

How far can we forecast?

Statistical tests of the predictive content

Jörg Breitung*
University of Cologne

Malte Knüppel†
Deutsche Bundesbank

October 10, 2017

Abstract

Forecasts are useless whenever the forecast error variance fails to be smaller than the unconditional variance of the target variable. This paper proposes tests for the null hypothesis that model based forecasts becomes uninformative beyond some limiting forecast horizon h^* . Following Diebold and Mariano (DM, 1995) we propose a test based on the comparison of the mean-squared error of the model forecast and the sample variance. We show that the resulting test does not possess a normal limiting distribution and suggest two simple modifications of the DM-type test with different limiting null distributions. Furthermore, a forecast encompassing test is developed that tends to better control the size of the test. In our empirical analysis, we apply our tests to macroeconomic forecasts from the survey of Consensus Economics. Our results suggest that forecasts of macroeconomic key variables are barely informative beyond 2–6 quarters ahead.

*Corresponding author: University of Cologne, Institute of Econometrics, 50923 Cologne, Germany.
Email: breitung@statistik.uni-koeln.de, malte.knueppel@bundesbank.de.

†The views expressed in this paper do not necessarily reflect the views of Deutsche Bundesbank or its staff.

1 Introduction

The choice of the largest forecast horizon appears to be an important issue for decision-makers. For example, in recent years, several central banks, including the Federal Reserve and the European Central Bank, decided to increase the horizon of their macroeconomic forecasts.¹ Yet, it is unclear whether the additional forecasts for these larger horizons provide valuable information, since their forecast error variance might be as large as the unconditional variance of the target variable. While statistical tools for such an assessment, based on the approach of Parzen (1981), have been proposed in the literature, formal statistical tests have not been available. The purpose of this paper is to develop such tests, thereby determining the largest informative forecast horizon.

The empirical literature reports few and differing results concerning the largest informative forecast horizon. The differences are at least partly due to different transformations, as pointed out by Galbraith and Tkacz (2007). For example, concerning quarterly GDP, they find that forecasts of quarter-on-quarter growth are barely informative beyond a forecast horizon of one quarter, but that for year-on-year forecasts this horizon increases to about 4 quarters. Concerning annual GDP growth, Isiklar and Lahiri (2007) find that forecasts are informative for horizons up to 6 quarters. Diebold and Kilian (2001) report even larger horizons for HP-filtered or linearly detrended GDP.

The tests provided in this paper are directly related to the predictability measures used in the studies mentioned. Diebold and Kilian (2001) develop a measure for predictability by comparing the loss function (say mean-squared error) of the short-run and long-run forecasts. Since our focus is on forecasting stationary time series subject to a quadratic loss function, i.e. on conditional mean forecasts, our benchmark is the unconditional mean of the time series, as proposed in Nelson (1976), Parzen (1981) and Clements and Hendry (1998).² We also discuss, however, how our approach can be applied to nonstationary time series.

The predictability measure suggested by Nelson (1976) and others is asymptotically equivalent to the R^2 from a regression of the realized values on their h step ahead forecasts and a constant. Accordingly, the null hypothesis of no predictive power is equivalent to the null hypothesis that the forecasts are not correlated with the actual values. Indeed, this is the null hypothesis underlying the encompassing version of our predictability test. In contrast, our Diebold-Mariano (1995) type test statistic directly compares the loss associated with the model-based forecast and the unconditional mean, where the unconditional mean is estimated by the mean of the evaluation sample. Therefore, the only data required for both tests are the forecasts and the actual values within the evaluation sample. This setup makes the tests applicable to forecasts from unknown models like survey forecasts.

¹See Knüppel (forthcoming) for a survey.

²Diebold and Kilian (2001) attribute this predictability measure to Granger and Newbold (1986).

This feature of the tests is important because such forecasts are often considered to be more accurate than other common forecasting approaches, as documented for inflation survey forecasts by Ang, Bekaert, and Wei (2007).

It is important to notice that, in general, the model-based forecast function nests a constant as a special case. Accordingly, comparing the model-based forecast and the unconditional mean should be treated as a nested forecast comparison in the spirit of Clark and McCracken (2001) and West (1996). It is therefore not surprising that the Diebold-Mariano type statistic has a nonstandard limiting distribution. To sidestep this difficulty we suggest a simple modification of the test statistic that results in an asymptotically χ^2 distributed random variable provided the null hypothesis is true. On the other hand, our encompassing variant of the test is equivalent to the (HAC) t -statistic from a regression of the actual values on the forecasts and a constant. We provide conditions for the standard limiting null distribution which provide a reasonable approximation in empirical practice.

The rest of this paper is organized as follows. In Section 2 we introduce our testing framework and alternative concepts of predictability are discussed in Section 3. The Diebold-Mariano type test and the encompassing test are analyzed in Sections 4 and 5. Section 6 investigates the small sample properties by means of Monte Carlo experiments and in Section 7 the proposed tests are applied to forecast of key macroeconomic variables as reported by Consensus Economics. Section 8 concludes.

2 Model framework

Let $\{y_{1+h}, \dots, y_{n+h}\}$ denote the set of n observed actual values corresponding to the model forecasts $\hat{y}_{t+h|t}$, $t = 1, \dots, n$, based on the relevant information set \mathcal{I}_t associated with time period t . We assume that y_t is generated by a stationary and ergodic stochastic process $\{Y_t\}$ and the model forecasts are realizations of the forecast generating process $Y_{t+h|t}^\theta$, where θ is the parameter vector of the forecasting model. As a simple example, assume that the target variable is generated by the univariate AR(1) process $Y_t = \alpha Y_{t-1} + u_t$ with $|\alpha| < 1$. In this example $Y_{t+h|t}^\theta = \alpha^h Y_t$ with $\theta = \alpha$. The actual forecast realization is denoted by $\hat{y}_{t+h|t} = y_{t+h|t}^{\hat{\theta}_t} = \hat{\alpha}_t^h y_t$, where $\hat{\theta}_t = \hat{\alpha}_t$ denotes some consistent estimate of θ based on the observations up to period t . Following West (1996) we distinguish two different estimation schemes. The recursive scheme fixes the starting point of the estimation sample at $t = -T + 1$ and adapts an increasing end point such that $\mathcal{S}_{T:t} = \{-T + 1, -T + 2, \dots, t\}$ and, hence, $\hat{\theta}_t$ indicates an estimate based on $t + T$ observations. The rolling-window estimation scheme fixes the size of the estimation samples to T time periods, that is, $\mathcal{S}_{t|T} = \{t - T + 1, t - T + 2, \dots, t\}$. It is important to note that we do not assume that the parameter estimates of the model and the sample size T are known. For our analysis we only need to observe the actual values $\{y_{h+1}, \dots, y_{n+h}\}$ and their h -step

ahead forecasts $\{\widehat{y}_{1+h}, \dots, \widehat{y}_{n+h}\}$.

The assumptions that characterize the process $\{Y_t\}$ are summarized in

Assumption 1 Let $Y_t = \mu + u_t$ with $u_t = \phi(L)\varepsilon_t$, $\phi(L) = 1 + \phi_1L + \phi_2L^2 + \dots$ is a lag polynomial with all roots outside the unit circle, $\sum_{i=1}^{\infty} |\phi_i| < \infty$ and ε_t is an i.i.d. white noise process with $\mathbb{E}(\varepsilon_t) = 0$ and $\mathbb{E}(\varepsilon_t^2) = \sigma_\varepsilon^2$. Furthermore $\mathbb{E}|\varepsilon_t|^{2+\delta} < \infty$ for some $\delta > 0$.

In the next section it is argued that if Y_t is integrated of order one (that is ΔY_t is stationary) the tests need to be applied to the differenced series $\Delta Y_t = Y_t - Y_{t-1}$. The assumptions of a linear process and constant variances are not essential and may be relaxed at the cost of a more demanding asymptotic framework. We are interested in testing the null hypothesis that the forecast function $Y_{t+h|t}^\theta$ is not informative for Y_{t+h} in the sense that

$$H_0 : \mathbb{E}(e_{t+h|t}^2) \geq \mathbb{E}(Y_{t+h} - \mu)^2 \quad (1)$$

where $e_{t+h|t} = Y_{t+h} - Y_{t+h|t}^\theta$ is the ‘‘theoretical’’ forecast error. Note that we can only observe the actual forecast error $\widehat{e}_{t+h|t} = y_{t+h} - \widehat{y}_{t+h|t}$. We define the *maximum forecast horizon* h^* as $h^* = \tilde{h} - 1$, where \tilde{h} the smallest value for which condition (1) is fulfilled.

It is interesting to analyze the implications of condition (1). Let us consider a simple dynamic regression model of the form

$$Y_t = \alpha + \beta X_{t-1} + e_t, \quad (2)$$

where X_t and e_t are independent white noise processes with expectation zero and with $\beta \neq 0$. If the forecast is correctly specified $Y_{t+1|t}^\theta = \alpha + \beta X_t$ and $Y_{t+h|t}^\theta = \alpha$ for $h > 1$. Thus, for $h = 1$ the condition (1) is violated but for $h = 2, 3, \dots$ the forecast becomes uninformative. Accordingly the maximum forecast horizon is $h^* = 1$.

More insight can be gained by decomposing the mean-squared prediction error (MSPE) as

$$\mathbb{E}(e_{t+h|t}^2) = \mathbb{E}(Y_{t+h} - \mu)^2 - 2\mathbb{E}(Y_{t+h} - \mu)(\widehat{Y}_{t+h|t} - \mu) + \mathbb{E}(\widehat{Y}_{t+h|t} - \mu)^2.$$

Obviously, a sufficient condition for the forecast being uninformative is $\mathbb{E}\widehat{Y}_{t+h|t} = \mu$. Another sufficient condition is that $\mathbb{E}(Y_{t+h} - \mu)(\widehat{Y}_{t+h|t} - \mu) = 0$ as $\mathbb{E}(\widehat{Y}_{t+h|t} - \mu)^2 \geq 0$. Furthermore for a rational forecast with $\mathbb{E}(e_{t+h|t} Y_{t+h|t}^\theta) = 0$ it follows that

$$\begin{aligned} \mathbb{E}(Y_{t+h} - \mu)(Y_{t+h|t}^\theta - \mu) &= \mathbb{E}(e_{t+h} + Y_{t+h|t}^\theta - \mu)(Y_{t+h|t}^\theta - \mu) \\ &= \mathbb{E}(Y_{t+h|t}^\theta - \mu)^2 \end{aligned}$$

Thus, any rational forecast with positive variance is informative. Accordingly, for a rational forecast the null hypothesis (1) is equivalent to

$$\text{cov}(Y_{t+h}, Y_{t+h|t}^\theta) = 0. \quad (3)$$

It is important to note that the maximum forecast horizon h^* can be identified by sequentially applying a consistent test for (1). The null hypothesis (1) is tested for $h = 1, 2, \dots$ until it is not rejected for the first time. Then, h^* is identified as the penultimate horizon tested. Provided that the tests are consistent, this identification is correct with probability approaching $1 - \alpha$ as $n \rightarrow \infty$ with α denoting the significance level of the test. Therefore, the significance level α of the test must tend to zero to achieve a consistent selection rule for h^* (see Remark 5 below). It should be noted that the error variance of rational forecasts are monotonously increasing with respect to the forecast horizon. Thus if some forecast is informative at some horizon h it must also be uninformative for any higher horizon and we therefore can stop the testing sequence whenever the test does not reject the first time.

3 Measuring predictability

For assessing the predictive content, Theil (1958) proposed (among other measures) the following inequality coefficient:

$$U2(h) = \sqrt{\frac{\sum_{t=1}^n (Y_{t+h} - \widehat{Y}_{t+h|t})^2}{\sum_{t=1}^n (Y_{t+h} - Y_{t+h}^0)^2}}$$

where Y_{t+h}^0 denotes some “naive forecast” (typically the no-change forecast). The model based forecast is uninformative whenever $U2(h)$ is close to unity. For a stationary variable the unconditional mean is a natural “naive” (resp. uninformative) forecast, whereas the no-change forecast is better suited for nonstationary (integrated) target variables; see e.g. Isiklar and Lahiri (2007) for an application to the survey of professional forecasters.

If the unconditional mean is employed as the benchmark, the inequality coefficient $U2(h)$ is related to the $R^2(h)$ measure proposed by Nelson (1976) and Diebold and Kilian (2001) given by

$$R^2(h) = 1 - \frac{\text{var}(\widehat{e}_{t+h|t})}{\text{var}(Y_{t+h})},$$

where $\widehat{e}_{t+h|t} = Y_{t+h} - \widehat{Y}_{t+h|t}$ denotes the model-based forecast error. In practice the variance is estimated by the sample variance such that $R^2(h) = 1 - U2(h)^2$ with $Y_{t+h}^0 = \bar{Y}_h = n^{-1} \sum_{t=1}^n Y_{t+h}$ yielding the sample analog

$$\widehat{R}^2(h) = 1 - \frac{\sum_{t=1}^n \widehat{e}_{t+h|t}^2}{\sum_{t=1}^n (Y_{t+h} - \bar{Y}_h)^2}.$$

Note however that this measure may become negative. An alternative measure with the usual properties is obtained as the R^2 from a (Mincer-Zarnowitz type) regression of Y_{t+h} on the forecast $\widehat{Y}_{t+h|t}$ yielding the square of the sample correlation between the actual values and the forecasts (see Section 5).

Diebold and Kilian (2001) proposed a generalized measure of predictability

$$Q(\mathcal{L}, h, k) = 1 - \frac{\mathbb{E}[\mathcal{L}(e_{t+h|t})]}{\mathbb{E}[\mathcal{L}(e_{t+k|t})]} \quad \text{for } k > h,$$

where $\mathcal{L}(\cdot)$ indicates the loss function and $e_{t+k|t}$ denotes the long-run prediction error. If (i) Y_t is stationary, (ii) the loss function is quadratic and (iii) $k \rightarrow \infty$, the Diebold-Kilian measure and the Nelson measure coincide. In what follows we focus on the Nelson measure and propose a test for the hypothesis $R^2(h) = 0$.

As argued by Diebold and Kilian (2001) their predictability measure is also valid for nonstationary variables, whereas in this case the Nelson measure tends to unity as $n \rightarrow \infty$, no matter of the predictive content.³ The latter approach remains valid, however, if it is applied to the differenced series. Since

$$\begin{aligned} e_{t+h|t} &= Y_{t+h} - Y_{t+h|t}^\theta \\ &= (Y_{t+h} - Y_t) - (Y_{t+h|t}^\theta - Y_t) \\ &= \left(\sum_{s=1}^h \Delta Y_{t+s} \right) - \left(\sum_{s=1}^h \Delta Y_{t+s|t}^\theta \right) \\ &= \sum_{s=1}^h \Delta e_{t+s|t}, \end{aligned}$$

with $\Delta e_{t+s|t} = \Delta Y_{t+s} - \Delta Y_{t+s|t}^\theta$ and $\Delta Y_{t+s|t}^\theta = Y_{t+h|t}^\theta - Y_{t+h-1|t}^\theta$, it follows that Y_{t+h} is not predictable whenever $\{\Delta Y_{t+1}, \dots, \Delta Y_{t+h}\}$ are jointly unpredictable. Hence, for nonstationary (integrated) time series the predictability tests are applied to the differenced series for $s \in \{1, \dots, h\}$.

4 Diebold-Mariano type test statistics

A natural test statistic for the hypothesis (1) is the DM statistic proposed by Diebold and Mariano (1995), which compares the sample MSPE of the model forecast $\widehat{Y}_{t+h|t}$ and an estimate of the sample mean used as the non-informative forecast. Since our test focuses on information of the evaluation sample we use $\bar{Y}_h = n^{-1} \sum_{t=h+1}^{n+h} Y_t$ as an estimator for $\mu = \mathbb{E}(Y_t)$. For the forecast functions $Y_{t+h|t}^\theta$ and $\widehat{Y}_{t+h|t} = Y_{t+h|t}^{\widehat{\theta}_t}$ we make the following assumption:

³More precisely, if y_t is $I(1)$ and the forecast error $e_{t+h|t}$ is $I(0)$ for fixed h , then the sample analog of the ratio $\text{var}(e_{t+h|t})/\text{var}(Y_{t+h})$ is $O_p(T^{-1})$ such that \widehat{R}^2 tends to unity.

Assumption 2 (i) Under the null hypothesis there exists some h^* such that $Y_{t+h|t}^\theta = \mu$ for all $h > h^*$. (ii) Under the null hypothesis, $u_{t+h} = Y_{t+h} - \mu$ is independent of the estimation error: $\mathbb{E}(u_{t+h}|\hat{\theta}_s - \theta) = 0$ for $s = t, t-1, t-2, \dots$. (iii) The parameters are estimated consistently with

$$\begin{aligned} a) \quad & \hat{\theta}_0 - \theta = O_p(T^{-1/2}) \\ b) \quad & \hat{\theta}_t - \hat{\theta}_0 = O_p\left(\frac{\sqrt{t}}{T}\right) \quad \text{for } t = 1, 2, \dots, n \end{aligned}$$

(iv) Let $D_{t+h}(\theta) = \partial Y_{t+h|t}^\theta / \partial \theta$ and $\bar{D}_h(\theta) = n^{-1} \sum_{t=1}^n D_{t+h}(\theta)$. Then

$$\begin{aligned} & \frac{1}{n} \sum_{t=1}^n (D_{t+h}(\theta) - \bar{D}_h)^2 \xrightarrow{p} \bar{D}^2 \quad \text{with } 0 < \bar{D}^2 < \infty \\ & \mathbb{E}|D_{t+h}(\theta)u_{t+h}|^{2+\delta} < \infty \quad \text{for some } \delta > 0 \text{ and all } t. \end{aligned}$$

(i) is the null hypothesis of the test. Part (ii) is an implication of the null hypothesis which claims that the time series is not predictable given the information set \mathcal{I}_t which includes the estimation error $\hat{\theta}_t - \theta$. Part (iii) a) suppose the usual convergence rate of the estimation error in the pre-evaluation sample $\{-T+1, \dots, 0\}$, whereas part (iii) b) limits the variation of the recursive estimation procedure.

To illustrate this assumption consider the forecast based on the regression model with $\hat{Y}_{t+h|t} = \hat{\alpha}_t + \hat{\beta}_t x_t$, where $\hat{\beta}_t$ is the least squares estimator based on the $T+t$ time periods $\{-T+1, \dots, t\}$. If x_t is stationary and, without loss of generality, $\mathbb{E}(x_t) = 0$, then $\mathbb{E}(\hat{\beta}_0 - \beta)^2 = \sigma_u^2 / (T\sigma_x^2)$, where $\sigma_x^2 = \mathbb{E}(x_t^2)$. Obviously, Assumption 2 (iii) a) is fulfilled. Since for $t < T$

$$\begin{aligned} \hat{\beta}_t - \hat{\beta}_0 &= \frac{-\left(\sum_{i=1}^t x_i^2\right)}{\left(\sum_{i=-T+1}^t x_i^2\right) \left(\sum_{i=-T+1}^0 x_i^2\right)} \left(\sum_{i=-T+1}^0 x_i u_i\right) + \frac{1}{\left(\sum_{i=-T+1}^t x_i^2\right)} \left(\sum_{i=1}^t x_i u_i\right) \\ &= \frac{O_p(t)}{O_p(T+t)O_p(T)} O_p(\sqrt{T}) + \frac{O_p(\sqrt{t})}{O_p(T+t)} \\ &= O_p\left(\frac{\sqrt{t}}{T} \cdot \sqrt{\frac{t}{T}}\right) + O_p\left(\frac{\sqrt{t}}{T}\right) = O_p(\sqrt{t}/T). \end{aligned}$$

also part (iii) b) is satisfied. In our simple example $\bar{D}^2 = \sigma_x^2$ and, thus, part (iv) is fulfilled as well.

Let us first consider a test statistic constructed in the spirit of Diebold and Mariano (1995):

$$d_h = \frac{1}{\hat{\omega}_\delta \sqrt{n}} \sum_{t=1}^n \delta_t^h, \quad (4)$$

where $\delta_t^h = \widehat{e}_{t+h|t}^2 - (Y_{t+h} - \bar{Y}_h)^2$ is the loss difference, $\bar{Y}_h = n^{-1} \sum_{t=1}^n Y_{t+h}$ denotes the sample mean, and $\widehat{\omega}_\delta^2$ denotes a consistent long-run variance estimator applied to δ_t^h . The following theorem presents the asymptotic distribution for the case that the number of observations T for estimating the parameters is large relative to the number of forecasts n :

Theorem 1 *Assume that Assumptions 1–2 hold. If $T \rightarrow \infty$, $n \rightarrow \infty$, $n/T \rightarrow 0$ we have*

$$d_{h^*} = \sqrt{n} \frac{|\bar{u}|}{2\widehat{\omega}_u} + O_p\left(\frac{n}{T}\right) \xrightarrow{d} \frac{|z|}{2},$$

where z is a standard normally distributed random variable and $\widehat{\omega}_u^2$ denotes the analogous estimator for the long-run variance $\omega_u^2 = \lim_{n \rightarrow \infty} \mathbb{E} \left(n^{-1} \sum_{t=1}^n u_{t+h} \right)^2$ of $u_{t+h} = Y_{t+h} - \mu$.

This finding gives rise to two variants of an adjusted Diebold-Mariano statistic:

Corollary 2 *Under Assumptions 1–2, $h > h^*$, $T \rightarrow \infty$, $n \rightarrow \infty$, $n/T \rightarrow 0$ it follows that*

$$2d_h \xrightarrow{d} |\mathcal{N}(0, 1)|$$

$$\widetilde{d}_h = \frac{1}{\widehat{\omega}_u^2} \sum_{t=1}^n \delta_t^h \xrightarrow{d} \chi_1^2$$

where $\widehat{\omega}_u^2$ is a consistent estimator for the long-run variance of $u_t = y_t - \mu$.

In contrast to most tests based on the χ^2 distribution, the null hypothesis is rejected for *small* values of the test statistic. For example, given a significance level of $\alpha = 0.05$, the critical value for \widetilde{d}_{h^*} equals 0.0039. The corresponding value for $2d_{h^*}$ is 0.0627.

REMARK 1: It is interesting to note that under the alternative with $\mathbb{E}(y_{t+h} - \widehat{y}_{t+h|t})^2 < \mathbb{E}(u_{t+h}^2)$ it follows that $d_h^* = O_p(\sqrt{n})$, whereas $\widetilde{d}_h^* = O_p(n)$. This is due to the fact that $\widehat{\omega}_\delta^2 = 4\bar{u}_h^2 \widehat{\omega}_u^2 + O_p(T^{-1/2}) + O_p(n^{-2})$. Since under a fixed alternative $\bar{u}_h^2 = O_p(1)$ (instead of $O_p(n^{-1})$ under the null hypothesis) the denominator of $2d_h$ changes the order of magnitude while the order of magnitude of the denominator of \widetilde{d}_h remains the same. This does not imply, however, that $2d_h$ is more powerful against alternatives in the vicinity of the null hypothesis. In fact under local alternatives both test statistics possess the same asymptotic power.

REMARK 2: To compare the small sample properties the two modifications we perform a Monte Carlo experiment, where the data are generated as $y_{t+1} = \mu + u_{t+1}$. Since the test statistic is not affected by the value of μ and the variance σ_u^2 , we set $\mu = 0$ and $\sigma_u^2 = 1$. The forecast is based on the model $\widehat{Y}_{t+1|t} = \widehat{a}_t + \widehat{b}_t X_t$, where $\theta_t = (\widehat{a}_t, \widehat{\beta}_t)'$ denotes the vector of OLS estimates from a simple regression of Y_{t+1} on X_t based on the sample

Table 1: Actual sizes for various n/T combinations ($\alpha = 0.05$)

T	$n = 25$		$n = 50$		$n = 100$		$n = 200$	
	$2d_1$	\tilde{d}_1	$2d_1$	\tilde{d}_1	$2d_1$	\tilde{d}_1	$2d_1$	\tilde{d}_1
50	0.094	0.090	0.070	0.066	0.047	0.044	0.029	0.027
100	0.112	0.107	0.094	0.090	0.070	0.066	0.047	0.044
200	0.121	0.115	0.111	0.105	0.091	0.087	0.070	0.066
500	0.125	0.118	0.122	0.115	0.115	0.110	0.099	0.094
1000	0.119	0.111	0.121	0.114	0.122	0.115	0.113	0.108
∞	0.049	0.049	0.049	0.049	0.049	0.049	0.051	0.051

Note: The limiting distribution of the statistics $2d_1$ and \tilde{d}_1 are presented in Corollary 2. For $T = \infty$ the test statistics are computed using the true parameter values. Results are based on 20,000 simulations.

$\{-T + 1, \dots, t\}$. X_t is standard normally distributed. The nominal size of all test is $\alpha = 0.05$ and 100,000 replications are used to compute the rejection rates.

From the empirical sizes for various combinations of n and T presented in Table 1 it turns out that for realistic sample sizes such as $n = 50$ and $T = 200$, say, the tests suffer from a substantial size bias. This is not surprising as the critical value for the χ_1^2 distribution is very close to zero, and, thus, a large amount of probability mass is located in the vicinity of the critical value. Accordingly, even small, asymptotically negligible terms may have large effects on the actual size.

REMARK 3: It is interesting to analyze the properties of the test for situations where part (ii) of Assumption 2 is violated. Assume that the model forecast is uninformative and biased with $E(Y_{t+h|t}) = m_t \neq \mu$. It follows from the proof of Theorem 1 that

$$\frac{1}{n} \sum_{t=h+1}^{n+h} \delta_t^h = \bar{u}_h^2 + \frac{1}{n} \sum_{t=1}^n (m_t - \mu)^2 + O_p(T^{-1}).$$

Therefore if the forecast is biased, the test becomes conservative in the sense that the rejection probability tends to zero as $n \rightarrow \infty$.

Next, assume that the uninformative forecast $Y_{t+h|t}^\theta$ is unbiased but different from a constant: $Y_{t+h|t}^\theta = \mu + v_t^h$, where $\text{var}(v_t^h) = \tau_h^2 > 0$. For example assume that the forecast $Y_{t+h|t}^\theta = \beta X_t$ does not account for the constant such that $\beta = \mu/\mathbb{E}(X_t)$ and $v_t^h = \beta[X_t - \mathbb{E}(X_t)]$. Note that this forecast violates Assumption 2 (ii) as $\mathbb{E}[(Y_{t+h} - \beta X_t)X_t] = \beta \text{var}(X_t) \neq 0$. It is not difficult to see that in this case $n^{-1} \sum_{t=h+1}^{n+h} \delta_t^h = \bar{u}_h^2 + \tau_h^2$ resulting again in a conservative test.

5 Encompassing tests

To overcome the small sample problems of the modified test statistic we consider a variant of the test based on the encompassing principle. As shown in section 2, a test of the null hypothesis (1) is equivalent to the null hypothesis

$$H'_0 : \mathbb{E}(Y_{t+h} - Y_{t+h|t}^\theta)(Y_{t+h|t}^\theta - \mu) = 0$$

whenever the forecast is rational with $\mathbb{E}(Y_{t+h} - Y_{t+h|t}^\theta | Y_{t+h|t}^\theta) = 0$. Accordingly, the null hypothesis can be rejected if the correlation between Y_{t+h} and $\widehat{Y}_{t+h|t}$ is positive. Another view on the relationship between this approach and the DM statistic emerge from the decomposition

$$\begin{aligned} \sum_{t=1}^n \delta_t^h &= \sum_{t=1}^n \left[Y_{t+h} - \bar{Y}_h - (\widehat{Y}_{t+h|t} - \bar{Y}_h) \right]^2 - (Y_{t+h} - \bar{Y}_h)^2 \\ &= \sum_{t=1}^n (\widehat{Y}_{t+h|t} - \bar{Y}_h)^2 - 2 \sum_{t=1}^n (Y_{t+h} - \bar{Y}_h)(\widehat{Y}_{t+h|t} - \bar{Y}_h). \end{aligned}$$

Since the first term on the right hand side is non-negative, it does not improve the power of the test and can be neglected. Thus, the DM type test statistic is mainly driven by the covariance between Y_{t+h} and $\widehat{Y}_{t+h|t}$.

This gives rise to a one-sided t -test of $\beta_{1,h} = 0$ vs. $\beta_{1,h} > 0$ in the Mincer-Zarnowitz regression

$$Y_{t+h} = \beta_{0,h} + \beta_{1,h} \widehat{Y}_{t+h|t} + \epsilon_{t+h}, \quad (5)$$

where the error is typically autocorrelated up to lag $h - 1$ due to the overlapping forecast horizon. This test admits an interpretation as a forecast encompassing test (cf. Chong and Hendry, 1986, and Clements and Hendry, 1993) that is based on a convex combination of the model based forecast $Y_{t+h|t}$ and the unconditional mean \bar{Y}_h :

$$\begin{aligned} Y_{t+h} &= \lambda_h \widehat{Y}_{t+h|t} + (1 - \lambda) \bar{Y}_{t+h} + \epsilon_{t+h} \\ Y_{t+h} - \bar{Y}_h &= \lambda_h (\widehat{Y}_{t+h|t} - \bar{Y}_h) + \epsilon_{t+h} \end{aligned}$$

and, therefore, a test of $\beta_{1,h} = 0$ is equivalent to a test of $\lambda_h = 0$.

In our asymptotic analysis we focus on the LM type test statistic:

$$Q_h = \frac{1}{\widehat{\omega}_\epsilon \sqrt{n}} \sum_{t=1}^n \xi_t^h \quad (6)$$

where

$$\xi_t^h = (Y_{t+h} - \bar{Y}_h)(\widehat{Y}_{t+h|t} - \bar{Y}_h)$$

Table 2: Actual sizes for various n/T combinations ($\alpha = 0.05$)

T	$n = 25$		$n = 50$		$n = 100$		$n = 200$	
	$\beta_{1,1}$	ϱ_1	$\beta_{1,1}$	ϱ_1	$\beta_{1,1}$	ϱ_1	$\beta_{1,1}$	ϱ_1
50	0.027	0.021	0.020	0.016	0.017	0.014	0.014	0.013
100	0.030	0.023	0.022	0.018	0.017	0.016	0.014	0.014
200	0.035	0.027	0.025	0.021	0.020	0.018	0.016	0.015
500	0.039	0.031	0.029	0.025	0.023	0.020	0.019	0.017
1000	0.045	0.037	0.033	0.030	0.027	0.025	0.022	0.021

Note: $\beta_{1,1}$ denotes the regression-based test using (5), ϱ_1 denotes the LM-type test based on (6). Results are based on 20,000 replication.

and $\widehat{\omega}_\xi^2$ denotes the corresponding long-run variance

$$\widehat{\omega}_\xi^2 = \widehat{\gamma}_0^\xi + 2 \sum_{j=1}^k w_j^k \widehat{\gamma}_j^\xi$$

$$\widehat{\gamma}_j^\xi = \frac{1}{n} \sum_{t=j+1}^n \xi_t^h \xi_{t-j}^h .$$

In practice any other asymptotically equivalent test such as the usual t -test of $\beta_{1,h} = 0$ in the regression (5) with appropriate HAC standard errors can be used.

A technical problem with a regression like (5) is that under the null hypothesis the regressor $\widehat{Y}_{t+h|t}$ and the constant are asymptotically collinear. To sidestep this problem we show in the proof of Theorem 3 that if $n/T \rightarrow 0$

$$\widehat{Y}_{t+h|t} - \bar{Y}_h \approx (\widehat{\theta}_0 - \theta)(D_{t+h}(\theta) - \bar{D}_h(\theta))$$

and, thus, the test of $\beta_{1,h} = 0$ in regression (5) is asymptotically equivalent to the test of $\beta_{1,h}^* = 0$ in the regressions

$$Y_{t+h} = \beta_0^* + \beta_{1,h}^* D_{t+h}(\theta) + \eta_{t+h}$$

or $Y_{t+h} = \beta_{0,h}^* + \beta_{1,h}^* D_{t+h}(\widehat{\theta}) + \widetilde{\eta}_{t+h}$,

where $\beta_{1,h}^* = (\widehat{\theta}_0 - \theta)\beta_{1,h}$. The details are provided in

Theorem 3 *Under Assumption 1–2, a recursive forecasting scheme with $h > h^*$, $T \rightarrow \infty$, $n \rightarrow \infty$ and $n/T \rightarrow 0$ we have $\varrho_h \xrightarrow{d} \mathcal{N}(0, 1)$, where ϱ_h is defined in (6).*

REMARK 4: In contrast to the DM-type test, the encompassing test turns out to be conservative for all combinations of n and T presented in Table 2. While the empirical

sizes for a given value of n vary depending on T , these variations are smaller than those of the DM-type tests.

REMARK 5: As mentioned above, a consistent selection rule for the maximal forecast horizon h^* requires that the size of the test tends to zero as $n \rightarrow \infty$. One possibility is to apply the critical value $\log(n)/2$. It is not difficult to see that under the alternative $\rho_h = O_p(n^{1/2})$ such that under the alternative $h \leq h^*$ we obtain $\lim_{n \rightarrow \infty} P(\rho_h < -\log(n)/2) = 1$, whereas for $h > h^*$ we have $\lim_{n \rightarrow \infty} P(\rho_h < -\log(n)/2) = 0$. Thus the decision rule that selects the last rejection in the sequence of tests with $h = 1, 2, \dots$ is weakly consistent. Note that for $n = 27$ the critical value $-\log(27)/2 = -1.65$ is similar to the one-sided 0.05 critical value of a standard normal distribution.

6 Local power

In order to gain some insight into the relative power of the two different tests we analyse local power against a suitable sequence of local alternatives. Consider the alternative of an informative forecast with

$$Y_{t+1} = \mu + \beta X_t + u_{t+1}$$

where X_t is a (demeaned) i.i.d. regressor with $\mathbb{E}(X_t) = 0$ and $\mathbb{E}(X_t^2) = \sigma_x^2 > 0$, u_t is white noise with $\mathbb{E}(u_t) = 0$, $\mathbb{E}(u_t^2) = \sigma_u^2$ and u_{t+1} is independent of $\{X_t, X_{t-1}, \dots\}$. As $T \rightarrow \infty$ we have $\hat{\theta}_t \xrightarrow{p} \theta$, $\mathbb{E}(Y_{t+1} - Y_{t+1|t}^\theta)^2 = \sigma_u^2$, where $Y_{t+1|t}^\theta = \mu + \beta X_t$ and $\mathbb{E}(Y_{t+1} - \mu)^2 = \sigma_u^2 + \beta^2 \sigma_x^2$. If $\beta \neq 0$ the forecast is informative and \tilde{d}_1 and $\hat{\varrho}_1$ are $O_p(n^{-1/2})$. Accordingly both tests are consistent against fixed alternatives $\beta \neq 0$. The asymptotic power of the test can be studied by considering a local alternative of the form $\beta = c/\sqrt{n}$. The asymptotic distribution of the two test statistics are presented in

Theorem 4 *Under the sequence of alternatives $\beta = c/\sqrt{n}$, $X_t \sim iid(0, \sigma_x^2)$, Assumptions 1 – 2 and $n/\sqrt{T} \rightarrow 0$ it follows that*


$$\tilde{d}_1 \xrightarrow{d} z_1^2 - 2\lambda z_2 - \lambda^2 \tag{7}$$

$$\hat{\varrho}_1 \xrightarrow{d} \text{sign}(c)z_2 + \lambda, \tag{8}$$

where $\lambda^2 = c^2 \sigma_x^2 / \sigma_u^2$ denotes the signal-to-noise ratio and z_1 and z_2 represent two independent standard normally distributed random variables.

Accordingly, the two test are not asymptotically equivalent; see Figure 1 for a comparison of the resulting local power curves. The DM type test is more powerful in the vicinity of the null hypothesis, whereas the relative power of the encompassing test increases when c

Figure 1: Local power curves



Note: Broken line: modified DM test. Solid line: encompassing test.

or the variance ratio σ_x^2/σ_u^2 gets large. It should also be noted that the power curves are symmetric with respect to the parameter c as the distribution remains the same if z_2 is replaced by $-z_2$ in (7) and (8).

7 Monte Carlo Results

In order to gain insights into the small sample behaviour of the tests, we conduct Monte Carlo experiments based on the cases displayed in Table 7. In the first three cases, univariate models are considered, whereas the last two cases refer to simple multivariate models. The first two forecast models are misspecified, because the data-generating processes (DGP) are MA(1) and MA(2) processes, but the forecast models assume an AR(1) process. Note that the respective null hypotheses $h^* = 1$ and $h^* = 2$ are nevertheless correct. Forecasts are made in a direct manner, i.e. for each forecast horizon, the target variable y_{t+h} is regressed on the explanatory variables known at time t . We calculate the standard errors according to Newey and West (1987) using the automatic lag length selection procedure proposed by Andrews (1991) and use the significance level $\alpha = 0.05$.

Table 4 displays the results for AR(1) forecasts, when the data is generated by an MA(1) model, hence $h^* = 1$. The evaluation sample includes $n = 50$ or $n = 100$ recursive forecasts and the initial estimation samples are based on $T = 100$ observations. The tests are conducted sequentially for the forecast horizons $h = 1, 2, 3, 4$. The last forecast horizon where the test rejects is identified as horizon h^* . If the test does not reject for any horizon, $h^* \geq 4$. In addition to the tests presented above, the classical DM-test

Table 3: Cases considered for Monte Carlo simulations

cases	DGP	forecast model
$MA(1)-AR(1)$	$y_t = \varepsilon_t + 0.5\varepsilon_{t-1}$	$\hat{y}_{t+h} = \hat{\theta}_1^h + \hat{\theta}_2^h y_t$
$MA(2)-AR(1)$	$y_t = \varepsilon_t + 0.5\varepsilon_{t-1} + 0.3\varepsilon_{t-2}$	$\hat{y}_{t+h} = \hat{\theta}_1^h + \hat{\theta}_2^h y_t$
$AR(1)-AR(1)$	$y_t = 0.8y_{t-1} + \varepsilon_t$	$\hat{y}_{t+h} = \hat{\theta}_1^h + \hat{\theta}_2^h y_t$
multivar. 1	$y_t = 0.5x_{t-1} + \varepsilon_t$	$\hat{y}_{t+h} = \hat{\theta}_1^h + \hat{\theta}_2^h x_t$
multivar. 2	$y_t = 0.5x_{t-1} + 0.3x_{t-2} + \varepsilon_t$	$\hat{y}_{t+h} = \hat{\theta}_1^h + \hat{\theta}_2^h x_t + \hat{\theta}_3^h x_{t-1}$

Note: ε_t and x_t are $iidN(0,1)$.

is considered by using a standard normal distribution for the test statistic specified in equation (4).

With $n = 50$, the DM-type and encompassing tests developed to detect the maximum forecast horizon have a power of at least 0.81 at $h = 1$, whereas the classical DM test attains 0.26 only. With $n = 100$, the power of the DM-type and encompassing tests is close to one. The classical DM test achieves 0.57 and its size is close to zero, whereas both DM-type tests are over-sized. The LM-type encompassing test based on (6) is slightly conservative, whereas the regression-based test using (5) has almost correct size and the highest power of all tests considered.

For $n = 50$ the regression-based encompassing test is also the most successful test in finding the maximum forecast horizon, identifying it correctly in 86% of the replications, while the other tests (except for the classical DM-test) do so in about 80% of the cases. With $n = 100$, both regression-based tests work very well, achieving a correct identification of h^* in more than 95% of the cases. Note that 95% would be attained by a correctly-sized test and an infinite evaluation sample. The DM-type tests find the correct h^* in about 90% of the replications. The estimators for h^* are nearly median unbiased for all DM-type and encompassing tests, but possess a small downward mean bias for $n = 50$ and a small upward mean bias for $n = 100$. The classical DM test estimates h^* with a strong downward bias due to its lack of power. With a rolling estimation window instead of a recursive estimation scheme, the results turn out to be almost unchanged (see Table 9 in Appendix C).

For the remaining Monte Carlo experiments, we do not report results for the DM-type test based on $2d_h$, because it performs very similar to \tilde{d}_h . Moreover, we focus on the regression-based encompassing test, because it has good size properties and more power

Table 4: Results for case ‘MA(1)-AR(1)’

forecast horizon h	0	1	2	3	4	0	1	2	3	4
	$T = 100, n = 50$					$T = 100, n = 100$				
MSPE / Variance	0.89	1.05	1.05	1.05		0.87	1.03	1.03	1.03	
rejections										
DM-type tests										
\tilde{d}_h	0.87	0.10	0.09	0.10		0.98	0.08	0.08	0.08	
$2d_h$	0.88	0.10	0.10	0.10		0.98	0.08	0.08	0.08	
encompassing tests										
$\beta_{1,h}$	0.92	0.06	0.05	0.05		1.00	0.04	0.04	0.03	
ϱ_h	0.81	0.03	0.02	0.02		0.99	0.03	0.02	0.02	
classical DM test	0.26	0.00	0.00	0.00		0.57	0.00	0.00	0.00	
\hat{h}^*										
DM-type tests										
\tilde{d}_h	0.13	0.79	0.05	0.02	0.01	0.02	0.90	0.05	0.02	0.00
$2d_h$	0.12	0.79	0.06	0.02	0.01	0.02	0.90	0.06	0.02	0.00
encompassing tests										
$\beta_{1,h}$	0.08	0.87	0.04	0.01	0.00	0.00	0.96	0.03	0.01	0.00
ϱ_h	0.19	0.80	0.01	0.00	0.00	0.01	0.97	0.02	0.00	0.00
classical DM test	0.74	0.26	0.00	0.00	0.00	0.43	0.57	0.00	0.00	0.00

Note: The values displayed in the category ‘rejections’ denote the percentage of rejections for each horizon h . The values displayed in the category ‘ \hat{h}^* ’ denote the percentage of cases in which h is identified as the maximum forecast horizon. The estimation is carried out recursively, and T denotes the number of observations used for the first parameter estimation. n is the number of observations for evaluation. The significance level is set to $\alpha = 0.05$. ‘classical DM test’ refers to the test statistic proposed by Diebold and Mariano (1995). ‘MSPE’ is the mean-squared prediction error. The in-sample variance in the MSPE-variance-ratio is calculated dividing by n . Bold entries refer to the true h^* . If a test rejects for all horizons, \hat{h}^* is set equal to the largest horizon $h = 4$. $2d_h$ denotes the test statistic distributed as $|\mathcal{N}(0,1)|$, \tilde{d}_h the test statistic distributed as χ_1^2 under the null. $\beta_{1,h}$ denotes the regression-based test using (5), ϱ_h denotes the LM-type test based on (6). Results are based on 20,000 simulations.

Table 5: Results of the remaining cases

forecast horizon h	0	1	2	3	4	0	1	2	3	4
	$T = 100, n = 50$					$T = 100, n = 100$				
	<i>MA(2)-AR(1)</i>									
MSPE / Variance	0.82	1.01	1.07	1.07		0.79	0.99	1.04	1.04	
rejections										
DM-type test	0.91	0.39	0.10	0.10		0.99	0.58	0.09	0.09	
encompassing test	0.95	0.34	0.07	0.06		1.00	0.55	0.05	0.05	
\hat{h}^*										
DM-type test	0.09	0.53	0.35	0.02	0.02	0.01	0.41	0.54	0.02	0.02
encompassing test	0.05	0.62	0.31	0.01	0.01	0.00	0.45	0.53	0.01	0.01
	<i>AR(1)-AR(1)</i>									
MSPE / Variance	0.44	0.72	0.90	1.01		0.40	0.65	0.82	0.92	
rejections										
DM-type test	0.99	0.85	0.62	0.42		1.00	0.98	0.86	0.68	
encompassing test	1.00	0.95	0.76	0.50		1.00	1.00	0.95	0.78	
\hat{h}^*										
DM-type test	0.01	0.14	0.23	0.20	0.42	0.00	0.02	0.12	0.19	0.68
encompassing test	0.00	0.05	0.20	0.26	0.50	0.00	0.00	0.05	0.18	0.77
	multivar. 1									
MSPE / Variance	0.83	1.04	1.04	1.04		0.82	1.02	1.02	1.02	
rejections										
DM-type test	0.93	0.09	0.09	0.09		0.99	0.07	0.07	0.06	
encompassing test	0.95	0.03	0.03	0.03		1.00	0.02	0.02	0.02	
\hat{h}^*										
DM-type test	0.07	0.84	0.07	0.01	0.00	0.01	0.93	0.06	0.01	0.00
encompassing test	0.05	0.93	0.03	0.00	0.00	0.00	0.98	0.02	0.00	0.00
	multivar. 2									
MSPE / Variance	0.78	0.98	1.05	1.05		0.77	0.96	1.03	1.03	
rejections										
DM-type test	0.96	0.60	0.11	0.11		1.00	0.75	0.07	0.07	
encompassing test	0.98	0.50	0.04	0.04		1.00	0.73	0.03	0.03	
\hat{h}^*										
DM-type test	0.04	0.36	0.52	0.05	0.03	0.00	0.25	0.68	0.04	0.02
encompassing test	0.02	0.49	0.47	0.02	0.01	0.00	0.27	0.71	0.02	0.01

Note: The DM-type test uses \tilde{d}_h , the encompassing test employs $\beta_{1,h}$. For further information, see Table 4.

than the LM-type encompassing test. Results for all cases except MA(1)-AR(1) (which are already presented in 4) are reported in Table 5. We do not consider the classical DM test, because of its problems in detecting h^* documented previously.

Concerning the case MA(2)-AR(1), the second MA-coefficient equals 0.3 only, making it relatively difficult for the tests to reject at $h = 2$. With $n = 50$, the ratio of the mean-squared prediction error (MSPE) to the evaluation-sample variance⁴ marginally exceeds 1, and hardly falls short of 1 even with $n = 100$. With $n = 50$, the DM-type test rejects in about 40% of the replications at $h = 2$. The encompassing test attains about 35%. With $n = 100$ all rejection frequencies increase by about 20 percentage points. Both tests tend to reject too often for $h \geq 3$, but the encompassing test has almost the correct size, whereas the DM-type test tends to reject too often. The maximum forecast horizon is identified correctly by both tests in about 30% of the replications with $n = 50$, whereas for more than 50% \hat{h}^* equals 1. With $n = 100$, \hat{h}^* equals h^* in about 50% of the simulations, yielding median-unbiased estimates of h^* . However, \hat{h}^* has a notable downward mean bias even with $n = 100$.

For an AR(1) process, h^* is not finite. The tests should therefore reject the null for as many horizons as possible. In this setting, the encompassing test performs best, with its rejection frequency falling from 1 at $h = 1$ to 0.5 at $h = 4$ for $n = 50$. With $n = 100$, the rejection frequency at $h = 4$ attains about 0.8. Based on the encompassing test, one would arrive at the correct conclusion that h^* is larger than 3 with a probability of about 0.5 for $n = 50$ and with a probability of about 0.8 for $n = 100$. With the DM-type test, these probabilities are about 10 percentage points lower.

The results in the multivariate case with one regressor are similar to those for the case MA(1)-AR(1). Again the two proposed test do a good job in estimating the forecast horizon h^* . As the *multivar. 1* setup is conformable with the local power analysis of Section 6 it is interesting to compare the results with the theoretical findings presented in Theorem 4. For the corresponding value $c = 0.5\sqrt{50} = 3.54$ and $\lambda = c \cdot 1$ the power for the two test is 0.9450 (for the DM-type test) and 0.9716 (for the encompassing test). These values correspond well to the empirical values of 0.93 and 0.95 in Table 5.

Qualitative similar results emerge when comparing the results of the multivariate case with two regressors (*multivar. 1*) to the univariate case MA(2)-AR(1) with $h^* = 2$.

8 Empirical Results

For the empirical application of the tests, we employ quarterly survey forecasts collected by Consensus Economics. The mean of the forecasts across all panelists is known to be a very accurate forecast, as documented, for example, by Ang, Bekaert, and Wei (2007) for

⁴The variance of the evaluation sample is calculated dividing by n , so that it equals the MSPE of the evaluation-sample mean.

inflation forecasts. We consider forecasts of real GDP growth and real private consumption growth, because these are the only quarter-on-quarter growth rates in the survey. For the indices of consumer prices, only forecasts for quarterly year-on-year rates are available. Given the importance of inflation forecasts, we also include these forecasts in our analysis. However, given the year-on-year definition, and denoting the forecast horizon for the current quarter, i.e. the nowcast by $h = 0$, we can expect to find $h^* \geq 2$. This is because knowledge about past values of the price index enables the forecasters to mechanically produce forecasts which have lower mean-squared prediction errors than the unconditional mean up to $h = 2$.⁵ In addition to these variables, we also investigate the forecasts of the end-quarter values of the 3-month interest rate. Since interest rates show signs of non-stationarity in the sample under study, we use the first differences of this variable. The countries under study are the United States, the euro area (labeled ‘Eurozone’ by Consensus Economics), Japan, Germany, the United Kingdom, Italy, Canada, and France.

Since Consensus Economics also provides data for recent quarters, we can employ this real-time data for the evaluation of the forecasts. We use the second vintage of all variables mentioned.

Considering forecasts for up to $h = 6$ quarters ahead, the balanced sample of forecasts and realizations starts in the third quarter of 1996 and ends in the first quarter of 2016, yielding a sample size of $n = 79$. However, the sample sizes for individual variables can be smaller, mainly due to changes in the survey. For example, in the beginning of the sample, the survey switched from asking for West German variables to variables for the reunified Germany, and we only consider the latter forecasts. The 3-month interest rate to be forecasted for Germany, Italy and France changed in the first quarter of 1999 from country-specific rates to the Euribor. Only the Euribor forecasts enter our analysis. While they can be expected to be similar across the three countries, differences might emerge, for instance, due to the smaller number of forecasters for Italy. For Japan, the target variable of the interest rate forecasts changed from the 3-month Yen certificate of deposit to the TIBOR in the second quarter of 2010, and we only use the former forecasts.

In some countries, several large changes of the value-added tax rate (VAT) occurred. Since these changes are commonly announced well in advance, their occurrence can have a major impact on the predictability of inflation. In addition, the growth rate of real private consumption tends to be strongly negatively correlated with the VAT rate change contemporaneously, and in addition, there is a strong positive correlation in the quarter prior to the VAT rate change. We use the following rule in order to limit the impact of

⁵The year-on-year rate for $h = 2$ equals the sum of the quarter-on-quarter rates for $h = -1, 0, 1, 2$. Using the observed quarter-on-quarter rate for $h = -1$ and the unconditional mean as the forecast of the quarter-on-quarter rates for the latter three horizons yields a MSPE for the year-on-year rate forecast for $h = 2$ which is lower than the variance of the year-on-year rates by construction. If information on the current quarter is available, the maximum forecast horizon must be equal to or larger than 3.

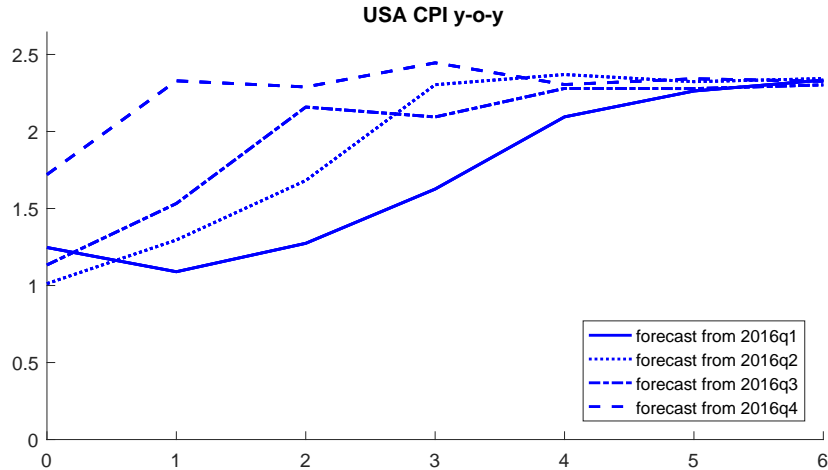


Figure 2: Forecasts for year-on-year US CPI inflation rates. The number on the x-axis denotes the forecast horizon in quarters with 0 being the nowcast.

these changes on our analysis: For countries with at least two VAT rate changes of at least 2 percentage points, we shorten the sample such that these changes are excluded. This rule leads to sample modifications for Japan and the UK. We are going to report results obtained without this modification in the text, but not in the following figures and tables. More details on the variables and samples entering our analysis are given in Appendix B.

The quarterly forecasts are usually gathered in the first half of the last month of a quarter. Therefore, the forecasters can be expected to have information about the variable of interest in the current quarter, i.e. for the forecast (resp. nowcast) horizon $h = 0$. Note that for the q-o-q growth rate of consumption in the UK the nowcast of the current forecast is uninformative and hence $h^* = -1$. Concerning inflation, at least the inflation rate for the first month of the current quarter should be known when the forecast is made. This implies that one can expect $h^* \geq 3$ for the year-on-year inflation rates.

As an example consider the inflation forecasts for the United States provided in 2016 as presented in 2. Forecasts for the current quarter are nowcasts, corresponding to $h = 0$. What is striking about the forecasts for longer horizons is that they tend to settle at a value of about 2.3 which is almost identical to the mean of inflation in the evaluation period, being equal to 2.27 percent.

The empirical maximum forecast horizons \hat{h}^* determined by the tests are shown in

Table 6: Maximum forecast horizons in quarters determined by DM-type and encompassing tests

	US	EA	JP	DE	UK	IT	CA	FR	median
GDP q-o-q									
DM-type test	2	2	1	1	3	1	1	2	1.5
encompassing test	2	2	1	2	3	5	1	4	2
CPI y-o-y									
DM-type test	3	5	3	2	2	3	4	3	3
encompassing test	3	3	4	3	3	3	4	3	3
PrivCons q-o-q									
DM-type test	3	1	0	0	-1	1	0	2	0.5
encompassing test	3	3	0	3	3	3	1	5	3
d(3m rate)									
DM-type test	1		0	3	2	2	1	2	2
encompassing test	2		0	2	6	3	1	2	2

Note: The DM-type test uses \tilde{d}_h , the encompassing test employs $\beta_{1,h}$. ‘GDP q-o-q’ denotes quarter-on-quarter growth rates of real GDP, ‘CPI y-o-y’ year-on-year growth rates of consumer prices, ‘PrivCons q-o-q’ quarter-on-quarter growth rates of real private consumption, and ‘d(3m rate)’ quarter-on-quarter changes of the end-quarter 3-month interest rate. The abbreviations used for the countries are ‘US’ for the United States, ‘EA’ for the euro area, ‘JP’ for Japan, ‘DE’ for Germany, ‘UK’ for the United Kingdom, ‘IT’ for Italy, ‘CA’ for Canada, and ‘FR’ for France. For further information on the variables, see the text and Appendix B.

Table 6. The sequential p -values of the tests giving rise to these values of \hat{h}^* are displayed in Figures 3 to 6. Notably, \hat{h}^* is virtually always smaller than the largest forecast horizon of $h = 6$. The encompassing test implies larger values of \hat{h}^* than the DM-type test in several cases. This may be due to the lower power of the latter test reported in the previous section.

For the quarter-on-quarter growth rates of real GDP growth, \hat{h}^* tends to range from 1 to 3 quarters. Only for Italy and France, the encompassing test leads to larger values. The median of \hat{h}^* is 1.5 quarters according to the DM-type test, and 2 quarters according to the encompassing test.

Concerning the year-on-year growth rates of the CPI, \hat{h}^* is mostly equal to 3 quarters, which is also the median across countries according to both tests. Only for Canada, both tests indicate a higher value of $\hat{h}^* = 4$. For the Euro area and Japan, at least one test indicates a larger value. Given the considerations with respect to year-on-year rates and nowcasts made at the end of $h = 0$, these results indicate substantial difficulties in making informative inflation forecasts. Using the full sample for Japan and the UK leaves the results for the UK unchanged, but leads to $\hat{h}^* \geq 6$ for Japan according to both tests.

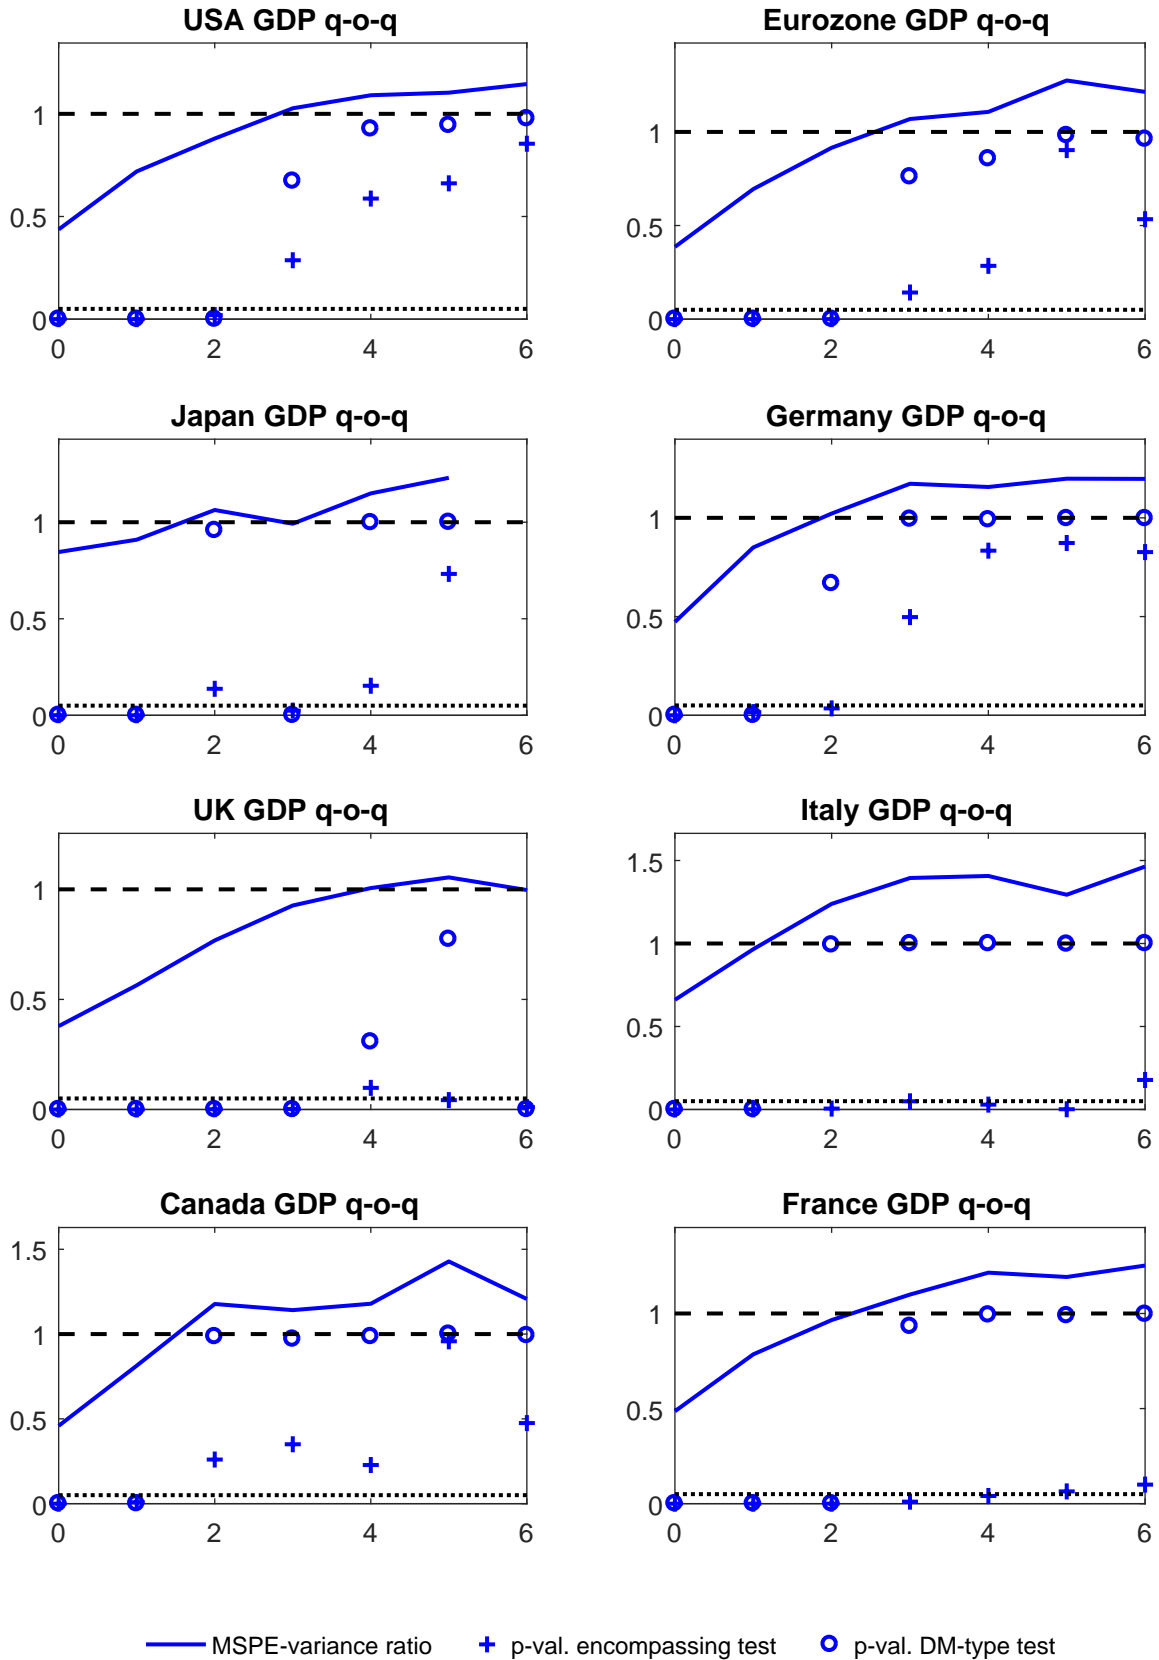


Figure 3: Test results for quarter-on-quarter growth rates of real GDP. The number on the x-axis denotes the forecast horizon in quarters with 0 being the nowcast. The dotted line is at 0.05, corresponding to the significance level of the tests. The dashed line is at 1. The solid line indicates the MSPE-variance ratio. The DM-type test uses \tilde{d}_h , the encompassing test employs $\beta_{1,h}$. The maximum forecast horizon \hat{h}^* identified by a test equals the horizon before the smallest horizon for which the p-value of the test exceeds 0.05.

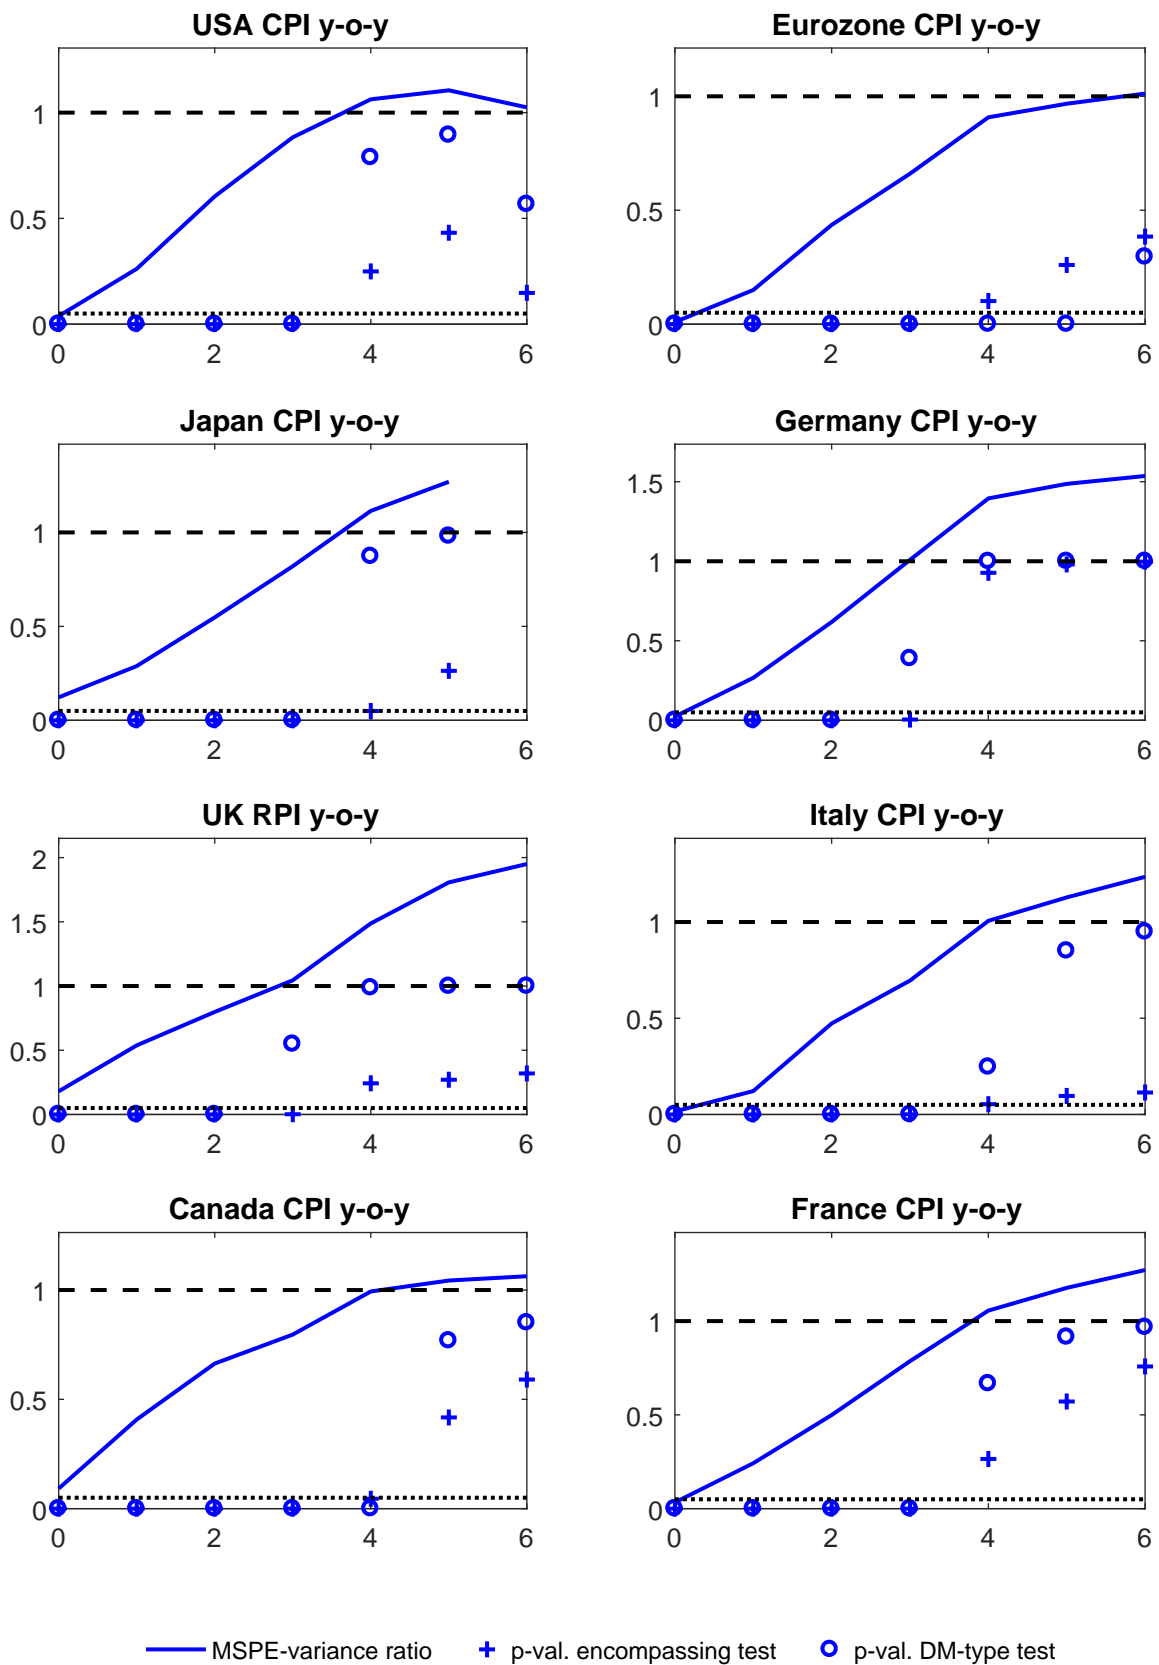


Figure 4: Test results for year-on-year growth rates of the CPI (the RPI in the case of the UK). For further explanations, see Figure 3.

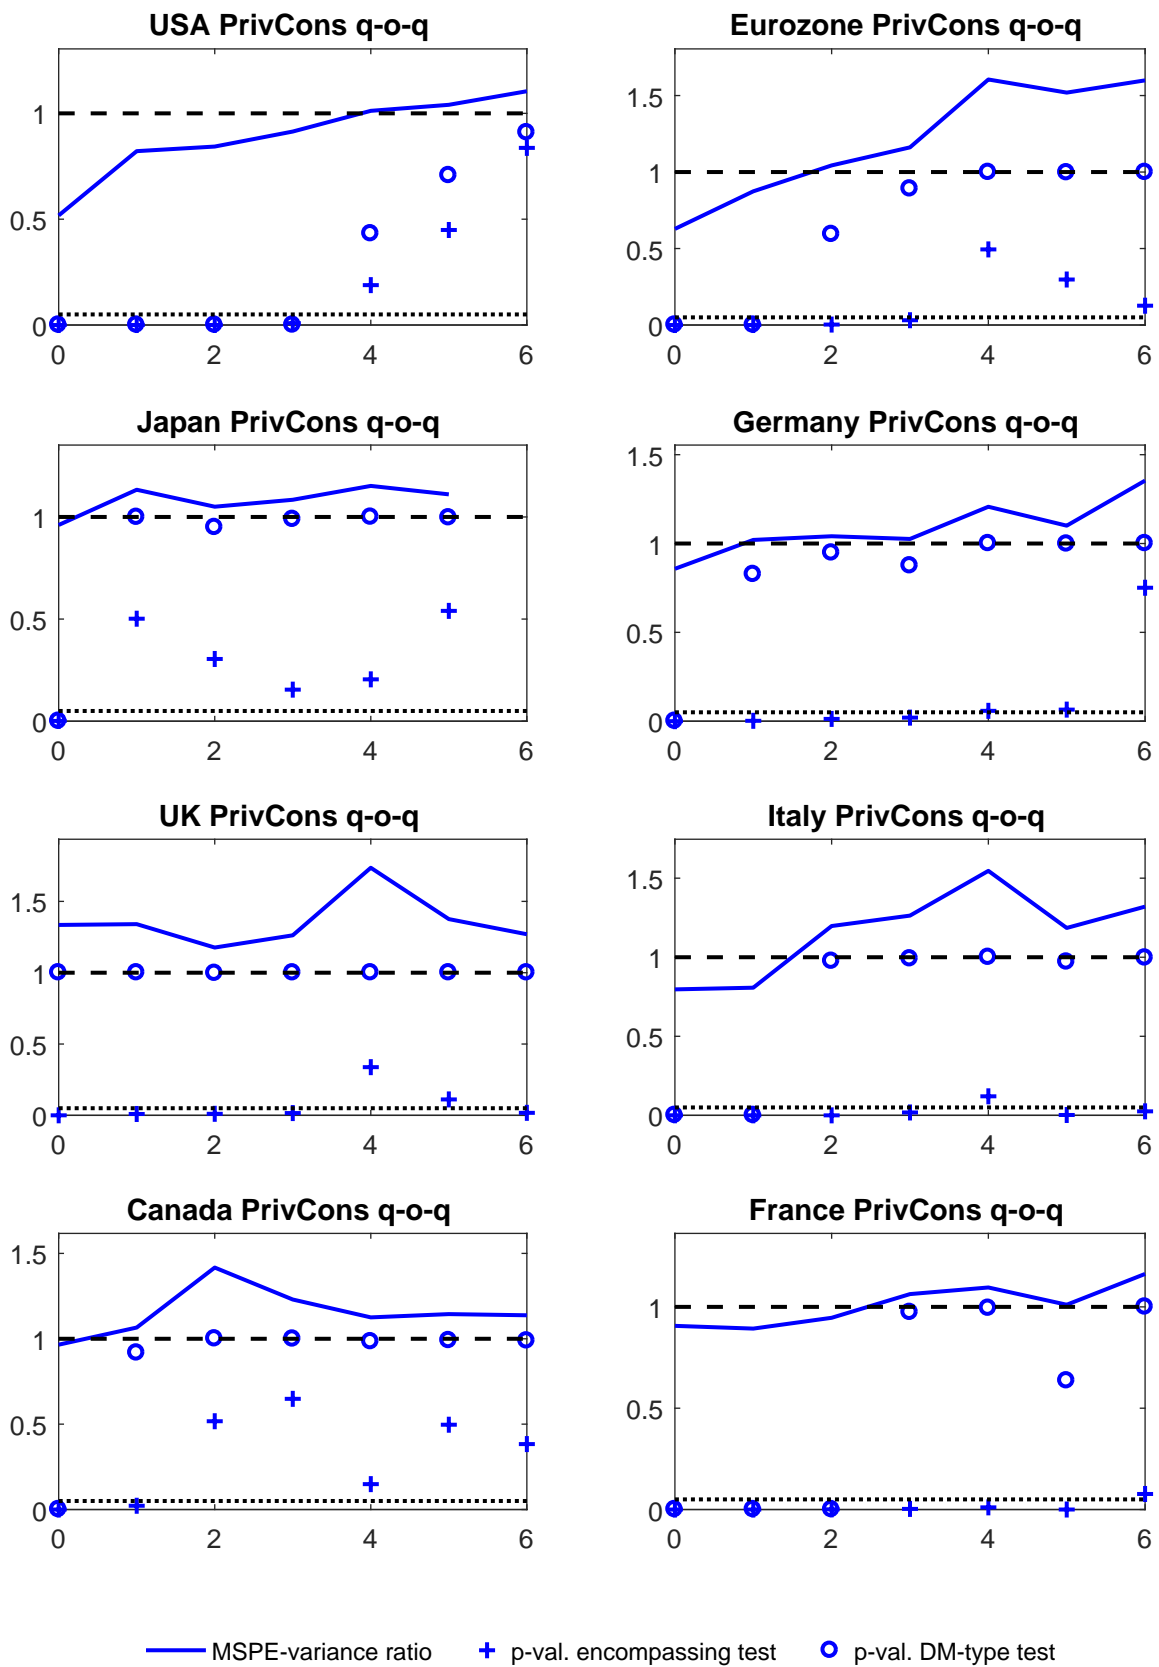


Figure 5: Test results for quarter-on-quarter growth rates of real private consumption. For further explanations, see Figure 3.

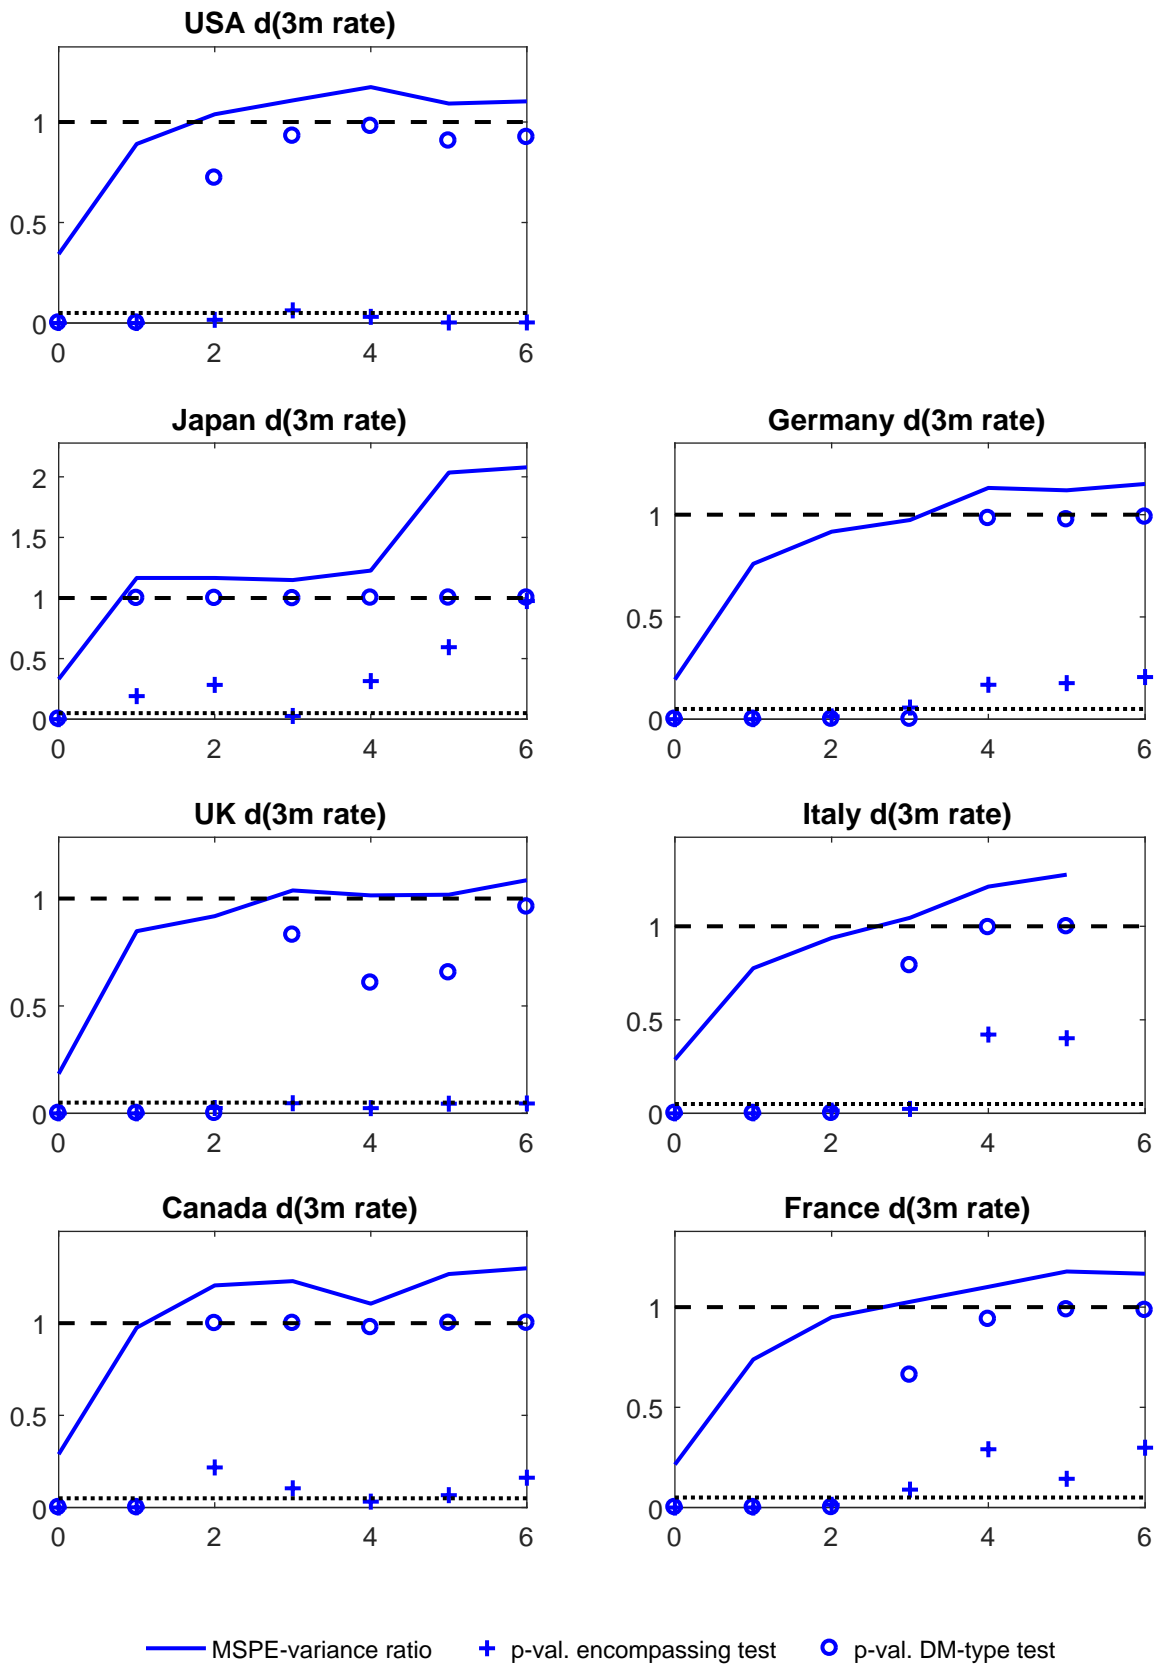


Figure 6: Test results for the change in the end-quarter 3-month interest rate. For further explanations, see Figure 3.

Thus, as to be expected, strong changes in the VAT rate which are announced well in advance can render inflation forecasts informative even at larger horizons.

The results for the quarter-on-quarter growth rates of real private consumption growth vary strongly across tests and countries. The encompassing test often implies pronouncedly larger values of \hat{h}^* than the DM-type test. Moreover, \hat{h}^* according to the encompassing test is mostly larger than in the case of real GDP growth. These results might be due to the facts that announced changes in the VAT rate are relatively important for private consumption, and that consumption forecasts often tend to be biased. Indeed, using the full sample for Japan and the UK leads to $\hat{h}^* \geq 6$ for both countries according to both tests. With the restricted sample, for the UK, the DM-type test indicates that not even the nowcast is informative.

The maximum forecast horizon for the change in the 3-month interest rate mostly varies between 1 quarter and 3 quarters. For the UK, the encompassing test finds informative forecasts at least up to $h = 6$. Except for the latter case, the two \hat{h}^* found by the tests for a given country do not differ by more than 1 quarter. For Japan, only the nowcasts turn out to be informative.

9 Concluding remarks

This paper develops a forecast evaluation framework for testing the null hypothesis that the forecast at some pre-specified horizon h is uninformative. The proposed tests can be applied sequentially to identify the maximum forecast horizon of predictive models. We show that due to the nested nature of the forecast comparison the standard Diebold-Mariano (DM) type test statistic has a nonstandard limiting distribution and suffers from a severe loss of power. To overcome this problem we adjust the test statistic and derive alternative tests from the encompassing principle that result in a Mincer-Zarnovitz regression. Our analysis of the local power reveals that the DM type test statistics is more powerful in the vicinity of the unit root hypothesis, whereas it performs similar to the encompassing test if the forecasts are more informative. Our Monte Carlo simulations suggest that the DM type test suffers from severe size distortions in reasonable sample sizes, whereas the regression variant of the encompassing test exhibits reliable sizes.

In our empirical analysis, we apply our tests to macroeconomic forecasts from the survey of Consensus Economics. Our results suggest that forecasts of macroeconomic key variables are hardly informative beyond 2–4 quarters ahead. Our results confirm earlier (anecdotal) findings from macroeconomic forecasting. The main contribution of our work is to provide statistical tests that allows the forecaster to assess the maximum forecast horizon of the forecast model of interest.

It is worth mentioning that our testing approach (as any other empirical methodology)

has two major limitations. First, the estimated maximum forecast horizon may be biased downwards if the predictive power is weak but not negligible. A similar caveat applies if the number of forecasts in the evaluation sample is small. Second, The estimated maximal forecast horizon depends on the specific model that generates the forecasts. If the model fails to exploit important information it may produce uninformative forecasts, while a richer forecasting model may result in informative forecasts. Accordingly, any qualification of the informative content is conditional on the forecasting model.

References

- ANDREWS, D. W. K. (1991): “Heteroskedasticity and Autocorrelation Consistent Covariance Matrix Estimation,” *Econometrica*, 59(3), 817–858.
- ANG, A., G. BEKAERT, AND M. WEI (2007): “Do macro variables, asset markets, or surveys forecast inflation better?,” *Journal of Monetary Economics*, 54(4), 1163–1212.
- CHONG, Y. Y., AND D. F. HENDRY (1986): “Econometric Evaluation of Linear Macroeconomic Models,” *Review of Economic Studies*, 53(4), 671–690.
- CLARK, T. E., AND M. W. MCCrackEN (2001): “Tests of equal forecast accuracy and encompassing for nested models,” *Journal of Econometrics*, 105(1), 85–110.
- CLEMENTS, M., AND D. HENDRY (1998): *Forecasting Economic Time Series*, no. 9780521634809 in Cambridge Books. Cambridge University Press.
- CLEMENTS, M. P., AND D. F. HENDRY (1993): “On the limitations of comparing mean square forecast errors,” *Journal of Forecasting*, 12(8), 617–637.
- DIEBOLD, F. X., AND L. KILIAN (2001): “Measuring predictability: theory and macroeconomic applications,” *Journal of Applied Econometrics*, 16(6), 657–669.
- DIEBOLD, F. X., AND R. S. MARIANO (1995): “Comparing Predictive Accuracy,” *Journal of Business & Economic Statistics*, 13(3), 253–63.
- GALBRAITH, J. W., AND G. TKACZ (2007): “Forecast content and content horizons for some important macroeconomic time series,” *Canadian Journal of Economics*, 40(3), 935–953.
- GRANGER, C. W. J., AND P. NEWBOLD (1986): *Forecasting Economic Time Series*, no. 9780122951831 in Elsevier Monographs. Elsevier.
- ISIKLAR, G., AND K. LAHIRI (2007): “How far ahead can we forecast? Evidence from cross-country surveys,” *International Journal of Forecasting*, 23(2), 167–187.
- KNÜPPEL, M. (forthcoming): “Forecast-Error-Based Estimation of Forecast Uncertainty When the Horizon Is Increased,” *International Journal of Forecasting*.
- MINCER, J. A., AND V. ZARNOWITZ (1969): “The Evaluation of Economic Forecasts,” in *Economic Forecasts and Expectations: Analysis of Forecasting Behavior and Performance*, ed. by J. A. Mincer, chap. 1, pp. 1–46. NBER.
- NELSON, C. R. (1976): “The Interpretation of R^2 in Autoregressive-Moving Average Time Series Models,” *The American Statistician*, 30(4), 175–180.

- NEWKEY, W. K., AND K. D. WEST (1987): "A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix," *Econometrica*, 55(3), 703–708.
- PARZEN, E. (1981): "Time Series Model Identification and Prediction Variance Horizon," in *Applied Time Series Analysis II*, ed. by D. F. Findley, pp. 415 – 447. Academic Press.
- WEST, K. D. (1996): "Asymptotic Inference about Predictive Ability," *Econometrica*, 64(5), 1067–1084.

Appendix A: Proofs

Proof of Theorem 1:

Let $h > h^*$. Applying a mean-value expansion of the form $Y_{t+h|t}^{\widehat{\theta}_t} = Y_{t+h|t}^\theta + D_{t+h}(\bar{\theta}_t)(\widehat{\theta}_t - \theta)$, where $\bar{\theta}_t$ denotes a value between $\widehat{\theta}_t$ and θ yields

$$\begin{aligned}\delta_t^h &= [u_{t+h} + D_{t+h}(\bar{\theta}_t)(\widehat{\theta}_t - \theta)]^2 - (u_{t+h} - \bar{u}_h)^2 \\ \delta_t^h &= u_{t+h}^2 - (u_{t+h} - \bar{u}_h)^2 + 2u_{t+h}D_{t+h}(\bar{\theta}_t)(\widehat{\theta}_t - \theta) + D_{t+h}(\bar{\theta}_t)^2(\widehat{\theta}_t - \theta)^2 \\ &= \bar{u}_h(2u_{t+h} - \bar{u}_h) + 2u_{t+h}D_{t+h}(\bar{\theta}_t)(\widehat{\theta}_t - \theta) + D_{t+h}(\bar{\theta}_t)^2(\widehat{\theta}_t - \theta)^2\end{aligned}$$

where $\bar{u}_h = n^{-1} \sum_{t=h+1}^{n+h} u_t$ and

$$\begin{aligned}\widehat{\theta}_t - \theta &= (\widehat{\theta}_t - \widehat{\theta}_0) + (\widehat{\theta}_0 - \theta) \\ &= O_p\left(\sqrt{\frac{t}{T}} \cdot T^{-1/2}\right) + O_p(T^{-1/2}) = O_p(T^{-1/2})\end{aligned}$$

due to Assumption 2 (iii). Furthermore, (ii) and (iv) imply

$$\begin{aligned}\frac{1}{n} \sum_{t=1}^n \delta_t^h &= \bar{u}_h^2 + 2\frac{1}{n} \sum_{t=1}^n u_{t+h}D_{t+h}(\bar{\theta}_t)(\widehat{\theta}_t - \theta) + \frac{1}{n} \sum_{t=1}^n D_{t+h}(\bar{\theta}_t)^2(\widehat{\theta}_t - \theta)^2 \\ &= \bar{u}_h^2 + O_p(T^{-1})\end{aligned}$$

and

$$\begin{aligned}\widehat{\gamma}_\delta(j) &= \frac{1}{n} \sum_{t=j+1}^n \delta_t^h \delta_{t-j}^h \\ &= \frac{1}{n} \bar{u}_h^2 \sum_{t=j+1}^n (2u_{t+h} - \bar{u}_h)(2u_{t+h-j} - \bar{u}_h) + O_p(T^{-1}) \\ &= \frac{1}{n} \bar{u}_h^2 \left[\left(\sum_{t=j+1}^n 4u_{t+h}u_{t+h-j} \right) - 3n\bar{u}_h^2 \right] + O_p(T^{-1}) \\ &= \bar{u}_h^2 \left(\frac{1}{n} \sum_{t=j+1}^n 4u_{t+h}u_{t+h-j} \right) + O_p(n^{-2}) + O_p(T^{-1}).\end{aligned}$$

Hence

$$\begin{aligned}\widehat{\omega}_\delta^2 &= \widehat{\gamma}_\delta(0) + 2 \sum_{j=1}^{h-1} j \widehat{\gamma}_\delta(j) \\ &= 4\bar{u}_h^2 \left(\widehat{\gamma}_u(0) + 2 \sum_{j=1}^{h-1} j \widehat{\gamma}_u(j) \right) + O_p(T^{-1}) + O_p(n^{-2}) \\ &= 4\bar{u}_h^2 \widehat{\omega}_u^2 + O_p(T^{-1}) + O_p(n^{-2}).\end{aligned}$$

Thus,

$$\begin{aligned}
d_h &= \frac{1}{\widehat{\omega}_\delta \sqrt{n}} \sum_{t=1}^h \delta_t^h \\
&= \frac{(\sqrt{n} \bar{u}_h)^2 + O_p(n/T)}{\sqrt{4n \bar{u}_h^2 \widehat{\omega}_u^2 + O_p(n/T) + O_p(n^{-1})}} \\
&= \frac{\sqrt{n} |\bar{u}_h|}{2 \widehat{\omega}_u} + O_p\left(\frac{n}{T}\right) \xrightarrow{d} \frac{|z|}{2}
\end{aligned}$$

where z is a standard normally distributed random variable.

Proof of Corollary 1:

The distribution of $2d_h$ follows directly from Theorem 1. As shown in Theorem 1 we have

$$\sum_{t=1}^n \delta_t^h = (\sqrt{n} \bar{u})^2 + O_p(n/T).$$

If $n/T \rightarrow 0$ and $\widehat{\omega}_u^2$ is a consistent estimator of the long-run variance of $u_{t+h} = Y_{t+h} - \mu$ then \widetilde{d}_h possesses a χ^2 limiting distribution with one degree of freedom.

Proof of Theorem 2

Consider some $h > h^*$. We first analyze

$$\sum_{t=1}^n (\widehat{Y}_{t+h|t} - \bar{Y}_h)(Y_{t+h} - \bar{Y}_h) = \sum_{t=1}^n \widehat{Y}_{t+h|t}(u_{t+h} - \bar{u}_h).$$

An important problem with analysing this expression is that the estimation error in $\widehat{Y}_{t+h|t} = Y_{t+h|t}^{\widehat{\theta}_t}$ is correlated with \bar{u}_h . To sidestep this difficulty we decompose the forecast into one component $Y_{t+h|t}^{\widehat{\theta}_0}$ that is independent of $\{u_{1+h}, \dots, u_{n+h}\}$ and show that the remaining component is asymptotically negligible. Applying a mean value expansion yields

$$\widehat{Y}_{t+h|t} = Y_{t+h|t}^{\widehat{\theta}_t} = Y_{t+h|t}^{\widehat{\theta}_0} + D_{t+h}(\bar{\theta}_t)(\widehat{\theta}_t - \widehat{\theta}_0)$$

where $D_{t+h}(\theta) = \partial Y_{t+h|t}^\theta / \partial \theta$ and $\bar{\theta}_t$ denotes some value between θ_0 and $\widehat{\theta}_t$. Note that by Assumption 2 $Y_{t+h|t}^{\widehat{\theta}_0}$ is uncorrelated with all $u_{h+1}, u_{h+2}, \dots, u_{n+h}$. Accordingly,

$$\sum_{t=1}^n \left[Y_{t+h|t}^{\widehat{\theta}_0} + D_{t+h}(\bar{\theta}_t)(\widehat{\theta}_t - \widehat{\theta}_0) \right] (u_{t+h} - \bar{u}_h) = A_{T,n} + B_{T,n}^1 + B_{T,n}^2$$

where

$$\begin{aligned} A_{T,n} &= \sum_{t=1}^n Y_{t+h|t}^{\hat{\theta}_0} (u_{t+h} - \bar{u}_h) \\ B_{T,n}^1 &= \sum_{t=1}^n D_{t+h}(\bar{\theta}_t) (\hat{\theta}_t - \hat{\theta}_0) u_{t+h} \\ B_{T,n}^2 &= \bar{u}_h \sum_{t=1}^n D_{t+h}(\bar{\theta}_t) (\hat{\theta}_t - \hat{\theta}_0). \end{aligned}$$

Another mean value expansion around the true value θ with $Y_{t+h|t}^\theta = \mu$ yields

$$\begin{aligned} A_{T,n} &= (\hat{\theta}_0 - \theta) \sum_{t=1}^n D_{t+h}(\bar{\theta}_0) u_{t+h} - (\hat{\theta}_0 - \theta) \bar{u}_h \sum_{t=1}^n D_{t+h}(\bar{\theta}_0) \\ &= A_{T,n}^1 + A_{T,n}^2 \end{aligned}$$

where $\bar{\theta}_0$ is some value between $\hat{\theta}_0$ and θ . Since $\hat{\theta}_0$ and $D_{t+h}(\bar{\theta}_0)$ are uncorrelated with u_{t+h} it follows that $A_{T,n}^1 = O_p(T^{-1/2})O_p(n^{1/2})$, whereas $A_{T,n}^2 = O_p(T^{-1/2})O_p(n^{-1/2})O_p(n)$. Thus, $A_{T,n}$ is $O_p(\sqrt{n/T})$. By Assumption 2 (iii) we have

$$\xi_{T,n,t} \equiv \hat{\theta}_t - \hat{\theta}_0 = a_{T,n,t} \frac{\sqrt{t}}{T},$$

where $a_{T,n,t}$ is $O_p(1)$. Furthermore under the null hypothesis $\xi_{T,n,t}$ and $D_{t+h}(\bar{\theta}_t)$ are uncorrelated with u_{t+h} . It follows that

$$\begin{aligned} \sum_{t=1}^n \xi_{T,n,t}^2 D_{t+h}(\bar{\theta}_t)^2 u_{t+h}^2 &= \sum_{t=1}^n \frac{a_{T,n,t}^2 t}{T^2} D_{t+h}(\bar{\theta}_t)^2 u_{t+h}^2 \\ &= \sum_{t=1}^n b_{T,n,t}^2 \frac{t}{T^2} = O_p\left(\frac{n^2}{T^2}\right), \end{aligned}$$

where $b_{T,n,t} = a_{T,n,t}^2 D_{t+h}(\bar{\theta}_t)^2 u_{t+h}^2 = O_p(1)$. Thus,

$$B_{T,n}^1 = \sum_{t=1}^n \xi_{T,n,t} D_{t+h}(\bar{\theta}_t) u_{t+h} = O_p\left(\frac{n}{T}\right).$$

Since

$$\begin{aligned} \sum_{t=1}^n \xi_{T,n,t} D_{t+h}(\bar{\theta}_t) &= \sum_{t=1}^n \frac{a_{T,n,t}^2 \sqrt{t}}{T} D_{t+h}(\bar{\theta}_t) \\ &= O_p\left(\frac{n^{3/2}}{T}\right) \end{aligned}$$

and, therefore, $B_{T,n}^2 = O_p(n^{-1/2})O_p(n^{3/2}/T) = O_p(n/T)$. As $n/T \rightarrow 0$ It follows that

$$\begin{aligned} \sqrt{\frac{T}{n}} \sum_{t=1}^n (\widehat{Y}_{t+h|t} - \overline{Y}_h)(Y_{t+h} - \overline{Y}_h) &= \sqrt{\frac{T}{n}} (A_{T,n} + B_{T,n}^1 + B_{T,n}^2) \\ &= \sqrt{\frac{T}{n}} A_{T,n} + O_p\left(\sqrt{\frac{n}{T}}\right) \\ &= \sqrt{T}(\widehat{\theta}_0 - \theta) \frac{1}{\sqrt{n}} \sum_{t=1}^n D_{t+h}(\overline{\theta}_0)(u_{t+h} - \overline{u}_h) + O_p\left(\sqrt{\frac{n}{T}}\right) \end{aligned}$$

Next we analyze

$$\sum_{t=1}^n (\widehat{Y}_{t+h|t} - \overline{Y}_h)^2 (Y_{t+h} - \overline{Y}_h)^2 = \sum_{t=1}^n (\widehat{Y}_{t+h|t} - \overline{Y}_h)^2 (u_{t+h} - \overline{u}_h)^2$$

Using the above mean value expansions we obtain

$$\begin{aligned} \widehat{Y}_{t+h|t} &= Y_{t+h|t}^{\widehat{\theta}_0} + D_{t+h}(\overline{\theta}_t)(\widehat{\theta}_t - \widehat{\theta}_0) \\ &= \mu + D_{t+h}(\overline{\theta}_0)(\widehat{\theta}_0 - \theta) + D_{t+h}(\overline{\theta}_t)(\widehat{\theta}_t - \widehat{\theta}_0) \\ \widehat{Y}_{t+h|t} - \overline{Y}_h &= \widetilde{D}_{t+h}(\overline{\theta}_0)(\widehat{\theta}_0 - \theta) + \widetilde{\Psi}_{t+h}(\widehat{\theta}_t, \widehat{\theta}_0) \end{aligned}$$

where

$$\begin{aligned} \widetilde{D}_{t+h}(\overline{\theta}_i) &= D_{t+h}(\overline{\theta}_i) - n^{-1} \sum_{t=1}^n D_{t+h}(\overline{\theta}_i) \\ \widetilde{\Psi}_{t+h}(\widehat{\theta}_t, \widehat{\theta}_0) &= D_{t+h}(\overline{\theta}_t)(\widehat{\theta}_t - \widehat{\theta}_0) - \frac{1}{n} \sum_{t=1}^n D_{t+h}(\overline{\theta}_t)(\widehat{\theta}_t - \widehat{\theta}_0) \end{aligned}$$

for $i = 0, 1, \dots$. It follows that

$$\begin{aligned} \sum_{t=1}^n (\widehat{Y}_{t+h|t} - \overline{Y}_h)^2 (Y_{t+h} - \overline{Y}_h)^2 &= (\widehat{\theta}_0 - \theta)^2 \sum_{t=1}^n \widetilde{D}_{t+h}(\overline{\theta}_0)^2 (u_{t+h} - \overline{u}_h)^2 \\ &\quad + \sum_{t=1}^n \widetilde{\Psi}_{t+h}(\widehat{\theta}_t, \widehat{\theta}_0)^2 (u_{t+h} - \overline{u}_h)^2 \\ &\quad + 2(\widehat{\theta}_0 - \theta) \sum_{t=1}^n \widetilde{D}_{t+h}(\overline{\theta}_0) \widetilde{\Psi}_{t+h}(\widehat{\theta}_t, \widehat{\theta}_0) (u_{t+h} - \overline{u}_h)^2 \\ &= C_{T,n}^0 + C_{T,n}^1 + C_{T,n}^2. \end{aligned}$$

For the leading term we obtain

$$C_{T,n}^0 = O_p(T^{-1})O_p(n) = O_p(n/T)$$

For the second term we note that

$$\sum_{t=1}^n a_{T,n,t}^2 \frac{t}{T^2} D_{t+h}(\bar{\theta}_t)^2 (u_{t+h} - \bar{u}_h)^2 = O_p(n^2/T^2).$$

Since the mean adjustment does not affect the order of magnitude we conclude that

$$C_{T,n}^1 = O_p(n^2/T^2)$$

For the last term we obtain

$$\sum_{t=1}^n a_{T,n,t} \frac{\sqrt{t}}{T} D_{t+h}(\bar{\theta}_0) D_{t+h}(\bar{\theta}_t) = O_p(n^{3/2}/T)$$

and, since the mean-adjustment does not affect the order of magnitude,

$$C_{T,n}^2 = O_p(n^{3/2}/T^{3/2}).$$

Combining these results yields

$$\frac{T}{n} \sum_{t=1}^n (\hat{Y}_{t+h|t} - \bar{Y}_h)^2 (Y_{t+h} - \bar{Y}_h)^2 = T(\hat{\theta}_0 - \theta)^2 \frac{1}{n} \sum_{t=1}^n \tilde{D}_{t+h}(\bar{\theta}_0)^2 (u_t - \bar{u}_h)^2 + O_p\left(\sqrt{\frac{n}{T}}\right)$$

In the same manner it can be shown that for $j = 1, 2, \dots, h$

$$\begin{aligned} & \frac{T}{n} \sum_{t=1+j}^n (\hat{Y}_{t+h|t} - \bar{Y}_h)(\hat{Y}_{t+h-j|t} - \bar{Y}_h)(Y_{t+h} - \bar{Y}_h)(Y_{t+h-j} - \bar{Y}_h) \\ &= T(\hat{\theta}_0 - \theta)^2 \frac{1}{n} \sum_{t=1}^n \tilde{D}_{t+h}(\bar{\theta}_0) \tilde{D}_{t+h-j}(\bar{\theta}_0) (u_{t+h} - \bar{u}_h)(u_{t+h-j} - \bar{u}_h) + O_p\left(\sqrt{\frac{n}{T}}\right). \end{aligned}$$

Define $\hat{V}_{n,T} = \hat{\Gamma}_{0,n,T} + 2 \sum_{j=1}^{h-1} \hat{\Gamma}_{j,n,T}$ where

$$\hat{\Gamma}_{j,n,T} = \frac{1}{n} \sum_{t=1+j}^n (\hat{Y}_{t+h|t} - \bar{Y}_h)(\hat{Y}_{t+h-j|t} - \bar{Y}_h)(Y_{t+h} - \bar{Y}_h)(Y_{t+h-j} - \bar{Y}_h).$$

It follows that

$$\frac{T}{n} \hat{V}_{n,T} = T(\hat{\theta}_0 - \theta)^2 \mathbb{E} \left[\frac{1}{n} \left(\sum_{t=1}^n \tilde{D}_{t+h}(\theta) (u_{t+h} - \bar{u}_h) \right)^2 \right] + o_p(1).$$

Applying a suitable version of the central limit theorem it follows that

$$\frac{1}{\sqrt{nV_{n,T}}} \sum_{t=1}^n (\hat{Y}_{t+h|t} - \bar{Y}_h)(u_{t+h} - \bar{u}_n) = \frac{\frac{1}{\sqrt{n}} \sum_{t=1}^n \tilde{D}_{t+h}(\theta) (u_{t+h} - \bar{u}_h)}{\sqrt{\mathbb{E} \left[\frac{1}{n} \left(\sum_{t=1}^n \tilde{D}_{t+h}(\theta) (u_{t+h} - \bar{u}_h) \right)^2 \right]}} + o_p(1)$$

$$\xrightarrow{d} \mathcal{N}(0, 1)$$

■

Proof of Theorem 4

Under the local alternative, we have

$$Y_{t+1} - \bar{Y}_1 = u_{t+1} - \bar{u}_1 + (c/\sqrt{n})(X_t - \bar{X})$$

and the model prediction error is given by $\hat{e}_{t+h|t} = u_{t+1} + O_p(T^{-1/2})$. Following the proof of Theorem 1 we obtain for $n/T \rightarrow 0$

$$\begin{aligned} \sum_{t=1}^n \delta_t^1 &= (\sqrt{n} \bar{u}_1)^2 - \frac{2c}{\sqrt{n}} \sum_{t=1}^n (X_t - \bar{X}) u_{t+1} - \frac{c^2}{n} \sum_{t=1}^n (X_t - \bar{X})^2 + O_p(n/T) \\ &= (\sqrt{n} \bar{u}_1)^2 - 2c\sigma_u\sigma_x R_n - c^2\sigma_x^2 + o_p(1) \\ &\xrightarrow{d} \sigma_u^2 z_1^2 - 2\sigma_u\sigma_x c z_2 - c^2\sigma_x^2, \end{aligned}$$

where

$$\begin{aligned} \frac{\sqrt{n} \bar{u}}{\sigma_u} &\xrightarrow{d} z_1 \stackrel{d}{=} \mathcal{N}(0, 1) \\ R_n &= \frac{1}{\sigma_u\sigma_x\sqrt{n}} \sum_{t=1}^n (X_t - \bar{X}) u_{t+1} \xrightarrow{d} z_2 \stackrel{d}{=} \mathcal{N}(0, 1) \end{aligned}$$

Accordingly, for the modified DM statistic we obtain

$$\tilde{d}_1 = \frac{1}{\hat{\sigma}_u^2} \sum_{t=1}^n \delta_t^1 \xrightarrow{d} z_1^2 - 2c \frac{\sigma_x}{\sigma_u} z_2 - c^2 \frac{\sigma_x^2}{\sigma_u^2}$$

where $\hat{\sigma}_u^2 = n^{-1} \sum_{t=1}^n (u_{t+1} - \bar{u}_1 + (c/\sqrt{n})X_t)^2 = \sigma_u^2 + O_p(n^{-1/2})$.

Using

$$\hat{Y}_{t+1|t} = \frac{c}{\sqrt{n}} X_t + O_p(T^{-1/2})$$

we have for $n/\sqrt{T} \rightarrow 0$

$$\begin{aligned} \sum_{t=1}^n \xi_t^1 &= \sum_{t=1}^n \left[u_{t+1} - \bar{u}_1 + \frac{c}{\sqrt{n}}(X_t - \bar{X}) \right] \hat{Y}_{t+1|t} \\ &= \frac{c}{\sqrt{n}} \sum_{t=1}^n u_{t+1}(X_t - \bar{X}) + \frac{c^2}{n} \sum_{t=1}^n (X_t - \bar{X})^2 + O_p(n/\sqrt{T}) \\ &\xrightarrow{d} c\sigma_u\sigma_x z_2 + c^2\sigma_x^2. \end{aligned}$$

Furthermore

$$\begin{aligned} n\hat{\omega}_\xi^2 &= \sum_{t=1}^n (\xi_t^1)^2 = \frac{c^2}{n} \sum_{t=1}^n (u_t - \bar{u}_1)^2 (X_t - \bar{X})^2 + o_p(1) \\ &\xrightarrow{p} c^2\sigma_u^2\sigma_x^2 \end{aligned}$$

and, thus,

$$\varrho_1 = \frac{\frac{1}{\sqrt{n}} \sum_{t=1}^n \xi_t^1}{\sqrt{\widehat{\omega}_\xi^2}} \xrightarrow{d} \text{sign}(c)z_2 + |c| \frac{\sigma_x}{\sigma_u}$$

with $\text{sign}(a) = 1$ if $a \geq 0$ and $\text{sign}(a) = -1$ for $a < 0$.

Appendix B - Data Descriptions

Concerning the quarter-on-quarter growth rates of real GDP, the only change that occurred in the sample is from West German GDP to the GDP of the reunified Germany in the fourth quarter of 1995. Forecasts for the euro area started being collected in the last quarter of 2002 for all variables except the 3-month interest rate.

The inflation measure used is the year-on-year growth rate of the index of consumer prices (CPI) for all countries except for the UK, where the retail price index (RPI) is used, because the sample of forecasts for the CPI does not start until 2004. Inflation forecasts for the reunified Germany started in the fourth quarter of 1996.

Private consumption is measured by the personal consumption expenditures in the US and Canada, by private consumption in Japan, Germany, and the Euro area, and by household consumption in France, the UK, and Italy. Private consumption forecasts for the reunified Germany started in the fourth quarter of 1995.

The 3-month interest rate is measured at the last day of the quarter. The interest rate used in the analysis is the 3-month treasury bill rate for the US and Canada, the 3-month Yen certificate of deposit rate for Japan, with the sample ending in the first quarter in 2010, the 3-month Euribor in Germany, Italy, and France, with the sample starting in the first quarter of 1999, and the 3-month interbank rate for the UK.

The changes in the VAT rates are listed in Table 7. Since Japan and the UK experienced two VAT rate changes of at least 2 percentage points, we adapt their samples of inflation and private consumption. For Japan, both samples start in the second quarter of 1997 and end in the fourth quarter of 2013, because in the first quarter of 2014, real private consumption already increased substantially due to the following VAT rate increase. For the UK, the samples continue to start in the first quarter of 1995, but end in the second quarter of 2008.

All resulting sample sizes can be found in Table 8.

Table 7: Changes in the value-added tax rates in percentage points

Country	Date	VAT rate in pp		
		from	to	change
Japan	Apr 97	3	5	2
Japan	Apr 14	5	8	3
Germany	Apr 98	15	16	1
Germany	Jan 07	16	19	3
France	Apr 00	20.6	19.6	-1
Italy	Oct 97	19	20	1
Italy	Sep 11	20	21	1
Italy	Oct 13	21	22	1
UK	Dec 08	17.5	15	-2.5
UK	Jan 10	15	17.5	2.5
UK	Jan 11	17.5	20	2.5
Canada	Jul 06	7	6	-1
Canada	Jan 08	6	5	-1

Table 8: Numbers of observations n

	US	EA	JP	DE	UK	IT	CA	FR
GDP q-o-q	79	48	79	76	79	79	79	79
CPI y-o-y	79	48	56	72	44	79	79	79
PrivCons q-o-q	79	48	56	76	44	79	79	79
d(3m rate)	79		53	63	79	63	79	63

Note: 'GDP q-o-q' denotes quarter-on-quarter growth rates of real GDP, 'CPI y-o-y' year-on-year growth rates of consumer prices, 'PrivCons q-o-q' quarter-on-quarter growth rates of real private consumption, and 'd(3m rate)' quarter-on-quarter changes of the end-quarter 3-month interest rate. The abbreviations used for the countries are 'US' for the United States, 'EA' for the euro area, 'JP' for Japan, 'DE' for Germany, 'UK' for the United Kingdom, 'IT' for Italy, 'CA' for Canada, and 'FR' for France.

Appendix C - Not for Publication

Table 9: Results for case ‘MA(1)-AR(1)’, rolling estimation

forecast horizon h	0	1	2	3	4	0	1	2	3	4
	$T = 100, n = 50$					$T = 100, n = 100$				
MSPE / Variance	0.89	1.06	1.06	1.06		0.87	1.04	1.04	1.04	
rejections										
DM-type tests										
\tilde{d}_h	0.87	0.11	0.10	0.10		0.98	0.08	0.07	0.07	
d_h	0.88	0.11	0.11	0.11		0.98	0.08	0.08	0.07	
encompassing tests										
$\beta_{1,h}$	0.90	0.06	0.05	0.05		0.99	0.04	0.04	0.03	
ϱ_h	0.79	0.03	0.02	0.02		0.98	0.02	0.02	0.02	
classical DM test	0.24	0.00	0.00	0.00		0.54	0.00	0.00	0.00	
\hat{h}^*										
DM-type tests										
\tilde{d}_h	0.13	0.78	0.05	0.02	0.01	0.02	0.91	0.05	0.02	0.01
d_h	0.12	0.78	0.06	0.03	0.01	0.02	0.90	0.05	0.02	0.01
encompassing tests										
$\beta_{1,h}$	0.10	0.86	0.03	0.01	0.00	0.01	0.95	0.03	0.01	0.00
ϱ_h	0.21	0.78	0.01	0.00	0.00	0.02	0.96	0.02	0.00	0.00
classical DM test	0.76	0.24	0.00	0.00	0.00	0.46	0.54	0.00	0.00	0.00

Note: For explanations, see Table 4. In contrast to Table 4, the estimation is carried out using a rolling window of length T .