

# Testing Forecast Optimality under Unknown Loss\*

Andrew J. Patton

*London School of Economics*

Allan Timmermann

*University of California, San Diego*

July 2006

Empirical tests of forecast optimality have traditionally been conducted under the assumption of mean squared error loss or some other known loss function. This paper establishes new testable properties that hold when the forecaster's loss function is unknown but testable restrictions can be imposed on the data generating process, trading off conditions on the data generating process against conditions on the loss function. We propose flexible estimation of the forecaster's loss function in situations where the loss depends not only on the forecast error but also on other state variables such as the level of the target variable. We apply our results to the problem of evaluating the Federal Reserve's forecasts of output growth. Forecast optimality is rejected if the Fed's loss only depends on the forecast error. However, the empirical findings are consistent with forecast optimality provided that over-predictions of output growth are costlier to the Fed than under-predictions, particularly during periods of low economic growth.

Keywords: Federal Reserve, output growth, forecast evaluation, loss function, forecast efficiency.

---

\*The authors would like to thank three anonymous referees for very helpful comments and suggestions. Jeremy Berkowitz, Sean Campbell, Carlos Capistran, Peter Christoffersen, Valentina Corradi, Frank Diebold, Graham Elliott, Raffaella Giacomini, Clive Granger, Peter Hansen, Mark Machina, Nour Meddahi, Jim Powell, Barbara Rossi, Paul Ruud, Kevin Sheppard, Hal White, Stanley Zin and seminar participants at Berkeley, Brown, CORE, Cornell, CREST, Erasmus University Rotterdam, LSE, McGill, Pennsylvania, Queen Mary, Stanford, UCLA, UC Riverside, UCSD, Warwick Business School and the 2002 EC<sup>2</sup> conference all provided useful comments. Patton: London School of Economics, Houghton Street, London WC2A 2AE, United Kingdom. Email: a.patton@lse.ac.uk. Timmermann: Department of Economics, University of California, San Diego, 9500 Gilman Drive, La Jolla CA 92093-0508, U.S.A. Email: atimmerm@ucsd.edu.

# 1 Introduction

Forecasts are used extensively to guide the decisions of individuals, businesses and macroeconomic policy-makers, so what constitutes a ‘good’ or optimal forecast is of great practical relevance. Indeed, knowledge of the properties possessed by an optimal forecast has been used, inter alia, in tests of the efficient market hypothesis in financial markets and in tests of the rationality of decision makers in a variety of economic applications. Almost without exception empirical work has relied on mean squared error (MSE) loss:

$$L(Y_{t+h}, \hat{Y}_{t+h,t}) = \left( Y_{t+h} - \hat{Y}_{t+h,t} \right)^2, \quad (1)$$

where  $L(\cdot)$  is the loss function,  $t$  is the time that the forecast is computed,  $h$  is the forecast horizon,  $Y_{t+h}$  is the predicted variable and  $\hat{Y}_{t+h,t}$  is the forecast based on information available at time  $t$ . Assuming covariance stationarity, under MSE loss the optimal forecast and the associated forecast error satisfy a set of standard conditions (see, e.g., Diebold and Lopez, 1996):

1. The optimal forecast of  $Y_{t+h}$  is its conditional expectation so the forecast error is conditionally (and unconditionally) unbiased;
2. The  $h$ -step forecast error exhibits zero serial covariance beyond lag  $(h - 1)$ ;
3. The unconditional variance of the forecast error is a non-decreasing function of the forecast horizon.

Unfortunately, such properties often do not provide a useful guide to empirical tests since they do not generally hold under other loss functions (Patton and Timmermann, 2004) and the forecaster’s loss function is unknown in most applications. This is particularly important since there are often good reasons to believe that loss depends asymmetrically on positive and negative forecast errors as discussed by, inter alia, Granger (1969), Varian (1974), Granger and Newbold (1986), Zellner (1986), Weiss (1996), West, Edison and Cho (1993) and Christoffersen and Diebold (1997).

As an important case in point, consider the one-quarter-ahead forecasts of real output (Gross Domestic Product or GDP) growth produced by the Board of Governors of the Federal Reserve – the so-called “Greenbook” forecasts – plotted in Figure 1 against realizations of annualized real GDP growth over the period 1968Q4 - 1999Q4, a total of 125 observations. The Federal Reserve

does not publish how the Greenbook forecasts are constructed (though these forecasts are known to be based on judgemental information as well as formal statistical models), nor does it publish the loss function used, either explicitly or implicitly, in the construction of the forecasts. However, it is known that considerable resources are devoted to the production of these forecasts and that they are used by the Fed in the setting of monetary policy. We should therefore reasonably expect that these forecasts contain valuable information about future GDP growth. In fact, due to the sensitivity of these forecasts—and the information they contain about monetary policy—the Fed only releases them with a five-year delay.

Figure 2 plots realized GDP growth against the forecast in a scatterplot. If the forecasts were optimal under MSE loss, it follows from property 1 above that they should fall around the 45-degree line. This property is commonly tested through a so-called Mincer-Zarnowitz (1969) regression of realized values on a constant and the forecast, the implication being that the associated coefficients should be zero and one, respectively. OLS parameter estimates (and standard errors) of this regression were:

$$\begin{aligned}
 Y_{t+1} &= \underset{(0.451)}{1.253} + \underset{(0.106)}{0.710} \hat{Y}_{t+1,t} + u_{t+1}. \\
 R^2 &= 0.218
 \end{aligned}$$

Under MSE loss, forecast optimality does not rule out heteroskedasticity in the residuals from this regression. The only implication of optimality is that the forecast errors should follow a martingale difference sequence with respect to the information set used by the forecaster. Autocorrelation in the residuals,  $u_{t+1}$ , is ruled out, however, since lagged forecast errors are generally assumed to be part of the forecaster's information set. This implication, and other implications, of forecast optimality is not tested in the Mincer-Zarnowitz regression which only uses information on the mean and correlation between the forecast,  $\hat{Y}_{t+1,t}$  and the realization,  $Y_{t+1}$ , to test for unbiasedness. As such, this regression should only be viewed as a weak (necessary but not sufficient) test of one implication of forecast optimality under MSE loss.

Even so, a Wald test of the joint restriction on the intercept and the slope yields a  $\chi^2_2$ -statistic of 8.57 and a  $p$ -value of 0.01. All standard errors in this paper use the heteroskedasticity and autocorrelation robust estimator of Newey and West (1987). The test statistic and  $p$ -value obtained when the usual OLS parameter covariance matrix is used are 8.90 and 0.01. In both cases we reject the null of forecast optimality under MSE loss at the 0.05 level.

This strong rejection of optimality is at first sight surprising given the incentives of the Fed to forecast this variable well and the considerable resources devoted to its prediction. It should be recalled, however, that the rejection was predicated on MSE loss whereas, in actuality, the Federal Reserve's loss function is unknown. In fact, as pointed out by Nobay and Peel (2003) and Capistran (2006), there are strong reasons to assume that the Federal Reserve has asymmetric loss over forecast errors rather than MSE loss.

These empirical findings therefore point towards the need for establishing tests of forecast optimality that apply when the loss function is unknown. To this end, our paper establishes new results that trade off restrictions on the forecaster's loss function against restrictions on the data generating process (DGP). For example, in situations where the conditional higher order moments of the predicted variable are constant, we show that although the optimal forecast may well be biased, optimal  $h$ -step-ahead forecast errors display serial dependence of at most order  $h - 1$ , implying that optimal one-step-ahead forecasts generate serially uncorrelated forecast errors. These results hold irrespective of the shape of the loss function and offer new ways to test forecast optimality that do not require knowledge of the specific loss function, but do require restrictions on the underlying DGP. They will be useful in the common situation where the shape of the loss function is unknown, but the restrictions on the DGP can be tested empirically. We present similar results that hold when the DGP exhibits heteroskedasticity of a general unknown form and mild restrictions are imposed on the shape of the loss function, using a new family of quantile-based tests. We are not aware of any existing tests of forecast optimality that are robust to the loss function of the forecaster.

An important restriction required for some of these new results is that the loss function only depends on the forecast error. However it is quite likely that the loss also depends on factors such as the level of the predicted variable or the value of some other state variable, c.f. Garratt et al. (2003). For example, it may be particularly costly to over-predict GDP growth when growth is already very low. In effect this would amount to the Federal Reserve signalling a false recovery which could lead to an overly tight monetary policy at exactly the wrong point in time. To deal with broader classes of loss functions, we present a new method to test forecast optimality based on a flexible model of the loss function, which is then tested via a set of over-identifying restrictions. Indeed, the empirical analysis suggests that this broader class of loss functions is required to maintain the assumption of rationality for the Federal Reserve's forecasts.

The outline of the paper is as follows. Section 2 derives testable properties of optimal forecasts when the loss function is unknown but testable restrictions can be imposed on the DGP. Section 3 extends our results to the case where the loss function does not exclusively depend on the forecast error, Section 4 reports size and power results from a Monte Carlo simulation experiment and Section 5 concludes. An appendix contains technical details and proofs.

## 2 Testable Implications under Unknown Loss

Suppose that a decision maker is interested in forecasting some univariate time series,  $Y \equiv \{Y_t; t = 1, 2, \dots\}$ ,  $h$  steps ahead given information at time  $t$ ,  $\mathcal{F}_t$ . Let  $\tilde{Z}_t$  be a  $(m \times 1)$  vector of predictor variables used by the forecaster, e.g. the Federal Reserve, and let  $\mathcal{F}_t$  be the  $\sigma$ -field generated by  $\left\{ \left( Y_{t-k}, \tilde{Z}'_{t-k} \right)'; k \geq 0 \right\}$ . We will denote a generic element of  $\mathcal{F}_t$  as  $Z_t$ , and denote the conditional distribution of  $Y_{t+h}$  given  $\mathcal{F}_t$  as  $F_{t+h,t}$ , and the conditional density as  $f_{t+h,t}$ . Point forecasts conditional on  $\mathcal{F}_t$  are denoted by  $\hat{Y}_{t+h,t}$  and belong to  $\mathcal{Y}$ , a subset of  $\mathbb{R}$ , while forecast errors are given by  $e_{t+h,t} = Y_{t+h} - \hat{Y}_{t+h,t}$ . In general the objective of the forecast is to minimize the expected value of some loss function,  $L(Y_{t+h}, \hat{Y}_{t+h,t})$ , which is a mapping from realizations and forecasts to the real line,  $L: \mathbb{R} \times \mathcal{Y} \rightarrow \mathbb{R}$ . An optimal forecast is thus defined as:

$$\hat{Y}_{t+h,t}^* \equiv \arg \min_{\hat{y} \in \mathcal{Y}} E [L(Y_{t+h}, \hat{y}) | \mathcal{F}_t]. \quad (2)$$

We will use  $E_t[\cdot]$  as shorthand notation for  $E[\cdot | \mathcal{F}_t]$ , the conditional expectation given  $\mathcal{F}_t$ . We also define the conditional variance,  $V_t[Y_{t+h}] = E[(Y_{t+h} - E[Y_{t+h} | \mathcal{F}_t])^2 | \mathcal{F}_t]$  and the unconditional equivalents,  $E[\cdot]$  and  $V[\cdot]$ . Finally, we define  $\mu_{t+h,t} \equiv E_t[Y_{t+h}]$  and  $\sigma_{t+h,t}^2 \equiv V_t[Y_{t+h}]$ .

To establish properties of optimal forecasts that may be tested when the loss function of the forecaster is unknown, restrictions on the DGP are required. We consider DGPs with dynamics in the conditional mean and conditional variance, but no dynamics in the remainder of the conditional distribution. These classes of DGPs are quite broad and include ARMA processes and non-linear regressions, possibly with dynamics in the conditional variance process.

### 2.1 Conditional Mean Dynamics Only

Consider first the class of DGPs that satisfy the following conditional homoskedasticity condition:

**Assumption D1:** *The DGP is such that  $Y_{t+h} = \mu_{t+h,t} + \varepsilon_{t+h}$ ,  $\varepsilon_{t+h} | \mathcal{F}_t \sim F_{\varepsilon,h} \left( 0, \sigma_{\varepsilon,h}^2 \right)$ , where*

$F_{\varepsilon,h} \left( 0, \sigma_{\varepsilon,h}^2 \right)$  is some distribution with mean zero and variance  $\sigma_{\varepsilon,h}^2$ , which may depend on  $h$ , but does not depend on  $\mathcal{F}_t$ .

The restriction of dynamics only in the conditional mean implies that the innovation term,  $\varepsilon_{t+h}$ , is drawn from some distribution,  $F_{\varepsilon,h}$ , which will generally depend on the forecast horizon, but is independent of  $\mathcal{F}_t$  and so is not denoted with a subscript  $t$ .

Following a long literature in economics and statistics, we concentrate initially on loss functions that satisfy the following assumption:

**Assumption L1:** *The loss function depends solely on the forecast error. That is,  $L(y, \hat{y}) = L(y - \hat{y}) = L(e)$ ,  $\forall (y, \hat{y}) \in \mathbb{R} \times \mathcal{Y}$ .*

Although assumption L1 rules out certain loss functions, many common loss functions are of this form, for example MSE, mean absolute error (MAE), lin-lin and linex loss. Cases where the loss depends not only on the size of the forecast error but also on the level of the predicted variable are ruled out, however. As we show in Section 3, such cases arise in economics if the costs of prediction errors are not the same in recessions and expansions. They could also occur in meteorology, where under- and over-predictions may be equally costly under normal weather conditions, but may differ under more extreme conditions. For example, underpredicting the strength of a hurricane or tornado could be particularly costly.

Under these two assumptions we obtain the following testable implications of forecast optimality:

**Proposition 1** *Suppose the loss is a function solely of the forecast error (assumption L1) while the DGP has dynamics only in the conditional mean (assumption D1). Then*

1. *The optimal forecast takes the form*

$$\hat{Y}_{t+h,t}^* = \mu_{t+h,t} + \alpha_h^*, \tag{3}$$

where  $\alpha_h^*$  is a constant that depends on the distribution  $F_{\varepsilon,h}$  and the loss function.

2. *The forecast error associated with the optimal forecast,  $e_{t+h,t}^*$ , is independent of all  $Z_t \in \mathcal{F}_t$ . In particular,  $\text{Cov} \left( e_{t+h,t}^*, e_{t+h-j,t-j}^* \right) = 0$  for all  $j \geq h$  and any  $h > 0$ .*
3. *The variance of the forecast error associated with the optimal forecast,  $V(e_{t+h,t}^*)$ , is non-decreasing in the forecast horizon,  $h$ .*

All proofs are contained in the appendix. This proposition shows that under a testable assumption on the DGP, and only one weak assumption on the loss function, the forecast errors

associated with the optimal forecast are serially uncorrelated at lags greater than or equal to the forecast horizon, for any error-based loss function. Given a sequence of realizations and forecasts,  $\left\{ \left( Y_{t+h}, \hat{Y}_{t+h,t} \right) \right\}_{t=1}^T$ , we may therefore test for forecast optimality *without* knowledge of the forecaster's loss function by testing, for example, the serial correlation properties of the forecast errors. For financial applications the assumption of constant higher-order conditional moments is clearly too strong, but in many economic applications the assumption that all dynamics are driven by the conditional mean may be reasonable. In these cases forecast optimality can be tested with a large degree of robustness to the loss function of the forecaster, e.g. by testing  $\beta = 0$  in regressions such as

$$e_{t+h,t} = \alpha + \beta' Z_t + u_{t+h}. \quad (4)$$

Note that no estimate of the conditional mean of  $Y_{t+h}$  is required to conduct this test. This is important since  $\mu_{t+h,t}$  is generally unknown. This test exploits the result that all variables known to have been in the forecaster's information set – except trivially for a constant – are orthogonal to the forecast error, which is observable without knowledge of  $\mu_{t+h,t}$ .

We focus on testing for correlation though of course more generally we could test for complete independence between  $e_{t+h,t}$  and any  $Z_t \in \mathcal{F}_t$ , or other implications of independence. For example, rather than using OLS to estimate the parameters in equation (4), we could use LAD or quantile regression, see Koenker and Bassett (1978).

The last part of the proposition reveals that optimality tests based on the variance of the forecast error being weakly increasing in the forecast horizon require restrictive assumptions either on the loss function (namely, that MSE loss applies) or on the DGP (i.e., dynamics in the conditional mean only).

## 2.2 Conditional Mean and Variance Dynamics

We next provide results for a more general class of conditional scale-location processes that satisfy the following assumption:

**Assumption D2:** *The DGP is such that  $Y_{t+h} = \mu_{t+h,t} + \sigma_{t+h,t}\eta_{t+h}$ ,  $\eta_{t+h}|\mathcal{F}_t \sim F_{\eta,h}(0,1)$ , where  $F_{\eta,h}(0,1)$  is some distribution with mean zero and unit variance, which may depend on  $h$ , but does not depend on  $\mathcal{F}_t$ .*

This class of DGPs is very broad and includes most common volatility processes, e.g. ARCH

and stochastic volatility, see Engle (1982) and Taylor (1982). It nests those of Assumption D1, at the cost that we must be more restrictive on the class of loss functions that we consider below. Specifically, we shall assume that the loss function is homogenous in the forecast error:

**Assumption L2:** *The loss function is a homogeneous function solely of the forecast error.*

This assumption implies that  $L(ae) = g(a)L(e)$  for some positive function  $g$ . Commonly used loss functions such as MSE, MAE, lin-lin, and asymmetric quadratic loss are all of this form. This requirement is needed to ensure that the units of the forecast (e.g., cents versus dollars) do not affect the optimal forecast beyond a scale adjustment. Although this holds for many loss functions encountered in economics and finance, the requirement does exclude cases such as linear-exponential loss.

With these two assumptions we obtain the following testable implications of forecast optimality:

**Proposition 2** *Suppose the loss function is homogenous in the forecast error (assumption L2) while the DGP can have conditional mean and variance dynamics (assumption D2). Define the standardized forecast error associated with the optimal forecast as  $d_{t+h,t}^* = e_{t+h,t}^*/\sigma_{t+h,t}$ . Then:*

1. *The optimal forecast takes the following form:*

$$\hat{Y}_{t+h,t}^* = \mu_{t+h,t} + \sigma_{t+h,t} \cdot \gamma_h^*, \quad (5)$$

where  $\gamma_h^*$  is a constant that depends only on the distribution  $F_{\eta,h}$  and the loss function.

2.  $d_{t+h,t}^*$  is independent of any element  $Z_t \in \mathcal{F}_t$ . In particular,  $\text{Cov}(d_{t+h,t}^{*r}, d_{t+h-j,t-j}^{*s}) = 0$  for all  $j \geq h$  and any  $h > 0$  and all  $r, s$  for which the covariance exists.

Although the forecast error associated with the optimal forecast will not, in general, be unbiased, serially uncorrelated or homoskedastic, the forecast error scaled by the conditional standard deviation will be independent of any  $Z_t \in \mathcal{F}_t$ . This implies that  $d_{t+h,t}^*$  will be serially uncorrelated at lags  $j \geq h$  and homoskedastic. Forecast optimality could therefore be tested by estimating non-parametric or flexible parametric models for the conditional mean and variance:

$$\begin{aligned} d_{t+h,t} &= \alpha_0 + g_1(Z_t; \theta_1) + u_{t+h} \\ u_{t+h} &= \sigma_{u,t+h} v_{t+h}, \quad v_{t+h} \sim (0, 1) \\ \sigma_{u,t+h}^2 &= \omega_0 + g_2(Z_t; \theta_2) \end{aligned} \quad (6)$$

If  $g_1$  and  $g_2$  are defined such that  $g_1(z; 0) = g_2(z; 0) = 0$  for all  $z$ , then a test of optimality may be obtained by testing  $H_0 : \omega_0 = 1 \cap \theta_1 = \theta_2 = 0$  so that the conditional mean of  $d_{t+h,t}$  is time-invariant and its conditional variance equals unity.

The above test is easily computed and does not require an estimate of  $\mu_{t+h,t}$  although it requires that an estimate of  $\sigma_{t+h,t}^2$  is available. Under certain conditions, a consistent estimate of this conditional variance can be obtained from the observed  $Y_t$  process either by means of parametric methods (e.g. using a GARCH-type model) or by non-parametric methods, using a realized volatility estimator (see Andersen, *et al.*, 2001, 2003) for example.

Researchers will not always have a reliable estimate of  $\sigma_{t+h,t}^2$  available, and so it would be particularly useful to establish under which conditions an optimality test can be based only on observables. Suppose we restrict the first and second moment dynamics to be linked in a manner consistent with the widely used “constant coefficient of variation” model:

**Assumption D2’:** *The DGP is such that  $Y_{t+h} = \beta\sigma_{t+h,t} + \sigma_{t+h,t}\eta_{t+h}$ ,  $\eta_{t+h}|\mathcal{F}_t \sim F_{\eta,h}(0, 1)$ , where  $\beta \in \mathbb{R}$  and  $F_{\eta,h}$  is some distribution with mean zero and unit variance, which may depend on  $h$ , but does not depend on  $\mathcal{F}_t$ .*

This “ARCH-in-mean” model is extensively used in financial applications, i.e. when the target variable is returns and expected returns are proportional to the level of risk as measured by the conditional standard deviation, see Engle, *et al.* (1987). Using this assumption, we obtain

**Corollary 1** *Suppose the loss function is homogenous in the forecast error (assumption L2), while the DGP satisfies assumption D2’ for  $\beta \neq -\gamma_h^*$ . Define  $\hat{d}_{t+h,t}^* \equiv (Y_{t+h} - \hat{Y}_{t+h,t}^*)/\hat{Y}_{t+h,t}^*$ . Then  $\hat{d}_{t+h,t}^*$  is independent of any element  $Z_t \in \mathcal{F}_t$ . In particular,  $\text{Cov}(\hat{d}_{t+h,t}^{*r}, \hat{d}_{t+h-j,t-j}^{*s}) = 0$  for all  $j \geq h$  and any  $h > 0$  and all  $r, s$  for which the covariance exists.*

Note that the assumption that  $\beta \neq -\gamma_h^*$  is easily checked: if this is true then the optimal forecast is identically zero for all  $t$ . Under the conditions of Corollary 1 we may test forecast optimality without specific knowledge of the loss function or any of the moments of the DGP, by testing that there is no serial correlation beyond lag  $h - 1$  in  $\hat{d}_{t+h,t} = (Y_{t+h} - \hat{Y}_{t+h,t})/\hat{Y}_{t+h,t}$ , and/or that the  $\hat{d}_{t+h,t}$  series is homoskedastic conditional on any  $Z_t \in \mathcal{F}_t$ . This could be done simply via a regression of powers of  $\hat{d}_{t+h,t}$  on a constant and lags, of order greater than or equal to  $h$ , of various powers of  $\hat{d}_{t+h,t}$ , or as in equation (6).

### 2.3 Quantile Tests

Under the conditions of Propositions 1 or 2 it is possible to show that we can express the optimal forecast as a conditional quantile of the variable of interest. The usefulness of this result lies in the surprising finding that the optimal forecast is the *same* quantile at all points in time, though the quantile may change with the forecast horizon and is typically unknown as it depends on the loss function. With such a representation we can obtain an alternative test of forecast optimality using tests of quantile forecasts, eliminating the need to estimate the conditional variance of the variable of interest.

**Proposition 3** *Suppose that either (i) the loss is a function solely of the forecast error (assumption L1) while the DGP has dynamics only in the conditional mean (assumption D1), or (ii) the loss function is homogenous in the forecast error (assumption L2) while the DGP has dynamics in the conditional mean and variance (assumption D2). Then:*

1. *The optimal forecast is such that, for all  $t$ ,*

$$F_{t+h,t}(\hat{Y}_{t+h,t}^*) = q_h^*, \quad (7)$$

where  $q_h^* \in (0, 1)$  depends only on the distribution  $F_{\eta,h}$  and the loss function. If  $F_{t+h,t}$  is continuous and strictly increasing, then we can alternatively express this as:

$$\hat{Y}_{t+h,t}^* = F_{t+h,t}^{-1}(q_h^*). \quad (8)$$

2. *Let*

$$I_{t+h,t}^* \equiv \mathbf{1}(Y_{t+h} \leq \hat{Y}_{t+h,t}^*),$$

where  $\mathbf{1}(\mathcal{A})$  equals one if  $\mathcal{A}$  is true and zero otherwise. Then  $I_{t+h,t}^*$  is independent of all  $Z_t \in \mathcal{F}_t$ . In particular,  $I_{t+h,t}^* - q_h^*$  is a martingale difference sequence with respect to  $\mathcal{F}_t$ .

Notice how assumptions on the loss function can be traded off against assumptions on the DGP. This proposition gives rise to a new test that is applicable even though  $q_h^*$  is unknown. The test simply projects the indicator function on elements in  $\mathcal{F}_t$  and an intercept and tests that  $\beta = 0$ :

$$I_{t+h,t}^* = \alpha + \beta' Z_t + u_{t+h}. \quad (9)$$

Alternatively, a logit model could be used, to better reflect the binary nature of the dependent variable, or the LR test of independence of Christoffersen (1998) could be employed to test for

serial dependence in  $I_{t+h,t}$ . If  $q_h^*$  is known, it can further be tested that  $\alpha = q_h^*$ . Even in the common case where  $q_h^*$  is unknown, the key point to notice is that the quantile test does not require knowledge of the true values of either  $\mu_{t+h,t}$  or  $\sigma_{t+h,t}$ . The test also does not require knowledge of the time-varying conditional distribution,  $F_{t+h,t}$ , which is unknown in practice. This can be compared to tests based on the probability integral transform of the data,  $F_{t+h,t}(Y_{t+h})$ , which should be Uniform(0,1) under the null that the forecasting model is correctly specified. Such tests can be difficult to implement given only a sequence of point forecasts, and without knowledge of the forecaster's loss function or the forecaster's information set, which is generally unobservable to the forecast evaluator. Due to this robustness property and the minimal information required for their implementation, we think that this new class of quantile tests of forecast optimality is likely to find widespread use in empirical work.

Of the existing tests of forecast optimality in the literature, the one that allows for the greatest flexibility with respect to the loss function is due to Elliott, et al. (2005). These authors derive tests of forecast optimality when the loss function is assumed to belong to the two-parameter family

$$L(e_{t+1}; \alpha, p) \equiv [\alpha + (1 - 2\alpha)\mathbf{1}(e_{t+1} < 0)] |e_{t+1}|^p, \quad (10)$$

with a positive exponent  $p$  and an asymmetry parameter  $\alpha$ ,  $0 < \alpha < 1$ . This is very different from the approach proposed in this section, which does not constrain the loss function to belong to a pre-specified parametric family of loss functions. In Section 3 we suggest an extension of the approach of Elliott, et al. (2005) for applications where assumptions L1 or L2 do not hold.

While Proposition 3 holds quite generally, it of course only applies to *point* forecasts. One possible extension of our work includes establishing results for density forecasting, an area recently considered by Garratt et al. (2003) and Campbell and Diebold (2005), although applications of such results to empirical data will be limited since, currently, the vast majority of published forecasts are point forecasts.

## 2.4 Non-homogeneous Loss or Dynamics in Higher-order Moments

The above results under unknown loss constitute an exhaustive set of testable properties in the following sense. First, suppose that the loss function is not homogeneous (L2). Allowing for conditional variance dynamics (as in assumption D2) then makes it difficult to obtain testable implications of forecast optimality that are robust to the loss function. Further, when there is

dynamics in third or higher order moments and assumption D2 fails to hold it is generally difficult to obtain results that are easy to test even if homogeneity is imposed on the loss function. To see this, consider the following more general DGP:

**Assumption D3:** *The DGP is such that  $Y_{t+h} = \mu_{t+h,t} + \sigma_{t+h,t}\eta_{t+h}$ ,  $\eta_{t+h}|\mathcal{F}_t \sim F_{\eta,t+h,t}(0,1)$ , where  $F_{\eta,t+h,t}(0,1)$  is some time-varying distribution, with mean zero and unit variance, that depends on  $\mathcal{F}_t$  and possibly on  $h$ .*

This class of DGPs nests those of Assumption D2, as we allow for a time-varying conditional mean, conditional variance and other properties of the distribution (e.g., time-varying conditional skew or kurtosis.) We obtain the following “impossibility” result:

**Proposition 4** 1. *Suppose the DGP has conditional mean and variance dynamics (assumption D2), but that the loss function is not homogenous in the forecast error (assumption L2 is violated). Then:*

$$\hat{Y}_{t+h,t}^* = \mu_{t+h,t} + \alpha_{t+h,t}^*, \quad (11)$$

where  $\alpha_{t+h,t}^*$  is a scalar that depends on the loss function,  $\sigma_{t+h,t}$ , and  $F_{\eta,h}$ .

2. *Suppose the loss function is homogenous in the forecast error (assumption L2), and that the DGP has dynamics beyond the conditional mean and variance (assumption D3). Then:*

$$\hat{Y}_{t+h,t}^* = \mu_{t+h,t} + \sigma_{t+h,t} \cdot \gamma_{t+h,t}^*, \quad (12)$$

where  $\gamma_{t+h,t}^*$  is a scalar that depends on the loss function and  $F_{\eta,t+h,t}$ .

Under the conditions for Part 1 or Part 2 of the above proposition, the forecast error associated with the optimal loss,  $e_{t+h,t}^* \equiv Y_{t+h} - \hat{Y}_{t+h,t}^*$ , its standardized equivalent,  $d_{t+h,t}^* = e_{t+h,t}^*/\sigma_{t+h,t}$ , and the indicator variable  $I_{t+h,t}^* \equiv \mathbf{1}(Y_{t+h} \leq \hat{Y}_{t+h,t}^*)$ , will all, in general, be correlated with some  $Z_t \in \mathcal{F}_t$  in a way that depends on the unknown loss function. Thus if the loss function is not homogeneous, or if there are higher-order dynamics in the DGP, these objects will not generally be useful for testing forecast optimality. In these cases it is difficult to obtain testable restrictions on the forecast error that do not require knowledge of the shape of the loss function, even if  $\sigma_{t+h,t}$  is known (or, more generally, even if  $F_{t+h,t}$  is known).

Of course, if the loss function and the data generating process are known, optimality properties of the forecast can be derived directly. For example, assuming linex loss and a first-order Markov

Gaussian mixture model, Patton and Timmermann (2004) derive the properties of the optimal forecast and forecast errors.

## 2.5 Empirical Results

We next return to the Federal Reserve Greenbook forecasts of output growth to demonstrate the theoretical results. This is the data shown in Figure 1. These forecasts are all one-quarter ahead forecasts, so  $h = 1$ .

Our realized values of GDP are taken as the 2006Q1 “vintage” of the real GDP growth figures over the sample period, obtained from the Federal Reserve Bank of Philadelphia’s web page. This data was studied in depth by Croushore and Stark (2001). The data and a detailed description of its construction are available at <http://www.phil.frb.org/econ/forecast/reaindex.html>. In Section 2.5.1 we discuss the results obtained using an earlier vintage of real GDP growth data, and show that they are similar to those obtained using the most recent vintage.

Under the joint assumption that the Federal Reserve’s forecasts are optimal, its loss is solely a function of the forecast error (assumption L1), and there are no dynamics beyond the conditional mean (assumption D1) of real GDP growth, Proposition 1 shows that the Greenbook forecast errors would be conditionally homoskedastic. We test for conditional heteroskedasticity in the Greenbook forecast errors using Engle’s (1982) test which is based on the null that  $\alpha_i = 0$  for all  $i \geq 1$  in the regression:

$$e_{t+1,t}^2 = \alpha_0 + \alpha_1 e_{t,t-1}^2 + \dots + \alpha_L e_{t-L+1,t-L}^2 + u_{t+1}.$$

A test with four lags generated a  $p$ -value of 0.03, indicating the presence of conditional heteroskedasticity. This suggests either that the Fed’s forecasts are sub-optimal, or that assumption D1 or L1 is violated.

Testing assumption D1 directly requires modeling the conditional mean of output growth. We used an ARMA(1,1) model for the conditional mean, which was found to remove all significant serial correlation, and then tested for serial correlation in the squared residuals from this model. The  $p$ -values from this test were 0.65, 0.09 and 0.01 for lags 1, 4 and 8. These results indicate the presence of conditional heteroskedasticity and hence suggest that D1 is an inappropriate assumption for quarterly output growth.

The previous results suggest that we need to allow for more general dynamics in the data

generating process for real GDP growth. For this purpose we tested assumption D2 directly on the real GDP growth series by estimating an ARMA(1,1) model for the conditional mean and a GARCH(1,1) model for the conditional variance, c.f. Bollerslev (1986). Simple, parsimonious, GARCH models have been shown to work well in numerous other studies of macroeconomic and financial time series. We tested for serial correlation in the first four powers of the standardized residuals, and found no evidence of any serial correlation out to eight lags, suggesting that D2 may be an appropriate assumption for this data. The generalized tests of Hong (1999) could also have been used to test Assumption D2.

To test the joint hypothesis that forecasts are optimal, that the loss function is homogenous in the forecast error, and that the more general assumption D2 (which allows for conditional heteroskedasticity) holds for our GDP data, we use the result from Proposition 2 that, under these assumptions, the forecast error associated with the optimal forecast will be of the following form:

$$\begin{aligned} e_{t+1,t} &= \gamma_1 \sigma_{t+1,t} + \sigma_{t+1,t} \eta_{t+1} \\ \eta_{t+1,t} | \mathcal{F}_t &\sim F_{\eta,1}(0, 1). \end{aligned} \tag{13}$$

Under the assumption of Corollary 1, the forecast  $\hat{Y}_{t+1,t}$  already incorporates information on the conditional mean of  $Y_{t+1}$  and so an independent estimate of  $\mu_{t+1,t}$  is not required in this test. If we knew  $\sigma_{t+1,t}$  we could construct  $d_{t+1,t} \equiv e_{t+1,t}/\sigma_{t+1,t}$ , which, under the conditions for Proposition 2, will be serially uncorrelated, conditionally homoskedastic, and have constant conditional higher-order moments. In order to implement a test we model the conditional variance as a simple GARCH(1,1) process, allowing for the GARCH-in-mean effects implied by Proposition 2. Of course, the possibility remains that this model is mis-specified and that this affects our conclusions. The GARCH-in-mean parameter,  $\gamma_1$ , was significant (with a  $p$ -value less than 0.01), presenting further evidence against the optimality of these forecasts under MSE loss. Simple tests for serial correlation in the third and fourth powers of the forecast errors standardized by the GARCH estimates of conditional variance did not reveal any evidence against assumption D2, conditional on the forecasts being optimal and assumption L2 being satisfied.

If we could further assume that assumption D2' holds for real GDP growth then we could employ Corollary 1 and use the Greenbook forecasts to standardize the forecast errors. Under the joint assumption that D2' and L2 holds and forecasts are optimal,  $\hat{d}_{t+1,t} = (Y_{t+1} - \hat{Y}_{t+1,t})/\hat{Y}_{t+1,t}$  should be independent of any element in the forecaster's information set, including past (standardized)

forecast errors and any transformation of these. We therefore tested for serial correlation in  $\hat{d}_{t+1,t}$  and  $\hat{d}_{t+1,t}^2$ . We found no evidence of serial correlation in  $\hat{d}_{t+1,t}$  at any lag up to eight, but we found significant evidence of serial correlation in  $\hat{d}_{t+1,t}^2$ : the  $p$ -values from a Ljung-Box test with four lags was 0.01, indicating significant serial correlation. This indicates that the joint hypothesis that D2' holds, that loss is homogeneous in the forecast error (L2) and that the Greenbook forecasts are optimal can be rejected.

Given the plausibility of assumption D2 for real GDP growth, and the violation of assumption D2', we now employ the test of forecast optimality based on Proposition 3. Rather than rely on a (possibly mis-specified) parametric estimator of the conditional variance of the real GDP growth series, we implement the test of forecast rationality under unknown loss using the indicator variable  $I_{t+1,t} \equiv \mathbf{1}(Y_{t+1} \leq \hat{Y}_{t+1,t})$ . This variable should be independent of all elements of  $\mathcal{F}_t$  under the null of forecast rationality, and does not require any specification of the mean or variance. We consider two parsimonious tests of this restriction, both obtained via OLS regression:

$$\begin{aligned} I_{t+1,t} &= \frac{0.346}{(0.071)} + \frac{0.036}{(0.016)} \hat{Y}_{t+1,t} + u_{t+1} \\ I_{t+1,t} &= \frac{0.334}{(0.095)} + \frac{0.039}{(0.015)} \hat{Y}_{t+1,t} + \frac{0.036}{(0.108)} I_{t,t-1} + u_{t+1}. \end{aligned}$$

The first regression reveals significant evidence of correlation between the indicator variable and the forecast: the  $t$ -statistic is 2.27, which is significant at the 0.05 level. In the second regression, when we include both the lagged indicator variable and the forecast, we find that the coefficient on the lagged indicator is not significantly different from zero, while the coefficient on the forecast is significant at the 0.05 level. A joint test that both coefficients are equal to zero yields a  $p$ -value of 0.06, which means that the null is not rejected at the 0.05 level but is rejected at the 0.10 level.

Since our tests revealed no evidence against assumption D2, these results constitute evidence against the optimality of the Greenbook forecasts under *any* loss function that is homogeneous in the forecast error. If, however, the Fed's loss function does not satisfy this restriction, then these forecasts may yet be optimal under a more general loss function. Overall, we take these regressions and the earlier results as evidence against the optimality of the Greenbook forecasts under unknown loss functions that are homogenous in the forecast error.

### 2.5.1 Real-time Data

In addition to the 2006Q1 “vintage” of real GDP growth figures, we also consider using “real-time” GDP figures, that is, the figures that historically would have been available to forecasters at each point in time. Diebold and Rudebusch (1991) and Croushore and Stark (2001) discuss the differences between using revised data versus real-time data. This data takes into account the possible effect of measurement errors on the forecasting performance as measured historically, and uses the fact that the Fed may have as its forecast target the initial release of GDP figures, rather than the underlying true GDP figures, as proxied by the latest vintage data. Following studies such as Romer and Romer (2000) and Capistran (2006), we use the second revision of the real GDP growth figures.

The results we obtain are consistent with those reported above for the final revision. For example, we find evidence against optimality under MSE loss using this data from the following regression:

$$e_{t+1,t} = 1.046 - 0.300\hat{Y}_{t+1,t} + 0.072e_{t,t-1} + u_{t+1}$$

(0.985)      (0.253)      (0.111)

While none of the coefficients in this model are individually significant at the 0.05 level, a joint test that all parameters equal zero yields a  $\chi^2_3$  statistic ( $p$ -value) of 8.56 (0.04), which implies a rejection of this restriction at the 0.05 level. We also find significant serial correlation in the real-time forecast errors, indicating either that assumptions D1 and L1 do not hold or that the forecasts are suboptimal. Furthermore, using the real-time data, the indicator variable regressions yielded

$$I_{t+1,t} = 0.362 + 0.027\hat{Y}_{t+1,t} + u_{t+1}$$

(0.067)      (0.019)

$$I_{t+1,t} = 0.255 + 0.028\hat{Y}_{t+1,t} + 0.250I_{t,t-1} + u_{t+1}$$

(0.076)      (0.019)      (0.087)

So for the real-time data the coefficient on the forecast is not significant while the coefficient on the lagged indicator is now highly significant. The  $\chi^2_2$  statistic and  $p$ -value from the second regression are 10.07 and 0.01, so again we strongly reject optimality in conjunction with assumptions D2 and L2.

### 2.5.2 The “Great Moderation”

A factor that may affect our results is the presence of a structural break in the variance of U.S. GDP growth, known as the “great moderation”, generally thought to have occurred around 1984Q1, see

Kim and Nelson (1999), McConnell and Perez-Quiros (2000) and Stock and Watson (2002). Breaks such as these are widely regarded as a key source of forecast failure, c.f. Clements and Hendry (1998, 2006). To explore this possibility, we carried out separate tests using data up to 1983Q4 (61 observations) and from 1984Q1 onwards (64 observations). The sub-sample results were very similar to those obtained from the full sample: the Mincer-Zarnowitz tests of forecast optimality under MSE loss yielded the following results:

$$\begin{aligned} \text{Pre-1984 sample } Y_{t+1} &= \begin{matrix} 0.953 \\ (0.655) \end{matrix} + \begin{matrix} 0.692 \\ (0.118) \end{matrix} \hat{Y}_{t+1,t} + u_{t+1}, & R^2 = 0.225 \\ \text{Post-1984 sample } Y_{t+1} &= \begin{matrix} 1.347 \\ (1.127) \end{matrix} + \begin{matrix} 0.782 \\ (0.381) \end{matrix} \hat{Y}_{t+1,t} + u_{t+1}, & R^2 = 0.144 \end{aligned}$$

The  $\chi^2_2$  statistics ( $p$ -values) that the intercepts equal zero and the slope coefficients equal one were 7.19 (0.03) and 7.66 (0.02) in the first and second sub-samples, respectively. Thus optimality under MSE loss is rejected in both sub-samples, as well as in the full sample.

The quantile-based test for optimality under unknown, homogeneous loss yielded the following results for the two sub-samples:

$$\begin{aligned} \text{Pre-1984 sample } I_{t+1,t} &= \begin{matrix} 0.556 \\ (0.108) \end{matrix} + \begin{matrix} 0.025 \\ (0.020) \end{matrix} \hat{Y}_{t+1,t} - \begin{matrix} 0.246 \\ (0.127) \end{matrix} I_{t,t-1} + u_{t+1} \\ \text{Post-1984 sample } I_{t+1,t} &= \begin{matrix} 0.059 \\ (0.171) \end{matrix} + \begin{matrix} 0.085 \\ (0.058) \end{matrix} \hat{Y}_{t+1,t} + \begin{matrix} 0.288 \\ (0.121) \end{matrix} I_{t,t-1} + u_{t+1} \end{aligned}$$

The  $\chi^2_2$  statistics ( $p$ -values) that the both slope coefficients equal zero in each regression were 6.09 (0.05) and 7.41 (0.02) in the first and second sub-samples. Again forecast optimality in conjunction with assumptions L2 and D2 is rejected in both sub-samples, as in the full sample. The similarity between the results in the two sub-samples and the results from the full sample of data, under both MSE and unknown loss, provides some assurance that a structural break in U.S. GDP growth volatility is not driving our findings.

### 3 Testable Implications under General Data Generating Processes

The rejection of optimality for the Greenbook forecasts reported in the previous section is surprising: the null hypothesis is very general and covers a large class of loss functions. One might expect that allowing for such a wide range of different loss functions would erode the power of our tests, particularly when applied to a sequence of forecasts that we expect ex-ante to be “good”, such as

those from the Federal Reserve. As shown in our simulation study, presented in Section 4 below, this turns out not to be the case.

However, it might realistically be the case that the Federal Reserve's loss function cannot be assumed to be solely a function of the forecast error. To a conservative policy maker, over-predictions of economic growth are likely not only to be more costly than under-predictions but also to be disproportionately costly in periods of low growth as it may incorrectly signal a recovery from a recession. This, in turn, is likely to lead to wrong choices in how monetary policy is set. This points to a need for considering forecast optimality in situations where not only the forecast error but also the level of the predicted variable matter. For such cases it is possible to construct a test based on a flexible parametric estimate of the first-derivative of the loss function with respect to  $\hat{y}$ . The first-order condition  $E_t \left[ \partial L \left( Y_{t+h}, \hat{Y}_{t+h,t}^* \right) / \partial \hat{y} \right] = 0$  implies  $E \left[ \partial L \left( Y_{t+h}, \hat{Y}_{t+h,t}^* \right) / \partial \hat{y} \cdot Z_t \right] = 0$  for any  $Z_t \in \mathcal{F}_t$ . For notational simplicity, let

$$\lambda(y, \hat{y}) \equiv \frac{\partial L(y, \hat{y})}{\partial \hat{y}}. \quad (14)$$

We may obtain a flexible parametric estimate of  $\lambda(y, \hat{y})$ , denoted  $\lambda(y, \hat{y}; \theta)$ , based on a linear spline model, for example. To see how a linear spline could be employed to approximate the function  $\lambda(y, \hat{y})$ , assume initially that  $\lambda = \partial L(e) / \partial e$ , let  $(\zeta_1, \dots, \zeta_K)$  be the nodes of the spline and impose that one of the nodes is zero. We impose that the spline is continuous, though not necessarily differentiable, except possibly at zero. We could allow discontinuities in  $\lambda$  at the cost of introducing more parameters to estimate.

With just a few nodes this class of loss functions is very flexible, nesting both MSE and MAE as special cases, as well as lin-lin, the symmetric, non-convex loss function of Granger (1969), and the class of loss functions used by Elliott, *et al.* (2005). If we further impose that the spline is continuous at zero, then MSE loss is nested in the interior of the parameter space, at the cost of the MAE and lin-lin loss functions not being nested. In this case the resulting estimated loss function is a piecewise quadratic spline, with piecewise linear derivative,  $\lambda(e; \theta)$ , which is continuous and differentiable everywhere (except at the  $K$  nodes):

$$\frac{\partial \lambda(e; \theta)}{\partial e} = \begin{cases} \gamma_1, & \text{for } e \leq \zeta_1 \\ \gamma_i, & \text{for } \zeta_{i-1} < e \leq \zeta_i, i = 2, \dots, K \\ \gamma_{K+1}, & \text{for } e > \zeta_K \end{cases}, \quad (15)$$

where  $\theta = [\gamma_1, \gamma_2, \dots, \gamma_{K+1}]'$ .  $\lambda(e; \theta)$  and  $L(e; \theta)$  are constructed from the above specification by imposing that  $\lambda(0; \theta) = L(0; \theta) = 0$  and that both  $\lambda(e; \theta)$  and  $L(e; \theta)$  are continuous in  $e$ . Since  $\lambda(e; \theta)$  is only identified up to a multiplicative constant, some normalization is required to identify the parameters; for example, we could impose that  $\sum_{i=1}^{K+1} \gamma_i = 1$ . Further, it is important to impose constraints on  $\theta$  so that the resulting estimate of  $\lambda$  satisfies the assumptions required for it to be the first derivative of some valid loss function, for example, that the loss function is weakly increasing in the absolute value of the forecast error.

In applications where we have reason to assume that the loss from a forecast is solely a function of the forecast error, i.e. assumption L1 is satisfied, the problem simplifies to approximating the function  $\lambda(y - \hat{y}) = \lambda(e)$ . In other applications, such a restriction may not be well-founded and so no such simplification is available. In this case we must employ a more flexible specification to approximate the function  $\lambda(y, \hat{y})$ , or, equivalently,  $\lambda(e, y)$ . Treating  $\lambda(e, y)$  rather than  $\lambda(y, \hat{y})$  makes it simpler to impose the required conditions on  $\lambda$ . We propose the following specification, which is structured so that MSE loss is obtained when all parameters are set equal to zero..

$$\frac{\partial \lambda(e, y; \theta)}{\partial e} = \begin{cases} \gamma_1(y) \equiv \Gamma(\varphi_{01} + \varphi_{11}y - \ln K), & \text{for } e \leq \zeta_1 \\ \gamma_i(y) \equiv \left(1 - \sum_{j=1}^{i-1} \gamma_j\right) \cdot \Gamma(\varphi_{0i} + \varphi_{1i}y - \ln K), & \text{for } \zeta_{i-1} < e \leq \zeta_i, i = 2, \dots, K \\ \gamma_{K+1}(y) = 1 - \sum_{j=1}^K \gamma_j, & \text{for } e > \zeta_K \end{cases} \quad (16)$$

where  $\Gamma(x) \equiv (1 + \exp(-x))^{-1}$  is the logistic transformation. This specification allows  $y$  to affect the slopes of  $\lambda$ , guarantees that all slopes are weakly positive, and that the sum of the slopes equals one. Alternative ways of restricting the  $\gamma_i$ 's to be non-negative are possible of course, e.g. by using a probit instead of the logistic function or by including additional powers of  $y$  at the cost of having to estimate more parameters. These issues can be addressed in empirical work by means of a sensitivity analysis with regard to the assumed form of  $\partial \lambda / \partial e$ .

Under standard regularity conditions the parameter vector of the approximating function can be estimated via the generalized method of moments (GMM):

$$\hat{\theta}_T \equiv \arg \min_{\theta \in \Theta} g_T(\theta)' W g_T(\theta) \quad (17)$$

$$g_T(\theta) \equiv \frac{1}{T} \sum_{t=1}^T \lambda(e_{t+h,t}, \hat{Y}_{t+h,t}; \theta) \cdot Z_t,$$

where  $W$  is a weighting matrix and  $\Theta$  is a compact set. A test of forecast optimality can be obtained from a test of over-identifying restrictions if we ensure that we have more moment restrictions,  $k$ ,

than parameters,  $p$ :

$$Tg_T(\hat{\theta}_T)' \hat{W}_{Tg_T}^* g_T(\hat{\theta}_T) \Rightarrow \chi_{k-p}^2, \text{ as } T \rightarrow \infty. \quad (18)$$

Here  $\hat{W}_T^*$  is a consistent estimate of the optimal weight matrix, c.f. Newey and McFadden (1994). This test of forecast optimality does not rely on any restrictions on the DGP, other than standard conditions required for GMM estimators to be consistent and asymptotically normal. It does, however, rely on the linear spline being sufficiently flexible to approximate the unknown loss function. Thus, a rejection of forecast optimality may be due either to a true failure of forecast optimality or to a failure of the approximation of the forecaster's loss function.

In contrast with the analysis in Elliott, *et al.* (2005), our tests do not assume that the loss function belongs to a two-parameter family, nor do we need to restrict the loss function only to depend on the forecast error. Indeed, when the latter restriction holds, we recommend using the quantile-based test presented in Proposition 3.

### 3.1 Empirical Results

We now apply the spline-based tests of forecast optimality to the Greenbook forecasts. As outlined above, we use a quadratic spline for the loss function. Initially we assume that the loss function is solely a function of the forecast error. We employ three nodes,  $[-2, 0, 2]$ , which correspond to the 0.17, 0.44 and 0.70 quantiles of the empirical distribution of the forecast errors. With just three nodes we require  $\gamma_i \geq 0$  for all  $i$  for  $\lambda$  to correspond to the first-derivative of a valid loss function. When the number of nodes exceeds three, non-negative loss and continuity of the derivatives does not rule out that some of the  $\gamma_i$  values associated with the middle segments are negative. However, the central (adjacent to zero) and outermost segments must have non-negative  $\gamma_i$  values so as not to violate these restrictions. In the estimation procedure we normalize the function by imposing  $\sum_{i=1}^4 \gamma_i = 1$ .

As instruments for the moment conditions we use a constant, the contemporaneous value of the forecast, and one lag each of the forecast error, realized GDP growth, and the derivative,  $\lambda(y, \hat{y})$ . We thus have five moment conditions and three free parameters. The estimated loss function is presented in Figure 3. This figure reveals that the estimated loss function is asymmetric, penalizing negative forecast errors (over-predictions) more than positive forecast errors (under-predictions). The average ratio of the loss from a negative forecast error to a positive error of the same magnitude,

i.e.  $L(-e)/L(e)$ , for errors in the empirically relevant range  $[0, 10]$  is 1.44, with minimum and maximum values of 0.52 and 2.76, respectively.

The test of forecast optimality under the estimated loss function, obtained by using the two over-identifying moment conditions, yields a  $\chi^2_2$ -statistic ( $p$ -value) of 5.57 (0.06). A rejection of optimality using this test is consistent with the results in the previous section, which held for all loss functions that are homogeneous in the forecast error. The spline-based loss functions considered here are generally not homogeneous, and so they relax that restriction, but they are still constrained to be functions solely of the forecast error. Nevertheless, the empirical results are conditional on our choice of a specific number of nodes (three) used in the spline function.

To test the sensitivity of these results to the assumption of a logit specification as in equation (16), we also implemented a probit-type specification, and obtained very similar results: the  $p$ -value from the  $J$ -test for the probit specification of the spline model, for example was 0.05, compared with 0.06 for the logit specification.

Finally, we estimate a more flexible specification of the Fed's loss function, allowing the loss function to depend on the forecast and the realization separately, rather than solely through the forecast error. As discussed above, we model  $\lambda$  as a function of  $(e, y)$  rather than  $(y, \hat{y})$  as it is simpler to impose the required constraints on the former than the latter. With three nodes, this model has a total of six free parameters, namely the intercept and slope parameters  $(\varphi_{01}, \varphi_{02}, \varphi_{03}, \varphi_{11}, \varphi_{12}, \varphi_{13})$  in the logistic specification in equation (16) compared with three in the simpler case,  $(\varphi_{01}, \varphi_{02}, \varphi_{03})$ . We used as additional instruments one extra lag of the forecast error, realized GDP growth and the derivative  $\lambda(y, \hat{y})$ . We plot the estimated loss function in Figure 4. To show the impact of realized GDP growth on the loss we plot the loss as a function of the forecast error, when realized GDP growth is fixed at its unconditional 0.25, 0.5 and 0.75 quantiles, corresponding to periods of low, average or high economic growth. This figure shows that the level of realized GDP growth has substantial impact on the degree of asymmetry in the loss function: during periods of high economic growth the loss function is approximately symmetric, whereas when realized GDP growth is at a lower level the asymmetry becomes more pronounced, and, as in the simpler model, over-predictions are penalized more heavily than under-predictions. The average ratio of the loss from a negative forecast error to a positive error of the same magnitude, i.e.  $L(-e, y)/L(e, y)$ , for errors in the range  $[0, 10]$  and for real GDP growth equal to its 0.25, 0.5 and 0.75 quantiles is 3.48, 1.97 and 1.26. Thus negative errors (over-predictions) are over three times as costly to the Fed

as positive errors (under-predictions) when GDP growth is low, whereas negative errors are only about 25% more costly than positive errors when GDP growth is high.

Using the two over-identifying moment conditions to test the optimality of the Greenbook forecasts based on this more flexible model yields a  $\chi^2$ -statistic ( $p$ -value) of 0.02 (0.99). To check the robustness of this result, we extended the flexible spline model to allow not only  $y$  but also  $y^2$  to affect the slopes of  $\lambda$ . We implemented this model, employing an extra four lags of the forecast errors as additional moment conditions, and found the estimated loss function to be very similar to those in Figure 4. The  $p$ -value from the  $J$ -test of the over-identifying restrictions for this more flexible model was 0.56, compared with 0.99 for the benchmark case.

These large  $p$ -values imply that we have no evidence against the optimality of the Federal Reserve's real GDP forecasts under this flexible loss function. This finding suggests that adequate modelling of the Federal Reserve's loss function requires allowing both the forecast and the realization of real GDP growth to enter separately into the loss function, rather than just through their difference. This is an important result for future researchers studying the forecast performance and objectives of the Federal Reserve.

Our most flexible estimated loss function is consistent with a scenario where the Federal Reserve seeks to issue 'conservative' estimates of future economic growth: under-estimates are penalized less heavily than over-estimates. This conservatism is particularly important when economic growth is moderate or low. Issuing conservative forecasts means that actual economic growth comes in above the forecast more often than not (actual GDP growth was greater than the predicted value in 56% of the quarters in our sample), leading to a positive 'surprise'.

Sorting on the predictions,  $\hat{y}$ , we found that for the lowest 25% of predicted values, there was a 65% probability that the actual comes in higher than the predicted value. In comparison, for the top 75% of predicted values, only 45% of the actual values come in above the predicted values. This is again consistent with the forecaster attempting to avoid overpredicting output growth in low-growth states of the world (as reflected in a current low estimate of  $\hat{y}$ ).

## 4 Simulation Results

To shed light on the finite-sample properties of the tests considered so far we present the results of a small simulation study tailored to capture properties of our data. Spline models of the loss

function are new to the literature, and GMM estimation is known to sometimes have problems in finite samples, see Hall (2005), and so a study of these tests in finite samples is potentially of value. We employ a simple but representative AR(1)-GARCH(1,1) model as the data generating process for the simulation:

$$\begin{aligned} Y_t &= 0.5Y_{t-1} + \sigma_t \varepsilon_t, \quad t = 1, 2, \dots, T \\ \sigma_t^2 &= 0.1 + 0.8\sigma_{t-1}^2 + 0.1\sigma_{t-1}^2 \varepsilon_{t-1}^2 \\ \varepsilon_t &\sim iid N(0, 1) \end{aligned} \tag{19}$$

We consider three sample sizes,  $T = 100, 250, 1000$ , and two loss functions, MSE loss and an asymmetric quadratic loss function:

$$L(e; a) = \begin{cases} ae^2, & e > 0 \\ e^2, & e \leq 0 \end{cases} \tag{20}$$

where we set  $a = 1.84$ . Both of these loss functions are homogeneous in the forecast error, and so by Proposition 2 we know that the optimal one-step ahead forecast will take the form

$$\hat{Y}_{t+1,t}^* = \mu_{t+1,t} + \sigma_{t+1,t} \gamma_1^* \tag{21}$$

Under MSE loss we have  $\gamma_1^* = 0$ , while under asymmetric quadratic loss we have  $\gamma_1^* = 0.25$ . To gauge the power of the tests we need to construct a forecast that is sub-optimal under *all* loss functions, not just MSE or some other given loss function. One such set of sub-optimal forecasts is obtained by simply adding independent noise to the optimal forecast:

$$\begin{aligned} \hat{Y}_{t+1,t} &= \hat{Y}_{t+1,t}^* + \xi \varepsilon_{t+1} \\ \varepsilon_{t+1} &\sim iid N(0, 1) \end{aligned} \tag{22}$$

This specification is representative of a number of sources of sub-optimality. For example, the noise could come from including irrelevant variables in a prediction model, or from random “judgemental adjustments” of a statistical forecasting model. We consider five values for the standard deviation of the noise:  $\xi = 0, 0.25, 0.5, 0.75$  and  $1$ . The case of  $\xi = 0$  corresponds, of course, to the forecast being optimal under MSE or asymmetric quadratic loss, and can thus be used to examine the finite-sample size of the tests.

We study five tests of forecast optimality. The first test is the standard Mincer-Zarnowitz (1969) test, based on the regression:

$$Y_{t+1} = \beta_0 + \beta_1 \hat{Y}_{t+1,t} + u_{t+1}$$

with the null hypothesis being  $\beta_0 = 0 \cap \beta_1 = 1$ . This test is appropriate for testing optimality under MSE loss, but may over-reject forecasts that are optimal under other loss functions. For this test, and the other two regression-based tests, we used Newey-West (1987) robust standard errors.

The second test is another test of optimality under MSE loss, this time based on the forecast errors,  $e_{t+1,t} \equiv Y_{t+1} - \hat{Y}_{t+1,t}$ ,

$$e_{t+1,t} = \beta_0 + \beta_1 \hat{Y}_{t+1,t} + \beta_2 e_{t,t-1} + u_{t+1}$$

In this case, the null hypothesis is that all parameters are equal to zero. The third test is our new test, based on the indicator variable  $I_{t+1,t} \equiv \mathbf{1}(Y_{t+1} \leq \hat{Y}_{t+1,t})$ . This is used to test optimality under some unknown homogeneous loss function, and is based on the regression:

$$I_{t+1,t} = \beta_0 + \beta_1 \hat{Y}_{t+1,t} + \beta_2 I_{t,t-1} + u_{t+1}$$

As shown in Proposition 2, the indicator variable  $I_{t+1,t}$  should be independent of all  $Z_t \in \mathcal{F}_t$ , although it has unknown mean. Thus the null hypothesis for this test is that  $\beta_1 = \beta_2 = 0$ .

The final two tests are based on spline-based estimation of the unknown loss function. The first test assumes  $L(y, \hat{y}) = L(e)$  and approximates  $L$  with a quadratic spline using three nodes, so there are three free parameters. The nodes used in the simulation are:  $\{\widehat{Median}[e_t | e_t < 0], 0, \widehat{Median}[e_t | e_t > 0]\}$ , where  $\widehat{Median}$  is the sample median. The instruments used in the moment conditions are a constant, the current value of the forecast, and the lagged values of the forecast error, the predicted variable and the derivative of the loss function,  $\lambda$  (yielding five moments conditions in total).

The second spline-based test allows  $L$  to depend on both  $y$  and  $\hat{y}$ , and approximates  $L$  with a quadratic spline using three nodes and the specification given in equation (16), implying six free parameters. The nodes used are the same as for the simple spline-based test. The instruments used in the moment conditions are a constant, the current value of the forecast, and two lags each of the forecast error, the predicted variable and the derivative of the loss function,  $\lambda$  (yielding eight moments in total). All tests are conducted at the 0.05 level, and the results are presented in Table 2. We generated 3000 replications of each simulation design.

Under MSE loss, we see that all five tests have generally satisfactory size properties. The rejection frequencies are slightly high for  $T = 100$ , though this is commonly observed in tests based on robust standard errors. The rejection frequencies improve as the sample size increases, though the size of the test based on the flexible spline model (“spline 2”) remains slightly high even for  $T = 1000$ .

Turning to the relative power of the tests, we find that, under MSE loss, the indicator-based test has approximately the same power as the Mincer-Zarnowitz (MZ) test. This happens even though the indicator-based test is designed to detect sub-optimality under any homogeneous loss function and may reflect, as noted in the introduction, that the MZ regression only tests a weak implication of forecast optimality. This interpretation is consistent with the finding that the test based on forecast errors (“ $e$  test”) has better power than both the MZ test and the indicator-based test since it uses additional instruments such as the lagged forecast error that can identify serial correlation in the forecast errors. This effect is also likely to explain the good power of the two spline tests which utilize even more conditioning variables as instruments.

Under asymmetric loss the results are quite different. Unsurprisingly, the MZ and  $e$  tests reject the null hypothesis of optimality far more frequently than the nominal size. Even with just 100 observations, the MZ and  $e$  test reject the null hypothesis over 60% of the time. The indicator-based test and the “spline 1” test both have good size and power properties. The more flexible spline-based test has poor size properties for  $T = 100$  and 250 (the empirical rejection frequencies are 0.45 and 0.35, compared to the nominal size of 0.05), and even for  $T = 1000$  the size is still large (0.17).

Overall, these simulation results indicate that if the true loss function is MSE, the MZ test and “ $e$  test” have reasonable size and power (with the latter test being more powerful than the former) in finite samples. Both tests are slightly oversized when  $T \leq 250$ . The indicator-based test and the simple spline test have reasonable properties in finite samples for both MSE and non-MSE loss. The flexible spline test appears to require large samples,  $T \geq 1000$ , before the size of this test is close to its nominal value, and thus rejections obtained using this test must be interpreted with caution.

With regard to our empirical application, the simulation results lend support to our interpretation of the reported findings. The reasonable finite sample size of the regression and quantile-based tests support our rejection of the joint hypothesis of forecast optimality under conditional

mean-variance dynamics (assumption D2) and loss that is homogenous in the forecast error (L2). Similarly, our rejection of the simple error-based spline model and our failure to reject the more general spline model that allows loss to depend on both the forecast error and the level of the outcome variable is not overturned by the tendency of tests based on the latter to be over-sized.

## 5 Conclusion

Motivated by the surprising rejection of the optimality of the Federal Reserve’s internal forecasts (so-called “Greenbook” forecasts) of real GDP growth under mean squared error loss, this paper proposed new tests of forecast optimality that are applicable when the forecaster’s loss function is unknown. Our first set of tests apply when the loss function is homogeneous in the forecast error, and the variable of interest has dynamics in the conditional mean and variance but constant higher-order moments. This restriction on the loss function is quite weak, while the necessary restrictions on the data generating process can be easily tested. In particular, we propose a new set of quantile tests that do not require knowledge of the conditional mean and variance of the predicted variable, both of which depend on the forecaster’s (unobserved) information set.

Our second set of tests are based on flexible parametric approximations to the unknown loss function. We found significant evidence against the optimality of the Federal Reserve forecasts when the loss function was assumed to be a function solely of the forecast error, but when the loss function was allowed to depend on the forecast and the realization separately we found no evidence against optimality. Our estimates of the loss function suggest that the Fed issues ‘conservative’ estimates of economic growth: under-estimates are penalized less heavily than over-estimates. This conservatism appears to be particularly important when economic growth is moderate or low. This is consistent with the Federal Reserve viewing over-predictions of economic growth as being not just more costly than under-predictions, but also disproportionately more costly in times of low growth, perhaps because such forecasts may incorrectly signal a recovery from a recession and could result in an overly tight monetary policy at a critical point in time.

## Appendix

**Proof of Proposition 1.** Part 1: This result was shown by Granger (1969) and Christoffersen and Diebold (1997). The proof is similar to our proof of Proposition 2 below.

Part 2: Given the representation result in Part 1, we know that the forecast error associated with the optimal forecast is  $e_{t+h,t}^* = Y_{t+h} - \hat{Y}_{t+h,t}^* = \varepsilon_{t+h} - \alpha_h^*$ , where  $\alpha_h^*$  solves  $\min_{\alpha} \int L(\varepsilon_{t+h} - \alpha) dF_{\varepsilon,h}$ . Since  $\alpha_h^*$  is constant for fixed  $h$ , we thus have  $e_{t+h,t}^*$  is independent of all  $Z_t \in \mathcal{F}_t$ .

Part 3: Consider  $h > 0$  and  $j > 0$ . Let

$$\begin{aligned} Y_{t+h+j} &= E_t [Y_{t+h+j}] + \eta_{t+h+j}, \quad \eta_{t+h+j} | \mathcal{F}_t \sim F_{\varepsilon,h+j} (0, \sigma_{\varepsilon,h+j}^2) \\ Y_{t+h+j} &= E_{t+j} [Y_{t+h+j}] + \varepsilon_{t+h+j}, \quad \varepsilon_{t+h+j} | \mathcal{F}_t \sim F_{\varepsilon,h} (0, \sigma_{\varepsilon,h}^2). \end{aligned}$$

Let  $\sigma_{\varepsilon,h}^2 < \infty$  and further assume that  $\sigma_{\varepsilon,h+j}^2 < \infty$ . Using again that  $\hat{Y}_{t+h,t}^* = E_t [Y_{t+h}] + \alpha_h^*$ , so  $e_{t+h+j,t}^* = \eta_{t+h+j} - \alpha_{h+j}^*$ , and  $e_{t+h+j,t+j}^* = \varepsilon_{t+h+j} - \alpha_{h+j}^*$ , where  $\alpha_h^*$  and  $\alpha_{h+j}^*$  are constants. Thus  $V_t [e_{t+h+j,t}^*] = V_t [\eta_{t+h+j}] = \sigma_{\varepsilon,h+j}^2$ , and  $V_t [e_{t+h+j,t+j}^*] = \sigma_{\varepsilon,h}^2$ , where these moments are independent of  $t$  by assumption D1. Note also that  $V [e_{t+h+j,t}^*] = E [E_t [\eta_{t+h+j}^2]] = \sigma_{\varepsilon,h+j}^2$ , and similarly  $V [e_{t+h+j,t+j}^*] = \sigma_{\varepsilon,h}^2$ . Now we seek to show that  $\sigma_{\varepsilon,h+j}^2 \geq \sigma_{\varepsilon,h}^2$ .

$$\begin{aligned} V [e_{t+h+j,t}^*] &= V_t [Y_{t+h+j} - E_t [Y_{t+h+j}]] \\ &= V_t [\varepsilon_{t+h+j} + (E_{t+j} [Y_{t+h+j}] - E_t [Y_{t+h+j}])] \\ &= \sigma_{\varepsilon,h+j}^2 + V_t [E_{t+j} [Y_{t+h+j}]] + 2Cov_t [\varepsilon_{t+h+j}, E_{t+j} [Y_{t+h+j}] - E_t [Y_{t+h+j}]] \\ &\geq \sigma_{\varepsilon,h+j}^2 \\ &= V [e_{t+h,t}^*]. \end{aligned}$$

The first equality follows from the equality of the conditional and unconditional variance of the forecast error under D1; the third equality follows from the fact that  $E_t [Y_{t+h+j}]$  is constant given  $\mathcal{F}_t$ ; the weak inequality follows from the non-negativity of  $V_t [E_{t+j} [Y_{t+h+j}]]$  and  $E_{t+j} [\varepsilon_{t+h+j} \cdot \phi(Z_{t+j})] = 0$ ; the final equality follows from the fact that  $F_{\varepsilon,h}$  does not change with  $t$ . The cases where  $h = 0$  and/or  $j = 0$  are trivial. Thus  $V [e_{t+h+j,t}^*] \geq V [e_{t+h,t}^*] \forall h, j \geq 0$  if  $V [e_{t+h+j,t}^*] < \infty$ . If  $\sigma_{\varepsilon,h}^2 < \infty$  but  $\sigma_{\varepsilon,h+j}^2$  is infinite then the proposition holds trivially. ■

**Proof of Proposition 2.** Part 1: By homogeneity,

$$\begin{aligned}
\hat{Y}_{t+h,t}^* &\equiv \arg \min_{\hat{y}} \int L(y - \hat{y}) dF_{t+h,t}(y) \\
&= \arg \min_{\hat{y}} \int \left[ g \left( \frac{1}{\sigma_{t+h,t}} \right) \right]^{-1} L \left( \frac{1}{\sigma_{t+h,t}} (y - \hat{y}) \right) dF_{t+h,t}(y) \\
&= \arg \min_{\hat{y}} \int L \left( \frac{1}{\sigma_{t+h,t}} (y - \hat{y}) \right) dF_{t+h,t}(y) \\
&= \arg \min_{\hat{y}} \int L \left( \frac{1}{\sigma_{t+h,t}} (\mu_{t+h,t} + \sigma_{t+h,t} \eta_{t+h} - \hat{y}) \right) dF_{\eta,h}(\eta)
\end{aligned}$$

Let us represent a forecast as  $\hat{Y}_{t+h,t} = \mu_{t+h,t} + \sigma_{t+h,t} \cdot \hat{\gamma}_{t+h,t}$ , so

$$\begin{aligned}
\hat{Y}_{t+h,t}^* &= \mu_{t+h,t} + \sigma_{t+h,t} \cdot \arg \min_{\hat{\gamma}} \int L \left( \frac{1}{\sigma_{t+h,t}} (\mu_{t+h,t} + \sigma_{t+h,t} \eta_{t+h} - \mu_{t+h,t} - \sigma_{t+h,t} \hat{\gamma}) \right) dF_{\eta,h}(\eta) \\
&= \mu_{t+h,t} + \sigma_{t+h,t} \cdot \arg \min_{\hat{\gamma}} \int L(\eta_{t+h} - \hat{\gamma}) dF_{\eta,h}(\eta) \\
&= \mu_{t+h,t} + \sigma_{t+h,t} \cdot \gamma_h^*,
\end{aligned}$$

where the last line follows from the fact that  $F_{\eta,h}$  is time-invariant under assumption D2.

Part 2: This result follows from noting that  $d_{t+h,t}^* = \eta_{t+h} - \gamma_h^*$ , where  $\gamma_h^*$  is a constant for fixed  $h$  and, by assumption D2,  $\eta_{t+h}$  is independent of all elements in  $\mathcal{F}_t$  and has unit variance. ■

**Proof of Corollary 1.** Under assumptions D2' and L2 we have from Proposition 2 that  $\hat{Y}_{t+h,t}^* = \sigma_{t+h,t}(\beta + \gamma_h^*)$ . Thus  $\hat{d}_{t+h,t}^* \equiv e_{t+h,t}^* / \hat{Y}_{t+h,t}^* = (\eta_{t+h} - \gamma_h^*) / (\beta + \gamma_h^*)$ , i.e., an affine transformation of  $\eta_{t+h}$ . The result follows by noting that  $\eta_{t+h}$  is independent of all  $Z_t \in \mathcal{F}_t$ . ■

**Proof of Proposition 3.** Part 1: Under assumptions D1 and L1, or D2 and L2, we know from Propositions 1 and 2 that  $Y_{t+h,t}^* = \mu_{t+h,t} + \sigma_{t+h,t} \cdot \gamma_h^*$ , with  $\sigma_{t+h,t}$  constant under assumption D1.  $\gamma_h^*$  depends only upon the loss function and  $F_{\eta,h}$ . Now notice that  $F_{t+h,t}(\hat{Y}_{t+h,t}^*) \equiv \Pr[Y_{t+h} \leq \hat{Y}_{t+h,t}^* | \mathcal{F}_t] = \Pr[\mu_{t+h,t} + \sigma_{t+h,t} \eta_{t+h} \leq \mu_{t+h,t} + \sigma_{t+h,t} \cdot \gamma_h^* | \mathcal{F}_t] = \Pr[\eta_{t+h} \leq \gamma_h^* | \mathcal{F}_t]$ , which is constant under assumption D2. If we denote  $q_h^* \equiv \Pr[\eta_{t+h} \leq \gamma_h^* | \mathcal{F}_t]$ , then  $\hat{Y}_{t+h,t}^*$  is the  $q_h^*$  conditional quantile of  $Y_{t+h} | \mathcal{F}_t \forall t$ . Note that  $q_h^*$  is only a function of the loss function and  $F_{\eta,h}$ .

Part 2: Since  $I_{t+h,t}^*$  is a binary random variable and  $\Pr[I_{t+h,t}^* = 1 | \mathcal{F}_t] = \Pr[Y_{t+h} \leq \hat{Y}_{t+h,t}^* | \mathcal{F}_t] = q_h^* \forall t$ , we thus have that  $I_{t+h,t}^*$  is independent of all  $Z_t \in \mathcal{F}_t$ . ■

**Proof of Proposition 4.** Part 1: Following the steps in the proof of Proposition 2, we find

$$\begin{aligned}\hat{Y}_{t+h,t}^* &= \mu_{t+h,t} + \arg \min_{\hat{\alpha}} \int L(\sigma_{t+h,t}(\eta_{t+h} - \hat{\alpha})) dF_{\eta,h}(\eta) \\ &\equiv \mu_{t+h,t} + \alpha_{t+h,t}^*\end{aligned}$$

where  $\alpha_{t+h,t}^*$  depends on the loss function,  $L$ , the conditional standard deviation of  $Y_{t+h}$ ,  $\sigma_{t+h,t}$ , and the distribution of the innovation,  $F_{\eta,h}$ .

2. Similarly,

$$\begin{aligned}\hat{Y}_{t+h,t}^* &= \mu_{t+h,t} + \sigma_{t+h,t} \cdot \arg \min_{\hat{\gamma}} \int L(\eta_{t+h} - \hat{\gamma}) dF_{\eta,t+h,t}(\eta) \\ &\equiv \mu_{t+h,t} + \sigma_{t+h,t} \cdot \gamma_{t+h,t}^*\end{aligned}$$

Hence  $\gamma_{t+h,t}^*$  will be a function of the loss function and  $F_{\eta,t+h,t}$ , the latter depending on time-varying properties of the conditional distribution of  $Y_{t+h}|\mathcal{F}_t$  beyond the conditional mean and variance. ■

## References

- [1] Andersen, T.G., T. Bollerslev, F.X. Diebold, and P. Labys, 2001, The Distribution of Realized Exchange Rate Volatility, *Journal of the American Statistical Association*, 96, 42-55.
- [2] Andersen, T.G., T. Bollerslev, F.X. Diebold and P. Labys, 2003, Modeling and Forecasting Realized Volatility, *Econometrica*, 71, 579-625.
- [3] Bollerslev, Tim, 1986, Generalized Autoregressive Conditional Heteroskedasticity, *Journal of Econometrics*, 31, 307-327.
- [4] Campbell, S.D. and F.X. Diebold, 2005, Weather Forecasting for Weather Derivatives, *Journal of the American Statistical Association*, 100, 6-16.
- [5] Capistran, C., 2006, Bias in Federal Reserve Inflation Forecasts: Is the Federal Reserve Irrational or just Cautious? Manuscript, Banco de Mexico.
- [6] Christoffersen, P.F., 1998, Evaluating Interval Forecasts, *International Economic Review*, 39, 841-862.
- [7] Christoffersen, P.F. and F.X. Diebold, 1997, Optimal prediction under asymmetric loss. *Econometric Theory* 13, 808-817.
- [8] Clements, M.P. and D.F. Hendry, 1998, *Forecasting Economic Time Series*, Cambridge University Press.
- [9] Clements, M.P. and D.F. Hendry, 2006, Forecasting with Breaks in Data Processes, in C.W.J. Granger, G. Elliott and A. Timmermann (eds.) *Handbook of Economic Forecasting*, 605-657, Amsterdam, North-Holland.

- [10] Croushore, Dean and Tom Stark, 2001, A Real-Time Data Set for Macroeconomists, *Journal of Econometrics*, 105, 111-130.
- [11] Diebold, F.X. and J. Lopez, 1996, Forecast Evaluation and Combination, in G.S. Maddala and C.R. Rao (eds.), *Handbook of Statistics*, 241-268, Amsterdam, North-Holland.
- [12] Diebold, F.X. and G.D. Rudebusch, 1991, Forecasting Output with the Composite Leading Index: A Real-Time Analysis, *Journal of the American Statistical Association*, 86, 603-610.
- [13] Elliott, G., I. Komunjer, and A. Timmermann, 2005, Estimation and Testing of Forecast Rationality under Flexible Loss, *Review of Economic Studies*, 72, 1107-1125.
- [14] Engle, R.F., 1982, Autoregressive Conditional Heteroskedasticity With Estimates of the Variance of U.K. Inflation, *Econometrica*, 50, 987-1008.
- [15] Engle, R.F., D.M. Lilien and R.P. Robins, 1987, Estimating Time Varying Risk Premia in the Term Structure: The ARCH-M Model, *Econometrica*, 55, 391-407.
- [16] Garratt, A., K. Lee, M. H. Pesaran, and Y. Shin, 2003, Forecast Uncertainties in Macroeconomic Modelling: An Application to the UK Economy. *Journal of the American Statistical Association*, 98, 464, 829-838.
- [17] Granger, C.W.J., 1969, Prediction with a Generalized Cost Function. *OR* 20, 199-207.
- [18] Granger, C.W.J. and P. Newbold, 1986, *Forecasting Economic Time Series*, 2nd Edition. Academic Press.
- [19] Hall, A.R., 2005, *Generalized Method of Moments*, Oxford University Press, Oxford.
- [20] Hong, Y., 1999, Hypothesis Testing in Time Series via the Empirical Characteristic Function: A Generalized Spectral Density Approach, *Journal of the American Statistical Association*, 94, 1201-1220.
- [21] Jarque, C.M. and A.K. Bera, 1987, A Test for Normality of Observations and Regression Residuals, *International Statistical Review*, 55, 163-172.
- [22] Kim, C-J. and C.R. Nelson, 1999, Has the U.S. Economy Become More Stable? A Bayesian Approach Based on a Markov-Switching Model of the Business Cycle, *Review of Economics and Statistics*, 81, 608-616.
- [23] Koenker, R., and Bassett, G., 1978, Regression Quantiles, *Econometrica*, 46, 33-50.
- [24] McConnel, M.M. and G. Perez-Quiros, 2000, Output Fluctuations in the United States: What Has Changed Since the Early 1980s?, *American Economic Review*, 90, 1464-1476.
- [25] Mincer, J. and V. Zarnowitz, 1969, The Evaluation of Economic Forecasts. In J. Mincer, ed., *Economic Forecasts and Expectations*. National Bureau of Economic Research, New York.
- [26] Newey, W.K., and D. McFadden, 1994, Large Sample Estimation and Hypothesis Testing, in R.F. Engle and D. McFadden, (eds.) *Handbook of Econometrics*, Vol. IV, 2111-2245, North-Holland, Amsterdam.

- [27] Newey, W.K. and K.D. West, 1987, A Simple Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix, *Econometrica*, 55, 703-708.
- [28] Nobay, A.R. and D.A. Peel, 2003, Optimal Discretionary Monetary Policy in a Model of Asymmetric Central Bank Preferences, *Economic Journal*, 113, 657-665.
- [29] Patton, A. J., and A. Timmermann, 2004, Properties of Optimal Forecasts under Asymmetric Loss and Nonlinearity, *Journal of Econometrics*, forthcoming.
- [30] Romer, C.D. and D.H. Romer, 2000, Federal Reserve Information and the Behavior of Interest Rates, *American Economic Review*, 90, 429-457.
- [31] Stock, J.H. and M.W. Watson, 2002, Has the Business Cycle Changed and Why?, in M. Gertler and K. Rogoff, eds., *National Bureau of Economic Research, Macroeconomics Annual 2002*, MIT Press, U.S.A.
- [32] Taylor, S.J., 1982, Financial Returns Modelled by the Product of Two Stochastic Processes: A Study of Daily Sugar Prices, 1961-79, in O.D. Anderson, ed., *Time Series Analysis: Theory and Practice 1*, North-Holland, Amsterdam, 203-226.
- [33] Varian, H. R., 1974, A Bayesian Approach to Real Estate Assessment. pp. 195-208 in S.E. Fienberg and A. Zellner (eds.) *Studies in Bayesian Econometrics and Statistics in Honor of Leonard J. Savage*, Amsterdam: North Holland.
- [34] Weiss, A.A., 1996, Estimating Time Series Models Using the Relevant Cost Function. *Journal of Applied Econometrics* 11, 539-560.
- [35] West, K.D., H.J. Edison and D. Cho, 1993, A Utility based Comparison of Some Models of Exchange Rate Volatility. *Journal of International Economics* 35, 23-46.
- [36] Zellner, A., 1986, Bayesian Estimation and Prediction Using Asymmetric Loss Functions. *Journal of the American Statistical Association* 81, 446-451.

**Table 1**  
**Summary statistics**

|                    | Realization | Forecast | Forecast error |
|--------------------|-------------|----------|----------------|
| Mean               | 3.095       | 2.594    | 0.501          |
| Standard deviation | 3.537       | 2.326    | 3.199          |
| Skewness           | -0.219      | -0.414   | 0.081          |
| Kurtosis           | 4.456       | 4.027    | 3.869          |
| Minimum            | -8.16       | -4.70    | -7.01          |
| Maximum            | 15.46       | 8.50     | 10.86          |
| Autocorrelation 1  | 0.257*      | 0.728*   | 0.010          |
| Autocorrelation 2  | 0.198*      | 0.552*   | 0.096          |
| Autocorrelation 3  | 0.075       | 0.446*   | -0.021         |
| Autocorrelation 4  | 0.045       | 0.341*   | 0.067          |
| JB statistic       | 12.044*     | 9.058*   | 3.530          |
| JB p-value         | 0.002       | 0.011    | 0.171          |

Notes: This table presents basic summary statistics on the realized real GDP growth, the Federal Reserve’s “Greenbook” forecasts of GDP growth, and the corresponding forecast errors, over the period 1968Q4 to 1999Q4. The rows labeled “Autocorrelation  $j$ ” report the  $j^{th}$ -order autocorrelation of the series. Autocorrelations that are significant at the 0.05 level are marked with an asterisk. The last two rows present the Jarque-Bera (1987) test statistic and p-value, testing the null hypothesis that the forecast errors are Normally distributed. A p-value of less than 0.05 indicates a rejection of the null hypothesis.

**Table 2**  
**Finite-sample size and power of the tests**

|              | MSE loss |          |          |          |          | Asymmetric loss |          |          |          |          |
|--------------|----------|----------|----------|----------|----------|-----------------|----------|----------|----------|----------|
|              | MZ test  | $e$ test | $I$ test | spline 1 | spline 2 | MZ test         | $e$ test | $I$ test | spline 1 | spline 2 |
| $T = 100$    |          |          |          |          |          |                 |          |          |          |          |
| $\xi = 0$    | 0.11     | 0.12     | 0.09     | 0.09     | 0.11     | 0.64            | 0.62     | 0.11     | 0.12     | 0.45     |
| $\xi = 0.25$ | 0.23     | 0.39     | 0.24     | 0.70     | 0.70     | 0.65            | 0.73     | 0.23     | 0.54     | 0.69     |
| $\xi = 0.5$  | 0.76     | 0.92     | 0.76     | 0.91     | 0.91     | 0.89            | 0.96     | 0.75     | 0.95     | 0.91     |
| $\xi = 0.75$ | 0.98     | 1.00     | 0.97     | 0.98     | 0.98     | 0.98            | 1.00     | 0.97     | 1.00     | 0.98     |
| $\xi = 1$    | 1.00     | 1.00     | 1.00     | 0.99     | 0.99     | 1.00            | 1.00     | 1.00     | 1.00     | 0.99     |
| $T = 250$    |          |          |          |          |          |                 |          |          |          |          |
| $\xi = 0$    | 0.07     | 0.09     | 0.07     | 0.06     | 0.10     | 0.95            | 0.93     | 0.07     | 0.08     | 0.35     |
| $\xi = 0.25$ | 0.28     | 0.56     | 0.34     | 0.89     | 0.88     | 0.96            | 0.96     | 0.31     | 0.87     | 0.88     |
| $\xi = 0.5$  | 0.95     | 0.99     | 0.97     | 1.00     | 1.00     | 0.99            | 1.00     | 0.96     | 1.00     | 0.99     |
| $\xi = 0.75$ | 1.00     | 1.00     | 1.00     | 1.00     | 1.00     | 1.00            | 1.00     | 1.00     | 1.00     | 1.00     |
| $\xi = 1$    | 1.00     | 1.00     | 1.00     | 1.00     | 1.00     | 1.00            | 1.00     | 1.00     | 1.00     | 1.00     |
| $T = 1000$   |          |          |          |          |          |                 |          |          |          |          |
| $\xi = 0$    | 0.06     | 0.06     | 0.06     | 0.06     | 0.09     | 1.00            | 1.00     | 0.06     | 0.05     | 0.17     |
| $\xi = 0.25$ | 0.66     | 0.97     | 0.80     | 1.00     | 1.00     | 1.00            | 1.00     | 0.79     | 1.00     | 1.00     |
| $\xi = 0.5$  | 1.00     | 1.00     | 1.00     | 1.00     | 1.00     | 1.00            | 1.00     | 1.00     | 1.00     | 1.00     |
| $\xi = 0.75$ | 1.00     | 1.00     | 1.00     | 1.00     | 1.00     | 1.00            | 1.00     | 1.00     | 1.00     | 1.00     |
| $\xi = 1$    | 1.00     | 1.00     | 1.00     | 1.00     | 1.00     | 1.00            | 1.00     | 1.00     | 1.00     | 1.00     |

Notes: This table reports the results of a Monte Carlo study of the finite-sample properties of the tests considered in this paper. The data generating process is a conditionally Gaussian AR(1)-GARCH(1,1) process, and the sample sizes considered are  $T = 100, 250, 1000$ . The forecasts tested are  $\hat{Y}_{t+1,t} = \hat{Y}_{t+1,1}^* + \xi \epsilon_{t+1}$ , where  $\hat{Y}_{t+1,1}^*$  is the true optimal forecast and  $\epsilon_{t+1}$  is a standard Normal noise term, independent of all other variables. When  $\xi = 0$  the forecast is truly optimal, and as  $\xi$  grows the forecast becomes increasingly contaminated with noise. Further details on the tests and simulation design are presented in Section 4. The number of replications was 3000.

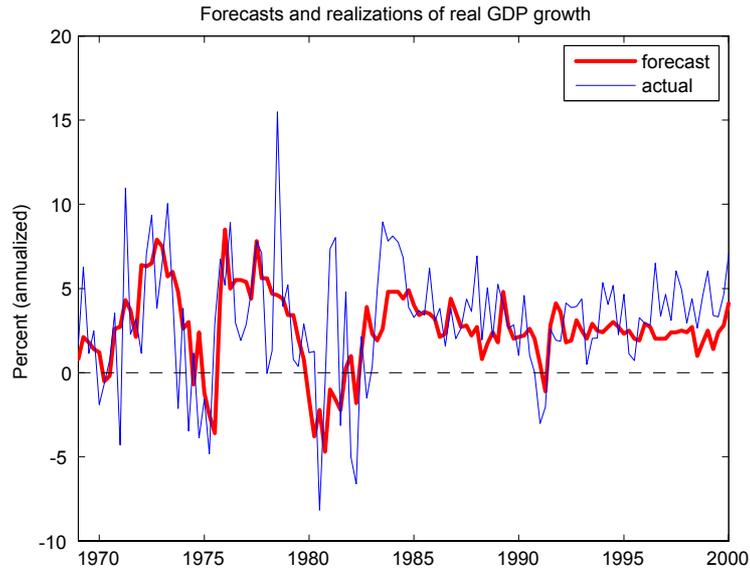


Figure 1: *Real GDP growth (annualized), and the Federal Reserve’s “Greenbook” forecasts of real GDP growth, over the period 1968Q<sub>4</sub> to 1999Q<sub>4</sub>.*

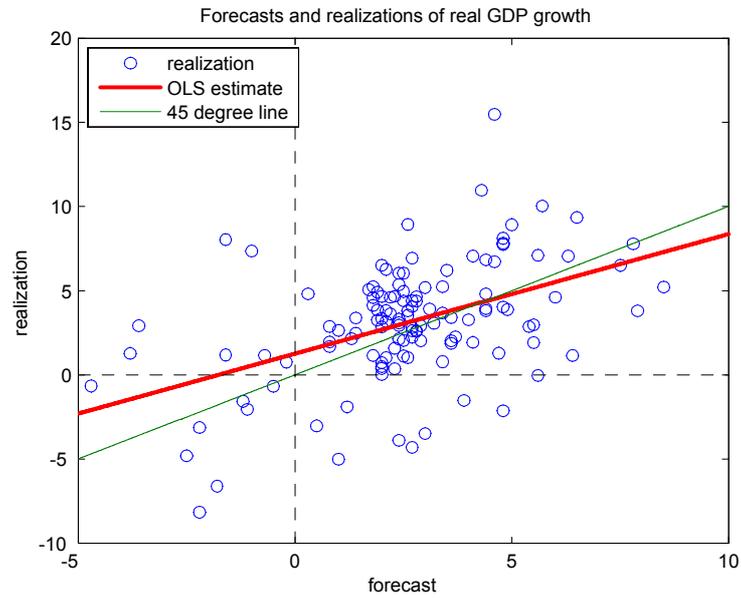


Figure 2: *Forecast and realized real GDP growth (annualized) over the period 1968Q<sub>4</sub> to 1999Q<sub>4</sub>, with OLS fitted line.*

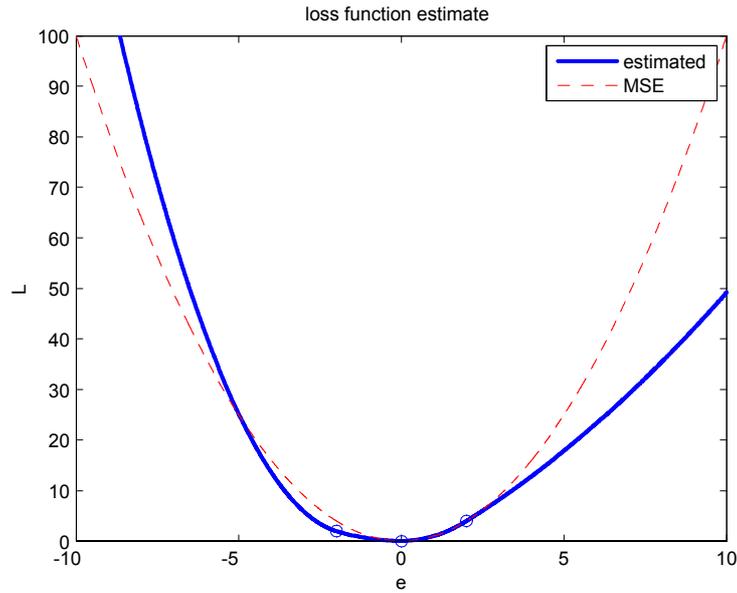


Figure 3: *The estimated loss function of the Federal Reserve for real GDP growth forecasts, based on a quadratic spline with nodes  $[-2, 0, 2]$ , imposing that the loss function is solely a function of the forecast error.*

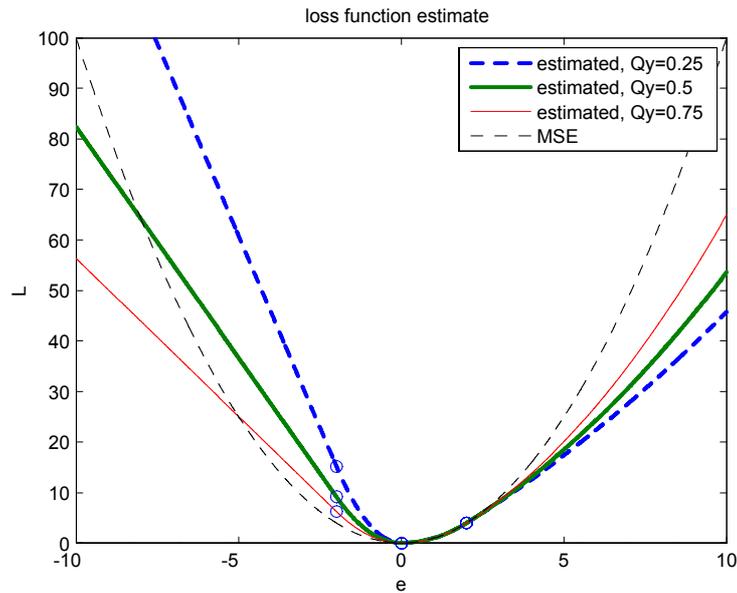


Figure 4: *The estimated loss function of the Federal Reserve for real GDP growth forecasts, based on quadratic splines with nodes  $[-2, 0, 2]$ . This model allows the level of GDP growth to also affect the loss function; the estimated loss function is evaluated for GDP growth equal to its 0.25, 0.5 and 0.75 quantiles.*