Australian National University

Crawford School of Public Policy

CAMA

Centre for Applied Macroeconomic Analysis

A State Space Approach to Evaluate Multi-horizon Forecasts

CAMA Working Paper 67/2017
November 2017

Thomas Goodwin
University of Tasmania

Jing Tian
University of Tasmania and
Centre for Applied Macroeconomic Analysis, ANU

Abstract

We propose a state space modeling framework to evaluate a set of forecasts that target the same variable but are updated along the forecast horizon. The approach decomposes forecast errors into three distinct horizon-specific processes, namely, bias, rational error and implicit error, and attributes forecast revisions to corrections for these forecast errors. We derive the conditions under which forecasts that contain error that is irrelevant to the target can still present the second moment bounds of rational forecasts. By evaluating multi-horizon daily maximum temperature forecasts for Melbourne, Australia, we demonstrate how this modeling framework analyzes the dynamics of the forecast revision structure across horizons. Understanding forecast revisions is critical for weather forecast users to determine the optimal timing for their planning decision.

Keywords

Rational forecasts, implicit forecasts, forecast revision structure, weather forecasts.

JEL Classification

C32; C53

Address for correspondence:

(E) cama.admin@anu.edu.au

ISSN 2206-0332

A State Space Approach to Evaluate Multi-horizon Forecasts *

Thomas Goodwin

University of Tasmania

Jing Tian †

University of Tasmania, CAMA

This version: 30 Oct 2017

Abstract

We propose a state space modeling framework to evaluate a set of forecasts that target the same variable but are updated along the forecast horizon. The approach decomposes forecast errors into three distinct horizon-specific processes, namely, bias, rational error and implicit error, and attributes forecast revisions to corrections for these forecast errors. We derive the conditions under which forecasts that contain error that is irrelevant to the target can still present the second moment bounds of rational forecasts. By evaluating multi-horizon daily maximum temperature forecasts for Melbourne, Australia, we demonstrate how this modeling framework analyzes the dynamics of the forecast revision structure across horizons. Understanding forecast revisions is critical for weather forecast users to determine the optimal timing for their planning decision.

JEL classification: C32; C53

Keywords: Rational forecasts, implicit forecasts, forecast revision structure, weather forecasts

*We thank Dean Croushore, Jan Jacobs and Mardi Dungey for detailed comments on the previous versions of the paper, and the seminar participants at University of Melbourne, Monash University, University of Adelaide, University of Queensland and Queensland University of Technology for helpful discussions. We also thank Andrew Patton for sharing the Matlab code for the rationality tests.

†Corresponding author: Jing Tian, Tasmanian School of Business and Economics, Private Bag 84, University of Tasmania, Hobart TAS 7001, Australia. Tel: +61 3 62262323, fax: +61 3 62267587. Email: jing.tian@utas.edu.au

1 Introduction

Forecasts for the same target are often provided multiple times before the target is realized. One of the most familiar examples relevant to everyday life is weather forecasts. The maximum and minimum temperatures and precipitation for a given day are forecast a few days in advance, and the public receives updates of these forecasts at least once each day before the target day. These multi-horizon weather forecasts play a significant role in decision-making processes related to a wide range of economic activities Dell et al. (2014). For instance, agricultural practices, such as irrigation schedules and timing for harvest operations, rely on updated weather forecasts. Electricity generators receive electricity demand forecasts multiple times before dispatch, and these demand forecasts are often updated as weather forecasts are updated. To maximize the contribution of multi-horizon weather forecasts to planning decisions, it is critical that forecast users understand the forecast revision structure across horizons. Wang and Cai (2009) show that incorporating 7-day weather forecasts may increase crop net profit by 20%, but that "perfect" two-week weather forecasts in the form of actual weather data can achieve a 42% profit increase. Their study acknowledges the economic significance of using longer-horizon weather forecasts in planning but also suggests that the consequent economic benefits may depend on the trade-offs between the length of the forecast horizon and the amount of relevant information contained in long-horizon forecasts. Therefore, knowledge of whether forecast revisions contain information and how that information evolves in a sequence of forecast revisions is crucial for forecast users to determine the optimal timing for their decision making.

This paper provides a framework for forecast users, who may have little knowledge about how forecasts are generated, to improve their understanding of the structure of forecast errors and revisions across forecast horizons. Our modeling approach describes a general form of multi-horizon forecast errors with three distinct horizon-specific processes: rational forecast

error based on Muth's (1961) rational expectation hypothesis, implicit forecast error based on Mill's (1957) implicit expectation hypothesis, and systematic forecast bias. Under this general forecast error structure, we show that forecast revisions along horizons involve reducing the rational forecast error by incorporating newly available information, adjusting for implicit forecast error that is uncorrelated with the target, and correcting for systematic bias.

In the literature, the quality of multi-horizon forecasts is often assessed through rationality tests. Nordhaus (1987) introduce the concept of weak forecast efficiency to evaluate whether forecasts for the same target are rational. Weak-form efficiency requires that multi-horizon forecasts have revisions that are independent of past revisions and past forecast errors. A number of testing approaches have been built on this definition. Clements (1997) considers the scenario where only a small number of fixed-events forecasts are made available and proposes pooling series of multi-horizon forecast across multiple target variables to conduct more powerful tests of weak-form efficiency. Clements and Taylor (2001) extend this approach to allow for non-normally distributed forecast revisions. Davies and Lahiri (1995, 1999); Davies et al. (2011) focus on a three-dimensional panel data approach where multi-horizon forecasts are produced by multiple forecasters and develop tests of rationality in a generalized method of moments framework.

In addition to forecast efficiency, rational multi-horizon forecasts also imply second moment bounds on the forecasts and forecast revisions. Patton and Timmermann (2012) propose a suite of inequality tests based on ten monotonic patterns of second moments implied by the rationality of multi-horizon forecasts. For example, mean squared rational forecasts should weakly increase as the forecast horizon shrinks because the conditional expectation on a larger information set at a shorter horizon has higher variance. Empirically, only a subset of the second moment bounds could be rejected, suggesting irrational multi-horizon forecasts, whereas the other bounds are retained, supporting the claim of rational forecasts. While the size and power properties of the inequality tests might play a role, we argue that rationality tests are

limited in gaining insight about the quality of multi-horizon forecasts from a mixed testing outcome.

Deviating from the commonly used testing approach for rationality, this paper proposes a modeling framework for evaluating a sequence of revised forecasts of the same target. We address the benefits of utilizing our modeling approach from the following aspects.

First, because the specifications of forecasts subject to a single type of error are nested by the specifications of multiple sources of error, model selection methods, such as information criteria and log likelihood tests, can be employed to identify the best fitted error structure of a given set of multi-horizon forecasts. When rational forecast error is estimated to be effectively the sole error type, the multi-horizon forecasts are seen to be rational since revisions are made purely for the purpose of adopting newly available information. Therefore, our modeling approach provides an alternative method to rationality tests to evaluate the overall performance of multi-horizon forecasts.

Second, the modeling approach helps to explain empirically why multi-horizon forecasts may present only a subset of the monotonic properties of the second moments of rational forecasts. We show that forecasts specified to contain only rational error satisfy all the second moment bounds addressed in Patton and Timmermann (2012). When a substantial amount of the forecast error is unrelated to the target (as opposed to future news that is relevant to the target), some of the monotonic patterns of rational forecasts, such as those of the covariances between the revised forecasts and the covariances between the revised forecasts and revisions, always hold. However, the upper variance bounds of the revisions and the monotonicity of the variances of the revisions may no longer be valid. We also derive the conditions under which all the monotonic properties of the rational forecasts hold for multi-horizon forecasts that contain implicit error components.

Furthermore, the estimation of our state space model that explicitly specifies an error structure provides the estimated magnitude of each source of forecast revision at all horizons.

Exploring the revision structure across horizons reveals the dynamics of how information contained in the forecast revisions evolves when approaching the target time. Forecast users are then able to identify the timing of the arrival of the largest amount of new information, which may help them to evaluate the trade-offs between early planning and information in long-horizon forecasts and to choose the optimal horizon forecast for making their planning decisions.

We demonstrate our model-based multi-horizon forecast evaluation approach using a real-time dataset of daily maximum temperature forecasts for Melbourne, Australia. Our results suggest that the weather forecasts revised at a daily frequency up to 14 days before the target day contain both rational and implicit errors. The composition of the sources for forecast revisions changes along the forecast horizon, and the incorporation of newly available information becomes the dominant attribute to revisions within a 7-day horizon. We find that the largest amount of information arrives in the revised forecasts at 6 days out, indicating that 6 days before the target may be the ideal time to consider temperature forecasts in planning decisions. We also illustrate the value of the upgrade of NCEP's Numerical Weather Prediction model on May 22, 2012. The upgrade provides more relevant information in the long-horizon forecasts and shifts the timing of the largest information arrival one day earlier in the forecasts made 7 days before the target.

The remainder of this paper is structured as follows. In Section 2, we propose a model of multi-horizon forecasts that contains multiple sources of forecast errors. In Section 3, we cast our models of various error structures in a state space form. Section 4 studies the internal consistency of the multi-horizon forecasts specified by our state space models. We evaluate multi-horizon weather forecasts as an empirical illustration of our approach in Section 5, and Section 6 concludes this paper.

2 Multifarious errors in multi-horizon forecasts

In this section, we develop a model of multi-horizon forecasts that contain multiple sources of error. We begin by outlining our assumptions and describing our most general model specification. In the subsections that follow, we explain how the target variable and three different types of forecast errors can be modeled as unobserved components.

2.1 A model of multi-horizon forecasts

Suppose forecast users are interested in evaluating multi-horizon forecasts of a stochastic univariate process $\tilde{y} \equiv \{\tilde{y}_t; t = 1, 2, \ldots\}$, where $\tilde{y}_t$ belongs to a general class of stochastic processes and may be non-stationary. Let the forecasts of $\tilde{y}_t$ made at h periods earlier, i.e., at time $t-h$, be denoted by $\hat{y}_{t|t-h}$. Suppose that the longest-horizon forecasts of $\tilde{y}_t$ are observed at horizon $h = H$, where $H > 1$. As the target date approaches, forecast users observe a sequence of forecast revisions, $d_{t|h-1,h} = \hat{y}_{t|t-h+1} - \hat{y}_{t|t-h}$, for the same target $\tilde{y}_t$.

Our aim is to characterize a sequence of multi-horizon forecasts $\left(\hat{y}_{t|t-H}, \hat{y}_{t|t-H+1}, ..., \hat{y}_{t|t-1}\right)'$ with a model that specifies one or more sources of forecast error. The target variable $\tilde{y}_t$ may be observed at time t, for instance, weather variables such as the maximum temperature. In this case, we assume the observed value $y_t = \tilde{y}_t$. The target also may not be observed at the target time t. For example, GDP series are typically published after the target time and are subject to data revisions during a long period after the target time.

We decompose forecasts $\hat{y}_{t|t-h}$ as

$$\hat{y}_{t|t-h} = \tilde{y}_t + \beta_h + \nu_{t|t-h} + \zeta_{t|t-h}. \tag{1}$$

Equation 1 indicates that the forecasting errors $\tilde{y}_t - \hat{y}_{t|t-h}$ nest three types of error: horizon specific bias, $-\beta_h$; accumulated unanticipated shocks during the forecast horizon after the forecasts are made, $-\nu_{t|t-h}$; and the error component unrelated to the target, $-\zeta_{t|t-h}$. We

discuss each type of error in the following subsections.

2.2 Rational forecast errors

The rational expectations hypothesis of Muth (1961) suggests that rational forecasters form their expectations by effectively using all the available information; hence, the only source of error is unanticipated information received after the forecasts are made. Let ω_{t-i} with $i = h-1, \dots 0$ denote the unanticipated shocks that occur at each time point $t-i$ after the forecasting time $t-h$ and before the target time t. Then, the target value of $\tilde{y}_t$ can be written as

$$\tilde{y}_t = \tilde{y}^*_{t,h} + \sum_{i=0}^{h-1} \omega_{t-i}, \tag{2}$$

where $\tilde{y}^*_{t,h}$ takes the target value at time t if there are zero unanticipated shocks over the forecast horizon. A rational forecaster who is able to correctly use all the available information at forecast time $t-h$ and has no bias in forming expectations will produce a rational forecast that is identical to $\tilde{y}^*_{t,h}$. Thus, the rational forecast

$$\hat{y}_{t|t-h} = \tilde{y}_t - \sum_{i=0}^{h-1} \omega_{t-i} = \tilde{y}_t + \nu_{t|t-h}, \tag{3}$$

and the rational forecast error is the accumulation of shocks ω_{t-i} over the horizon, that is, $-\nu_{t|t-h}$.

We now analyze the properties of the rational forecast error. Assume each unanticipated shock $\omega_{t-i} = \sigma_{\omega_i}\eta_{\omega_i,t}$, where $\eta_{\omega_i,t} \sim i.i.d.N(0,1)$. Then, $\nu_{t|t-h}$ has a zero expectation and possesses a number of distinct properties. First, equation (2) indicates that rational forecast error is correlated with the target value, i.e., $cov(\tilde{y}_t, \nu_{t|t-h}) \neq 0$. Second, the unanticipated shocks that occur after the forecasts are made are future information yet to be incorporated into the rational forecasts; therefore, $cov(\hat{y}_{t|t-h}, \nu_{t|t-h}) = 0$. Moreover, the variance of the

rational forecast error, computed as $\mathrm{E}(-\nu^2_{t|t-h}) = \sum_{i=0}^{h-1} \sigma^2_{\omega_{t-i}}$, is non-decreasing as the forecast horizon h increases. This last property is intuitively appealing as we expect there to be less relevant information available to forecasters at longer horizons. A decline in relevant information will cause an increase in forecast uncertainty and an associated increase in the variation of rational forecast errors. Lastly, on the basis of the first two properties, the variance of the target exceeds the variance of the rational forecast.

2.3 Implicit forecast error

Forecasts may deviate from rational forecasts. Mills (1957) introduces the concept of "implicit" expectation to describe the motivation of firms to hold inventories that deviate from rational expectations of demand. This hypothesis was empirically tested by Lovell (1986). In the scenario where rational forecasts of future demand are declining, a firm may still face a relatively high demand forecast and hold inventories in order to reduce fluctuations in production. Such forecast errors are formed due to consideration of the high costs associated with rapid changes in production. They are uncorrelated with the actual demand but correlated with the demand forecasts (as they are introduced by forecasters). The covariance properties of implicit forecast error are opposite to those possessed by rational forecast error.

Implicit forecast error $\zeta_{t|t-h}$ is horizon specific but is typically not a function of horizon h. It is modeled by

$$\zeta_{t|t-h} = \sigma_{\zeta_h} \eta_{\zeta_h,t}, \tag{4}$$

where $\eta_{\zeta_h,t} \sim i.i.d.\ \mathrm{N}(0,1)$. The following covariance assumptions enable us to differentiate them from the rational forecast error: 1) $\mathrm{Cov}(\hat{y}_{t|t-h}, \zeta_{t|t-h}) \neq 0$ and 2) $\mathrm{cov}(\tilde{y}_t, \zeta_{t|t-h}) = 0$. If multi-horizon forecasts are only subject to implicit forecast error, the variance of the target must be less than the variance of the implicit forecasts. In contrast to the monotonic pattern of rational forecast error, the variance of implicit forecast error, $\mathrm{E}(\zeta^2_{t|t-h}) = \sigma^2_{\zeta_h}$, may either

increase or decrease as the forecast horizon shortens.

Forecasters may introduce implicit error into their forecasts with or without intention. For instance, information used by forecasters may contain measurement error, which prevents forecasters from using the information efficiently (Lovell, 1986). Depending on their experience, some forecasters may underreact or overreact to the arrival of new information without intention (Isiklar and Lahiri, 2007). Nordhaus (1987) discusses a scenario in which forecasters may wish to cultivate a reputation for producing stable forecasts and hence intentionally introduce errors that are irrelevant to the target.

2.4 Bias

Both rational and implicit forecast errors have zero expectations; however, a large body of research has empirically found non-zero bias in a wide range of forecasts, such as financial analysts' earning forecasts, gross debt forecasts and inflation forecasts. It is now well accepted that forecasts can remain rational in the presence of bias if the forecasters' loss function is asymmetric (See Christoffersen and Diebold (1997), Lim (2001) and Patton and Timmermann (2007)).

Therefore, we allow for a non-zero systematic bias as the third type of forecast error. We assume that the bias β_h is time-invariant but horizon specific [1] so that empirically, bias can easily be distinguished from rational and implicit forecast errors.

[1] Note that although the forecasters' learning process may indicate time-varying bias, in this paper, we restrict the bias to be constant to avoid identification issues in the state space model estimation. The time-invariant feature of bias is also consistent with Davies and Lahiri (1995).

3 A state space representation of multi-horizon forecasts

3.1 A general structure

We cast the model for multi-horizon forecasts that contain multiple sources of error in a state space form. The time-invariant state space model consists of a set of measurement equations and a set of transition equations, that is,

$$\boldsymbol{y}_t = \boldsymbol{Z}\boldsymbol{\alpha}_t \tag{5}$$

$$\boldsymbol{\alpha}_t = \boldsymbol{T}\boldsymbol{\alpha}_{t-1} + \boldsymbol{R}\boldsymbol{\eta}_t. \tag{6}$$

The measurement vector, $\boldsymbol{y}_t = [\hat{y}_{t|t-H}, \hat{y}_{t|t-(H-1)}, \ldots, y_t]'$, stacks observed multi-horizon forecast variables on the observed target variable, assuming H is the longest horizon at which forecast users are provided forecasts for y_t. In our most general model, the state vector is partitioned as follows

$$\boldsymbol{\alpha}_t = \begin{bmatrix} \tilde{y}_t & \boldsymbol{\Phi}_{\boldsymbol{t}}' & \boldsymbol{\beta_h}' & \boldsymbol{\nu}'_{t|t-h} & \boldsymbol{\zeta}'_{t|t-h} \end{bmatrix}', \tag{7}$$

where $\boldsymbol{\Phi_t}$ is related to the dynamics of the target, which will be explained in detail in subsection 3.3. Suppose the dimensions of $\boldsymbol{\Phi_t}$ are $b \times 1$ and that the state vector $\boldsymbol{\alpha}_t$ has length $(1 + b + H + H + H)$. The associated measurement equation is

$$\boldsymbol{y}_t = \begin{bmatrix} \boldsymbol{Z}_1 & \boldsymbol{Z}_2 & \boldsymbol{Z}_3 & \boldsymbol{Z}_4 & \boldsymbol{Z}_5 \end{bmatrix} \cdot \begin{bmatrix} \tilde{y}_t \\ \boldsymbol{\Phi_t} \\ \boldsymbol{\beta_h} \\ \boldsymbol{\nu}_{t|t-h} \\ \boldsymbol{\zeta}_{t|t-h} \end{bmatrix}, \tag{8}$$

where $\boldsymbol{Z} = [\boldsymbol{Z}_1 \quad \boldsymbol{Z}_2 \quad \boldsymbol{Z}_3 \quad \boldsymbol{Z}_4]$ is a partitioned matrix conforming to the unobserved compo-

nents of the state vector; $\boldsymbol{Z}_1 = \mathbf{1}_{(H+1)}$ (which is an $(H+1) \times 1$ vector of ones) is related to the target variable component; and $\boldsymbol{Z}_2$ is $\mathbf{0}_{(H+1)\times b}$ for the dynamic component of the target. $\boldsymbol{Z}_3$, $\boldsymbol{Z}_4$ and $\boldsymbol{Z}_5$ are each $[\boldsymbol{I}_H, \mathbf{0}'_{1\times H}]'$ (an $H \times H$ identity matrix, atop a conformably defined vector of zeros) and are related to the bias component, the rational error component and the implicit error component, respectively. This measurement equation reflects the forecast decomposition of equation (1) and also equates the observed value of the target variable y_t with the "true" unobserved target value $\tilde{y}_t$ [2].

Transition equations describe the dynamics of the unobserved components in terms of the state vector

$$\begin{bmatrix} \tilde{y}_t \\ \boldsymbol{\Phi}_t \\ \boldsymbol{\beta}_h \\ \boldsymbol{\nu}_{t|t-h} \\ \boldsymbol{\zeta}_{t|t-h} \end{bmatrix} = \begin{bmatrix} T_{11} & \boldsymbol{T}_{12} & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \boldsymbol{T}_{21} & \boldsymbol{T}_{22} & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \boldsymbol{T}_3 & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \boldsymbol{T}_4 & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \boldsymbol{T}_5 \end{bmatrix} \cdot \begin{bmatrix} \tilde{y}_{t-1} \\ \boldsymbol{\Phi}_{t-1} \\ \boldsymbol{\beta}_h \\ \boldsymbol{\nu}_{t-1|t-1-h} \\ \boldsymbol{\zeta}_{t-1|t-1-h} \end{bmatrix}$$
$$+ \begin{bmatrix} R_{11} & \boldsymbol{R}_{12} & \mathbf{0} & \boldsymbol{R}_4 & \mathbf{0} \\ \boldsymbol{R}_{21} & \boldsymbol{R}_{22} & \mathbf{0} & \boldsymbol{R}_{24} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \boldsymbol{R}_3 & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & -\boldsymbol{U} \cdot \mathrm{diag}(\boldsymbol{R}_4) & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathrm{diag}(\boldsymbol{R}_5) \end{bmatrix} \cdot \begin{bmatrix} \eta_{\xi,t} \\ \boldsymbol{\eta}_{\Phi,t} \\ \boldsymbol{\eta}_{\beta_h,t} \\ \boldsymbol{\eta}_{\omega_h,t} \\ \boldsymbol{\eta}_{\zeta_h,t} \end{bmatrix}, \tag{9}$$

where $\boldsymbol{U}$ is an $H \times H$ matrix with zeros below the main diagonal and each of the remaining elements equal to one, and $[\eta_{\xi,t}, \boldsymbol{\eta}'_{\Phi,t}, \boldsymbol{\eta}'_{\beta_h,t}, \boldsymbol{\eta}'_{\omega_h,t}, \boldsymbol{\eta}'_{\zeta_h,t}]' \sim i.i.d.N(\mathbf{0}, \boldsymbol{I})$. Looking at the partitions of the transition coefficient $\boldsymbol{T}$, T_{11} is a scalar; $\boldsymbol{T}_{12}$, $\boldsymbol{T}_{21}$ and $\boldsymbol{T}_{22}$ are $1 \times b$, $b \times 1$ and $b \times b$,

[2] In a scenario where the target value is unobserved, such as GDP series, the last row of the measurement equation related to y_t can be removed. This measurement equation can also be modified by adding measurement error to $\tilde{y}_t$ in the last row to accommodate the scenario in which the observed target value is known to be greatly affected by measurement error.

respectively; and $\boldsymbol{T}_3$, $\boldsymbol{T}_4$ and $\boldsymbol{T}_5$ are all $H \times H$. $\mathbf{0}$ is an $H \times H$ null matrix. The partitions of the loading coefficients $\boldsymbol{R}$ have similar dimensions to those in $\boldsymbol{T}$, and $\boldsymbol{R}_{24}$,$\boldsymbol{R}_4$ and $\boldsymbol{R}_5$ are all $1 \times H$.

3.2 Specifications for forecasting errors

We first explain the blocks for the unobserved forecast errors. Since the bias is a horizon-specific constant that does not vary over time, we set $\boldsymbol{T}_3$ as an $H \times H$ identity matrix and $\boldsymbol{R}_3$ as an $H \times H$ null matrix. On the basis of the discussion in section 2.2, at each horizon $\nu_{t|t-h} = \sum_{i=0}^{h-1} \omega_{t-i} = \sum_{i=0}^{h-1} \sigma_{\omega_i} \eta_{\omega_i,t}$; therefore, we let $\boldsymbol{T}_4$ be an $H \times H$ null matrix and $\boldsymbol{R}_4 = [\sigma_{\omega_{H-1}}, \dots \sigma_{\omega_0}]$. The rational forecast error component is then given by

$$\boldsymbol{\nu}_{t|t-h} = -\boldsymbol{U} \cdot \text{diag}(\boldsymbol{R}_2) \cdot \boldsymbol{\eta}'_{\nu,t} = \begin{bmatrix} -\sigma_{\omega_{H-1}} & -\sigma_{\omega_{H-2}} & \dots & -\sigma_{\omega_1} & -\sigma_{\omega_0} \\ 0 & -\sigma_{\omega_{H-2}} & \dots & -\sigma_{\omega_1} & -\sigma_{\omega_0} \\ \vdots & \ddots & \ddots & -\sigma_{\omega_1} & -\sigma_{\omega_0} \\ 0 & \dots & \dots & 0 & -\sigma_{\omega_0} \end{bmatrix} \cdot \begin{bmatrix} \eta_{\omega_{H-1}t} \\ \eta_{\omega_{H-2}t} \\ \vdots \\ \eta_{\omega_1} \\ \eta_{\omega_0 t} \end{bmatrix}. \quad (10)$$

Jacobs and Van Norden (2011) use this restricted specification of the loading matrix to model the rational data revision process.

The implicit forecast error is horizon specific and uncorrelated across horizons. Thus, the matrix $\boldsymbol{T}_5$ is a null matrix, and the dynamics of the implicit component are completely described by diag($\boldsymbol{R}_5$), where $\boldsymbol{R}_5$ is a row vector of standard deviations of the implicit forecast error, i.e., $[\sigma_{\zeta_H}, \sigma_{\zeta_{H-1}}, ..., \sigma_{\zeta_1}]$. The equation below describes the process of implicit forecast

error:

$$\boldsymbol{\zeta_{t|t-h}} = \text{diag}(\boldsymbol{R}_5) \cdot \boldsymbol{\eta}'_{\zeta,t} = \begin{bmatrix} \sigma_{\zeta_H} & 0 & \dots & 0 \\ 0 & \sigma_{\zeta_{H-1}} & \ddots & \vdots \\ \vdots & \ddots & \ddots & 0 \\ 0 & \dots & 0 & \sigma_{\zeta_1} \end{bmatrix} \cdot \begin{bmatrix} \eta_{\zeta_H t} \\ \eta_{\zeta_{H-1} t} \\ \vdots \\ \eta_{\zeta_1 t} \end{bmatrix}. \tag{11}$$

The most general model represented by equations (8) and (9) nests several simpler forecast error structures. For example, by removing the blocks for bias $\boldsymbol{\beta}_h$, multi-horizon forecasts that contain both rational and implicit errors are represented by

$$\boldsymbol{y}_t = \begin{bmatrix} \boldsymbol{Z}_1 & \boldsymbol{Z}_2 & \boldsymbol{Z}_4 & \boldsymbol{Z}_5 \end{bmatrix} \cdot \begin{bmatrix} \tilde{y}_t \\ \boldsymbol{\Phi}_t \\ \boldsymbol{\nu}_{t|t-h} \\ \boldsymbol{\zeta}_{t|t-h} \end{bmatrix}, \text{ and} \tag{12}$$

$$\begin{bmatrix} \tilde{y}_t \\ \boldsymbol{\Phi}_t \\ \boldsymbol{\nu}_{t|t-h} \\ \boldsymbol{\zeta}_{t|t-h} \end{bmatrix} = \begin{bmatrix} T_{11} & \boldsymbol{T}_{12} & \boldsymbol{0} & \boldsymbol{0} \\ \boldsymbol{T}_{21} & \boldsymbol{T}_{22} & \boldsymbol{0} & \boldsymbol{0} \\ \boldsymbol{0} & \boldsymbol{0} & \boldsymbol{T}_4 & \boldsymbol{0} \\ \boldsymbol{0} & \boldsymbol{0} & \boldsymbol{0} & \boldsymbol{T}_4 \end{bmatrix} \cdot \begin{bmatrix} \tilde{y}_{t-1} \\ \boldsymbol{\Phi}_{t-1} \\ \boldsymbol{\nu}_{t-1|t-1-h} \\ \boldsymbol{\zeta}_{t-1|t-1-h} \end{bmatrix} + \begin{bmatrix} R_{11} & \boldsymbol{R}_{12} & \boldsymbol{R}_4 & \boldsymbol{0} \\ R_{21} & \boldsymbol{R}_{22} & \boldsymbol{R}_{24} & \boldsymbol{0} \\ \boldsymbol{0} & \boldsymbol{0} & -\boldsymbol{U} \cdot \text{diag}(\boldsymbol{R}_4) & \boldsymbol{0} \\ \boldsymbol{0} & \boldsymbol{0} & \boldsymbol{0} & \text{diag}(\boldsymbol{R}_5) \end{bmatrix} \cdot \begin{bmatrix} \eta_{\xi,t} \\ \boldsymbol{\eta}_{\Phi,t} \\ \boldsymbol{\eta}_{\omega,t} \\ \boldsymbol{\eta}_{\zeta,t} \end{bmatrix}. \tag{13}$$

A state space form for multi-horizon forecasts subject purely to rational forecast error can then be obtained by further restricting $\boldsymbol{Z}_5$ and $\boldsymbol{R}_5$ to be null matrices, that is,

$$\boldsymbol{y}_t = \begin{bmatrix} \boldsymbol{Z}_1 & \boldsymbol{Z}_2 & \boldsymbol{Z}_4 \end{bmatrix} \cdot \begin{bmatrix} \tilde{y}_t \\ \boldsymbol{\Phi}_t \\ \boldsymbol{\nu}_{t|t-h} \end{bmatrix}, \text{ and} \tag{14}$$

$$\begin{bmatrix} \tilde{y}_t \\ \boldsymbol{\Phi}_t \\ \boldsymbol{\nu}_{t|t-h} \end{bmatrix} = \begin{bmatrix} T_{11} & \boldsymbol{T}_{12} & \mathbf{0} \\ \boldsymbol{T}_{21} & \boldsymbol{T}_{22} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \boldsymbol{T}_4 \end{bmatrix} \cdot \begin{bmatrix} \tilde{y}_{t-1} \\ \boldsymbol{\Phi}_{t-1} \\ \boldsymbol{\nu}_{t-1|t-1-h} \end{bmatrix} + \begin{bmatrix} R_{11} & \boldsymbol{R}_{12} & \boldsymbol{R}_4 \\ \boldsymbol{R}_{21} & \boldsymbol{R}_{22} & \boldsymbol{R}_{24} \\ \mathbf{0} & \mathbf{0} & -\boldsymbol{U} \cdot \text{diag}(\boldsymbol{R}_4) \end{bmatrix} \cdot \begin{bmatrix} \eta_{\xi,t} \\ \boldsymbol{\eta}_{\Phi,t} \\ \boldsymbol{\eta}_{\omega,t} \end{bmatrix}. \quad (15)$$

Alternatively, in the absence of unanticipated new information over the forecast horizon, multi-horizon forecasts may only consist of forecast error uncorrelated with the target. In this case, we set $\boldsymbol{Z}_4$ in equation (12) to be a null matrix, and let $\boldsymbol{R}_4$ be a $1 \times H$ vector of zeros in equation (13). The pure implicit forecasts are then given by

$$\boldsymbol{y}_t = \begin{bmatrix} \boldsymbol{Z}_1 & \boldsymbol{Z}_2 & \boldsymbol{Z}_5 \end{bmatrix} \cdot \begin{bmatrix} \tilde{y}_t \\ \boldsymbol{\Phi}_t \\ \boldsymbol{\zeta}_{t|t-h} \end{bmatrix}, \text{ and} \quad (16)$$

$$\begin{bmatrix} \tilde{y}_t \\ \boldsymbol{\Phi}_t \\ \boldsymbol{\zeta}_{t|t-h} \end{bmatrix} = \begin{bmatrix} T_{11} & \boldsymbol{T}_{12} & \mathbf{0} \\ \boldsymbol{T}_{21} & \boldsymbol{T}_{22} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \boldsymbol{T}_5 \end{bmatrix} \cdot \begin{bmatrix} \tilde{y}_{t-1} \\ \boldsymbol{\Phi}_{t-1} \\ \boldsymbol{\nu}_{t-1|t-1-h} \end{bmatrix} + \begin{bmatrix} R_{11} & \boldsymbol{R}_{12} & \mathbf{0} \\ \boldsymbol{R}_{21} & \boldsymbol{R}_{22} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \text{diag}(\boldsymbol{R}_5) \end{bmatrix} \cdot \begin{bmatrix} \eta_{\xi,t} \\ \boldsymbol{\eta}_{\Phi,t} \\ \boldsymbol{\eta}_{\zeta,t} \end{bmatrix}. \quad (17)$$

3.3 Specifications for the target variable

We now discuss the dynamics of the target $\tilde{y}_t$. Starting from equation (2) in section 2.2, we specify the value of the target under zero unanticipated shocks over the longest forecast horizon H as

$$\tilde{y}^*_{t,H} = \tilde{y}_{t-H} + \xi_t, \quad (18)$$

where $\xi_t = \sigma_\xi \eta_{\xi,t}$. Since ξ_t is a part of the rational forecast of every horizon, it can be interpreted as the anticipated change in $\tilde{y}_t$ compared with $\tilde{y}_{t-H}$. Equation (2) can be rewritten as

$$\tilde{y}_t = \tilde{y}_{t-H} + \xi_t + \sum_{i=0}^{H-1} \omega_{t-i}. \tag{19}$$

Note that this equation is analogous to a random walk process. For forecast users who have no knowledge of the true data generating process of the target variable, it is natural to expect that the target value at time t is the same as the target value at the time when the forecasts are made and that the actual deviations are driven by the accumulation of unanticipated shocks that occur over the forecast horizon.

We specify the unobserved state variable $\mathbf{\Phi}_t$ to be $\tilde{y}_{t-H+1}$, which similarly to equation (19), is equal to the value of the target variable in the previous period $t-H$ plus an unanticipated shock relevant to the target during the current period $t-H+1$, as well as a remainder, namely, ϵ_{t-H+1}. The blocks for the target variable in equation (9) are hence given as

$$\begin{bmatrix} \tilde{y}_t \\ \mathbf{\Phi}_t \end{bmatrix} = \begin{bmatrix} \tilde{y}_t \\ \tilde{y}_{t-H+1} \end{bmatrix} = \begin{bmatrix} 0 & 1 \\ 0 & 1 \end{bmatrix} \cdot \begin{bmatrix} \tilde{y}_{t-1} \\ \tilde{y}_{t-H} \end{bmatrix} + \begin{bmatrix} \sigma_\xi & 0 & \sigma_{\omega_{H-1}} & \sigma_{\omega_{H-2}} & \dots & \sigma_{\omega_1} & \sigma_{\omega_0} \\ 0 & \sigma_\epsilon & \sigma_{\omega_{H-1}} & 0 & \dots & 0 & 0 \end{bmatrix} \cdot \begin{bmatrix} \eta_{\xi,t} \\ \eta_{\epsilon,t-H+1} \\ \boldsymbol{\eta}_{\omega,t} \end{bmatrix}.$$

Specifically, the transition coefficients $T_{11} = \boldsymbol{T}_{21} = 0$ and $\boldsymbol{T}_{12} = \boldsymbol{T}_{22} = 1$, the loading coefficients $R_{11} = \sigma_\xi$, $\boldsymbol{R}_{22} = \sigma_\epsilon$, and $\boldsymbol{R}_{12} = \boldsymbol{R}_{21} = 0$ and $\boldsymbol{R}_{24}$ is a $1 \times H$ vector of $[\sigma_{\omega_{H-1}}, 0, ..., 0]$.

3.4 Estimation and analysis of the forecast revision structure

We check the various state space representations of multi-horizon forecasts against the sufficient conditions of controllability and observability provided by Harvey (1989) and Jacobs and van Norden (2007). The parameters in the models discussed in the previous subsection are identified.

The maximum likelihood estimator with the Kalman filter is used to estimate the unobserved component models. Since the specifications of the multi-horizon error structure are nested, model selection methods, such as log likelihood ratio tests and the conventional information criteria, may be implemented to determine the best fitted error structure.

The estimation of multifarious error models helps to gain insight into the sources of forecast revisions across horizons. Suppose that the multi-horizon forecasts consist of all three types of forecast errors and are modeled by equations (8) and (9). The differences between $\hat{y}_{t|t-(h-1)}$ and $\hat{y}_{t|t-h}$, namely, the marginal revisions made between two updating points $t-(h-1)$ and $t-h$, are derived as follows,

$$\begin{aligned} d_{t|h-1,h} &= \hat{y}_{t-(h-1)} - \hat{y}_{t|t-h} \\ &= \beta_{h-1} - \beta_h + \sigma_{\omega_{h-1}} \eta_{\omega_{h-1},t} + \zeta_{t|t-(h-1)} - \zeta_{t|t-h}. \end{aligned} \tag{20}$$

Therefore, the mean squared forecast revisions (MSFR) can be decomposed as

$$\begin{aligned} MSFR_{t|h-1,h} &= E(d^2_{t|h-1,h}) \\ &= (\beta_{h-1} - \beta_h)^2 + \sigma^2_{\omega_{h-1}} + \sigma^2_{\zeta_{h-1}} + \sigma^2_{\zeta_h}. \end{aligned} \tag{21}$$

Equation (21) indicates that marginal revisions are made to correct bias (first term), to adopt newly available information (second term) and to adjust implicit errors that are uncorrelated with the target (the last two terms).

4 Internal Consistency

In this section, we discuss the internal consistency of the multi-horizon forecasts specified by our state space representations. We focus on the monotonicity properties of the second moment across forecast horizons. Patton and Timmermann (2012) analyze ten monotonic patterns of the second moments of rational multi-horizon forecasts and utilize them to design

inequality constraints to test rationality. Empirically, some of these inequality constraints are often retained while others are rejected. Although a joint testing procedure, such as a Bonferroni bound, can provide a pragmatic solution, rationality tests in general have limitations for gaining insights into why forecasts empirically may only present a subset of the properties for rational forecasts.

In the following subsections, we show that by decomposing forecast errors and revisions into multiple types, our modeling approach for evaluating multi-horizon forecasts helps to explain the mixed rationality testing outcomes of Patton and Timmermann (2012). We focus on three specifications (assuming no bias [3]), namely, rational forecasts (see equations (14) and (15)), rational and implicit forecasts (see equations (12) and (13)), and implicit forecasts (see equations (12) and (13)). We investigate whether each type of forecast satisfies the monotonicity properties of the second moments of rational forecasts, and if not, what conditions are required to validate the patterns. The results are summarized in Table 1, and the proofs related to rational forecasts, rational and implicit forecasts, and implicit forecasts are included in Appendices A.1, A.2 and A.3, respectively.

[Insert Table 1]

The first column of Table 1 lists the second moment bounds for rational multi-horizon forecasts in relation to forecasts, forecast errors and forecast revisions. The first three rows present the monotonicity of the variances. Rows four to six include the monotonicity for the covariance between forecasts, the covariance between the target and forecast errors, and the covariance between the target and forecasts. The covariance patterns and the upper variance bounds related to forecast revisions are listed in the last four rows. The symbol "$\surd$" indicates that the internal consistency presented in a particular row holds for the forecasts specified by the state space representation named in the columns. We use the symbol "X" if a particular

[3]The assumption of no bias simplifies the proofs, and since we specify the bias to be horizon specific but time-invariant, the internal consistency defined by the second moments is not affected.

aspect of internal consistency does not hold.

4.1 For rational forecasts

The second column, titled " Rational", is for the multi-horizon forecasts that consist of only rational forecast error. We have checked against each of Patton and Timmermann's second moment bounds for rational forecasts and conclude that the multi-horizon forecasts specified by our rational forecast model possess all of the internal consistency properties defined by the forecast rationality tests of Patton and Timmermann (2012).

4.2 For rational and implicit forecasts

For unbiased multi-horizon forecasts that have both rational and implicit forecast errors, we report the results in the third column, titled "Rational & Implicit". Regardless of the relative sizes of the rational and implicit forecast errors, some of the monotonic patterns, such as those for the covariances between forecasts, the two covariances related to the target and the covariance between forecasts and revisions, always hold. This result indicate that even when multi-horizon forecasts contain a substantial amount of forecast error irrelevant to the target, some inequality constraints featuring rational forecasts can be retained by the rationality tests of Patton and Timmermann (2012).

Furthermore, Table 1 indicates that when the compositions of the rational and implicit errors satisfy certain conditions, multi-horizon forecasts that contain both types of error present all the second moment bounds of the rational forecasts. With the exception of the upper variance bound for forecast revisions, these conditions compare the variances of the newly acquired information between two updating points with the differences in the variances of the implicit forecast errors. The following analysis of the compositions of the mean squared forecast errors (MSFE) may help to understand these conditions.

The MSFE of the rational and implicit forecasts made at a long horizon $h = l$ is given by

$$MSFE_{t|t-l} = E(e^2_{t|t-l}) = \sum_{i=0}^{l-1} \sigma^2_{\omega_i} + \sigma^2_{\zeta_l}, \tag{22}$$

and similarly, for the rational and implicit forecasts made at a shorter horizon $h = s$, we have

$$MSFE_{t|t-s} = \sum_{i=0}^{s-1} \sigma^2_{\omega_i} + \sigma^2_{\zeta_s}. \tag{23}$$

The changes in $MSFE$ thus take the form

$$\Delta MSFE_{t|s,l} = MSFE_{t|t-s} - MSFE_{t|t-l} = -\sum_{i=s}^{l-1} \sigma^2_{\omega_i} + \sigma^2_{\zeta_s} - \sigma^2_{\zeta_l}. \tag{24}$$

This equation explicitly shows that the changes in MSFE are due to a reduction in the rational forecast error (the first term) and correction of the implicit forecast error (the last two terms). The relative size of these two sources determines whether MSFE can be lowered by forecast revisions, specifically, when $\sum_{i=s}^{l-1} \sigma^2_{\omega_i} \geq \sigma^2_{\zeta_s} - \sigma^2_{\zeta_l}$.

There are two possible forecast error structures that imply the rational component of revisions dominates the implicit component of revisions and hence smaller MSFE values in shorter horizon forecasts. First, if the rational forecast error outweighs the implicit error across all forecast horizons, the variance of revisions $d_{t|s,h}$ increases and the variance of forecasts $\hat{y}_{t|t-h}$ decreases as the forecast horizon h extends further from the target date. Alternatively, if revisions introduce more implicit error or fail to reduce the variance of the implicit error in the previous long-horizon forecasts, then the differences in the variance of implicit error between a short horizon and a long horizon must be no greater than zero, and hence less than the (positive) variance of the rational component of the revisions.

In the last two rows, we report that the upper variance bound of the forecast revisions for rational forecasts can also be applied to rational and implicit forecasts if and only if the

variance of the newly available information is no less than the sum of the variances of the implicit forecast errors made at the two forecast times. This inequality condition relates to the compositions of the MSFR. For example, comparing the forecasts made at point $t-l$ and $t-s$, we have

$$MSFR_{t|s,l} = E\left[(\hat{y}_{t|t-s} - \hat{y}_{t|t-l})^2\right] = \sum_{i=s}^{l-1} \sigma_{\omega_i}^2 + \sigma_{\zeta_s}^2 + \sigma_{\zeta_l}^2. \tag{25}$$

The first term in $MSFR_{t|s,l}$, i.e. $\sum_{i=s}^{l-1} \sigma_{\omega_i}^2$, is the same as the left side of the inequality condition for the variance bound of $d_{t|s,l}$. It measures the portion of MSFR due to utilizing the new information available between two updating points. The right side of the inequality condition $\sigma_{\zeta_s}^2 + \sigma_{\zeta_l}^2$ measures the MSFR due to adjustment for implicit errors that are irrelevant to the target. This inequality restriction implies that if the new information accounts for more than 50% of the overall MSFR, the variance bounds of the revisions for rational forecasts are still valid, even if the forecasts contain implicit error.

Note that by using equation (24) and (25), we are certain that $MSFR_{t|s,l} > -\Delta MSFE_{t|s,l}$ as long as the short-horizon forecasts $\hat{y}_{t|t-s}$ consist of implicit errors, i.e., $\sigma_{\zeta_s} \neq 0$. The magnitude of $MSFR_{t|s,l}$ can be interpreted as the amount of effort that forecasters made in the revision process, and $-\Delta MSFE_{t|s,l}$ measures the reward of the forecast revisions in terms of forecast accuracy. Isiklar and Lahiri (2007) compare the effort and reward to assess whether rational forecasters react to news in an optimal way. By allowing for an implicit error component contained in the multi-horizon forecasts, we extend their interpretation of the comparison. If the revised forecasts still contain error unrelated to the target, the revision effort is always greater than the revision reward, and the revisions are seen to be suboptimal. The revision effort can be fully compensated for if revisions eliminate the implicit errors; in this case, the forecast revisions are optimal.

4.3 For implicit forecasts

The last column of Table 1 shows that multi-horizon forecasts subject purely to implicit forecast error can possess some second moment bounds for rational forecasts, despite the fact that forecast revisions do not incorporate any new information. We derive the conditions required and find that the monotonic pattern in the variance of implicit forecast error plays the key role.

For example, the third row indicates that if the variance of the implicit forecast error decreases as the forecast horizon shrinks, the variance of the revisions, and hence the value of $MSFR_{t|s,h}$, also decreases. This property implies that even when multi-horizon forecasts fail to reject the rationality tests based on MSFE and MSFR, we cannot conclude that the forecasts are rational across horizons or that the forecast revisions are optimal.

Note that the conditions for monotonicity in the first two rows are opposite to each other. This result has a clear intuition. The non-increasing variance of rational forecasts builds upon the fact that the conditional expectations using a larger information set at a shorter forecast horizon are associated with larger variance. Implicit forecasts, however, are not formed based on relevant information and hence are subject to only irrelevant errors. Therefore, if the variance of the implicit forecast error increases with the forecast horizon, the variance of the forecasts must also increase. This difference in the monotonicity of forecast variance is key to distinguishing between implicit forecasts and rational forecasts.

The two covariances related to the target, i.e., the covariance between the target and the forecast errors and the covariance between the target and the forecasts, are constant across forecast horizons for pure implicit forecasts. In practice, when the target values are observable, these two covariance bounds may help to differentiate implicit forecasts from rational forecasts. When the target is unobservable, we can focus on the constant covariance between forecasts at different horizons and the constant covariance between forecasts and revisions made at two

forecast dates (see the fourth row and the eighth row). Since these constant covariances are subsets of the inequality constraints under the null hypothesis of Patton and Timmermann (2012), pure implicit forecasts that have errors that are entirely irrelevant to the target can still lead to retaining the monotonicity of these covariances featured by rational forecasts.

The last two rows of Table 1 show that the upper variance bounds for revisions are likely to be invalid for implicit forecasts (when there exists non-zero implicit forecast error) because the covariance between the target and the revisions and the covariance between the short-horizon forecasts and previous revisions are both zero. The only scenario where $var(d_{t|s,h}) = cov(\tilde{y}_t, d_{t|s,h})$ or $var(d_{t|m,l}) = cov(\hat{y}_{t|t-s}, d_{t|m,l})$ is when the implicit forecast error $\zeta_{t|t-h}$ is the same across horizons. Consequently, there is no revision; therefore, the variance of the revisions is zero.

To summarize, the internal consistency discussed by Patton and Timmermann (2012) can be observed in multi-horizon forecasts that consist of implicit forecast error irrelevant to the target. The practical implication is that forecasts that fail to reject the rationality tests based on the monotonicity of the second moments may not be rational forecasts. In other words, multi-horizon forecasts may be internally consistent but not rational, and revisions deviate from being optimal. The state space approach proposed in this paper decomposes forecast errors into multiple types at each forecast horizon and hence is able to provide greater insight into the composition of the forecast errors as the forecasting horizon approaches the target time.

5 Evaluating Multi-horizon Weather Forecasts

In this section, we evaluate multi-horizon forecasts of the daily maximum temperature (degrees Celsius) for Melbourne, Australia to demonstrate how our model-based evaluation approach

extracts the various types of forecast errors and forecast revision components and to suggest how these results could assist decision makers in planning.

During the past few decades, meteorological services in Australia have produced increasingly accurate weather forecasts at ever increasing forecast horizons (Stern, 2008; Stern and Davidson, 2015). For decision makers who rely on weather conditions, the technological improvement of weather forecasts provides an opportunity to choose between using long-horizon weather forecasts and short-horizon weather forecasts. We aim to provide insight into the revision process for daily maximum temperature forecasts over 14-day horizons so that decision makers can optimally time their planning decisions conditional on the weather forecast revision process.

5.1 Data

We retrieve the data from `http://www.weather-climate.com`, which consists of an experimental daily maximum temperature forecast series generated at multiple 14-day horizons and an observed maximum daily temperature series for Melbourne, Australia. The sample period runs from February 1, 2009 to December 31, 2014, comprising a total of 2159 days, with forecasts available at 14,13,...,2,1 days out from each observation date.

These experimental daily maximum temperature forecasts were produced in real-time using a forecast combination algorithm, as documented in Stern (2007) and Stern and Davidson (2015). A number of data sources are used to produce the combined forecasts, including the official forecasts from the Australian Bureau of Meteorology (BOM), the previous day's maximum temperature forecasts, statistical forecasts,[4] and climatological forecasts [5]. Table 2 describes the forecast combination weightings used to generate the daily maximum tem-

[4]Stern and Davidson (2015) provide a brief explanation of the statistical forecasts. These forecasts are computed for local weather based on the output of the long-range numerical weather prediction (NWP) models provided by the National Center for Environmental Prediction (NCEP) of the National Oceanic and Atmospheric Administration (NOAA). See Wilks (2011) for examples of statistical forecast methods for meteorological variables.

[5]Climatological forecasts are the averages of historical observations over many years.

perature forecasts. The combination weights vary depending on the length of the forecast horizon. For example, BOM does not publish forecasts of maximum temperature for more than 7 days out from the target date; hence heavier weights are imposed on the statistical and climatological forecasts that are based on long-run and large-scale weather forecasts.

[Insert Table 2]

Figure 1 plots the daily maximum temperature observations (black line) and the corresponding meteorological forecasts generated 7 days (blue line) and 14 days (red line) out from the target observation date, spanning the whole sample. This figure presents three main features. First, the observed maximum temperature series is more volatile than both forecast series, and the forecasts made 7 days before the target are more volatile than the forecasts made 14 days before the target. These observations meet Patton and Timmermann (2012)'s necessary conditions for rational forecasts, that is, the variance of short-horizon forecasts is no less than the variance of long-horizon forecasts and is bounded by the variance of the realization. However, as discussed in Section 4, other than the accumulation of new information as the forecast horizon shrinks, the inclusion of error irrelevant to the target could also increase the variance of the forecasts. Second, there is a permanent increase in the variations of the two forecast series beginning from mid-2012. Stern and Davidson (2015) discuss an improvement in forecast skill beginning from mid-2012, which they attribute to a major upgrade of the NWP models on 22 May, 2012. In addition, the observed temperature series and the two forecast temperature series exhibit more variability on warmer days than on cooler days. Stern and Davidson (2015) note that the competing influence of warm dry winds from the Australian interior and cool moist winds from the Southern Ocean make temperature forecasting for Melbourne particularly challenging, and this influence is strongest during the warmer months.

[Insert Figure 1]

We report both whole-sample and sub-sample means and the standard deviations of the maximum daily temperature observations and all fourteen horizon forecasts in Table 3[6]. The statistics confirm our observations regarding Figure 1.

[Insert Table 3]

5.2 Results

5.2.1 Forecast decomposition

We focus on three model specifications: the rational model that specifies unanticipated new information over the forecast horizon as the only source of forecast error; the rational and implicit model that allows some forecast errors to be unrelated to the target; and the bias+rational+implicit model that adds a horizon-specific but time-invariant systematic bias [7].

Table 4 reports the estimation results of the three alternative models. The top panel presents the estimated values of $\sigma_{\omega_{h-1}}$, which represents the marginal increase in information content owing to forecast revisions made at a shorter horizon $h-1$ compared to the forecast of horizon h. The estimated standard deviations of the implicit errors at each horizon h, denoted by σ_{ζ_h}, are reported in the middle panel. These values capture the size of the noise uncorrelated with the targeted maximum daily temperature. The bottom panel shows the Kalman smoothed estimates of the horizon-specific forecast bias. We also report the log likelihood values and the Akaike and Baysian Information Criteria for each model.

[Insert Table 4]

[6]For the Southern hemisphere, we denote the period from September 21 to March 20 as warm days and the period from March 21 to September 20 as cool days.

[7]The estimation results of other alternative models, including the pure implicit model, the bias and rational model, and the bias and implicit model, can be provided upon request. The log likelihood values suggest that these three models are less preferred to the rational model, the rational and implicit model and the bias+rational+implicit model reported here.

Since the rational model is nested by both multiple-error models, we can apply either log likelihood ratio tests or information criterion to identify the empirically preferred model(s). Both approaches reach the conclusion that multi-horizon forecasts of the maximum daily temperature for Melbourne up to 14 days out are subject to multiple types of forecast errors. Between the two multiple-error models that have different state variables but the same number of unknown parameters, we can simply compare the maximum values of log likelihood. Since the rational and implicit model achieves a higher log likelihood value than that of the model with systematic bias, we prefer this bi-error structure and continue our further analysis based on the estimates of the rational and implicit model [8].

The estimated value of $\sigma_{\omega_{h-1}}$ starts low in the revised forecast at the horizon of 13 days and then gradually increases in the subsequent revisions until it peaks at the horizon of 6 days. Afterwards, as the horizon shrinks further, the estimated $\sigma_{\omega_{h-1}}$ declines, suggesting decreasing marginal information content adoption within 5 days before the target day.

The estimated standard deviation of the implicit forecast error σ_{ζ_h} exhibits a different pattern. It is highest at the longest forecast horizon and gradually declines as the forecast time approaches 7 days before the target day. Revisions of the maximum temperature forecasts within a week of the target day leads to trivial implicit errors.

We use Figure 2 to depict the dynamics of the forecast revision structure across forecast horizons. With a modification of equation (21) for the rational and implicit forecasts by removing the bias term, we can decompose the value of $MSFR_{t|h-1,h}$ into $\sigma^2_{\omega_{h-1}}$ plus $\sigma^2_{\zeta_h}$ and $\sigma^2_{\zeta_{h-1}}$. Figure 2 shows how the sizes of these components for the revised daily maximum temperature forecasts evolve from from 13 days to 1 day before the realization. The total length of each bar represents the size of $MSFR_{t|h-1,h}$. Within each bar, the red color represents the

[8] Note that for the bias+rational+implicit model, the estimated biases over all 14 horizons are significantly negative, but the estimates related to the rational revision and implicit forecast errors are virtually the same as those from the rational and implicit model. Therefore, our following analysis, specifically the change in marginal information content in the revisions over forecast horizons, is not affected by our modeling choice.

variation of newly adopted information in the updated forecast made at horizon $h-1$, and the blue and green colors represent the variations of irrelevant implicit error made at horizons $h-1$ and h, respectively.

[Insert Figure 2]

In general, the MSFR between two adjacent horizons decreases as the forecast day approaches the target day. Furthermore, the proportional contribution of each source of forecast revision varies over horizons. For example, at horizons longer than 8 days before the target date, adjusting for irrelevant noise accounts for most of the MSFR values. When the forecast day is within a week of the target day, more than 90% of the MSFR values is accounted for by the incorporation of newly available information. This alteration of the main source of forecast revisions can be explained by how the multi-horizon maximum temperature forecasts for each horizon are constructed. The weighting structure reported in Table 2 shows that at horizons longer than 7 days out, forecasts are a combination of statistical and climatological forecasts derived from long-run and large-scale mathematical models. The large amount of variation in the combination forecasts that is irrelevant to the target reflects the inaccuracy of these sourcing forecasts for local maximum daily temperature when made more than 7 days out from the target day. For horizons within a week of the target, BOM's official forecasts contribute 50% of the combination forecasts. The fact that information adoption becomes the sole source of short-term forecast revisions indicates BOM's very high short-term prediction skill for maximum daily temperature.

An understanding the forecast revision structure across horizons can potentially help decision makers to choose an optimal horizon forecast for their planning decisions. As analyzed above, the daily maximum temperature forecast revisions made in the second week before the target day are associated with minimal information content and hence are less likely to result in significant economic benefits for the forecast users. If early planning is preferred, they may

be better off using 12-day-out revised forecasts rather than 9-day-out revised forecasts. By waiting for a few days longer and bearing the opportunity costs of delayed planning, decision makers are provided with revised weather forecasts that incorporate a substantially larger amount of information. Since the information content reaches its highest level in the 6-day-out revised forecasts and then decreases as the forecast horizon shrinks, the value added by using 6-day-out forecasts may be higher than that using forecasts at horizons shorter than 6 days.

5.2.2 Rationality tests

We now compare the evaluation results of our modeling approach with the rationality testing approach proposed by Patton and Timmermann (2012). Table 5 reports the p-values of seven second moment bound tests for forecasts up to 7 days out and up to 14 days out from the target day. The null hypothesis of the monotonic properties in the second moments for rational multi-horizon forecasts are listed in the first column. The p-values in all inequality tests are large, so the null hypothesis that the maximum temperature forecasts at both 1-week and 2-week horizons are rational is retained.

[Insert Table 5]

The discussions in Section 4.2 suggest that it is possible for multi-horizon forecasts that consist of errors that are irrelevant to the target to fail to reject the rationality properties. We check the results against the conditions for which the rational and implicit forecasts can be identified as rational forecasts. We focus our discussion on the daily maximum temperature forecasts made in the second week out (when h is 8 to 14 days) since these long-horizon forecasts contain a substantial amount of implicit error that is irrelevant to the target. The estimated variance of the implicit forecast error is non-decreasing with increasing horizon h, and the difference in these variances across horizons is smaller than the difference in the variances

of the rational errors. The estimation results guarantee that long-horizon daily maximum temperature forecasts have non-increasing mean squared forecasts $MSF_{t|t-h}$, non-decreasing $MSFE_{t|t-h}$, non-decreasing $MSFR_{t|s,h}$ and non-decreasing covariance $cov(e_{t|t-h}, d_{t|s,h})$ (see rows 1,2,3 and 7 of Table 1). The accumulation of newly relevant information arriving between the forecast dates must be larger than the total variance of the irrelevant implicit error made in the forecasts so that the variance of the revisions is bounded by the covariance of the revisions and the target or the covariances of the revisions and the subsequently updated forecasts (see rows 9 and 10 in Table 1). Our estimation results show that when two updating points are close, there is little relevant new information but large irrelevant noise in the revisions; hence the above condition does not hold. As the forecasts are further revised, information accumulation increases and the implicit error component declines. The upper variance bounds may still hold for forecasts made in the second week out.

5.3 Subsample evaluations

5.3.1 Effects of the NWP upgrade

The difference in forecast variabilities before and after May 22, 2012 is evident in Figure 1. On May 22, 2012, the NCEP's operational system, including the long-range numerical weather prediction models that serve as an input for the maximum temperature forecasts, were upgraded. In this section, we examine the effect of the upgrade on the forecast and revision structure.

We estimate the rational and implicit model while allowing the values of $\sigma_{\omega_{h-1}}$ and σ_{ζ_h} in matrix $\boldsymbol{R}$ to be different in the subsamples pre- and post-May 22, 2012. We then calculate the estimated $MSFR_{t|h-1,h}$ and its components in the two subsamples over the forecast horizons and present the results in Figure 3.

[Insert Figure 3]

The sizes of $MSFR_{t|h-1,h}$ at horizons longer than 7 days after the NWP model upgrade are approximately four times larger than those prior to the upgrade. However, the upgrade does not have much impact on the $MSFR_{t|h-1,h}$ at short horizons within 5 days before the target day is reached. The increase in $MSFR_{t|h-1,h}$ for long horizons is associated with higher $\sigma^2_{\omega_{h-1}}$ and higher $\sigma^2_{\zeta_h}$ and $\sigma^2_{\zeta_{h-1}}$, suggesting that following the NWP model upgrades, the forecast revisions contain more relevant information but also more irrelevant noise. The upgrade does not affect the composition of the forecast revisions across horizons, and information adoption becomes the dominant source of revisions for forecasts made within 7 days of the target.

Before the NWP model upgrade, the $MSFR_{t|h-1,h}$ peaks at 6 days out, where new information adoption between two adjacent updating points is also the highest. After the upgrade, the highest marginal information adoption occurs one day earlier, at 7 days out, and the $MSFR_{t|h-1,h}$ again exhibits a declining pattern as forecasts are made closer to the target. The NWP model upgrade provides users of maximum daily temperature forecasts an optimal forecast horizon one day earlier than the optimal horizon prior to the upgrade. Given the same amount of information adoption in forecast revisions, since both optimal horizon forecasts contain a similar amount of newly available information, incorporating longer-horizon forecasts in planning may result in higher profits.

5.3.2 Effects of seasons

The geographical location of Melbourne leads to wide variation in the maximum temperature during warm months, making it difficult to forecast. In this section, we study whether the composition of forecast revisions across horizons is consistent over seasons. We estimate a rational and implicit model allowing for different $\sigma_{\omega_{h-1}}$ and σ_{ζ_h} between warm months and cool months. Figure 4 illustrates the values and the compositions of $MSFR_{t|h-1,h}$ over the forecast horizons, with the top panel for the warm months and the bottom panel for the cool months.

[Insert Figure 4]

The $MSFR_{t|h-1,h}$ values across all horizons in warm months are much higher than those in cool months, consistent with the fact that the forecast volatilities at all horizons are higher in warmer months. In general, the values of $MSFR_{t|h-1,h}$ decrease as the forecast horizon shrinks, regardless of the season, except that the marginal forecast revisions made at 6 days out provides the highest $MSFR_{t|h-1,h}$.

The shares of the sources for forecast revisions are consistent between seasons. Regardless of whether the target days are in the warm or cool season, the long-horizon forecast revisions (made in the second week before the target date) are mainly due to adjustment of implicit forecast errors that are irrelevant to the target, suggesting inefficiencies in the long-horizon maximum temperature forecasts. Starting from 7-day-out forecasts, newly available information becomes the dominant attribute of the revisions, and these short-horizon maximum temperature forecasts are effective rational forecasts.

6 Conclusion

The availability of multi-horizon forecasts of the same target offers forecast users an opportunity to investigate the revision structure across forecast horizons. This paper proposes a state space modeling approach that decomposes multi-horizon forecast errors into several unobserved components, including 1) rational forecast errors that occur due to unanticipated information related to the target, 2) implicit sources of forecast errors that are irrelevant to the target, and 3) horizon-specific bias that captures systematic under- or over-forecasts. By using this modeling approach, forecast users can explore the best fitted forecast error structure for the whole set of multi-horizon forecasts and study the key attributes of forecast revisions at each horizon. Understanding the dynamics of forecast revision structure across horizons may help forecast users to identify the most desirable horizon forecast for their planning decisions.

We investigate the internal consistency of the multi-horizon forecasts specified by various unobserved component models. Focusing on the monotonicity properties of the second moment proposed by Patton and Timmermann (2012), we derive the conditions under which forecasts that consist of implicit error irrelevant to the target can posses the internal consistency of rational forecasts. For instance, when multi-horizon forecasts are subject to implicit error with substantially large variance, as long as the variance of implicit error increases with increasing forecast horizon, the mean squared forecast errors must be an increasing function of the forecast horizon. This monotonicity is a well-known property for rational forecasts. Our approach may be used to provide an explanation for potential mixed testing outcomes from the suite of Patton and Timmermann (2012) inequality tests.

In our application, we use the state space modeling approach to evaluate maximum daily temperature forecasts for Melbourne, Australia. Using the forecasts up to 14 days before the target, which ranges from February 1, 2009 to December 31, 2014, we find that these multi-horizon weather forecasts contain both rational forecast error and implicit forecast error. The variance of each type of forecast error changes along the forecast horizon, with the short-horizon revised forecasts (made up to 7 days before the target) containing more information and less irrelevant noise, than the long-horizon forecasts (made in the second week before the target).

An important application of our modeling approach is to analyze the sources of forecast revisions and how the composition of these sources changes along the forecast horizon. By decomposing the value of the mean squared temperature forecast revisions between two adjacent updating points into a rational component due to adopting newly available information and an implicit component due to adjusting irrelevant noise, we show that marginal information adoption accounts for a similarly small proportion of forecast revisions made from 13 days to 9 days before the target. Information adoption becomes (effectively) the single attribute of the forecast revisions as the forecast horizon shrinks to within 7 days. The largest amount of

information adoption in MSFR over the whole sample period is achieved at the horizon of 6 days before declining as the target date is approached.

Our state space modeling approach provides a means to select the best decision-making point. Early planning is beneficial but generally suffers from little relevant information being incorporated in revised forecasts made at long horizons. Our results of the forecast revision structure of the maximum daily temperature show that the forecast horizon of 6 days, when the highest amount of information is found in the revisions (using the whole sample period), could be the ideal horizon for decision making. We also find that the upgrade of the NWP models moves the occurrence of the highest marginal information adoption to the 7 day horizon and hence shifts the best decision-making time one day earlier. This improvement shows the value of investment to upgrade weather prediction systems since early planning of economic activities related to future weather conditions, including crop irrigation and harvest and energy supply, results in potential profits for the decision makers.

References

Christoffersen, P. F. and Diebold, F. X. (1997). Optimal prediction under asymmetric loss. *Econometric Theory*, 13:808–817.

Clements, M. P. (1997). Evaluating the rationality of fixed-event forecasts. *Journal of Forecasting*, 16(4):225–239.

Clements, M. P. and Taylor, N. (2001). Robust evaluation of fixed-event forecast rationality. *Journal of Forecasting*, 20(54):285–295.

Davies, A. and Lahiri, K. (1995). A new framework for analyzing survey forecasts using three-dimensional panel data. *Journal of Econometrics*, 68:205–227.

Davies, A. and Lahiri, K. (1999). Re-examining the rational expectations hypothesis using panel data on multi-period forecasts. In Hsiao, C., Lahiri, K., Lee, H.-F., and Pesaran, H. M., editors, *Analysis of panels and limited dependent variable models*, pages 226–254. Cambridge University Press.

Davies, A., Lahiri, K., and Sheng, X. (2011). Analyzing three-dimensional panel data of forecasts. In Clements, M. P. and Hendry, D. F., editors, *The Oxford handbook of economic forecasting*, pages 473–495. Oxford University Press.

Dell, M., Jones, B. F., and Olken, B. A. (2014). What do we learn from the weather? The new climate–economy literature. *Journal of Economic Literature*, 52(3):740–798.

Harvey, A. C. (1989). *Forecasting, Structural Time Series Models and the Kalman Filter*. Cambridge: Cambridge University Press.

Isiklar, G. and Lahiri, K. (2007). How far ahead can we forecast? evidence from cross-country surveys. *International Journal of Forecasting*, 23:167–187.

Jacobs, J. P. and van Norden, S. (2007). Appendix to modelling data revisions: Measurement error and dynamics of true values. Mimeo, University of Groningen.

Jacobs, J. P. and Van Norden, S. (2011). Modeling data revisions: Measurement error and dynamics of true values. *Journal of Econometrics*, 161(2):101–109.

Lim, T. (2001). Rationality and analysts' forecast bias. *Journal of Finance*, 56(1):369–385.

Lovell, M. C. (1986). Tests of the rational expectations hypothesis. *The American Economic Review*, 76:110–124.

Mills, E. S. (1957). The theory of inventory decisions. *Econometrica*, 25:222–238.

Muth, J. F. (1961). Rational expectations and the theory of price movements. *Econometrica*, 29:315–335.

Nordhaus, W. D. (1987). Forecasting efficiency: concepts and applications. *The Review of Economics and Statistics*, 69(4):667–674.

Patton, A. and Timmermann, A. (2007). Properties of optimal forecasts under asymmetric loss and nonlinearity. *Journal of Econometrics*, 140(2):884–918.

Patton, A. J. and Timmermann, A. (2012). Forecast rationality tests based on multi-horizon bounds. *Journal of Business & Economic Statistics*, 30:1–17.

Stern, H. (2007). Improving forecasts with mechanically combined predictions. *Bulletin of the American Meteorological Society*, 88:850–851.

Stern, H. (2008). The accuracy of weather forecasts for Melbourne, Australia. *Meteorological Applications*, 15(1):65–71.

Stern, H. and Davidson, N. E. (2015). Trends in the skill of weather prediction at lead times of 1–14 days. *Quarterly Journal of the Royal Meteorological Society*, 141(692):2726–2736.

Wang, D. and Cai, X. (2009). Irrigation scheduling: Role of weather forecasting and farmers' behavior. *Journal of Water Resource Planning and Management*, 135(5):364–373.

Wilks, D. S. (2011). *Statistical methods in the atmospheric sciences*, volume 100. Academic Press.

Table 1: Summary of the monotonicity patterns in the second moments of different types of multi-horizon forecasts

	Rational	**Rational & Implicit**	**Implicit**
$var(\hat{y}_{t\|t-s}) \geq var(\hat{y}_{t\|t-l})$	$\surd$	$\surd$ iff $\sum_{i=s}^{l-1} \sigma^2_{\omega_i} \geq \sigma^2_{\zeta_l} - \sigma^2_{\zeta_s}$	$\surd$ iff $\sigma_{\zeta_s} \geq \sigma_{\zeta_l}$
$var(e_{t\|t-s}) \leq var(e_{t\|t-l})$	$\surd$	$\surd$ iff $\sum_{i=s}^{l-1} \sigma^2_{\omega_i} \geq \sigma^2_{\zeta_s} - \sigma^2_{\zeta_l}$	$\surd$ iff $\sigma_{\zeta_s} \leq \sigma_{\zeta_l}$
$var(d_{t\|s,m}) \leq var(d_{t\|s,l})$	$\surd$	$\surd$ iff $\sum_{i=m}^{l-1} \sigma^2_{\omega_i} \geq \sigma^2_{\zeta_m} - \sigma^2_{\zeta_l}$	$\surd$ iff $\sigma_{\zeta_m} \leq \sigma_{\zeta_l}$
$cov(\hat{y}_{t\|t-m}, \hat{y}_{t\|t-s}) \geq cov(\hat{y}_{t\|t-l}, \hat{y}_{t\|t-s})$	$\surd$	$\surd$	$\overline{cov} = var(\tilde{y}_t)$
$cov(\tilde{y}_t, e_{t\|t-s}) \leq cov(\tilde{y}_t, e_{t\|t-l})$	$\surd$	$\surd$	$\overline{cov} = 0$
$cov(\tilde{y}_t, \hat{y}_{t\|t-s}) \geq cov(\tilde{y}_t, \hat{y}_{t\|t-l})$	$\surd$	$\surd$	$\overline{cov} = var(\tilde{y}_t)$
$cov(e_{t\|t-m}, d_{t\|s,m}) \leq cov(e_{t\|t-l}, d_{t\|s,l})$	$\surd$	$\surd$ iff $\sum_{i=m}^{l-1} \sigma^2_{\omega_i} \geq \sigma^2_{\zeta_m} - \sigma^2_{\zeta_l}$	$\surd$ iff $\sigma_{\zeta_m} \leq \sigma_{\zeta_l}$
$cov(\hat{y}_{t\|t-s}, d_{t\|s,m}) \leq cov(\hat{y}_{t\|t-s}, d_{t\|s,l})$	$\surd$	$\surd$	$\overline{cov} = \sigma^2_{\zeta_s}$
$var(d_{t\|s,l}) \leq 2cov(\tilde{y}_t, d_{t\|s,l})$	$\surd$	$\surd$ iff $\sum_{i=s}^{l-1} \sigma^2_{\omega_i} \geq \sigma^2_{\zeta_s} + \sigma^2_{\zeta_l}$	X iff $\sigma_{\zeta_s} \neq 0$ and/or $\sigma_{\zeta_l} \neq 0$, or $\zeta_{t\|t-s} \neq \zeta_{t\|t-l}$
$var(d_{t\|m,l}) \leq 2cov(\hat{y}_{t\|t-s}, d_{t\|m,l})$	$\surd$	$\surd$ iff $\sum_{i=m}^{l-1} \sigma^2_{\omega_i} \geq \sigma^2_{\zeta_m} + \sigma^2_{\zeta_l}$	X iff $\sigma_{\zeta_m} \neq 0$ and/or $\sigma_{\zeta_l} \neq 0$, or $\zeta_{t\|t-m} \neq \zeta_{t\|t-l}$

Notes: We use h to represent a generic forecast horizon and $h = H, H-1, ..., l, l-1, ..., m, m-1, ...s, s-1, ..., 0$, where H is the longest forecast horizon, and horizons $l \geq m \geq s$. The second moment bounds listed in the first column are properties possessed by rational multi-horizon forecasts (see Patton and Timmermann (2012)). The symbol $\surd$ indicates that forecasts specified by a particular type of multi-horizon error structure satisfy the second moment bound listed in the row. The symbol X indicates the opposite. We use $\overline{cov}$ to mark the constant covariances.

Table 2: Forecast combination weightings for the multi-horizon daily maximum temperature forecasts in Melbourne, Australia

Horizon h	Official	Previous	Statistical	Climatology
1 to 7 days	0.50	0.25	0.25	
8 to 13 days		0.25	0.50	0.25
14 days			0.50	0.50

Notes: This table lists the weightings of the multi-horizon real-time forecasts of the daily maximum temperature of Melbourne. The weightings are provided by Stern and Davidson (2015), which explains that the forecast combination is based on official forecasts from the Australian Bureau of Meteorology, previous day's forecasts, statistical forecasts, and climatological forecasts.

Table 3: Descriptive statistics of daily maximum temperature observations and multi-horizon forecasts for Melbourne, Australia

	Full sample		Pre-NWP upgrade		Post-NWP upgrade		Warm months		Cool months	
Horizon h	Mean	Std. Dev.	Mean	Std. Dev.	Mean	Std. Dev.	Mean	Std. Dev.	Mean	Std. Dev.
14	20.53	4.64	20.79	4.51	20.19	4.79	24.07	3.11	17.13	3.06
13	20.57	4.72	20.79	4.50	20.30	4.98	24.14	3.29	17.17	3.08
12	20.60	4.79	20.78	4.52	20.37	5.10	24.21	3.37	17.15	3.09
11	20.61	4.82	20.78	4.53	20.40	5.17	24.24	3.42	17.14	3.12
10	20.62	4.84	20.81	4.54	20.38	5.18	24.27	3.41	17.13	3.10
9	20.60	4.79	20.83	4.55	20.31	5.07	24.20	3.40	17.15	3.09
8	20.62	4.77	20.82	4.54	20.36	5.05	24.18	3.40	17.21	3.13
7	20.74	4.95	20.82	4.59	20.65	5.37	24.31	3.78	17.33	3.23
6	20.88	5.30	21.04	5.06	20.68	5.58	24.53	4.40	17.39	3.38
5	20.93	5.45	21.11	5.24	20.70	5.69	24.57	4.71	17.45	3.47
4	20.96	5.56	21.12	5.37	20.75	5.78	24.62	4.90	17.46	3.52
3	20.97	5.62	21.11	5.44	20.79	5.85	24.62	5.04	17.48	3.56
2	20.98	5.70	21.10	5.54	20.84	5.89	24.66	5.16	17.48	3.59
1	20.99	5.76	21.13	5.62	20.81	5.93	24.67	5.28	17.47	3.60
0	21.11	6.01	21.24	5.89	20.94	6.15	24.71	5.73	17.66	3.86

Notes: This table reports the means and standard deviations of daily maximum temperature observations and their multi-horizon forecasts for Melbourne, Australia. The full sample covers 2159 target days from February 1, 2009 to December 31, 2014. The pre-NWP upgrade period spans February 1, 2009 to May 21, 2012 before the Numerical Weather Prediction models experienced a major upgrade on May 22, 2012. The post-NWP update period is from May 22, 2012 to December 31, 2014. For the southern hemisphere, we denote the days from September 21 to March 20 as warm months, and the cool months cover March 21 to September 20. The horizon h indicates the number of days before the target day that the forecasts are made. The last row, $h = 0$, represents the observed daily maximum temperature.

Table 4: Estimation results of the multi-horizon forecasts of daily maximum temperature for Melbourne, Australia between February 1, 2009 and December 31, 2014

		Alternative models					
		Rational		**Rational + Implicit**		**Bias +Rational + Implicit**	
Forecast Errors		Estimate	Std err	Estimate	Std err	Estimate	Std err
Rational revision	$\sigma_{\omega_{13}}$	1.568	(0.037)	0.373	(0.062)	0.374	(0.061)
	$\sigma_{\omega_{12}}$	1.654	(0.040)	0.551	(0.063)	0.549	(0.063)
	$\sigma_{\omega_{11}}$	1.564	(0.039)	0.496	(0.066)	0.499	(0.065)
	$\sigma_{\omega_{10}}$	1.534	(0.039)	0.583	(0.056)	0.584	(0.056)
	σ_{ω_9}	1.472	(0.037)	0.633	(0.053)	0.634	(0.053)
	σ_{ω_8}	1.421	(0.037)	0.981	(0.043)	0.979	(0.043)
	σ_{ω_7}	1.416	(0.036)	1.315	(0.037)	1.308	(0.037)
	σ_{ω_6}	1.560	(0.037)	1.560	(0.037)	1.554	(0.035)
	σ_{ω_5}	1.095	(0.036)	1.095	(0.035)	1.094	(0.035)
	σ_{ω_4}	0.973	(0.023)	0.973	(0.024)	0.973	(0.024)
	σ_{ω_3}	0.820	(0.022)	0.820	(0.022)	0.820	(0.022)
	σ_{ω_2}	0.728	(0.020)	0.728	(0.020)	0.727	(0.020)
	σ_{ω_1}	0.697	(0.019)	0.697	(0.019)	0.697	(0.019)
	σ_{ω_0}	1.690	(0.032)	1.691	(0.032)	1.687	(0.032)
Implicit error	$\sigma_{\zeta_{14}}$			1.193	(0.032)	1.192	(0.032)
	$\sigma_{\zeta_{13}}$			1.223	(0.030)	1.224	(0.030)
	$\sigma_{\zeta_{12}}$			1.134	(0.032)	1.134	(0.032)
	$\sigma_{\zeta_{11}}$			1.078	(0.036)	1.077	(0.036)
	$\sigma_{\zeta_{10}}$			1.040	(0.035)	1.039	(0.035)
	σ_{ζ_9}			0.891	(0.036)	0.892	(0.036)
	σ_{ζ_8}			0.524	(0.059)	0.527	(0.058)
	σ_{ζ_7}			0.027	(0.010)	0.026	(0.011)
	σ_{ζ_6}			0.018	(0.005)	0.017	(0.005)
	σ_{ζ_5}			0.014	(0.003)	0.013	(0.003)
	σ_{ζ_4}			0.016	(0.006)	0.016	(0.005)
	σ_{ζ_3}			0.011	(0.003)	0.011	(0.003)
	σ_{ζ_2}			0.010	(0.002)	0.010	(0.002)
	σ_{ζ_1}			0.017	(0.005)	0.017	(0.006)
Bias	β_{14}					-0.581	(0.082)
(smoothed states)	β_{13}					-0.531	(0.082)
	β_{12}					-0.506	(0.081)
	β_{11}					-0.494	(0.080)
	β_{10}					-0.487	(0.078)
	β_9					-0.509	(0.076)
	β_8					-0.489	(0.072)
	β_7					-0.361	(0.065)
	β_6					-0.226	(0.056)
	β_5					-0.176	(0.050)
	β_4					-0.148	(0.046)
	β_3					-0.138	(0.042)
	β_2					-0.121	(0.039)
	β_1					-0.118	(0.036)
Log likelihood		-53,621		-51,583		-51,717	
Akaike Info Criterion		107,275		103,226		103,493	
Baysian Info Criterion		107,366		103,397		103,663	

Notes: The sample covers 2159 target days from February 1, 2009 to December 31, 2014. For the daily maximum temperature of each target date, meteorological forecasts are made at a daily frequency at horizons from 14 days to 1 day before the target date. The standard errors of the estimates are in parentheses.

Table 5: Results of the rationality tests proposed by Patton and Timmermann (2012).

	Horizons	
Test	$h = 1$ to $h = 7$	$h = 8$ to $h = 14$
$\text{Var}(\hat{y}_{t\|t-s}) \geq \text{Var}(\hat{y}_{t\|t-l})$	0.96	0.85
$\text{Var}(e_{t\|t-s}) \leq \text{Var}(e_{t\|t-l})$	0.95	0.85
$\text{Var}(d_{t\|s,m}) \leq \text{Var}(d_{t\|s,l})$	0.91	0.95
$\text{Cov}(\hat{y}_{t\|t-s}, \tilde{y}_t) \geq \text{Cov}(\hat{y}_{t\|t-l}, \tilde{y}_t)$	0.95	0.93
$\text{Cov}(\hat{y}_{t\|t-m}, \hat{y}_{t\|t-s}) \geq \text{Cov}(\hat{y}_{t\|t-l}, \hat{y}_{t\|t-s})$	0.93	0.95
$\text{Var}(d_{t\|s,l}) \leq 2\text{Cov}(\tilde{y}_t, d_{t\|s,l})$	0.78	0.50
$\text{Var}(d_{t\|m,l}) \leq 2\text{Cov}(\hat{y}_{t\|t-s}, d_{t\|m,l})$	0.77	0.13

Notes: We use h to represent a generic forecast horizon and $h = H, H-1, ..., l, l-1, ..., m, m-1, ...s, s-1, ..., 0$, where H is the longest forecast horizon, and horizons $l \geq m \geq s$. The inequality relations in the first column are the null hypotheses of the Patton and Timmermann (2012) tests for forecast rationality. The values reported in this table are p-values from each test. We use Matlab code provided by Andrew Patton for the inequality tests, and partition the forecast horizon into $h = 1$ to $h = 7$ and $h = 8$ to $h = 14$ and test each partition separately.

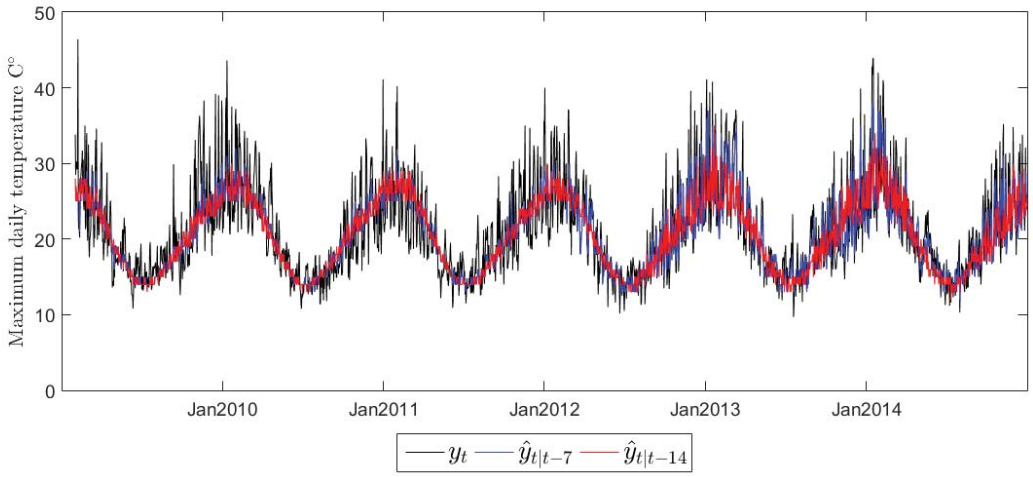

Figure 1: Daily maximum temperature observations and multi-horizon forecasts for Melbourne, Australia from Feb 01, 2009 to Dec 31, 2014

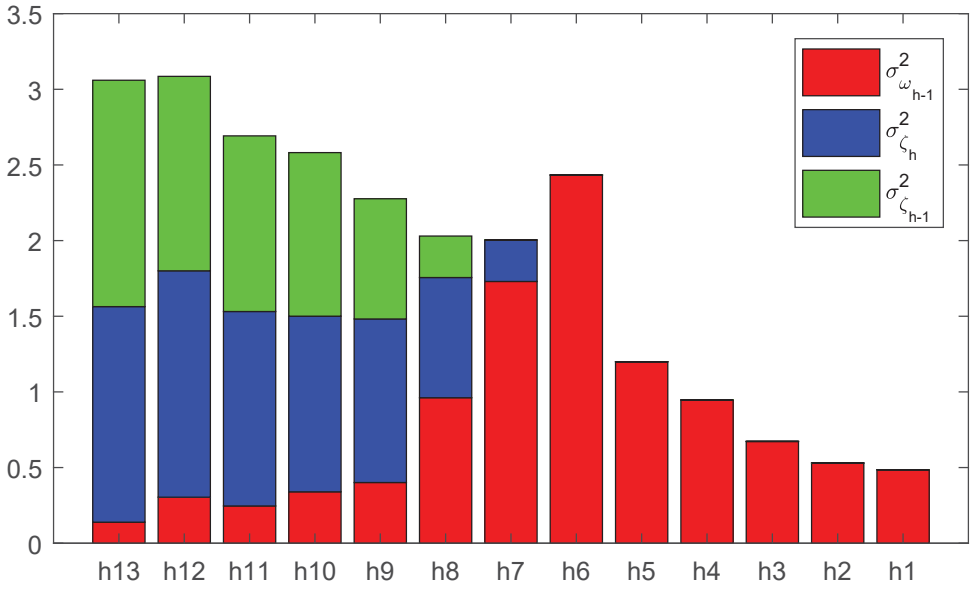

Figure 2: $MSFR_{t|h-1,h}$ and the components

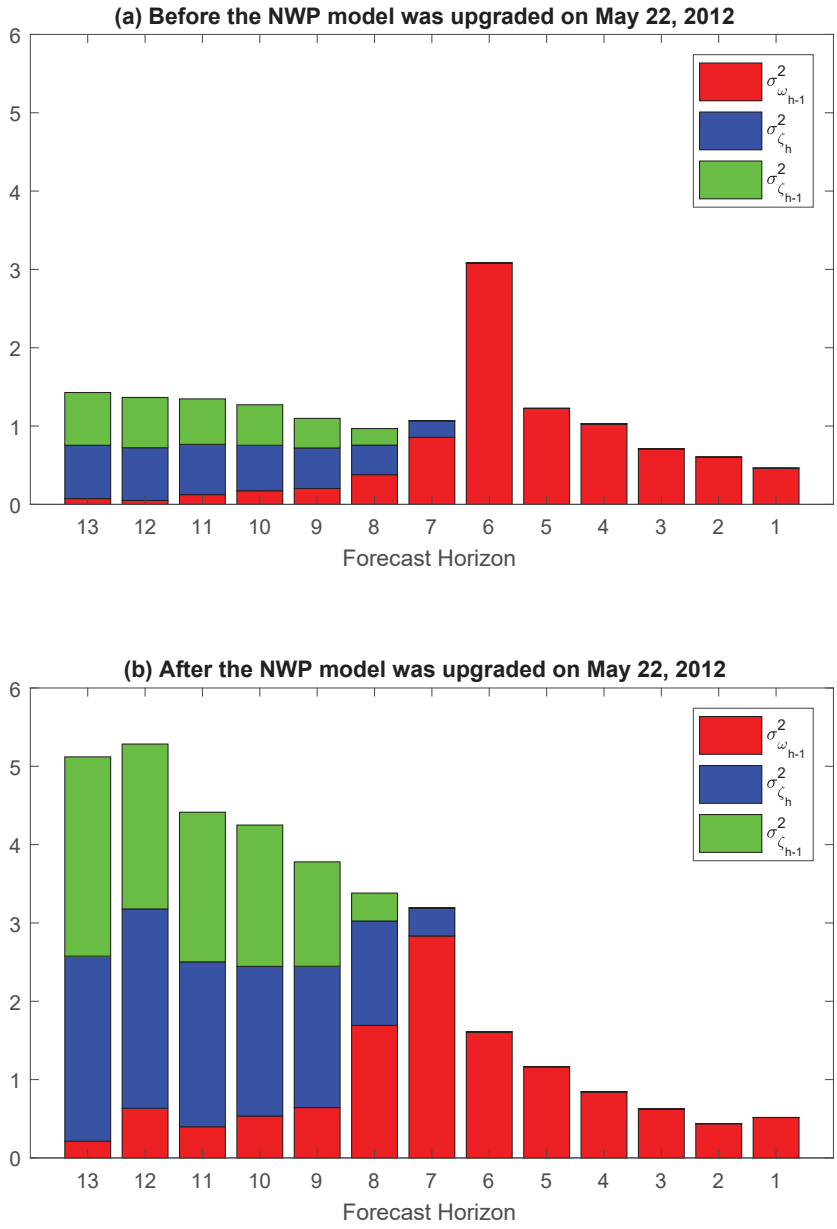

Figure 3: $MSFR_{t|h-1,h}$ and the components before and after May 20, 2012

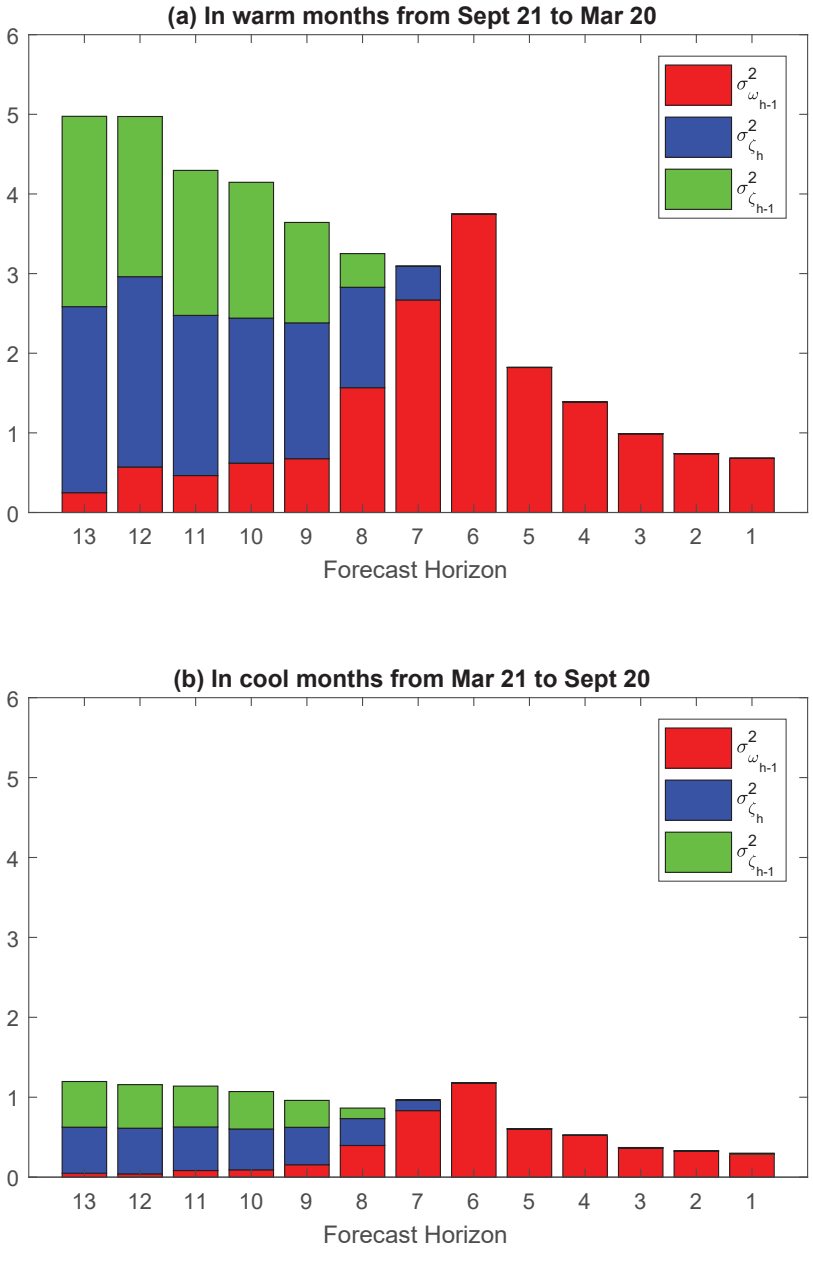

Figure 4: $MSFR_{t|h-1,h}$ and the components in warm and cool months

A Appendix: Proofs for Internal Consistency

This appendix provides a formal connection between two approaches to multi-horizon forecast evaluation: our model-based approach and the variance bound test approach proposed by Patton and Timmermann (2012). We list the variables defined by our models, which will be used in the proofs in the following subsections, in Table 6.

Table 6: Expression of the multi-horizon forecasts, forecast errors and revisions (with no bias)

Variable	Expression
Target	$\tilde{y}_t = \tilde{y}_{t-H} + \xi_t + \sum_{i=0}^{H-1} \sigma_{\omega_i} \eta_{\omega_i,t}$
Rational forecasts	$\hat{y}_{t\|t-h} = \tilde{y}_t - \sum_{i=0}^{h-1} \sigma_{\omega_i} \eta_{\omega_i,t}$
Rational+implicit forecasts	$\hat{y}_{t\|t-h} = \tilde{y}_t - \sum_{i=0}^{h-1} \sigma_{\omega_i} \eta_{\omega_i,t} + \sigma_{\zeta_h} \eta_{\zeta_h,t}$
Implicit forecasts	$\hat{y}_{t\|t-h} = \tilde{y}_t + \sigma_{\zeta_h} \eta_{\zeta_h,t}$
Rational forecast errors	$e_{t\|t-h} = \sum_{i=0}^{h-1} \sigma_{\omega_i} \eta_{\omega_i,t}$
Rational+implicit forecast errors	$e_{t\|t-h} = \sum_{i=0}^{h-1} \sigma_{\omega_i} \eta_{\omega_i,t} - \sigma_{\zeta_h} \eta_{\zeta_h,t}$
Implicit forecast errors	$e_{t\|t-h} = -\sigma_{\zeta_h} \eta_{\zeta_h,t}$
Rational forecast revisions between $h = m$ and $h = s$	$d_{t\|s,m} = \sum_{i=s}^{m-1} \sigma_{\omega_i} \eta_{\omega_i,t}$
Rational forecast revisions between $h = l$ and $h = s$	$d_{t\|s,l} = \sum_{i=s}^{l-1} \sigma_{\omega_i} \eta_{\omega_i,t}$
Rational+implicit forecast revisions between $h = m$ and $h = s$	$d_{t\|s,m} = \sum_{i=s}^{m-1} \sigma_{\omega_i} \eta_{\omega_i,t} + \sigma_{\zeta_s} \eta_{\zeta_s,t} - \sigma_{\zeta_m} \eta_{\zeta_m,t}$
Rational+implicit forecast revisions between $h = l$ and $h = s$	$d_{t\|s,l} = \sum_{i=s}^{l-1} \sigma_{\omega_i} \eta_{\omega_i,t} + \sigma_{\zeta_s} \eta_{\zeta_s,t} - \sigma_{\zeta_m} \eta_{\zeta_m,t}$
Implicit forecast revisions between $h = m$ and $h = s$	$d_{t\|s,m} = \sigma_{\zeta_s} \eta_{\zeta_s,t} - \sigma_{\zeta_m} \eta_{\zeta_m,t}$
Implicit forecast revisions between $h = l$ and $h = s$	$d_{t\|s,l} = \sigma_{\zeta_s} \eta_{\zeta_s,t} - \sigma_{\zeta_l} \eta_{\zeta_l,t}$

Notes: We use h to represent a generic forecast horizon and $h = H, H-1, ..., l, l-1, ..., m, m-1, ...s, s-1, ..., 0$, where H is the longest forecast horizon, and horizons $l \geq m \geq s$. For rational-implicit forecasts $\hat{y}_{t|t-h}$, the rational error component $\nu_{t|t-h} = -\sigma_{i=0}^{h-1} \sigma_{\omega_i} \eta_{\omega_i,t}$ where $\eta_{\omega_i,t} \sim i.i.d.N(0,1)$, and the implicit error component $\zeta_{t|t-h} = \sigma_{\zeta_h} \eta_{\zeta_h,t}$ where $\eta_{\zeta_h,t} \sim i.i.d.N(0,1)$.

A.1 Rational Forecasts

A.1.1 Proof for the monotonicity of the variances of the forecasts, forecast errors and revisions

Given that rational forecasts $\hat{y}_{t|t-h} = \tilde{y}_t - \sum_{i=0}^{h-1} \sigma_{\omega_i} \eta_{\omega_i,t}$, the variance of the target is $var(\tilde{y}_t) = var(\hat{y}_{t|t-h}) + \sum_{i=0}^{h-1} \sigma_{\omega_i}^2$. For two forecast horizons $h = l$ and $h = s$, $var(\tilde{y}_t) = var(\hat{y}_{t|t-l}) + \sum_{i=0}^{l-1} \sigma_{\omega_i} = var(\hat{y}_{t|t-s}) + \sum_{i=0}^{s-1} \sigma_{\omega_i}$; thus, $var(\hat{y}_{t|t-s}) - var(\hat{y}_{t|t-l}) = \sum_{i=s}^{l-1} \sigma_{\omega_i} \geq 0$. Note that since the mean squared forecasts $MSF_{t|t-h} = var(\hat{y}_{t|t-h}) + (E[\tilde{y}_t])^2$, the MSF is non-increasing as the forecast horizon increases, that is $MSF_{t|t-s} \geq MSF_{t|t-l}$.

Given the expression of the rational forecast errors in Table 6, $var(e_{t|t-s}) - var(e_{t|t-l}) = -\sum_{i=s}^{l-1} \sigma_{\omega_i} \leq 0$, and the mean squared forecast errors $MSFE_{t|t-s} \leq MSFE_{t|t-l}$.

Using the expressions of the rational forecast revisions, we have $var(d_{t|s,m}) = \sum_{i=s}^{m-1} \sigma_{\omega_i} \leq var(d_{t|s,l}) = \sum_{i=s}^{l-1} \sigma_{\omega_i}$ for $m \leq l$. As $E[d_{t|t-h}] = 0$, the mean squared forecast revisions $MSFR_{t|s,m} \leq MSFR_{t|s,l}$.

A.1.2 Proof for the monotonicity of the covariance between the forecasts, the forecasts and the target and the forecast errors and the target

The covariances between the forecasts for the same target at different horizons can be written as $cov(\hat{y}_{t|t-m}, \hat{y}_{t|t-s}) = cov(\hat{y}_{t|t-m}, \hat{y}_{t|t-m} + d_{t|s,m}) = var(\hat{y}_{t|t-m})$ since $cov(\hat{y}_{t|t-m}, d_{t|s,m}) = 0$. Similarly, we derive $cov(\hat{y}_{t|t-l}, \hat{y}_{t|t-s}) = var(\hat{y}_{t|t-l})$. Therefore, $cov(\hat{y}_{t|t-m}, \hat{y}_{t|t-s}) \geq cov(\hat{y}_{t|t-l}, \hat{y}_{t|t-s})$.

The covariances between the forecasts and the target $cov(\tilde{y}_t, \hat{y}_{t|t-h}) = cov(\hat{y}_{t|t-h} + \sum_{i=0}^{h-1} \sigma_{\omega_i} \eta_{\omega_i,t}, \hat{y}_{t|t-h}) = var(\hat{y}_{t|t-h})$. Since $var(\hat{y}_{t|t-h})$ weakly decrease as h increases, $cov(\tilde{y}_t, \hat{y}_{t|t-s}) \geq cov(\tilde{y}_t, \hat{y}_{t|t-l})$.

The covariances between the forecast errors and the target $cov(\tilde{y}_t, e_{t|t-h}) = cov(\hat{y}_{t|t-h} + e_{t|t-h}, e_{t|t-h}) = var(e_{t|t-h}) = \sum_{i=0}^{h-1} \sigma_{\omega_i}^2$. Therefore, $cov(\tilde{y}_t, e_{t|t-s}) \leq cov(\tilde{y}_t, e_{t|t-l})$.

A.1.3 Proof for the monotonicity of the covariances with forecast revisions

We first check the monotonicity of the covariances between the revised forecasts at a short horizon and the amount of forecast revision, i.e., $cov(\hat{y}_{t|t-s}, d_{t|s,h})$. For $h = m$, $cov(\hat{y}_{t|t-s}, d_{t|s,m}) = cov(\hat{y}_{t|t-m} + d_{t|s,m}, d_{t|s,m}) = var(d_{t|s,m})$ since $cov(\hat{y}_{t|t-m}, d_{t|s,m}) = 0$. Similarly, for $h = l$, $cov(\hat{y}_{t|t-s}, d_{t|s,l}) = cov(\hat{y}_{t|t-l} + d_{t|s,l}, d_{t|s,l}) = var(d_{t|s,l})$. Because $var(d_{t|s,m}) \leq var(d_{t|s,l})$, $cov(\hat{y}_{t|t-s}, d_{t|s,m}) \leq cov(\hat{y}_{t|t-s}, d_{t|s,l})$.

Given the expressions for rational forecasting errors and rational forecasting revisions provided in Table 6 , $cov(e_{t|t-m}, d_{t|s,m}) = \sum_{i=s}^{m-1} \sigma^2_{\omega_i}$ and $cov(e_{t|t-l}, d_{t|s,l}) = \sum_{i=s}^{l-1} \sigma^2_{\omega_i}$. Therefore, $cov(e_{t|t-m}, d_{t|s,m}) \leq cov(e_{t|t-l}, d_{t|s,l})$.

To prove the covariance bound, i.e., $var(d_{t|s,l}) \leq 2cov(\tilde{y}_t, d_{t|s,l})$, we follow Patton and Timmermann (2012), starting from $var(e_{t|t-l}) \geq var(e_{t|t-s})$. We then have $var(\tilde{y}_t - \hat{y}_{t|t-l}) \geq var(\tilde{y}_t - \hat{y}_{t|t-s})$, and

$$\begin{aligned}
var(\tilde{y}_t) - 2cov(\tilde{y}_t, \hat{y}_{t|t-l}) + var(\hat{y}_{t|t-l}) &\geq var(\tilde{y}_t) - 2cov(\tilde{y}_t, \hat{y}_{t|t-s}) + var(\hat{y}_{t|t-s}) \\
-2cov(\tilde{y}_t, \hat{y}_{t|t-l}) + var(\hat{y}_{t|t-l}) &\geq -2cov(\tilde{y}_t, \hat{y}_{t|t-l} + d_{t|s,l}) + var(\hat{y}_{t|t-l} + d_{t|s,l}) \\
-2cov(\tilde{y}_t, \hat{y}_{t|t-l}) + var(\hat{y}_{t|t-l}) &\geq -2cov(\tilde{y}_t, \hat{y}_{t|t-l}) - 2cov(\tilde{y}_t, d_{t|s,l}) + var(\hat{y}_{t|t-l}) + var(d_{t|s,l}) \\
var(d_{t|s,l}) &\leq 2cov(\tilde{y}_t, d_{t|s,l}).
\end{aligned}$$

We can write $cov(\hat{y}_{t|t-s}, d_{t|m,l})$ as $cov(\tilde{y}_t - \sum_{i=0}^{s-1} \sigma^2_{\omega_i}, d_{t|m,l}) = cov(\tilde{y}_t, d_{t|m,l})$ since $cov(-e_{t|t-s}, d_{t|m,l}) = 0$. Therefore, based on the covariance bound proved above, $var(d_{t|m,l}) \leq 2cov(\tilde{y}_t, d_{t|m,l}) = 2cov(\hat{y}_{t|t-s}, d_{t|m,l})$.

A.2 Rational and Implicit Forecasts

A.2.1 Patterns in the variances of the forecasts, forecast errors and revisions

The variance of rational and implicit forecasts is given by $var(\hat{y}_{t|t-h}) = var(\tilde{y}_t) + var(\nu_{t|t-h}) + var(\zeta_{t|t-h}) + 2cov(\tilde{y}_t, \nu_{t|t-h})$. Since $cov(\tilde{y}_t, \nu_{t|t-h}) = -var(\nu_{t|t-h})$, assuming that $cov(\tilde{y}_{t-H}, \nu_{t|t-h}) =$

0, $var(\hat{y}_{t|t-h}) = var(\tilde{y}_t) - var(\nu_{t|t-h}) + var(\zeta_{t|t-h})$. When $h = s$, $var(\hat{y}_{t|t-s}) = var(\tilde{y}_t) - \sum_{i=0}^{s-1}\sigma^2_{\omega_i} + \sigma^2_{\zeta_s}$, and when $h = l$, $var(\hat{y}_{t|t-l}) = var(\tilde{y}_t) - \sum_{i=0}^{l-1}\sigma^2_{\omega_i} + \sigma^2_{\zeta_l}$. Therefore, $var(\hat{y}_{t|t-s}) - var(\hat{y}_{t|t-l}) = \sum_{i=s}^{l-1}\sigma^2_{\omega_i} + \sigma^2_{\zeta_s} - \sigma^2_{\zeta_l}$. If and only if $\sum_{i=s}^{l-1}\sigma^2_{\omega_i} \geq \sigma^2_{\zeta_l} - \sigma^2_{\zeta_s}$, we have $var(\hat{y}_{t|t-s}) \geq var(\hat{y}_{t|t-l})$, and $MSF_{t|t-s} \geq MSF_{t|t-l}$.

Using the expression of the rational+implicit forecast errors in Table 6, the variances of forecasting errors $var(e_{t|t-h}) = \sum_{i=0}^{i=h-1}\sigma^2_{\omega_i} + \sigma^2_{\zeta_h}$. For two forecasts made at horizons s and l, $var(e_{t|t-s}) - var(e_{t|t-l}) = -\sum_{i=s}^{l-1}\sigma^2_{\omega_i} + \sigma^2_{\zeta_s} - \sigma^2_{\zeta_l}$. Therefore, if and only if $\sum_{i=s}^{l-1}\sigma^2_{\omega_i} \geq \sigma^2_{\zeta_s} - \sigma^2_{\zeta_l}$, $var(e_{t|t-s}) \leq var(e_{t|t-l})$, and $MSFE_{t|t-s} \leq MSFE_{t|t-l}$.

Using the expressions of the rational+implicit forecast revisions in Table 6, $var(d_{t|s,m}) = \sum_{i=s}^{m-1}\sigma^2_{\omega_i} + \sigma^2_{\zeta_s} + \sigma^2_{\zeta_m}$, and $var(d_{t|s,l}) = \sum_{i=s}^{l-1}\sigma^2_{\omega_i} + \sigma^2_{\zeta_s} + \sigma^2_{\zeta_l}$. The differences in two variances $var(d_{t|s,m}) - var(d_{t|s,l}) = -\sum_{i=m}^{l-1}\sigma^2_{\omega_i} + \sigma^2_{\zeta_m} - \sigma^2_{\zeta_l}$. Therefore, if and only if $\sum_{i=m}^{l-1}\sigma^2_{\omega_i} \geq \sigma^2_{\zeta_m} - \sigma^2_{\zeta_l}$, $var(d_{t|s,m}) \leq var(d_{t|s,l})$, and $MSFR_{t|t-s} \leq MSFE_{t|t-l}$.

A.2.2 Patterns in the covariances between the forecasts, the forecasts and the target and the forecast errors and the target

The covariances between two forecasts $cov(\hat{y}_{t|t-m}, \hat{y}_{t|t-s}) = cov(\hat{y}_{t|t-m}, \hat{y}_{t|t-m} + d_{t|s,m}) = var(\hat{y}_{t|t-m}) + cov(\hat{y}_{t|t-m}, d_{t|s,m})$. Using the expressions in Table 6, we have $cov(\hat{y}_{t|t-m}, d_{t|s,m}) = -\sigma^2_{\zeta_m}$. Therefore, $cov(\hat{y}_{t|t-m}, \hat{y}_{t|t-s}) = var(\hat{y}_{t|t-m}) - \sigma^2_{\zeta_m}$. Similarly, $cov(\hat{y}_{t|t-l}, \hat{y}_{t|t-s}) = var(\hat{y}_{t|t-l}) - \sigma^2_{\zeta_l}$. Comparing with these two covariances, $cov(\hat{y}_{t|t-m}, \hat{y}_{t|t-s}) - cov(\hat{y}_{t|t-l}, \hat{y}_{t|t-s}) = var(\hat{y}_{t|t-m}) - \sigma^2_{\zeta_m} - var(\hat{y}_{t|t-l}) + \sigma^2_{\zeta_l}$. Using the derivation of the variances of forecasts in section A.2.1, we then have $cov(\hat{y}_{t|t-m}, \hat{y}_{t|t-s}) - cov(\hat{y}_{t|t-l}, \hat{y}_{t|t-s}) = -\sum_{i=0}^{m-1}\sigma^2_{\omega_i} + \sum_{i=0}^{l-1}\sigma^2_{\omega_i} = \sum_{i=m}^{l-1}\sigma^2_{\omega_i} \geq 0$.

The covariances between the forecasts $\hat{y}_{t|t-h}$ and the target $\tilde{y}_t$ are given by $cov(\tilde{y}_t, \hat{y}_{t|t-h}) = var(\tilde{y}_t) + cov(\tilde{y}_t, \nu_{t|t-h}) = var(\tilde{y}_t) - var(\nu_{t|t-h}) = var(\tilde{y}_t) - \sum_{i=0}^{h-1}\sigma^2_{\omega_i}$. As h increases, $cov(\tilde{y}_t, \hat{y}_{t|t-h})$ weakly decreases, and $cov(\tilde{y}_t, \hat{y}_{t|t-s}) \geq cov(\tilde{y}_t, \hat{y}_{t|t-l})$.

The covariances between the target and forecasting errors $cov(\tilde{y}_t, e_{t|t-h}) = \sum_{i=0}^{h-1}\sigma^2_{\omega_i}$.

Therefore, $cov(\tilde{y}_t, e_{t|t-s}) \leq cov(\tilde{y}_t, e_{t|t-l})$.

A.2.3 Patterns in the covariances of the forecast revisions

We first derive $cov(\hat{y}_{t|t-s}, d_{t|s,m})$ using the expression of $d_{t|s,m}$ in Table 6.

$$\begin{aligned} cov(\hat{y}_{t|t-s}, d_{t|s,m}) &= cov(\hat{y}_{t|t-m} + d_{t|s,m}, d_{t|s,m}) \\ &= cov(\hat{y}_{t|t-m}, d_{t|s,m}) + var(d_{t|s,l}) \\ &= -var(\zeta_{t|t-m}) + var(d_{t|s,m}) \\ &= -\sigma^2_{\zeta_m} + \textstyle\sum_{i=s}^{m-1} \sigma^2_{\omega_i} + \sigma^2_{\zeta_s} + \sigma^2_{\zeta_m} \\ &= \textstyle\sum_{i=s}^{m-1} \sigma^2_{\omega_i} + \sigma^2_{\zeta_s}. \end{aligned}$$

Similarly, $cov(\hat{y}_{t|t-s}, d_{t|s,l}) = \sum_{i=s}^{l-1} \sigma^2_{\omega_i} + \sigma^2_{\zeta_s}$. Therefore, $cov(\hat{y}_{t|t-s}, d_{t|s,m}) \leq cov(\hat{y}_{t|t-s}, d_{t|s,l})$.

When $h = m$, the covariances between forecast errors and forecasting revisions $cov(e_{t|t-m}, d_{t|s,m}) = cov(\sum_{i=0}^{m-1} \sigma_{\omega_i}\eta_{\omega_i,t} - \sigma_{\zeta_m}\eta_{\zeta_m,t}, \sum_{i=s}^{m-1} \sigma_{\omega_i}\eta_{\omega_i,t} + \sigma_{\zeta_s}\eta_{\zeta_s,t} - \sigma_{\zeta_m}\eta_{\zeta_m,t}) = \sum_{i=s}^{m-1} \sigma^2_{\omega_i} + \sigma^2_{\zeta_m}$. Additionally, $cov(e_{t|t-l}, d_{t|s,l}) = \sum_{i=s}^{l-1} \sigma^2_{\omega_i} + \sigma^2_{\zeta_l}$. Therefore, $cov(e_{t|t-m}, d_{t|s,m}) - cov(e_{t|t-l}, d_{t|s,l}) = -\sum_{i=m}^{l-1} \sigma^2_{\omega_i} + \sigma^2_{\zeta_m} - \sigma^2_{\zeta_l}$, and if and only if $\sum_{i=m}^{l-1} \sigma^2_{\omega_i} \geq \sigma^2_{\zeta_m} - \sigma^2_{\zeta_l}$, we have $cov(e_{t|t-m}, d_{t|s,m}) \leq cov(e_{t|t-l}, d_{t|s,l})$.

The variance of rational and implicit forecast revisions between $h = s$ and $h = l$ is $var(d_{t|s,l}) = \sum_{i=s}^{l-1} \sigma^2_{\omega_i} + \sigma^2_{\zeta_l} + \sigma^2_{\zeta_s}$. The covariance between the target and $d_{t|s,l}$ is $cov(\tilde{y}_t, d_{t|s,l}) = \sum_{i=s}^{l-1} \sigma^2_{\omega_i}$. Therefore, if and only if $\sum_{i=s}^{l-1} \sigma^2_{\omega_i} \geq \sigma^2_{\zeta_l} + \sigma^2_{\zeta_s}$, the covariance bound of Patten and Timmermann (2012) holds, that is, $var(d_{t|s,l}) \leq 2cov(\tilde{y}_t, d_{t|s,l})$.

A.3 Implicit Forecasts

A.3.1 Patterns in the variances of the forecasts, forecast errors and revisions

Given that implicit forecasts $\hat{y}_{t|t-h} = \tilde{y}_t + \zeta_{t|t-h}$ and $cov(\tilde{y}_t, \zeta_{t|t-h}) = 0$, we have $var(\hat{y}_{t[|t-h}) = var(\tilde{y}_t) + \sigma^2_{\zeta_h}$. Therefore, if and only if $\sigma_{\zeta_s} \geq \sigma_{\zeta_l}$, then $var(\hat{y}_{t|t-s}) \geq var(\hat{y}_{t|t-l})$ and $MSF_{t|t-s} \geq$

$MSF_{t|t-l}$.

The variance of the implicit forecast errors is $var(e_{t|t-h}) = \sigma^2_{\zeta_h}$, and if and only if $\sigma_{\zeta_s} \leq \sigma_{\zeta_l}$, we will have $var(e_{t|t-s}) \leq var(e_{t|t-l})$ and $MSFE_{t|t-s} \leq MSFE_{t|t-l}$.

To check any monotonicity pattern in the mean squared forecasting revisions, we compare $var(d_{t|s,m})$ and $var(d_{t|s,l})$. The variance of the implicit forecast revisions $var(d_{t|s,m}) = \sigma^2_{\zeta_s} + \sigma^2_{\zeta_m}$ and $var(d_{t|s,l}) = \sigma^2_{\zeta_s} + \sigma^2_{\zeta_l}$. Therefore, if and only if $\sigma_{\zeta_m} \leq \sigma_{\zeta_l}$, $var(d_{t|s,m}) \leq var(d_{t|s,l})$ and $MSFR_{t|s,m} \leq MSFR_{t|s,l}$.

A.3.2 Patterns in the covariances between the forecasts, the forecasts and the target and the forecast errors and the target

Since $cov(\tilde{y}_t, \zeta_{t|t-h}) = 0$ and implicit forecast errors are uncorrelated across horizons, we have $cov(\hat{y}_{t|t-m}, \hat{y}_{t|t-s}) = cov(\hat{y}_{t|t-l}, \hat{y}_{t|t-s}) = var(\tilde{y}_t)$, $cov(\tilde{y}_t, e_{t|t-s}) = cov(\tilde{y}_t, e_{t|t-l}) = 0$, and $cov(\tilde{y}_t, \hat{y}_{t|t-s}) = cov(\tilde{y}_t, \hat{y}_{t|t-l}) = var(\tilde{y}_t)$.

A.3.3 Patterns in the covariances of the forecast revisions

Now we derive for $cov(e_{t|t-m}, d_{t|s,m})$ and $cov(e_{t|t-l}, d_{t|s,l})$. Using the expressions in Table 6, we have $cov(e_{t|t-m}, d_{t|s,m}) = cov(-\zeta_{t|t-m}, \zeta_{t|t-s} - \zeta_{t|t-m}) = var(\zeta_{t|t-m}) = \sigma^2_{\zeta_m}$, and similarly $cov(e_{t|t-l}, d_{t|s,l}) = \sigma^2_{\zeta_l}$. Therefore, $cov(e_{t|t-m}, d_{t|s,m}) \leq cov(e_{t|t-l}, d_{t|s,l})$ if and only if $\sigma_{\zeta_m} \leq \sigma_{\zeta_l}$.

To see the pattern in the covariances of short-horizon forecasts $\hat{y}_{t|t-s}$ and forecast revisions $d_{t|s,h}$, we first let $h = m$, and then $cov(\hat{y}_{t|t-s}, d_{t|s,m}) = cov(\tilde{y}_t + \zeta_{t|t-s}, \zeta_{t|t-s} - \zeta_{t|t-m}) = var(\zeta_s) = \sigma^2_{\zeta_s}$. Letting $h = l$, $cov(\hat{y}_{t|t-s}, d_{t|s,l}) = \sigma^2_{\zeta_s}$. We see that these covariances between forecasts at short horizons and revisions are always constant for all forecast horizons.

We now prove that the upper variance bound for revisions as a function of covariances between the target and revisions does not hold for implicit multi-horizon forecasts. Since $var(d_{t|s,l}) = var(\zeta_{t|t-s} - \zeta_{t|t-l}) = \sigma^2_{\zeta,s} + \sigma^2_{\zeta,l}$ and $cov(\tilde{y}_t, d_{t|s,l}) = 0$, we have $var(d_{t|s,l}) > 2cov(\tilde{y}_t, d_{t|s,l})$ if $\sigma_{\zeta_h} \neq 0$.

Lastly, $var(d_{t|m,l}) = \sigma^2_{\zeta_m} + \sigma^2_{\zeta_l}$, and $cov(\hat{y}_{t|t-s}, d_{t|m,l}) = cov(\tilde{y}_t + \zeta_{t|t-s}, \zeta_{t|t-m} - \zeta_{t|t-l}) = 0$. Therefore, $var(d_{t|m,l}) > 2cov(\hat{y}_{t|t-s}, d_{t|m,l})$ if $\sigma_{\zeta_h} \neq 0$.