

Testing Quantile Forecast Optimality*

Jack Fosten^{†‡}

Daniel Gutknecht[§]

Marc-Oliver Pohle[¶]

May 3, 2023

Abstract

Quantile forecasts made across multiple horizons have become an important output of many financial institutions, central banks and international organisations. This paper proposes misspecification tests for such quantile forecasts that assess optimality over a set of multiple forecast horizons and/or quantiles. The tests build on multiple Mincer-Zarnowitz quantile regressions cast in a moment equality framework. Our main test is for the null hypothesis of autocalibration, a concept which assesses optimality with respect to the information contained in the forecasts themselves. We provide an extension that allows to test for optimality with respect to larger information sets and a multivariate extension. Importantly, our tests do not just inform about general violations of optimality, but may also provide useful insights into specific forms of sub-optimality. A simulation study investigates the finite sample performance of our tests, and two empirical applications to financial returns and U.S. macroeconomic series illustrate that our tests can yield interesting insights into quantile forecast sub-optimality and its causes.

JEL Classification: C01, C12, C22, C52, C53

Keywords: Forecast evaluation, forecast rationality, multiple horizons, quantile regression.

*We are grateful to Maximo Camacho, Mike Clements, José Olmo, Emese Lazar, Adam Lee, and Xiaohan Xue for helpful comments on the paper as well as participants of the SAEe 2022 in Valencia, the International Symposium on Forecasting 2021, the Statistische Woche 2022 and 2021 and seminar attendees at the ICMA Centre, University of Reading, 2022, and the University of Hohenheim and Goethe University Frankfurt, 2023. Marc-Oliver Pohle is grateful for support by the Klaus Tschira Foundation, Germany.

[†]King's Business School, King's College London, WC2B 4BG London, UK, E-mail: jack.fosten@kcl.ac.uk

[‡]Data Analytics in Finance and Macro (DAFM) Research Centre, King's College London, UK

[§]Faculty of Economics and Business, Goethe University Frankfurt, 60629 Frankfurt am Main, Germany. E-mail: gutknecht@wiwi.uni-frankfurt.de.

[¶]Heidelberg Institute for Theoretical Studies, 69118 Heidelberg, Germany. E-mail: marc-oliver.pohle@h-its.org.

1 Introduction

Economic and financial forecasters have become increasingly interested in making quantile predictions, often across different quantile levels and at multiple horizons into the future. In financial markets, for instance, such multi-step quantile predictions are produced due to the 10-day value-at-risk (VaR) requirements of the Basel Committee on Banking Supervision.¹ In the growth-at-risk (GaR) literature on the other hand, Adrian et al. (2019) propose quantile models to predict downside risks to real gross domestic product (GDP) growth at horizons ranging from one quarter ahead to one year ahead. These methods are now widely implemented in academic research (Plagborg-Møller et al., 2020; Brownlees and Souza, 2021) and in international institutions like the IMF (Prasad et al., 2019), and are typically applied across various quantile levels. This trend for multi-horizon quantile forecasts has also developed into a growing literature in nowcasting GaR that typically uses several intra-period nowcast horizons (e.g., Antolin Diaz et al., 2021; Ferrara et al., 2021; Carriero et al., 2020). Finally, it is common for central banks, such as the Bank of England, to produce fan charts of key economic variables such as GDP growth, unemployment or the Consumer Price Index (CPI) inflation rate across several quantile levels and horizons.

However, despite the expansion in empirical and methodological research, there is currently very little statistical guidance for assessing whether a set of multi-step ahead, multi-quantile forecasts are consistent with respect to the outcomes observed. This consistency is often referred to as ‘optimality’, ‘rationality’, or ‘calibration’ in the literature, with ‘full optimality’ referring to optimality relative to the information set known to the forecaster, while a weaker form of optimality known as ‘autocalibration’ is defined with respect to the information contained in the forecasts themselves (see Gneiting and Ranjan, 2013; Tsyplakov, 2013). This paper aims to fill this gap in the literature by proposing various (out-of-sample) optimality tests for quantile forecasts that can accommodate predictions either derived from known econometric forecasting models, or from external sources like institutional or professional forecasters. Specifically, we develop tests that assess optimality of quantile forecasts over multiple forecast horizons *and* multiple quantiles simultaneously.

The main test of this paper is a joint test of autocalibration for quantile forecasts obtained across different horizons and quantile levels. The test is based on a series of quantile Mincer-Zarnowitz (MZ) regressions (see Gaglianone et al., 2011) across all quantile levels and horizons, which are in turn used to construct a test statistic for the null hypothesis of autocalibration across horizons and quantiles using a set of moment equalities (e.g., Romano and Shaikh, 2008, 2010; Andrews and Soares, 2010). We suggest a block bootstrap procedure to obtain critical values for the test. The bootstrap is simple to implement and avoids the need to estimate a large variance-covariance matrix that would be required in a more standard Wald-type test. We establish the first-order asymptotic validity of these bootstrap critical values.

The test of autocalibration based on MZ regressions can provide valuable information to forecasters. In particular, failure to reject the null hypothesis of autocalibration suggests that the forecaster may proceed to use the forecasts as they are without the need to ‘re-calibrate’ them. On the other hand, if the null hypothesis is rejected, the test hints at directions for improvement of the forecasts. That is, it informs the forecaster about the horizons, quantiles or horizon-quantile combinations that contributed strongest to the rejection of the null, and thus an improvement of the forecasts is warranted.

¹See for instance: <https://www.bis.org/publ/bcbs148.pdf> [Last Accessed: 18/12/20]

In addition, the estimated MZ regression can be used to infer about the nature of the deviations from autocalibration, when the forecasts are plotted alongside the realisations. The estimated MZ coefficients may also be used to perform a re-calibration of the original forecasts, as has been suggested in the case of mean forecasts in the recent work of Clements (2022).

We provide two extensions of this test for autocalibration. The first extension allows for additional predictors in the MZ regressions, which we call the augmented quantile Mincer-Zarnowitz test. This test is operationally similar to that of the first test, but may provide richer information to the forecaster. It tests a stronger form of optimality relative to a larger information set than autocalibration. If autocalibration is not rejected, but the null hypothesis of the augmented test is rejected, it indicates that the additional variables used in the MZ regression carry additional informational content which should be used in making the forecasts themselves. The second extension allows to test optimality for multiple time series variables and not just for a single variable. Testing multiple time series variables simultaneously may be useful in cases where we are interested in testing whether one type of model delivers optimal forecasts for multiple macroeconomic variables, or across different financial asset returns, e.g. from the same industry.

Finally, as a separate contribution, we also outline in the Appendix a test for monotonically non-decreasing expected quantile loss as the forecast horizon increases. This extends the result of Patton and Timmermann (2012) to the quantile case whereas they focussed on the mean squared forecast error (MSFE) case for optimal multi-horizon mean forecasts. The test makes use of empirical moment inequalities using the Generalised Moment Selection (GMS) procedure of Andrews and Soares (2010). This test can also be seen as complementary to monotonicity tests used in the nowcasting literature for the MSFE of mean nowcasts (see Fosten and Gutknecht, 2020, and references therein).

We assess the finite sample properties of the MZ and augmented MZ tests using Monte Carlo simulations which confirms that both tests have good finite sample properties across various sample sizes and bootstrap block lengths. We then provide two empirical applications of our methodology. The first application applies the basic MZ test to classical VaR forecasts for S&P 500 returns constructed from a GARCH(1,1) model via the GARCH bootstrap (Pascual et al., 2006). We test jointly over the quantile levels 0.01, 0.025 and 0.05 and horizons from 1 to 10 trading days. Autocalibration is rejected overall and the miscalibration of the forecasts gets stronger for larger forecast horizons and more extreme quantiles. Furthermore, a clear pattern emerges over all quantiles and horizons regarding the conditional quantile bias: the VaR forecasts tend to underestimate risk in calmer times, but overestimate it in more stressful periods.

The second application applies the test in the spirit of the emerging GaR literature, where we focus on the extensions of our test using the augmented MZ test and the test with multiple time series. We expand on the work of Adrian et al. (2019) to formally investigate the performance of simple quantile regression models using financial conditions indicators in predicting a range of U.S. macroeconomic series. Interestingly, we find that the forecasts across four different series and a range of quantile levels and horizon are sub-optimal in that they are not autocalibrated. However, further analysis of the results shows that this sub-optimality is present only in inflation-type series and not in real series like industrial production and employment growth. We also find poorer calibration at the most extreme quantile under consideration.

In relation to the existing literature, this paper extends the work on quantile forecast optimality or, in other words, absolute evaluation of quantile forecasts. The focus of this literature has been on

single-horizon prediction at a single quantile, which mainly stems from the extensive body of research on backtesting VaR, such as Christoffersen (1998), Engle and Manganelli (2004), Escanciano and Olmo (2010, 2011), Gaglianone et al. (2011) and Nolde and Ziegel (2017). Our work also complements the literature on testing the relative forecast performance of conditional quantile models such as Giacomini and Komunjer (2005), Manzan (2015) or more recently Corradi et al. (2021). Finally, as our focus lies on testing for optimality across horizons, the paper also relates to Quaadvlieg (2021), who emphasized the importance of multi-horizon forecast evaluation to avoid multiple testing issues in the context of relative evaluation of mean forecasts. The only work on multi-horizon optimality testing we are aware of is Patton and Timmermann (2012), who consider the case of mean forecasts as well and discuss several implications of optimality specific to the multi-horizon context and how to construct tests for them, most notably the monotonicity of expected loss over horizons, which we extend to the quantile case in the Appendix.

The rest of the paper is organised as follows. Section 2 lays out the notion of quantile forecast optimality that will provide the foundation of our tests. Section 3 then introduces the test for auto-calibration via MZ regression, along with the bootstrap methodology and theory. Section 4 extends the test to the augmented MZ test and the test for multiple variables. Section 5 provides the results of the Monte Carlo study, while Section 6 gives the two empirical applications of our methods. Finally, section 7 concludes the paper. All proofs can be found in the Appendix along with the monotonicity test and some additional empirical results and graphs.

2 Quantile Forecast Optimality

Consider a multivariate stochastic process $\{\mathbf{V}_t\}_{t \in \mathbb{Z}}$, where \mathbf{V}_t is a random vector which contains a response variable of interest y_t and other observable predictors. We denote the forecaster's information set at time t by $\mathcal{F}_t = \sigma(\mathbf{V}_s; s \leq t)$, where $\sigma(\cdot)$ denotes the σ -algebra generated by a set of random variables. Assuming a continuous outcome y_t for the rest of the paper, our target functional is the conditional τ -quantile of y_t given \mathcal{F}_{t-h} :

$$q_t(\tau | \mathcal{F}_{t-h}) = F_{y_t | \mathcal{F}_{t-h}}^{-1}(\tau),$$

where $F_{y_t | \mathcal{F}_{t-h}}(\cdot)$ is the cumulative distribution function of y_t conditional on \mathcal{F}_{t-h} . We denote an h -step ahead forecast at time $t-h$ for this τ -quantile $q_t(\tau | \mathcal{F}_{t-h})$ by $\hat{y}_{\tau,t,h}$, and assume that we observe these forecasts $\hat{y}_{\tau,t,h}$ for each target period t at multiple horizons, $h \in \mathcal{H} = \{1, \dots, H\}$, and multiple quantile levels, $\tau \in \mathcal{T} = \{\tau_1, \dots, \tau_K\}$, for some finite integers K and H . That is, at each time point t we have a matrix of forecasts, $(\hat{y}_{\tau,t,h})_{\tau=\tau_1, \dots, \tau_K, h=1, \dots, H}$. In addition, throughout the paper, we will assume strict stationarity of $\{\mathbf{V}_t\}_{t \in \mathbb{Z}}$ and finite first moments of the forecasts $\hat{y}_{\tau,t,h}$ and y_t itself, see Assumptions A1 and A2 in Section 3.

Since our focus lies on the evaluation of quantile forecasts, the loss function used for evaluation in this context is the 'tick' or 'check' loss which is well-known from quantile regression:

$$L_\tau(y_{t+h} - \hat{y}_{\tau,t,h}) = \rho_\tau(y_{t+h} - \hat{y}_{\tau,t,h}),$$

where $\rho_\tau(u) = u(\tau - 1\{u < 0\})$ and where $1\{\cdot\}$ denotes the indicator function giving a value of one when the expression is true and zero otherwise.

While relative forecast evaluation deals with comparing different forecasting methods or models, mainly by ranking them via their expected loss, the subject of this paper is absolute forecast evaluation across different quantile levels and/or forecasting horizons, in other words the assessment of a *particular* forecasting model or method in terms of absolute evaluation criteria for multiple quantile levels and horizons. These evaluation criteria are usually different forms of optimality (or ‘rationality’/‘calibration’). We start by defining and discussing various forms of quantile forecast optimality before showing how to operationalise the latter for testing.

Definition 1 (Optimality). *An h -step ahead forecast $\hat{y}_{\tau,t,h|\mathcal{I}_{t-h}}^*$ for the τ -quantile is optimal relative to an information set $\mathcal{I}_{t-h} \subset \mathcal{F}_{t-h}$ if:*

$$\hat{y}_{\tau,t,h|\mathcal{I}_{t-h}}^* = \arg \min_{\hat{y}_{\tau,t,h}} \mathbb{E} [L_{\tau}(y_t - \hat{y}_{\tau,t,h}) | \mathcal{I}_{t-h}].$$

We simply call it optimal and denote it by $\hat{y}_{\tau,t,h}^*$ if $\mathcal{I}_{t-h} = \mathcal{F}_{t-h}$, i.e. if it is optimal relative to the full information set:

$$\hat{y}_{\tau,t,h}^* \equiv \hat{y}_{\tau,t,h|\mathcal{F}_{t-h}}^*.$$

Analogous to the case of mean forecasts (Granger, 1969), an optimal quantile forecast relative to an information set can alternatively be characterised as being equal to the respective conditional quantile provided the information set is sufficiently large and includes the forecasts themselves. Specifically, since ‘tick’ loss is a strictly consistent scoring function for the corresponding quantile (see Definition 1 and Proposition 1 in Gneiting, 2011), it holds that an h -step ahead forecast $\hat{y}_{\tau,t,h}$ for the τ -quantile is optimal relative to any information (sub-)set \mathcal{I}_{t-h} satisfying $\sigma(\hat{y}_{\tau,t,h}) \subset \mathcal{I}_{t-h} \subset \mathcal{F}_{t-h}$, where $\sigma(\hat{y}_{\tau,t,h})$ denotes the sigma algebra spanned by the forecast itself, if and only if

$$\hat{y}_{\tau,t,h} = q_{\tau}(\tau | \mathcal{I}_{t-h}).^2 \tag{1}$$

While interest often lies in testing the null hypothesis of (full) optimality relative to the information set \mathcal{F}_{t-h} , which amounts to testing if the forecast, $\hat{y}_{\tau,t,h}$, is equal to its target, $q_{\tau}(\tau | \mathcal{F}_{t-h})$, the possibly large and generally unknown information set \mathcal{F}_{t-h} usually makes direct tests of this hypothesis difficult in practice. Thus, we next discuss weaker forms of optimality that will form the basis of our test(s) in Sections 3 and 4 below. In fact, in Appendix A (see Lemma A.1) we show formally that these weaker forms of optimality may always be viewed as a direct implication of optimality with respect to the ‘full’ information set \mathcal{F}_{t-h} . That is, any h -step ahead forecast optimal with respect to the full information set \mathcal{F}_{t-h} , is also optimal relative to any ‘smaller’ information (sub-)set $\mathcal{I}_t \subset \mathcal{F}_t$.

A special case of this ‘weaker’ form of optimality is optimality with respect to the information contained in the forecast itself, $\sigma(\hat{y}_{\tau,t,h})$, or autocalibration, a term first coined by Tsyplakov (2013) and Gneiting and Ranjan (2013) in the context of probabilistic forecasts.³

Definition 2 (Autocalibration). *An h -step ahead forecast $\hat{y}_{\tau,t,h}$ for the τ -quantile is autocalibrated if it holds that:*

$$\hat{y}_{\tau,t,h} = q_{\tau}(\tau | \sigma(\hat{y}_{\tau,t,h})).$$

²Note that this result continues to hold for any quantile loss from the class of generalised piecewise linear loss functions, or in fact for any consistent scoring function if the τ -quantile is substituted for the corresponding statistical functional for which this scoring function is consistent.

³See also Krüger and Ziegel (2021) and Pohle (2020) for a discussion of autocalibration in the context of mean forecasts and of more general statistical functionals, respectively.

On the one hand, autocalibration may be regarded as a direct implication of full optimality that is particularly suitable for testing as it only relies upon the forecasts themselves and does not require any assumptions on the information set \mathcal{F}_{t-h} or a selection of variables from it. On the other hand, however, autocalibration may also be viewed as a criterion for absolute forecast evaluation in its own right for several reasons. Firstly, the concept has a clear interpretation since a forecast user provided with autocalibrated forecasts should use them as they are and not transform or ‘recalibrate’ them. Secondly, only involving forecasts and observations and no information set that depends on other quantities, it comes closest to the idea of forecast calibration as a concept of consistency between forecasts and observations (see Gneiting et al., 2007). Thirdly, autocalibration might often be a more reasonable criterion to demand from forecasts than full optimality, which is often hard to fulfill in practice. Finally, the Murphy decomposition of expected loss (Pohle, 2020) shows that autocalibration is a fundamental property of forecasts in that expected loss is driven by only two forces: deviations from autocalibration and the information content of the forecasts.

The next section will outline how to test autocalibration, viewed either as an implication of full optimality or as a forecast property in its own right, across multiple quantile levels and horizons simultaneously.

3 Quantile Mincer-Zarnowitz Test

3.1 Null Hypothesis and Quantile Mincer-Zarnowitz Regressions

While autocalibration testing has a long tradition in econometrics through the use Mincer-Zarnowitz regressions for mean forecasts (Mincer and Zarnowitz, 1969), the latter may also be used directly for the case of quantiles (see Gaglianone et al., 2011). Definition 2 in fact suggests that a natural test for autocalibration of an h -step ahead forecast for the τ -level quantile may be based on checking whether, for a given sample of outcomes and quantile forecasts at level τ_k and horizon h , it holds that:

$$q_t(\tau|\widehat{y}_{\tau,t,h}) = \alpha_h^\dagger(\tau_k) + \widehat{y}_{\tau_k,t,h}\beta_h^\dagger(\tau_k) = \widehat{y}_{\tau_k,t,h}$$

almost surely. More generally, since our goal is to test for autocalibration over multiple forecast horizons and quantile levels jointly, we specify such a linear quantile regression model for every horizon $h \in \mathcal{H}$ and $\tau_k \in \mathcal{T}$ as follows:

$$y_t = \alpha_h^\dagger(\tau_k) + \widehat{y}_{\tau_k,t,h}\beta_h^\dagger(\tau_k) + \varepsilon_{t,h}(\tau_k) = \mathbf{X}'_{\tau_k,t,h}\boldsymbol{\beta}_h^\dagger(\tau_k) + \varepsilon_{t,h}(\tau_k), \quad (2)$$

where $\mathbf{X}_{\tau_k,t,h} = (1, \widehat{y}_{\tau_k,t,h})'$ and $\boldsymbol{\beta}_h^\dagger(\tau_k) = (\alpha_h^\dagger(\tau_k), \beta_h^\dagger(\tau_k))'$, and $\mathbb{E}[1\{\varepsilon_{t,h}(\tau_k) \leq 0\} - \tau_k] = 0$ holds by construction. Here, the population coefficient vector $\boldsymbol{\beta}_h^\dagger(\tau_k) = (\alpha_h^\dagger(\tau_k), \beta_h^\dagger(\tau_k))'$ of this linear quantile regression model is defined as:

$$\boldsymbol{\beta}_h^\dagger(\tau_k) = \arg \min_{\mathbf{b} \in \mathcal{B}} \mathbb{E} [\rho_{\tau_k}(y_t - \mathbf{X}'_{\tau_k,t,h}\mathbf{b})]. \quad (3)$$

The composite null hypothesis is given by:

$$H_0^{\text{MZ}} : \{\alpha_h^\dagger(\tau_k) = 0\} \cap \{\beta_h^\dagger(\tau_k) = 1\} \quad (4)$$

for all $h \in \mathcal{H}$ and $\tau_k \in \mathcal{T}$ versus:

$$H_1^{\text{MZ}} : \{\alpha_h^\dagger(\tau_k) \neq 0\} \quad \text{and/or} \quad \{\beta_h^\dagger(\tau_k) \neq 1\} \quad (5)$$

for at least some $h \in \mathcal{H}$ and $\tau_k \in \mathcal{T}$. Testing the null hypothesis in (4) not only yields a multi-horizon, multi-quantile test of autocalibration, but also provides us with an idea about possible deviations from the null. In particular, examining the contributions of single horizons, quantiles or horizon-quantile combinations to the overall test statistic, which will be introduced in Subsection 3.2 below, may also be informative about deviations from autocalibration. Moreover, the empirical counterpart of $q_t(\tau|\widehat{y}_{\tau,t,h}) = \alpha_h^\dagger(\tau_k) + \widehat{y}_{\tau_k,t,h}\beta_h^\dagger(\tau_k)$ may be interpreted as autocalibrated forecasts such that, for a specific value of the forecast $\widehat{y}_{\tau_k,t,h}$, the deviations between the regression line (or recalibrated forecast) and the forecasts themselves, $\widehat{y}_{\tau_k,t,h} - q_t(\tau|\widehat{y}_{\tau,t,h})$ can be interpreted as the quantile version of a conditional bias. The direction and size of this conditional quantile bias informs about deficiencies of forecasts in certain situations, a point that we will come back to and illustrate in the applications in Section 6.

3.2 Test Statistic and Bootstrap

In what follows, assume that we observe an evaluation sample of size P , in other words a scalar-valued time series of observations starting at some point in time $R+1 \in \mathbb{Z}$, $\{y_t\}_{t=R+1}^T$, and a matrix-valued time series of forecasts,

$$\left\{ (\widehat{y}_{\tau,t,h})_{\tau=\tau_1,\dots,\tau_K, h=1,\dots,H} \right\}_{t=R+1}^T$$

Moreover, we may also observe a vector of additional variables \mathbf{Z}_{t-h} from the forecaster's information set \mathcal{F}_{t-h} . We will write this additional sample of vector-valued time series as $\{\mathbf{Z}_t\}_{t=R+1-H}^{T-1}$.

Using the taxonomy of Giacomini and Rossi (2010), forecasts $\widehat{y}_{\tau_k,t,h}$ may stem either from ‘forecasting methods’ or from ‘forecasting models’. In the former case, we are typically without knowledge about the underlying model such as with forecasts from the Survey of Professional Forecasters (SPF), or may use forecasts that depend on parameters estimated in-sample using so-called limited-memory estimators based on a finite rolling estimation window (see Giacomini and White, 2006). In the case of ‘forecasting models’ on the other hand, we need to account for the contribution of estimation uncertainty to the asymptotic distribution of the statistic. However, since the focus of this paper lies on detecting systematic forecasting bias rather than dealing with specific forms of estimation error, we consider the latter only under the recursive scheme with a ‘large’ in sample estimation window. Specifically, for forecasting models, we assume that we also observe R additional observations of y_t prior to $R+1$ that may be used as estimation window (note that the R in-sample observations also comprise H observations that are used to produce the initial out-of-sample forecast for period $R+1$), and that $P/R \rightarrow 0$ as $P, R \rightarrow \infty$. This allows us to abstract from estimation error in the analysis and to focus on systematic features of the forecasts.⁴

The parametric models we consider in this paper take the form $m(\mathbf{W}_{t-h}; \boldsymbol{\theta}_{\tau_k,h}^\dagger)$, where \mathbf{W}_{t-h} denotes a vector of predictor variable(s) and we assume for simplicity that \mathbf{W}_{t-h} is a subset of \mathbf{V}_{t-h} . Moreover, $\boldsymbol{\theta}_{\tau_k,h}^\dagger$ is a population parameter vector that needs to be estimated in a first step, while the function $m(\mathbf{W}_{t-h}; \cdot)$ on the other hand is assumed to be a ‘smooth’ function of the parameter vector

⁴In addition, it also allows us to resample forecasts directly in the bootstrap.

in the sense of Assumption A6 below. For instance, $m(\mathbf{W}_{t-h}; \boldsymbol{\theta}_{\tau_k, h}^\dagger)$ could itself take the form of a linear quantile regression model:

$$q_{t,h}(\tau_k | \mathbf{W}_{t-h}) = m(\mathbf{W}_{t-h}; \boldsymbol{\theta}_{\tau_k, h}^\dagger) = \mathbf{W}'_{t-h} \boldsymbol{\theta}_{\tau_k, h}^\dagger.$$

Alternatively, the model could also take the form of a nonlinear location scale model:

$$q_{t,h}(\tau_k | \mathbf{W}_{t-h}) = m(\mathbf{W}_{t-h}; \boldsymbol{\theta}_{\tau_k, h}^\dagger) = m_\mu(\mathbf{W}_{t-h}; \boldsymbol{\theta}_{h, \mu}^\dagger) + \sigma(\mathbf{W}_{t-h}; \boldsymbol{\theta}_{h, \sigma}^\dagger) q_{t,h, \epsilon}(\tau_k),$$

where $\boldsymbol{\theta}_{\tau_k, h}^\dagger = (\boldsymbol{\theta}_{h, \mu}^{\dagger'}, \boldsymbol{\theta}_{h, \sigma}^{\dagger'}, q_{t,h, \epsilon}(\tau_k))'$ and $q_{t,h, \epsilon}(\tau_k)$ denotes the unconditional τ_k quantile of the error term of the location scale model.

To accommodate both ‘forecasting methods’ and ‘forecasting models’, we adopt a more generic notation in what follows and let $\mathbf{X}_{\tau_k, t, h}(\boldsymbol{\theta}_{\tau_k, h}^\dagger)$ stand either for the vector of stemming from a forecasting method or for the population vector of Mincer-Zarnowitz regressors stemming from a corresponding forecasting model. On the contrary, when forecasts have been generated from a model that has been estimated through the recursive scheme using the first $t-h$ observations with $t = R+1, R+2, \dots$, we write $\mathbf{X}_{\tau_k, t, h}(\widehat{\boldsymbol{\theta}}_{\tau_k, t, h})$ to denote the dependence on the estimated parameter vector $\widehat{\boldsymbol{\theta}}_{\tau_k, t, h}$. Note that even though we focus on the recursive scheme in light of the applications, the rolling estimation scheme whereby a window of the last R observations is used (running from $t-h-R+1$ to $t-h$) or the fixed scheme where the parameter vector is estimated only once, i.e. $\widehat{\boldsymbol{\theta}}_{\tau_k, t, h} = \widehat{\boldsymbol{\theta}}_{\tau_k, R+1-h, h}$, are equally compatible with our set-up and the assumptions below could be adapted in a straightforward manner. Finally, ‘forecasting methods’ can be accommodated by assuming $\widehat{\boldsymbol{\theta}}_{\tau_k, t, h} = \boldsymbol{\theta}_{\tau_k, h}^\dagger$ almost surely.

Thus, to implement the test for the null hypothesis in (4) versus the alternative in (5), we first estimate the coefficient vector as:

$$\widehat{\boldsymbol{\beta}}_h(\tau_k) = \begin{pmatrix} \widehat{\alpha}_h(\tau_k) \\ \widehat{\beta}_h(\tau_k) \end{pmatrix} = \arg \min_{\mathbf{b} \in \mathcal{B}} \frac{1}{P} \sum_{t=R+1}^T \rho_\tau(y_t - \mathbf{X}_{\tau_k, t, h}(\widehat{\boldsymbol{\theta}}_{\tau_k, t, h})' \mathbf{b}) \quad (6)$$

for each h and τ_k . With these estimates at hand, different possibilities to construct a suitable test statistic exist. More specifically, since the number of elements in \mathcal{H} and \mathcal{T} is finite, one option is to construct a Wald-type test based on the estimates in (6) together with a suitable estimator of the variance-covariance matrix. However, when interest lies in testing (4) against (5) for a larger number of quantile levels and horizons, constructing a Wald test involves estimating a large variance-covariance matrix, which can be difficult in practice and may lead to a poor finite sample performance. On the other hand, as we argue below, using a moment equality based test in combination with the nonparametric bootstrap does not suffer from this drawback. In fact, the moment equality framework can be extended easily to other set-ups which give rise to an even larger number of equalities (see Section 4).

To see the possibility of a moment equality based test, note that under H_0^{MZ} and the Assumptions A1 to A7 outlined below it holds that:

$$\sqrt{P} \left(\widehat{\boldsymbol{\beta}}_h^\dagger(\tau_k) - \begin{pmatrix} \alpha_h^\dagger(\tau_k) \\ \beta_h^\dagger(\tau_k) \end{pmatrix} \right) \xrightarrow{d} N \left(\mathbf{0}, \tau_k(1-\tau_k) \mathbf{J}_h(\tau_k)^{-1} \mathbf{E} \left[\mathbf{X}_{\tau_k, t, h}(\boldsymbol{\theta}_{\tau_k, h}^\dagger) \mathbf{X}_{\tau_k, t, h}(\boldsymbol{\theta}_{\tau_k, h}^\dagger)' \right] \mathbf{J}_h(\tau_k)^{-1} \right)$$

pointwise in h and τ_k , where the matrix $\mathbf{J}_h(\tau_k)$ is given by:

$$\mathbf{J}_h(\tau_k) \equiv \mathbb{E} \left[f_{t,h} \left(\mathbf{X}_{\tau_k,t,h}(\boldsymbol{\theta}_{\tau_k,h}^\dagger)' \boldsymbol{\beta}^\dagger(\tau_k) \right) \mathbf{X}_{\tau_k,t,h}(\boldsymbol{\theta}_{\tau_k,h}^\dagger) \mathbf{X}_{\tau_k,t,h}(\boldsymbol{\theta}_{\tau_k,h}^\dagger)' \right] \quad (7)$$

and $f_{t,h}(\cdot)$ is defined in Assumption A4. In fact, under H_0^{MZ} , in Appendix B we establish the linear Bahadur representation:

$$\begin{aligned} & \sqrt{P} \left(\widehat{\boldsymbol{\beta}}_h(\tau_k) - \begin{pmatrix} 0 \\ 1 \end{pmatrix} \right) \\ &= \mathbf{J}_h(\tau_k)^{-1} \left(\frac{1}{\sqrt{P}} \sum_{t=R+1}^T \mathbf{X}_{\tau_k,t,h}(\boldsymbol{\theta}_{\tau_k,h}^\dagger) \left(1 \{ y_t \leq \mathbf{X}_{\tau_k,t,h}(\boldsymbol{\theta}_{\tau_k,h}^\dagger)' \boldsymbol{\beta}_{\tau_k,h}^\dagger(\tau_k) \} - \tau_k \right) \right) + o_{\text{Pr}}(1). \end{aligned}$$

The above representation motivates the use of a moment equality type statistic for a test of autocalibration. Thus, define the set:

$$\mathcal{C}^{\text{MZ}} = \{(h, \tau_k, j) : h \in \mathcal{H}, \tau_k \in \mathcal{T}, j \in \{1, 2\}\},$$

and let $|\mathcal{C}^{\text{MZ}}| = \kappa$ denote the cardinality of \mathcal{C}^{MZ} , while $s = 1, \dots, \kappa$ is a generic element from \mathcal{C}^{MZ} . For the test statistic, define \widehat{m}_s either as $\widehat{\alpha}_h(\tau_k)$ or as $(\widehat{\beta}_h(\tau_k) - 1)$ for a specific τ_k and h . A test statistic for the null hypothesis in (4) is then given by:

$$\widehat{U}_{\text{MZ}} = \sum_{s=1}^{\kappa} \left(\sqrt{P} \widehat{m}_s \right)^2. \quad (8)$$

Note that the non-studentised statistic in (8) above does not require an estimation of the asymptotic variance, and consequently will be non-pivotal as its asymptotic distribution does depend on the full variance-covariance matrix. That is, under H_0^{MZ} and conditions A1 to A7 below:

$$\sqrt{P} \left(\begin{pmatrix} \widehat{\alpha}_1(\tau_1) \\ \widehat{\beta}_1(\tau_1) \\ \vdots \\ \widehat{\alpha}_H(\tau_K) \\ \widehat{\beta}_H(\tau_K) \end{pmatrix} - \begin{pmatrix} 0 \\ 1 \\ \vdots \\ 0 \\ 1 \end{pmatrix} \right) \xrightarrow{d} N(\mathbf{0}, \boldsymbol{\Sigma}), \quad (9)$$

where $\boldsymbol{\Sigma}$ is the variance-covariance matrix, which is unknown in practice and depends on features of the data generating process (DGP). Of course, this nuisance parameter problem can be taken into account by using a suitable bootstrap procedure (e.g., White, 2000; Hansen et al., 2011). In particular, we will generate bootstrap critical values using the moving block bootstrap (MBB) of Künsch (1989), whose formal validity for quantile regression with time series observations has recently been established by Gregory et al. (2018). In doing so, we will resample the forecasts directly as the limiting distribution will be derived under the condition that $P/R \rightarrow \pi = 0$, implying that forecast estimation error does not feature into the asymptotic distribution of the test statistic (cf. West, 1996). We do so to abstract from the dependence on a particular estimator, which would require further details about the underlying forecasting models.

We generate bootstrap samples of length P consisting of K_b blocks of length l such that $P = K_b l$.

We draw the starting index I_j of each block $1, \dots, K_b$, $\{I_j, I_{j+1}, \dots, I_{j+l}\}$, from a discrete random uniform distribution on $[R+1, T-l]$. These indices are used to resample from $\{y_t, \widehat{y}_{\tau,t,h}\}_{t=R+1}^T$ jointly for each $\tau = \tau_1, \dots, \tau_K$ and $h = 1, \dots, H$. This way we generate B bootstrap samples, each with $\{y_t^b, \widehat{y}_{\tau,t,h}^b\}_{t=R+1}^T$ for all $\tau = \tau_1, \dots, \tau_K$ and $h = 1, \dots, H$. For each bootstrap sample, we construct bootstrap equivalents of (6) and then the corresponding bootstrap statistic:

$$\widehat{U}_{\text{MZ}}^b = \sum_{s=1}^{\kappa} \left(\sqrt{P}(\widehat{m}_s^b - \widehat{m}_s) \right)^2. \quad (10)$$

The critical value is then given by the $(1 - \alpha)$ quantile of the empirical bootstrap distribution of $\widehat{U}_{\text{MZ}}^b$ over B draws, say $c_{B,P,(1-\alpha)}$.

3.3 Assumptions and Asymptotic Validity

For the asymptotic validity of this procedure, we make the following assumptions:

A1: The outcome variable y_t is strictly stationary, satisfies $\mathbb{E}[|y_t|] < \infty$ for all t , and is α -mixing with α -mixing coefficients satisfying $\sum_{k=1}^{\infty} \alpha(k)^{\frac{\epsilon}{2+\epsilon}} < \infty$, where ϵ is defined in A2.

A2: For all $h \in \mathcal{H}$ and $\tau_k \in \mathcal{T}$, and any $\boldsymbol{\theta} \in \boldsymbol{\Theta}$, it holds that $\mathbb{E} \left[\|\mathbf{X}_{\tau_k,t,h}(\boldsymbol{\theta})\|^{4+\epsilon} \right] < \infty$, $\epsilon > 0$, where $\|\cdot\|$ denotes the Euclidean norm and $\boldsymbol{\Theta}$ is defined in A6 below. The distribution of $\mathbf{X}_{\tau_k,t,h}(\boldsymbol{\theta})$ is absolutely continuous with Lebesgue density.

A3: For every $\tau_k \in \mathcal{T}$ and $h \in \mathcal{H}$, assume that the parameter space of $\boldsymbol{\beta}_h(\tau_k)$, \mathcal{B} , is a compact and convex set. The coefficient vector $\boldsymbol{\beta}_h^\dagger(\tau_k)$ from (3) lies in the interior of \mathcal{B} for all values $\tau_k \in \mathcal{T}$ and $h \in \mathcal{H}$.

A4: For all $h \in \mathcal{H}$, the conditional distribution function of y_t (given \mathcal{F}_{t-h}), $F_{t,h}(\cdot) \equiv F_t(\cdot | \mathcal{F}_{t-h})$, admits a continuous Lebesgue density, $f_{t,h}(\cdot) \equiv f_t(\cdot | \mathcal{F}_{t-h})$, which is bounded away from zero and infinity for all u in $\mathcal{U} = \{u : 0 < F_{t,h}(u) < 1\}$. The density $f_{t,h}(\cdot)$ is integrable uniformly over \mathcal{U} .

A5: For all $h \in \mathcal{H}$ and $\tau_k \in \mathcal{T}$, assume that the matrix $\mathbf{J}_h(\tau_k)$ defined in (7) is positive definite.

A6: Assume that $\boldsymbol{\Theta}$ is compact and that, for each $\tau_k \in \mathcal{T}$ and $h \in \mathcal{H}$, $\boldsymbol{\theta}_{\tau_k,h}^\dagger$ lies in its interior. For all $\boldsymbol{\theta}_1, \boldsymbol{\theta}_2 \in \boldsymbol{\Theta}$, it holds uniformly over $t \geq R+1$ that:

$$\|\mathbf{X}_{\tau_k,t,h}(\boldsymbol{\theta}_1) - \mathbf{X}_{\tau_k,t,h}(\boldsymbol{\theta}_2)\| \leq B(\mathbf{X}_{\tau_k,t,h}) \|\boldsymbol{\theta}_1 - \boldsymbol{\theta}_2\|$$

for some positive, real-valued function $B(\mathbf{X}_{\tau_k,t,h})$ satisfying $\mathbb{E}[\sup_{t \geq R+1} B(\mathbf{X}_{\tau_k,t,h})] < \infty$. In addition, assume that for every $h \in \mathcal{H}$ and $\tau_k \in \mathcal{T}$, the estimator $\widehat{\boldsymbol{\theta}}_{\tau_k,t,h}$ satisfies:

$$\sup_{t \geq R+1} \|\widehat{\boldsymbol{\theta}}_{\tau_k,t,h} - \boldsymbol{\theta}_{\tau_k,h}^\dagger\| = O_{\text{Pr}} \left(\frac{1}{\sqrt{R}} \right).$$

A7: Assume that $R, P, l \rightarrow \infty$ as $T \rightarrow \infty$ with $P/R \rightarrow \pi = 0$ and $l/P \rightarrow 0$.

Assumption A1 imposes some mild restrictions on the time dependence of the data that are in turn linked to the existence and finiteness of corresponding moments in A2. On the other hand, the continuity of $\mathbf{X}_{\tau_k,t,h}(\boldsymbol{\theta})$ for any given τ_k , h , and $\boldsymbol{\theta}$ only serves the purpose to simplify some of the arguments in the Appendix, and could be relaxed at the expense of more cumbersome notation.

Assumptions A3-A5 are required to derive the limiting distribution of the linear quantile regression estimator (e.g., Koenker and Xiao, 2006; Galvao et al., 2011). In fact, A3 and A4 represent standard assumptions on the parameter space and the smoothness of the (conditional) distribution of y_t , while A5 ensures asymptotic normality of the quantile regression estimator. Assumption A6 on the other hand is only needed when the focus lies on forecasting models. Specifically, it places restrictions on the underlying parametric models, but is in fact compatible with many commonly used nonlinear location scale (e.g., GARCH) or linear quantile regression models that satisfy the Lipschitz condition in A6 and that can be estimated at rate \sqrt{R} . Finally, Assumption A7 governs the rates at which P and R as well as the block length l may grow to infinity. In particular and in analogy to West (1996), we require $\pi = 0$ for estimation error to be ignorable asymptotically and to be able to resample directly from the forecasts (rather than to resample from the realised predictors). In turn, this allows us to focus on miscalibration as a structural feature of the models.

We are now ready to derive the asymptotic properties of the statistic under the null hypothesis:

Theorem 1. *Assume that A1 to A7 hold. Then under H_0^{MZ} :*

$$\lim_{T, B \rightarrow \infty} \Pr \left(\widehat{U}_{MZ} > c_{B, P, (1-\alpha)} \right) = \alpha.$$

Theorem 1 establishes the asymptotic size control of the moment equality test. It is easy to implement using the standard block bootstrap, and, unlike the Wald test, can be modified to accommodate more complex set-ups with a larger number of moment equalities that we will outline next.

4 Extensions

In the following two subsections, we will describe two extensions of the test from Section 3, firstly to accommodate additional predictors \mathbf{Z}_{t-h} in the Mincer-Zarnowitz regression to test different forms of optimality, and secondly to test for autocalibration of quantile forecasts for multiple time series simultaneously.

4.1 Augmented Quantile Mincer-Zarnowitz Test

Recalling the characterisation of optimality relative to any information set $\mathcal{I}_{t-h} \subset \mathcal{F}_{t-h}$ from (1), it becomes clear that the Mincer-Zarnowitz set-up from the previous section may also be used to test stronger forms of optimality with respect to larger information sets than $\sigma(\widehat{y}_{\tau_k, t, h})$. More precisely, while H_0^{MZ} vs. H_1^{MZ} is a test of autocalibration, it does not check if all available valuable information from \mathcal{F}_{t-h} was incorporated into the forecasting model or taken into account by the forecaster. We therefore suggest the idea of augmented quantile Mincer-Zarnowitz regressions, where a vector of additional regressors $\mathbf{Z}_{t-h} \in \mathcal{F}_{t-h}$ is added to the regression model in (2) to test for optimality relative to $\sigma(\widehat{y}_{\tau_k, t, h}, \mathbf{Z}_{t-h})$, see also Elliott and Timmermann (2016, chapter 15) for a discussion of augmented Mincer-Zarnowitz regressions in the context of mean forecasts.

That is, in analogy to the previous section, we again specify a linear quantile regression model for every horizon $h \in \mathcal{H}$ and $\tau_k \in \mathcal{T}$ as follows:

$$y_t = \alpha_h^\dagger(\tau_k) + \widehat{y}_{\tau_k, t, h} \beta_h^\dagger(\tau_k) + \mathbf{Z}'_{t-h} \boldsymbol{\gamma}_h^\dagger(\tau_k) + \varepsilon_{t, h}(\tau_k), \quad (11)$$

where with slight abuse of notation we use the same symbols as in the previous section for the first two regression coefficients as well as the error term, and suppress the possible dependence of the forecasts on some parameter vector $\boldsymbol{\theta}^\dagger$. If the coefficients of \mathbf{Z}_{t-h} in the population augmented Mincer-Zarnowitz regression are non-zero, i.e. $\boldsymbol{\gamma}_h^\dagger \neq \mathbf{0}$, there is valuable information in \mathbf{Z}_{t-h} that has not been incorporated into the forecasts yet. As a result, those variables or a subset thereof should be included into the model to improve forecast accuracy.

Formally, the null hypothesis we test is given by:

$$H_0^{\text{AMZ}} : \{\alpha_h^\dagger(\tau_k) = 0\} \cap \{\beta_h^\dagger(\tau_k) = 1\} \cap \{\boldsymbol{\gamma}_h^\dagger(\tau_k) = \mathbf{0}\} \quad (12)$$

for all $h \in \mathcal{H}$ and $\tau_k \in \mathcal{T}$ versus:

$$H_1^{\text{AMZ}} : \{\alpha_h^\dagger(\tau_k) \neq 0\} \quad \text{and/or} \quad \{\beta_h^\dagger(\tau_k) \neq 1\} \quad \text{and/or} \quad \{\boldsymbol{\gamma}_h^\dagger(\tau_k) \neq \mathbf{0}\} \quad (13)$$

for at least some $h \in \mathcal{H}$ and $\tau_k \in \mathcal{T}$.

In contrast to a standard Mincer-Zarnowitz quantile regression, the augmented version requires a choice of variables from \mathcal{F}_{t-h} . These variables included in \mathbf{Z}_{t-h} have to be chosen *a priori* and may, in some situations, suggest themselves naturally as illustrated in our macroeconomic application. In other situations, however, it might be hard to pick those variables from a potentially very large information set. In those cases, regularised or factor-augmented Mincer-Zarnowitz regressions could be used instead of (11). We leave this extension to future research and state the asymptotic size result of a moment equality based test of H_0^{AMZ} vs. H_1^{AMZ} with corresponding test statistic, say \widehat{U}_{AMZ} , as a direct corollary of Theorem 1.

Corollary 1. *Assume that A1 to A7 hold. Then under H_0^{AMZ} :*

$$\lim_{T, B \rightarrow \infty} \Pr \left(\widehat{U}_{\text{AMZ}} > c_{B, P, (1-\alpha)} \right) = \alpha.$$

4.2 Multivariate Quantile Mincer-Zarnowitz Test

While Section 3 establishes a test for autocalibration of quantile forecasts of a single time series y_t , researchers may sometimes be interested in testing this property across several time series $\mathbf{y}_t = (y_{1,t}, \dots, y_{G,t})'$ using an array of h step ahead τ_k -level forecasts $(\widehat{\mathbf{y}}_{\tau,t,h})_{\tau=\tau_1, \dots, \tau_K, h=1, \dots, H}$, where $\widehat{\mathbf{y}}_{\tau,t,h} = (\widehat{y}_{1,\tau_k,t,h}, \dots, \widehat{y}_{G,\tau_k,t,h})'$, where $h \in \mathcal{H}$, $\tau_k \in \mathcal{T}$, and finite $G \in \mathbb{N}$. For instance, we may be interested in testing for forecast autocalibration jointly across different industries or sectors; components of GDP growth like consumption or export growth; or different macro series like in our application in Section 6.2. It is important to distinguish this situation with many univariate quantile forecasts issued by the same forecaster or institution, which often arises in practice, from the situation where truly multivariate forecasts are issued in the form of a multidimensional density or CDF (see Knüppel et al., 2022).

We now sketch the extension of the autocalibration test to such a multivariate set-up. To this end, consider again the following linear quantile regression model:

$$y_{i,t} = \alpha_{i,h}(\tau_k) + \widehat{y}_{i,\tau_k,t,h} \beta_{i,h}(\tau_k) + \varepsilon_{i,t,h}(\tau_k), \quad h \in \mathcal{H}, \quad \tau_k \in \mathcal{T}, \quad i = 1, \dots, G. \quad (14)$$

Here, the coefficient vector $\beta_{h,G}(\tau_k) = (\alpha_{1,h}(\tau_k), \dots, \alpha_{G,h}(\tau_k), \beta_{1,h}(\tau_k), \dots, \beta_{G,h}(\tau_k))'$ is defined as:

$$\beta_{h,G}(\tau_k) = \arg \min_{\mathbf{b}_G \in \mathcal{B}} \mathbb{E} \left[\sum_{i=1}^G \left(\rho_{\tau_k} \left(y_{i,t} - \mathbf{X}_{i,\tau_k,t,h}(\boldsymbol{\theta}_{i,\tau_k,h}^\dagger)' \mathbf{b}_G \right) \right) \right] \quad (15)$$

with $\mathbf{X}_{i,\tau_k,t,h}(\boldsymbol{\theta}_{i,\tau_k,h}^\dagger)' = (1\{i=1\}, \dots, 1\{i=G\}, 1\{i=1\} \times \widehat{y}_{i,\tau_k,t,h}, \dots, 1\{i=G\} \times \widehat{y}_{i,\tau_k,t,h})'$.⁵ Therefore, the set-up in (14) allows for miscalibration also at the individual time series level (e.g., industry or sector) if for some i, h , and τ_k it holds that $\alpha_{i,h}(\tau_k) \neq 0$ and/or $\beta_{i,h}(\tau_k) \neq 1$, respectively.

The sample analogue of (15), using the evaluation sample of observations $\{y_t\}_{t=R+1}^T$ and array-valued forecasts

$$\left\{ (\widehat{\mathbf{y}}_{\tau,t,h})_{\tau=\tau_1, \dots, \tau_K, h=1, \dots, H} \right\}_{t=R+1}^T$$

is given by:

$$\widehat{\beta}_{h,G}(\tau_k) = \arg \min_{\mathbf{b}_G \in \mathcal{B}} \frac{1}{P} \sum_{s=R+1}^T \sum_{i=1}^G \left(\rho_{\tau} \left(y_{i,s} - \mathbf{X}_{i,\tau_k,s,h}(\widehat{\boldsymbol{\theta}}_{i,\tau_k,s,h})' \mathbf{b}_G \right) \right). \quad (16)$$

As in Section 3, we are interested in testing the composite null hypothesis:

$$H_0^{\text{MMZ}} : \{\alpha_{i,h}(\tau_k) = 0\} \cap \{\beta_{i,h}(\tau_k) = 1\} \quad (17)$$

for all $h \in \mathcal{H}$, $\tau_k \in \mathcal{T}$, $i = 1, \dots, G$, versus:

$$H_1^{\text{MMZ}} : \{\alpha_{i,h}(\tau_k) \neq 0\} \quad \text{and/or} \quad \{\beta_{i,h}(\tau_k) \neq 1\} \quad (18)$$

for at least some h , τ_k , and i . Note that, since G is finite, for a given horizon h and quantile level τ_k (with A1-A7 adapted to hold for $y_{i,t}$ and $\mathbf{X}_{i,\tau_k,t,h}(\boldsymbol{\theta}_{i,\tau_k,h}^\dagger)$, $i = 1, \dots, G$), the estimator in (16) is consistent for $\beta_{h,G}(\tau_k)$. The test statistic for the null hypothesis in (17) versus (18) is therefore given by:

$$\widehat{U}_{\text{MMZ}} = \sum_{s=1}^{\kappa} \left(\sqrt{P} \widehat{m}_{s,G} \right)^2,$$

where either $\widehat{m}_{s,G} = \widehat{\alpha}_{i,h}(\tau_k)$ or $\widehat{m}_{s,G} = \widehat{\beta}_{i,h}(\tau_k) - 1$, respectively, and in a similar way as above, κ denotes the total number of moment conditions which in this case accounts for the G different series being used in the test.

In order to construct a suitable bootstrap statistic in analogue to Subsection 3.2, we construct bootstrap analogues $\widehat{\beta}_{h,G}^b(\tau_k)$ of (16) from bootstrap samples of length $P = K_b l$ from K_b blocks of length l by resampling again from the series of forecasts-observation pairs, where the forecasts are array-valued. The bootstrap procedure does not just take the horizon, but also the group structure as given, which ensures that the dependence of the original data across horizons h as well as across series $i = 1, \dots, G$ is maintained. More precisely, we again draw the starting index I_j of each block of forecasts and observations $1, \dots, K_b$ from a discrete random uniform distribution on $[R+1, T-l]$. These indices are used to resample from $\left\{ y_t, (\widehat{\mathbf{y}}_{\tau,t,h})_{\tau=\tau_1, \dots, \tau_K, h=1, \dots, H} \right\}_{t=R+1}^T$. This way we generate B bootstrap samples, each with $\left\{ y_t^b, (\widehat{\mathbf{y}}_{\tau,t,h}^b)_{\tau=\tau_1, \dots, \tau_K, h=1, \dots, H} \right\}_{t=R+1}^T$.

⁵Note that we use $\boldsymbol{\theta}_i^\dagger$ to denote the possible dependence of the forecast for series i on a forecasting model.

We then construct a corresponding bootstrap estimator given by:

$$\widehat{\beta}_{h,G}^b(\tau_k) = \arg \min_{b_G \in \mathcal{B}} \frac{1}{P} \sum_{s=R+1}^T \sum_{i=1}^G \left(\rho_\tau \left(y_{i,s}^b - \mathbf{X}_{i,\tau_k,s,h}(\widehat{\theta}_{i,\tau_k,s,h}^{b'}) \mathbf{b}_G \right) \right).$$

The final bootstrap statistic becomes:

$$\widehat{U}_{MMZ}^b = \sum_{s=1}^{\kappa} \left(\sqrt{T}(\widehat{m}_{s,G}^b - \widehat{m}_{s,G}) \right)^2,$$

where $\widehat{m}_{s,G}^b$ is equal to $\widehat{\alpha}_{i,h}^b(\tau_k)$ or $\widehat{\beta}_{i,h}^b(\tau_k) - 1$, respectively. Constructing critical values on the basis of \widehat{U}_{MMZ}^b , $b = 1, \dots, B$, as in Subsection 3.2, the following corollary holds:

Corollary 2. *Assume that Assumptions A1 to A7 hold for every $i = 1, \dots, G$, with y_t , $\mathbf{X}_{\tau_k,t,h}(\theta_{\tau_k,h}^\dagger)$, and $\widehat{\theta}_{\tau_k,t,h}$ replaced by $y_{i,t}$, $\mathbf{X}_{i,\tau_k,t,h}(\theta_{i,\tau_k,h}^\dagger)$, and $\widehat{\theta}_{i,\tau_k,t,h}$, respectively. Then, under H_0^{MMZ} :*

$$\lim_{T,B \rightarrow \infty} \Pr \left(\widehat{U}_{MMZ} > c_{B,P,(1-\alpha)} \right) = \alpha.$$

5 Monte Carlo Simulations

In this section we will present results to explore the finite sample properties of the Mincer-Zarnowitz and augmented Mincer-Zarnowitz tests described above using a variety of set-ups, DGPs and sample sizes.

5.1 Mincer-Zarnowitz Simulations

We first develop a set-up which is designed to assess the size and power of the Mincer-Zarnowitz test. Here we explore a case where the target variable is generated according to a simple AR(1) process. The forecasts, on the other hand, will be computed using a potentially incorrect value for the AR(1) parameter which will drive cases of mis-calibration. Specifically, the DGP for y_t is:

$$y_t = by_{t-1} + \varepsilon_t \tag{19}$$

where we generate ε_t as an i.i.d. normal variable with mean zero and variance $1 - b^2$ which fixes the variance of y_t to unity. The optimal h -step ahead quantile forecast is therefore $q_{t+h}(\tau|\mathcal{F}_t) = b^h y_t + \sqrt{1 - b^{2h}} \Phi^{-1}(\tau)$.

The actual quantile forecasts for quantile level τ are generated with the same functional form as the optimal forecast but using a (potentially) incorrect value \tilde{b} for the AR(1) parameter instead of b . In other words:

$$\widehat{q}_{t+h}^{inc}(\tau) = \tilde{b}^h y_t + \sqrt{1 - \tilde{b}^{2h}} \Phi^{-1}(\tau). \tag{20}$$

In this specification, the forecasts are autocalibrated when we have $\tilde{b} = b$ whereas we have miscalibration when $\tilde{b} \neq b$. This gives us a simple way to assess the size and power of the Mincer-Zarnowitz type test: in the former case we expect the slope coefficient in the Mincer-Zarnowitz regression (2) to be equal to unity and we are under the null hypothesis in (4), whereas in the latter case it is not

equal to unity and we are under the alternative hypothesis in (5).⁶ We will use values of $b = 0.6$ and $\tilde{b} = \{0.6, 0.8\}$, where the case with $\tilde{b} = 0.6$ allows us to assess the size performance and $\tilde{b} = 0.8$ will assess power.

We perform the simulations for a range of evaluation sample sizes $P \in \{120, 240, 480\}$ and we let there be a total of $H = 4$ horizons. We will generate a sample of size $P + H + 1$ for this AR(1) case which includes the single initial condition (drawn from the stationary distribution of the AR(1) process) and allows for up to H -step ahead forecasts to be made for all P periods. We will consider multiple quantile levels $\mathcal{T} = \{0.25, 0.5, 0.75\}$. For the bootstrap we will use various block lengths $l \in \{4, 8, 12\}$ and we will perform $B = 1999$ Monte Carlo replications, saving a single bootstrap draw each time as in the Warp Speed bootstrap of Giacomini et al. (2013). We compute rejection rates for a nominal size of $\alpha = 0.05$.

We now present the size and power results for the moment equality Mincer-Zarnowitz test in (8). Table 1 presents the rejection rates of the test with \tilde{b} set to 0.6 which assesses the size of the test. The results are displayed for the different bootstrap block lengths and sample sizes discussed above. Then the power of the test is displayed in Table 2 for the case when $\tilde{b} = 0.8$ which results in mis-calibration.

Table 1: MZ Test - Size - $\tilde{b} = 0.6$

	$P = 120$	$P = 240$	$P = 480$
$l = 4$	0.037	0.051	0.055
$l = 8$	0.053	0.038	0.034
$l = 12$	0.039	0.044	0.045

Table 2: MZ Test - Power - $\tilde{b} = 0.8$

	$P = 120$	$P = 240$	$P = 480$
$l = 4$	0.792	0.970	1.000
$l = 8$	0.747	0.959	1.000
$l = 12$	0.740	0.961	1.000

The results in Table 1 show that the test has fairly good size properties for a reasonable block length selection. For instance, we see rejection rates very close to nominal size (5.1% and 5.5%) for block length $l = 4$ for the two larger samples size $P = 240$ and $P = 480$. Even in smaller samples the size is fairly close to nominal size. Regarding the power of the test, Table 2 shows very good power properties with rejection rates roughly 75% or above, even at the lowest sample size $P = 120$. These rejection rates rise to above 95% as the sample size increases to $P = 240$ and are equal to unity when $P = 480$.

5.2 Augmented Mincer-Zarnowitz Simulations

In order to display the properties of the augmented Mincer-Zarnowitz test, we take the simplest possible case where \mathbf{Z}_t comprises a single variable z_t which is to be included in the regression described

⁶Note that an analytical expression for the Mincer-Zarnowitz coefficients is simple to derive in the case of making mean predictions. In that case, with MSFE being the criterion used to assess the forecasts, it can be shown that the slope of the Mincer-Zarnowitz regression is $(\beta/b)^h$ which is different from unity when $b \neq \beta$. An analogous result will hold in the quantile case.

in (11). We also generate this variable according to an AR(1) process:

$$z_t = dz_{t-1} + \nu_t \quad (21)$$

where ν_t is similarly drawn from an i.i.d. normal distribution with mean zero and variance $1 - d^2$, and for simplicity we set the AR parameter d equal to that in (19), in other words $d = b$.

We can easily assess the size of the augmented MZ test if we let (19) be the true DGP and we use the augmented MZ regression (11) with the additional regressor z_t which will enter the MZ equation with a coefficient of zero. If we use the forecasts described in (20) with $\tilde{b} = b = 0.6$ we are therefore under the null hypothesis in (12). On the other hand, to display the power of the augmented MZ test, we need to use an additional DGP for y_t which involves z_t and a set of forecasts which omit this additional variable. As such we propose an ADL(1,1) process for y_t :

$$y_t = by_{t-1} + cz_{t-1} + \varepsilon_t \quad (22)$$

where we use the same parameter values and random draws b and ε_t as in (19) with the addition of z_t which is drawn according to (21) and we set $c = 0.5$. If we make the forecasts in the same way as (20) above, which only makes use of the variable y_t then the coefficient on z_t in the augmented MZ regression is non-zero which violates the condition in the null in (12). For the \tilde{b} coefficient used in making the forecast, we use the projection coefficient of y_t on y_{t-1} , which leads to autocalibrated forecasts, that is, the null of the basic MZ test is not violated. The null of the augmented MZ test is, however, violated, and both coefficients (β and γ) differ from their values under the null in (12). For the parameter values described above, the projection coefficient of y_t on y_{t-1} under (22) is roughly $\tilde{b} = 0.70$.

The results for size and power are displayed in Tables 3 and 4 below. The size results show that the augmented MZ test has good size properties with rejection rates fairly close to the nominal size of 5% across sample sizes and block lengths. This shows that the augmented version of the test, indeed, has similar size properties to the standard MZ test. For the power, we also see results in line with what is expected. In Table 4 we see that power is above 60% in the small sample size $P = 120$. This rapidly improves to above 95% as the sample size rises to $P = 240$ and to unity when $P = 480$.

6 Empirical Applications

In this section, we provide two empirical applications. The first is a finance application applying the MZ test to test the optimality of GARCH predictions of the tail quantiles of financial returns. The second application uses the MZ test extensions to assess the optimality of GaR forecasts made across a range of U.S. macroeconomic variables.

6.1 Empirical Application 1: Financial Returns

Forecasts of lower tail quantiles of returns distributions (or upper tail quantiles of loss distributions) play an important role in financial risk management as the most prominent risk measures are either themselves tail quantiles or defined in terms of tail quantiles (see He et al. (2022) for a recent overview).

Table 3: Augmented MZ Test - Size - $\tilde{b} = 0.6$

	$P = 120$	$P = 240$	$P = 480$
$l = 4$	0.053	0.058	0.052
$l = 8$	0.034	0.055	0.058
$l = 12$	0.039	0.038	0.052

Notes: DGP for size is equation (19).

Table 4: Augmented MZ Test - Power - $\tilde{b} = 0.70$

	$P = 120$	$P = 240$	$P = 480$
$l = 4$	0.667	0.979	1.000
$l = 8$	0.646	0.983	1.000
$l = 12$	0.638	0.976	1.000

Notes: DGP for power is equation (22).

For instance, the VaR at level τ is just the τ -quantile of the return distribution y_t ,

$$VaR(\tau) = q(\tau),$$

where usually τ is chosen to be either 0.05 or 0.01. Expected shortfall, which is a coherent risk measure and more sensitive to the shape of the tail distribution beyond the VaR, is just the expectation of this tail distribution and can be written as an integral over all quantiles with levels below τ :

$$ES(\tau) = \frac{1}{\tau} \int_0^\tau VaR(\alpha) d\alpha.$$

Median shortfall, a natural alternative to expected shortfall, is the median of the tail distribution and just a quantile itself:

$$MS(\tau) = VaR\left(\frac{\tau}{2}\right).$$


Producing and backtesting VaR forecasts is a central task in financial risk management. Therefore, as discussed above, the majority of contributions to quantile forecast optimality testing was motivated by this problem. The evaluation of expected shortfall forecasts is more difficult as they cannot be evaluated on their own due to the non-identifiability and non-elicitability of expected shortfall (Fissler and Ziegel, 2016; Fissler et al., 2015). Bayer and Dimitriadis (2022) exploit the joint identifiability of expected shortfall and VaR to jointly test their optimality for a given quantile level and horizon, while Kratz et al. (2018) propose to test optimality of expected shortfall forecasts implicitly by testing forecasts for multiple quantiles of the tail distribution for a given horizon. The latter is possible naturally in our multi-quantile evaluation framework, and so our tests can give guidance regarding the optimality of forecasts for all three risk measures. In addition, note that risk management requires forecasts of risk measures over multiple horizons, e.g. for one day ahead and cumulative losses over the next ten trading days. Nevertheless, extant evaluation methods focus on a single horizon and consequently one-day-ahead forecasts are usually evaluated. Our tests solve this problem as they enable joint evaluation of risk measure forecasts over multiple horizons.

To illustrate the use of the Mincer-Zarnowitz test for the evaluation of forecasts of financial risk

measures, we apply it to multi-horizon, multi-quantile forecasts for daily S&P 500 returns. We consider horizons from $h = 1$ through $h = 10$ in terms of trading days and three quantile levels $\tau \in \{0.01, 0.025, 0.05\}$. The classic model for return volatility and VaR forecasting is the GARCH(1,1) model (Bollerslev, 1986). As no closed-form formula for multi-period-ahead GARCH quantile forecasts is available, except for the case of Gaussian innovations (Abadir et al., 2022), we use the GARCH bootstrap of Pascual et al. (2006). It draws standardised residuals from the estimated one-period-ahead model to simulate draws multiple periods in the future, from which quantiles can be obtained.⁷ We employ 1000 bootstrap draws for the generation of each forecast and choose student- t errors for the estimation of the model.

Our sample consists of daily S&P 500 returns from January 3rd 2000 to June 27th 2022, amounting to 5634 observations.⁸ We use recursive pseudo-out-of-sample forecasting with an initial estimation window of size 3000 for $h = 10$ or, in other words, $R = 3009$, leading to an evaluation sample of size $P = 2625$. Figure 1 displays the one-day ahead forecasts for the three quantiles and the realisations. The forecasts for the other horizons look very similar, but are expectedly a bit wider.

Figure 1: Quantile Forecasts and Realisations, $h = 1$



We first use our Mincer-Zarnowitz test over the three quantiles, $\mathcal{T} = \{0.01, 0.025, 0.05\}$, and ten horizons, $\mathcal{H} = \{1, \dots, 10\}$. We use $B = 1000$ bootstrap draws and a block length of $l = 10$. Table 5 presents the results. With a p -value of 0.01 there is clear evidence against the null of autocalibration. Alternative block length choices of $l = 5$ and $l = 20$ lead to very similar p -values (0.011 and 0.015) which is promising in that the results are insensitive to block length.

Table 5: Mincer-Zarnowitz Test Results, Finance Application

Stat	90%	95%	99%	p-value
9834.131	5009.153	6569.754	9821.539	0.01

As the test statistic from (8) can directly be interpreted as an empirical distance from the null,

⁷We use the implementation of the GARCH bootstrap from the `rugarch` package in R (Ghalanos, 2022).

⁸Data taken from the Oxford-Man Realized Library: <https://realized.oxford-man.ox.ac.uk/data/download> [Last accessed: 05/07/22]

consisting of squared deviations of all the Mincer-Zarnowitz regression coefficients from their values under the null, we can also look at the individual contributions to this statistic from single quantiles and single horizons or single quantile-horizon combinations. Table 6 contains these individual contributions. That is, for a certain quantile-horizon combination it contains the scaled (by \sqrt{P}) sum of the squared intercept and the squared difference between the slope and unity. From this table a clear picture emerges. The outer quantiles and the longer forecast horizons contribute more to the test statistic, and thus show stronger evidence for miscalibration.

Table 6: Individual Contributions to Test Statistic, Mincer-Zarnowitz Test, Finance Application


	$\tau = 0.01$	$\tau = 0.025$	$\tau = 0.05$	Sum
$h = 1$	427.463	81.455	39.217	548.135
$h = 2$	439.467	195.77	50.126	685.363
$h = 3$	672.67	266.256	127.524	1066.45
$h = 4$	591.907	265.84	99.559	957.306
$h = 5$	549.574	431.091	141.886	1122.551
$h = 6$	553.68	431.926	114.26	1099.866
$h = 7$	149.554	291.555	230.722	671.831
$h = 8$	258.922	298.486	223.656	781.063
$h = 9$	560.313	405.563	402.132	1368.008
$h = 10$	497.402	562.498	473.658	1533.558
Sum	4700.952	3230.439	1902.74	9834.131

Since risk management is typically concerned about the performance of a certain risk model such as the GARCH(1,1) across a range of quantile levels or horizons, this also demonstrates that a common practice to evaluate those models only for a specific choice of the latter may lead to wrong conclusions about the overall performance of the prediction model. In fact, Table 9 in part D of the Appendix, which contains the p-values for individual autocalibration tests at given values τ and h , illustrates that such ‘telescoping’ practice may indeed be misleading. For instance, we observe that for the quantile level 0.05 and at smaller forecast horizons there is no strong evidence against autocalibration from those individual test results.

The tests may also convey information about how models could be improved by a closer look at the Mincer-Zarnowitz regression lines themselves. Specifically, Tables 10 and 11 in part D of the Appendix report the estimated intercepts and slopes of these regression lines for all horizons and quantiles considered. Remarkably, all estimated intercepts are negative and all slopes smaller than one. We use $h = 1$ and $\tau = 0.01$ as an illustrative and representative example in Figure 2. The figure contains a version of a quantile calibration plot or quantile reliability diagram (Pohle, 2020; Gneiting and Resin, 2021; Gneiting et al., 2022). It shows the scatter plot of forecast-observation pairs alongside the estimated Mincer-Zarnowitz regression line and the diagonal. The latter represents the population regression line under H_0^{MZ} , in other words when $\alpha_1^\dagger(0.01) = 0$ and $\beta_1^\dagger(0.01) = 1$, respectively. The discrepancy between the Mincer-Zarnowitz regression line and the diagonal thus suggests that the forecasts are in fact mis-calibrated. Contrasting the two, it becomes clear that in calmer times (when the forecasts and realizations are less extreme, i.e. closer to 0) the GARCH(1,1) forecasts tend to under-predict the actual risk, in other words the quantile forecasts are not extreme enough, while in more volatile times the forecasts tend to overestimate risk.

Finally, as a robustness check, we executed our analysis with a shorter sample ending in 2019 to

Figure 2: Scatter Plot of Forecast-Realisation Pairs with Mincer-Zarnowitz Regression Line (red) and Diagonal (orange) for $h = 1$ and $\tau = 0.01$



exclude the COVID-19 period. The results did not change qualitatively, but the p -values tended to be even smaller, that is, the rejections of autocalibration were even more pronounced. Furthermore, we also experimented with the GJR-GARCH model (Glosten et al., 1993), but results remained virtually unchanged.

6.2 Empirical Application 2: U.S. Macro Series

In this section we use our tests to explore the optimality of model-based forecasts of various U.S. macroeconomic series. The analysis of quantile forecasts for macroeconomic series has become widespread since studies like Manzan (2015). More recently, the GaR literature has emerged to provide a tool to monitor downside risk to economic growth using quantile predictions. This approach typically analyses quarterly real GDP growth using financial conditions indicators (see Adrian et al., 2019), and has been subsequently applied to other quarterly macro series like employment and inflation by Adams et al. (2021).

However, in spite of the increasing interest in quantile forecasting in macroeconomics, none of these papers subject their models to the type of forecast optimality test we develop in this paper. We aim to fill this gap in the empirical literature, applying our tests to shed light on the optimality of commonly-used models in predicting various macro series.

Instead of using quarterly data we propose the use of monthly variables (also used recently in similar contexts by Chavleishvili and Manganello, 2019, and Corradi et al., 2021) and we will focus on the same four target variables analysed in Manzan (2015). These series, all transformed to stationarity using the growth rate, are the Consumer Price Index for All Urban Consumers (CPIAUCSL), Industrial Production: Total Index (INDPRO), All Employees, Total Nonfarm (PAYEMS) and Personal Consumption Expenditures Excluding Food and Energy (Chain-Type Price Index) (PCEPILFE).⁹ These series are very close in nature to the quarterly series analysed in Adams et al. (2021). These will be regressed on an autoregressive term and the Chicago Fed National Financial Conditions Index

⁹All series in the study are taken from the Federal Reserve Economic Data (FRED). Url: <https://fred.stlouisfed.org/> [Last accessed: 08/03/22]

(NFCI) as in Adrian et al. (2019).

Specifically, we use the direct forecasting scheme to generate quantile forecasts at quantile levels τ_k for $k = 1, \dots, K$ and horizons $h = 1, \dots, H$ as follows:

$$\hat{y}_{\tau_k, t, h} = \hat{\gamma}_{0, h, t}(\tau_k) + \hat{\gamma}_{1, h, t}(\tau_k)y_{t-h} + \hat{\gamma}_{2, h, t}(\tau_k)x_{t-h} \quad (23)$$

where y_{t-h} is the autoregressive term corresponding to one of the four target variables mentioned above and x_{t-h} is the NFCI. The parameter estimates are obtained by the standard quantile regression estimator and are indexed both by τ_k and h to denote that a separate quantile regression is run at each quantile and horizon as in the direct scheme, as well as by t as the forecasts are generate in a pseudo out-of-sample fashion as mentioned below. In essence, equation (23) boils down to a forecast made by a quantile autoregressive distributed lag (QADL) model (Galvao Jr. et al., 2013) using the direct forecasting scheme.

The data series span the period 1984M1 to 2019M12, giving a total number of $T = 432$ monthly observations. We use the recursive out-of-sample scheme and split the sample into equal portions for the initial estimation sample and the evaluation sample, $R = P = 216$. This gives an evaluation sample size, P , around the middle of the range of Monte Carlo simulations above. In making forecasts using (23) we will use horizons $h = 1, \dots, 12$ and quantile levels $\tau_k \in \{0.1, 0.25, 0.5\}$. The use of these quantile levels allows us to focus on the left part of the distribution, as is common in GaR studies such as Adams et al. (2021), but also includes the median as an important case of predicting the centre of the distribution. For the bootstrap implementation we use $B = 1000$ bootstrap draws and employ a block length of $l = 4$ as this is seen to work well in the simulation study.

The results in Table 7 display the results of the Mincer-Zarnowitz tests for autocalibration. To give some graphical insight into the behaviour of the out-of-sample predictions, some time series plots for $h = 1$ are given in section E.1 in the Appendix. We first analyse the joint Mincer-Zarnowitz test (“Joint”) which works on multiple time series, as described above, where in this context we have $G = 4$ target variables and we jointly test for autocalibration across all series to avoid the multiple testing problem. The results in the first row of Table 7 show that there is indeed some evidence against autocalibration when looking across all four macro series. The p -value of 0.07 indicates that there is evidence at the 10% significance level that the QADL-type model does not produce well-calibrated forecasts jointly across these four series, for forecast horizons $h = 1, \dots, 12$ and quantile levels $\tau_k \in \{0.1, 0.25, 0.5\}$.

Table 7: Mincer-Zarnowitz Test Results

	Stat	90%	95%	99%	p -value
Joint	38264.280	28908.454	45259.085	86531.304	0.067
CPIAUCSL	18269.966	18033.852	32452.813	66594.353	0.099
INDPRO	4258.078	7578.204	11224.918	24413.160	0.222
PAYEMS	871.704	1574.085	2060.305	4994.712	0.308
PCEPILFE	14864.532	2316.907	2792.387	3678.394	0.000

With this in mind, it is useful to dig further into the individual series to see which of them are likely to be causing the rejection of the joint null of autocalibration. The remainder of Table 7 displays the Mincer-Zarnowitz test when performed individually for each series. For the two real series, industrial

production and employment, we see little evidence against the null. On the other hand, for the two price-type series we see somewhat different results with a clear rejection in the case of PCEPILFE and a p -value just under 10% for CPIAUCSL. This suggests that the QADL-type approach suggested by Adrian et al. (2019) does appear appropriate for real macroeconomic series but less-so for price series.

As a next step of the analysis, for the price series one can also consider digging even further into the test statistics for each series to try and isolate which specific horizons and/or quantile levels are contributing most to the rejection of the null hypothesis. We report these results in Table 12 in the Appendix. For both PCEPILFE and CPIAUCSL we find that the smallest contribution to the test statistic comes from quantile level $\tau_k = 0.5$, whereas there is a much larger contribution from the $\tau_k = 0.1$ quantile level. Except for $h = 1$, where the contributions to the test statistic are much smaller, there is no systematic conclusion when looking across horizons $h = 2, \dots, 12$. This seems to indicate that further study should consider investigating the types of series which might deliver better-calibrated predictions in the far-left tail of the distribution of inflation-type series. We could dig even further by analysing the single MZ regression coefficients and drawing similar forecast-realisation scatter plots as in the previous application. We content ourselves with providing one example of such plots in Appendix E.2 for $h = 1$ and $\tau = 0.1$.

One final exercise we perform is to apply the augmented MZ test where we use additional predictors in the MZ regression. Table 8 displays the results for each of the four series above, where in each case the remaining three variables were used as the augmenting regressors. This serves as a simple check to see if any of these other variables would have been able to improve the forecasts if they were added to the forecasting model, especially for the real variables for which the weaker null of autocalibration was not rejected. However, the results in Table 8 are similar to the non-augmented version of the test. As expected, the stronger null is rejected as well for the inflation type series, which already showed rejections for the weaker null of autocalibration. More interestingly, for the real variables, we still get no rejections. This suggests that we are not able to improve these forecasts by the addition of inflation type variables to the forecasting model.

Table 8: Augmented Mincer-Zarnowitz Test Results

	Stat	90%	95%	99%	p -value
CPIAUCSL	21984.030	19794.203	29896.138	57657.304	0.085
INDPRO	5194.690	8722.551	12596.841	27604.813	0.224
PAYEMS	723.354	1494.399	2011.985	4470.360	0.350
PCEPILFE	15648.207	2455.174	2938.071	3801.048	0.000

7 Conclusion

This paper deals with the absolute evaluation of quantile forecasts in situations where prediction are made over multiple horizons and possibly multiple quantile levels. We propose multi-horizon, multi-quantile tests for optimality employing quantile Mincer-Zarnowitz regressions and a moment equality framework with a bootstrap methodology which avoids the estimation of a large covariance matrix. The main quantile Mincer-Zarnowitz test is of the null hypothesis of autocalibration, which is a fundamental property of forecast consistency. We also provide two extensions. The first extension tests a stronger null hypothesis, which allows us to add further important variables to the information

set with respect to which optimality is tested. This augmented quantile Mincer-Zarnowitz test thus makes it possible to examine if the information contained in those variables was used optimally by the forecaster. The second extension is a multivariate quantile Mincer-Zarnowitz test and allows us to check autocalibration of forecasts for multiple time series at possibly multiple horizons and quantiles.

Our tests allow for an overall decision about the quality of a forecasting approach, whether it is a single model used over multiple horizons and quantiles or a mix of different models and expert judgement employed by an institution. Crucially, it avoids the multiple testing problem inherent to most practical situations, where many forecasts are made over horizons, quantiles or multiple variables. Importantly, our testing framework is constructive in that it does not only provide a formal procedure to reach this overall decision, but may also provide valuable feedback about possible weaknesses of the forecasting approach under consideration and how it could be improved.

There are many possible future avenues arising from our work, for instance the evaluation of distributional or probabilistic forecasts that have attracted substantial interest in recent years (Tay and Wallis, 2000; Gneiting and Katzfuss, 2014). Since these distributional forecasts are considered quantile calibrated (Gneiting and Resin, 2021) when the corresponding quantile forecasts for all quantiles are autocalibrated, one future extension of our work may look into optimality testing across many quantiles requiring the use of many moment equality tests (Chernozhukov et al., 2021).

References

- Abadir, K. M., A. Luati, and P. Paruolo (2022). GARCH density and functional forecasts. *Journal of Econometrics*, forthcoming.
- Adams, P. A., T. Adrian, N. Boyarchenko, and D. Giannone (2021). Forecasting macroeconomic risks. *International Journal of Forecasting* 37(3), 1173–1191.
- Adrian, T., N. Boyarchenko, and D. Giannone (2019). Vulnerable growth. *American Economic Review* 109(4), 1263–1289.
- Andrews, D. (1994). Empirical process methods in econometrics. In R. Engle and D. McFadden (Eds.), *Handbook of Econometrics*, Volume 4, Chapter 37, pp. 228–2294. Elsevier.
- Andrews, D. and P. Guggenberger (2009). Validity of subsampling and ‘plug-in asymptotic’ inference for parameters defined by moment inequalities. *Econometric Theory* 25, 669–709.
- Andrews, D. and G. Soares (2010). Inference for parameters defined by moment inequalities using generalized moment selection. *Econometrica* 78, 119–157.
- Antolin Diaz, J., T. Drechsel, and I. Petrella (2021). Advances in nowcasting economic activity. *Mimeo*.
- Bayer, S. and T. Dimitriadis (2022). Regression-based expected shortfall backtesting. *Journal of Financial Econometrics* 20(3), 437–471.
- Bollerslev, T. (1986). Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics* 31(3), 307–327.
- Brownlees, C. T. and A. Souza (2021). Backtesting Global Growth-at-Risk. *Journal of Monetary Economics* 118, 312–330.
- Carriero, A., T. E. Clark, and M. Marcellino (2020). Nowcasting Tail Risks to Economic Activity with Many Indicators. *Federal Reserve Bank of Cleveland, Working Paper No. 20-13*.
- Chavleishvili, S. and S. Manganelli (2019). Forecasting and stress testing with quantile vector autoregression. *ECB Working Paper 2330*.
- Chernozhukov, V., D. Chetverikov, and K. Kato (2021). Inference on causal and structural parameters using many moment inequalities. *The Review of Economic Studies* 86(5), 1867–1900.
- Christoffersen, P. F. (1998). Evaluating interval forecasts. *International economic review*, 841–862.
- Clements, M. P. (2022). Forecaster Efficiency, Accuracy, and Disagreement: Evidence Using Individual-Level Survey Data. *Journal of Money, Credit and Banking* 54(2-3), 537–568.
- Corradi, V., J. Fosten, and D. Gutknecht (2021). Conditional Quantile Coverage: an Application to Growth-at-Risk. Available at <https://ssrn.com/abstract=3670575>.
- Elliott, G. and A. Timmermann (2016). *Economic Forecasting*. Princeton University Press.

- Engle, R. F. and S. Manganelli (2004). CAViaR: Conditional autoregressive value at risk by regression quantiles. *Journal of Business & Economic Statistics* 22(4), 367–381.
- Escanciano, J. C. and J. Olmo (2010). Backtesting Parametric Value-at-Risk With Estimation Risk. *Journal of Business & Economic Statistics* 28(1), 36–51.
- Escanciano, J. C. and J. Olmo (2011, jun). Robust Backtesting Tests for Value-at-risk Models. *Journal of Financial Econometrics* 9(1), 132–161.
- Ferrara, L., M. Mogliani, and J.-G. Sahuc (2021). High-frequency monitoring of growth at risk. *International Journal of Forecasting*.
- Fissler, T. and J. F. Ziegel (2016). Higher order elicibility and Osband’s principle. *The Annals of Statistics* 44(4), 1680–1707.
- Fissler, T., J. F. Ziegel, and T. Gneiting (2015). Expected shortfall is jointly elicitable with value at risk-implications for backtesting. *arXiv preprint arXiv:1507.00244*.
- Fosten, J. and D. Gutknecht (2020). Testing Nowcast Monotonicity with Estimated Factors. *Journal of Business & Economic Statistics* 38(1), 107–123.
- Gaglianone, W. P., L. R. Lima, O. Linton, and D. R. Smith (2011). Evaluating Value-at-Risk Models via Quantile Regression. *Journal of Business & Economic Statistics* 29(1), 150–160.
- Galvao, A., G. Montes-Rojas, and J. Olmo (2011). Threshold quantile autoregressive models. *Journal of Time Series Analysis* 32, 253–267.
- Galvao Jr., A. F., G. Montes-Rojas, and S. Y. Park (2013). Quantile autoregressive distributed lag model with an application to house price returns. *Oxford Bulletin of Economics and Statistics* 75(2), 307–321.
- Ghalanos, A. (2022). rugarch: Univariate GARCH models. *R package version 1.4-9*.
- Giacomini, R. and I. Komunjer (2005). Evaluation and Combination of Conditional Quantile Forecasts. *Journal of Business & Economic Statistics* 23(4), 416–431.
- Giacomini, R., D. N. Politis, and H. White (2013). A Warp-Speed Method for Conducting Monte Carlo Experiments Involving Bootstrap Estimators. *Econometric Theory* 29(3), 567–589.
- Giacomini, R. and B. Rossi (2010). Forecast comparisons in unstable environments. *Journal of Applied Econometrics* 25(4), 595–620.
- Giacomini, R. and H. White (2006). Tests of Conditional Predictive Ability. *Econometrica* 74(6), 1545–1578.
- Glosten, L. R., R. Jagannathan, and D. E. Runkle (1993). On the relation between the expected value and the volatility of the nominal excess return on stocks. *The Journal of Finance* 48(5), 1779–1801.
- Gneiting, T. (2011). Making and evaluating point forecasts. *Journal of the American Statistical Association* 106(494), 746–762.

- Gneiting, T., F. Balabdaoui, and A. E. Raftery (2007). Probabilistic forecasts, calibration and sharpness. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 69(2), 243–268.
- Gneiting, T. and M. Katzfuss (2014). Probabilistic forecasting. *Annual Review of Statistics and Its Application* 1, 125–151.
- Gneiting, T. and R. Ranjan (2013). Combining predictive distributions. *Electronic Journal of Statistics* 7, 1747–1782.
- Gneiting, T. and J. Resin (2021). Regression diagnostics meets forecast evaluation: Conditional calibration, reliability diagrams, and coefficient of determination. *arXiv preprint arXiv:2108.03210*.
- Gneiting, T., D. Wolfram, J. Resin, K. Kraus, J. Bracher, T. Dimitriadis, V. Hagenmeyer, A. I. Jordan, S. Lerch, K. Phipps, and M. Schienle (2022). Model diagnostics and forecast evaluation for quantiles. *Annual Review of Statistics and Its Application* 10.
- Goncalves, S. and H. White (2002). The bootstrap of the mean for dependent and heterogeneous arrays. *Econometric Theory* 18, 1367–1384.
- Goncalves, S. and H. White (2004). Maximum likelihood and the bootstrap for nonlinear dynamic models. *Journal of Econometrics* 119, 199–219.
- Götze, F. and H. Künsch (1996). Second-order correctness of the blockwise bootstrap for stationary observations. *The Annals of Statistics* 24(5), 1914–1933.
- Granger, C. W. (1969). Prediction with a generalized cost of error function. *Journal of the Operational Research Society* 20(2), 199–207.
- Gregory, K., S. Lahiri, and D. Nordman (2018). A smooth block bootstrap for quantile regression with time series. *Annals of Statistics* 46(3), 1138–1166.
- Hansen, P. R., A. Lunde, and J. M. Nason (2011). The Model Confidence Set. *Econometrica* 79(2), 453–497.
- He, X. D., S. Kou, and X. Peng (2022). Risk measures: robustness, elicibility, and backtesting. *Annual Review of Statistics and Its Application* 9, 141–166.
- Knüppel, M., F. Krüger, and M.-O. Pohle (2022). Score-based calibration testing for multivariate forecast distributions. *arXiv preprint arXiv:2211.16362*.
- Koenker, R. and Z. Xiao (2006). Quantile autoregression. *Journal of the American Statistical Association* 101(475), 980–990.
- Kratz, M., Y. H. Lok, and A. J. McNeil (2018). Multinomial VaR backtests: A simple implicit approach to backtesting expected shortfall. *Journal of Banking & Finance* 88, 393–407.
- Krüger, F. and J. F. Ziegel (2021). Generic conditions for forecast dominance. *Journal of Business & Economic Statistics* 39(4), 972–983.
- Künsch, H. (1989). The jackknife and the bootstrap for general stationary observations. *Annals of Statistics* 17(3), 1217–1241.

- Manzan, S. (2015). Forecasting the distribution of economic variables in a data-rich environment. *Journal of Business & Economic Statistics* 33(1), 144–164.
- Mincer, J. A. and V. Zarnowitz (1969). The evaluation of economic forecasts. In *Economic forecasts and expectations: Analysis of forecasting behavior and performance*, pp. 3–46. NBER.
- Nolde, N. and J. F. Ziegel (2017). Elicitability and backtesting: Perspectives for banking regulation. *The Annals of Applied Statistics* 11(4), 1833–1874.
- Oodaira, H. and K. Yoshihara (1971). The law of the iterated logarithm for stationary processes satisfying mixing conditions. *Kodai Mathematical Seminar Reports* 23(3), 311–334.
- Pascual, L., J. Romo, and E. Ruiz (2006). Bootstrap prediction for returns and volatilities in GARCH models. *Computational Statistics & Data Analysis* 50(9), 2293–2312.
- Patton, A. and A. Timmermann (2012). Forecast rationality tests based on multi-horizon bounds. *Journal of Business and Economic Statistics* 30(1).
- Plagborg-Møller, M., L. Reichlin, G. Ricco, and T. Hasenzagl (2020). When is growth at risk? *BPEA Conference Draft, Spring*.
- Pohle, M.-O. (2020). The Murphy Decomposition and the Calibration-Resolution Principle: A New Perspective on Forecast Evaluation. *arXiv preprint arXiv:2005.01835*.
- Prasad, A., S. Elekdag, P. Jeasakul, R. Lafarguette, A. Alter, A. X. Feng, and C. Wang (2019). Growth at Risk: Concept and Application in IMF Country Surveillance. IMF Working Paper 19/36, International Monetary Fund.
- Quaedvlieg, R. (2021). Multi-Horizon Forecast Comparison. *Journal of Business & Economic Statistics* 39(1), 40–53.
- Romano, J. and A. Shaikh (2008). Inference for identifiable parameters in partially identified econometric models. *Journal of Statistical Inference and Planning* 138, 2786–2807.
- Romano, J. and A. Shaikh (2010). Inference for the identified set in partially identified econometric models. *Econometrica* 78, 169–211.
- Tay, A. S. and K. F. Wallis (2000). Density forecasting: a survey. *Journal of forecasting* 19(4), 235–254.
- Tsyplakov, A. (2013). Evaluation of probabilistic forecasts: proper scoring rules and moments. *Available at SSRN 2236605*.
- West, K. D. (1996). Asymptotic Inference about Predictive Ability. *Econometrica* 64(5), 1067–1084.
- White, H. (2000). A Reality Check for Data Snooping. *Econometrica* 68(5), 1097–1126.

Appendix

A Characterisations of Optimality

Lemma A.1. *Let $\widehat{y}_{\tau,t,h}^*$ be an optimal h -step ahead forecast for the τ -quantile (cf. Definition 1). Then $\widehat{y}_{\tau,t,h}^*$ is optimal relative to $\mathcal{I}_t \subset \mathcal{F}_t$ as well.*

Proof of Lemma A.1: By optimality it holds for $\widehat{y}_{\tau,t,h}^*$ that

$$\mathbb{E} [L_{\tau} (y_t - \widehat{y}_{\tau,t,h}^*) | \mathcal{F}_{t-h}] \leq \mathbb{E} [L_{\tau} (y_t - \widehat{y}_{\tau,t,h}) | \mathcal{F}_{t-h}]$$

for all possible forecasts $\widehat{y}_{\tau,t,h}$ based on \mathcal{F}_{t-h} . By monotonicity of (conditional) expectation and the law of iterated expectation, this implies

$$\mathbb{E} [L_{\tau} (y_t - \widehat{y}_{\tau,t,h}^*) | \mathcal{I}_{t-h}] \leq \mathbb{E} [L_{\tau} (y_t - \widehat{y}_{\tau,t,h}) | \mathcal{I}_{t-h}],$$

which proves the claim. ■

B Proofs of Lemmas and Theorems

This section provides an asymptotic linear Bahadur representation for the quantile regression estimator and a stochastic equicontinuity result in Lemma B.1, which is subsequently used in the proof of Theorem 1. The proof of the auxiliary Lemma B.1 follows at the end of this section.

Lemma B.1. *Under Assumptions A1 and A7, it holds that:*

(i) *For each $\tau \in \mathcal{T}$ and $h \in \mathcal{H}$:*

$$\left\| \widehat{\beta}_h(\tau) - \beta_h^{\dagger}(\tau) \right\| = o_{\text{Pr}}(1),$$

where $\beta_h^{\dagger}(\tau)$ is defined in Equation (3) and $\widehat{\beta}_h(\tau)$ in Equation (6).

(ii) *The empirical process:*

$$\frac{1}{\sqrt{P}} \sum_{t=R+1}^T (\mathbf{X}_{\tau,t,h}(\boldsymbol{\theta}) (1 \{y_t \leq \mathbf{X}_{\tau,t,h}(\boldsymbol{\theta})' \boldsymbol{\beta}\} - \tau) - \mathbb{E} (\mathbf{X}_{\tau,t,h}(\boldsymbol{\theta}) (1 \{y_t \leq \mathbf{X}_{\tau,t,h}(\boldsymbol{\theta})' \boldsymbol{\beta}\} - \tau)))$$

is stochastically equicontinuous in $\boldsymbol{\theta} \in \Theta$, $\boldsymbol{\beta} \in \mathcal{B}$, and $\tau \in \mathcal{T}$ w.r.t. the L_2 pseudo-metric:

$$\begin{aligned} & \rho_{\Theta \times \mathcal{B} \times \mathcal{T}}((\boldsymbol{\theta}_1, \boldsymbol{\beta}_1, \tau_1), (\boldsymbol{\theta}_2, \boldsymbol{\beta}_2, \tau_2))^2 \\ &= \max_{l \in d} \mathbb{E} \left((X_{l,t,h}(\boldsymbol{\theta}_1) (1 \{y_t \leq \mathbf{X}_{\tau,t,h}(\boldsymbol{\theta}_1)' \boldsymbol{\beta}\} - \tau) - X_{l,t,h}(\boldsymbol{\theta}_2) (1 \{y_t \leq \mathbf{X}_{\tau,t,h}(\boldsymbol{\theta}_2)' \boldsymbol{\beta}\} - \tau_2))^2 \right) \end{aligned}$$

where $X_{l,t,h}(\boldsymbol{\theta})$ denotes the l -th element of $\mathbf{X}_{\tau,t,h}(\boldsymbol{\theta})$ and d is the dimension of $\mathbf{X}_{\tau,t,h}(\boldsymbol{\theta})$.

(iii) *For each $\tau \in \mathcal{T}$ and $h \in \mathcal{H}$:*

$$\begin{aligned} & \sqrt{P} \left(\widehat{\beta}_h(\tau) - \beta_h^{\dagger}(\tau) \right) \\ &= \mathbf{J}_h(\tau)^{-1} \left(\frac{1}{\sqrt{P}} \sum_{t=R+1}^T \mathbf{X}_{\tau,t,h}(\boldsymbol{\theta}_{\tau,h}^{\dagger}) \left(1 \{y_t \leq \mathbf{X}_{\tau,t,h}(\boldsymbol{\theta}_{\tau,h}^{\dagger})' \boldsymbol{\beta}^{\dagger}(\tau)\} - \tau \right) \right) + o_{\text{Pr}}(1), \end{aligned}$$

where $\mathbf{J}_h(\tau)$ is defined in Assumption A3.

Proof of Theorem 1: Under A1-A7 and Lemma B.1, it follows under the null hypothesis that:

$$\sqrt{P}\widehat{m}_s \xrightarrow{d} N(0, \Omega_s),$$

for each given $s \in \{1, \dots, \kappa\}$, where Ω_s is a diagonal element of the 2×2 matrix:

$$\tau_k(1 - \tau_k)\mathbf{J}_h(\tau_k)^{-1} \mathbb{E} \left(\mathbf{X}_{\tau_k, t, h}(\boldsymbol{\theta}_h^\dagger) \mathbf{X}_{\tau_k, t, h}(\boldsymbol{\theta}_h^\dagger)' \right) \mathbf{J}_h(\tau_k)^{-1}$$

for a given h and τ_k , see also Corollary 2 of Koenker and Xiao (2006). Thus, under H_0^{MZ} and by continuous mapping:

$$\sum_{s=1}^{\kappa} \left(\sqrt{P}\widehat{m}_s \right)^2 \xrightarrow{d} \sum_{s=1}^{\kappa} Z_s^2,$$

where Z_s is an element of:

$$\mathbf{Z} = \begin{pmatrix} Z_1 \\ \vdots \\ Z_\kappa \end{pmatrix} \sim N(\mathbf{0}, \boldsymbol{\Sigma}).$$

and $\boldsymbol{\Sigma}$ is the asymptotic variance-covariance matrix of:

$$\sqrt{P} \left(\begin{pmatrix} \widehat{\alpha}_1(\tau_1) \\ \widehat{\beta}_1(\tau_1) \\ \vdots \\ \widehat{\alpha}_H(\tau_K) \\ \widehat{\beta}_H(\tau_K) \end{pmatrix} - \begin{pmatrix} 0 \\ 1 \\ \vdots \\ 0 \\ 1 \end{pmatrix} \right).$$

Note also that the first stage estimation error does not feature into this asymptotic variance when $\pi = 0$ as shown in the proof of Lemma B.1. As a result, we may directly resample from the generated forecasts. That is, under A1 to A7 and the block length condition, the first order validity of the MBB follows from Theorem 2 of Gregory et al. (2018) as a special case of the SETBB using untapered blocks for each τ_k and h . That is, pointwise in $\tau_k \in \mathcal{T}$, $h \in \mathcal{H}$, and for any $\nu > 0$:

$$\Pr \left(\sup_{x \in \mathbb{R}^2} \left| \Pr^b \left(\sqrt{P}(\widehat{\beta}_h^b(\tau_k) - \widehat{\beta}_h(\tau_k)) \leq x \right) - \Pr \left(\sqrt{P}(\widehat{\beta}_h(\tau_k) - \beta_h^\dagger(\tau_k)) \leq x \right) \right| > \nu \right) \rightarrow 0$$

Then, since \widehat{m}_s^b either $\widehat{m}_s^b = \widehat{\alpha}_h^b(\tau_k)$ or $\widehat{m}_s^b = \widehat{\beta}_h^b(\tau_k) - 1$, the result follows again by continuous mapping. \blacksquare

Proof of Lemma B.1:

(i) First, observe that for any value $\boldsymbol{\beta} \in \mathcal{B}$, $\tau \in \mathcal{T}$, and $h \in \mathcal{H}$:

$$Q_P(\widehat{\boldsymbol{\theta}}_{\tau, j, h}; \tau, \boldsymbol{\beta}) = Q_P(\boldsymbol{\theta}_{\tau, h}^\dagger; \tau, \boldsymbol{\beta}) - \left(Q_P(\widehat{\boldsymbol{\theta}}_{\tau, j, h}; \tau, \boldsymbol{\beta}) - Q_P(\boldsymbol{\theta}_{\tau, h}^\dagger; \tau, \boldsymbol{\beta}) \right) \quad (24)$$

where:

$$\widehat{Q}_P(\widehat{\boldsymbol{\theta}}_{\tau, j, h}; \tau, \boldsymbol{\beta}) \equiv \frac{1}{P} \sum_{j=R+1}^T \rho_\tau(y_j - \mathbf{X}_j(\widehat{\boldsymbol{\theta}}_{\tau, j, h})' \boldsymbol{\beta})$$

and

$$Q_P(\boldsymbol{\theta}_{\tau,h}^\dagger; \tau, \boldsymbol{\beta}) \equiv \frac{1}{P} \sum_{j=R+1}^P \rho_\tau(y_j - \mathbf{X}_j(\boldsymbol{\theta}_{\tau,h}^\dagger)' \boldsymbol{\beta}).$$

Next, note that:

$$\begin{aligned} & \sup_{\widehat{\boldsymbol{\theta}}_{\tau,j,h}; j \geq R+1} \sup_{\boldsymbol{\beta} \in \mathcal{B}} \left| \left(Q_P(\widehat{\boldsymbol{\theta}}_{\tau,j,h}; \tau, \boldsymbol{\beta}) - Q_P(\boldsymbol{\theta}_{\tau,h}^\dagger; \tau, \boldsymbol{\beta}) \right) \right| \\ & \leq \sup_{\widehat{\boldsymbol{\theta}}_{\tau,j,h}; j \geq R+1} \sup_{\boldsymbol{\beta} \in \mathcal{B}} \left(\frac{1}{P} \sum_{j=R+1}^T \left| \rho_\tau(y_j - \mathbf{X}_j(\widehat{\boldsymbol{\theta}}_{\tau,j,h})' \boldsymbol{\beta}) - \rho_\tau(y_j - \mathbf{X}_j(\boldsymbol{\theta}_{\tau,h}^\dagger)' \boldsymbol{\beta}) \right| \right) \\ & \leq \sup_{\widehat{\boldsymbol{\theta}}_{\tau,j,h}; j \geq R+1} \left(\frac{1}{P} \sum_{j=R+1}^T \left\| \mathbf{X}_j(\widehat{\boldsymbol{\theta}}_{\tau,j,h}) - \mathbf{X}_j(\boldsymbol{\theta}_{\tau,h}^\dagger) \right\| \right) \times \sup_{\boldsymbol{\beta} \in \mathcal{B}} \|\boldsymbol{\beta}\| \times \max\{\tau, 1 - \tau\} \\ & \leq \frac{1}{P} \sum_{j=R+1}^T \sup_{\widehat{\boldsymbol{\theta}}_{\tau,j,h}; j \geq R+1} \left\| \mathbf{X}_j(\widehat{\boldsymbol{\theta}}_{\tau,j,h}) - \mathbf{X}_j(\boldsymbol{\theta}_{\tau,h}^\dagger) \right\| \times \sup_{\boldsymbol{\beta} \in \mathcal{B}} \|\boldsymbol{\beta}\| \\ & \leq C \cdot O_{\text{Pr}} \left(R^{-\frac{1}{2}} \right) \end{aligned}$$

for some generic constant $C > 0$, where the second inequality follows from the properties of the ‘check function’, the third one from Jensen’s inequality, and the last one from Markov’s inequality and Assumptions A3 and A6. The rest of the proof can now proceed by standard arguments from the quantile regression literature. In particular, observe that the definition of $Q_P(\boldsymbol{\theta}_{\tau,h}^\dagger; \tau, \boldsymbol{\beta})$ (and $Q_P(\widehat{\boldsymbol{\theta}}_{\tau,j,h}; \tau, \boldsymbol{\beta})$) may be changed to:

$$Q_P(\boldsymbol{\theta}_{\tau,h}^\dagger; \tau, \boldsymbol{\beta}) \equiv \frac{1}{P} \sum_{j=R+1}^T \left(\rho_\tau(y_j - \mathbf{X}_j(\boldsymbol{\theta}_{\tau,h}^\dagger)' \boldsymbol{\beta}) - \rho_\tau(y_j - \mathbf{X}_j(\boldsymbol{\theta}_{\tau,h}^\dagger)' \boldsymbol{\beta}_h^\dagger(\tau)) \right)$$

without affecting minimisation or any of the preceding arguments. In addition, note that by A2, A3, and A4, it follows that:

$$Q_\infty(\tau, \boldsymbol{\beta}) \equiv \text{E} \left[\left(\rho_\tau(y_j - \mathbf{X}_j(\boldsymbol{\theta}_{\tau,h}^\dagger)' \boldsymbol{\beta}) - \rho_\tau(y_j - \mathbf{X}_j(\boldsymbol{\theta}_{\tau,h}^\dagger)' \boldsymbol{\beta}_h^\dagger(\tau)) \right) \right] = O(1).$$

As a result, pointwise in $\boldsymbol{\beta}$, τ , h , $Q_P(\boldsymbol{\theta}_{\tau,h}^\dagger; \tau, \boldsymbol{\beta}) \xrightarrow{\text{Pr}} Q_{\infty, \boldsymbol{\theta}^\dagger}(\tau, \boldsymbol{\beta})$ by A1, A3, and McLeish’s law of large numbers for strong mixing processes. Consistency of $\widehat{\boldsymbol{\beta}}$ for $\boldsymbol{\beta}_h^\dagger(\tau)$ pointwise in τ and h then follows since $Q_\infty(\tau, \boldsymbol{\beta})$ is uniquely minimised at $\boldsymbol{\beta}_h^\dagger(\tau)$ by the stated assumptions.

(ii) To establish stochastic equicontinuity w.r.t. the pseudo-metric $\rho_{\Theta \times \mathcal{B} \times \mathcal{T}}(\cdot, \cdot)$, note that $\mathbf{X}_j(\boldsymbol{\theta})$ belongs to a type II class function as defined in Andrews (1994) with envelope $\max_{R+1 \leq j \leq T} \|\mathbf{X}_j(\boldsymbol{\theta})\|$, where by A2:

$$\begin{aligned} \Pr \left(\max_{R+1 \leq j \leq T} \|\mathbf{X}_j(\boldsymbol{\theta})\| > P^{\frac{1}{2}} \right) & \leq \sum_{j=R+1}^T \Pr \left(\|\mathbf{X}_j(\boldsymbol{\theta})\| > P^{\frac{1}{2}} \right) \\ & \leq \text{E} \left[\|\mathbf{X}_j(\boldsymbol{\theta})\|^{4+\epsilon} \right] / P^{\frac{4+\epsilon}{2}} = o(1). \end{aligned}$$

In addition, note that:

$$(1\{y_j \leq \mathbf{X}_j(\boldsymbol{\theta})' \boldsymbol{\beta}\} - \tau)$$

can be shown to belong to a type IV class function as defined in Andrews (1994) by A2, A3, A4, and A6. Now, given the α -mixing condition in Assumption A1, stochastic equicontinuity in $(\boldsymbol{\theta}, \boldsymbol{\beta}, \tau) \in \Theta \times \mathcal{B} \times \mathcal{T}$ follows from repeated applications of Theorems 5 and 6 together with Theorem 4 in Andrews (1994).

(iii) By Equation (6) and A2 (see e.g. Gregory et al., 2018), it holds that:

$$\begin{aligned} & \left\| \frac{1}{\sqrt{P}} \sum_{j=R+1}^T \mathbf{X}_j(\widehat{\boldsymbol{\theta}}_{\tau,t,h}) \left(1 \{y_j \leq \mathbf{X}_j(\widehat{\boldsymbol{\theta}}_{\tau,t,h})' \widehat{\boldsymbol{\beta}}_P(\tau)\} - \tau \right) \right\| \\ & \leq d \max_{R+1 \leq j \leq T} \left\| \mathbf{X}_j(\boldsymbol{\theta}_{\tau,h}^\dagger) \right\| \frac{1}{\sqrt{P}} \sum_{j=R+1}^T 1 \{y_j = \mathbf{X}_j(\widehat{\boldsymbol{\theta}}_{\tau,t,h})' \widehat{\boldsymbol{\beta}}_P(\tau)\} + o_{\text{Pr}}(1) = o_{\text{Pr}}(1), \end{aligned} \quad (25)$$

where the inequality follows from A6, and the last equality follows again from the fact that for every $\tau \in \mathcal{T}$ and $h \in \mathcal{H}$:

$$\max_{R+1 \leq j \leq T} \left\| \mathbf{X}_j(\boldsymbol{\theta}_{\tau,h}^\dagger) \right\| = o_{\text{Pr}}(P^{\frac{1}{2}}).$$

Moreover, denoting $E_T[\cdot]$ the expectation conditional on the original sample observations, we have that:

$$\begin{aligned} & \frac{1}{\sqrt{P}} \sum_{j=R+1}^T \mathbf{X}_j(\widehat{\boldsymbol{\theta}}_{\tau,t,h}) \left(1 \{y_j \leq \mathbf{X}_j(\widehat{\boldsymbol{\theta}}_{\tau,t,h})' \widehat{\boldsymbol{\beta}}_P(\tau)\} - \tau \right) \\ &= \frac{1}{\sqrt{P}} \sum_{j=R+1}^T \mathbf{X}_j(\boldsymbol{\theta}_{\tau,h}^\dagger) \left(1 \{y_j \leq \mathbf{X}_j(\boldsymbol{\theta}_{\tau,h}^\dagger)' \widehat{\boldsymbol{\beta}}_P(\tau)\} - \tau \right) + \\ & \quad \left(\frac{1}{\sqrt{P}} \sum_{j=R+1}^T \left(\mathbf{X}_j(\widehat{\boldsymbol{\theta}}_{\tau,t,h}) \left(1 \{y_j \leq \mathbf{X}_j(\widehat{\boldsymbol{\theta}}_{\tau,t,h})' \widehat{\boldsymbol{\beta}}_P(\tau)\} - \tau \right) - E_T \left[\mathbf{X}_j(\widehat{\boldsymbol{\theta}}_{\tau,t,h}) \left(1 \{y_j \leq \mathbf{X}_j(\widehat{\boldsymbol{\theta}}_{\tau,t,h})' \widehat{\boldsymbol{\beta}}_P(\tau)\} - \tau \right) \right] \right) \right. \\ & \quad \left. - \frac{1}{\sqrt{P}} \sum_{j=R+1}^T \left(\mathbf{X}_j(\boldsymbol{\theta}_{\tau,h}^\dagger) \left(1 \{y_j \leq \mathbf{X}_j(\boldsymbol{\theta}_{\tau,h}^\dagger)' \widehat{\boldsymbol{\beta}}_P(\tau)\} - \tau \right) - E_T \left[\mathbf{X}_j(\boldsymbol{\theta}_{\tau,h}^\dagger) \left(1 \{y_j \leq \mathbf{X}_j(\boldsymbol{\theta}_{\tau,h}^\dagger)' \widehat{\boldsymbol{\beta}}_P(\tau)\} - \tau \right) \right] \right) \right) \\ & \quad + \sqrt{P} E_T \left[\mathbf{X}_j(\widehat{\boldsymbol{\theta}}_{\tau,t,h}) \left(1 \{y_j \leq \mathbf{X}_j(\widehat{\boldsymbol{\theta}}_{\tau,t,h})' \widehat{\boldsymbol{\beta}}_P(\tau)\} - \tau \right) - \mathbf{X}_j(\boldsymbol{\theta}_{\tau,h}^\dagger) \left(1 \{y_j \leq \mathbf{X}_j(\boldsymbol{\theta}_{\tau,h}^\dagger)' \widehat{\boldsymbol{\beta}}_P(\tau)\} - \tau \right) \right] \\ &= I_{1,P} + I_{2,P} + I_{3,P} \end{aligned}$$

By part (ii) of this lemma as well as A6 and A7, we have that, conditional on the original sample, $|I_{2,P}| = o_{\text{Pr}}(1)$. In addition, note that:

$$\begin{aligned} |I_{3,P}| &\leq \sup_{\boldsymbol{\beta} \in \mathcal{B}} \left| \sqrt{P} E_T \left[\left((\mathbf{X}_j(\widehat{\boldsymbol{\theta}}_{\tau,t,h}) - \mathbf{X}_j(\boldsymbol{\theta}_{\tau,h}^\dagger))' \boldsymbol{\beta} \right) (1 \{y_j \leq \mathbf{X}_j(\boldsymbol{\theta}_{\tau,h}^\dagger)' \boldsymbol{\beta}\} - \tau) \right] \right| \\ &\quad + \sup_{\boldsymbol{\beta} \in \mathcal{B}} \left| \sqrt{P} E_T \left[\left(y_j - \mathbf{X}_j(\boldsymbol{\theta}_{\tau,h}^\dagger)' \boldsymbol{\beta} \right) (1 \{y_j \leq \mathbf{X}_j(\widehat{\boldsymbol{\theta}}_{\tau,t,h})' \boldsymbol{\beta}\} - 1 \{y_j \leq \mathbf{X}_j(\boldsymbol{\theta}_{\tau,h}^\dagger)' \boldsymbol{\beta}\}) \right] \right| \\ &\quad + \sup_{\boldsymbol{\beta} \in \mathcal{B}} \left| \sqrt{P} E_T \left[\left((\mathbf{X}_j(\widehat{\boldsymbol{\theta}}_{\tau,t,h}) - \mathbf{X}_j(\boldsymbol{\theta}_{\tau,h}^\dagger))' \boldsymbol{\beta} \right) (1 \{y_j \leq \mathbf{X}_j(\widehat{\boldsymbol{\theta}}_{\tau,t,h})' \boldsymbol{\beta}\} - 1 \{y_j \leq \mathbf{X}_j(\boldsymbol{\theta}_{\tau,h}^\dagger)' \boldsymbol{\beta}\}) \right] \right| \\ &= \sup_{\boldsymbol{\beta} \in \mathcal{B}} |A_{1,P}(\tau, \boldsymbol{\beta})| + \sup_{\boldsymbol{\beta} \in \mathcal{B}} |A_{2,P}(\tau, \boldsymbol{\beta})| + \sup_{\boldsymbol{\beta} \in \mathcal{B}} |A_{3,P}(\tau, \boldsymbol{\beta})| \end{aligned}$$

Defining $\boldsymbol{\delta}_\beta = \boldsymbol{\beta} - \boldsymbol{\beta}_h^\dagger(\tau)$ and $\varepsilon_j = y_j - \mathbf{X}_j(\boldsymbol{\theta}_{\tau,h}^\dagger)' \boldsymbol{\beta}_h^\dagger(\tau)$, note that by iterated expectations, A3, A6,

for some constants $\tilde{C}, \tilde{C} > 0$:

$$\begin{aligned} \sup_{\boldsymbol{\beta} \in \mathcal{B}} |A_{1,P}(\tau, \boldsymbol{\beta})| &\leq \tilde{C} \sqrt{\frac{P}{R}} \int B(\mathbf{X}_{\tau_k, j, h}) \sup_{\boldsymbol{\beta} \in \mathcal{B}} \left| F_{\varepsilon_j | X}(\mathbf{X}_j(\boldsymbol{\theta}_{\tau, h}^\dagger)' \boldsymbol{\delta}_\beta | \mathbf{X}_j) - \tau \right| f_X(\mathbf{X}_j) d\mathbf{X}_j \\ &\leq \tilde{C} \sqrt{\frac{P}{R}} \sup_{t \geq R+1} \mathbb{E} [B(\mathbf{X}_{\tau_k, t, h})] = o(1). \end{aligned}$$

Turning to $A_{2,P}(\tau, \boldsymbol{\beta})$, observe that:

$$\begin{aligned} &\sup_{\boldsymbol{\beta} \in \mathcal{B}} |A_{2,P}(\tau, \boldsymbol{\beta})| \\ &= \sup_{\boldsymbol{\beta} \in \mathcal{B}} \left| \sqrt{P} \mathbb{E}_T \left[\left(y_j - \mathbf{X}_j(\boldsymbol{\theta}_{\tau, h}^\dagger)' \boldsymbol{\beta} \right) \left(1\{\mathbf{X}_j(\hat{\boldsymbol{\theta}}_{\tau, t, h})' \boldsymbol{\beta} \leq y_j \leq \mathbf{X}_j(\boldsymbol{\theta}_{\tau, h}^\dagger)' \boldsymbol{\beta}\} \right. \right. \right. \\ &\quad \left. \left. \left. + 1\{\mathbf{X}_j(\boldsymbol{\theta}_{\tau, h}^\dagger)' \boldsymbol{\beta} \leq y_j \leq \mathbf{X}_j(\hat{\boldsymbol{\theta}}_{\tau, t, h})' \boldsymbol{\beta}\} \right) \right] \right| \end{aligned}$$

By Cauchy-Schwarz:

$$\begin{aligned} &\sup_{\boldsymbol{\beta} \in \mathcal{B}} |A_{2,P}(\tau, \boldsymbol{\beta})| \\ &\leq \sup_{\boldsymbol{\beta} \in \mathcal{B}} \left(\mathbb{E} \left[\left(y_j - \mathbf{X}_j(\boldsymbol{\theta}_{\tau, h}^\dagger)' \boldsymbol{\beta} \right)^2 \right] \right)^{\frac{1}{2}} \\ &\quad \times \sup_{\boldsymbol{\beta} \in \mathcal{B}} \left(\sqrt{P} \mathbb{E}_T \left[\left(\left(1\{\mathbf{X}_j(\hat{\boldsymbol{\theta}}_{\tau, t, h})' \boldsymbol{\beta} \leq y_j \leq \mathbf{X}_j(\boldsymbol{\theta}_{\tau, h}^\dagger)' \boldsymbol{\beta}\} + 1\{\mathbf{X}_j(\boldsymbol{\theta}_{\tau, h}^\dagger)' \boldsymbol{\beta} \leq y_j \leq \mathbf{X}_j(\hat{\boldsymbol{\theta}}_{\tau, t, h})' \boldsymbol{\beta}\} \right)^2 \right] \right)^{\frac{1}{2}} \end{aligned}$$

The first term on the Right Hand Side (RHS) is of order $O(1)$ by A2, A3, and A4. Focusing on $1\{\mathbf{X}_j(\boldsymbol{\theta}_{\tau, h}^\dagger)' \boldsymbol{\beta} \leq y_j \leq \mathbf{X}_j(\hat{\boldsymbol{\theta}}_{\tau, t, h})' \boldsymbol{\beta}\}$ (the other term and the cross-product follow by similar arguments) and recalling the definition of $\boldsymbol{\delta}_\beta = \boldsymbol{\beta} - \boldsymbol{\beta}_h^\dagger(\tau)$ and ε_j , we have again conditional on the sample and by iterated expectations:

$$\sqrt{P} \sup_{\boldsymbol{\beta} \in \mathcal{B}} \int \left| F_{\varepsilon_j | X}(\mathbf{X}_j(\boldsymbol{\theta}_{\tau, h}^\dagger)' \boldsymbol{\delta}_\beta + V_j(\mathbf{X}_j; \hat{\boldsymbol{\theta}}_{\tau, t, h}) | \mathbf{X}_j) - F_{\varepsilon_j | X}(\mathbf{X}_j(\boldsymbol{\theta}_{\tau, h}^\dagger)' \boldsymbol{\delta}_\beta | \mathbf{X}_j) \right| f_X(\mathbf{X}_j) d\mathbf{X}_j$$

with:

$$V_j(\mathbf{X}_j; \hat{\boldsymbol{\theta}}_{\tau, t, h}) = (\mathbf{X}_j(\hat{\boldsymbol{\theta}}_{\tau, t, h}) - \mathbf{X}_j(\boldsymbol{\theta}_{\tau, h}^\dagger)' \boldsymbol{\beta}_h^\dagger(\tau) + (\mathbf{X}_j(\hat{\boldsymbol{\theta}}_{\tau, t, h}) - \mathbf{X}_j(\boldsymbol{\theta}_{\tau, h}^\dagger)' \boldsymbol{\delta}_\beta)$$

Using A3, A4, A6 together with a mean value expansion around $\mathbf{X}_j(\boldsymbol{\theta}_{\tau, h}^\dagger)' \boldsymbol{\delta}_\beta$, the last term can be bounded by:

$$\sqrt{\frac{P}{R}} C \sup_{\hat{\boldsymbol{\theta}}_{\tau, t, h}; t \geq R+1} \mathbb{E} [\|B(\mathbf{X}_{\tau, j, h})\|] \text{diam}(\mathcal{B}) = o(1)$$

for some constant $C > 0$, where $\text{diam}(\mathcal{B})$ denotes the diameter of \mathcal{B} and the equality follows from A7.

Finally, the cross-product term $A_{3,P}(\tau, \boldsymbol{\beta})$ can be shown to be $o(1)$ uniformly in $\boldsymbol{\beta} \in \mathcal{B}$ by Cauchy-Schwarz and similar arguments to before, which establishes that for the second term on the RHS of (24) it holds that:

$$|I_{3,P}| = O\left(\frac{\sqrt{P}}{\sqrt{R}}\right) = o(1),$$

conditional on the sample. Now, using parts (i) and (ii), stochastic equicontinuity again yields that:

$$\begin{aligned} & \frac{1}{\sqrt{P}} \sum_{j=R+1}^T \left(\mathbf{X}_j(\boldsymbol{\theta}_{\tau,h}^\dagger) \left(1 \left\{ y_j \leq \mathbf{X}_j(\boldsymbol{\theta}_{\tau,h}^\dagger)' \widehat{\boldsymbol{\beta}}_P(\tau) \right\} - \tau \right) - \mathbb{E}_T \left[\mathbf{X}_j(\boldsymbol{\theta}_{\tau,h}^\dagger) \left(1 \left\{ y_j \leq \mathbf{X}_j(\boldsymbol{\theta}_{\tau,h}^\dagger)' \widehat{\boldsymbol{\beta}}_P(\tau) \right\} - \tau \right) \right] \right) \\ &= \frac{1}{\sqrt{P}} \sum_{j=R+1}^T \left(\mathbf{X}_j(\boldsymbol{\theta}_{\tau,h}^\dagger) \left(1 \left\{ y_j \leq \mathbf{X}_j(\boldsymbol{\theta}_{\tau,h}^\dagger)' \boldsymbol{\beta}^\dagger(\tau) \right\} - \tau \right) - \underbrace{\mathbb{E}_T \left[\mathbf{X}_j(\boldsymbol{\theta}_{\tau,h}^\dagger) \left(1 \left\{ y_j \leq \mathbf{X}_j(\boldsymbol{\theta}_{\tau,h}^\dagger)' \boldsymbol{\beta}^\dagger(\tau) \right\} - \tau \right) \right]}_{=0} \right) \\ &+ o_{\text{Pr}}(1). \end{aligned}$$

Combining this result with (25), we obtain:

$$\begin{aligned} & \sqrt{P} \mathbb{E}_T \left[\mathbf{X}_j(\boldsymbol{\theta}_{\tau,h}^\dagger) \left(1 \left\{ y_j \leq \mathbf{X}_j(\boldsymbol{\theta}_{\tau,h}^\dagger)' \widehat{\boldsymbol{\beta}}_P(\tau) \right\} - \tau \right) \right] \\ &= \frac{1}{\sqrt{P}} \sum_{j=R+1}^T \mathbf{X}_j(\boldsymbol{\theta}_{\tau,h}^\dagger) \left(1 \left\{ y_j \leq \mathbf{X}_j(\boldsymbol{\theta}_{\tau,h}^\dagger)' \boldsymbol{\beta}^\dagger(\tau) \right\} - \tau \right) + o_{\text{Pr}}(1). \end{aligned}$$

Applying again iterated expectations and a mean value expansion of the left hand side around $\boldsymbol{\beta}^\dagger(\tau)$, we have by A4 and A5 that:

$$\begin{aligned} & \sqrt{P} \mathbb{E}_T \left[\mathbf{X}_j(\boldsymbol{\theta}_{\tau,h}^\dagger) \left(1 \left\{ y_j \leq \mathbf{X}_j(\boldsymbol{\theta}_{\tau,h}^\dagger)' \widehat{\boldsymbol{\beta}}_P(\tau) \right\} - \tau \right) \right] \\ &= \sqrt{P} \left(\int f_{\varepsilon_j|X}(\bar{U}_j | \mathbf{X}_j) \mathbf{X}_j(\boldsymbol{\theta}_{\tau,h}^\dagger) \mathbf{X}_j(\boldsymbol{\theta}_{\tau,h}^\dagger)' f_X(\mathbf{X}_j) d\mathbf{X}_j \right) \left(\widehat{\boldsymbol{\beta}}(\tau) - \boldsymbol{\beta}^\dagger(\tau) \right), \end{aligned}$$

where the intermediate value \bar{U}_j lies between 0 and $\mathbf{X}_j(\boldsymbol{\theta}^\dagger)' \left(\widehat{\boldsymbol{\beta}}(\tau) - \boldsymbol{\beta}^\dagger(\tau) \right)$. Therefore:

$$\frac{1}{\sqrt{P}} \sum_{j=R+1}^T \mathbf{X}_j(\boldsymbol{\theta}^\dagger) \left(1 \left\{ y_j \leq \mathbf{X}_j(\boldsymbol{\theta}^\dagger)' \boldsymbol{\beta}^\dagger(\tau) \right\} - \tau \right) = \mathbf{J}(\tau) \sqrt{P} \left(\widehat{\boldsymbol{\beta}}(\tau) - \boldsymbol{\beta}^\dagger(\tau) \right) + o_{\text{Pr}}(1)$$

where $\mathbf{J}(\tau) = \mathbf{J}_h(\tau_k)$. Since $\mathbf{J}(\tau)$ is positive definite, the Bahadur representation for $\sqrt{P} \left(\widehat{\boldsymbol{\beta}}(\tau) - \boldsymbol{\beta}^\dagger(\tau) \right)$ follows. ■

C Horizon Monotonicity Test

In this section, we outline a further possibility to test for quantile forecast optimality, namely across forecast horizons when forecasts are reported at multiple horizons. The test exploits the fact that, under optimality, expected quantile loss is monotonically non-decreasing in the forecast horizon. More specifically, let $\widehat{y}_{\tau,t,h}^*$, $h \in \mathcal{H} = \{1, \dots, H\}$, be optimal forecasts for the τ -quantile. It then holds that

$$\mathbb{E} \left[L_\tau \left(y_{t,h_i} - \widehat{y}_{\tau,t,h_i}^* \right) \right] \leq \mathbb{E} \left[L_\tau \left(y_{t,h} - \widehat{y}_{\tau,t,h_j}^* \right) \right]$$

for all $h_i, h_j \in \mathcal{H}$ with $h_i < h_j$, which follows from strict stationarity, iterated expectations, and the monotonicity of conditional expectations. The result is akin to Patton and Timmermann (2012) who demonstrated that the mean squared forecast error (MSFE) of optimal multi-horizon conditional mean forecasts does not decrease with the forecast horizon, and similar tests have in fact been constructed

for mean forecasts (see Fosten and Gutknecht, 2020, and references therein). Intuitively, forecasting becomes more difficult the longer the forecast horizon gets as the information set gets smaller, so optimal forecasts should display non-decreasing expected loss. A decrease in expected loss from a certain horizon to the next therefore indicates non-optimal forecasts, meaning that forecasts further into the future are systematically more accurate than the forecasts at the shorter horizon. This can arise, for instance, if different models are used for short and long-term forecasting or if different specifications of models are selected and estimated at different horizons. The latter is often the case in quantile forecasting where the direct scheme is typically used, as opposed to the iterative scheme, and e.g. relevant variables may be excluded at shorter horizons, or modelled and processed in a way such that non-monotonicity in the information content over the horizons arises.

The notation is as in the previous sections, although we will treat estimation error here as implicit since arguably the most relevant applications (unless different models are used at different horizons) are to situations where forecasts are reported without model like in the case of the Survey of Professional Forecasters. We are interested in the following hypothesis:

$$H_0^{\text{MH}} : \mathbb{E}[L_{\tau_k}(y_t - \hat{y}_{\tau_k,t,h_i}) - L_{\tau_k}(y_t - \hat{y}_{\tau_k,t,h_j})] \geq 0 \quad (26)$$

for all $h_i, h_j \in \mathcal{H}$ s.t. $h_i < h_j$ and $\tau_k \in \mathcal{T}$ versus:

$$H_1^{\text{MH}} : \mathbb{E}[L_{\tau_k}(y_t - \hat{y}_{\tau_k,t,h_i}) - L_{\tau_k}(y_t - \hat{y}_{\tau_k,t,h_j})] < 0 \quad (27)$$

for at least some $h_i, h_j \in \mathcal{H}$ and $\tau_k \in \mathcal{T}$. As before, define the set:

$$\mathcal{C}^{\text{MH}} = \{(h_i, h_j, \tau_k) : (h_i, h_j) \in \mathcal{H} \text{ s.t. } h_i < h_j, \tau_k \in \mathcal{T}\}.$$

We use the same test statistic as in Andrews and Soares (2010). Hereafter, denote

$$\mathcal{F}_0 = \{F_0 : H_0^{\text{MH}} \text{ holds}\}$$

as the set of null DGPs such that (26) and Assumptions C1 to C3 below hold. Also, denote, with some abuse of notation, $|\mathcal{C}^{\text{MH}}| = \kappa$ the cardinality of the set \mathcal{C}^{MH} and define $L_{\tau_k,t,h_j} \equiv L_{\tau_k}(y_t - \hat{y}_{\tau_k,t,h_j})$ as well as $L_{\tau_k,t,h_i} \equiv L_{\tau_k}(y_t - \hat{y}_{\tau_k,t,h_i})$, where $L_{s,t}$ stands for the difference $L_{\tau_k,t,h_j} - L_{\tau_k,t,h_i}$ for a specific τ_k , h_i , and h_j combination. That is, under the null hypothesis, for every $s \in \{1, \dots, \kappa\}$:

$$\mathbb{E}_{F_0} [L_{s,t}] \equiv \bar{m}_s \geq 0$$

We also let (with some abuse of notation):

$$\hat{m}_s = \frac{1}{P} \sum_{t=R+1}^T L_{s,t}$$

denote the empirical moment condition. As mentioned above, we keep implicit the (possible) depen-

dence of $L_{s,t}$ on estimated parameters. The test statistic is then given by:

$$\widehat{U}_{\text{MH}} = \sum_{s=1}^{\kappa} \left[\frac{\sqrt{P} \widehat{m}_s}{\widehat{\sigma}_{s,s}} \right]_{-}^2, \quad (28)$$

where $[x]_{-} = x1\{x < 0\}$ and $\widehat{\sigma}_{s,s}^2$ is an estimator of the diagonal element of $\overline{\Sigma}$, the variance-covariance matrix of the moment inequalities. That is, the statistic in (28) gives positive weight only to those empirical moment inequalities which are indeed violated. A corresponding HAC estimator of $\overline{\Sigma}$ may be chosen as:

$$\widehat{\Sigma}_T = \frac{1}{P} \sum_{t=R+s_T}^{T-s_T} \sum_{k=-s_T}^{s_T} \lambda_{k,t} \left(\mathbf{m}_t - \widehat{\mathbf{m}}_T \right) \left(\mathbf{m}_{t+k} - \widehat{\mathbf{m}}_T \right)'$$

where:

$$\mathbf{m}_t = \begin{pmatrix} L_{1,t} \\ \vdots \\ L_{\kappa,t} \end{pmatrix} \quad \text{and} \quad \widehat{\mathbf{m}}_T = \begin{pmatrix} \frac{1}{P} \sum_{t=R+1}^T L_{1,t} \\ \vdots \\ \frac{1}{P} \sum_{t=R+1}^T L_{\kappa,t} \end{pmatrix}.$$

Note, however, that the statistic in (28) in principle only requires estimates of the diagonal elements, which may be an advantage when \mathcal{H} and \mathcal{T} are large. The downside, of course, is that the limiting distribution is again non-pivotal as it depends on the (unknown) correlation structure between the different moment (in-)equalities. Thus, as in the Mincer-Zarnowitz tests, we will generate bootstrap critical values using the MBB of Künsch (1989) with the resampling as explained before. For each bootstrap sample, we obtain the bootstrap equivalent of $L_{s,t}$, say $L_{s,t}^b$, to construct the bootstrap statistic as:

$$\widehat{U}_{\text{MH}}^b = \sum_{s=1}^{\kappa} \left[\frac{\sqrt{P}(\widehat{m}_s^b - \widehat{m}_s)}{\widehat{\sigma}_{s,s}^b} \right]_{-}^2 1 \left\{ \frac{\widehat{m}_s}{\widehat{\sigma}_{s,s}} \leq \sqrt{2 \ln(\ln(P))/P} \right\}, \quad (29)$$

where \widehat{m}_s^b is defined as before but using the bootstrap series instead of the original sample, i.e.:

$$\widehat{m}_s^b = \frac{1}{P} \sum_{t=R+1}^T L_{s,t}^b.$$

The term $\left(\widehat{\sigma}_{s,s}^b \right)^2$, on the other hand, is the bootstrap variance estimator, which requires more careful consideration for first order validity of the bootstrap variance estimator (Götze and Künsch, 1996; Goncalves and White, 2004). More specifically, we follow Goncalves and White (2004) and use the following bootstrap variance estimator:

$$\left(\widehat{\sigma}_{s,s}^b \right)^2 = \frac{1}{K_b} \sum_{k=1}^{K_b} \frac{1}{l} \left(\sum_{i=1}^l \left(L_{s, I_k+i}^b - \frac{1}{P} \sum_{t=R+1}^T L_{s,t}^b \right) \right)^2$$

Finally, note that the second term on the right hand side of (29) implements the Generalised Moment Selection (GMS) procedure introduced by Andrews and Soares (2010) that uses information about the slackness of the sample moment conditions to infer which population moment conditions are most likely to be binding, and thus will enter into the limiting distribution. The critical value will then be based on the $(1 - \alpha)$ quantile of the empirical bootstrap distribution of $\widehat{U}_{\text{MH}}^b$ over B draws, which we

denote by $\bar{c}_{B,P,(1-\alpha)}$. We make the following additional assumptions:

C1: For all $s \in \{1, \dots, \kappa\}$ and $F_0 \in \mathcal{F}_0$, $L_{s,t}$ is strictly stationary with α -mixing coefficients satisfying the mixing condition from A1 and $\mathbb{E}_{F_0}[|L_{s,t}|^{3r}] \leq M < \infty$ with $r = 2 + \epsilon$, $\epsilon > 0$.

C2: For all $s \in \{1, \dots, \kappa\}$, $F_0 \in \mathcal{F}_0$, and some $0 < \delta \leq 2$, it holds that:

$$\frac{1}{P} \sum_{t=R+1}^T \left| \mathbb{E}_{F_0} [L_{s,t}] - \frac{1}{P} \sum_{t=R+1}^P \mathbb{E}_{F_0} [L_{s,t}] \right|^{2+\delta} = o(l^{-1-\delta/2}).$$

C3: The variance-covariance matrix $\bar{\Sigma}$ is positive definite for any $F_0 \in \mathcal{F}_0$. Moreover, it holds that:

$$\widehat{\mathbf{D}}_T^{-\frac{1}{2}} \widehat{\Sigma}_T \widehat{\mathbf{D}}_T^{-\frac{1}{2}} \xrightarrow{\text{Pr}_{F_0}} \mathbf{D}^{-\frac{1}{2}} \Sigma \mathbf{D}^{-\frac{1}{2}},$$

where $\widehat{\mathbf{D}}_T = \text{diag}(\widehat{\sigma}_{1,1}, \dots, \widehat{\sigma}_{\kappa,\kappa})$ and $\mathbf{D} = \text{diag}(\bar{\sigma}_{1,1}, \dots, \bar{\sigma}_{\kappa,\kappa})$ are $(\kappa \times \kappa)$ diagonal matrices and:

$$\widehat{\sigma}_{s,s} / \bar{\sigma}_{s,s} \xrightarrow{\text{Pr}_{F_0}} 1$$

for all $s \in \{1, \dots, \kappa\}$.

Assumption C1 corresponds to Assumption 2.1 in Goncalves and White (2002) and, together with C2, ensures the first order validity of the block bootstrap procedure in our set-up. It entails, for expositional simplicity, a homogeneity assumption across $s \in \{1, \dots, \kappa\}$, which is stronger than what is required in Goncalves and White (2002). The latter allows for considerable heterogeneity across the series. Assumption C2 on the other hand is identical Assumption A.2.2' in Goncalves and White (2004). We obtain the following result:

Theorem C.1. *Assume that C1 to C3 hold, and $T \rightarrow \infty$, $B \rightarrow \infty$, $l \rightarrow \infty$, $\frac{l}{\sqrt{T}} \rightarrow 0$. Then, under H_0^{MH} :*

$$\limsup_{T,B \rightarrow \infty} \sup_{F_0 \in \mathcal{F}_0} \text{Pr}_{F_0} \left(\widehat{U}_{\text{MH}} > \bar{c}_{B,P,(1-\alpha)} \right) \leq \alpha.$$

Theorem C.1 states that the monotonicity test proposed in this section has asymptotic size at most equal α . As pointed out in Andrews and Soares (2010), the test is non-conservative whenever some weak inequality hold with equality. Note also that Theorem C.1 required a tightening of the block length condition for the (first order) validity of the bootstrap variance estimator.

In analogy to the extension of the autocalibration test in Section 3 to a multivariate set-up, the Horizon Monotonicity test may also be extended to a group of time series. More specifically, we may be interested in testing the following null hypothesis:

$$H_0^{\text{MMH}} : \mathbb{E}[L_{\tau_k}(y_{i,t} - \widehat{y}_{i,\tau_k,t,h_i}) - L_{\tau_k}(y_{i,t} - \widehat{y}_{i,\tau_k,t,h_j})] \geq 0 \quad (30)$$

for all $h_l, h_j \in \mathcal{H}$ s.t. $h_l < h_j$, $\tau_k \in \mathcal{T}$, $i = 1, \dots, G$, versus:

$$H_1^{\text{MMH}} : \mathbb{E}[L_{\tau_k}(y_{i,t} - \widehat{y}_{i,\tau_k,t,h_i}) - L_{\tau_k}(y_{i,t} - \widehat{y}_{i,\tau_k,t,h_j})] < 0 \quad (31)$$

for at least some $h_l, h_j \in \mathcal{H}$, $\tau_k \in \mathcal{T}$, $i = 1, \dots, G$. As before, define the set:

$$\mathcal{C}^{\text{MMH}} = \{(i, h_l, h_j, \tau_k) : i \in \{1, \dots, G\}, (h_l, h_j) \in \mathcal{H} \text{ s.t. } h_l < h_j, \tau_k \in \mathcal{T}\},$$

with κ denoting the cardinality again. Then, for every difference $L_{i,\tau_k,t,h_j} - L_{i,\tau_k,t,h_l}$ with $L_{i,\tau_k,t,h_j} \equiv L_{i,\tau_k}(y_{i,t} - \hat{y}_{i,\tau_k,t,h_j})$ and $L_{i,\tau_k,t,h_l} \equiv L_{i,\tau_k}(y_{i,t} - \hat{y}_{i,\tau_k,t,h_l})$, which we denote by $\bar{L}_{s,t}$, $s \in \{1, \dots, \kappa\}$, we may test:

$$\mathbb{E} [\bar{L}_{s,t}] \equiv \bar{m}_{s,G} \geq 0, \quad s \in \{1, \dots, \kappa\}.$$

This population moment inequality can be replaced by its sample analogue:

$$\hat{m}_{s,G} = \frac{1}{P} \sum_{t=R+1}^T \bar{L}_{s,t}.$$

The test statistic \hat{U}_{MMH} and the bootstrap statistic \hat{U}_{MMH}^b is then constructed in analogy to before, with the only difference consisting in the fact that for $t = R + 1, \dots, T$, the series to be resampled comes again in array form as in the multivariate extension of Subsection 4.2. Under Assumptions C1-C3, the bootstrap distribution of \hat{U}_{MMH}^b , $b = 1, \dots, B$, provides asymptotically valid critical values that yield a test of size at most α by Theorem C.1.

Proof of Theorem C.1: Firstly, let:

$$h_s = \lim_{P \rightarrow \infty} \sqrt{P} \mathbb{E}_{F_0} [L_{s,t}].$$

Then, under the null hypothesis and Assumptions C1 and C3, it holds by arguments from the proof of Theorem 1 in Andrews and Guggenberger (2009) that for a given $F_0 \in \mathcal{F}_0$:

$$\sum_{s=1}^{\kappa} \left[\frac{\sqrt{P} \hat{m}_s}{\hat{\sigma}_{s,s}} \right]_-^2 \xrightarrow{d} \sum_{s=1}^{\kappa} \left[\sum_{j=1}^{\kappa} \omega_{s,j} Z_s + h_s \right]_-^2, \quad (32)$$

where Z_s is an element of:

$$\mathbf{Z} = \begin{pmatrix} Z_1 \\ \vdots \\ Z_\kappa \end{pmatrix} \sim N(\mathbf{0}, \mathbf{I}), \quad (33)$$

and $\omega_{s,j}$ is the square root of a generic element of the correlation matrix:

$$\mathbf{\Omega} = \mathbf{D}^{-\frac{1}{2}} \bar{\mathbf{\Sigma}} \mathbf{D}^{-\frac{1}{2}},$$

with $\mathbf{D} = \text{Diag}(\bar{\mathbf{\Sigma}})$ and $\bar{\mathbf{\Sigma}}$ denoting the population variance-covariance matrix of the moment inequalities.¹⁰ Next we need to show that the $(1 - \alpha)$ percentile of this limiting distribution is accurately approximated by the corresponding $(1 - \alpha)$ percentile of the bootstrap (limiting) distribution. To this end, note that for all $s \in \{1, \dots, \kappa\}$ it holds by Assumptions C1 and the law of iterated logarithms (Oodaira and Yoshihara, 1971, Theorem 5) that:

$$\lim_{P \rightarrow \infty} \sup \left(\frac{P}{2 \ln(\ln(P))} \right)^{\frac{1}{2}} \left(\frac{\hat{m}_s}{\hat{\sigma}_{s,s}} \right) = 1$$

¹⁰Note that in the case where $h_s = -\infty$, it follows by definition that $Z_s + h_s = -\infty$.

a.s. under \Pr_{F_0} when $E_{F_0} [L_{s,t}] = 0$, while

$$\lim_{P \rightarrow \infty} \left(\frac{P}{2 \ln(\ln(P))} \right)^{\frac{1}{2}} \left(\frac{\widehat{m}_s}{\widehat{\sigma}_{s,s}} \right) > 1$$

a.s. under \Pr_{F_0} when $E_{F_0} [L_{s,t}] > 0$. Moreover, note that by C10 for all s it holds that $\frac{\widehat{\sigma}_{s,s}}{\sigma_{s,s}} \xrightarrow{\Pr_{F_0}} 1$, so that by standard arguments:

$$\lim_{P \rightarrow \infty} \Pr_{F_0} \left(\left(\frac{P}{2 \ln(\ln(P))} \right)^{\frac{1}{2}} \left(\frac{\widehat{m}_s}{\widehat{\sigma}_{s,s}} \right) > 1 \right) = 0$$

when $E_{F_0} [L_{s,t}] = 0$, while

$$\lim_{P \rightarrow \infty} \Pr_{F_0} \left(\left(\frac{P}{2 \ln(\ln(P))} \right)^{\frac{1}{2}} \left(\frac{\widehat{m}_s}{\widehat{\sigma}_{s,s}} \right) > 1 \right) = 1$$

when $E_{F_0} [L_{s,t}] > 0$. Thus, for sufficiently large T only moment conditions holding with equality will contribute to the bootstrap limiting distribution, and the probability of eliminating a binding moment equality approaches zero as $T \rightarrow \infty$.

Moreover, from Theorem 2.2 of Goncalves and White (2002), it follows by C1 and C2 that for any $\nu > 0$ and $F_0 \in \mathcal{F}_0$:

$$\Pr_{F_0} \left(\sup_{x \in \mathbb{R}} \left| \Pr_{F_0}^b \left(\sqrt{T}(\widehat{m}_s^b - \widehat{m}_s) \leq x \right) - \Pr_{F_0} \left(\sqrt{T}(\widehat{m}_s - \overline{m}_s) \leq x \right) \right| > \nu \right) \rightarrow 0$$

for all $s \in \{1, \dots, \kappa\}$, where $\Pr_{F_0}^b$ denotes the probability measure induced by the bootstrap under F_0 . Likewise, for a given $F_0 \in \mathcal{F}_0$, by Lemma B1 in Goncalves and White (2004) we have that for any $\epsilon > 0$:

$$\Pr_{F_0} \left(\Pr_{F_0}^b \left(|\widehat{\sigma}_{s,s}^{b2} - \widehat{\sigma}_{s,s}^2| > \nu \right) > \nu \right) \rightarrow 0$$

This suggests that for any $\nu > 0$ and a given F_0 :

$$\Pr_{F_0} \left(\sup_{x \in \mathbb{R}} \left| \Pr_{F_0}^b \left(\frac{\sqrt{T}(\widehat{m}_s^b - \widehat{m}_s)}{\widehat{\sigma}_{s,s}^b} \leq x \right) - \Pr_{F_0} \left(\frac{\sqrt{T}(\widehat{m}_s - \overline{m}_s)}{\widehat{\sigma}_{s,s}} \leq x \right) \right| > \nu \right) \rightarrow 0$$

for all $s \in \{1, \dots, \kappa\}$.

Now, let $\bar{c}_{B,P,(1-\alpha)}$ be the $(1 - \alpha)$ critical value of $\widehat{U}_{\text{MH}}^b$ based on B bootstrap replications. Also, consider a sequence $\{\gamma_P\}_{P=1}^\infty$ with $\gamma_P = (\gamma_{1,P}, \dots, \gamma_{\kappa,P})$ and each $\gamma_P \in \mathcal{F}_0$ such that $\sqrt{P}\gamma_P \rightarrow \mathbf{h} = (h_1, \dots, h_\kappa)'$ and $(\ln(\ln(P)))^{-1}\sqrt{P}\gamma_P \rightarrow \boldsymbol{\xi}$ where $\mathbf{h}, \boldsymbol{\xi} \in \mathbb{R}_{-\infty}^\kappa$ with $\mathbb{R}_- = \{x \in \mathbb{R} : x \leq 0\}$ and $\mathbb{R}_{-\infty} = \mathbb{R}_- \cup \{-\infty\}$. Then, let $\bar{c}_{P,1-\alpha}$ be the $(1 - \alpha)$ critical values of $\widehat{U}_{\text{MH},\gamma}$ defined as:

$$\widehat{U}_{\text{MH},\gamma} = \sum_{s=1}^{\kappa} \left[\frac{\sqrt{P}(\widehat{m}_s - \gamma_{s,P})}{\widehat{\sigma}_{s,s}} \right]_-^2.$$

By Lemma 2(a) in the supplement of Andrews and Soares (2010), $\bar{c}_{P,1-\alpha}^b \leq \bar{c}_{P,(1-\alpha)}$ almost surely for all P for a sequence such that $\bar{c}_{P,1-\alpha}^b \xrightarrow{\Pr_{F_0}} \bar{c}_{1-\alpha}^b = \lim_{B,T \rightarrow \infty} \bar{c}_{B,P,1-\alpha}^b$ noting that the assumptions

together with the HAC estimator $\widehat{\sigma}_{s,s}$ satisfy conditions (A.2) and (A.3) of Andrews and Soares (2010) for dependent data. Also, under the drifting sequence $\{\gamma_P, P \geq 1\}$, $\lim_{P \rightarrow \infty} \bar{c}_{P,(1-\alpha)} = c_{1-\alpha}^\dagger$ which is the $(1 - \alpha)$ critical value of the limiting distribution of \widehat{U}_{MH} in Theorem C.1. The result then follows from subsequence arguments analogous to the ones used in the proof of Theorem 1(i)-(ii) in Andrews and Soares (2010). ■

D Additional Results - Empirical Application 1

D.1 Table: Individual p-values

Table 9: Individual p-values, Mincer-Zarnowitz Test, Finance Application

	$\tau = 0.01$	$\tau = 0.025$	$\tau = 0.05$	all
$h = 1$	0.000	0.092	0.294	0.006
$h = 2$	0.000	0.011	0.161	0.002
$h = 3$	0.008	0.011	0.131	0.001
$h = 4$	0.002	0.023	0.176	0.008
$h = 5$	0.005	0.001	0.141	0.010
$h = 6$	0.010	0.032	0.236	0.012
$h = 7$	0.312	0.073	0.069	0.137
$h = 8$	0.228	0.065	0.091	0.125
$h = 9$	0.030	0.113	0.029	0.028
$h = 10$	0.122	0.044	0.021	0.010
all	0.013	0.011	0.063	0.010

D.2 Tables: Intercepts and Slopes of MZ Regressions

Table 10: Intercepts, Mincer-Zarnowitz Regressions, Finance Application


	$\tau = 0.01$	$\tau = 0.025$	$\tau = 0.05$
h=1	-0.009	-0.003	-0.001
h=2	-0.009	-0.004	-0.002
h=3	-0.011	-0.006	-0.003
h=4	-0.011	-0.006	-0.002
h=5	-0.011	-0.007	-0.003
h=6	-0.011	-0.007	-0.003
h=7	-0.006	-0.006	-0.004
h=8	-0.008	-0.007	-0.004
h=9	-0.012	-0.007	-0.005
h=10	-0.011	-0.009	-0.005



Table 11: Slopes, Mincer-Zarnowitz Regressions, Finance Application

	$\tau = 0.01$	$\tau = 0.025$	$\tau = 0.05$
h=1	0.597	0.824	0.878
h=2	0.591	0.727	0.862
h=3	0.494	0.682	0.78
h=4	0.525	0.682	0.805
h=5	0.543	0.595	0.768
h=6	0.541	0.594	0.791
h=7	0.761	0.667	0.704
h=8	0.686	0.663	0.708
h=9	0.538	0.607	0.609
h=10	0.565	0.537	0.575

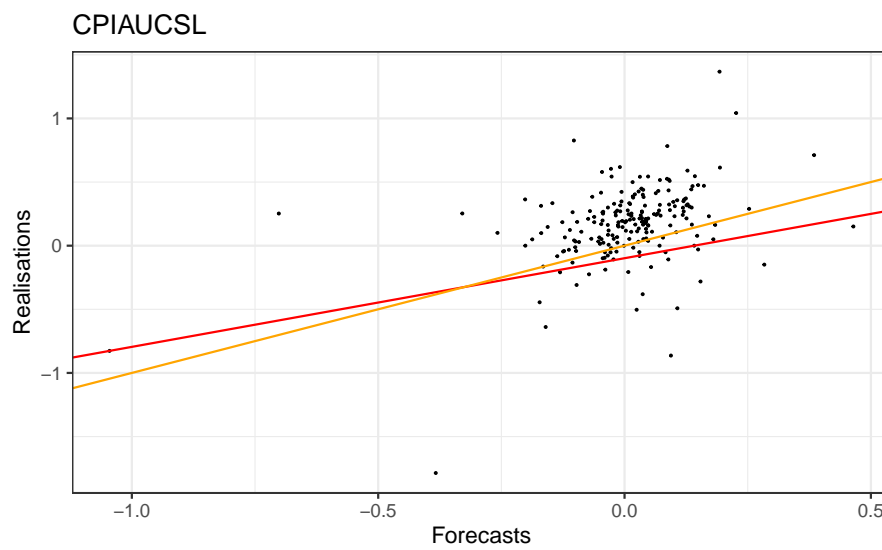
E Additional Results - Empirical Application 2

E.1 Figures: Out-of-sample Predictions for $h = 1$







E.2 Figures: MZ Regression Lines for $h = 1$ and $\tau = 0.1$




INDPRO



PAYEMS



PCEPILFE



E.3 Table: Horizon and Quantile Contributions

Table 12: Contributions to the Test Statistic by Horizon and Quantile Level

CPIAUCSL				INDPRO			
	$\tau = 0.1$	$\tau = 0.25$	$\tau = 0.5$		$\tau = 0.1$	$\tau = 0.25$	$\tau = 0.5$
$h = 1$	20.848	0.441	0.778	$h = 1$	5.720	4.196	18.968
$h = 2$	647.848	900.753	168.374	$h = 2$	1.340	6.251	4.954
$h = 3$	603.497	544.596	356.875	$h = 3$	6.800	1.778	9.367
$h = 4$	743.235	768.847	93.374	$h = 4$	3.014	1.558	21.279
$h = 5$	386.648	544.058	264.158	$h = 5$	7.245	52.469	31.338
$h = 6$	781.774	720.688	314.995	$h = 6$	43.711	15.228	96.149
$h = 7$	498.539	1070.123	325.264	$h = 7$	9.559	46.382	284.909
$h = 8$	1133.387	1758.585	721.687	$h = 8$	99.920	41.575	288.573
$h = 9$	1349.857	857.559	510.451	$h = 9$	10.912	215.386	288.707
$h = 10$	217.077	174.061	226.312	$h = 10$	108.957	330.662	310.437
$h = 11$	361.884	230.551	252.698	$h = 11$	221.439	404.278	318.836
$h = 12$	174.499	232.623	313.021	$h = 12$	226.706	406.279	313.196

PAYEMS				PCEPILFE			
	$\tau = 0.1$	$\tau = 0.25$	$\tau = 0.5$		$\tau = 0.1$	$\tau = 0.25$	$\tau = 0.5$
$h = 1$	0.280	0.316	0.046	$h = 1$	49.667	83.078	124.573
$h = 2$	0.297	0.483	0.067	$h = 2$	442.858	127.057	279.434
$h = 3$	0.836	4.267	9.897	$h = 3$	622.015	529.844	245.981
$h = 4$	0.022	1.628	20.176	$h = 4$	350.324	281.311	281.12
$h = 5$	0.170	11.645	26.387	$h = 5$	755.340	483.032	275.909
$h = 6$	0.007	9.931	38.356	$h = 6$	311.629	312.997	299.373
$h = 7$	29.247	53.612	31.944	$h = 7$	1226.937	687.718	306.836
$h = 8$	4.018	40.909	48.416	$h = 8$	711.826	431.873	313.763
$h = 9$	19.371	52.964	53.312	$h = 9$	787.800	326.569	279.552
$h = 10$	27.079	51.491	34.155	$h = 10$	388.266	337.273	365.069
$h = 11$	12.010	80.984	36.712	$h = 11$	885.383	300.752	321.414
$h = 12$	6.047	92.947	71.674	$h = 12$	659.990	337.922	340.049