WIFO ■ WORKING PAPERS 710/2025

Enhancing Macroeconomic Forecasts with Uncertainty-Informed Intervals

Christian Glocker
Serguei Kaniovski

Enhancing Macroeconomic Forecasts with Uncertainty-Informed Intervals

Christian Glocker, Serguei Kaniovski

Research assistance: Astrid Czaloun

WIFO Working Papers 710/2025
September 2025

Abstract

We propose a methodology for constructing confidence intervals for macroeconomic forecasts that directly incorporate quantitative measures of uncertainty – such as survey-based indicators, stock market volatility, and policy uncertainty. By allowing the width of confidence intervals to vary systematically with prevailing uncertainty conditions, this approach yields more informative and context-sensitive intervals than traditional, static methods relying solely on past forecast errors. An empirical application using Austrian data demonstrates that uncertainty measures significantly explain the variation in forecast errors, underscoring the value of integrating these indicators for improved communication and analytical robustness of economic projections.

E-Mail: christian.glocker@wifo.ac.at, serguei.kaniovski@wifo.ac.at

2025/1/W/0

Media owner (publisher), producer: Austrian Institute of Economic Research
1030 Vienna, Arsenal, Objekt 20 | Tel. (43 1) 798 26 01 0 | https://www.wifo.ac.at
Place of publishing and production: Vienna

ENHANCING MACROECONOMIC FORECASTS WITH UNCERTAINTY-INFORMED INTERVALS

CHRISTIAN GLOCKER AND SERGUEI KANIOVSKI

Abstract. We propose a methodology for constructing confidence intervals for macroeconomic forecasts that directly incorporate quantitative measures of *uncertainty* – such as survey-based indicators, stock market volatility, and policy uncertainty. By allowing the width of confidence intervals to vary systematically with prevailing uncertainty conditions, this approach yields more informative and context-sensitive intervals than traditional, static methods relying solely on past forecast errors. An empirical application using Austrian data demonstrates that uncertainty measures significantly explain the variation in forecast errors, underscoring the value of integrating these indicators for improved communication and analytical robustness of economic projections.

JEL codes: C32; C53; C40; E37;
Key words: Confidence intervals; Forecast errors; Uncertainty; SUR

1. Introduction

We propose an approach for integrating quantitative uncertainty measures – such as survey-based direct uncertainty indicators, stock market volatility, and policy uncertainty metrics – into the construction of confidence intervals for macroeconomic forecasts. In this framework, the width of the confidence intervals changes with the prevailing levels of uncertainty indicators, resulting in wider intervals during periods of heightened uncertainty and narrower intervals when uncertainty is low.

The motivation for this methodology is rooted in the persistent challenges that characterize economic forecasting for both researchers and practitioners. Forecast accuracy is frequently compromised by a range of factors, such as

---

This version: August 2025.

C. Glocker: Austrian Institute of Economic Research (Wifo), Arsenal Objekt 20, 1030 Vienna, Austria. Phone: +43 (0) 1 789 26 01-303, E-mail: *christian.glocker@wifo.ac.at*
S. Kaniovski: Austrian Institute of Economic Research (Wifo), Arsenal Objekt 20, 1030 Vienna, Austria. Phone: +43 (0) 1 789 26 01-231, E-mail: *serguei.kaniovski@wifo.ac.at*

The authors would like to thank Maximilian Böck, Gabriel Felbermayr, Werner Hölzl, Philipp Wegmüller, Stefan Schiman-Vukan, and Thomas Url for valuable comments and helpful discussions. Excellent research assistance by Astrid Czaloun and Alexandros Charos is gratefully acknowledged.

conflicting signals from various survey instruments, unforeseen and severe exogenous shocks, the endogenous policy responses that forecasts themselves may provoke, as well as data revisions and broader data quality concerns. These complexities often result in substantial forecast errors, which tend to attract considerable media attention, particularly during periods of acute economic crisis – precisely when accurate forecasts are most critical and errors are most consequential.

In light of these challenges, many forecasting institutions have adopted the practice of attaching confidence intervals to their economic projections.[1] Typically, the construction of such intervals relies exclusively on the distribution of past forecast errors, which are then appended, often mechanically, to the current forecast to produce fan charts. These charts provide probabilities associated with different ranges of potential outcomes for the forecast variable. The immediate objective of this approach is to contextualize the current forecast within a probabilistic range derived from historical forecast errors (Razi and Loke, 2017). Ultimately, the purpose of these intervals is to convey the inherent uncertainty faced by forecasters when generating their predictions.

Despite being intended as a tool for quantifying forecast uncertainty, conventional confidence intervals neglect to incorporate any quantitative measure of uncertainty as perceived by economic agents, such as households or firms. This omission is particularly surprising for at least two reasons. First, forecasters indeed utilize household and business survey data in the process of generating their predictions. Second, in many countries, these surveys feature specific questions regarding uncertainty (see Glocker and Hölzl, 2022, for instance). For example, since 1996, the Austrian Institute of Economic Research (WIFO) has included in its business survey a question explicitly addressing firms' perceptions of (subjective) uncertainty:

Die zukünftige ENTWICKLUNG unserer Geschäftslage ist:

- *leicht abschätzbar*
- *einigermaßen leicht abschätzbar*
- *einigermaßen schwer abschätzbar*
- *schwer abschätzbar*

[1]See, for example, the Bank of England, the Institut national de la statistique et des études économiques du Grand-Duché de Luxembourg (STATEC) as outlined in Kaniovski (2019), the Czech National Bank, the Central Bank of Mexico, the Bank of Russia, the Bank of Spain, or the Central Bank of Chile, among others.

> Translation: *The future DEVELOPMENT of our business situation is: easy to predict / moderately easy to predict / moderately difficult to predict / difficult to predict.*

Intuitively, low uncertainty can arise both during expansions and recessions: in the former, firms predominantly anticipate improvement; in the latter, they largely expect deterioration. Business survey indicators – especially assessments of the current situation and short-term expectations – are central inputs to economic forecasting. Their informational content, however, hinges critically on the uncertainty embedded in respondents' evaluations. Elevated uncertainty diminishes the informational content of these indicators and enlarges forecast errors, whereas low uncertainty enhances their informativeness and improves forecast accuracy. Consequently, uncertainty measures provide a systematic means to assess the quality of information conveyed by situation and expectation indicators. Given the pivotal role of these indicators in forecasting, incorporating uncertainty measures can sharpen the delineation of forecast corridors and more credibly capture the range of potential economic outcomes.

In this context, we propose that the width of the confidence interval accompanying an economic forecast should be systematically adjusted to reflect the prevailing level of subjective uncertainty as reported by households and firms. To this end, we introduce a methodology for constructing confidence intervals whose width varies with a quantitative measure of perceived uncertainty. This approach allows the forecast to be more appropriately contextualized by incorporating information on the uncertainty environment at the time the forecast is produced.

Literature. Contributions in this respect have been made by, among others, Kannan and Elekdag (2009) who develop a procedure for incorporating market-based information into the construction of fan charts, and Hsieh (2016) who too proposes the use of external indicators in adjusting the width of forecast-related confidence intervals. Further related contributions include Biswas (2019); Chourou et al. (2021); Morikawa (2023) all of which demonstrate that forecast errors significantly vary with measures of uncertainty. Furthermore, Wang et al. (2023) present an interval forecasting methodology in the context of wind power, where the width of prediction intervals adjusts dynamically to time-varying correlations. Other contributions are, for instance, Turner (2017) who proposes a method to construct confidence intervals which are parameterized on the basis of the historical forecasting track record, but distinguish between a "safe" and "downturn-risk" regime.

Our approach extends prior contributions along several dimensions. First, we introduce a framework for constructing uncertainty-dependent confidence intervals that explicitly account for the forecast horizon, allowing the width of the intervals to vary with the forecast horizon. Second, our methodology incorporates the timing of forecast production, thus capturing the sensitivity of forecast uncertainty to the specific point in the year when forecasts are generated (given that many institutions, for instance the IMF, the OECD, etc., produce at least two forecasts a year). Third, we systematically evaluate a range of uncertainty measures and assess their informational value for explaining the historical variation in forecast errors. This allows to subsequently determine their suitability as indicators for modulating the confidence interval widths. Fourth, we investigate potential non-linearities in the relationship between uncertainty and forecast errors, and implement an estimation approach that accommodates correlations among forecast errors across successive forecast vintages and forecast horizons.

To empirically illustrate our approach, we focus on the Austrian economy, leveraging forecasts published by the Austrian Institute of Economic Research (Wifo). This context is particularly apt for two reasons. First, since 1996, the institute has conducted (monthly) business surveys containing an explicit question on uncertainty, offering a long time series of a direct subjective uncertainty measure (see also Glocker and Hölzl, 2022). Second, over the same period, four times a year, the institute has produced short-term economic forecasts for both the current and following year (for the Austrian economy), yielding a long time series of forecast errors. Historical forecast errors for Austria have been comprehensively evaluated in Fortin et al. (2020). This unique confluence of data enables us to rigorously investigate the statistical relationship between forecast errors and measures of uncertainty.

The remainder of the text is organized as follows. Section 2 introduces the methodology for integrating uncertainty measures into the construction of forecast confidence intervals and presents empirical results. Section 3 extends the baseline approach by addressing non-linearities, evaluating alternative uncertainty measures, and relaxing parameter constraints. Section 4 discusses further extensions (some of which are implemented in the Appendix) and Section 5 provides a use case for the proposed methodology. Section 6 concludes the paper.

2. Methodological framework

Let $y_{t+h|t}$ denote the forecast of variable y for period $t+h$ made at time t, and let y_{t+h} represent the realization of y in period $t+h$. The forecast error $\varepsilon_{t+h|t}$ for horizon h at time t ($t = 1, \ldots, T$) is then defined as:

$$\varepsilon_{t+h|t} = y_{t+h|t} - y_{t+h} \tag{1}$$

and the variance of the forecast error at horizon h over the sample period is calculated by:

$$\sigma_h^2 = \frac{1}{T} \sum_{t=1}^{T} \varepsilon_{t+h|t}^2 \tag{2}$$

where we implicitly assume that $\frac{1}{T}\sum_{t=1}^{T} \varepsilon_{t+h|t} = 0$, that is, the forecast errors are mean zero.

The prevailing practice among forecasting institutions is to utilize the historical forecast error variance at horizon h to compute the corresponding standard deviation (σ_h), which is then used to construct a confidence interval for the forecast. Specifically, the $(1-\alpha)\%$ confidence interval (CI) for the forecast $y_{t+h|t}$ is given by:

$$\mathrm{CI}_{\mathrm{L}}^{\mathrm{U}}(h) = y_{t+h|t} \pm \kappa\, \sigma_h \tag{3}$$

where $\kappa \geq 0$ is a critical value corresponding to the desired confidence level. As indicated in equation (3), this approach determines the width of the confidence interval solely based on the distribution of historical forecast errors, disregarding potentially relevant external information available at time t when the forecast is produced, that may influence forecast uncertainty and subsequent forecast error magnitudes.

The central innovation of this study is to introduce a quantifiable measure of uncertainty – denoted ζ_t – to incorporate this measure into the construction of the confidence interval. We therefore extend equation (3) as follows:

$$\mathrm{CI}_{\mathrm{L}}^{\mathrm{U}}(h) = y_{t+h|t} \pm \kappa\, \sigma_h(\zeta_t) \tag{4}$$

where $\sigma_h(\zeta_t)$ explicitly denotes that the width of the confidence interval is a function of the observed level of uncertainty at time t when the forecast is produced.

2.1. **Squared forecast errors and subjective uncertainty.** To establish a relationship between the forecast error variance σ_h^2 and the uncertainty indicator ζ_t, we proceed in two steps. The key point here is that, rather than using the average squared forecast error ($\sigma_{h,m}^2$) as of equation (2), we employ the time-specific squared forecast error ($\varepsilon_{t+h|t,m}^2$). In this respect, we first substitute σ_h^2 for the squared forecast error, $\varepsilon_{t+h|t}^2$, recognizing that $\varepsilon_{t+h|t}^2$ may be systematically related to the level of uncertainty ζ_t. To capture this relationship empirically, we posit (as a starting point) a linear specification:

$$\varepsilon_{t+h|t}^2 = \alpha_h + \beta(\zeta_t - \bar{\zeta}) + u_t \tag{5}$$

where $\bar{\zeta}$ denotes the sample mean of the uncertainty indicator and u_t is a mean-zero error term. The subscript on the parameter α_h indicates that its value depends on the forecast horizon h considered in $\varepsilon_{t+h|t}^2$. Equation (5) can be directly associated to the variance of the historical forecast errors put forth in equation (2). To see this, we use equation (5) and operate with $\frac{1}{T}\sum_{t=1}^{T}$ on both sides to obtain:

$$\frac{1}{T}\sum_{t=1}^{T} \varepsilon_{t+h|t}^2 = \alpha_h = \sigma_h^2 \tag{6}$$

where the last equality follows from equation (2). This allows to rewrite equation (5) as follows:

$$\varepsilon_{t+h|t}^2 = \sigma_h^2 + \beta(\zeta_t - \bar{\zeta}) + u_t \tag{7}$$

Equation (7) defines a standard linear regression model that can be estimated using standard techniques. Letting $\hat{\sigma}_h^2$ and $\hat{\beta}$ denote the point estimates of the regression coefficients (constant term and the slope coefficient), the predicted value for the squared forecast error is:

$$\hat{\varepsilon}_{t+h|t}^2 = \hat{\sigma}_h^2 + \hat{\beta}(\zeta_t - \bar{\zeta}) \tag{8}$$

The predicted values $\hat{\varepsilon}^2_{t+h|t}$ can be readily used in equation (3). In particular, we substitute the mean of the squared forecast errors ($\frac{1}{T}\sum_{t=1}^{T}\varepsilon^2_{t+h|t}$) with the predicted values of the squared forecast error at a particular point in time t ($\hat{\varepsilon}^2_{t+h|t}$) as of equation (8). This adds a time-varying element to the width of the confidence interval and equation (3) is now given by:

$$\mathrm{CI}_{\mathrm{L}}^{\mathrm{U}}(h) = y_{t+h|t} \pm \kappa\sqrt{\hat{\sigma}^2_h + \hat{\beta}(\zeta_t - \bar{\zeta})} \tag{9}$$

where $\hat{\sigma}^2_h \geq 0$ and $\hat{\beta} \geq 0$ are expected. The latter reflects the positive relationship between higher uncertainty and increased forecast errors. As equation (9) makes explicit, when the observed uncertainty ζ_t exceeds its historical mean $\bar{\zeta}$, the confidence interval widens, reflecting greater ex ante forecast uncertainty. Conversely, when ζ_t falls below $\bar{\zeta}$, the confidence interval narrows accordingly.

Most importantly, this approach yields confidence intervals whose widths vary deterministically over time t with realized levels of uncertainty, thus providing a more context-sensitive and informative characterization of forecast uncertainty than is possible using traditional, static methods relying solely on the variance of historical forecast errors (σ^2_h).

2.2. **Practical implementation.** In the following, we outline a practical approach for the implementation of the proposed methodology, focusing in particular on equation (7), which forms the basis for the construction of confidence intervals as detailed in equation (9). We illustrate our approach using the WIFO forecasting process as an example. Given that multiple forecasts are produced at distinct points throughout each year, we start to demonstrate this procedure in the most parsimonious, that is, simplest, way possible. This serves as the baseline model, providing a clear and tractable starting point. In subsequent sections, this specification is progressively extended.

Each year, WIFO releases four forecasts (in March, June, September, and December) for both the current year and the subsequent year. In principle, this structure allows us to estimate equation (7) separately for each forecast vintage, resulting in up to four potentially different estimates for the parameter β for each forecast horizon h. However, in order to obtain a single, coherent estimate for β across all forecast vintages, we extend the model to explicitly account for the month m in which the forecast is made.

Let us redefine the forecast error as:

$$\varepsilon_{t+h|t,m} = y_{t+h|t,m} - y_{t+h} \tag{10}$$

where $m = \{M, J, S, D\}$ indexes the forecast vintage (March, June, September, December) and t denotes annual data. Here, (t, m) signifies the instance (month m of year t) in which a forecast for the year $t + h$ is produced. In the context of WIFO forecasts, we consider two horizons: $h = 0$ (current year) and $h = 1$ (one year ahead) and we consider the first (official) data release for y_{t+h}.[2]

We now generalize equation (7) to incorporate the month m of the forecast vintage:

$$\varepsilon^2_{t+h|t,m} = \sigma^2_{h,m} + \beta(\zeta_{t,m} - \bar{\zeta}) + u_{t,m} \tag{11}$$

where $\sigma^2_{h,m}$ denotes the unconditional forecast error variance for horizon h and month m, and $u_{t,m}$ is a mean-zero error term. The unconditional forecast error variance is hence allowed to vary over the forecast horizon h and the forecast vintage m and we expect that, for instance, $\sigma^2_{0,D} \leq \sigma^2_{0,M}$, that is, the forecast error variance for the current year's forecast ($h = 0$) is smaller in December ($m = D$) than in March ($m = M$).[3] The uncertainty measure $\zeta_{t,m}$ employed in this analysis is constructed from the WIFO monthly business survey,[4] as described earlier. This measure directly captures subjective uncertainty of firms. We extend it to four separate time series at an annual frequency for each forecast vintage $m = \{M, J, S, D\}$; for instance, $\zeta_{t,M}$ denotes the uncertainty prevailing in March ($m = M$) of year t, with this value being the average over the current and the two preceding months. The formulation proposed by equation (11) allows the model to capture heterogeneity in the forecast error variance that varies both across forecast horizons and forecast vintages, and links it systematically to an uncertainty measure. For parsimony, we restrict β to be invariant with respect to the forecast vintage and forecast horizon (subsequently relaxed).

[2]Consider the Appendix for a discussion on the benefit of using first release data.
[3]More generally, we expect that $\sigma^2_{h,m'} \leq \sigma^2_{h,m}$ for $m' > m$ and any h.
[4]Consider the appendix for further information on the direct subjective business uncertainty measure.

To jointly estimate the parameters across all forecast vintages (m), we cast equation (11) in a multivariate system that encapsulates the four forecasts released each year. For each forecast horizon (h), we consider the following four-dimensional system:

$$\boldsymbol{\varepsilon}_{h,t} = \boldsymbol{\sigma}_h + \beta \boldsymbol{\zeta}_t + \boldsymbol{u}_{h,t}. \tag{12}$$

where the vectors are given by:

$$\boldsymbol{\varepsilon}_{h,t} = \begin{bmatrix} \varepsilon^2_{t+h|t,M} \\ \vdots \\ \varepsilon^2_{t+h|t,D} \end{bmatrix}, \quad \boldsymbol{\sigma}_h = \begin{bmatrix} \sigma^2_{h,M} \\ \vdots \\ \sigma^2_{h,D} \end{bmatrix}, \quad \boldsymbol{\zeta}_t = \begin{bmatrix} \zeta_{t,M} - \bar{\zeta} \\ \vdots \\ \zeta_{t,D} - \bar{\zeta} \end{bmatrix}$$

and $\boldsymbol{u}_{h,t}$ is assumed to follow a multivariate normal distribution with zero mean and (possibly full) covariance matrix $\boldsymbol{\Sigma}_h$. The dimension of the system is defined by the number of forecasts produced per year (that is, the number of forecast vintages, which is four in our case). Since equation (12) applies to any forecast horizon h, we consider the stacked system for both $h = 0$ and $h = 1$:

$$\begin{bmatrix} \boldsymbol{\varepsilon}_{0,t} \\ \boldsymbol{\varepsilon}_{1,t} \end{bmatrix} = \begin{bmatrix} \boldsymbol{\sigma}_0 \\ \boldsymbol{\sigma}_1 \end{bmatrix} + \beta \left(\mathbf{1}_2 \otimes \boldsymbol{\zeta}_t \right) + \begin{bmatrix} \boldsymbol{u}_{0,t} \\ \boldsymbol{u}_{1,t} \end{bmatrix} \tag{13}$$

where $\mathbf{1}_2$ is a two-dimensional unit vector. The joint error term is assumed to satisfy

$$\begin{bmatrix} \boldsymbol{u}_{0,t} \\ \boldsymbol{u}_{1,t} \end{bmatrix} \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}), \quad \text{where} \quad \boldsymbol{\Sigma} = \begin{bmatrix} \boldsymbol{\Sigma}_0 & \boldsymbol{\Sigma}_{01} \\ \boldsymbol{\Sigma}_{10} & \boldsymbol{\Sigma}_1 \end{bmatrix} \tag{14}$$

with $\boldsymbol{\Sigma}$ full and unrestricted.

A key advantage of the multivariate framework is its ability to easily impose parameter restrictions. For instance, it enables the slope coefficient β to be held constant across forecast horizons h and forecast vintages m, whereas in a purely univariate setting, this coefficient would vary across h and m giving rise to $\beta_{h,m}$. In the following, we detail the estimation procedure for equation (13) and highlight further advantages of the multivariate framework.

2.2.1. *Estimation results.* The system specified in equation (13) is estimated using the Seemingly Unrelated Regressions (SUR) estimator, applied to data spanning from 1996 to 2019.[5] This implies that we have 24 observations for the forecast errors $\varepsilon_{t+h|t,m}$ for each data vintage ($m = \{M, J, S, D\}$) and forecast horizon ($h = 0, 1$) resulting in a total of 192 ($= 24 \times 4 \times 2$) observations. The SUR estimator is well-suited for this multivariate set-up, as it captures the strong contemporaneous correlation among error terms $u_{t,m}$ across equations – reflecting both the temporal dependence across forecast vintages and across forecast horizons. This is a notable feature in our context, given that the forecast error associated with a particular month m in which a forecast is produced for horizon h is likely to be carried over, at least in part, into the subsequent month $m + 1$ in which the next forecast for horizon h is produced. The same applies to the cross-correlation in the error term over the forecast horizons $h = 0$ and $h = 1$. In fact, empirical correlations among the error terms $u_{t,m}$ often exceed 0.9, in case $\mathbf{\Sigma}$ is restricted to be diagonal, highlighting the efficiency and necessity of the SUR methodology for this multivariate setting to account for the strong contemporaneous correlations among the error terms. Finally, the multivariate framework enhances the efficiency of estimating β relative to a univariate setting, as in the former, all 192 observations are utilized simultaneously.

We begin by discussing the results pertaining to the unconditional forecast error variance, σ_h^2, as defined in equation (2), and its corresponding SUR-based estimates from equation (13), which are directly comparable. The former is reported in the first numerical column[6] of Table 1, while the latter appears in the second column. As shown, the estimated variance $\sigma_{h,m}^2$ increases with the forecast horizon h but decreases across forecast vintages m; that is, $\sigma_{1,m}^2 > \sigma_{0,m}^2$ for any m, and $\sigma_{h,m}^2 > \sigma_{h,m+1}^2$ for any h. Overall, forecast errors and their variances tend to increase with forecasting horizon. Forecasts for the current year become more accurate as the year progresses, as quarterly and monthly data gradually become available, while forecasts for the next year become more accurate as we approach the next year.

[5]We have excluded Covid-19 and post-Covid-19 observations from the estimation sample due to their excessive variance.

[6]In light of equation (6), which estimates the variance σ_h^2 as the sample mean of ε_t^2, a simple t-test on the mean estimate indicates the following: $\sigma_{0,M}^2$, $\sigma_{0,J}^2$, $\sigma_{0,S}^2$, and $\sigma_{0,D}^2$ are statistically different from zero at the 1% significance level; $\sigma_{1,S}^2$ and $\sigma_{1,D}^2$ are statistically significant at the 5% level; and $\sigma_{1,M}^2$ and $\sigma_{1,J}^2$ are statistically significant at the 10% level.

TABLE 1. Estimates for $\sigma^2_{h,m}$ and $\beta_{h,m}$ using firm uncertainty

	Sample	Baseline $\hat{\beta}_{h,m} = \hat{\beta}$	Restrict h $\hat{\beta}_{h,m} = \hat{\beta}_m$	Restrict m $\hat{\beta}_{h,m} = \hat{\beta}_h$	Unrestricted $\hat{\beta}_{h,m}$
$\hat{\sigma}^2_{0,M}$	0.62	0.69**	0.70**	0.62***	0.61***
$\hat{\sigma}^2_{0,J}$	0.32	0.41	0.43	0.33***	0.31***
$\hat{\sigma}^2_{0,S}$	0.21	0.31	0.32	0.23***	0.22***
$\hat{\sigma}^2_{0,D}$	0.15	0.27	0.21*	0.17***	0.17***
$\hat{\sigma}^2_{1,M}$	2.47	2.54**	2.55**	2.65**	2.74**
$\hat{\sigma}^2_{1,J}$	2.37	2.47**	2.49**	2.60***	2.64***
$\hat{\sigma}^2_{1,S}$	1.82	1.90**	1.90**	2.04***	1.99**
$\hat{\sigma}^2_{1,D}$	1.08	1.18***	1.12**	1.34***	1.17***
$\hat{\beta}_{0,M}$		0.048***	0.054***	0.011***	0.009
$\hat{\beta}_{0,J}$		0.048***	0.058***	0.011***	0.001
$\hat{\beta}_{0,S}$		0.048***	0.050***	0.011***	0.008**
$\hat{\beta}_{0,D}$		0.048***	0.026***	0.011***	0.009***
$\hat{\beta}_{1,M}$		0.048***	0.054***	0.107***	0.156***
$\hat{\beta}_{1,J}$		0.048***	0.058***	0.107***	0.126***
$\hat{\beta}_{1,S}$		0.048***	0.050***	0.107***	0.084***
$\hat{\beta}_{1,D}$		0.048***	0.026***	0.107***	0.044*

Test of coefficient ($\hat{\beta}$) restrictions:

H_0	Prob
$\hat{\beta}_{h,m} = \hat{\beta}_m \quad \forall\, h$	0.005***
$\hat{\beta}_{h,m} = \hat{\beta}_h \quad \forall\, m$	0.097*
$\hat{\beta}_{h,m} = \hat{\beta} \quad \forall\, h \wedge m$	0.030**

Note: The table presents estimates of the unconditional variance ($\hat{\sigma}^2_{h,m}$) of forecast errors and the slope-estimate ($\hat{\beta}$) of distinct models. The estimation uses a direct survey-based uncertainty measure collected among manufacturing firms. Statistical significance levels are: 10% (*), 5% (**), and 1% (***).

Table 1 presents the estimated values of β across a range of model specifications. In the baseline model (as specified in equation (13)), the estimate of β is 0.048 and is statistically significantly different from zero at the one percent level. This finding indicates that uncertainty explains part of the variation in the squared forecast error; equivalently, forecast errors exhibit a tendency to increase with rising levels of uncertainty. This result is broadly consistent with the findings of Morikawa (2023). The central question that follows is the extent to which this relationship alters the width of confidence intervals associated with the forecasts, which we turn to next.

2.2.2. *Implications for the confidence intervals.* The β estimate obtained for the baseline model (0.048) yields the following marginal effect of uncertainty on the width of the confidence interval:

$$\frac{\partial \mathrm{CI}_{\mathrm{L}}^{\mathrm{U}}(h,m)}{\partial \zeta_{t,m}} = \pm\, \kappa \, \frac{\hat{\beta}}{2\sqrt{\hat{\sigma}_{h,m}^{2} + \hat{\beta}(\zeta_{t,m} - \bar{\zeta})}} \tag{15}$$

which evaluates to ± 0.04 when $\kappa = 1$, $\zeta_t = \bar{\zeta}$, and $h = 0$ for forecasts produced in September ($m = S$). In practical terms, this result implies that the width of the confidence interval for the current year's GDP forecast (produced in September) increases by 0.08 percentage points when the uncertainty indicator rises by one index point.[7]

Figure 1 illustrates the application of uncertainty-dependent confidence intervals to a (hypothetical) GDP growth forecast. In this example, the baseline growth rate is set equal to the potential output growth rate of 0.5 percent p.a., reported across three periods: period $t-1$ (the most recent year with available data), and periods t and $t+1$, for which forecasts are produced (corresponding to forecast horizons $h = 0$ and $h = 1$). The forecasts are assumed to be produced in September ($m = S$). The figure compares the resulting confidence intervals across two distinct uncertainty scenarios: a low-uncertainty case, defined as the 25th percentile of the empirical distribution of the uncertainty indicator ($\zeta = 39.4$; black dashed lines, left panel), and a high-uncertainty case, corresponding to the 75th percentile ($\zeta = 48.2$; black dashed lines, right panel). As a benchmark, the figure also presents confidence intervals for the scenario in which the uncertainty measure equals its historical mean ($\zeta_t = \bar{\zeta}$ where $\bar{\zeta} = 44.3$), resulting in an interval based solely on the standard deviation of historical forecast errors (black solid lines).

The figure underscores three key aspects. First, confidence intervals widen as the forecast horizon h increases, a pattern that holds for any level of uncertainty ζ_t (by construction). Second, the confidence intervals are sensitive to changes in the level of uncertainty. Specifically, at low uncertainty, the intervals are narrower than the benchmark intervals based exclusively on historical forecast errors (see the left panel, where dashed lines are compared to solid black lines). Conversely, under conditions of elevated uncertainty, the intervals become noticeably broader than the benchmark (right panel). Third, the influence of uncertainty is more pronounced for the current-year forecast

[7]The uncertainty measure (ζ_t) attains a minimum of 24.8, a maximum of 76.3, a mean of 44.3 and a standard deviation of 11.0 over the period 1996–2019.

Figure 1. Confidence intervals for GDP growth forecasts (Baseline)

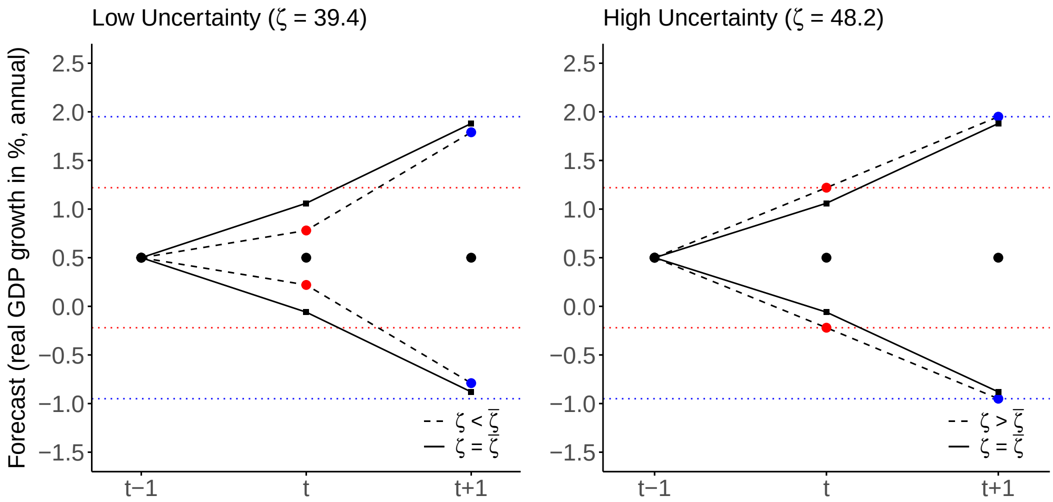

Note: The black dashed lines in the panels show confidence intervals for GDP forecast (centered at the potential output growth rate) for t ($h = 0$, red dots) and $t + 1$ ($h = 1$, blue dots) under low (left panel) and high (right panel) uncertainty (25th and 75th percentiles). The solid black lines indicate confidence intervals at the mean uncertainty level ($\zeta_t = \bar{\zeta} = 44.3$), which yields the benchmark confidence interval according to equation (3).

($h = 0$) than for the one-year-ahead forecast ($h = 1$). This differential impact reflects the following

$$\frac{\hat{\sigma}^2_{0,m}}{\hat{\beta}(\zeta_{t,m} - \bar{\zeta})} < \frac{\hat{\sigma}^2_{1,m}}{\hat{\beta}(\zeta_{t,m} - \bar{\zeta})} \tag{16}$$

with the consequence that variations in the uncertainty measure $\zeta_{t,m}$ cause a quantitatively larger effect on the width of the confidence interval for $h = 0$ than for $h = 1$.

3. Extensions

The preceding analysis was designed to present the conceptual framework in a manner that is intentionally simple and transparent. However, achieving this level of clarity necessitated the adoption of several strong and, at times, implausible assumptions and restrictions. In what follows, we systematically relax them and address three main elements: first, the parameter restrictions regarding the partial effect of the uncertainty measure on the squared forecast error; second, the linearity assumption embedded in equation (5); and third, the exclusive reliance on a single uncertainty measure, namely business-level uncertainty, to the exclusion of additional measures such as household or economic policy uncertainty.

3.1. **Restricted versus unrestricted estimates for the uncertainty coefficient.** The baseline approach assumes a fully restricted model in which $\beta_{h,m} = \beta$, that is, the effect of uncertainty on the forecast error is assumed to be identical for both the current year ($h = 0$) and the next year ($h = 1$), and invariant with respect to the forecast vintage (m). However, this restriction may not be warranted in practice. Against this background, we extend the analysis along two dimensions. First, we investigate whether the partial effect of the uncertainty measure varies across the forecasting horizon, allowing for $\beta \to \beta_h$, that is, distinct effects for the current-year and one-year-ahead forecasts. Second, we consider whether the partial effect differs by the month m in which the forecast is made, so that $\beta \to \beta_m$.

To empirically assess these possibilities, we estimate the most general model, which allows for variation in both dimensions, that is, $\beta \to \beta_{h,m}$. Under this specification, equation (13) becomes:

$$\begin{bmatrix} \boldsymbol{\varepsilon}_{0,t} \\ \boldsymbol{\varepsilon}_{1,t} \end{bmatrix} = \begin{bmatrix} \boldsymbol{\sigma}_{0} \\ \boldsymbol{\sigma}_{1} \end{bmatrix} + \text{vec}(\boldsymbol{B}) \odot (\mathbf{1}_2 \otimes \boldsymbol{\zeta}_t) + \begin{bmatrix} \boldsymbol{u}_{0,t} \\ \boldsymbol{u}_{1,t} \end{bmatrix}, \tag{17}$$

$$\text{where} \quad \boldsymbol{B} = \begin{bmatrix} \beta_{0,M} & \beta_{1,M} \\ \beta_{0,J} & \beta_{1,J} \\ \beta_{0,S} & \beta_{1,S} \\ \beta_{0,D} & \beta_{1,D} \end{bmatrix}$$

Our primary interest lies in the estimates of the elements of the matrix $\boldsymbol{B}$. Estimation is again carried out using the SUR framework, with results reported in Table 1 under the column labeled "Unrestricted". Turning first to the unconditional forecast error variances, we observe that the estimates for $\sigma^2_{h,m}$ show only minor deviations from those in the baseline model. Importantly, all variance estimates in the unrestricted specification are now statistically significantly different from zero.

Turning to the estimates for the partial effects ($\beta_{h,m}$), we find that all point estimates are positive, as expected, but that their magnitudes vary considerably across both forecast horizon h and forecast vintage m. For example, uncertainty appears to have little explanatory power for the (squared) forecast errors of current-year forecasts ($h = 0$) produced in March ($m = M$) or June ($m = J$). However, for forecasts for the next year ($h = 1$) made in these same months, uncertainty significantly accounts for a portion of the variation in the forecast errors. To formally assess these patterns, we employ a Wald test of the null hypothesis that $\beta_{h,m} = \beta$ for all h and m. The results, shown in the lower part of Table 1, indicate that this hypothesis is rejected at the five percent level, supporting the more general model in which the effect of uncertainty varies with both the forecast horizon h and the forecast vintage m.

We further explore the structure of the partial effects by considering two intermediate cases: variation only across forecast vintages (m) and variation only across forecast horizons (h).

3.1.1. *Variation over the forecast vintage (m).* If the partial effects are allowed to vary across the forecast vintages but not across the forecast horizons, so that $\beta \to \beta_m$ for $m \in \{M, J, S, D\}$, equation (17) simplifies to:

$$\begin{bmatrix} \boldsymbol{\varepsilon}_{0,t} \\ \boldsymbol{\varepsilon}_{1,t} \end{bmatrix} = \begin{bmatrix} \boldsymbol{\sigma}_0 \\ \boldsymbol{\sigma}_1 \end{bmatrix} + (\mathbf{1}_2 \otimes \boldsymbol{b}) \odot (\mathbf{1}_2 \otimes \boldsymbol{\zeta}_t) + \begin{bmatrix} \boldsymbol{u}_{0,t} \\ \boldsymbol{u}_{1,t} \end{bmatrix}, \quad \text{where} \quad \boldsymbol{b} = \begin{bmatrix} \beta_M \\ \beta_J \\ \beta_S \\ \beta_D \end{bmatrix} \tag{18}$$

The lower part of Table 1 reports the results of a Wald test that compares the above specification with the fully unrestricted model given by equation (17). The corresponding null hypothesis, $\beta_{h,m} = \beta_m$ for all h, is rejected at the one percent significance level. The test therefore rejects the homogeneity of the coefficients in the forecasting horizons h.

3.1.2. *Variation over the forecast horizon (h).* When allowing the partial effects to vary by forecast horizon, that is, $\beta \to \beta_h$ for $h \in \{0, 1\}$, equation (17) reduces to:

$$\begin{bmatrix} \boldsymbol{\varepsilon}_{0,t} \\ \boldsymbol{\varepsilon}_{1,t} \end{bmatrix} = \begin{bmatrix} \boldsymbol{\sigma}_0 \\ \boldsymbol{\sigma}_1 \end{bmatrix} + \begin{bmatrix} \beta_0 & 0 \\ 0 & \beta_1 \end{bmatrix} \otimes \boldsymbol{\zeta}_t + \begin{bmatrix} \boldsymbol{u}_{0,t} \\ \boldsymbol{u}_{1,t} \end{bmatrix} \tag{19}$$

The corresponding Wald test results, reported in Table 1, examines the null hypothesis, $\beta_{h,m} = \beta_h$ for all m. This hypothesis is rejected at the ten percent level of statistical significance. The test therefore rejects the homogeneity of the coefficients between forecasting vintages m.

The two Wald tests reject restrictions on the coefficient of the full model given by equation (17). The fact that the restrictions on the vintages are only weakly rejected suggests that there is more heterogeneity across forecasting horizons than across forecasting vintages.

3.2. **The role of non-linearities.** The linear regression model proposed in equation (5) carries the drawback that its predicted values for the squared forecast error – a variable that is necessarily non-negative – may become negative. This limitation motivates the exploration of non-linear alternative specifications that ensure the predicted values conform to the non-negativity constraint.

To address this, we consider a log-linear reformulation of equation (5) which applies a logarithmic transformation of the squared forecast error. This yields an exponential specification for the predicted quantity of interest. Specifically, we model the log of the squared forecast error as:

$$\ln\left(\varepsilon^2_{t+h|t,m}\right) = \alpha_{h,m} + \beta(\zeta_{t,m} - \bar{\zeta}) + u_{t,m} \tag{20}$$

where u_t is a again a mean-zero error term. The logarithmic transformation guarantees that, when exponentiated, the predicted values are strictly non-negative.

We reformulate the model within the multivariate system structure of equation (13) and re-estimate its parameters using the SUR approach to appropriately capture cross-equation correlation among the errors. The estimation results, summarized in Table 2, indicate that the coefficient β is positive and statistically significant at the one percent level. While the magnitude of β is not directly comparable to its counterpart from the linear specification in Table 2, the estimates for α_h can be usefully contrasted. Specifically, when $\zeta_{t,m} = \bar{\zeta}$, it follows that $\hat{\varepsilon}^2_{t+h|t} = e^{\hat{\alpha}_{h,m}} = \hat{\sigma}^2_{h,m}$. The findings reveal that the estimates are quantitatively smaller in size but they still exhibit consistent variation across forecast horizons and months; namely, we observe $\sigma^2_{1,m} > \sigma^2_{0,m}$ for all m, and $\sigma^2_{h,m} > \sigma^2_{h,m+1}$ for any given h, mirroring the patterns detected in the original linear framework.

Finally, we examine the implications for the width of the confidence intervals, as illustrated in Figure 2. This figure is constructed analogously to Figure 1,

Table 2. Estimates of $\sigma^2_{h,m} = e^{\alpha_{h,m}}$ and β using the log-linear specification

$\hat{\sigma}^2_{0,M}$	$\hat{\sigma}^2_{0,J}$	$\hat{\sigma}^2_{0,S}$	$\hat{\sigma}^2_{0,D}$	$\hat{\sigma}^2_{1,M}$	$\hat{\sigma}^2_{1,J}$	$\hat{\sigma}^2_{1,S}$	$\hat{\sigma}^2_{1,D}$	$\hat{\beta}$
0.20***	0.13***	0.12***	0.08***	0.81	0.74	0.38**	0.30***	0.034***

Note: The table presents estimates for the log-linear specification with $\hat{\sigma}^2_{h,m} = e^{\hat{\alpha}_{h,m}}$. The estimation uses a direct survey-based uncertainty measure collected among manufacturing firms. Statistical significance levels are: 10% (*), 5% (**), and 1% (***).

thereby enabling a direct comparison of the resulting intervals for the GDP forecasts under both specifications. Our results show that the width of the confidence interval is substantially smaller with the log-linear model compared to the baseline. Most notably, although the uncertainty measure continues to influence the width of the confidence interval, its impact is markedly reduced in the log-linear specification. As a result, the unconditional forecast error variance ($\sigma^2_{h,m}$) emerges as the dominant factor shaping the width of the confidence intervals in this setting.[8]

Nevertheless, these findings underscore the robustness of the relationship between uncertainty and forecast error magnitude across alternative functional forms. The log-linear transformation, in particular, is advantageous in practical applications, as it ensures the non-negativity of the dependent variable, while on the other hand, the effect of the uncertainty measure on the width of the confidence interval is limited.

3.3. **Further uncertainty measures.** The preceding models rely exclusively on a firm-specific uncertainty measure derived directly from a business survey. We now broaden the empirical framework by incorporating alternative sources of uncertainty: a household uncertainty measure derived from consumer survey data, the VSTOXX and VIX indices as indicators of implied stock market volatility and hence uncertainty as perceived by financial market participants, and a policy uncertainty measure for Germany – Austria's principal trading partner – as constructed by Baker et al. (2016).[9] Each metric

[8]We further explored a non-linear extension of equation (20) by including a quadratic term, $\tilde{\beta}(\zeta_t - \bar{\zeta})^2$, on the right-hand side. Although the estimated coefficient for $\tilde{\beta}$ is statistically significantly different from zero, its magnitude is exceedingly small. As a result, the corresponding confidence intervals are visually indistinguishable from those depicted in Figure 2. Thus, while the quadratic term formally influences the width of the confidence interval, its quantitative effect is negligible.

[9]We employ the German policy uncertainty index owing to the lack of equivalent measures specifically for Austria.

FIGURE 2. Confidence intervals for GDP growth forecasts (log-linear)

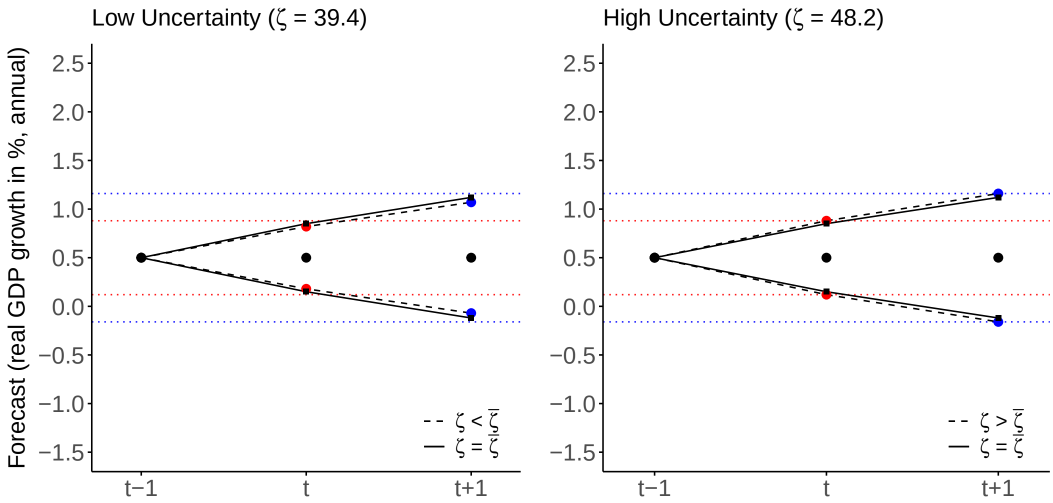

Note: The black dashed lines in the panels show confidence intervals for GDP forecast (centered at the potential output growth rate) for $h = 0$ (red dots) and $h = 1$ (blue dots) under low (left panel) and high (right panel) uncertainty (25th and 75th percentiles). The solid black lines indicate confidence intervals at the mean uncertainty level ($\zeta_t = \bar{\zeta} = 44.3$).

reflects a distinct, yet empirically relevant, dimension of economic uncertainty, all of which have documented effects on both real GDP growth and the accuracy of macroeconomic forecasts (Camacho and Garcia-Serrador, 2014). The household-specific measure is based on the perceived unemployment risk,[10] as obtained from a particular question in monthly consumer survey (European Commission, 2025).

We estimate the extended model – incorporating these additional uncertainty measures alongside the baseline measure – which gives rise to the following new specification of equation (11):

$$\varepsilon^2_{t+h|t,m} = \sigma^2_{h,m} + \sum_{k=1}^{5} \beta_k(\zeta_{t,m,k} - \bar{\zeta}_k) + u_{t,m} \tag{21}$$

where the $\zeta_{t,m,k}$ terms ($\forall\ k = 1, ..., 5$) refer to the five uncertainty measures used. We estimate the extended model again using the SUR approach. The estimation results are presented in Table 3. The estimates indicate that the unconditional forecast error variances ($\sigma^2_{h,m}$) exhibit only negligible changes

[10]Exact Formulation of the Question on Uncertainty derived from the perceived risk of job loss (as per the official harmonized EU Consumer Survey): "How do you expect the number of people unemployed in this country to change over the next 12 months? The number will: increase sharply / increase slightly / remain the same / fall slightly / fall sharply / Don't know (N)"

Table 3. Estimates of $\sigma^2_{h,m}$ and β using various uncertainty measures

Unconditional variance ($\hat{\sigma}^2_{h,m}$)		Slope coefficients ($\hat{\beta}_k$)	
$\hat{\sigma}^2_{0,M}$	0.72**	Firm uncertainty	0.043**
$\hat{\sigma}^2_{0,J}$	0.44	Household uncertainty	0.005
$\hat{\sigma}^2_{0,S}$	0.33	VSTOXX	-0.035
$\hat{\sigma}^2_{0,D}$	0.27	VIX	0.072**
$\hat{\sigma}^2_{1,M}$	2.58**	Policy uncertainty (DE)	0.008**
$\hat{\sigma}^2_{1,J}$	2.50***		
$\hat{\sigma}^2_{1,S}$	1.91**		
$\hat{\sigma}^2_{1,D}$	1.18**		

Note: The table presents estimates for the baseline specification ($\hat{\beta}_{h,m,k} = \hat{\beta}_k \; \forall \; k$) using the baseline uncertainty measure ("firm uncertainty") and four additional ones: household uncertainty, stock market uncertainty (VIX and VSTOXX) and economic policy uncertainty. Statistical significance levels are: 10% (*), 5% (**), and 1% (***).

compared to the baseline model. Importantly, in the extended specification, all variance estimates have statistical significance similar to the baseline model.

Of greater importance are the estimated slope coefficients of the various uncertainty measures. The partial effect of the baseline uncertainty measure remains statistically significant at the five percent level and preserves its size (by and large). In contrast, the partial effects of the household uncertainty measure and the VSTOXX are not statistically distinguishable from zero, with the latter even exhibiting a coefficient with an unexpected sign. Strikingly, the estimates for the coefficients for the VIX index and the German economic policy uncertainty index are positive and statistically significantly different from zero at the five percent level. Taken together, these results underscore both the relevance of uncertainty measures in explaining movements in (squared) forecast errors and the practical value of adjusting the width of confidence intervals attached to forecasts in accordance with prevailing uncertainty conditions.

4. Alternative modeling approaches

Beyond the primary framework proposed in this study, alternative econometric approaches are available for incorporating evolving uncertainty measures into the modeling of forecast error variance and consequently the width of confidence intervals. Two prominent methodologies are extensions of GARCH models with exogenous regressors – commonly referred to as GARCH-X models

– and stochastic volatility models with latent factor structures (see Hauzenberger et al., 2018, for instance) designed to reflect both observed and unobserved sources of uncertainty.

GARCH-X models extend the standard GARCH specification by including contemporaneous external indicators directly in the conditional variance equation. This enables a parsimonious and interpretable framework for capturing the influence of, in our case, exogenous uncertainty shocks on forecast error volatility. The mechanics and practical implementation of the GARCH-X approach are outlined in detail in the Appendix.

Alternatively, models with stochastic volatility (SV) combined with factor structures offer a highly flexible environment for modeling volatility dynamics. Such models can parsimoniously represent both common and idiosyncratic volatility components associated with multiple uncertainty sources and allow for richer volatility features than standard GARCH processes. The Bayesian estimation techniques central to these approaches, as demonstrated in Kastner and Hosszejni (2021), permit the modeling of latent volatility factors and their attribution to macroeconomic or financial uncertainty proxies. This factor-driven SV approach excels in capturing persistent and common movements in volatility across series and in accommodating highly flexible latent volatility dynamics.

Despite their conceptual appeal, a central impediment to the reliable implementation of both the GARCH-X and factor SV models is the limited number of available observations. Macroeconomic forecasting datasets often span only a few decades at best, and with multiple forecast vintages or series involved, the resulting time series are short relative to the number of parameters to be estimated – especially for the computationally intensive factor SV models. This makes parameter estimation challenging, raises risks of overfitting, and underscores the need for model parsimony or hierarchical shrinkage techniques in practical work.

It should be emphasized that this list of modeling extensions is by no means exhaustive. Additional promising avenues for future research include, for example, the construction of asymmetric confidence intervals (Kaniovski, 2019), which may better capture the asymmetric risks associated with severe economic shocks.

Figure 3. Use case: GDP (annual growth rate) forecasts of March 2022 and March 2023

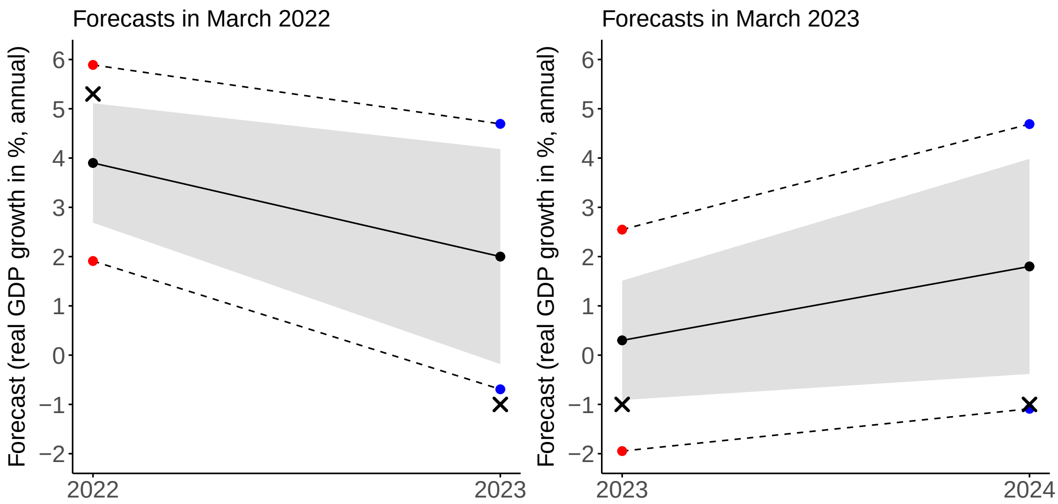

Note: The figure displays the actual forecasts of annual GDP growth rates produced by Wifo in March 2022 (covering 2022 and 2023) and March 2023 (covering 2023 and 2024), indicated by black dots. Corresponding realized growth rates (first releases) are marked by x-shaped symbols. The gray shaded areas depict benchmark confidence intervals derived from the variance of historic forecast errors, while the red and blue dots along black dashed lines represent uncertainty-dependent confidence intervals constructed using values for the three statistically significant uncertainty measures (see, Table 3) available at the time of each forecast.

5. Use case: Post-Covid-19 forecasting and uncertainty in practice

We consider a use case to demonstrate the operation and implications of uncertainty-dependent confidence intervals. The post-COVID-19 period provides an especially relevant context for this analysis, as it was characterized by an extraordinarily high level of uncertainty. This uncertainty was shaped by several factors, including geopolitical tensions, economic policy ambiguity, and doubts regarding the steady supply of intermediate goods and energy. To capture these dynamics, we analyze Wifo forecasts produced in March 2022 and March 2023. In each instance, forecasts were generated for the current year as well as the subsequent year.

Figure 3 presents the forecasts for GDP growth: Black points plotted on solid black lines indicate the point forecasts for 2022 and 2023 based on the March 2022 forecast (left panel), and for 2023 and 2024 based on the March 2023 forecast (right panel). The gray shaded area represents the benchmark confidence interval constructed from the variance of (historic) forecast errors, as detailed in the first numerical column of Table 1. The x-shaped markers in each subplot denote the first official release of GDP data (annual growth rate). Notably, the realized GDP growth rate falls outside the benchmark confidence interval in all four cases. This indicates that the magnitude of forecast errors was sufficiently large that confidence intervals based solely on historic forecast error variances failed to capture the true economic outcomes.

In contrast, the uncertainty-dependent confidence intervals, depicted by red dots (current year's forecast, $h = 0$) and blue dots (next year's forecast, $h = 1$) along dashed black lines, paint a different picture. These intervals derive from the approach presented in Section 3.3 and use the three statistically significant uncertainty measures (see, Table 3). The modified intervals successfully capture the realized GDP growth in three out of four cases: consistently for the current year's forecast, and once for the next year's forecast.

5.1. **Discussion.** This improved performance is attributable to the model's accommodation of elevated uncertainty levels, stemming from firm-level volatility, economic policy unpredictability, and heightened stock market participants' perceptions. By adapting the width of confidence intervals to contemporaneous uncertainty measures, the uncertainty-adjusted intervals provide a more realistic representation of the inherent risks in economic forecasting during turbulent periods. Such enhanced contextualization is particularly advantageous for the communication of forecasts in times of heightened uncertainty, as it tempers expectations and informs policy and market participants of the broader range of plausible outcomes. Furthermore, this approach encourages a more prudent interpretation of forecast data, reducing the risk of overconfidence and facilitating better-informed decision-making under uncertainty.

5.2. **Why not adjusting κ?** A straightforward alternative would be to increase the value of κ in the traditional approach (see, Equation (3)) uniformly to widen the confidence intervals. Although this may improve coverage of realizations, the approach has notable drawbacks:

- The widening is non-specific, i.e., independent of the actual uncertainty level at the forecast time. During periods of low uncertainty – and thus relatively low forecast difficulty – this results in unnecessarily wide intervals, diminishing the forecast's informational value. Excessively broad intervals in stable periods may create an impression of imprecision, undermining the forecast's credibility.

- A uniform adjustment of κ ignores the empirically observed heterogeneity of forecast errors associated with differing uncertainty levels, thus discarding valuable time-specific information.
- Unlike the uncertainty-dependent approach, merely increasing κ lacks a dynamic adaptation to changing economic conditions and fails to achieve context-sensitive communication of forecast uncertainty.

For these reasons, an uncertainty-adjusted confidence interval is clearly superior to a blanket adjustment of the critical value κ: it enables a more precise, context-appropriate, and thus more credible representation of forecast uncertainty. We regard this as the key advantage of the approach: it allows for a clear and interpretable presentation of the inherent uncertainty forecasters face and thereby significantly contributes to the assessment of the reliability of economic forecasts.

6. Conclusion

This study introduces a simple yet effective procedure for integrating exogenous uncertainty measures into the construction of confidence intervals for macroeconomic forecasts. The proposed framework results in intervals that appropriately expand or contract with prevailing uncertainty, thereby enhancing the interpretability and credibility of forecasts. Empirical evidence from Austrian data confirms that accounting for uncertainty – especially via survey-based and policy indicators – explains the variation in forecast errors.

Importantly, the implementation of this approach is not confined to the Austrian context. The increasing availability of high-frequency uncertainty data, including survey-based indicators, in Germany and many other countries (Rossi, 2020; Lautenbacher et al., 2021) enables direct application and empirical evaluation of this methodology in a wide range of settings. This broad applicability underlines the potential of uncertainty-dependent confidence intervals as a transparent and robust tool for communicating the reliability of macroeconomic forecasts across diverse institutional environments.

The results highlight how uncertainty-dependent intervals can transparently reveal the limitations and confidence levels of economic forecasts, giving users clearer insight into prediction reliability.

Finally, it should be noted that uncertainty-dependent confidence intervals are not meant to make inaccurate forecasts appear more favorable by artificially widening the intervals. Rather, they are meant to quantify the current uncertainty and link it directly to the forecast. The endeavor to produce a point forecast that is as close as possible to the eventual outcome, based on the available data, and whose deviations (forecast errors) do not exhibit any systematic patterns, should remain the central aspect of forecasting.

References

Baker, Scott, Nicholas Bloom, and Steven Davis (2016) 'Measuring economic policy uncertainty.' *The Quarterly Journal of Economics* 131(4), 1593–1636

Biswas, Rumpa (2019) 'Does economic policy uncertainty affect analyst forecast accuracy?' Technical Report, University of New Orleans, January

Camacho, Maximo, and Agustin Garcia-Serrador (2014) 'The Euro-Sting Revisited: The Usefulness of Financial Indicators to Obtain euro area GDP Forecasts.' *Journal of Forecasting* 33(3), 186–197

Chourou, Lamia, Lynnette Purda, and Samir Saadi (2021) 'Economic policy uncertainty and analysts' forecast characteristics.' *Journal of Accounting and Public Policy* 40(4), 106775

European Commission (2021) 'European business cycle indicators, 3rd quarter 2021: Special topic: New survey-based measures of economic uncertainty.' Technical Paper 051, European Commission Directorate-General for Economic and Financial Affairs

——— (2025) 'The Joint Harmonised EU Programme of Business and Consumer Surveys: User Guide.' Technical Report, Directorate-General for Economic and Financial Affairs (DG ECFIN), European Commission

Fortin, Ines, Sebastian P. Koch, and Klaus Weyerstrass (2020) 'Evaluation of economic forecasts for Austria.' *Empirical Economics* 58(1), 107–137

Glocker, Christian, and Werner Hölzl (2022) 'A direct measure of subjective business uncertainty.' *German Economic Review* 23(1), 121–155

Hauzenberger, Niko, Maximilian Böck, Michael Pfarrhofer, Anna Stelzer, and Gregor Zens (2018) 'Implications of macroeconomic volatility in the euro area.' Department of Economics Working Paper Series 261, WU Vienna University of Economics and Business

Hsieh, Pei-Hua (2016) 'Prediction intervals for macroeconomic variables based on uncertainty measures.' PhD dissertation, Tilburg University

Hölzl, Werner, Jürgen Bierbaumer, Michael Klien, and Agnes Kügler (2025) 'Economic Assessments Remained Cautious in Spring. Results of the WIFO-Konjunkturtest Quarterly Survey of April 2025.' *WIFO Reports on Austria, No. 8*

Kaniovski, Serguei (2019) 'Probability intervals for STATEC forecasts.' Working papers du STATEC 107/2019, STATEC, May

Kannan, Prakash, and Selim A Elekdag (2009) 'Incorporating market information into the construction of the fan chart.' IMF Working Papers 2009/178, International Monetary Fund, August

Kastner, Gregor, and Darjus Hosszejni (2021) 'Modeling Univariate and Multivariate Stochastic Volatility in R with stochvol and factorstochvol.' *Journal of Statistical Software* 100(12), 1–34

Lautenbacher, Stefan, Stefan Sauer, and Klaus Wohlrabe (2021) 'How difficult is it for managers to predict their future business development? A new measure of perceived business uncertainty.' *CESifo Forum* 22(01), 40–42

Morikawa, Masayuki (2023) 'Firms' subjective uncertainty and forecast errors: Survey evidence from Japan.' *Applied Economics Letters* 30(1), 33–36

Razi, Ahmad, and Po Ling Loke (2017) 'Fan chart: The art and science of communicating uncertainty.' In *Statistical implications of the new financial landscape,* ed. Bank for International Settlements, vol. 43 (Bank for International Settlements)

Rossi, Luca (2020) 'Indicators of uncertainty: a brief user's guide.' Questioni di Economia e Finanza (Occasional Papers) 564, Bank of Italy, Economic Research and International Relations Area, Jun

Rülke, Jan-Christoph, Maria Silgoner, and Julia Wörz (2016) 'Herding behavior of business cycle forecasters.' *International Journal of Forecasting* 32(1), 23–33

Turner, David (2017) 'Designing fan charts for GDP growth forecasts to better reflect downturn risks.' OECD Economics Department Working Papers 1428, OECD, November

Wang, Yanwen, Yanying Sun, Yalong Li, Chen Feng, and Peng Chen (2023) 'Interval Forecasting Method of Aggregate Output for Multiple Wind Farms Using LSTM Networks and Time-Varying Regular Vine Copulas.' *Processes* 11(5), 1530

Appendix A. GARCH modeling of forecast error variance with an uncertainty measure

In the following, we consider a completely different approach to capture the time-varying nature of forecast error volatility. We extend the baseline framework by modeling the forecast error ($\varepsilon_{t+h|t}$) with a conditional heteroskedasticity process. Specifically, we employ a GARCH(1) specification augmented by the external uncertainty measure ζ_t, thereby allowing for a direct estimation of the partial effect of uncertainty on the conditional variance (denoted by $\sigma^2_{t+h|t}$).

Let $\varepsilon_{t+h|t}$ denote the forecast error at time t for horizon h, as previously defined. We specify the following GARCH-model:

$$\varepsilon_{t+h|t} = \sigma_{t+h|t} z_t, \quad z_t \sim \text{i.i.d. } N(0,1) \tag{22}$$

$$\sigma^2_{t+h|t} = \omega + \rho\sigma^2_{t+h-1|t-1} + \gamma(\zeta_t - \bar{\zeta}) \tag{23}$$

where $\omega > 0$ is a constant, $\rho \geq 0$ is the GARCH parameter measuring the persistence of conditional variance, and γ captures the direct partial effect of the (centered) uncertainty measure $(\zeta_t - \bar{\zeta})$ on the conditional forecast error variance.

Equation (23) extends the standard GARCH(1) model by adding an exogenous regressor, $(\zeta_t - \bar{\zeta})$, to the conditional variance equation, referred to as GARCH(1)-X. The term $\sigma^2_{t+h|t}$ denotes the conditional variance of the forecast error at horizon h given information available up to time t. It captures the expected magnitude of the squared forecast error, allowing this expectation to evolve dynamically over time according to historical volatility and contemporaneous economic uncertainty. This setup allows the conditional volatility of the forecast error to systematically respond to contemporaneous fluctuations in observed uncertainty, after controlling for its own past shocks and past volatility. This time-varying variance comprises an alternative approach for constructing uncertainty dependent confidence intervals within the GARCH approach.

The parameter γ directly quantifies the marginal contribution of changes in the uncertainty measure to the (conditional) forecast error volatility. A positive and statistically significant estimate for γ would suggest that heightened uncertainty – whether due to survey signals, financial market volatility, or policy risks – translates immediately into higher forecast error variance and, consequently, broader confidence intervals.

Estimation of this model can proceed via maximum likelihood, treating the conditional variance equation as a standard GARCH(1) with exogenous covariates. The resulting time-varying conditional standard deviation, $\hat{\sigma}_{t+h|t}$, can then be used in equation (3) to construct confidence intervals.

This approach provides both a flexible and empirically grounded method for capturing the partial impact of uncertainty on forecast reliability through the lens of dynamic volatility modeling.

A.1. **Practical implementation.** We again use the forecast errors of Wifo $\varepsilon_{t+h|t,m}$ for $h = 0, 1$ and $m = \{M, J, S, D\}$. To accurately reflect the structure of the data, both the forecast error and the associated uncertainty measure are indexed by the forecast horizon h and forecast vintage m.

As before, let $\varepsilon_{t+h|t,m}$ denote the forecast error for period $t+h$, made at time t in month $m \in \{M, J, S, D\}$, and let $\zeta_{t,m}$ represent the value of the uncertainty indicator for the corresponding month and year. The univariate GARCH(1)-X model with an exogenous uncertainty regressor is then specified as follows:

$$\varepsilon_{t+h|t,m} = \sigma_{t+h|t,m}\, z_t, \qquad z_t \sim \text{i.i.d. } N(0,1) \tag{24}$$

$$\sigma^2_{t+h|t,m} = \omega_{h,m} + \rho_{h,m}\, \sigma^2_{t+h-1|t-1,m} + \gamma_{h,m}\, (\zeta_{t,m} - \bar{\zeta}_m) \tag{25}$$

Here, $\omega_{h,m}$ and $\rho_{h,m}$ are GARCH parameters which may vary by h and m and $\gamma_{h,m}$ measures the partial effect of the uncertainty measure, assumed flexible across h and m.

TABLE 4. Estimates of $\sigma^2_{h,m}$ and $\beta_{h,m}$ using GARCH(1)-X

Unconditional variance ($\hat{\sigma}^2_{h,m}$)		Slope coefficients ($\hat{\beta}_{h,m}$)	
$\hat{\sigma}^2_{0,M}$	0.75	$\hat{\beta}_{0,M}$	0.048
$\hat{\sigma}^2_{0,J}$	0.36	$\hat{\beta}_{0,J}$	0.025
$\hat{\sigma}^2_{0,S}$	0.14	$\hat{\beta}_{0,S}$	0.014
$\hat{\sigma}^2_{0,D}$	0.10	$\hat{\beta}_{0,D}$	0.007
$\hat{\sigma}^2_{1,M}$	2.10**	$\hat{\beta}_{1,M}$	0.067**
$\hat{\sigma}^2_{1,J}$	2.10*	$\hat{\beta}_{1,J}$	0.067*
$\hat{\sigma}^2_{1,S}$	1.61**	$\hat{\beta}_{1,S}$	0.053**
$\hat{\sigma}^2_{1,D}$	1.04	$\hat{\beta}_{1,D}$	0.045

Note: The table presents estimates for the GARCH(1)-X model using the baseline uncertainty measure (direct survey-based measure collected among manufacturing firms). The values refer to the following two transformed parameters: $\sigma^2_{h,m} = \frac{\omega_{h,m}}{1-\rho_{h,m}}$ and $\beta_{h,m} = \frac{\gamma_{h,m}}{1-\rho_{h,m}}$. The estimation is based on 24 observations for each forecast vintage and horizon (1996 to 2019). Statistical significance levels are: 10% (*), 5% (**), and 1% (***) and refer to the estimates of $\omega_{h,m}$ for $\sigma^2_{h,m}$ and $\gamma_{h,m}$ for $\beta_{h,m}$.

We estimate the time-varying conditional standard deviation, $\widehat{\sigma}_{t+h|t,m}$, for each forecast horizon h and forecast vintage m to construct confidence intervals, as specified in equation (3).

The estimated coefficients are reported in Table 4. These coefficients have been transformed to allow for direct comparison with the corresponding estimates from the unrestricted baseline model, as presented in the last column of Table 1. Specifically for the GARCH(1)-X model, we obtain $\sigma^2_{h,m} = \omega_{h,m}/(1-\rho_{h,m})$ and $\beta_{h,m} = \gamma_{h,m}/(1-\rho_{h,m})$. The estimates for $\sigma^2_{h,m}$ are generally in line with those of the unrestricted baseline model. However, within the GARCH(1)-X framework, most estimated parameters are not statistically significantly different from zero (statistical significance indicated in Table 4 is based solely on the estimates of $\omega_{h,m}$ and $\gamma_{h,m}$), and the same holds for the $\beta_{h,m}$ estimates. The reduced statistical significance likely reflects the smaller effective sample size available for GARCH model estimation relative to the system-wide estimation employed in the baseline model.

Figure 4 displays the resulting confidence intervals. The width of the benchmark confidence interval (where $\zeta_t = \bar{\zeta}$) closely resembles that of the baseline

Figure 4. Confidence intervals for GDP growth forecasts (GARCH(1)-X)

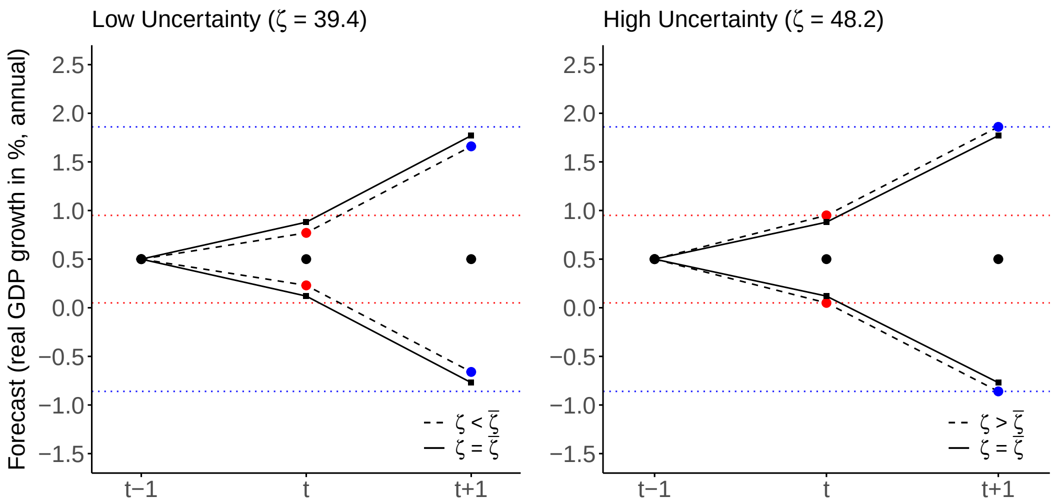

Note: The black dashed lines in the panels show confidence intervals for GDP forecast (centered at the potential output growth rate) for $h = 0$ (red dots) and $h = 1$ (blue dots) under low (left panel) and high (right panel) uncertainty (25th and 75th percentiles). The solid black lines indicate confidence intervals at the mean uncertainty level ($\zeta_t = \bar{\zeta} = 44.3$).

model in Figure 1. Nevertheless, within the GARCH(1)-X framework, the influence of the uncertainty measure on the width of the interval remains modest for both the current-year forecast ($h = 0$) and the one-year-ahead forecast ($h = 1$).

Despite these limitations, the GARCH(1)-X modeling approach confers an important advantage: it allows the volatility of forecast errors, and thus the width of prediction intervals, to respond dynamically to both past dynamic of forecast error volatility (the GARCH-effect) and the contemporaneous level of exogenous uncertainty (the X-effect). This underscores the usefulness of the GARCH-X model as an alternative methodology.

Appendix B. Further remarks

B.1. **First release data versus final release data for GDP.** An alternative to relying on the first release of GDP data to compute forecast errors, would be to use, for example, the final available data vintage (Rülke et al., 2016). However, this approach is associated with several important drawbacks. Most notably, in the context of GDP and National Accounts statistics, truly final data do not exist, as periodic revisions can extend many years or even decades into the past. As a result, the reference point for "final" data is itself inherently fluid.

Moreover, the first release of GDP data – typically published around nine months after the end of the respective year – offers key methodological advantages. In particular, the use of first-release data facilitates the timely incorporation of definitional or methodological changes to GDP and the broader data-generating process. This is a salient consideration, as substantial revisions to National Accounts occurred in 2014, altering both the level and growth rates of GDP. By benchmarking forecasts against first-release data, it is possible (in most cases) to ensure that forecast evaluation is consistent with the definitions and classifications that were in place at the time the forecasts were produced.

In contrast, employing the most recent data vintage for forecast evaluation disregards contemporaneous changes in definitions and classifications, thus undermining the comparability of forecast errors over time. Furthermore, forecasts cannot anticipate the effects of future changes to the System of National Accounts (SNA), as these rely by necessity on the information available at the time of forecast production. For these reasons, the use of first-release data represents a more logically consistent and practically feasible basis for forecast evaluation in the presence of evolving statistical standards.

B.2. **The direct measure of subjective business uncertainty.** The question on uncertainty in the business survey was reworded in 2014. The original wording was: "Die zukünftige ENTWICKLUNG unserer Geschäftslage ist • in gewissem Maße abschätzbar • wenig abschätzbar • sehr unsicher • unsicherer als je zuvor." (Translation: The future development of our business situation can be assessed to a certain degree / is difficult to assess / is very uncertain / is more uncertain than ever.) Given the limited length of the post-2014 subsample, it is not feasible to conduct an econometric analysis based solely on the period covered by the revised wording. To construct a consistent long-run series, we therefore merge the uncertainty indicators based on the old and new survey formulations. Specifically, each measure is standardized within its respective sub-period, after which the standardized series are concatenated to form a unified time series of direct subjective uncertainty, spanning from the first quarter of 1996 to the second quarter of 2025. Further methodological details are provided in Glocker and Hölzl (2022).

An equivalent uncertainty measure for manufacturing now exists for all EU countries which is obtained within the EU's monthly business survey. However, the specific question on business uncertainty was introduced for all EU countries only in May 2021 (see European Commission, 2021, for instance), meaning the resulting uncertainty indicator for all EU countries is available for just over four years at this point (from May 2021 onward), although it is available for some countries starting from April 2019. By contrast, Austria's uncertainty indicator is based on survey data that has been collected since 1996, offering almost 30 years of historical coverage up to 2025 (Hölzl et al., 2025).