDP fan charts generally provide improved estimates of the rest of the GDP growth distribution, where this is particularly pronounced for the centre of the distribution (i.e. at $\tau = 0.5$). These results are intuitive given central bank forecasts have traditionally focused on the central projection, and given the development of GDP-at-risk models in both policy and research spheres specifically to estimate the left-tail of the GDP growth distribution

We next turn to comparisons which rely on an estimate of the full conditional GDP growth density. To ease comparison here we fit a two-piece normal distribution to the quantile regression estimates. The method we use to do so follows [Adrian et al. \(2019\)](#). We choose the three parameters (μ , σ , γ) of the two-piece normal distribution to minimize the squared distance between our estimated quantiles and the quantiles of the two-piece normal distribution from Equation 1 to match the 5, 25, 50, 75, and 95 percent quantiles.⁸

⁸[Adrian et al. \(2019\)](#) follow this approach assuming a skewed- t density and fit to the 5, 25, 75, and 95th per-

Parameter Estimates Building on Figures 1 and 2, in Figures 3 and 4 we compare the estimated parameters of the two-piece normal distribution from our GDP-at-risk model with the Bank of England GDP fan charts across horizons.

A few key results stand out. First, consistently across horizons, the GDP-at-risk model estimates notably larger uncertainty and more downside skew in the run-up to the global financial crisis than the Bank of England fan charts. In part, this reflects rapid credit growth in the UK during these years which the GDP-at-risk model associates with greater variance and more downside skew.

Second, particularly at medium-term horizons, the GDP-at-risk model consistently estimates a higher mode than the Bank of England estimates following the global financial crisis. This may reflect, for example, the fact the GDP-at-risk model misses the slowing in UK potential GDP following the crisis.

PDF Estimates In Figure 5, we compare estimated PDFs from our GDP-at-risk model with the Bank of England GDP fan charts across horizons.

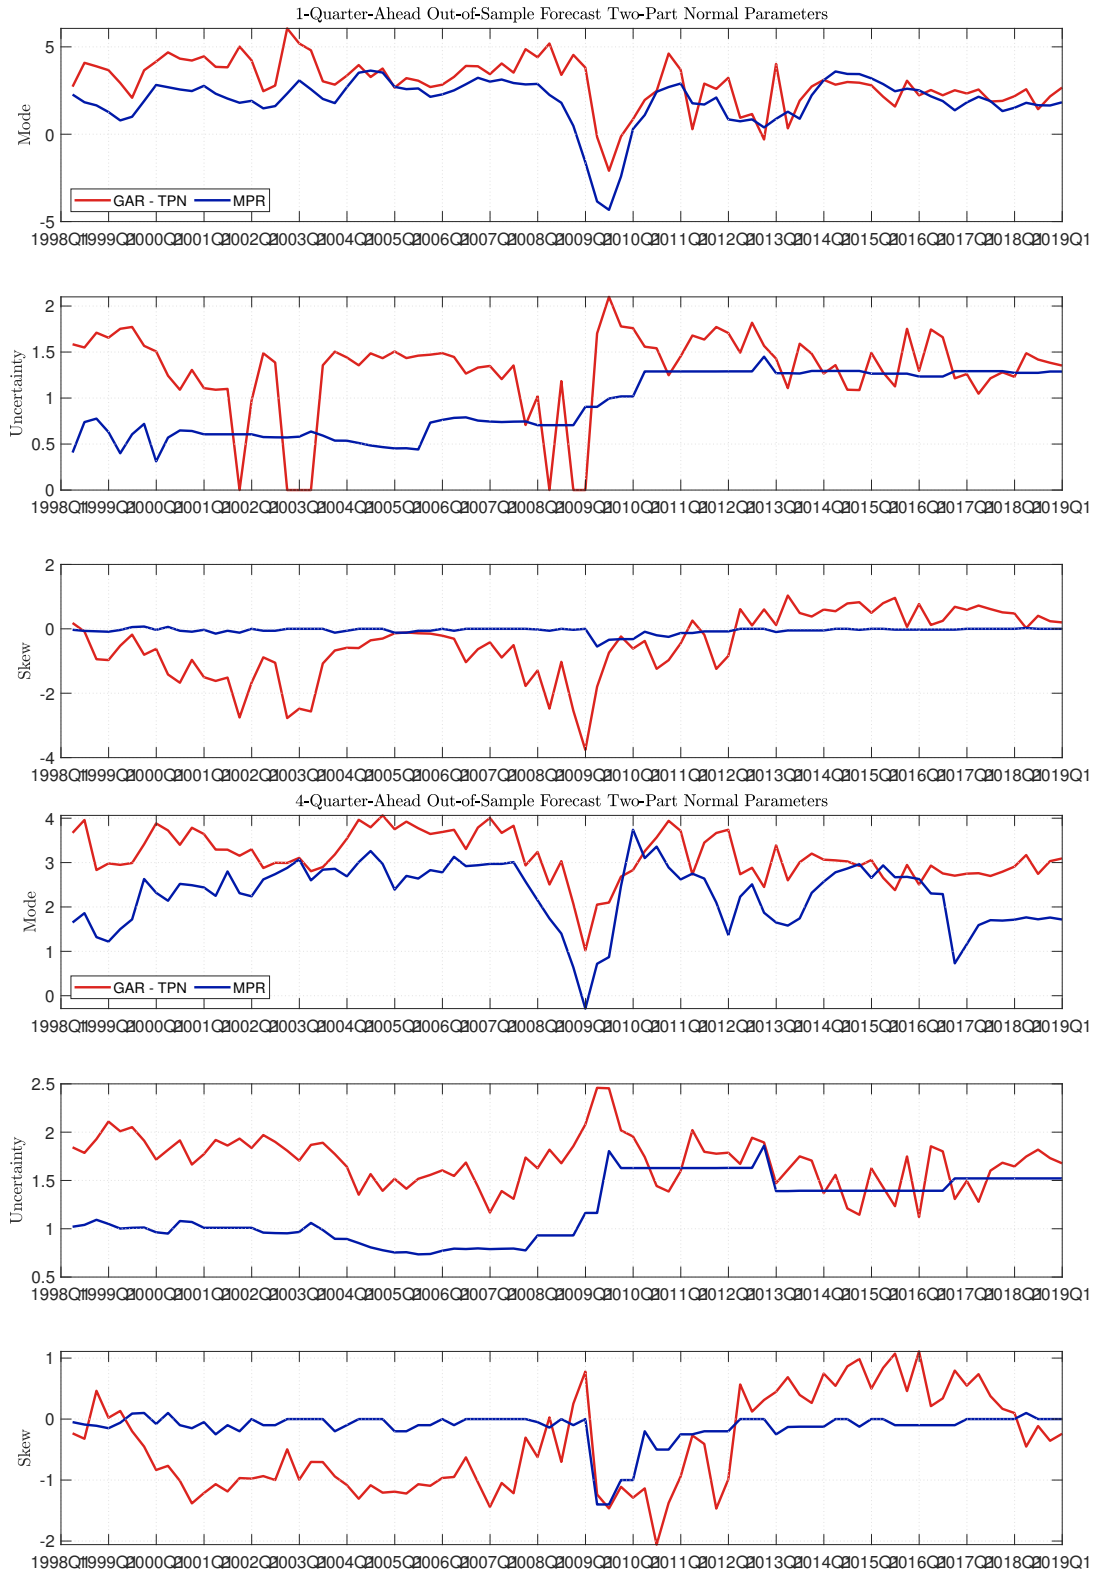
Probability Integral Transform Next, we analyse the calibration of the predictive densities from the Bank of England GDP fan charts with our GDP-at-risk model by computing the empirical cumulative distribution of the probability integral transform (PIT). This measures the percentage of observations that are below any given quantile. The closer the empirical cumulative distribution of the PITs is to the 45-degree line, the better calibrated the model is. We present results in Figure 6.

The figures illustrate that the Bank of England fan charts reflect robust predictive distributions, consistently lying within dashed confidence bands—constructed following the approach of Rossi and Sekhposyan (2019)—across all horizons. In contrast, we find some evidence of mis-specification for the GDP-at-risk model at the 1-quarter and 12-quarter horizon with estimates exceeding the confidence bands, albeit at upper quantiles. Over the whole distribution, the fan charts appear to be better calibrated generally than those from the GDP-at-risk model, consistently lying closer to the 45-degree line across each horizon. However, differences are less marked at the left-tail of the predictive GDP densities.

Discussion In general, our GDP-at-risk model seems to perform worse than the Bank of England fan chart estimates over the whole distribution. On the whole, quantile scores are higher

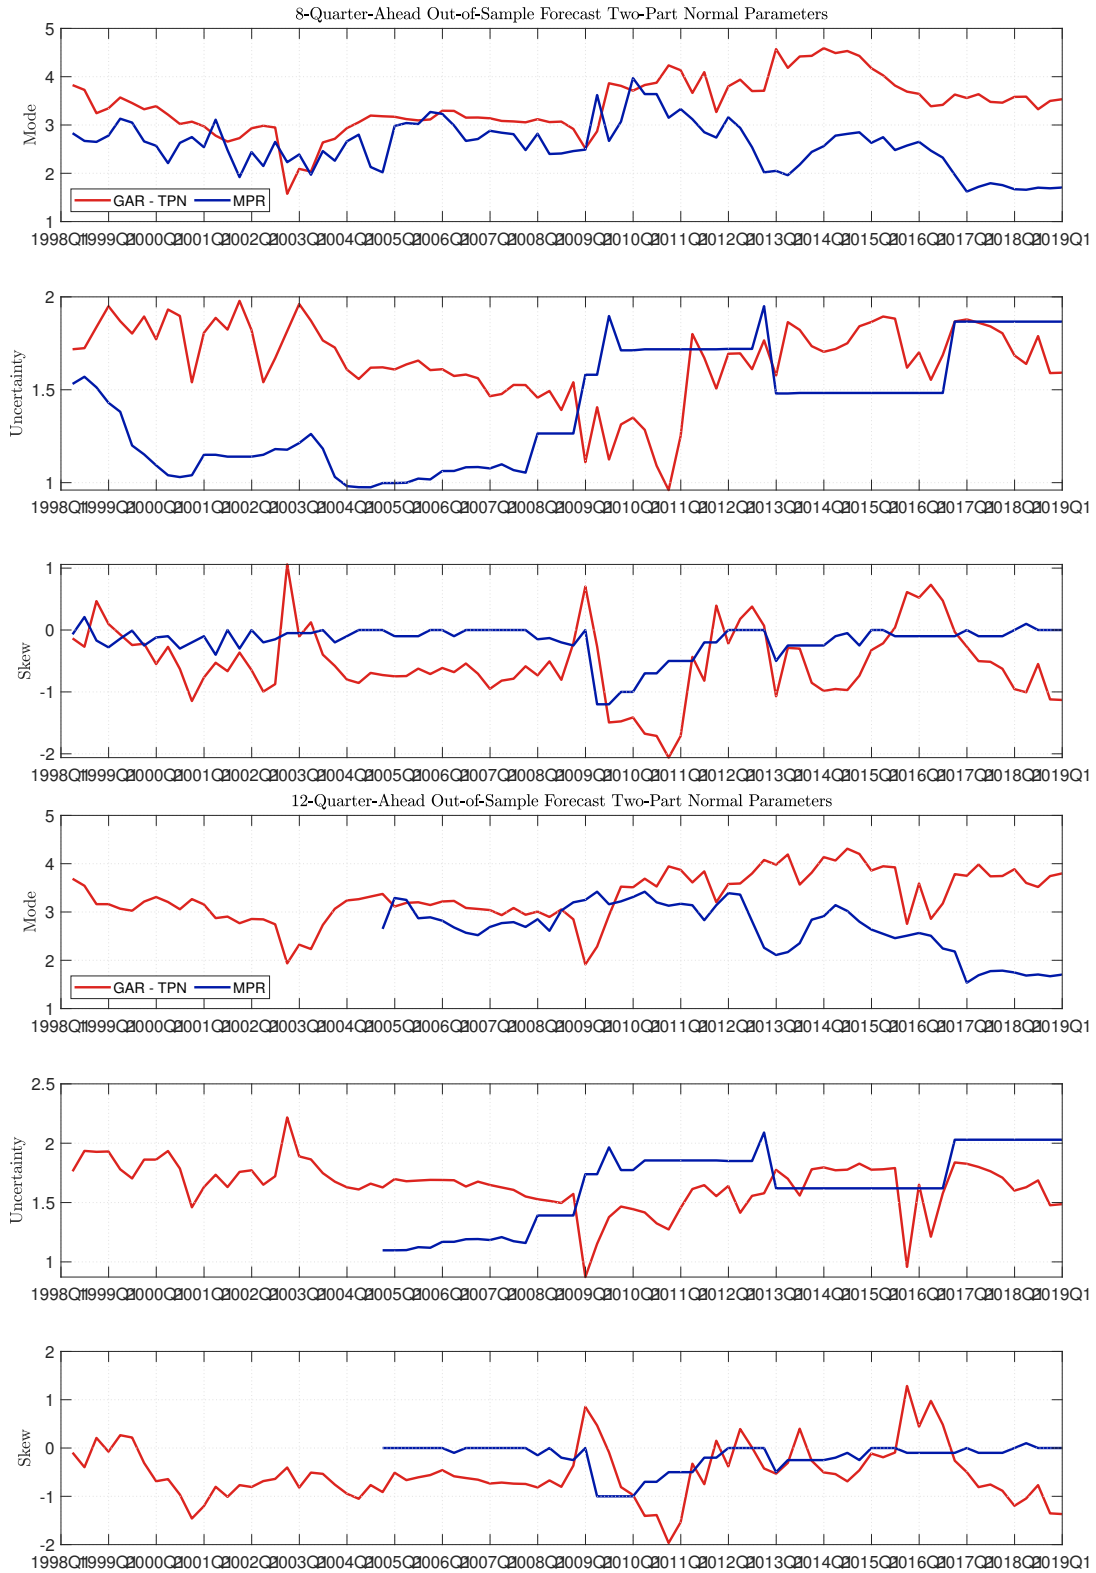
centiles. Our results are broadly unchanged when choosing to fit to other quantiles (e.g. fitting instead to the four quantiles chosen by Adrian et al. (2019)).

Figure 3: Two-Piece Normal Parameter Estimates from Fan Charts and GDP-at-risk Model



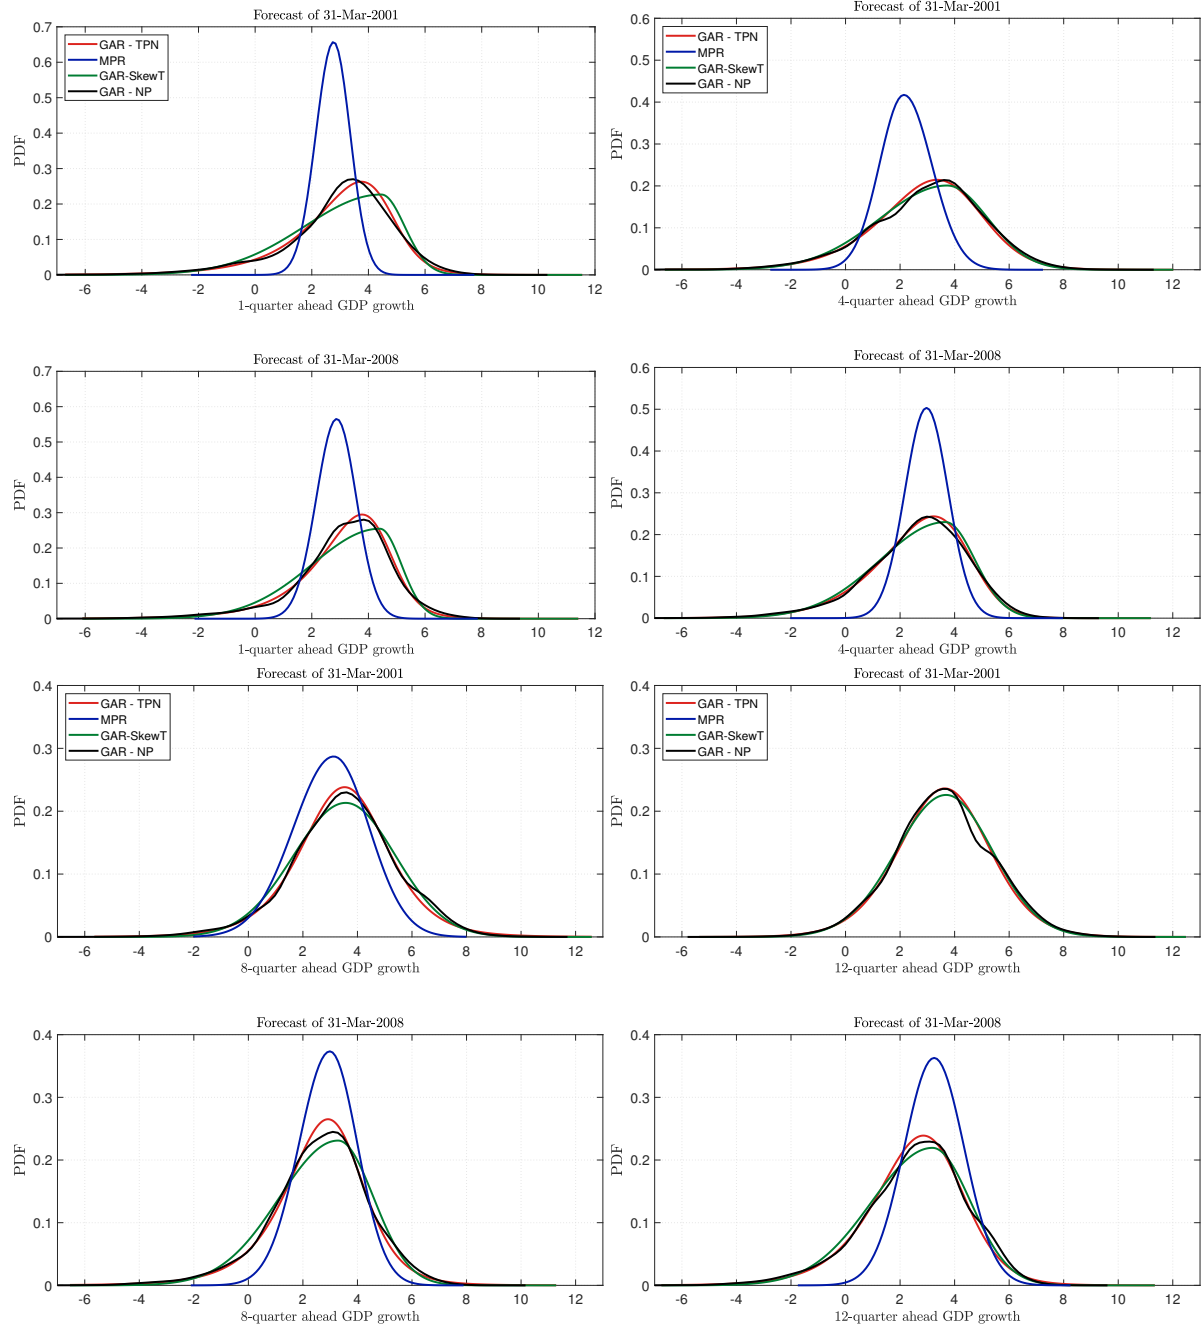
Note: Comparison between GDP-at-risk model out-of-sample predicted two-piece normal parameters from Section 3 (in red), and MPR fan chart predicted quantiles (in blue) across different horizons over time.

Figure 4: Two-Piece Normal Parameter Estimates from Fan Charts and GDP-at-risk Model



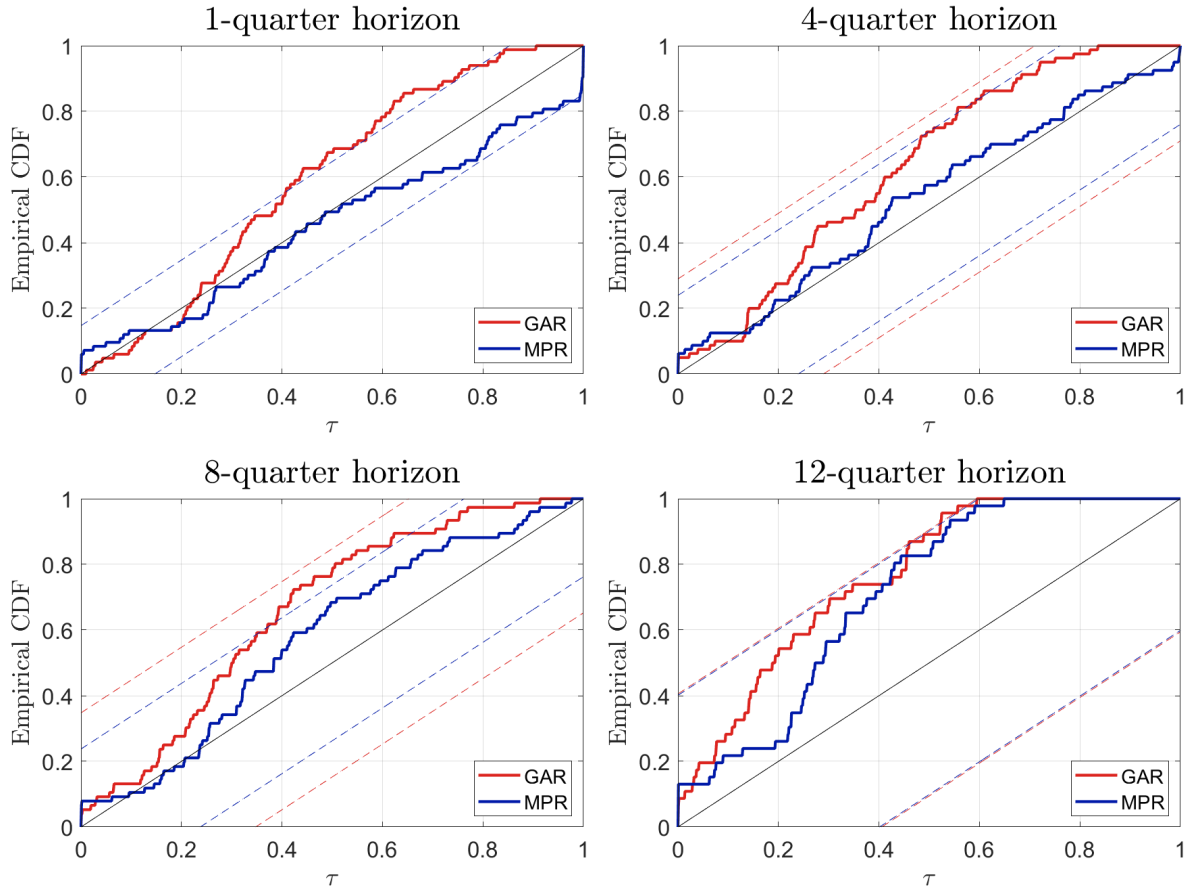
Note: Comparison between GDP-at-risk model out-of-sample predicted two-piece normal parameters from Section 3 (in red), and MPR fan chart predicted quantiles (in blue) across different horizons over time.

Figure 5: Two-Piece Normal PDF Estimates From Bank of England and GDP-at-risk Model



Note: Comparison between GDP-at-risk model out-of-sample predicted two-piece normal PDFs from Section 3 (in red), and MPR fan chart predicted quantiles (in blue) across different horizons over time.

Figure 6: PITs from Fan Charts and GDP-at-risk Model



Note: Empirical cumulative distribution of the probability integral transform (PIT) of out-of-sample estimates for the UK GDP at the 1-, 4-, 8- and 12-quarter-ahead horizons. Blue line shows the estimates from the Bank of England's MPR, while red line shows the estimates from the GDP-at-risk model. Dashed lines denote 95 confidence intervals, obtained using the method of Rossi and Sekhposyan (2019).

and PIT tests indicate that Bank of England estimates are better calibrated. However, this conclusion is not a general one and, in particular, the GDP-at-risk model appears to be particularly powerful at the left tail of the predictive distribution. This suggests that combination methods may be effective, since GDP-at-risk estimates seem to better predict left-tail of GDP growth distribution in particular. Additionally, Figures 3 and 4 highlight the potential usefulness of the GDP-at-risk model in informing estimates of higher moments of the GDP growth distribution, especially in the run-up to the global financial crisis as it is able to pick-up the effects of heightened macro-financial vulnerabilities during these years.

In next section, we consider whether we can improve estimates of GDP growth distribution by *combining* estimates from our GDP-at-risk model with Bank of England fan charts.

Table 4: Relative Quantile Scores for Combined Model Relative to MPR Forecasts

	$\tau = 0.05$	$\tau = 0.25$	$\tau = 0.5$	$\tau = 0.75$	$\tau = 0.95$
$h = 1$	0.82	1.20	1.24	1.08	0.67
$h = 4$	0.88	1.17	1.09	0.99	0.87
$h = 8$	0.77	1.03	0.98	1.03	1.07
$h = 12$	0.75	1.01	0.99	0.96	0.99

Note: Relative quantile score for combined model relative to MPR. Statistics below 1 denote greater forecast accuracy of GDP-at-risk model relative to MPR.

4.2 Combination Methods

The analysis in the previous section begs the question: could GDP predictive density estimates from the Bank of England and from our GDP-at-risk model be improved by combining these two predictions? An extensive literature has considered a range of methods for optimally combining density forecasts (see, e.g. [Aastveit, Mitchell, Ravazzolo, and van Dijk, 2018](#), for a review).

Before turning to these methods, we begin with a simple combination approach. In particular, we take the modal estimate from the MPR, and combine with the estimates of the skew and uncertainty from the GDP-at-risk model (as presented in Figures 3 and 4). While this approach is unlikely to be optimal from a forecasting perspective, it potentially offers a simple and practical proposal for central banks to use GDP-at-risk models to inform GDP fan charts.

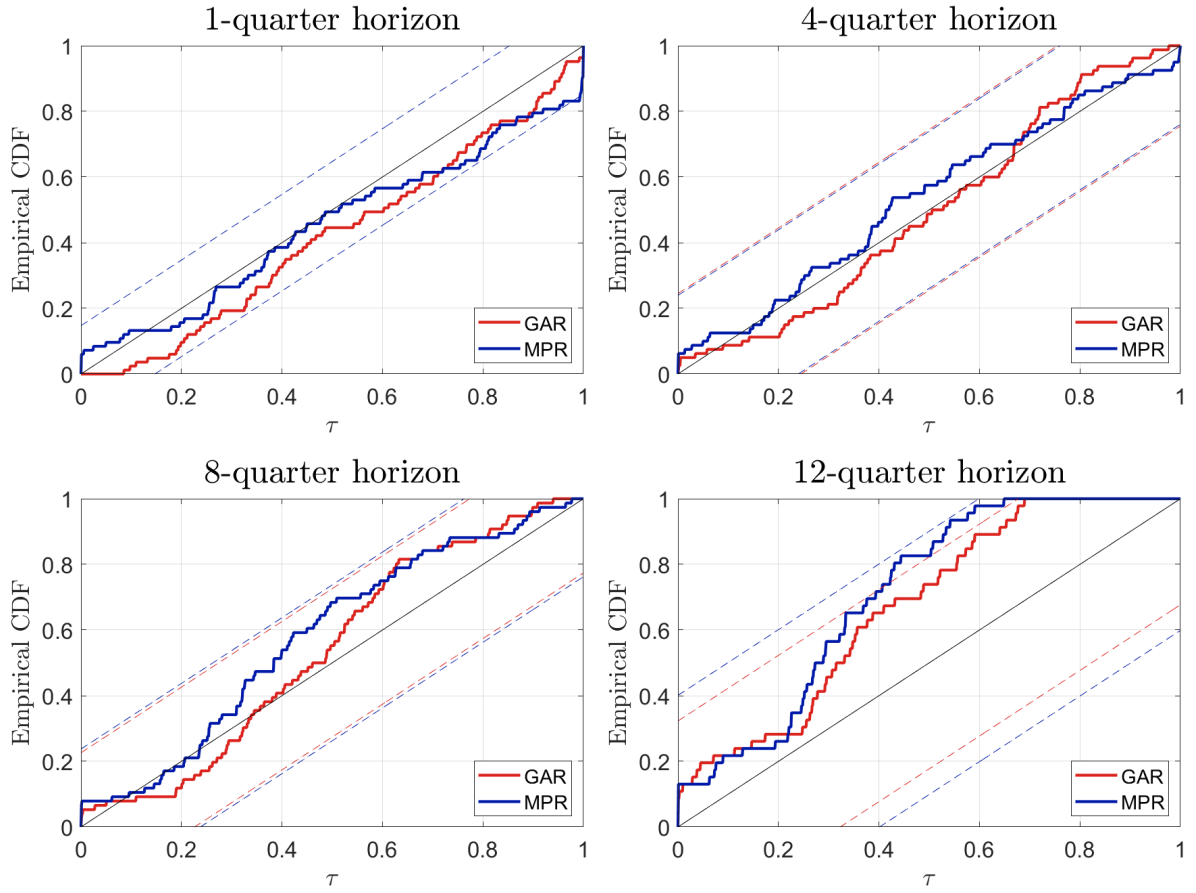
Figure 7 highlights that even this simple approach can improve calibration of the Bank of England’s GDP fan charts. In particular, at the 4-, 8-, and 12-quarter-ahead horizon the combination method notably improves calibration of the fan charts, with the empirical cumulative distribution of PITs lying closer to the 45-degree line than the estimates from the Bank of England.

This is further confirmed by the quantile scores from this combination method show in Table 4. In general, across quantiles the combination method performs better than the Bank of England estimates, where this is particularly noticeable at the 5th percentile, and at medium-term horizons.

5 Conclusion

Both monetary and macroprudential policymakers use conditional density forecasts of GDP to inform their policy decisions. The former tend to use ‘fan charts’ that convey uncertainty around central projections from linearised macroeconomic models. The latter draw on estimates of ‘growth-at-risk’, typically estimated using quantile-regression techniques, reflecting

Figure 7: PITs from Combined Forecast Densities



Note: Empirical cumulative distribution of the probability integral transform (PIT) of out-of-sample estimates for the UK GDP at the 1-, 4-, 8- and 12-quarter-ahead horizons. Blue line shows the estimates from the Bank of England’s MPR, while red line shows the estimates from combining GDP-at-risk model-estimates of higher-order moments with the Bank of England’s MPR estimate of the mode. Dashed lines denote 95 confidence intervals, obtained using the method of Rossi and Sekhposyan (2019).

their focus on the tails of the GDP distribution. Focusing on the UK, we study how the fan charts constructed by the Bank of England’s MPC compare to predictive densities derived from statistical growth-at-risk models. We find that GDP-at-risk models provide improved estimates of the left-tail of the GDP-growth distribution, in particular at medium-term horizons compared to the Bank of England’s GDP fan charts. However, GDP-at-risk models generally perform worse than the fan charts at the centre of the distribution. Combining forecast densities in a parsimonious manner provides the best forecasts, with limited losses compared to optimal density combination methods. Doing so offers central banks the opportunity to improve GDP density forecasts, and unify the narrative provided by monetary and macroprudential policymakers about the economic outlook and uncertainty around it.

References

- AASTVEIT, K. A., J. MITCHELL, F. RAVAZZOLO, AND H. VAN DIJK (2018): “The Evolution of Forecast Density Combinations in Economics,” Tinbergen Institute Discussion Papers 18-069/III, Tinbergen Institute.
- ADAMS, P. A., T. ADRIAN, N. BOYARCHENKO, AND D. GIANNONE (2021): “Forecasting macroeconomic risks,” *International Journal of Forecasting*, 37, 1173–1191.
- ADRIAN, T., N. BOYARCHENKO, AND D. GIANNONE (2019): “Vulnerable Growth,” *American Economic Review*, 109, 1263–1289.
- (2021): “Multimodality In Macrofinancial Dynamics,” *International Economic Review*, 62, 861–886.
- ADRIAN, T. AND F. DUARTE (2018): “Financial Vulnerability and Monetary Policy,” CEPR Discussion Papers 12680, C.E.P.R. Discussion Papers.
- ADRIAN, T., F. GRINBERG, N. LIANG, S. MALIK, AND J. YU (2022): “The Term Structure of Growth-at-Risk,” *American Economic Journal: Macroeconomics*, Forthcoming.
- AIKMAN, D., J. BRIDGES, S. BURGESS, R. GALLETTY, I. LEVINA, C. O’NEILL, AND A. VARADI (2018): “Measuring risks to UK financial stability,” Bank of England working papers 738, Bank of England.
- AIKMAN, D., J. BRIDGES, S. HACIOGLU HOKE, C. O’NEILL, AND A. RAJA (2019): “Credit, capital and crises: a GDP-at-Risk approach,” Bank of England working papers 824, Bank of England.
- BRITTON, E., P. FISHER, AND J. WHITLEY (1998): “Home The Inflation Report projections: understanding the fan chart,” *Bank of England Quarterly Bulletin*, Q1.
- BROWNLEES, C. AND A. B. SOUZA (2021): “Backtesting global Growth-at-Risk,” *Journal of Monetary Economics*, 118, 312–330.
- BURGESS, S., E. FERNANDEZ-CORUGEDO, C. GROTH, R. HARRISON, F. MONTI, K. THEODORIDIS, AND M. WALDRON (2013): “The Bank of England’s forecasting platform: COMPASS, MAPS, EASE and the suite of models,” Bank of England working papers 471, Bank of England.
- CANAY, I. A. (2011): “A simple approach to quantile regression for panel data,” *The Econometrics Journal*, 14, 368–386.

- CARNEY, M. (2020): “The Grand Unifying Theory (and practice) of Macroprudential Policy,” Speech at Logan Hall, University College London, Bank of England.
- CLEMENTS, M. (2004): “Evaluating the Bank of England Density Forecasts of Inflation,” *Economic Journal*, 114, 844–866.
- DOWD, K. (2007): “Too good to be true? The (In)credibility of the UK inflation fan charts,” *Journal of Macroeconomics*, 29, 91–102.
- ELDER, R., G. KAPATANIOS, T. TAYLOR, AND T. YATES (2005): “Assessing the MPC’s fan charts,” *Bank of England Quarterly Bulletin*, 45.
- GALVAO, A. F., J. GU, AND S. VOLGUSHEV (2020): “On the unbiased asymptotic normality of quantile regression with fixed effects,” *Journal of Econometrics*, 218, 178–215.
- GALVÃO, A., A. GARRATT, AND J. MITCHELL (2021): “Does judgment improve macroeconomic density forecasts?” *International Journal of Forecasting*, 37, 1247–1260.
- GIGLIO, S., B. KELLY, AND S. PRUITT (2016): “Systemic risk and the macroeconomy: An empirical evaluation,” *Journal of Financial Economics*, 119, 457–471.
- JORDÀ, O. (2005): “Estimation and Inference of Impulse Responses by Local Projections,” *American Economic Review*, 95, 161–182.
- KATO, K., A. F. GALVAO, AND G. V. MONTES-ROJAS (2012): “Asymptotics for panel quantile regression models with individual effects,” *Journal of Econometrics*, 170, 76–91.
- KNÜPPEL, M. AND G. SCHULTEFRANKENFELD (2012): “How Informative Are Central Bank Assessments of Macroeconomic Risks?” *International Journal of Central Banking*, 8, 87–139.
- (2019): “Assessing the uncertainty in central banks’ inflation outlooks,” *International Journal of Forecasting*, 35, 1748–1769.
- KOOP, G. AND D. KOROBILIS (2014): “A new index of financial conditions,” *European Economic Review*, 71, 101–116.
- LAMARCHE, C. (2021): “Quantile Regression for Panel Data and Factor Models,” *Oxford Research Encyclopedia of Economics and Finance*.
- LLOYD, S., E. MANUEL, AND K. PANCHEV (2022): “Foreign vulnerabilities, domestic risks: the global drivers of GDP-at-Risk,” Bank of England working papers 940, Bank of England.

- MITCHELL, J. AND S. HALL (2005): “Evaluating, Comparing and Combining Density Forecasts Using the KLIC with an Application to the Bank of England and NIESR ‘Fan’ Charts of Inflation*,” *Oxford Bulletin of Economics and Statistics*, 67, 995–1033.
- PLAGBORG-MØLLER, M., L. REICHLIN, G. RICCO, AND T. HASENZAGL (2020): “When Is Growth at Risk?” *Brookings Papers on Economic Activity*, 167–229.
- ROSSI, B. AND T. SEKHPOSYAN (2019): “Alternative tests for correct specification of conditional predictive densities,” *Journal of Econometrics*, 208, 638–657.
- SCHULARICK, M. AND A. M. TAYLOR (2012): “Credit Booms Gone Bust: Monetary Policy, Leverage Cycles, and Financial Crises, 1870-2008,” *American Economic Review*, 102, 1029–1061.
- SUAREZ, J. (2021): “Growth-at-risk and macroprudential policy design,” ESRB Occasional Paper Series 19, European Systemic Risk Board.
- WALLIS, K. (2004): “An Assessment of Bank of England and National Institute Inflation Forecast Uncertainties,” *National Institute Economic Review*, 189, 64–71.