



The World's Largest Open Access Agricultural & Applied Economics Digital Library

This document is discoverable and free to researchers across the globe due to the work of AgEcon Search.

Help ensure our sustainability.

Give to AgEcon Search

AgEcon Search

<http://ageconsearch.umn.edu>

aesearch@umn.edu

*Papers downloaded from **AgEcon Search** may be used for non-commercial purposes and personal study only. No other use, including posting to another Internet site, is permitted without permission from the copyright owner (not AgEcon Search), or as allowed under the provisions of Fair Use, U.S. Copyright Act, Title 17 U.S.C.*

No endorsement of AgEcon Search or its fundraising activities by the author(s) of the following work or their employer(s) is intended or implied.

EVALUATING THE RATIONALITY OF FIXED-EVENT FORECASTS

Michael P. Clements

No.457

WARWICK ECONOMIC RESEARCH PAPERS



DEPARTMENT OF ECONOMICS

EVALUATING THE RATIONALITY OF FIXED-EVENT FORECASTS

Michael P. Clements
Department of Economics
University of Warwick
Coventry CV4 7AL
England

No.457

March 1996

This paper is circulated for discussion purposes only and its contents should be considered preliminary.

Evaluating the rationality of fixed-event forecasts

Michael P. Clements*
Department of Economics,
University of Warwick.

March 19, 1996

Abstract

A test of forecast rationality based on the weak efficiency of fixed-event forecasts has recently been proposed by Nordhaus (1987). This paper considers the scope for pooling fixed-event forecasts across ‘events’ to deliver more powerful tests of the weak-efficiency hypothesis. In an empirical illustration we demonstrate the usefulness of this approach when only a small number of forecasts are available. We also suggest an interpretation of the rejection of the null hypothesis of weak efficiency in favour of negative autocorrelation in series of revisions to fixed-event forecasts.

The relationship between weak efficiency and rationality when loss functions are asymmetric and prediction error variances are time-varying is also considered.

JEL classification: 132

Key words:

fixed-event forecasts, weak efficiency, forecast rationality, pooling.

1 Introduction

There is now a large literature on assessing the *ex post* rationality of forecasts in the tradition of the realization-forecast analysis of Theil (1958) (see Wallis, 1995 for a recent review). An early application of this approach was Ball (1962), and since then progress has been made in improving the econometric implementation of realization-forecast analysis and the testing of various forecast optimality conditions. The focus has typically been on what can be termed rolling-event forecasts, so that the properties of a series of forecasts of fixed length h are analysed. If we denote by $P_{t+h|t}$ a forecast made at period t about the value of the process in period $t + h$ (denoted A_{t+h}), and let $e_{t+h|t}$ denote the forecast error ($e_{t+h|t} \equiv A_{t+h} - P_{t+h|t}$), then we might run a regression of the form:

$$A_{t+h} = \alpha + \beta P_{t+h|t} + u_{t+h|t} \quad (1)$$

where the sample consists of $\{A_{t+h}, P_{t+h|t}\}$ for $t = T, \dots, T + H - h$, say, where T is the first forecast origin and we have realizations and forecasts of the process up to period $T + H$. Tests of unbiasedness and efficiency then relate to the estimated values of α and β , see, for example, Mincer and Zarnowitz (1969), Zarnowitz (1985), Holden and Peel (1990) and Kirchgässner

*I am grateful to David Hendry and Ken Wallis for helpful discussions and comments. Financial support under ESRC grant L116251015 is acknowledged.

(1993). When the forecast horizon h exceeds the sampling frequency forecasts overlap in the sense that they are made before the error in the previous forecast is known. In which case it is necessary to allow for a moving-average process for the u_{t+h} of order $h - 1$, so that consistent estimates of the covariance matrix of the coefficient estimates require the use of the Hansen and Hodrick (1980) correction (see, for example, Brown and Maital, 1981). Other problems arise when forecasts are pooled over individuals (see Zarnowitz, 1985) but in general the econometric implementation of tests of rationality based on fixed-event forecasts is reasonably well developed.

An alternative type of forecast that has received scant attention in the literature are fixed-event forecasts, where we consider the properties of forecasts of a given event (say, A_t at $t = \tau$) made at a variety of different times (variable h). Formally, the forecasts are denoted by $P_{\tau|\tau-h}$ where $h \in (1, \dots, \kappa)$ gives 1-step up to κ -step forecasts of the value of the process in period τ . Fixed-event forecasts were discussed by Nordhaus (1987), and in contrast to rolling-event forecasts, are a relatively unexploited resource for testing various notions of rationality. The motivation for developing a framework for analyzing fixed-event forecasts is that in many cases forecasts of the same event are made at a number of points in time. For example, many of the major economic forecasting agencies produce forecasts either every month or every quarter of the state of the economy two or three years ahead. Based on the analysis of Nordhaus (1987), we wish to investigate whether such forecasts are rational in the sense that new information is immediately and fully reflected in the latest forecasts. Since only a small number of forecasts are typically available of the value of a variable at a given future date, the framework is extended to allow the pooling of forecasts for a particular variable over a number of different target dates, and the econometric issues that arise are discussed. Finally, some new interpretations of the rejection of the weak efficiency version of rationality are offered.

In section 2 we briefly describe the notion of weak efficiency due to Nordhaus (1987), and suggest extensions to the econometric implementation of tests of forecast rationality based on fixed-event forecasts. Section 3 tackles one of the main obstacles to the more widespread use of fixed-event forecasts in assessing the rationality of economic forecasts, which is to do with the relatively short spans of data typically available (κ small). We investigate the possibility of pooling forecasts for a particular variable over different events (τ), and of using the generalized linear regression model. We discuss the extent to which this requires knowledge of the model used to generate the sample of forecasts.

Section 4 reports an empirical illustration of some of these ideas, and interprets the finding of negative autocorrelation in series of revisions as evidence of an absence of ‘significant news’ over the period. Section 5 shows that the test for weak efficiency is robust to asymmetric loss assuming a constant prediction error variance, but in general a rejection of weak efficiency could be consistent with either irrationality or optimally biased forecasts in the presence of time-varying prediction error variance. Section 6 concludes.

2 Analysis of fixed-event forecasts

Weak efficiency (due to Nordhaus, 1987) is the requirement that forecasts efficiently incorporate new information as it becomes available. This is testable from series of forecast revisions $v_{\tau|j}$, defined as:

$$v_{\tau|j} = P_{\tau|j-1} - P_{\tau|j},$$

which denotes (minus) the difference between the $(\tau - j)$ -step ahead forecast of the value of the process in the future period τ , made at j , and the $(\tau - j + 1)$ -step ahead forecast of the process at τ , made at $j - 1$. Thus it is the revision from period $j - 1$ to period j . Moreover, $P_{\tau|\tau} \equiv A_{\tau}$, and for a given τ we have a series of revisions $\{v_{\tau|j}\}$ where $j = \kappa + 1, \dots, \tau$. $v_{\tau|\tau}$ is the 1-step ahead forecast error, and $v_{\tau|\kappa+1}$ defines the first available forecast of period τ ($P_{\tau|\kappa}$) as having a length of $\tau - \kappa$. Then one implication of weak efficiency is that the revision at j is independent of all revisions up to $j - 1$:

$$E[v_{\tau|j} \mid v_{\tau|j-1}, \dots, v_{\tau|\kappa+1}] = 0, \quad j = \kappa + 2, \dots, \tau. \quad (2)$$

(2) implies that forecast revisions should be a white noise series, $P_{\tau|j-1} - P_{\tau|j} = -\varepsilon_j$, where $\varepsilon_j \sim \text{ID}(0, \sigma_{\varepsilon}^2)$, say. Thus, efficient forecasts follow a (possibly heteroscedastic) random walk¹ with drift, where the drift term $P_{\tau|\kappa}$ is the first available forecast:

$$P_{\tau|j} = P_{\tau|\kappa} + \sum_{s=\kappa+1}^j \varepsilon_s.$$

If we assume the ε_j are approximately normal, then an obvious test for weak efficiency is to calculate the first-order autocorrelation coefficient and test whether it differs significantly from zero (using standard t -tables, assuming revisions are a stationary $I(0)$ series). Equivalently, run the regression:

$$v_{\tau|j} = \alpha v_{\tau|j-1} + \zeta_{\tau,j}, \quad (3)$$

where the sample is over all the available j , and test the null that $\alpha = 0$ against $\alpha \neq 0$. In section 4 we discuss possible interpretations of rejecting the null in favour of $\alpha < 0$.

More generally, if the $v_{\tau|j}$ are found to have non-zero autocorrelations at any lag then this contradicts the white noise hypothesis required for rationality. Sufficiently long time series on revisions would allow more general forms of dependence to be investigated.

The ‘random walk’ hypothesis can be tested more directly than in (3). A general model that embeds (3) is:

$$P_{\tau|j-1} - P_{\tau|j} = \gamma_0 + \gamma_1 P_{\tau|j-1} + \gamma_2 P_{\tau|j-2} + \zeta_{\tau,j} \quad (4)$$

where we allow an intercept. Testing the null that $\alpha = 0$ in (3) is equivalent to testing that $\gamma_1 (= -\alpha) = 0$ and $\gamma_2 (= \alpha) = 0$ against $\gamma_1 + \gamma_2 = 0$ in (4), with $\gamma_0 = 0$ under both the null and alternative hypotheses. The rationality hypothesis is essentially that the martingale property, or the ‘law of iterated expectations’, applies to fixed-event forecasts. That is, $E_s[P_{\tau|j}] = P_{\tau|s}$, $s < j$ - my expectation today of what my expectation will be tomorrow of the outcome of the football match at the end of the week is simply the expectation I hold today of the outcome of the match. This can be tested within (4) if we set $\gamma_2 = 0$. The null is that $\gamma_1 = 0$ and the alternative is that $\gamma_1 \neq 0$ and we allow a non-zero intercept. If the null is rejected then I know today that the forecast I will make tomorrow will differ in a known direction from the forecast I hold today, contrary to the law of iterated expectations. Conducting the test within

¹ The ideas here parallel the random walk model of stock market prices, which states ‘that price changes in a stock (or index) cannot be predicted from earlier price changes in this stock (or index).’ Granger and Morgenstern (1970), p.96. Granger and Morgenstern (1970) attribute the first statement of the random walk model of stock prices to Bachelier in 1900, and reference some of the early contributions, including Samuelson (1965) (see their section 3.3). More recently the random walk model has become influential in the consumption literature, beginning with Hall (1978), and surveyed by Deaton (1992).

(3) has the advantage that both the regressand and regressor are $I(0)$ so that the t -statistic has a standard distribution, although critical values have been tabulated by Fuller (1976) for testing for a unit root. In certain circumstances testing the random walk hypothesis within (4) may be more natural if we harbour the suspicion that fixed-event forecasts may themselves be white noise (rather than this being an attribute of the series of revisions to fixed-event forecasts, see section 4).

The form of the maintained hypothesis in (3) requires that the relationship between adjacent revisions is constant and in particular does not depend upon $l = \tau - j$, the length of the forecast. It may be that 'forecast stickiness' is more pronounced for large l , particularly if it is expected that future events (or 'news') may suggest offsetting changes. With 'credibility costs' (loss of credibility from often changing forecasts) some degree of stickiness may be rational.

As for rolling-event forecasts, we can define a notion of strong efficiency which requires that revisions are an innovation against the forecaster's information set. Tests of strong efficiency can only ever refute the null since the information set of the forecaster, against which revisions should be an innovation, is clearly open-ended. Failure to reject against a specific variable does not rule out rejection for the next series tried. If forecasters efficiently utilise the latest available information, then we would expect to find $\gamma = 0$ in:

$$v_{\tau|j} = \gamma' Z_{j-s} + u_{\tau|j} \quad (5)$$

where Z_{j-s} is known at period $j - 1$ for $s \geq 1$. Rejection of $\gamma = 0$ is consistent with 'partial rationality', defined as efficient use of the agent's information, if that information is not complete (some elements of Z_{j-s} are not in the information set) and also with the inefficient use of complete information.

We end this section by briefly reviewing the literature related to fixed-event forecast evaluation.

Nordhaus (1987) finds evidence of positive correlations in revisions to forecasts of nuclear capacity, energy consumption, and real GNP growth for four major U.S forecasting services. A number of reasons as to why this might be expected to arise in general are given. These include: the bureaucratic nature of forecast agencies that need to achieve a consensus view, so that there is a degree of inertia in revising forecasts; forecasts are more credible and consumers happier if forecasts change only slowly; and 'a tilt toward wishing to hold on to old notions that are relatively familiar rather than to adjust to surprises.' (p. 673).

Mankiw and Shapiro (1986) test the revisions of the BEA's estimates of GNP for forecast efficiency. However, they regress revisions of the estimates on earlier revisions, where in terms of (3) j is fixed and the observations are generated by considering different τ . Thus their approach is in the traditional vein of evaluating fixed length forecasts. They find no evidence of positive correlation of revisions.

Swidler and Ketcher (1990) also consider rationality in terms of forecasts (rather than forecast revisions) by running regressions of the form:

$$A_{\tau} = \alpha_l + \beta_l P_{\tau|l} + v_{\tau|l}. \quad (6)$$

They implement (6) by running the regression across τ for a fixed l : this is the traditional approach. They obtain estimates of the α 's and β 's for a number of l , and argue that since information accrues with time, later forecasts should be more accurate than earlier ones, so that the R^2 's of regressions should be inversely related to the lead time l . They consider forecasts of real GNP growth and price inflation, and find that for both later forecasts appear to be more accurate.

Kirchgässner (1993) considers testing predictions with different time horizons for weak rationality. For example, he tests the null hypothesis $H_0 : \mu = 0, \Gamma = I$, in the equation:

$$e_{t+k|t} = \mu + \Gamma e_{t+k|t+j} + \epsilon_t$$

allowing for autocorrelation of up to $j - 1$ in the regression residuals (Kirchgässner, 1993 equation 7, my notation). Although the regressand and regressor are the errors in forecasts of the same target date ($t + k$) made at different origins (t and $t + j$) the analysis is within the rolling-event tradition since the sample is over t for fixed k and j .

Finally, Clements (1995) runs regressions of the form (3) on fixed-event forecasts of UK GDP growth and inflation, and fails to find positive autocorrelation of revisions.

3 Pooling fixed-event forecasts

Consider pooling the $v_{\omega|\iota}$ over ω in (3). Then we can write this as a system of equations:

$$\begin{bmatrix} \mathbf{v}_{\tau-(S-1)|\iota} \\ \vdots \\ \mathbf{v}_{\tau|\iota} \end{bmatrix} = \begin{bmatrix} \alpha_1 \mathbf{I}_J & \mathbf{0} & \cdots & \mathbf{0} \\ \mathbf{0} & \ddots & \mathbf{0} & \mathbf{0} \\ \vdots & & \ddots & \vdots \\ \mathbf{0} & \mathbf{0} & \cdots & \alpha_S \mathbf{I}_J \end{bmatrix} \begin{bmatrix} \mathbf{v}_{\tau-(S-1)|\iota-1} \\ \vdots \\ \mathbf{v}_{\tau|\iota-1} \end{bmatrix} + \begin{bmatrix} \zeta_1 \\ \vdots \\ \zeta_S \end{bmatrix} \quad (7)$$

where $\mathbf{v}_{\omega|\iota} = [v_{\omega|j-(J-1)} \ v_{\omega|j-(J-2)} \ \dots \ v_{\omega|j-1} \ v_{\omega|j}]'$, and $\zeta_i = [\zeta_{i1} \ \dots \ \zeta_{iJ}]'$. The formulation in (7) assumes that we have forecasts made at J different periods $\iota = (j - (J - 1), \dots, j)$ of the value of the process at S target dates $\omega = (\tau - (S - 1), \dots, \tau)$ (a total of $S \times J$ forecasts). Thus the shortest forecast horizon is $(\tau - (S - 1)) - j$, and the longest is $\tau - (j - (J - 1))$. The design is balanced in the sense that for each target data ω we have a forecast made at each origin ι .²

Forecast rationality in the sense of weak efficiency requires that $\alpha_i = 0, i \in (1, \dots, S)$. Separate regressions could be run for each value of ω , but for small J this may generate imprecise results, and pooling the observations in some way may yield efficiency gains. We could simply pool the observations over different ω by imposing $\alpha_i = \alpha, i \in (1, \dots, S)$ and then estimate (7) by OLS, and base a test of weak efficiency on the t -value of $H_0: \alpha = 0$. However this assumes that $\Omega \equiv E[\zeta \zeta']$ is diagonal ($\sigma_\zeta^2 \mathbf{I}_{S \times J}$), where $\zeta = [\zeta_1' \ \dots \ \zeta_S']'$. This is likely to be a poor assumption since revisions to forecasts of different target dates but made at the same time are likely to be highly correlated: news at period j that leads to a revision in the forecast of the value of the variable at period $\tau - 1$ is also likely to lead to a revision in the forecast of the value of the variable at period τ .

In order to estimate (7) by generalized least squares we need to be able to reduce the number of free parameters in Ω substantially. Below we derive the structure of Ω assuming that the forecasting model is known and that the forecasts being assessed are unadulterated model-based forecasts. Under the null hypothesis of interest ($\alpha_i = 0 \ \forall i$), the covariance matrix of errors in the forecast revision regressions (Ω from (7)) is simply the covariance matrix of forecast revisions.

² In practice this is unlikely, and the empirical example in section 4 is unbalanced. However, as indicated in section 7, the necessary adjustments are straightforward. The simpler assumption is made here for expositional convenience.

3.1 The covariance structure of forecast revisions for an ARMA process

Suppose the process has the ARMA(p, q) representation:

$$\phi(L) y_t = \theta(L) \epsilon_t \quad (8)$$

where $\epsilon_t \sim \text{IN}(0, \sigma_\epsilon^2)$, and $\phi(L) = 1 - \phi_1 L - \dots - \phi_p L^p$, $\theta(L) = 1 - \theta_1 L - \dots - \theta_q L^q$. The infinite MA representation is:

$$y_t = \epsilon_t + \sum_{j=1}^{\infty} \psi_j \epsilon_{t-j}. \quad (9)$$

From (9) we can obtain expressions for the fixed event forecasts:

$$P_{\tau|\tau-h} = \sum_{j=h}^{\infty} \psi_j \epsilon_{\tau-j}$$

and forecast revisions:

$$v_{\tau|\tau-h+1} \equiv P_{\tau|\tau-h} - P_{\tau|\tau-h+1} = -\psi_{h-1} \epsilon_{\tau-h+1}. \quad (10)$$

Consider now the autocovariance function for the forecast revisions. Let:

$$\gamma_{\tau_1, j; \tau_2, i} = \text{E} [v_{\tau_1|j} v_{\tau_2|i}]$$

then:

$$\gamma_{\tau_1, j; \tau_2, i} = \begin{cases} 0 & j \neq i \\ \sigma_\epsilon^2 \psi_{\tau_1-j} \psi_{\tau_2-j} & j = i \end{cases}$$

from which we can deduce that Ω is symmetric and that each of the S^2 blocks Ω_{ii} :

$$\Omega = \begin{bmatrix} \Omega_{11} & \cdots & \Omega_{1,S} \\ \vdots & & \vdots \\ \Omega_{S,1} & \cdots & \Omega_{S,S} \end{bmatrix} \quad (11)$$

is diagonal. The i^{th} diagonal element of the r, c block, Ω_{rc} is given by:

$$\sigma_\epsilon^2 \psi_{\tau+1-r-i} \psi_{\tau+1-c-i}, i = j - (J - 1), \dots, j.$$

For the special case of an AR(1) model ($p = 1, q = 0$ in (8)) then $\psi_j = \phi_1^j$, and so for example Ω_{rc} has diagonal elements given by:

$$\sigma_\epsilon^2 \phi^{2(\tau+1-i)-r-c}, i = j - (J - 1), \dots, j.$$

Consider the case of $S = 2$, so that there are only two target dates or fixed-events for which we have forecasts, $\omega = \tau - 1, \tau$, and consider the AR(1) model again. Then (7) has the form:

$$\begin{bmatrix} v_{\tau-1|j-(J-1)} \\ \vdots \\ v_{\tau-1|j} \\ \vdots \\ v_{\tau|j-(J-1)} \\ \vdots \\ v_{\tau|j} \end{bmatrix}_{2J} = \begin{bmatrix} \alpha_1 \mathbf{I}_J & \mathbf{0} \\ \mathbf{0} & \alpha_2 \mathbf{I}_J \end{bmatrix}_{2J \times 2J} \begin{bmatrix} v_{\tau-1|j-J} \\ \vdots \\ v_{\tau-1|j-1} \\ \vdots \\ v_{\tau|j-J} \\ \vdots \\ v_{\tau|j-1} \end{bmatrix} + \begin{bmatrix} \zeta_1 \\ \zeta_2 \end{bmatrix} \quad (12)$$

and (11) is given by $\Omega = \begin{bmatrix} \Omega_{11} & \Omega_{12} \\ \Omega_{12} & \Omega_{22} \end{bmatrix}$:

$$\Omega = \sigma_\epsilon^2 \begin{bmatrix} \begin{bmatrix} \phi^{2(l+J-1)} & & & \\ & \ddots & & \\ & & \phi^{2(l+1)} & 0 \\ & 0 & & \phi^{2l} \end{bmatrix} & \begin{bmatrix} \phi^{2(l+J-\frac{3}{2})} & & & \\ & \ddots & & 0 \\ & & \phi^{2(l-\frac{1}{2})} & \\ \phi^{2(l+J-2)} & 0 & & \phi^{2(l-\frac{1}{2})} \end{bmatrix} \\ \begin{bmatrix} \phi^{2(l+J-\frac{3}{2})} & & & \\ & & 0 & \\ & & & \ddots \\ 0 & & & \phi^{2(l-\frac{1}{2})} \end{bmatrix} & \begin{bmatrix} \phi^{2(l+J-2)} & & & \\ & \ddots & & 0 \\ & & \phi^{2l} & \\ 0 & & & \phi^{2(l-1)} \end{bmatrix} \end{bmatrix} \quad (13)$$

where $l = \tau - j$.

3.2 Informational requirements

The above model-based analysis appears to run counter to the spirit of realization-forecast analysis as discussed in section 1 which requires knowledge only of the predictions and realizations. In many cases we may wish to analyse forecasts which are not model-based, or only loosely model-based, where we do not know the model being used. However, it is evident from the analysis of the general ARMA process that the error-covariance structure of the pooled regressions is likely to have a fairly simple form, and we conjecture that it may be possible to approximate this structure sufficiently accurately for there to be benefits to pooling when the forecast model is not known.

Moreover, a standard result is that the OLS estimator of α will be consistent in the presence of untreated heteroscedasticity, but the properties of tests of the significance of α will depend on the assumptions made concerning the covariance structure. Using the above analysis we can impose sufficient restrictions on Ω to enable the free parameters to be estimated from the OLS residuals of the pooled regression.

3.3 Testing for weak efficiency by pooling

The above discussion suggests a number of possible approaches, which we now briefly summarise:

- A** The first method is to run OLS on the $S \times J$ observations in (7) assuming $\alpha_i = \alpha, \forall i$, and test $H_0 : \alpha = 0$. This implies Ω proportional to the identity matrix, and serves as the benchmark case.
- B** Next, we consider GLS whereby each block Ω_{rc} is diagonal with equal elements for that block, but varying over blocks. Thus, in terms of the example with $S = 2$ we obtain:

$$\Omega = \begin{bmatrix} \sigma_{11} \mathbf{I}_J & \sigma_{12} \mathbf{I}_J \\ \sigma_{12} \mathbf{I}_J & \sigma_{22} \mathbf{I}_J \end{bmatrix}.$$

If the diagonal elements are equal across blocks ($\sigma_{11} = \sigma_{12} = \sigma_{22}$) Ω is singular. Nordhaus (1987) footnote 5. p.669 argues that heteroscedasticity is unlikely to be a problem in the fixed-event model if large forecast revisions are associated with major external events (which are randomly spaced) and not with the distance from the target. The assumption of

equal elements within blocks corresponds to the assumption of homoscedasticity in Nordhaus' analysis of a single fixed event at a time. However, when the time frame is relatively short and forecasts are made at short intervals, say every month, then relatively few major events may occur, and the assumption of homoscedasticity may be more questionable. Moreover, monthly revisions may be greater in the months that quarterly national accounts data appear.

- C We could allow for this possibility by imposing the diagonal structure on each Ω_{rc} , to capture the effect of the same news item leading to revisions in the forecasts of all target dates, but also attempt to model the heteroscedasticity within blocks in a way which reflects the 'exponential decay' in (13) for the AR(1) process but that is likely to arise more generally for this type of model. The heteroscedasticity within blocks will arise whenever the information content of news is inversely related to the length of the lead time ($\omega - \iota$). These considerations suggest approximating Ω for $S = 2$ by:

$$\Omega = \begin{bmatrix} \begin{bmatrix} \rho_{11}^J & 0 \\ & \ddots \\ 0 & \rho_{11} \\ \rho_{12}^J & 0 \\ & \ddots \\ 0 & \rho_{12} \end{bmatrix} & \begin{bmatrix} \rho_{12}^J & 0 \\ & \ddots \\ 0 & \rho_{12} \\ \rho_{22}^J & 0 \\ & \ddots \\ 0 & \rho_{22} \end{bmatrix} \end{bmatrix}$$

where $|\rho_{ij}| < 1$, $i, j = 1, 2$.

4 Empirical illustration

As an illustration, we consider forecasts made by the National Institute of Economic and Social Research of GDP growth and consumer price inflation in the UK for 1993 and 1994. Forecasts of these variables are given every quarter in the 'Medium term projections' table of issues of the National Institute Economic Review. For 1993, we have forecasts made in 90:1 to 93:4 (exploiting the lag in data availability), and for 1994, forecasts made in 91:1 through to 94:4. Appendix 7 sketches the necessary changes to the 'balanced' case of section 3.

Series of forecasts of economic variables made every month are available, although it was felt that revisions to forecasts made every quarter would be more likely to embody useful information. Using forecasts made at higher frequencies (such as every month) is appealing as it generates many more observations to analyse, and this is the approach adopted in Clements (1995). A readily available source of economic forecasts made on a monthly basis by a large number of forecasting groups is the HM Treasury's monthly publication 'Forecasts for the UK economy: A comparison of independent forecasts'.³

Table 1 reports the results of testing the National Institute forecasts for weak efficiency by running regressions of the form of (3).

Without pooling, there are 14 forecast revisions in each case, and we obtain negative values of α , although only one of these has an absolute t -value in excess of 2. Nevertheless there is

³ Although this reports the latest forecasts of the main macro aggregates for a number of city and non-city forecasters, the horizon is only up to two years ahead, so that even selecting agencies that produced new forecasts nearly every month, there are generally fewer than 20 forecasts by an agency of a particular event.

Table 1 Testing the weak efficiency of National Institute forecasts of GDP growth and inflation.

	Estimated α coefficient	Standard Error	t -value
GDP growth			
OLS: 1993	-0.483	0.240	-2.015
OLS: 1994	-0.324	0.262	-1.236
OLS: Pooled	-0.441	0.171	-2.578
GLS ₁	-0.412	0.178	-2.315
GLS ₂	-0.411	0.178	-2.315
CPI inflation			
OLS: 1993	-0.261	0.219	-1.192
OLS: 1994	-0.171	0.267	-0.641
OLS: Pooled	-0.219	0.169	-1.296
GLS ₁	-0.284	0.142	-2.003
GLS ₂	-0.225	0.170	-1.326

no evidence of revisions being ‘smooth’ as would be suggested by positive correlations. The rows entitled ‘OLS: Pooled’ pool the observations assuming Ω is proportional to the identity matrix, as in [A] in section 3.3. We then calculate two variants of GLS. The first, GLS₁ is [B] in section 3.3. The precise form of Ω is given in the Appendix, but essentially we allow the diagonal elements of Ω_{11} and Ω_{22} (relating to the 1993 and 1994 target years, respectively) to differ, and these are estimated from the residuals of [A] at 0.119 and 0.052, for GDP, and 0.378 and 0.500, for inflation. We also allow Ω_{12} ($= \Omega_{21}$) to be diagonal, with elements given by 0.007 for GDP, and 0.391 for inflation. The rows entitled GLS₂ are as for GLS₁ but with $\Omega_{12} = \Omega_{21} = 0$.

There are a number of interesting features to comment upon. For GDP growth the correlation in the errors for the revisions to forecasts for 1993 and 1994 made at the same time is negligible, and so GLS₁ and GLS₂ are very similar. Relative to OLS on the pooled observations, GLS results in a larger standard error on α , but the t -value is still significant, suggesting α is negative. For inflation the correlation of the errors for different targets at the same forecast date is large, and neglecting this as in GLS₂ suggests we are unable to reject α being zero. For GLS₁ the t -value is borderline but larger than OLS because the coefficient is larger (more negative) and its associated standard error smaller. Given the small number of observations we did not attempt to allow for heteroscedasticity as in [C].

Nordhaus (1987) only found evidence of positive autocorrelation in revisions, which can readily be interpreted as slow adjustment to new information, for whatever reason, and did not consider the meaning of rejecting the null of weak efficiency because of negative autocorrelation. One possibility is the following. Suppose little relevant information becomes available between forecasts made at time j_{i-1} and j_i where $[j_{i-1} \ j_i)$ is a time interval so that:

$$P_{\tau|j} = P_{\tau}^{\{j\}} + \nu_j \quad (14)$$

and for $j \in [j_{i-1} \ j_i)$, $P_{\tau}^{\{j\}} = P_{\tau,i}$, where the period during which forecasts are made ($j = \kappa, \dots, \tau - 1$, say) consists of I (possibly unequal) time intervals. In the extreme case, let $I = 1$, and let $P_{\tau}^{\{j\}} = P_{\tau}$, say, then $P_{\tau|j} = P_{\tau} + \nu_j$, $j = \kappa, \dots, \tau - 1$, so that each forecast has in common P_{τ} and a specific white noise element ν_j . This might be a reasonable characterisation if the forecast agency takes an initial view at $j = \kappa$, denoted $P_{\tau|\kappa}$, and then sees no reason to

change it in subsequent forecasts, save for some white noise component. But then the series of forecast revisions are differenced white noise:

$$v_{\tau|j} = -\Delta\nu_j$$

for which the theoretical first-order autocorrelation coefficient is $-\frac{1}{2}$. In terms of testing the random walk hypothesis within (4) we would find $\gamma_1 = 1$, indicating that $P_{\tau|j}$ follows an AR(0), that is, is white noise. If the forecasts are white noise then the formulation in (3) induces a negative autocorrelation in revisions by ‘over-differencing’.

More generally, if we suppose that forecasts follow the AR(1) process:

$$P_{\tau|j} = \gamma_0 + \gamma_1^* P_{\tau|j-1} + \zeta_{\tau,j}$$

($\gamma_1^* = (1 - \gamma_1)$ in (4)) where $-1 \leq \gamma_1^* \leq 1$, then the value of α in (3) will be given by:

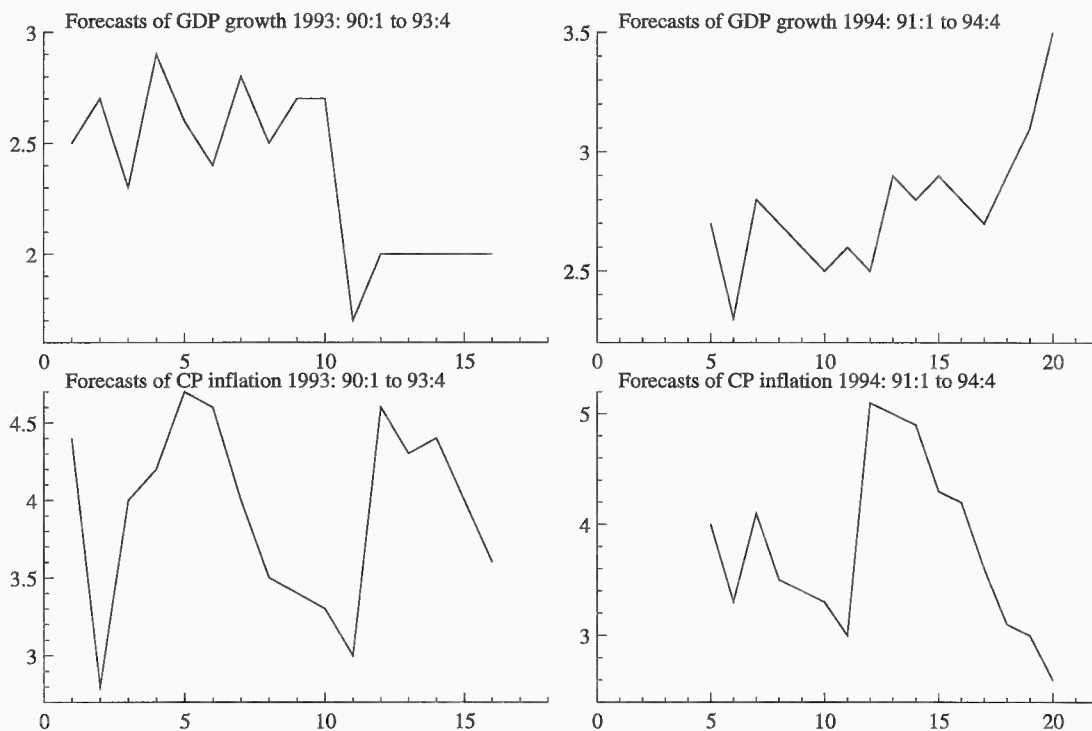
$$\alpha = \left(2\gamma_1^* - \gamma_1^{*2} - 1\right) (2 - \gamma_1^*)^{-1}$$

so that α takes on values in the range $[-\frac{4}{3}, 0]$. As described above, $\alpha = 0$ when $\gamma_1^* = 1$ and $\alpha = -\frac{1}{2}$ when $\gamma_1^* = 0$. The former is the random walk hypothesis and the latter may also be viewed as being consistent with weakly-efficient behaviour. For the National Institute forecasts of GDP growth and inflation we were never able to reject both the hypothesis that $\gamma_1^* = 0$ and that $\gamma_1^* = 1$, but this must be partly due to the small sample size (only 14 observations).

Finally, if $\sigma_{\nu_j}^2 = 0$ in (14) so that the same value is forecast at each point j (assuming $I = 1$) then all revisions are identically equal to zero. ν_j could be interpreted as a form of measurement error.

The forecasts are shown in figure 1. Notice the sharp rise in the forecasts of inflation for 1993 and 1994 made in the fourth quarter 1992 (observation 12) following the departure of sterling from the European Monetary System on the 16 September 1992. The last points in the graphs are either the forecasts made in 93:4 of the percentage rate of change of 1993 over 1992, or in 94:4 for the annual rate of growth in 1994 relative to 1993.

Figure 1: Fixed-event forecasts of GDP growth and consumer price inflation.



5 Asymmetric loss functions

Zellner (1986) has criticised the interpretation of rejecting unbiasedness as implying irrational behaviour. The argument is that the conditional mean predictor is not optimal for asymmetric loss functions. Thus the finding of biased predictions based on regressions such as (1) for rolling-event forecasts could either be evidence of irrationality or of optimal biased forecasts being used (see the comments in Zellner, 1986 on the evidence for biased inflation and real output growth forecasts in Zarnowitz, 1985).

Granger (1969) has shown that under certain circumstances the optimal predictor is the conditional expectation with a simple fixed adjustment, which depends only on the form of the loss function and the forecast error variance. The circumstances are that the process is Gaussian, and that the loss function is defined over forecast errors, rather than the realizations and forecasts more generally. More recently, Christoffersen and Diebold (1994) have extended Granger's theorem to allow for processes which are only *conditionally* Gaussian, which include processes which exhibit autoregressive conditional heteroscedasticity (ARCH)⁴ and its generalizations. They show that if the prediction error variance is time-varying, then the adjustment to the conditional expectation required to deliver the optimal predictor will no longer be constant, since it depends on the conditional error variance.

In this section we consider the implications of asymmetry in the loss function and time-varying prediction error variances for the interpretation of weak efficiency tests of fixed-event forecasts. We proceed by a specific example. Following Christoffersen and Diebold (1994),

⁴ ARCH was introduced by Engle (1982). There are a number of good surveys of the literature on ARCH and related models: Engle and Bollerslev (1987), Bollerslev, Chou and Kroner (1992), Bera and Higgins (1993), Shephard (1996). Engle (1995) contains an edited selection of some of the key papers.

who extend the analysis of Zellner (1986), we consider the ‘linex’ loss function of Varian (1975), defined by:

$$C(e_h) = b [\exp(ae_h) - ae_h - 1], \quad a \neq 0, b \geq 0$$

where $e_h = y_\tau - \hat{y}_{\tau|\tau-h}$ is the forecast error and A_τ and $P_{\tau|\tau-h}$ are replaced by y_τ and $\hat{y}_{\tau|\tau-h}$ for notational convenience. When $a > 0$ the loss function is approximately linear for $e_h < 0$, ‘over-predictions’, and exponential for $e_h > 0$, ‘under-predictions’. For small a , the loss function is approximately quadratic:

$$C(e) \simeq \frac{ba}{2} e_h^2$$

(from the first two terms of the Taylor series expansion). The advantage of the linex loss function is that the optimal predictor is easily obtained - in general analytic solutions are not available for arbitrary loss functions and it is necessary to resort to numerical methods (see Christoffersen and Diebold, 1994). The optimal predictor h -steps ahead solves:

$$\arg \min_{y_{\tau|\tau-h}^*} E_{\tau-h} \{b [\exp(ae_h) - ae_h - 1]\} \quad (15)$$

where $y_{\tau|\tau-h}^*$ is the optimal predictor, and is assumed to have the form:

$$y_{\tau|\tau-h}^* = \mu_{\tau|\tau-h} + \delta_{\tau|\tau-h}$$

where the process is conditionally Gaussian, $y_\tau | \mathcal{I}_{\tau-h} \sim N(\mu_{\tau|\tau-h}, \sigma_{\tau|\tau-h}^2)$, so that $\mu_{\tau|\tau-h}$ is the conditional mean predictor and $\sigma_{\tau|\tau-h}^2$ is the conditional variance. Substituting for $e_h = y_\tau - (\mu_{\tau|\tau-h} + \delta_{\tau|\tau-h})$ in (15) and using the result that:

$$E_{\tau-h} [\exp(ay_\tau)] = \exp \left[a\mu_{\tau|\tau-h} + \frac{a^2 \sigma_{\tau|\tau-h}^2}{2} \right]$$

gives:

$$\arg \min_{\delta_{\tau|\tau-h}} E_{\tau-h} \left\{ b \left[\exp \left(\frac{a^2 \sigma_{\tau|\tau-h}^2}{2} - a\delta_{\tau|\tau-h} \right) - a(y_\tau - \mu_{\tau|\tau-h} - \delta_{\tau|\tau-h}) - 1 \right] \right\}. \quad (16)$$

The first-order condition is satisfied by:

$$\delta_{\tau|\tau-h} = \frac{a}{2} \sigma_{\tau|\tau-h}^2$$

so that the optimal predictor becomes:

$$y_{\tau|\tau-h}^* = \mu_{\tau|\tau-h} + \frac{a}{2} \sigma_{\tau|\tau-h}^2. \quad (17)$$

To interpret (17), assume that $a \gg 0$ so that a is not close to zero and being positive implies that costs to under-prediction are weighted more heavily than those to over-prediction. The optimal predictor then exceeds the conditional expectation, so that the conditionally expected error will on average be negative, that is, there will be a tendency to over-predict. The greater the conditionally expected variation around the conditional expectation the greater the tendency to over-predict. As the degree of asymmetry lessens ($a \rightarrow 0$) so the optimal predictor approaches the conditional expectation.

Consider now the effect of the additional term in (17) on the properties of the series of forecast revisions defined by (10). Denoting the revision to the optimal forecast under linear loss by $\tilde{v}_{\tau|\tau-h+1}$, then from (10) and (17) we obtain:

$$\tilde{v}_{\tau|\tau-h+1} = v_{\tau|\tau-h+1} + \frac{a}{2} (\sigma^2_{\tau|\tau-h} - \sigma^2_{\tau|\tau-h+1}) \quad (18)$$

and hence (2) only holds when the prediction error variance is not time-varying. The test for weak efficiency is robust to asymmetric loss assuming a constant prediction error variance, but in general a rejection of weak efficiency might indicate either irrationality or optimally biased forecasts in the face of time-varying prediction error variances.

6 Conclusions

We have applied tests of weak efficiency to series of fixed-event forecasts of GDP growth and consumer price inflation and explored the possibility of pooling the series of revisions over target dates to obtain more powerful tests. An explanation has been offered of the finding of negative serial correlation in forecast revisions. Finally, we argue that asymmetric loss may in principle divorce the concept of weak efficiency from that of rationality.

References

- Ball, R. J. (1962). The prediction of wage-rate changes in the United Kingdom economy 1957-60. *Economic Journal*, **72**, 27-44.
- Bera, A. K., and Higgins, M. L. (1993). ARCH models: Properties, Estimation and Testing. *Journal of Economic Surveys*, **7**, 305-366.
- Bollerslev, T., Chou, R. S., and Kroner, K. F. (1992). ARCH modelling in finance – A review of the theory and empirical evidence. *Journal of Econometrics*, **52**, 5-59.
- Brown, B. W., and Maital, S. (1981). What do economists know? An empirical study of experts' expectations. *Econometrica*, **49**, 491-504.
- Christoffersen, P. F., and Diebold, F. X. (1994). Optimal prediction under asymmetric loss. Discussion paper, Department of Economics, University of Pennsylvania.
- Clements, M. P. (1995). Rationality and the role of judgement in macroeconomic forecasting. *Economic Journal*, **105**, 410-420.
- Deaton, A. S. (1992). *Understanding Consumption*. Oxford: Oxford University Press.
- Engle, R. F. (1982). Autoregressive conditional heteroscedasticity, with estimates of the variance of United Kingdom inflations. *Econometrica*, **50**, 987-1007.
- Engle, R. F. (1995). *ARCH*. Oxford: Oxford University Press.
- Engle, R. F., and Bollerslev, T. (1987). Modelling the persistence of conditional variances. *Econometric Reviews*, **5**, 1-50.
- Fuller, W. A. (1976). *Introduction to Statistical Time Series*. New York: John Wiley & Sons.
- Granger, C. W. J. (1969). Investigating causal relations by econometric models and cross-spectral methods. *Econometrica*, **37**, 424-438.
- Granger, C. W. J., and Morgenstern, O. (1970). *Predictability of Stock Market Prices*. Lexington, Massachusetts: D. C. Heath and Company.
- Hall, R. E. (1978). Stochastic implications of the life cycle-permanent income hypothesis: Evidence. *Journal of Political Economy*, **86**, 971-987.
- Hansen, L. P., and Hodrick, R. J. (1980). Forward exchange rates as optimal predictors of future spot rates: An econometric analysis. *Journal of Political Economy*, **88**, 829-853.
- Holden, K., and Peel, D. A. (1990). On testing for unbiasedness and efficiency of forecasts. *Manchester School*, **58**, 120-127.
- Kirchgässner, G. (1993). Testing Weak Rationality of Forecasts with Different Time Horizons. *Journal of Forecasting*, **12**, 541-558.
- Mankiw, N. G., and Shapiro, M. D. (1986). News or noise: An analysis of GNP revisions. *Survey of Current Business* (May 1986), US Department of Commerce, Bureau of Economic Analysis, 20-25.
- Mincer, J., and Zarnowitz, V. (1969). The evaluation of economic forecasts. In Mincer, J. (ed.), *Economic Forecasts and Expectations*. New York: National Bureau of Economic Research.
- Nordhaus, W. D. (1987). Forecasting efficiency: Concepts and applications. *Review of Economics and Statistics*, **69**, 667-674.

- Samuelson, P. A. (1965). A Proof That Properly Anticipated Prices Fluctuate Randomly. *Industrial Management Review*, **6**, 41–49.
- Shephard, N. G. (1996). Statistical aspects of ARCH and stochastic volatility. In Cox, D. R., Hinkley, D. V., and Barndorff-Nielsen, O. E. (eds.), *Time Series Models In econometrics, finance and other fields*, pp. 1–67. London: Chapman and Hall.
- Swidler, S., and Ketcher, D. (1990). Economic forecasts, rationality, and the processing of new information over time. *Journal of Money, Credit, and Banking*, **22**, 65–76.
- Theil, H. (1958). *Economic Forecasts and Policy*. Amsterdam: North-Holland Publishing Company. (2nd edition 1961).
- Varian, H. R. (1975). A Bayesian approach to real estate assessment. In Fienberg, S. E., and Zellner, A. (eds.), *Studies in Bayesian econometrics and statistics in honor of Leonard J. Savage*, pp. 195–208. Amsterdam: North Holland.
- Wallis, K. F. (1995). Large-scale macroeconomic modelling. In Pesaran, M. H., and Wickens, M. R. (eds.), *Handbook of Applied Econometrics: Macroeconomics*: Basil Blackwell.
- Zarnowitz, V. (1985). Rational expectations and macroeconomic forecasts. *Journal of Business and Economic Statistics*, **3**, 293–311.
- Zellner, A. (1986). Biased predictors, rationality and the evaluation of forecasts. *Economics Letters*, **21**, 45–48.

7 Appendix

For the analysis of the National Institute forecasts in section 4 the regressand and regressor in (7) are given by \mathbf{Y} and \mathbf{X} below, as is the structure of the covariance matrix $\mathbf{\Omega}$:

$$\mathbf{Y} = \begin{bmatrix} v_{93|90:3} \\ \vdots \\ v_{93|93:4} \\ \dots \\ v_{94|91:3} \\ \vdots \\ v_{94|94:4} \end{bmatrix}, \mathbf{X} = \begin{bmatrix} v_{93|90:2} \\ \vdots \\ v_{93|93:3} \\ \dots \\ v_{94|91:2} \\ \vdots \\ v_{94|94:3} \end{bmatrix}, \mathbf{\Omega} = \begin{bmatrix} \begin{bmatrix} \mathbf{I}_4 \sigma_{11}^2 & \mathbf{0} \\ \mathbf{0} & \mathbf{I}_{10} \sigma_{11}^2 \end{bmatrix} & \begin{bmatrix} \mathbf{0} & \mathbf{0} \\ \mathbf{I}_{10} \sigma_{12}^2 & \mathbf{0} \end{bmatrix} \\ \begin{bmatrix} \mathbf{0} & \mathbf{I}_{10} \sigma_{12}^2 \\ \mathbf{0} & \mathbf{0} \end{bmatrix} & \begin{bmatrix} \mathbf{I}_{10} \sigma_{22}^2 & \mathbf{0} \\ \mathbf{0} & \mathbf{I}_4 \sigma_{22}^2 \end{bmatrix} \end{bmatrix} \quad (19)$$

The GLS estimator and the asymptotic covariance matrix are given by $\left(\mathbf{X}'\hat{\mathbf{\Omega}}^{-1}\mathbf{X}\right)^{-1}\left(\mathbf{X}'\hat{\mathbf{\Omega}}^{-1}\mathbf{Y}\right)$ and $\left(\mathbf{X}'\hat{\mathbf{\Omega}}^{-1}\mathbf{X}\right)^{-1}$, where $\hat{\mathbf{\Omega}}$ is $\mathbf{\Omega}$ with the σ_{ij}^2 replaced by $\hat{\sigma}_{ij}^2$ estimated from the OLS residuals of the pooled regressions: $\mathbf{e} = \mathbf{Y} - \mathbf{X}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'\mathbf{Y}$. That is: $\hat{\sigma}_{11}^2 = \frac{1}{13}\mathbf{e}[1:14]'\mathbf{e}[1:14]$, $\hat{\sigma}_{22}^2 = \frac{1}{13}\mathbf{e}[15:28]'\mathbf{e}[15:28]$ and $\hat{\sigma}_{12}^2 = \frac{1}{9}\mathbf{e}[5:14]'\mathbf{e}[15:24]$, and $[x:y]$ denotes elements x through to y of the vector.