

Boeck, Maximilian; Pfarrhofer, Michael

**Article — Published Version**

## Belief Shocks and Implications of Expectations About Growth-at-Risk

Journal of Applied Econometrics

**Provided in Cooperation with:**

John Wiley & Sons

*Suggested Citation:* Boeck, Maximilian; Pfarrhofer, Michael (2025) : Belief Shocks and Implications of Expectations About Growth-at-Risk, Journal of Applied Econometrics, ISSN 1099-1255, Wiley, Hoboken, NJ, Vol. 40, Iss. 3, pp. 341-348, <https://doi.org/10.1002/jae.3117>

This Version is available at:

<https://hdl.handle.net/10419/319328>

### Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

### Terms of use:

*Documents in EconStor may be saved and copied for your personal and scholarly purposes.*

*You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.*

*If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.*



<http://creativecommons.org/licenses/by/4.0/>

## REPLICATION

## OPEN ACCESS



# Belief Shocks and Implications of Expectations About Growth-at-Risk

Maximilian Boeck<sup>1</sup> | Michael Pfarrhofer<sup>2</sup><sup>1</sup>Friedrich-Alexander-University Erlangen-Nuremberg, Nuremberg, Germany | <sup>2</sup>Vienna University of Economics and Business, Vienna, Austria**Correspondence:** Maximilian Boeck ([maximilian.boeck@fau.de](mailto:maximilian.boeck@fau.de))**Received:** 13 February 2024 | **Revised:** 22 November 2024 | **Accepted:** 23 November 2024**Funding:** Pfarrhofer received financial support from the Jubiläumsfonds of the Oesterreichische Nationalbank (OeNB, project 18765).**Keywords:** density forecasting | nowcast | quantile regression | tail risk

## ABSTRACT

This paper revisits the question of how shocks to expectations of market participants can cause business cycle fluctuations. We use a vector autoregression to estimate dynamic causal effects of belief shocks which are extracted from nowcast errors about output growth. In a first step, we replicate and corroborate the findings of Enders, Kleemann, and Müller (2021). The second step computes nowcast errors about growth-at-risk at various quantiles. This involves both recovering the quantiles of the nowcast distribution of output growth from the Survey of Professional Forecasters, and, since the true quantiles of output growth are unobserved, estimating them with quantile regressions. We document a lack of distinct patterns in response to shocks arising from nowcasts misjudging macroeconomic risk. Although the differences are statistically insignificant, belief shocks about downside risk seem to produce somewhat sharper business cycle fluctuations.

**JEL Classification:** C22, C32, E32, E71

## 1 | Introduction

This paper revisits the question of how shocks to expectations of market participants can cause business cycle fluctuations (see, e.g., Beaudry and Portier 2006, 2014). We build upon the empirical framework of Enders, Kleemann, and Müller (2021) who discuss the identification of belief shocks. Usually, *expectations* about variables enter theoretical models as expected values conditional on some information set. In empirical models, they are often operationalized as mean-based predictions. That is, they are summary statistics of predictive distributions. This potentially disregards higher-order moments and abstracts from implied judgements of economic agents about macroeconomic risks.<sup>1</sup>

We diverge partly from the related literature and explicitly leverage full predictive distributions. Our first contribution is to

replicate the findings of Enders, Kleemann, and Müller (2021), which are based on using the median consensus nowcast for US real gross domestic product (GDP) growth from the Survey of Professional Forecasters (SPF). They use this data to compute real-time nowcast errors (NEs, defined as the difference between realized output growth and nowcasts). Their identification scheme then discriminates between belief and non-belief shocks via sign restrictions in a vector autoregression (VAR). They impose the belief shock to induce a negative co-movement between output and NEs, while the non-belief shock causes a positive co-movement. This reflects the notion that a favorable belief shock (e.g., an overly optimistic outlook due to noisy signals observed by markets) shifts predictions beyond actual growth. This results in a negative NE while increasing output growth. The second contribution is to move beyond the point prediction and to recover the probabilistic distribution of nowcasted output growth in the United States. This aspect of our work

This is an open access article under the terms of the [Creative Commons Attribution](https://creativecommons.org/licenses/by/4.0/) License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

© 2025 The Author(s). *Journal of Applied Econometrics* published by John Wiley & Sons Ltd.

is motivated by the recent interest in macroeconomic (tail) risk due to the influential paper by Adrian, Boyarchenko, and Giannone (2019). A common definition of macroeconomic risk is to consider the conditional quantiles of some underlying series of interest. The quantiles of output growth, for instance, are referred to as growth-at-risk (GaR) at some pre-defined probability. Related work finds time-varying macroeconomic risk to be at least partly predictable (see, e.g., Adams et al. 2021; Clark et al. 2023).

Evaluating and processing nowcasts of tail risk comes with two main challenges. First, the SPF does not contain probabilistic predictions in a format required for our analysis. Our solution is to rely on the ensemble methods proposed by Krüger and Nolte (2016) to recover the implied predictive distribution from individual point forecasts of SPF participants. Second, the quantiles of the dynamic process governing output growth are not observed. Here, we rely on time-varying parameter quantile regression (TVP-QR as in Pfarrhofer 2022) to estimate the real-time quantiles of output growth (see also Loria, Matthes, and Zhang 2024, for a related approach). These two ingredients are used to compute reduced form GaR-NEs (which measure an inaccurate assessment of potential best/worst-case scenarios about the economic outlook; in this context, see, e.g., the [Federal Reserve Bank of New York Outlook-at-Risk dashboard](#)), which we then employ to study the dynamic effects of belief shocks.

Our empirical results can be summarized as follows. First, we demonstrate the replicability of the results of Enders, Kleemann, and Müller (2021) in a narrow sense. Replicability in this context refers to extending the sampling period and considering alternative specifications in addition to the original implementations. Second, in a wide sense, we investigate whether belief shocks about GaR induce distinct business cycle fluctuations. The answer to this question is no. Indeed, we find very similarly shaped responses. This is true both when comparing our quantile-based estimates to the original framework, but also, when we compare, for example, downside to upside risk. We conjecture that this is because the nowcast distribution of output growth, which we recover from the SPF, turns out to be rather symmetric for most of the sample. This notion relates to the nowcast setting (as opposed to forecasts), which are typically the focus of studies on macroeconomic risk (see Adrian et al. 2022). Although the differences are statistically insignificant, belief shocks about downside risk seem to produce somewhat sharper business cycle fluctuations.

The paper proceeds as follows. In Section 2, we describe the SPF data, the original framework to extract belief shocks, and our extensions to recovering a suitable probabilistic predictive distribution. Section 3 discusses the VAR framework and identification procedure for belief shocks, before moving on to our main empirical results. Section 4 concludes.

## 2 | Survey and Nowcast Distributions

### 2.1 | The Survey of Professional Forecasters

Predictions about the current state of the economy lie at the heart of our paper. The SPF is a quarterly survey of such macroeconomic predictions, maintained by the [Federal Reserve Bank of Philadelphia](#). It collects projections, from a changing panel of participants, which are submitted in form of a single number which is the point forecast of the target variable by the respective forecaster. For many applications, it is sufficient to aggregate these by computing the mean or median consensus prediction at any given point in time for any desired forecast horizon. This produces a sequence of (point) forecasts and yields a single time series based on an equal-weighted combination. Indeed, this is the default format for downloading SPF data and what Enders, Kleemann, and Müller (2021) use in their original paper.

In our paper, as an extension, we intend to identify belief shocks related to the full predictive distribution. They are extracted from NEs about GaR. While the SPF in principle collects probabilistic forecasts, these come in the form of probability bins for a specific horizon and transformation of output, which we thus cannot use for our purposes. For this reason, we recover nowcast distributions from point predictions of individual forecasters as follows. Suppose that the predictive density  $y_t^{(f,i)}$  of forecaster  $i$  at time  $t$  has mean  $\mu_{it}$  and variance  $\varsigma_{it}^2$ . We observe the point prediction  $\mu_{it}$  (made at time  $t$ , when realizations were not yet available because GDP is released with a lag) of individual forecaster  $i = 1, \dots, N_t$ , but we do not observe the variance  $\varsigma_{it}^2$ .

Due to the design of the SPF, with coverage and number of forecasters  $N_t$  varying over time, we aim to construct an ensemble forecast using participants that are exchangeable. Define the average mean forecast  $\bar{\mu}_t = N_t^{-1} \sum_{i=1}^{N_t} \mu_{it}$  and forecaster disagreement  $s_t^2 = (N_t - 1)^{-1} \sum_{i=1}^{N_t} (\mu_{it} - \bar{\mu}_t)^2$  as the cross-sectional mean and variance of the point predictions across all survey forecasts. What is missing here is the unpredictable randomness (encoded in  $\varsigma_{it}^2$ ) of the target series surrounding the nowcasts. To solve this issue, we follow Krüger and Nolte (2016) and use an ensemble method, which has been shown to work well for the SPF.

In particular, we assume that the ensemble nowcast distribution  $y_t^{(f)}$  can be written as an equal-weighted mixture of Gaussians with a common (but unknown) variance:

$$y_t^{(f)} = \frac{1}{N_t} \sum_{i=1}^{N_t} \mathcal{N}(\mu_{it}, \varsigma^2), \quad (1)$$

where  $\varsigma^2$  is the sole parameter to be estimated. Note that this implies that the joint forecast distribution  $y_t^{(f)}$  has a mean equal to the average across forecasters,  $\bar{\mu}_t$ , and its variance is given by  $\bar{\varsigma}_t^2 = s_t^2 + \varsigma^2$ . The first variance component measures dispersion over the cross-section of forecasters, while the second reflects the unpredictable part of GDP. Note that even though the individual components in the sum of Equation (1) are Gaussian, the mixture allows for nowcast distributions with highly non-Gaussian features such as skewness, multi-modality or heavier-than-normal tails (see, e.g., Frühwirth-Schnatter 2006, for a detailed discussion). Estimates for  $\varsigma^2$  are obtained from a

rolling window of  $\tau = 12$  quarters, that is, 3 years worth of quarterly nowcasts. We optimize the parameter using the continuous ranked probability score (CRPS) as predictive loss that we seek to minimize for the pool of forecasters.<sup>2</sup>

It remains to explicitly define the minimization problem. We use  $y_t^{(r)}$  to denote the real-time realization of GDP growth and assume normally distributed nowcasts, in line with Equation (1):  $\hat{\zeta}_\tau^2 = \min_{\zeta^2} \left( \sum_{t=(\tau-w)}^{\tau-1} \sum_{i=1}^{N_t} \text{CRPS}(\mathcal{N}(\mu_{it}, \zeta^2), y_t^{(r)}) \right)$ . Moving the rolling window forward yields a sequence of estimates  $\hat{\zeta}_t^2$ . Two features are worth mentioning. First, this introduces time-varying variances, as we have estimates for each  $t$ , and  $w$  governs the persistence of these estimates. Second, even though the  $\hat{\zeta}_t^2$ 's are dated at time  $t$ , we only use information up to  $t - 1$  (i.e., we do not mix information sets). We may use this procedure to obtain Monte Carlo samples or quantiles from the ensemble predictive distribution in Equation (1):

$$\frac{1}{N_t} \sum_{i=1}^{N_t} \mathcal{N}(\mu_{it}, \hat{\zeta}_t^2). \quad (2)$$

Indeed, this is how we obtain the full predictive distribution of the SPF ensemble nowcast whose  $p$ th quantile,  $y_{pt}^{(f)}$ , is the ensemble nowcast for GaR at quantile  $p$  at time  $t$ . The consensus nowcast is given by  $\bar{y}_t^{(f)} = \bar{\mu}_t$ .

## 2.2 | Estimating Real-Time Quantiles of Output Growth

The final challenge is that the *true* quantiles of output growth are not observed. Each observation of GDP is just a single realization of an underlying stochastic process, and comparing these realizations to our GaR expectations is futile. Next, we discuss how we estimate the real-time latent quantiles of the GDP growth process, which we require to compute NEs quantile-by-quantile (by contrast, having access to explicit measures of quantiles is not necessary when the focus is on out-of-sample predictive inference and designing scoring rules used for model selection, see, e.g., Gneiting and Ranjan 2011).

Our baseline framework is similar to Loria, Matthes, and Zhang (2024); that is, to estimate the unobserved quantiles of output, we use a variant of quantile regression. In particular, given its substantial degree of flexibility, our implementation is based on a Bayesian time-varying parameter quantile regression (TVP-QR) as in Pfarrhofer (2022). Let  $\{y_t\}_{t=1}^T$  denote a scalar dependent variable,  $\{\mathbf{x}_t\}_{t=1}^T$  comprises  $K$  predictors at time  $t = 1, \dots, T$ , and  $q_p(\mathbf{x}_t) = \mathbf{x}_t' \boldsymbol{\beta}_{pt}$  is the  $p$ th quantile function of  $y_t$  given  $\mathbf{x}_t$  for  $p \in (0, 1)$ . We use a model of the form  $y_t = \mathbf{x}_t' \boldsymbol{\beta}_{pt} + \varepsilon_t$  with  $\int_{-\infty}^0 f_p(\varepsilon_t) d\varepsilon_t = p$ , that is, the  $p$ th quantile of the error distribution  $f(\bullet)$  is equal to zero. Specifically, we assume  $\varepsilon_t$  to follow an asymmetric Laplace (AL) distribution with scale  $\sigma_p^2$ . The  $\boldsymbol{\beta}_{pt}$ 's are quantile-specific vectors of size  $K \times 1$  which collect the parameters that vary over time. We assume an independent random walk state equation for each of these parameters and rely on a dynamic shrinkage prior for regularization.

This framework consists of two crucial ingredients. First, it allows coefficients to vary at quantile  $p$ , allowing for heterogeneous

effects across specific parts of the distribution of  $y_t$ . Second, the magnitudes of these effects are allowed to vary over time. The former reflects the literature on measuring tail risks of GDP growth, following Adrian, Boyarchenko, and Giannone (2019), while the latter allows for another layer of nonlinearity that has been found to improve accuracy (see, e.g., the corresponding discussion in Clark et al. 2024). The target variable  $y_t$  is the respective first released vintage of the annualized growth rate of real GDP. These data are from the *Real-Time Data Set for Macroeconomists*. We choose these vintages such that our target variable is as close as possible to the conceptual variable that the participants of the SPF were asked to forecast at the time.

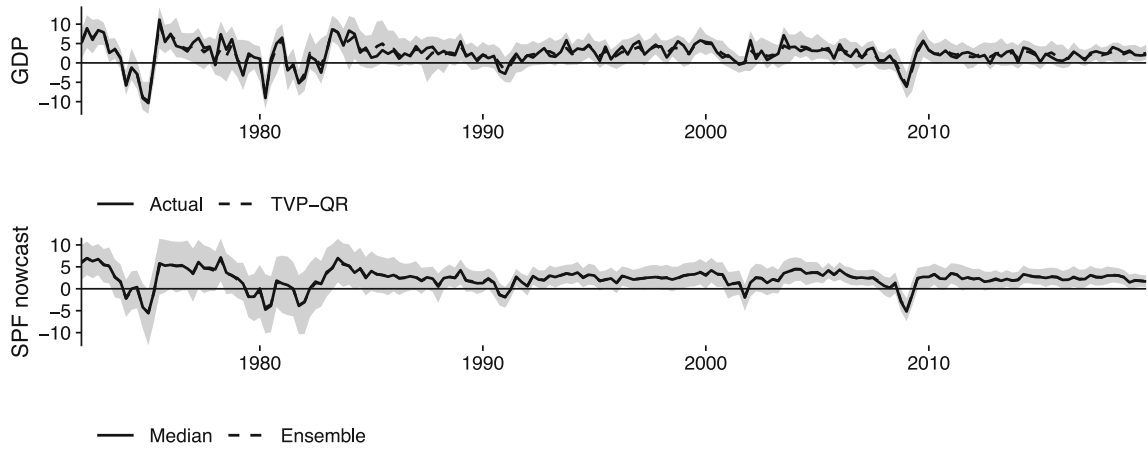
Our vector of predictors contains  $Q$  common factors  $\mathbf{f}_t$  that drive economic fluctuations in the US economy. In particular, we use a macroeconomic real-time dataset of 80 variables such that it resembles the potential information set a forecaster of the SPF has access to.<sup>3</sup> We extract  $Q = 4$  factors following Stock and Watson (2002). Additionally, we add lags of the National Financial Conditions Index (NFCI), labeled  $z_t$ , which has been identified as an important variable that shifts the quantiles of GDP, lags of the dependent variable and an intercept term. We use  $P = 4$  lags such that  $\mathbf{x}_t = (1, y_{t-1}, \dots, y_{t-P}, z_{t-1}, \dots, z_{t-P}, \mathbf{f}_{t-1}', \dots, \mathbf{f}_{t-P}')'$ .

Running our algorithm produces estimates for the quantiles of GDP growth for each point in time, that is, the fitted “realized” values  $\hat{y}_{pt}^{(r)} = \mathbf{x}_t' \hat{\boldsymbol{\beta}}_{pt}$ , which we interpret as the best possible estimate of the *true real-time* quantile (we stress that these quantiles are subject to potential measurement errors, and some of our results below thus must be interpreted with caution). Note that our procedure is based on an expanding window of observations, such that the information sets of the quantile model and the one available to the SPF forecasters is consistent. Our implementation is fully Bayesian, which implies that we obtain a posterior distribution for the fitted quantiles. We summarize this distribution by taking the posterior median at a particular quantile of interest.

## 2.3 | Empirical Estimates for Growth-at-Risk Nowcast Errors

Our full sample runs from 1968Q4 to 2019Q4. Figure 1 shows real GDP growth and the quantiles estimated with TVP-QR in the upper panel. The shaded areas range between the 10th and 90th percentile. The solid line is actual GDP growth, while the dashed line indicates the 50th percentile estimated using TVP-QR. The lower panel is a chart of the SPF nowcast distribution. The shaded area again reflects the 80 percent credible set, “Median” indicates the default SPF aggregation as used in Enders, Kleemann, and Müller (2021) whereas “Ensemble” is the point nowcast arising from using the methods of Krüger and Nolte (2016) as described above. These two approaches coincide for the mean/median, but our implementation yields a full predictive distribution.

Next, we formally define the versions of NEs that we consider descriptively and in our structural application. We previously denoted the consensus nowcast with  $y_t^{(f)}$  and the realization as  $y_t^{(r)}$ . These serve as the basis for the variant that replicates Enders, Kleemann, and Müller (2021). Recall that our estimates from TVP-QR yield the quantile-based counterpart for observed GDP,



**FIGURE 1** | Real GDP growth, estimated quantiles, and SPF nowcast distribution. Shaded areas show the 10th and 90th percentiles around the median; “Median” refers to the default SPF.

$\hat{y}_{pt}^{(r)}$ , and that the quantiles of the distribution in Equation (2) define our SPF nowcast of the  $p$ th quantile,  $y_{pt}^{(f)}$ . The NEs are

$$ne_t = y_t^{(r)} - \bar{y}_t^{(f)}, \quad (3)$$

$$ne_{pt} = \hat{y}_{pt}^{(r)} - y_{pt}^{(f)}. \quad (4)$$

We purge the NEs of any remaining predictable components by running ARIMA models with automatic lag selection.<sup>4</sup> Figure 2 shows the resulting NEs; “Actual” refers to those for the consensus nowcast and actual GDP observations, as in Equation (3), while the colored lines mark those for GaR at the indicated quantile based on Equation (4). The baseline NEs are identical (up to a scaling factor) to those shown in Figure 1 of Enders, Kleemann, and Müller (2021). The dynamics of our quantile-based versions are similar. Indeed, all our versions of NEs are positively correlated at varying strengths (additional results are provided in the Supporting Information). For the median, the correlation exceeds 0.8, and we conclude that our framework to extract quantile-based nowcast errors yields reasonable results.

### 3 | Belief Shocks and Their Effects on the Business Cycle

We now turn to the specification and identification of the structural VAR that we use to recover dynamic causal effects of belief shocks. This section contains the main empirical results: those for our narrow replication of the original implementation and our extension focusing on nowcasts about GaR.

#### 3.1 | The Vector Autoregression and Structural Identification

As laid out in Enders, Kleemann, and Müller (2021), belief shocks can be extracted from the NE. We use a bivariate VAR that features the NE ( $ne_t$  for the plain version, and  $ne_{pt}$  when we consider GaR) and output, labeled  $y_t$ , as endogenous variables. We

stack these in the vector  $\mathbf{x}_t = (ne_t, y_t)'$  and estimate the following reduced form VAR:

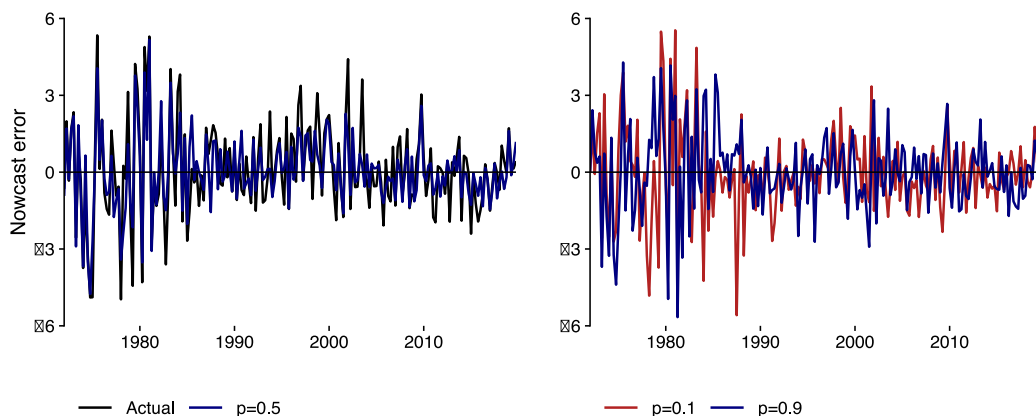
$$\mathbf{x}_t = \sum_{l=1}^P \mathbf{A}_l \mathbf{x}_{t-l} + \mathbf{B} \mathbf{d}_t + \mathbf{u}_t, \quad \mathbf{u}_t \sim \mathcal{N}(\mathbf{0}, \Sigma). \quad (5)$$

$\mathbf{A}_l$  is the dynamic coefficient matrix for lag  $l = 1, \dots, P$ ,  $\mathbf{B}$  comprises the parameters associated with deterministic components  $\mathbf{d}_t = (1, t, t^2)'$  and  $\mathbf{u}_t = (u_{ne,t}, u_{y,t})'$  are reduced form errors following a multivariate Gaussian distribution with zero mean and covariance matrix  $\Sigma$ . The structural *belief* (b) and *non-belief* (nb) shocks comprise the vector  $\epsilon_t = (\epsilon_{nb,t}, \epsilon_{b,t})'$ . They are uncorrelated and their variance is normalized such that  $\epsilon_t \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ . To achieve identification, we need to pin down the elements of the matrix  $\mathbf{A}_0$ , which maps structural shocks to reduced form innovations  $\mathbf{u}_t = \mathbf{A}_0 \epsilon_t$ . This implies that  $\Sigma = \mathbf{A}_0 \mathbf{A}_0'$  and presents a well-known identification problem.

Further restrictions are necessary to give economic meaning to our structural shocks, which we introduce as follows. The belief shock causes a negative co-movement between the nowcast error and output, because a negative (positive) nowcast error means that the consensus survey expectation of growth is higher (lower) than current output growth in real-time. Hence, agents are overly optimistic which causes an outward (inward) shift of the demand curve and subsequently output increases (decreases). The non-belief shock differs from the belief shock insofar as it causes co-movement between the nowcast error and output. This gives rise to sign restrictions on the elements of  $\mathbf{A}_0$ :

$$\begin{bmatrix} u_{ne,t} \\ u_{y,t} \end{bmatrix} = \begin{bmatrix} + & + \\ + & - \end{bmatrix} \begin{bmatrix} \epsilon_{nb,t} \\ \epsilon_{b,t} \end{bmatrix},$$

which set-identifies the structural shocks. For estimation and structural inference we rely on Bayesian methods. We use variants of Bayesian VARs with weakly informative Minnesota-type priors, and alternative implementations are noted when applicable. We use a Gibbs sampler and draw 15,000 times from the posterior distribution while discarding the first 5000 draws as burn-in. For the sign restrictions, we rely on the algorithm proposed by Rubio-Ramirez, Waggoner, and Zha (2010).<sup>5</sup>



**FIGURE 2** | NEs computed with actual realizations of real GDP growth and the indicated selected GaR probabilities.

### 3.2 | Revisiting the Original Belief Shocks

We replicate the main findings of Enders, Kleemann, and Müller (2021) and corroborate their results along several dimensions. First, we consider two sampling periods. Specifically, we apply our framework to the original sample which ranges from 1968Q4 to 2014Q4, but also use an extended version which runs through 2019Q4.<sup>6</sup> Second, we consider different specifications of the VARs. On the one hand, we vary the number of lags  $P \in \{2, 4, 6\}$ . On the other hand, we consider both a flat prior (for the “exact” replication of Enders, Kleemann, and Müller 2021) and the weakly informative prior implementation mentioned above.<sup>7</sup>

Impulse response functions are computed for a horizon of 20 quarters and shown in Figure 3 for the respective subsamples in panels (a) and (b). Non-belief and belief shocks are in the respective left and right columns; the upper rows refer to the NE, while the bottom row depicts the response of output growth. The recovered dynamic responses are virtually identical to those in Enders, Kleemann, and Müller (2021, compare their Figure 5 to our Figure 3a), for all considered model specifications. The imposed sign restrictions result in positive co-movement between the NE and output in response to the non-belief shock, and negative co-movement for the belief shock on impact. Note that propagation dynamics are left unconstrained by this identification scheme. The response of the NE is short-lived for both shocks and insignificant for all horizons apart from the impact. The output response by contrast is fairly persistent. Depending on the level of statistical significance, it turns indistinguishable from zero after about 12–15 quarters. A negative NE corresponds to excessively optimistic beliefs. In this case, survey expectations exceed actual real-time output growth, because agents have an optimistic outlook. This optimism about current output growth causes actual output growth to increase.

Varying the number of lags and introducing modest shrinkage via a Minnesota-type prior has minor implications for the persistence of our posterior median estimates. But these differences are statistically insignificant. Comparing the extended sample in Figure 3b to the original period in Figure 3a, as in Enders, Kleemann, and Müller (2021), indicates that this extension has no discernible consequences for the results. To sum up this narrow replication study, we find that the original results are robust to alternative

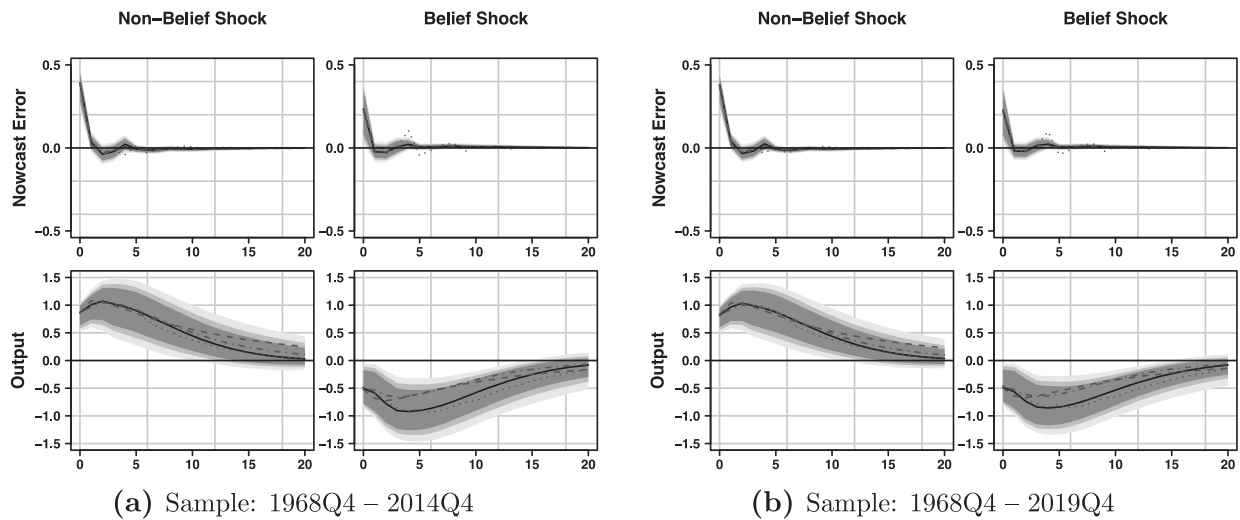
specifications and implementations of the baseline econometric framework.

### 3.3 | Nowcast Errors About Growth-at-Risk

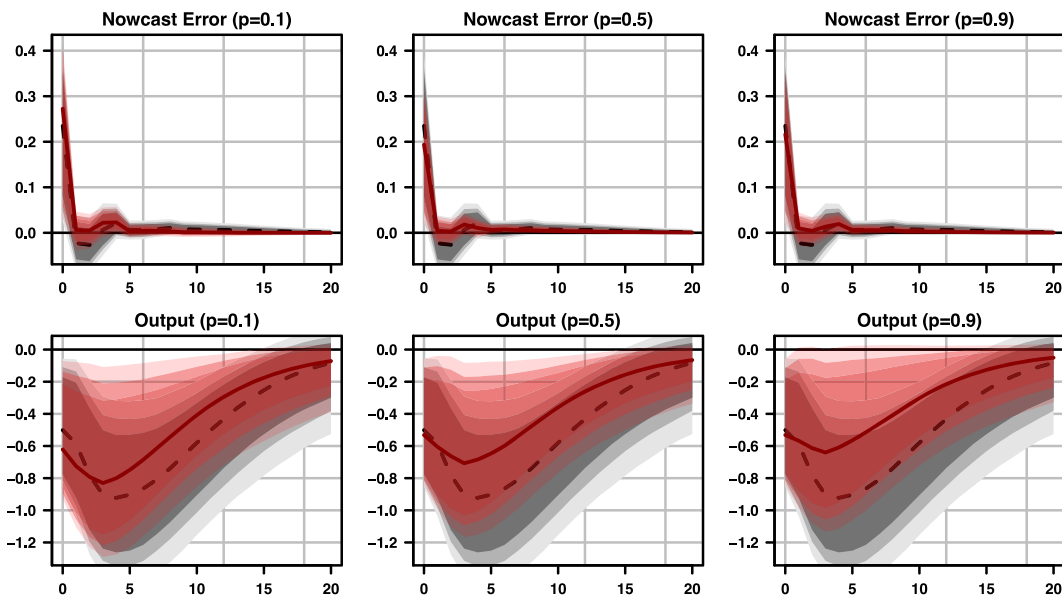
In this subsection, we investigate whether nowcast errors about GaR can be used as an alternative reduced form measure to identify belief shocks. These nowcast errors can be interpreted as misjudgments of macroeconomic risk in real-time. The resulting belief shocks potentially differ from the mean-based ones discussed above (e.g., through fundamental macroeconomic or financial shocks asymmetrically affecting the objective versus subjective nowcast distribution of output). The empirical findings of this section thus also relate to those of Loria, Matthes, and Zhang (2024), who measure quantile-specific responses of output growth to several fundamental (mean-based or externally identified) macroeconomic and financial shocks. By contrast, we use the NEs that originate in the quantiles to pin down shocks in a mean-based linear VAR framework.

Given the robustness of the original results that we established in the preceding section, we limit ourselves to using the full sample ranging from 1968Q4 to 2019Q4, and use a lag length of  $P = 4$  in the Bayesian VAR. As pointed out earlier, we now use the full predictive distribution from the SPF and rely on the NEs as defined in Equation (4) to capture these aspects.<sup>8</sup> The main results from this exercise are presented in Figure 4, in form of the red-colored impulse response functions. The columns now report the dynamic effects of belief shocks for three different quantiles  $p \in \{0.1, 0.5, 0.9\}$ . Since the nowcast errors are a reduced form measure, we identify belief shocks via sign restrictions using measures of downside, median, and upside risk. For ease of reference, the effects measured in the narrow replication are shown in shades of gray, and the dashed line marks the posterior median estimate from before.

Peak response effects occur slightly earlier and they are also a bit subdued when compared to the original framework. But the credible sets are inflated when considering the quantiles, and the structural VAR yields similar (and statistically indistinguishable) dynamic effects. Overall we thus conclude that considering belief shocks in different parts of the nowcast distribution does not cause any noteworthy differences when the focus is on explaining



**FIGURE 3** | Impulse response functions to a non-belief and belief shock. Estimation of bivariate Bayesian VAR(4) identified with sign restrictions. Solid black lines are posterior median responses, while gray shaded areas depict the 68%/80%/90% credible sets of the flat prior version of this model. NEs and output are measured in deviations from trend in percentage points. Dashed and dotted lines refer to posterior medians of alternative specifications.



**FIGURE 4** | Impulse response functions to belief shocks using GaR-NEs. Estimation of bivariate Bayesian VAR(4) identified with sign restrictions. Black/red lines denote the posterior median responses while gray/red shaded areas depict the 68%/80%/90% credible sets. Nowcast error and output are measured in deviations from trend in percentage points. Responses in black/gray denote the original model of Enders, Kleemann, and Müller (2021), while the responses in red denote the impulse responses to the belief shocks arising from the tails of nowcast distributions.

business cycle fluctuations. The mean-based original implementation is sufficient to induce the characteristic effects which are mostly homogeneous across GaR quantiles. This corroborates the results of Enders, Kleemann, and Müller (2021) in a wide sense. And this finding can, at least in part, be traced back to the notion that SPF nowcasts for output growth in our sampling period are mostly unimodal and symmetric.

But while differences between effects at different probabilities of GaR are statistically insignificant, some interesting heterogeneity still emerges. For instance, zooming into the belief shock using an upside risk NE at  $p = 0.9$  we find that the corresponding 90

percent credible set includes zero for all horizons apart from the impact. It is also worth mentioning that the median response is less clearly hump-shaped and flatter, particularly when contrasted with the one for downside risk at  $p = 0.1$ . This pattern appears monotonically when transitioning from the upper to the lower quantiles (using finer grids of GaR).

We again stress that none of these differences are significant in a statistical sense, but they point towards the notion that the overall effects of belief shocks are at least to some extent driven by misperceptions about adverse economic dynamics (in the lower tails). This is in line with the literature on GaR, which

has indeed almost exclusively focused on the lower tails of output growth. Given the linearity of the model, positive and negative shocks yield symmetric effects, so this also implies a somewhat stronger expansion in response to benign belief shocks about downside risk.

## 4 | Closing Remarks

In this paper, we replicate the study of belief shocks and their implications by Enders, Kleemann, and Müller (2021). Their results are robust in a narrow sense concerning data sourcing, econometric specification, and software implementation. In a wide sense, we also investigate whether belief shocks differ when using nowcast errors from the tails of the nowcast distribution of output growth. Our findings suggest that the originally proposed approach is sufficient to measure the overall effects of belief shocks on business cycle fluctuations adequately. Distinct patterns in dynamic responses arising from considering the full nowcast distribution are negligible for the most part.

### Acknowledgments

We thank Zeno Enders for useful suggestions and comments. Pfarrhofer gratefully acknowledges financial support by the Jubiläumsfonds of the Oesterreichische Nationalbank (OeNB, project 18765). Open Access funding enabled and organized by Projekt DEAL.

### Data Availability Statement

Data and code is available at [github.com/mboeck11/replication-belief-tailrisk](https://github.com/mboeck11/replication-belief-tailrisk). The data is also available at the online repository of the JAE, as linked in the text: <https://doi.org/10.15456/jae.2025006.1147228743>.

### Open Research Badges



This article has been awarded Open Data Badge for making publicly available the digitally-shareable data necessary to reproduce the reported results. Data is available at <https://doi.org/10.15456/jae.2025006.1147228743>.

### Endnotes

<sup>1</sup> The literature investigating belief distortions/wedges or disagreement usually focuses on mean outcomes and the cross-sectional dispersion of beliefs or expectations (Lahiri and Sheng 2010; Dovern, and Enders 2023; Adam, Matveev, and Nagel 2021; Bianchi, Ludvigson, and Ma 2022; Bhandari, Borovička, and Ho 2022; Boeck 2023; Born et al., 2012; Pei 2024).

<sup>2</sup> The  $CRPS(f(\bullet|\theta), w) = \int_{-\infty}^{\infty} (F(z|\theta) - \mathbb{I}(w \leq z))^2 dz$ , where  $f(\bullet|\theta)$  refers to the probability density function of some (predictive) distribution with parameter vector  $\theta$ ,  $F(z|\theta) = \int_{-\infty}^z f(w|\theta) dw$  is the corresponding cumulative distribution function, and  $w$  is the realized value. In our application,  $f(\bullet|\theta)$  is Gaussian, and  $\theta$  comprises a known mean and unknown variance. See Gneiting and Raftery (2007) and Gneiting and Ranjan (2011) for details and discussions of the favorable properties of the CRPS as a scoring rule.

<sup>3</sup> The dataset is described in more detail in Appendix B. Our results are robust to relying on the most recent data vintage of the FRED-QD database. In this case, we explicitly exploit the most recently available information to measure the quantiles of the first release vintage of GDP for each period.

<sup>4</sup> For some quantiles the NEs exhibit a modest amount of persistence, which we eliminate with this procedure. We also estimated the VAR model in Section 3 without purging the NEs from predictable components. In this case, the main results for output are very similar, but there is some persistence in the responses of the nowcast errors.

<sup>5</sup> This algorithm is based on a QR decomposition to draw uniformly from the space of orthonormal matrices to construct  $A_0$  that satisfies the sign restriction. The original implementation of Enders, Kleemann, and Müller (2021) uses Givens rotation matrices to draw uniformly from the space of orthonormal matrices. Hence, these approaches draw from the same space of orthonormal matrices to construct  $A_0$ . An alternative approach that introduces identification information explicitly is due to Baumeister and Hamilton (2015).

<sup>6</sup> We thus estimate the model excluding the post-Covid period. Extending the sample further but downweighting/dropping the pandemic observations (see, e.g., Lenza and Primiceri 2022) yields qualitatively similar results.

<sup>7</sup> Note that Enders, Kleemann, and Müller (2021) carried out their computations in MATLAB and relied on a frequentist approach to estimation and inference. By contrast, we have independently compiled the dataset, and use a Bayesian VAR implemented in R. This provides robustness from a data, econometric, and software perspective.

<sup>8</sup> The NEs are based on real-time predictions of GaR. In a robustness check, we also estimate the true quantile processes based on full-sample information for the final vintage data of the predictors. The results are robust to this choice.

### References

- Adam, K., D. Matveev, and S. Nagel. 2021. “Do Survey Expectations of Stock Returns Reflect Risk Adjustments?” *Journal of Monetary Economics* 117: 723–740.
- Adams, P. A., T. Adrian, N. Boyarchenko, and D. Giannone. 2021. “Forecasting Macroeconomic Risks.” *International Journal of Forecasting* 37, no. 3: 1173–1191.
- Adrian, T., N. Boyarchenko, and D. Giannone. 2019. “Vulnerable Growth.” *American Economic Review* 109, no. 4: 1263–1289.
- Adrian, T., F. Grinberg, N. Liang, S. Malik, and J. Yu. 2022. “The Term Structure of Growth-At-Risk.” *American Economic Journal: Macroeconomics* 14, no. 3: 283–323.
- Baumeister, C., and J. D. Hamilton. 2015. “Sign Restrictions, Structural Vector Autoregressions, and Useful Prior Information.” *Econometrica* 83, no. 5: 1963–1999.
- Beaudry, P., and F. Portier. 2006. “Stock Prices, News, and Economic Fluctuations.” *American Economic Review* 96, no. 4: 1293–1307.
- Beaudry, P., and F. Portier. 2014. “News-Driven Business Cycles: Insights and Challenges.” *Journal of Economic Literature* 52, no. 4: 993–1074.
- Bhandari, A., J. Borovička, and P. Ho. 2022. “Survey Data and Subjective Beliefs in Business Cycle Models.” SSRN 2763942.
- Bianchi, F., S. C. Ludvigson, and S. Ma. 2022. “Belief Distortions and Macroeconomic Fluctuations.” *American Economic Review* 112, no. 7: 2269–2315.
- Boeck, M. 2023. “Belief Distortions in Risk Premia.” SSRN 4010599.
- Born, B., J. Dovern, and Z. Enders. 2023. “Expectation Dispersion, Uncertainty, and the Reaction to News.” *European Economic Review* 154: 104440.
- Clark, T. E., F. Huber, G. Koop, M. Marcellino, and M. Pfarrhofer. 2023. “Tail Forecasting With Multivariate Bayesian Additive Regression Trees.” *International Economic Review* 64, no. 3: 979–1022.

- Clark, T. E., F. Huber, G. Koop, M. Marcellino, and M. Pfarrhofer. 2024. "Investigating Growth-At-Risk Using a Multicountry Nonparametric Quantile Factor Model." *Journal of Business & Economic Statistics* 42, no. 4: 1302–1317.
- Dovern, J., U. Fritsche, and J. Slacalek. 2012. "Disagreement Among Forecasters in G7 Countries." *Review of Economics and Statistics* 94, no. 4: 1081–1096.
- Enders, Z., M. Kleemann, and G. J. Müller. 2021. "Growth Expectations, Undue Optimism, and Short-Run Fluctuations." *Review of Economics and Statistics* 103, no. 5: 905–921.
- Frühwirth-Schnatter, S. 2006. *Finite Mixture and Markov Switching Models*. New York: Springer.
- Gneiting, T., and A. E. Raftery. 2007. "Strictly Proper Scoring Rules, Prediction, and Estimation." *Journal of the American Statistical Association* 102, no. 477: 359–378.
- Gneiting, T., and R. Ranjan. 2011. "Comparing Density Forecasts Using Threshold-and Quantile-Weighted Scoring Rules." *Journal of Business & Economic Statistics* 29, no. 3: 411–422.
- Krüger, F., and I. Nolte. 2016. "Disagreement Versus Uncertainty: Evidence From Distribution Forecasts." *Journal of Banking & Finance* 72: S172–S186.
- Lahiri, K., and X. Sheng. 2010. "Measuring Forecast Uncertainty by Disagreement: The Missing Link." *Journal of Applied Econometrics* 25, no. 4: 514–538.
- Lenza, M., and G. E. Primiceri. 2022. "How to Estimate a Vector Autoregression After March 2020." *Journal of Applied Econometrics* 37, no. 4: 688–699.
- Loria, F., C. Matthes, and D. Zhang. 2024. "Assessing Macroeconomic Tail Risk." *Economic Journal* 135, no. 655: 264–284.
- Pei, G. 2024. "Pessimism, Disagreement, and Economic Fluctuations." *Journal of the European Economic Association* 22, no. 3: 1177–1227.
- Pfarrhofer, M. 2022. "Modeling Tail Risks of Inflation Using Unobserved Component Quantile Regressions." *Journal of Economic Dynamics and Control* 143: 104493.
- Rubio-Ramirez, J. F., D. F. Waggoner, and T. Zha. 2010. "Structural Vector Autoregressions: Theory of Identification and Algorithms for Inference." *The Review of Economic Studies* 77, no. 2: 665–696.
- Stock, J. H., and M. W. Watson. 2002. "Macroeconomic Forecasting Using Diffusion Indexes." *Journal of Business & Economic Statistics* 20, no. 2: 147–162.

## Supporting Information

Additional supporting information can be found online in the Supporting Information section.