LCERPA Working paper No. 2023-6

December 2023

Seize the Last Day: Period-End-Point Sampling for Forecasts of Temporally Aggregated Data

Reinhard Ellwanger
Bank of Canada

Stephen Snudden
Wilfrid Laurier University

Lenin Arango-Castillo
Bank of Canada

Seize the Last Day: Period-End-Point Sampling for Forecasts of Temporally Aggregated Data*

Reinhard Ellwanger
Bank of Canada

Stephen Snudden[‡]
Wilfrid Laurier University

Lenin Arango-Castillo
Bank of Mexico

First Version: Dec. 2023
Updated: Dec. 2024

Abstract

Forecasts of temporal aggregates, such as monthly or quarterly averages of daily data, are often constructed using the aggregated data due to difficulties in disaggregating estimation. However, we show that when the daily data is persistent, forecasts constructed with aggregated data are inefficient and can nearly double forecast error. We propose a new forecasting technique, Period-End-Point Sampling (PEPS), that corrects for the information loss, and allows models to maintain the lower frequency of the forecast target. PEPS uses end-of-period data to construct point-in-time forecasts equal to the period average forecast. Applications to real-time forecasts of 10-year bond yields and the price of copper show that PEPS rivals the accuracy of bottom-up approaches.

JEL classification: C1, C53, E47, F37, Q47
Keywords: Temporal Aggregation, Forecasting Methods, Interest Rates, Commodity Prices

*The views expressed in this paper are those of the authors; no responsibility for them should be attributed to the Bank of Canada or the Bank of Mexico.

[†]Corresponding author. Email: ssnudden@wlu.ca

[‡]We thank Lutz Kilian, Ron Alquist, Christiane Baumeister, Edouard Djeutem, Konrad Zmitrowicz, Pierre Siklos, Wing Chan, Willy Chetwin and Luis Uzeda, as well as seminar and conference participants at the Bank of Canada, the Bank of Mexico, the European Economic Association Annual Congress, and the International Symposium on Forecasting for their detailed comments and suggestions. This research was supported by the Social Sciences and Humanities Research Council (SSHRC) grant 430-2020-01202. This work supersedes analysis in the paper previously distributed with the title "Predictability of Temporally Aggregated Real Series."

1 Introduction

Time series are often temporally aggregated, that is, averages or sums of higher-frequency data, such as monthly or quarterly averages of daily data. Forecasts of such series, especially when expressed in real terms, play a key role in expectation formation, central bank projections, and investment decisions.[1] Unfortunately, temporal aggregation may introduce a loss of information contained in higher-frequency data, making forecasts constructed with temporally aggregated data inefficient (Zellner and Montmarquette, 1971; Amemiya and Wu, 1972; Tiao, 1972; Wei, 1978; Kohn, 1982; Lütkepohl, 1986). The main approach to avoid this information loss is the bottom-up (BU) approach, which consists of computing forecasts for the underlying high-frequency data and averaging the forecasts ex-post (Zellner and Montmarquette, 1971; Lütkepohl, 1986).[2] However, in practice, forecast and projection models are commonly implemented with averaged data due to the challenges of altering the frequency of the forecast model.

In this paper, we propose a general method for forecasting temporally aggregated data with disaggregated observations that allows the forecast model to maintain the same frequency as the target variable. This method, which we call Period-End-Point Sampling (PEPS), consists of adjusting point forecasts constructed with end-of-period observations so that they are equal to the period average forecast.

A major advantage of PEPS is that it allows forecasters to maintain all variables within the same lower frequency as the forecasted series. This is often desirable in structural and multivariate models, which include other variables that are observed only at a lower frequency. In contrast, existing approaches such as the bottom-up approach or mixed-frequency techniques require introducing the higher-frequencies into the forecast model. PEPS is straightforward to implement and, in the simplest case, only requires substituting aggregated observations for end-of-period observations when constructing forecasts.

Using simulation analysis, we find that both the PEPS and BU approaches substantially improve on the forecasts constructed with averaged data, with up to 46% improvements in the mean-squared forecast error (MSFE) and directional accuracy at the one-step-ahead prediction. The

[1]For example, forecasts of average prices might be necessary to accurately predict total costs or revenues. More generally, average observations more closely reflect average economic conditions over a certain period and are common in forecasts of real effective exchange rates and quarterly energy prices (see, e.g., Christiano and Eichenbaum, 1987; Ellwanger and Snudden, 2023b).

[2]An alternative approach relying on information from higher frequency data is Mixed Data Sampling (MIDAS, see, e.g. Ghysels et al., 2007; Andreou et al., 2013). The PEPS can be applied in this context, but herein our primary focus is on recursive forecasts, for which the bottom-up approach is the principle existing method.

PEPS forecasts rival the efficiency of bottom up forecasts and work particularly well for monthly and quarterly data.

In our empirical application, we examine the real-time performance of alternative forecasts methods for the nominal yield on 10-year U.S. bonds and the real price of copper. For all series, we find large and robust improvements in forecast accuracy from employing disaggregated approaches at short horizons. For the one-month-ahead forecast, the MSFEs of the disaggagregated approaches are up to 45% lower than for forecasts constructed with averaged data, which is unprecedented in existing forecast applications for these series. Confirming our simulation results, the PEPS forecasts perform similarly at short horizons and can do better at longer horizons compared to the BU forecasts. A practical advantage of PEPS is that the data can be backcasted, which can result in additional forecast improvements at longer forecast horizons.

A key contribution of our analysis is to examine temporal aggregation of daily data. Prior investigations comparing BU and aggregated approaches have focused on aggregation from monthly to quarterly or from quarterly to yearly data (see, e.g., Zellner and Montmarquette, 1971; Wei, 1978; Lütkepohl, 1986; Athanasopoulos et al., 2011). In our simulations and empirical applications, the loss in forecast accuracy resulting from daily to monthly aggregation is much greater than the loss resulting from monthly to quarterly aggregation. This occurs because daily-to-monthly aggregation moves from a state of no aggregation to a substantive aggregation (typically 21 business days). In contrast, monthly to quarterly aggregations represent additional aggregations of already aggregated series. As such, using information from daily data results in forecast gains that are much larger than currently understood.

Our results complement a related body of existing work on the effect of aggregation on forecasts and the dynamics of economic series (see, e.g., Working, 1960; Brewer, 1973; Weiss, 1984; Rossana and Seater, 1995; Marcellino, 1999). Theoretical results for ARIMA models show that forecasts constructed with aggregated time series are less efficient than forecasts constructed with the BU approach (Amemiya and Wu, 1972; Tiao, 1972; Wei, 1978; Lütkepohl, 1986). This holds even when accounting for the effect of temporal aggregation on the ARIMA structure (Kohn, 1982) and under parameter uncertainty (Lütkepohl, 1986). We extend these results by showing that suitably constructed point forecasts can correct for the information loss from temporal aggregation and can be as efficient as the BU approach in forecasting averaged series.

2 Equivalence of Point and Bottom-Up Forecasts

2.1 Point Sampling and Temporal Aggregation

Consider a continuous time series $y_{t,i}$ such that $t = 1, 2, \ldots, T$ indicates a lower-frequency period and $i \in [0, n]$ denote sub-periods within the period so that the last observation of period t is $y_{t,n}$. For example, this series may represent prices of a continuously traded asset within day i in month t.

In practice, it is common to point-sample and record daily data at a natural time unit, such as closing values at the end of each day. In this case, $y_{t,i}$ is now discretized such that i takes only integer values, $i = 1, 2, \ldots, n$, and the first observation of the next period is $y_{t+1,1}$.[3] In practice, daily closing prices are the basis of many period average macroeconomic series such exchange rates, interest rates, and commodity prices.

Consider a forecaster whose objective is to use the information available at the end of period T, to predict the k-periods-ahead average observation, $\bar{y}_{T+k}$. The period-t average is given by

$$\bar{y}_t = \frac{1}{n} \sum_{i=1}^{n} y_{t,i}. \tag{1}$$

Such period average forecasts are common place in macroeconomics and for investment decisions. For example, forecasts of average prices might be necessary to accurately predict total costs or revenues. More generally, average observations more closely reflect average economic conditions over a certain period and are common in forecasts of effective exchange rates and interest rates (see, e.g., Meese and Rogoff, 1983; Christiano and Eichenbaum, 1987).

Another way to reduce the frequency of daily observations is point sampling at the period t frequency. This may consist of recording observations from a specific day i of each month. A point-sampled series particularly relevant is the series of end-of-period observations, $y_{1,n}, y_{2,n}, \ldots, y_{T,n}$. For example, asset prices are commonly reported as end of period values, and bilateral exchange rates are commonly reported as both period average and end-of-period values. In effect, a series of end-of-period values are observations of daily closing prices that have been sampled by skipping every $n - 1$ daily observations.

The form of sampling plays a subtle but critical role in determining the stochastic properties of the resulting series. In the previous literature, this role has been investigated separately for point sampling and temporal aggregation. In this paper, we exploit the advantages of selective sampling

[3] This form of sampling has also been called *skip sampling, systematic sampling, point-in-time sampling,* or *selective sampling.*

while maintaining the goal of forecasting period averages.

The first advantage is that relative to period averaging, period-end point sampling tends to be less distortionary for the original data-generating process. For (V)AR processes, averaging always introduces an additional MA-component into the process, while this is not necessarily the case for point sampling (Brewer, 1973; Weiss, 1984; Wei, 1981; Marcellino, 1999). For example, when the underlying series is generated by a random walk, averaging produces a moving average process (Working, 1960), while point sampling maintains the random-walk properties at the lower frequency.

The second advantage, which accrues to point sampling of end-of-period observations, is that the last observed instantaneous (non-averaged) observation can reflect the full information available to economic agents at the time the forecast is formed. This property is particularly relevant if the underlying series has a forward-looking component, such as prices of storable goods and assets, as famously argued by Fama (1970). Averaging dilutes the latest information contained in the last observed observation by averaging over past observations containing stale information. As a result, computations of forecasts with averaged data will generally result in a deterioration in forecast accuracy, which is not necessarily the case for forecasts computed with point-sampled data (Kohn, 1982).

2.2 Optimal Point Forecasts

The main approach to avoid the information loss in forecasting is the bottom-up (BU) approach, which consists of computing forecasts for the point sampled daily data and averaging the forecasts ex-post to the desired period frequency (Zellner and Montmarquette, 1971; Lütkepohl, 1986).

More formally, let $\hat{\bar{y}}_{T+k}$, $k = 1, 2, \ldots, K$, be the period average forecast of $\bar{y}_{T+k}$. Let the model-implied conditional expectations at the end of period T be denoted by $\mathbf{E}_{T,n}[\cdot]$. The BU forecast constructs a point-in-time forecast, $\hat{y}_{T+k,i}$, for all days i in forecast period $T + k$. Then the period-average forecast is given by the simple average of the daily forecasts. For example, for the k-period-ahead forecast

$$\mathbf{E}_{T,n}[\bar{y}_{T+k}] = \frac{1}{n}\sum_{i=1}^{n}\hat{y}_{T+k,i}. \tag{2}$$

We now want to understand the relationship between point forecasts of i within t and the bottom up forecast, $\hat{\bar{y}}_{T+k}$.

Theorem 1. *Let $\hat{y}_{T+k,i}$ be a continuous function on an interval $i \in [0, n]$, observed when i takes* *on integer values, $i = 1, 2, \ldots, n$. For $n \geq 2$, there always exists at least one point-in-time i^* in $T+k$ such that there is a point forecast, $\hat{y}_{T+k,i^*}$, that is equal to the bottom-up forecast $\hat{\bar{y}}_{T+k}$.*

Proof. Assume without loss of generality that $\hat{y}_{T+k,1}$ and $\hat{y}_{T+k,n}$ are not both the maximum or both the minimum of $\hat{y}_{T+k,i}$. If they are, the series is constant, and the average is equal to $\hat{y}_{T+k,i}$ for every value of i. There exists at least one maximum value $\hat{y}_{T+k,i^{max}}$ and one minimum value $\hat{y}_{T+k,i^{min}}$ within the time series, and by definition $\hat{y}_{T+k,i^{min}} \leq \hat{\bar{y}}_{T+k} \leq \hat{y}_{T+k,i^{max}}$. By the Intermediate Value Theorem (IVT) there exists at least one point i^* between $\hat{y}_{T+k,1}$ and $\hat{y}_{T+k,n}$ for which $\hat{y}_{T+k,i^*} = \hat{\bar{y}}_{T+k}$. See IVT in appendix A1.1. $\square$

Consider some implications of Theorem 1. First, there always exists at least one point forecast in $T+k$ that is equal to the forecast of the average. Forecasting this point-in-time is equivalent to forecasting the period average. Second, there may exist more than one point forecast within a period that is equal to the period average forecast. Finally, even if the bottom-up forecasts consist of point sampled daily data, the point i^* does not need to be discrete for the equivalence to hold.

We refer to the use of a point forecasts within t which are constructed with data that is sampled at the end of each period and equal to the forecast of the period average as Period-End-Point Sampling (PEPS).

Definition 1. *Definition of PEPS: Let $\hat{\bar{y}}_{T+k}$, $k = 1, 2, \ldots, K$, be the forecast of $\bar{y}_{T+k}$. We construct a point-in-time forecast for day i^* in forecast period $T+k$ where the point forecast is equal to the period average forecast in $T+k$,*

$$\hat{\bar{y}}_{T+k} = \mathbf{E}_{T,n}\left[y_{T+k,i}\right], \textit{ for } i^* \textit{ of period } T+k.$$

In practice, there are several possible ways to derive a point forecast that in expectation is equal to the period average forecast. Herein, we explore a few approaches, but the discovery of additional methods presents a promising avenue for further research.

The most obvious method is to utilize the known mappings between the sampled data and the underlying higher-frequency process. Existing works have provided the mapping between the point sampled end-of-period data and the underlying higher-frequency process. For example, Wei (1981) and Weiss (1984) provide the mapping between for AutoRegressive Integrated Moving-Average (ARIMA) processes and Marcellino (1999) does so for Vector AutoRegressive Integrated Moving-Average (VARIMA) processes. Tsai and Chan (2005) and Man and Tiao (2006) consider

AutoRegressive Fractionally Integrated Moving-Average (ARFIMA) processes. The objective is to recover the underlying process, and then use the mapping to construct a point forecast for i^* in $T+k$. For example, when the mapping is known, we can use observations $y_{1,n}, y_{2,n}, \ldots, y_{T,n}$ to estimate models using point sampled observations, which can then be used to construct $\hat{y}_{T+k,i^*}$ and forecast $\hat{\bar{y}}_{T+k}$. In this case, model-based forecasts are estimated with end-of-period observations to be consistent with the forecaster using available information at the end of the period. An example of this direct mapping approach as shown in the next sections.

For some processes, the end-of-period forecasts may naturally coincide with the period average forecasts. For example, the period average real price of crude oil was found to naturally coincide with the end-of-period forecast by Benmoussa et al. (2023). In addition, for stationary processes, the end of period and underlying forecasts are equal to the period average forecasts as they converge to the mean at longer forecast horizons (see for example, Telser, 1967; Brewer, 1973; Weiss, 1984; Marcellino, 1999). Moreover, as n becomes large, $\bar{y}_t$ becomes white noise, and the forecasts quickly approach the mean (Tiao, 1972; Stram and Wei, 1986). Both the example in section 2.3.3 and the simulations in section 3 quantify that this convergence can occur quickly. These are three cases where end-of-period and period average forecasts already coincide without any alterations needed.

Note that in cases where the daily data is observed, we can numerically derive estimates of i^*. This involves estimating the underlying process and calculating i^* directly. For example, a daily model can be estimated, used to construct forecasts $\hat{y}_{T+k,i}$, and y_{T+k,i^*} is then given where the bottom-up forecast $\bar{y}_{T+k}$ and the daily forecast $\hat{y}_{T+k,i}$ intersect. Once i^* is known, one could employ forecasts of this point in time as the period average forecast. For example, one could use market-based expectations of y_{T+k,i^*}, using financial derivatives such as futures contracts, to derive a forecast of the period average expectations (see e.g Farag et al., 2024).

Finally, it is possible to estimate the period average forecast through combinations of end-of-period forecasts. This can be analytically by again mapping between the end-of-period and period average forecast. This can also be done numerically by relying on the daily data or estimating the relation between end-of-period and period average data. For example, forecast weighting methods can be used to construct the period average forecast using period-end point forecasts. Consider a simple case of a continuous forecast within $\hat{y}_{T+k-1,n}$ and $\hat{y}_{T+k,n}$ for all i within $T+k$.

Claim 1. *For $n \geq 2$, if the period average lies within $\hat{y}_{T+k-1,n}$ and $\hat{y}_{T+k,n}$, then there exists is at least one weight, ω^*, such that a piecewise linear interpolation of adjacent end-of-period forecasts, $\hat{y}_{T+k-1,n}$ and $\hat{y}_{T+k,n}$, $\dot{y}_{T+k,i} = (1-\omega)\hat{y}_{T+k-1,n} + \omega\hat{y}_{T+k,n}$ is equal to the bottom-up forecast $\hat{\bar{y}}_{T+k}$*

at point $i^{\#}$.

This again follows from the IVT, see proof in appendix A1.2. This means that the period average forecast can be derived as an interpolation, $\dot{y}_{T+k,i}$, of the adjacent end-of-period forecasts $\hat{y}_{T+k,n}$. We refer to methods that interpolate end-of-period forecasts to forecast the average as Period-End-Point Sampled Interpolation (PEPSI).

Definition 2. *Definition of PEPSI: Let $\hat{\bar{y}}_{T+k}$, $k = 1, 2, \ldots, K$, be for forecast of $\bar{y}_{T+k}$. Using forecast averaging, we construct a point-in-time forecast for $i^{\#}$ in forecast period $T + k$ where the point forecast of $\dot{y}_{T+k,i}$ is equal to the period average forecast in $T + k$,*

$$\hat{\bar{y}}_{T+k} = \mathbf{E}_{T,n}\left[\dot{y}_{T+k,i}\right] \text{ , for } i^{\#} \text{ of period } T + k.$$

For a piecewise linear interpolation, the requirement that the period average lies within $\hat{y}_{T+k-1,n}$ and $\hat{y}_{T+k,n}$ is not particular demanding, and, for example, applies to any forecast exhibiting monotonicity. While the optimal weighting can be solved analytically for many stochastic processes, the weights can be solved numerically by analyzing the bottom-up forecasts if the observations exist, see section 2.3. Alternatively, the weight can be solved numerically by optimizing the weight using out-of-sample forecasts. Such numerical optimization is common place, and is akin to solving for parameter weights in forecast combination. In section 4 we illustrate how this can be implemented in practice.

Both PEPS and PEPSI provide alternative techniques to the bottom-up approach to construct period average forecasts. The PEPS approach can, and the PEPSI approach does, rely on the use of end-of-period forecasts to forecast the period average. The construction of end-of-period forecasts and the use of end-of-period data in estimation is straightforward and already common place in practice. In many applications, this is simple to implement and merely requires replacing period average with end of period data. This is particularly advantages in situations such as policy and multivariate models where it is impractical to introduce estimations at the daily frequency.

The use of end-of-period data in estimation also has potential practical advantages. First, while daily data has typically only been available since the 80's, monthly data is typically available over longer periods and can be used for backcasting. Second, because the values of autoregressive coefficients approach zero as n increases, the smaller autoregressive coefficients could lessen the downward bias in estimation when using end-of-period versus daily data (Tiao, 1972; Stram and Wei, 1986). Such gains would have to be weighed against losses from the use of fewer observations but could make PEPS even more efficient than the bottom-up approach given a sufficient sample size when the parameter values need to be estimated.

2.3 Intuition: The AR(1) Case

The easiest way to provide intuition for ability of point forecasts to be equal to period average point forecasts is to consider the case when the daily series is generated by a stationary[4] autoregressive model of order one, AR(1),

$$y_{t,i} = \rho y_{t,i-1} + \epsilon_{t,i}\,, \qquad \text{for } i = 1, 2, ..., n;\; t = 1, 2, ...T. \tag{3}$$

Here, n is the number of daily observations within a period, $|\rho| < 1$, and $\epsilon_{t,i} \sim iid(0, \sigma_\epsilon)$ denotes the daily innovation.

The goal is to construct a forecast of a temporally aggregated series, $\bar{y}_t$, which is obtained by averaging across $n \geq 2$ non-overlapping observations. Forecasts are constructed using the entire information set up to period T, including daily observations. We assume that ρ is known to abstract from estimation uncertainty.

2.3.1 PEPS: AR(1) case

In this setting, the forecast for the observation on day i in period $T + k$ is given by

$$\mathbf{E}_{T,n}\left[y_{T+k,i}\right] = \rho^{(k-1)n+i} y_{T,n}\,. \tag{4}$$

Under the BU approach, the period-average forecast is given by the simple average of the daily forecasts. For example, for the one-month-ahead forecast ($k = 1$)

$$\mathbf{E}_{T,n}\left[\bar{y}_{T+1}\right] = \frac{1}{n}\sum_{i=1}^{n} \rho^i y_{T,n}. \tag{5}$$

Figure 1 illustrates the equivalence of the point and period-average forecasts, for a monthly average ($n = 21$) of daily forecasts for an AR(1) model with $\rho = 0.95$ and $y_{T,n} = 1$. The monthly average forecast is constructed using the BU approach, and thus is a simple average over the daily forecasts.

[4] As is standard, we assume that stationarity can be achieved by differencing (Wei, 1978; Kohn, 1982; Lütkepohl, 1984, 1986, 2006).

Figure 1. Intersection of Point and Average Forecast

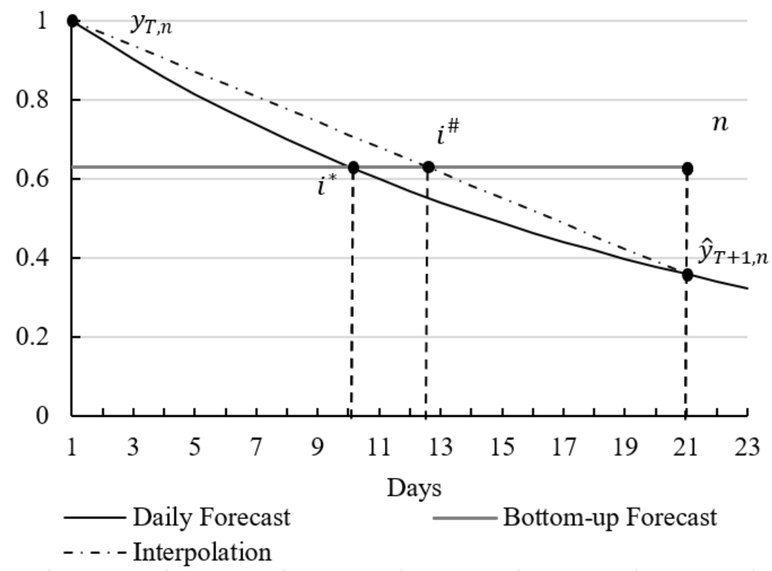

Note: Monthly forecast and monthly average for an AR(1) model with $\rho = 0.95$. Assumes $n = 21$ days in a month.

Since the daily forecast is monotonically declining, it must be that the monthly average forecast intercepts the daily forecast at some point-in-time i^*, with $0 < i^* < n$. For stationary AR process, arbitrary k, $0 < \rho < 1$, i^* can be found by solving

$$\rho^{(k-1)n+i^*} y_{T,n} = \frac{1}{n} \sum_{i=1}^{n} \rho^{(k-1)n+i} y_{T,n}, \tag{6}$$

which yields

$$i^* = \frac{\ln \left[\frac{\rho^{(k-1)n+1}(\rho^n - 1)}{n(\rho - 1)} \right]}{\ln(\rho)} - (k-1)n. \tag{7}$$

This is the closed form solution to determine a point forecast i^* within period t that is equal to the period-average forecasts. Figure 2 graphs the values of i^* for alternative values of ρ for arbitrary $n \geq 2$ for the one-period ahead forecast. It indicates that $1 < i^* < (n-1)/2$. That is, for low values of ρ, forecasts for a point within the beginning of the period will be closest to the period average forecasts, while for values of ρ close to one, values closer to the middle of the month coincide with period average forecasts.

Figure 2. Day i^* for which the Average and Point Forecasts Intercept

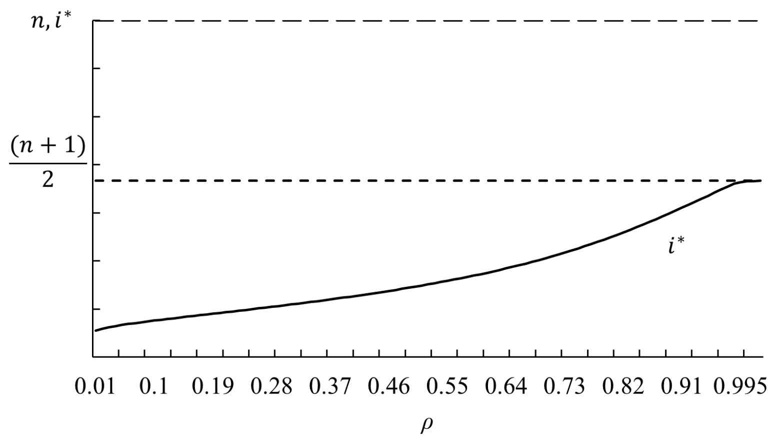

Note: Forecast of point i^* for which the point forecast is equal to the period average forecast. Calculations for a daily AR(1) process with $0 < \rho < 1$.

2.3.2 PEPSI: AR(1) case

Alternatively, as illustrated in Figure 1 the period average forecast can be derived as a weighted average of the adjacent end-of-period forecast. For arbitrary k this interpolation is given by:

$$\dot{y}_{T+k,i} = (1-\omega)\hat{y}_{T+k-1,n} + \omega\hat{y}_{T+k,n}. \tag{8}$$

In this case, the objective is to find ω^* such that a weighted average of the adjacent end-of-period forecasts, $\dot{y}_{T+k,i}$, is equal to the bottom-up forecast $\hat{\bar{y}}_{T+k}$ at some point $i^{\#}$ within t. For arbitrary k, ω^* can be found by solving

$$(1-\omega)\rho^{(k-1)n}y_{T,n} + \omega\rho^{kn}y_{T,n} = \frac{1}{n}\sum_{i=1}^{n}\rho^{(k-1)n+i}y_{T,n}, \tag{9}$$

which yields

$$\omega^* = \frac{\rho(\rho^n - 1) - n(\rho - 1)}{n(\rho - 1)(\rho^n - 1)}. \tag{10}$$

Thus, it is possible to derive the exact weighting of end-of-period forecasts that will be equal to the period average forecast. As illustrated in Figure 3 for most values of ρ the weight is primarily given to the end-of-period forecast in period t. Only when the process is very persistent, do the end-of-period forecasts approach an equal weighting (as $\rho \to 1$, $\omega^* \to 0.5$) .

Figure 3. Weight ω^* for which the Average and Point Forecasts Intercept

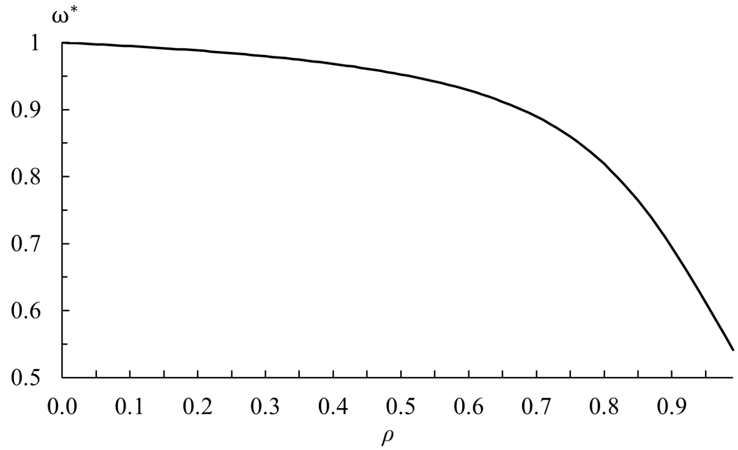

Note: Weight ω^* for which the weighted end-of-period forecast is equal to the period average forecast. Calculations for a daily AR(1) process with $0 \leq \rho < 1$.

2.3.3 End-of-Period Forecasts: AR(1) case

Suppose that instead of using the optimal point forecast or the optimal weighting of end-of-period forecasts, a forecaster uses only the end-of-period forecast as a forecast of the period average. For an AR(1), as shown in the last section, this point forecast should result in the largest forecast error from the true period average for any point forecast within t, for arbitrary k, n. However, the potential use of end-of-period forecast has important practical applications. For example, it informs when forecasts conclusions derived for end-of-period forecasts are also likely to hold for forecasts of period averages. So, when is an end-of-period forecast a good forecast of the period average?

Proposition 1. *Under the BU approach, the forecast error for the temporally aggregated AR(1) data with $n \geq 2$ and forecast horizon k is given by*

$$\sum_{i=1}^{n} \left[y_{T+k,i} - \mathbf{E}_{T,n} \left(y_{T+k,i} \right) \right] = \sum_{i=1}^{n} \sum_{j=1}^{i} \sum_{l=1}^{k} \rho^{(k-l)n-j+i} \varepsilon_{T+l,j} \, .$$

Proof. See appendix A1.3 □

Proposition 2. *Using the end-of-period forecast as the forecast of temporally aggregated AR(1) data, for $n \geq 2$ and forecast horizon k, the forecast error is given by*

$$\sum_{i=1}^{n} \left[y_{T+k,i} - \mathbf{E}_{T,n} \left(y_{T+k,n} \right) \right] = \sum_{i=1}^{n} y_{T,n} \left(\rho^{(k-1)n+i} - \rho^{kn} \right) + \sum_{i=1}^{n} \sum_{j=1}^{i} \sum_{l=1}^{k} \rho^{(k-l)n-j+i} \varepsilon_{T+l,j} \, .$$

Proof. See appendix A1.4 □

Claim 2. *For $n \geq 2$, the end-of-period forecast converges to the bottom-up forecast of temporally aggregated AR(1) data as*

(a) k becomes large and as

(b) $\rho \to 1$.

Proof. The result is immediate from Proposition 1 and Proposition 2 as the difference in the two forecasts is given by the first term on the right-hand side of Proposition 2. □

Claim 2(a) indicates that the difference between point and average forecasts is primarily an issue for short-horizon forecasts, as the forecasts converge at longer horizons. Moreover, claim 2(b) indicates that independent of the forecast horizon, end-of-period point forecasts converge to BU forecasts as the persistence of the data increases. This is noteworthy because this is precisely the situation in which aggregation matters most, in the sense that forecasts constructed from models estimated with average data perform poorly relative to the BU approach (Amemiya and Wu, 1972; Tiao, 1972). As such, end of period forecasts are expected to be useful precisely when recursive forecasts constructed with aggregate data are performing the worst. Due to the high persistence often found in economic series, this case is also particularly relevant in practice.

3 Simulated Forecast Performance

We now use simulation analysis to quantify the forecast efficiency of alternative methods to forecast period averages. For this purpose, we assume that the underlying data, $y_{t,i}$, is observed for each day i in period t. Consistent with the existing simulation analysis examining forecast losses from aggregation (Amemiya and Wu, 1972; Tiao, 1972), we consider an autoregressive process of order one:

$$y_{t,i} = \rho y_{t,i-1} + \epsilon_{t,i}, \qquad \text{for } i = 0, 1, 2, ..., n; \; t = 1, 2, ...T. \tag{11}$$

where $\epsilon_{t,i} \sim N(0, 1)$, and n is the number of daily observations in t, which we allow to be months, or quarters with $n = 21$, and 63, respectively. These assumptions extend the existing simulation analysis of Amemiya and Wu (1972) and Tiao (1972) who only examine up to $n = 4$. In addition, we extend existing analysis by quantifying values of ρ between 0.9 and 1 which is the operational range for most economic daily data.

The objective is to forecast the k-period-ahead average series, $\hat{\bar{y}}_{T+k}$, where $\bar{y}_t = n^{-1}\sum_{i=1}^{n} y_{t,i}$. As a baseline, we simulate 40 years worth of data in addition to burning the first 500 daily observations. We use the first 75 percent of the sample for estimation and the remaining 25 percent as the forecast evaluation sample. This setup reflects common applications for daily financial data, which are typically available since the early 1980s. The setup is also consistent with the applications to macroeconomic variables in section 4.

We assume that the model structure is known but allow for parameter uncertainty. The autoregressive parameter is reestimated at every period with an expanding window, and forecasts are computed out-of-sample. Under point sampling of end-of-period data, the sampled data remains an AR(1) process, whereas the monthly average data is best approximated by an ARMA(1,1) (Weiss, 1984). Consequently, forecasts estimated with period average observations rely on an ARMA(1,1).

We report two common forecast criteria, the MSFE ratio and the success ratio for directional accuracy. Both criteria are expressed relative to the period-average no-change forecast. This no-change benchmark is commonly used in forecasting applications for aggregated data and is most suitable to highlight the gains from using disaggregated forecast approaches.

The MSFE ratio for the k-steps-ahead forecast, $MSFE_k^{ratio}$, is calculated as the ratio of the MSFE of the model-based forecast to the MSFE of the period-average no-change forecast:

$$MSFE_k^{ratio} = \frac{\sum_{q=1}^{Q}(\bar{y}_{q+k} - \hat{\bar{y}}_{q+k|q})^2}{\sum_{q=1}^{Q}(\bar{y}_{q+k} - \bar{y}_q)^2}, \tag{12}$$

where $q = 1, 2, 3, \ldots, Q$ denotes all periods of the forecast evaluation sample, and $\hat{\bar{y}}_{q+k|q}$ is the conditional forecast for the k-step-ahead aggregated observation, $\bar{y}_{q+k}$.

The directional accuracy is assessed using the mean directional accuracy, referred to as the success ratio. It describes the fraction of times the forecasting model can correctly predict the change in the direction of the series of interest:

$$SR_k = \frac{1}{Q}\sum_{q=1}^{Q} \mathbb{1}[sgn(\bar{y}_{q+k} - \bar{y}_q) = sgn(\hat{\bar{y}}_{q+k|q} - \bar{y}_q)], \tag{13}$$

where $\mathbb{1}[\cdot]$ is an indicator function with 1 if true and 0 otherwise, and $sgn(\cdot)$ is a sign function with

$$sgn(x) = \begin{cases} 1 & x > 0 \\ -1 & x \leq 0 \end{cases}. \tag{14}$$

Under this definition, the success ratio equals one half ($SR_k = 0.5$) when there is no directional accuracy, while success ratios greater than one half ($SR_k > 0.5$) indicate directional predictability.

3.1 Forecast Models

To quantify the accuracy of alternative forecasting techniques, we compare the accuracy of five alternative forecasts.

The first approach that we consider is the BU approach. In the current setting, it involves first estimating the AR(1) model at the daily frequency and then ex post averaging the forecasts to the period average. Specifically, the estimated model using the daily data, $y_{t,i}$, is:

$$y_{t,i} = \hat{\rho} y_{t,i-1} + \hat{\epsilon}_{t,i}, \quad \forall\, t \leq T\,. \tag{15}$$

The parameter $\hat{\rho}$ is used to construct recursive model-based forecasts of the daily data, $\hat{y}_{T,n+h|T}$, where $h \geq 1$ is the forecast horizons in days. Then, the daily forecasts, $\hat{y}_{T+k,i|T}$, $k \geq 1$, are averaged to the period average:

$$\hat{\bar{y}}_{T+k|T} = \frac{1}{n} \sum_{i=1}^{n} \hat{y}_{T+k,i|T}, \quad \forall\, k \geq 1. \tag{16}$$

The second approach is the aggregated approach. In the current setting, it involves estimating an ARMA(1,1) model with period-average data to construct forecasts of $\bar{y}_{T+k}$. Specifically, the estimate using the monthly average data, $\bar{y}_t$, is:

$$\bar{y}_t = \tilde{\rho} \bar{y}_{t-1} + \tilde{\epsilon}_t + \tilde{\alpha} \tilde{\epsilon}_{t-1}, \quad \forall\, t \leq T\,. \tag{17}$$

The parameters $\tilde{\rho}$ and $\tilde{\alpha}$ are used to construct recursive model-based forecasts of the monthly average, $\tilde{\bar{y}}_{T+k|T}$.

We next examine an application of period-end-point sampling (PEPS). The model is estimated with end-of-period data, and then point forecasts are constructed and used as the forecasts of the period average. Specifically, the estimated AR(1) model at the period-t frequency using a time series of end-of-period values is:

$$y_{t,n} = \check{\rho} y_{t-1,n} + \check{\epsilon}_t, \quad \forall\, t \leq T\,. \tag{18}$$

The end-of-period point forecasts are used to construct the point forecast for i^* derived in section 2, denoted $\check{y}_{T+k,i^*|T}$, which is equivalent to the period average forecast. For an AR(1), point sampling implies that $\check{\rho} = \rho^n$ (Zellner and Montmarquette, 1971). Thus, the point forecast is equal to the BU forecast using the estimated parameter $\check{y}_{T+k,i^*|T} = \check{\rho}^{(k-1)+i^*/n}\hat{y}_{T,n}$. Then, following PEPS, the i^* point forecasts are used as the forecasts of the period averages $\hat{\bar{y}}_{T+k} = \check{y}_{T+k,i^*|T}$.

PEPSI forecasts are obtained using a linear interpolation of the end-of-month forecasts of two adjacent months:

$$\dot{y}_{T+k,i^{\#}|T} = \omega\check{y}_{T+k,n|T} + (1-\omega)\check{y}_{T+k-1,n|T}, \qquad \forall\ k \geq 1\,, \tag{19}$$

where ω is the forecast averaging weight, which use the estimates of $\check{\rho}$ in equation 10. Then, the $i^{\#}$ point forecasts are used as the forecasts of the period averages $\hat{\bar{y}}_{T+k} = \dot{y}_{T+k,i^{\#}|T}$.

In addition, we also quantify the use of the end-of-period (EoP) forecast as the forecast of the period average. In this case, the model-based forecasts of the end-of-period $\check{y}_{T+k,n|T}$. are used as the forecasts of the period average $\check{\bar{y}}_{T+k} = \check{y}_{T+k,n|T}$. The purpose of including this forecast is to examine the trade-off between the forecast gains from using point-sampled data with the forecast losses from not using the optimal point forecast.

Finally, we also consider the no-change (NC) forecast of the daily data $\hat{\bar{y}}_{T+k} = y_{T,n}$ which is given by the end-of-period no-change. This no-change forecast is used to test against the random walk hypothesis, and when, $\rho = 1$, the end-of-month no-change forecast outperforms the period-average no-change forecast for all n, h (Ellwanger and Snudden, 2023a).

3.2 Simulated Performance

The comparison of the alternative forecasts of monthly average observations for alternative values of ρ is provided in Table 1. Values of the MSFE ratio less than one indicate mean-squared forecast improvements relative to the monthly average no-change forecast. Values of the success ratio above 0.5 indicate improvements in directional accuracy above random chance.

The last column of Table 1 shows the relative performance of the no-change forecast constructed with end-of-month values. When $\rho > 0.9$, the end-of-month no-change forecast is a more accurate forecast of the monthly average than the monthly average no-change forecast, consistent with Tiao (1972). For large values of ρ, the forecast gains from the end-of-month no-change forecast relative to the monthly average no-change forecast can be substantial, with MSFE reductions approaching 46 percent. In terms of directional accuracy, the end-of-month no-change forecast outperforms the

Table 1. One-Month-Ahead Forecast Performance of Alternative Forecast Approaches

Method	Aggregate	Bottom-Up	PEPS	PEPSI	EoP	No-Change
Data	Average	Daily	EoM	EoM	EoM	Daily
ρ			MSFE Ratio			
1.00	0.94	0.54	0.54	0.54	0.55	0.54
	(0.044)	(0.065)	(0.065)	(0.065)	(0.066)	(0.064)
0.995	0.89	0.54	0.54	0.54	0.56	0.56
	(0.043)	(0.064)	(0.064)	(0.064)	(0.070)	(0.068)
0.99	0.90	0.54	0.54	0.54	0.57	0.58
	(0.044)	(0.064)	(0.064)	(0.064)	(0.072)	(0.073)
0.98	0.85	0.54	0.54	0.54	0.60	0.63
	(0.043)	(0.063)	(0.063)	(0.063)	(0.076)	(0.081)
0.95	0.75	0.53	0.54	0.53	0.63	0.77
	(0.041)	(0.060)	(0.060)	(0.061)	(0.077)	(0.111)
0.90	0.65	0.53	0.53	0.53	0.63	1.04
	(0.044)	(0.056)	(0.061)	(0.056)	(0.072)	(0.166)
ρ			Success Ratio			
1.00	0.58	0.74	0.74	0.74	0.74	0.74
	(0.045)	(0.039)	(0.039)	(0.039)	(0.039)	(0.039)
0.995	0.61	0.74	0.74	0.74	0.73	0.73
	(0.041)	(0.039)	(0.039)	(0.039)	(0.040)	(0.039)
0.99	0.61	0.74	0.74	0.74	0.73	0.73
	(0.041)	(0.039)	(0.039)	(0.039)	(0.039)	(0.040)
0.98	0.63	0.74	0.74	0.74	0.72	0.72
	(0.038)	(0.039)	(0.039)	(0.039)	(0.039)	(0.040)
0.95	0.67	0.74	0.74	0.74	0.72	0.69
	(0.037)	(0.039)	(0.039)	(0.039)	(0.038)	(0.042)
0.90	0.70	0.74	0.74	0.74	0.72	0.65
	(0.036)	(0.038)	(0.044)	(0.038)	(0.038)	(0.044)

Note: Comparison of monthly forecasts for 10000 simulations of an AR(1) model at the daily frequency when estimated with alternative methods (standard deviation of the ratios in brackets). All MSFE and success ratios are expressed relative to the monthly average no-change forecast. Values of the MSFE ratio less than one indicate improvements over the monthly average no-change forecast. Values of the success ratio greater than 0.5 indicate improvements over random chance. End-of-period (EoP) uses end-of-period forecasts as the forecast of the average. The last column presents the forecasts from the no-change forecast based on the end-of-month (EoM) observation.

monthly average no-change forecast even for $\rho = 0.9$, demonstrating that temporal disaggregation gains arise for directional accuracy over a wider parameter space than for MSFE precision.

The forecasts constructed with aggregated data are shown in the first column of Table 1. These forecasts outperform the period-average no-change forecast for all values of $\rho \leq 1$. However, for $\rho > 0.95$, the forecasts constructed with aggregated data perform worse than the end-of-month no-change forecast. This example shows that comparisons with the period-average no-change forecast can lead to spurious predictability, and illustrates the importance of comparing forecasts of aggregated data to the random forecast of the daily data to evaluate the usefulness of the forecasts.

In contrast, the relative performance of the BU forecasts (column three of Table 1) does not depend as much on the autocorrelation of the underlying data. For all $\rho < 1$, the bottom-up approach outperforms the end-of-month no-change forecast. The gains in forecast accuracy from PEPS and PEPSI, shown in column four and five of Table 1, are very similar to that of the BU forecasts, consistent section 2. The magnitude of the forecast gains from the PEPS and BU forecasts demonstrates the substantial advantages from temporal disaggregation of daily data in forecasts.

Finally, for MSFE precision, the end-of-month forecasts outperform the end-of-month no-change forecast outperforms for all $\rho < 0.99$. These results are fascinating as they suggest that the forecast gains from using point-sampled data far exceed the forecast losses from not using the optimal point forecast within t. In fact, the end-of-month forecasts substantially outperform period average no-change forecasts and model-based forecasts constructed using period averages in all cases.

The exercises demonstrate that substantial gains in forecast accuracy can be obtained by using information from the underlying daily data for forecasts of the monthly average data. These results also suggest that PEPS point forecasts rival the efficiency of the bottom up approach, and provide effective forecasts of monthly average data.

3.3 Longer Horizons

The loss in relative forecast accuracy for forecasts computed with temporally aggregated data is generally largest at the one-step-ahead prediction, and decreases for longer forecast horizons (Amemiya and Wu, 1972; Tiao, 1972). The intuition for this result is that the aggregation bias is constant across horizons, whereas the forecast error resulting from unpredictable innovations increases with the forecast horizon. For the same reason, the relative performances of the BU forecasts, PEPS and PEPSI forecasts, and end-of-period forecasts, and forecasts computed with aggregated data converge at longer horizons.

This pattern is confirmed in Table 2, which shows the performance of the five forecasts (equations 15–18) at the 1-, 3-, 6-, and 12-month horizon. For the 3-month horizon and beyond, the performance of all model-based forecasts is similar for the case of $\rho = 0.95$. The performance of PEPS is very close to the BU forecast beyond the one-month-ahead horizon. In contrast, the use of aggregated data results in the largest forecast errors at medium-run horizons when the daily data is persistent ($\rho = 0.995$). These results indicate that the relative performance of the BU and PEPS approaches quickly convergence at longer horizons, reinforcing the idea that the loss in forecast

Table 2. Performance of Forecasts at Alternative Horizons

Model	ρ=0.995						ρ=0.95					
Method	Aggregate	BU	PEPS	PEPSI	EoP	NC	Aggregate	BU	PEPS	PEPSI	EoP	NC
Data	Average	Daily	EoM	EoM	EoM	Daily	Average	Daily	EoM	EoM	EoM	Daily
Horizon						MSFE Ratio						
1	0.92	0.54	0.54	0.54	0.56	0.56	0.75	0.53	0.54	0.53	0.63	0.77
	(0.043)	(0.064)	(0.064)	(0.064)	(0.070)	(0.068)	(0.041)	(0.060)	(0.060)	(0.061)	(0.077)	(0.111)
3	0.88	0.81	0.81	0.81	0.81	0.92	0.54	0.54	0.54	0.54	0.54	1.18
	(0.062)	(0.063)	(0.076)	(0.076)	(0.076)	(0.052)	(0.060)	(0.059)	(0.062)	(0.062)	(0.062)	(0.103)
6	0.79	0.77	0.77	0.77	0.77	0.98	0.51	0.51	0.51	0.51	0.51	1.20
	(0.105)	(0.100)	(0.111)	(0.111)	(0.111)	(0.041)	(0.062)	(0.062)	(0.062)	(0.062)	(0.062)	(0.099)
12	0.67	0.66	0.66	0.66	0.66	1.01	0.51	0.51	0.51	0.51	0.51	1.20
	(0.154)	(0.152)	(0.157)	(0.157)	(0.157)	(0.034)	(0.065)	(0.065)	(0.065)	(0.065)	(0.065)	(0.103)
						Success Ratio						
1	0.60	0.74	0.74	0.74	0.73	0.73	0.67	0.74	0.74	0.74	0.72	0.69
	(0.042)	(0.039)	(0.039)	(0.039)	(0.040)	(0.039)	(0.037)	(0.039)	(0.039)	(0.039)	(0.038)	(0.042)
3	0.62	0.65	0.65	0.65	0.65	0.60	0.74	0.74	0.74	0.74	0.74	0.55
	(0.048)	(0.043)	(0.044)	(0.044)	(0.044)	(0.040)	(0.038)	(0.038)	(0.038)	(0.038)	(0.038)	(0.042)
6	0.66	0.67	0.67	0.67	0.67	0.56	0.75	0.75	0.75	0.75	0.75	0.54
	(0.059)	(0.054)	(0.055)	(0.055)	(0.055)	(0.040)	(0.038)	(0.038)	(0.038)	(0.038)	(0.038)	(0.042)
12	0.70	0.70	0.70	0.70	0.70	0.53	0.75	0.75	0.75	0.75	0.75	0.54
	(0.070)	(0.068)	(0.068)	(0.068)	(0.068)	(0.041)	(0.040)	(0.040)	(0.040)	(0.040)	(0.040)	(0.043)

Note: Comparison of monthly forecasts at alternative forecast horizons for 10,000 simulations of an AR(1) model at the daily frequency when estimated with alternative methods (standard deviation of the ratios in brackets). All MSFE and success ratios are expressed relative to the monthly average no-change forecast. Values of the MSFE ratio less than one indicate improvements over the monthly average no-change forecast. Values of the success ratio greater than 0.5 indicate improvements over random chance. End-of-period (EoP) uses end-of-period forecasts as the forecast of the average. The last column presents the forecasts from the no-change forecast based on the end-of-month (EoM) observation.

accuracy from using aggregated data is the largest at the one-step-ahead forecast.

3.4 Quarterly Data

Table 3 reports the one-period ahead relative forecast gains for quarterly aggregation frequencies and values of ρ. The gains from disaggregated methods are even greater at the quarterly frequency. This indicates that PEPS is useful for aggregation frequencies commonly encountered in practice.

Table 3. Performance of Forecasts at Quarterly Sampling Frequency

Aggregation	Quarterly					
Method	Aggregate	BU	PEPS	PEPSI	EoP	NC
Data	Average	Daily	EoM	EoM	EoM	Daily
ρ			MSFE Ratio			
1.00	0.96	0.53	0.53	0.53	0.55	0.53
	(0.085)	(0.116)	(0.116)	(0.116)	(0.123)	(0.114)
0.995	0.88	0.53	0.53	0.53	0.58	0.59
	(0.082)	(0.114)	(0.115)	(0.115)	(0.138)	(0.138)
0.99	0.83	0.53	0.53	0.53	0.61	0.66
	(0.080)	(0.112)	(0.113)	(0.113)	(0.142)	(0.163)
0.98	0.74	0.52	0.53	0.53	0.64	0.80
	(0.080)	(0.108)	(0.109)	(0.111)	(0.142)	(0.218)
0.95	0.62	0.52	0.53	0.54	0.62	1.30
	(0.084)	(0.100)	(0.101)	(0.171)	(0.132)	(0.409)
0.90	0.57	0.52	0.53	0.60	0.58	2.21
	(0.084)	(0.092)	(0.097)	(0.231)	(0.120)	(0.741)
ρ			Success Ratio			
1.00	0.58	0.75	0.75	0.75	0.74	0.75
	(0.078)	(0.067)	(0.067)	(0.067)	(0.068)	(0.067)
0.995	0.62	0.75	0.74	0.73	0.73	0.73
	(0.068)	(0.068)	(0.068)	(0.068)	(0.068)	(0.070)
0.99	0.64	0.75	0.75	0.72	0.72	0.71
	(0.065)	(0.067)	(0.067)	(0.068)	(0.068)	(0.070)
0.98	0.67	0.75	0.75	0.72	0.72	0.69
	(0.063)	(0.067)	(0.067)	(0.065)	(0.065)	(0.073)
0.95	0.72	0.75	0.75	0.72	0.72	0.65
	(0.062)	(0.065)	(0.065)	(0.063)	(0.063)	(0.078)
0.90	0.73	0.75	0.75	0.73	0.73	0.61
	(0.060)	(0.063)	(0.063)	(0.062)	(0.062)	(0.080)

Note: Comparison of one-period ahead quarterly forecasts for 10,000 simulations of an AR(1) model at the daily frequency when estimated with alternative methods (standard deviation of the ratios in brackets). All MSFE and success ratios are expressed relative to the monthly average no-change forecast. Values of the MSFE ratio less than one indicate improvements over the monthly average no-change forecast. Values of the success ratio greater than 0.5 indicate improvements over random chance. End-of-period (EoP) uses end-of-period forecasts as the forecast of the average. The last column presents the forecasts from the no-change forecast based on the end-of-month (EoM) observation.

4 Application to Real-time Forecasts

4.1 Data

The empirical applications consider real-time monthly average forecasts of two macroeconomic series: the nominal yield on 10-year U.S. government bonds and the real price of copper. Daily closing spot prices of copper are obtained from the London Metal Exchange. Daily data for the yield on the 10-year Treasury bonds is obtained from FRED. The price index used to deflate the nominal price of copper is the seasonally adjusted U.S. consumer price index (CPI) from the real-time database of the Philadelphia Federal Reserve. The CPI is published with a one-month delay and is nowcasted using the average historical growth rate. A detailed description of the data sources and CPI nowcasting is provided in appendix A2.

Table 4. Descriptive Statistics

Real Series	Date Range	PACF(1)	N	Mean	Std. Dev.	Data Construction
Monthly Average						
Copper	1986.04–2021.01	0.9908	418	2094.32	899.27	Monthly average
Backcasted Copper	1973.01–2021.01	0.9890	577	2153.15	929.22	Monthly average
10-year Treasury Yield	1973.01–2021.01	0.9956	577	6.19	3.21	Monthly average
End of Month						
Copper	1986.04–2021.01	0.9908	418	2103.78	908.33	Last trading day
Backcasted Copper	1973.01–2021.01	0.9890	577	2158.74	934.38	Last trading day
10-year Treasury Yield	1973.01–2021.01	0.9944	577	6.18	3.23	Last trading day
Daily						
Copper	1986.04.01–2021.01.31	0.9991	8837	2106.94	905.61	Closing prices
10-year Treasury Yield	1973.01.02–2021.01.29	0.9998	12075	6.16	3.22	Closing prices
Monthly Index						
Consumer Price Inflation	1973.01–2021.01	0.9922	565	3.86	3.03	Monthly index
OECD CLI	1973.01–2021.01	0.9758	577	0.00	1.42	Monthly index

Note: Copper prices are reported in real terms, 10-year bond yields in nominal terms. Consumer price inflation is reported in year-over-year growth rates in percent. The number of monthly or daily observations is denoted N. The 2021.04 vintage is reported for the consumer price index. *PACF(1)* denotes the partial autocorrelation coefficient of the first lag.

The descriptive statistics for the end-of-month and monthly average series are reported in Table 4. Monthly averages are the simple average of daily closing prices. End-of-month observations are the closing price on the last trading day of each month. For all of our series, the standard deviation of the end-of-period observations are very similar to that of the aggregated observations. The partial autocorrelations coefficient of the monthly real data indicates that all the series are highly persistent.

One advantage of using monthly data is that the monthly average and end-of-month prices can be backcasted using monthly average data. This is useful because the backcasting may provide a better estimate of the long-run mean, which could improve the forecast accuracy at longer horizons. The usefulness of backcasting is an empirical question and is quantified for the real price of copper. End-of-month and monthly average copper prices are backcasted to 1973M1 using the World Bank Monthly Average Commodities Price.

All high-frequency observations are available in real time and are not subject to historical revisions. Models are estimated using the available data, which starts in 1973M1 for yields and backcasted copper prices and in 1986M4 for copper prices. The real end-of-period observations, $R_{t,n}$, are constructed by applying the CPI index to the last observation of the period using $R_{t,n|t} = p_{t,n}/CPI_{t|t}$. This measure of end-of-month real prices is common for forecasts of end-of-month real prices, such as bilateral exchange rates (Meese and Rogoff, 1983) and primary commodities (West and Wong, 2014).

4.2 Forecasts

All forecasts are computed out-of-sample using real-time methods. Specifically, we use historical vintages of (nowcasted) data available in the month of the forecast, and the models are re-estimated at every monthly step with an expanding window. The forecast evaluation period is 2000M1-2021M1.

Augmented Dickey-Fuller tests are used to test for unit roots and determine the appropriateness of estimating the models in levels or differences. The results suggest that the backcasted monthly real prices of copper are weakly stationary in log real levels, but the nominal daily data and the non-backcasted monthly data are only weakly stationary in differences. Moreover, both the daily and monthly yields are stationary in differences. Accordingly, the models for nominal interest rate are estimated in differences, and the models for the real backcasted copper are estimated in log levels, and real and non-backcasted copper are estimated in growth rates.

For models estimated in monthly averages in real log levels, $\bar{r}_t = \ln(\bar{R}_t)$, the ARMA$(p,q)$ with p autoregressive and q moving-average parameters, and estimated innovations $\tilde{\epsilon}_t$ is given by:

$$\tilde{a}(L)\bar{r}_t \quad = \quad \tilde{c} + \tilde{b}(L)\tilde{\epsilon}_t \,, \qquad \forall\; t \leq T \,, \tag{20}$$

where $\tilde{c}$ is the estimated constant and L is the lag operator such that $Ly_t = y_{t-1}$, $\tilde{b}(L) = (1 +$

$\tilde{\alpha}_1 L + \cdots + \tilde{\alpha}_q L^q)$, and $\tilde{a}(L) = (1 - \tilde{\rho}_1 L - \cdots - \tilde{\rho}_p L^p)$. These estimated parameters are used to construct recursive model-based forecasts, $\tilde{\bar{r}}_{T+k|T}$, which are converted back into real prices in levels, $\tilde{\bar{R}}_{T+k|T} = \exp\left(\tilde{\bar{r}}_{T+k|T}\right)$.

Forecasts of the monthly average nominal yields, $\bar{S}_t$, are estimated in differences $\bar{d}_t = \bar{S}_t - \bar{S}_{t-1}$. Then, the level forecasts are constructed using the model-implied difference, such that

$$\tilde{\bar{S}}_{T+k|T} = \tilde{\bar{d}}_{T+k|T} + \tilde{\bar{S}}_{T+k-1|T}, \; \forall \; k \geq 1 \, .$$

Similarly, forecasts constructed using monthly average real growth rates, $\bar{g}_t = \bar{R}_t / \bar{R}_{t-1} - 1$, are converted back into real-level forecasts using the model-implied net growth rate

$$\tilde{\bar{R}}_{T+k|T} = (1 + \tilde{\bar{g}}_{T+k|T}) \cdot \tilde{\bar{R}}_{T+k-1|T} \; \forall \; k \geq 1 \, .$$

Since all nominal daily prices are difference stationary, we estimate the model using the growth rate of nominal daily prices, $g_{t,i} = S_{t,i}/S_{t,i-1}$. The estimated ARMA$(p, q)$ model is given by

$$\hat{a}(L) g_{t,i} \quad = \quad \hat{c} + \hat{b}(L)\hat{\epsilon}_{t,i}, \qquad \forall \; i, \; t \leq T. \tag{21}$$

Consistent with Rossana and Seater (1995), we employ information criteria for model selection to select autoregressive terms.[5] The estimated parameters are used to construct recursive model-based forecasts of the growth rate of daily nominal prices, $\hat{g}_{T,n+h|T}$, where $h \geq 1$ is the forecast horizons in days. The forecasts for the level of the nominal price on day i of month $T + k$, given month T information, are based on the model-implied net growth rate:

$$\hat{S}_{T,n+h|T} = (1 + \hat{g}_{T,n+h|T}) \cdot \hat{S}_{T,n+h-1|T} \, , \quad \forall \; h \geq 1. \tag{22}$$

For the BU forecasts of monthly average real data, one issue to overcome is that the CPI is only available at the monthly frequency, which may help explain why this approach has been overlooked in applications to real macroeconomic variables that are aggregated from a daily frequency (see for example, Box et al., 2015). However, we can simply average the nominal daily forecasts to the

[5] Our results are quantitatively robust to the ARMA(1,1) benchmark, as shown in section 4.5.

monthly frequency and deflate into real prices using the expected CPI deflator:

$$\hat{\bar{R}}_{T+k|T} = \frac{n^{-1}\sum_{i=1}^{n}\hat{S}_{T+k,i|T}}{\mathbf{E}_{T,n}[CPI_{T+k|T}]}, \quad \forall\ k \geq 1. \tag{23}$$

The expected CPI, $\mathbf{E}_{T,n}[CPI_{T+k|T}]$, is constructed using the standard practice of expanding the (nowcasted) current CPI observations with the average historical rate of inflation.

Yield forecasts are constructed similarly, except that the nominal daily data is transformed into differences, $d_{t,i} = S_{t,i} - S_{t,i-1}$, instead of growth rates. In this case, the forecasts of the series in levels are constructed by summing over the forecasted differences:

$$\hat{S}_{T,n+h|T} = \hat{d}_{T,n+h|T} + \hat{S}_{T,n+h-1|T}, \quad \forall\ h \geq 1. \tag{24}$$

The daily forecasts of the nominal series in levels are then averaged to the monthly frequency:

$$\hat{\bar{S}}_{T+k|T} = \frac{1}{n}\sum_{i=1}^{n}\hat{S}_{T+k,i|T}, \quad \forall\ k \geq 1. \tag{25}$$

Implementing the end-of-month forecasts straightforward, as it merely requires replacing the time series of monthly average prices with the time series of real end-of-period observations, $R_{t,n}$, during the model estimation. Specifically, the ARMA(p,q) model is estimated at the monthly frequency with time series of end-of-month prices in log-real levels and expressed as:

$$\breve{a}(L)r_{t,n} \quad = \quad \breve{c} + \breve{b}(L)\breve{\epsilon}_t, \quad \forall\ t \leq T. \tag{26}$$

As before we, report the recursive model-based forecasts of the end-of-month values $\breve{r}_{T+k,n|T}$, which are converted back into real prices in levels, $\breve{R}_{T+k,n|T} = \exp\left(\breve{r}_{T+k,n|T}\right)$. The end-of-month forecasts are used as forecasts of the corresponding monthly average value, $\breve{\bar{R}}_{T+k|T} = \breve{R}_{T+k,n|T}$. End-of-month forecasts estimated in nominal differences are converted into nominal levels following the approach outlined above.

We explore two disaggregated PEPSI approaches, which are obtained using a linear interpolation of the end-of-month forecasts of two adjacent months using equation 19.

The first PEPSI application uses a fixed ω numerically selected using the daily data from the pre-forecast evaluation sample. Figure 4 reports the in-sample Pearson autocorrelations for both the daily real price of copper and the nominal 10-year treasury bonds. For both copper and

Figure 4. Intersection of Point and Average Forecast

(a) Real Price of Copper (b) Nominal 10-Year Treasury Bonds

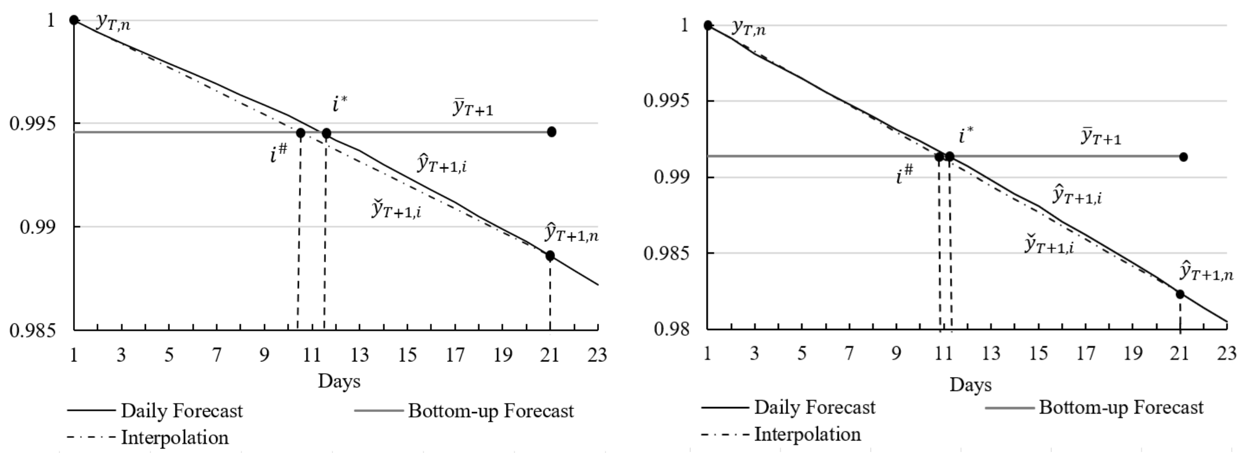

Note: In sampled estimates with $\hat{y}_{T+1,i}$ calculated using daily Pearson autocorrelations. $\dot{y}_{T+1,i}$ is the linearly interpolated end-of-period values, and $\bar{y}_{T+1}$ is the bottom-up forecast calculated as the simple average of the daily forecasts.

yields, points forecasts near the middle-of-month are found to be equal to the period average, and the optimal weight, ω, for PEPSI is given by 0.524 and 0.5095, respectively.[6] The period average PEPSI forecasts are constructed using these weights in equation 19 for the entire forecast evaluation sample.

The second PEPSI application estimates ω using forecast averaging. The parameter is estimated with constrained linear regressions (Davidson, 1993) with an expanding window in real-time. To provide an estimation sample, we utilize forecast outcomes beginning in 1992.01.[7]

Finally, to demonstrate the usefulness of PEPSI in a multivariate setting, we examine forecasts in a two-variable VAR estimated at the monthly frequency that includes the real price of copper and the Organisation for Economic Co-operation and Development's Composite Leading Indicator (CLI) for the G20 countries. The VAR(p) model with p autoregressive parameters can be expressed as:

$$(1 - \check{\mathbf{a}}(L))\mathbf{g}_t \quad = \quad \check{\mathbf{e}}_t, \qquad \forall\ t < T\,, \tag{27}$$

where $\mathbf{g}_t$, is a 2x1 vector of growth rates, $\check{\mathbf{a}}(L)$ is the autoregressive parameter matrix of order p, and $\check{\mathbf{e}}_t$ is a 2x1 vector of innovations. The estimated parameters are used to construct recursive

[6]Interestingly, a grid search of alternative values of ω to two decimal places confirm that these parameter values minimizes forecast error over the forecast evaluation sample.

[7]These forecasts are only used to estimate the weights. We continue to report forecast evaluation criteria for 2000M1–2021M1.

model-based forecasts of the growth rates, $\check{\bar{g}}_{T+k|T}$ which are converted back into real prices in levels in the same way as described above for the ARIMA forecasts.

It is not possible to apply the BU approach to this VAR as the CLI is only observed at the monthly frequency. However, applying PEPSI to a VAR model only involves replacing the monthly average growth rates, $\bar{g}_t$, in the vector $\mathbf{g}_t$, with the end-of-month growth rates $g_{t,n}$. Again, the estimated parameters are used to construct recursive forecasts of the end-of-period growth rates, which are converted back into levels following the approaches described for the ARIMA model.

4.3 Forecast Criteria

For forecast evaluation, we report the MSFE ratio and the success ratio expressed relative to the end-of-month no-change forecast to test the null hypothesis that the future period averages are conditionally unpredictable (Ellwanger and Snudden, 2023a).

The MSFE ratio for the k-steps-ahead forecast, $MSFE_k^{ratio}$, is calculated as the ratio of the MSFE of the model-based forecast to the MSFE of the end-of-month no-change forecast:

$$MSFE_k^{ratio} = \frac{\sum_{q=1}^{Q}(\bar{R}_{q+k} - \hat{\bar{R}}_{q+k|q})^2}{\sum_{q=1}^{Q}(\bar{R}_{q+k} - R_{q,n|q})^2}, \tag{28}$$

where $q = 1, 2, \ldots, Q$ denotes all periods of the evaluation sample, $\hat{\bar{R}}_{q+k|q}$ is the conditional, model-implied real-time forecast for the k-month-ahead observation, $\bar{R}_{q+k}$, and $R_{q,n}$ is the time q end-of-period observation.

The null hypothesis of an equal MSFE for the model-based forecast relative to the no-change forecasts is tested following Diebold and Mariano (1995) and compared against standard normal critical values. P-values for tests relative to the monthly average are reported in parentheses.[8]

Directional accuracy is assessed via the success ratio, SR_k, indicating the fraction of times the forecasting model correctly predicts the change in direction of the series of interest:

$$SR_k = \frac{1}{Q}\sum_{q=1}^{Q} \mathbb{1}[sgn(\bar{R}_{q+k} - R_{q,n}) = sgn(\hat{\bar{R}}_{q+k|q} - R_{q,n|q})], \tag{29}$$

The null hypothesis that the success ratio is 0.5 (corresponding to the case that the directional prediction is completely random) is tested following Pesaran and Timmermann (2009), with the

[8] Under nested models, real-time data that is subject to revisions, and estimation uncertainty, the assumptions underlying Diebold and Mariano (1995) are not met (Diebold, 2015). As is standard, the tests are still reported with this caveat in mind.

corresponding p-values reported in parentheses.

4.4 Results

4.4.1 Nominal Yields on 10-year Treasury Bonds

Table 5 presents the performance of forecasts of the monthly average nominal yields on 10-year U.S. government bonds. It also shows a loss in forecast accuracy at short horizons from using aggregated data, with an MSFE precision that is 81 percent worse than the end-of-month no-change forecast at the one-month-ahead horizon. However, at longer horizons, the forecasts based on aggregated data outperform the end-of-month no-change forecast. These improvements are significant at the 5 percent level for the MSFE at the two-year horizon, and for directional accuracy at the six-month horizon.

Table 5. Real-Time Model-Based Forecasts of the Nominal 10-Year Treasury Bonds

Method	Aggregate	Bottom-up	PEPSI	PEPSI(est)	EoP
Data	Average	Daily	EoM	EoM	EoM
Horizon			MSFE Ratio		
1	1.81 (1.000)	1.02 (0.926)	1.02 (0.784)	**0.99 (0.160)**	1.11 (0.996)
3	1.06 (0.927)	1.02 (0.926)	1.00 (0.427)	1.01 (0.670)	1.01 (0.579)
6	1.03 (0.743)	1.01 (0.666)	**0.98 (0.269)**	**0.99 (0.355)**	**0.99 (0.433)**
12	**0.98 (0.337)**	**0.97 (0.231)**	**0.92 (0.097)**	1.12 (0.991)	**0.92 (0.090)**
24	**0.84 (0.006)**	**0.96 (0.198)**	**0.82 (0.004)**	1.11 (0.998)	**0.81 (0.004)**
			Success Ratio		
1	0.50 (0.460)	0.46 (0.767)	0.49 (0.444)	**0.54 (0.064)**	0.49 (0.444)
3	0.49 (0.524)	0.47 (0.848)	**0.56 (0.031)**	0.46 (0.943)	**0.57 (0.025)**
6	**0.57 (0.027)**	**0.51 (0.880)**	**0.56 (0.121)**	0.49 (0.780)	**0.58 (0.048)**
12	**0.54 (0.431)**	**0.62 (0.932)**	**0.56 (0.398)**	0.47 (0.742)	**0.59 (0.221)**
24	**0.58 (0.161)**	**0.57 (1.000)**	**0.59 (0.209)**	0.45 (0.844)	**0.59 (0.214)**

Note: Real-time, out-of-sample forecasts of the nominal monthly average 10-year Treasury bonds in levels, 2000M1–2021M1. End-of-period (EoP) uses end-of-period forecasts as the forecast of the average. AR(12) selected using AIC. Forecast criteria expressed relative to the end-of-month no-change forecast, with p-values reported in parentheses. Bold values indicate forecast gains relative to the no-change forecast for the MSFE ratio and gain in directional accuracy for the success ratio.

Relative to the forecasts computed with aggregated data, the BU approach improves the MSFE accuracy but fares worse in terms of directional accuracy at short horizons. This arises because directional criteria are impervious to the magnitude of the forecast error. Moreover, the BU approach does not show significant gains in terms of MSFE or directional accuracy beyond the 6-month horizon, highlighting the disadvantages of the BU approach at longer forecast horizons.

At the one-month-ahead horizon, the PEPSI forecasts perform similarly to the BU forecasts in terms of the MSFE. The PEPSI with the estimated weights "PEPSI(est)" exhibits the best one step ahead forecast performance, and the gains for directional accuracy are significant at the 10 percent level. The PEPSI approach with the constant weight performs similar well at longer horizons, which further corroborates the usefulness of the disaggregated approaches utilizing end-of-period data at both short- and long-horizon forecasts.

4.4.2 Real Price of Copper

The forecasts of the monthly average real price of copper are reported in Table A3. In terms of MSFE, the forecasts constructed with monthly average data perform very poorly at short horizons. The use of backcasted monthly average copper prices results in small gains at short horizons. With backcasted data, the long-horizon forecasts outperform the end-of-month no-change forecast in terms of directional accuracy at the two-year horizon. These results illustrate that the longer sample periods are a comparative advantage of using lower-frequency data for longer horizon forecasts.

Table 6. Real-Time Model-Based Forecasts of the Real Price of Copper

Method	Aggregate	Bottom-up	PEPSI	PEPSI(est)	EoP
Data	Average	Daily	EoM	EoM	EoM
Horizon			MSFE Ratio		
1	1.70 (0.999)	**0.99 (0.299)**	**0.93 (0.089)**	**0.96 (0.297)**	**0.97 (0.376)**
3	1.11 (0.809)	1.03 (0.872)	**0.96 (0.326)**	1.01 (0.560)	1.00 (0.495)
6	1.09 (0.708)	1.06 (0.906)	1.01 (0.527)	1.03 (0.626)	1.05 (0.619)
12	1.07 (0.662)	1.14 (0.958)	1.04 (0.590)	1.04 (0.683)	1.06 (0.647)
24	1.09 (0.743)	1.29 (0.998)	1.05 (0.661)	**0.96 (0.339)**	1.07 (0.708)
			Success Ratio		
1	**0.55 (0.057)**	**0.54 (0.081)**	**0.52 (0.309)**	**0.52 (0.309)**	**0.52 (0.309)**
3	**0.52 (0.175)**	0.50 (0.726)	**0.56 (0.053)**	**0.51 (0.343)**	**0.54 (0.092)**
6	**0.54 (0.121)**	0.50 (0.720)	**0.52 (0.210)**	0.50 (0.442)	**0.52 (0.237)**
12	**0.58 (0.111)**	0.42 (0.977)	**0.58 (0.101)**	**0.56 (0.167)**	**0.58 (0.101)**
24	**0.62 (0.040)**	0.38 (1.000)	**0.61 (0.047)**	**0.62 (0.030)**	**0.61 (0.047)**

Note: Real-time, out-of-sample forecasts of the monthly average real price of copper in levels, 2000M1–2021M1. End-of-period (EoP) uses end-of-period forecasts as the forecast of the average. AR(12) selected using AIC. Monthly value begin in 1973 and are backcast. Forecast criteria expressed relative to the end-of-month no-change forecast, with p-values reported in parentheses. Bold values indicate forecast gains relative to the no-change forecast for the MSFE ratio and gain in directional accuracy for the success ratio. "Backcast" data refers to backcasting monthly series using monthly average data to 1973M1.

The third column reports the forecasts constructed using the BU approach. The forecasts improve upon the average forecasts in terms of MSFE precision at short horizons, but not in terms of directional accuracy. Moreover, the forecasts fail to outperform the end-of-month no-change forecast at any horizon and are especially poor at longer horizons. In fact, at horizons beyond six months, the BU approach is worse than the forecasts computed with backcasted average data. Although the BU approach is efficient at short horizons, the inability to backcast daily data suggests that the approach fails to provide effective estimates of the long-run mean. This highlights a relative disadvantage of the BU approach in practice.

The PEPSI forecasts provide the largest MSFE gains at short horizons, and the gains from the constant weight "PEPSI" exhibits significance at the 10 percent level. The approaches utilizing the backcasted end of period data perform exceptionally well at longer horizons. This further corroborates the usefulness of the disaggregated approaches utilizing end-of-period data for both short- and long-horizon forecasts.[9]

Table 7. Real-Time VAR Forecasts of the Real Price of Copper

Method	Aggregate	PEPSI	PEPSI(est)	EoP
Data	Average	EoM	EoM	EoM
Horizon		MSFE Ratio		
1	1.65 (0.999)	**0.94 (0.221)**	**0.85 (0.073)**	**1.04 (0.599)**
3	1.06 (0.609)	**0.93 (0.341)**	**0.78 (0.102)**	**0.98 (0.475)**
6	**0.96 (0.434)**	**0.91 (0.345)**	**0.73 (0.119)**	**0.93 (0.393)**
12	**0.95 (0.399)**	**0.92 (0.343)**	**0.85 (0.177)**	**0.95 (0.412)**
24	**0.84 (0.149)**	**0.84 (0.123)**	1.07 (0.656)	**0.82 (0.103)**
		Success Ratio		
1	**0.56 (0.028)**	**0.54 (0.090)**	**0.57 (0.016)**	**0.54 (0.090)**
3	**0.56 (0.013)**	**0.61 (0.001)**	**0.61 (0.000)**	**0.61 (0.001)**
6	**0.63 (0.000)**	**0.63 (0.000)**	**0.67 (0.000)**	**0.63 (0.000)**
12	**0.67 (0.000)**	**0.70 (0.000)**	**0.71 (0.000)**	**0.70 (0.000)**
24	**0.70 (0.000)**	**0.71 (0.000)**	**0.63 (0.009)**	**0.71 (0.000)**

Note: Real-time, out-of-sample forecasts of the monthly average real price of copper in levels, 2000M1–2021M1. End-of-period (EoP) uses end-of-period forecasts as the forecast of the average. AR(12) selected using AIC. Monthly value begin in 1973 and are backcast. Forecast criteria expressed relative to the end-of-month no-change forecast, with p-values reported in parentheses. Bold values indicate forecast gains relative to the no-change forecast for the MSFE ratio and gain in directional accuracy for the success ratio. "Backcast" data refers to backcasting monthly series using monthly average data to 1973M1.

Table 7 reports the VAR forecasts which use the OCED CLI and estimated at the monthly frequency. The results again show that the model-based forecasts constructed using period average data results in poor forecast performance. However, the monthly average forecasts improve substantially merely by replacing the monthly average with the end-of-month copper data in the VAR. Like for the ARIMA models, both the PEPSI forecasts with the constant and estimated weights continue to improve forecast accuracy. The PEPSI(est) forecasts show significant improvements in both MSFE precision and direction accuracy at the one-step ahead at the 10 and 5 percent level,

[9]When the series are not backcasted, shown in appendix Table A3.3, all methods fail to show gains at longer horizons.

respectively.

The VAR example shows that it is easy to implement PEPS even in situations where altering the frequency of the model is not viable, which is often the case for central bank projection models or other low-frequency macroeconomic frameworks. By replacing the period average observations with the end-of-period observations and ex-post adjusting the forecasts to be equal to the period average, a forecaster can substantially improve the forecast accuracy of period-average observations.

4.5 Robustness

We find that any averaging of the end-of-month observations, such as two-day or weekly averages, systemically reduces accuracy of both the no-change forecasts and model-based forecasts that are estimated with the PEPS approach, see appendix A3.2. This pattern supports the idea that the last available information contained in the closing price reflects all available information about future levels.

The results presented in this paper are remarkably robust to alternative modeling choices. For example, similar results are obtained when model lags are selected using the Schwarz (1978) information criterion, as done by Rossana and Seater (1995). Moreover, the results remain qualitatively and, for the most part, quantitatively unchanged, when ARMA(1,1) models are used instead of the AIC criterion, see appendix A3.4.

Model-based forecasts using disaggregated approaches are also superior for forecasts of quarterly average prices. This is consistent with the simulation evidence in section 3, showing that aggregation over more observations increases the information loss (see also Amemiya and Wu, 1972; Wei, 1978).

Finally, the results hold for alternative assumptions for CPI. For example, using ex-post revised data instead of real-time data does not affect the conclusions. This robustness is expected, as fluctuations in the CPI deflator are generally small compared to fluctuations in prices and tend to have minimal impact on forecasts.

5 Conclusion

We have proposed the method of period-end-point sampling (PEPS) to forecast persistent temporally aggregated data. It consists of using end-of-period observations to construct point forecasts that are equal to the period average forecasts. We have shown that PEPS avoids the loss of information that is induced by aggregation and generally yields superior accuracy to forecasts computed from aggregated data. The gains are sizeable at short horizons and have practical advantages that improve long-run forecast accuracy. The forecast gains rival the BU approach.

A major advantage of PEPS is that it allows the forecast model to maintain the same frequency as the target variable. We have shown that it is straightforward to combine higher-frequency information with information from lower-frequency variables without modifying the estimation frequency. Forecasters who want to maintain models at a lower frequency, which is often the case for central bank projections and other macroeconomic forecasts, should no longer exclude themselves from achieving substantial gains in forecast accuracy using disaggregated techniques.

Data availability statement

The data that support the findings of this study are available in the public domain. See data appendix for details.

Disclosure statement

The authors report there are no competing interests to declare.

References

Amemiya, T. and Wu, R. Y. (1972). The effect of aggregation on prediction in the autoregressive model. *Journal of the American Statistical Association*, 67(339):628–632.

Andreou, E., Ghysels, E., and Kourtellos, A. (2013). Should macroeconomic forecasters use daily financial data and how? *Journal of Business & Economic Statistics*, 31(2):240–251.

Athanasopoulos, G., Hyndman, R. J., Song, H., and Wu, D. C. (2011). The tourism forecasting competition. *International Journal of Forecasting*, 27(3):822–844.

Benmoussa, A. A., Ellwanger, R., and Snudden, S. (2023). Carpe diem: Can daily oil prices improve model-based forecasts of the real price of crude oil? *LCEPRA Working Papers, 2023-5.*

Box, G. E., Jenkins, G. M., Reinsel, G. C., and Ljung, G. M. (2015). *Time series analysis: Forecasting and control*. John Wiley & Sons.

Brewer, K. R. (1973). Some consequences of temporal aggregation and systematic sampling for ARMA and ARMAX models. *Journal of Econometrics*, 1(2):133–154.

Christiano, L. J. and Eichenbaum, M. (1987). Temporal aggregation and structural inference in macroeconomics. *Carnegie-Rochester Conference Series on Public Policy*, 26:63–130.

Davidson, R. (1993). Estimation and inference in econometrics.

Diebold, F. X. (2015). Comparing predictive accuracy, twenty years later: A personal perspective on the use and abuse of diebold–mariano tests. *Journal of Business & Economic Statistics*, 33(1):1–1.

Diebold, F. X. and Mariano, R. S. (1995). Comparing predictive accuracy. *Journal of Business & Economic Statistics*, 13(3):253–263.

Ellwanger, R. and Snudden, S. (2023a). Forecasts of the real price of oil revisited: Do they beat the random walk? *Journal of Banking and Finance*, 154:1–8.

Ellwanger, R. and Snudden, S. (2023b). Futures prices are useful predictors of the spot price of crude oil. *The Energy Journal*, 44(4).

Fama, E. F. (1970). Efficient capital markets: A review of theory and empirical work. *The Journal of Finance*, 25(2):383–417.

Farag, M., Snudden, S., and Upton, G. (2024). Can futures prices predict the real price of primary commodities? *USAEE Working Paper No. 24-629*.

Ghysels, E., Sinko, A., and Valkanov, R. (2007). Midas regressions: Further results and new directions. *Econometric Reviews*, 26(1):53–90.

Kohn, R. (1982). When is an aggregate of a time series efficiently forecast by its past? *Journal of Econometrics*, 18(3):337–349.

Lütkepohl, H. (1984). Forecasting contemporaneously aggregated vector ARMA processes. *Journal of Business & Economic Statistics*, 2(3):201–214.

Lütkepohl, H. (1986). Forecasting temporally aggregated vector ARMA processes. *Journal of Forecasting*, 5(2):85–95.

Lütkepohl, H. (2006). Chapter 6 forecasting with VARMA models. volume 1 of *Handbook of Economic Forecasting*, pages 287–325. Elsevier.

Man, K. and Tiao, G. (2006). Aggregation effect and forecasting temporal aggregates of long memory processes. *International Journal of Forecasting*, 22(2):267–281.

Marcellino, M. (1999). Some consequences of temporal aggregation in empirical analysis. *Journal of Business & Economic Statistics*, 17(1):129–136.

Meese, R. A. and Rogoff, K. (1983). Empirical exchange rate models of the seventies: Do they fit out of sample? *Journal of International Economics*, 14(1-2):3–24.

Pesaran, M. H. and Timmermann, A. (2009). Testing dependence among serially correlated multicategory variables. *Journal of the American Statistical Association*, 104(485):325–337.

Rossana, R. J. and Seater, J. J. (1995). Temporal aggregation and economic time series. *Journal of Business & Economic Statistics*, 13(4):441–451.

Schwarz, G. (1978). Estimating the dimension of a model. *The Annals of Statistics*, pages 461–464.

Stram, D. O. and Wei, W. W. (1986). Temporal aggregation in the ARIMA process. *Journal of Time Series Analysis*, 7(4):279–292.

Telser, L. G. (1967). Discrete samples and moving sums in stationary stochastic processes. *Journal of the American Statistical Association*, 62(318):484–499.

Tiao, G. C. (1972). Asymptotic behaviour of temporal aggregates of time series. *Biometrika*, 59(3):525–531.

Tsai, H. and Chan, K.-S. (2005). Temporal aggregation of stationary and nonstationary discrete-time processes. *Journal of Time Series Analysis*, 26(4):613–624.

Wei, W. W. (1978). Some consequences of temporal aggregation in seasonal time series models. In *Seasonal analysis of economic time series*, pages 433–448. NBER.

Wei, W. W. (1981). Effect of systematic sampling on arima models. *Communications in Statistics-Theory and Methods*, 10(23):2389–2398.

Weiss, A. A. (1984). Systematic sampling and temporal aggregation in time series models. *Journal of Econometrics*, 26(3):271–281.

West, K. D. and Wong, K.-F. (2014). A factor model for co-movements of commodity prices. *Journal of International Money and Finance*, 42:289–309.

Working, H. (1960). Note on the correlation of first differences of averages in a random chain. *Econometrica*, 28(4):916–918.

Zellner, A. and Montmarquette, C. (1971). A study of some aspects of temporal aggregation problems in econometric analyses. *The Review of Economics and Statistics*, 53(4):335–342.

Online Appendix (Not intended for publication)

A1 Detailed Proofs

A1.1 Proof of IVT in Theorem 1

Intermediate Value Theorem for Forecasts: If $\hat{y}_{T+k,i}$ is a continuous function on the interval $i \in [1, n]$, and $\hat{\hat{y}}_{T+k}$ is any number between $\hat{y}_{T+k,1}$ and $\hat{y}_{T+k,n}$ inclusive, then there exists at least one i^* in the interval $[1, n]$ such that $\hat{y}_{T+k,i^*} = \hat{\hat{y}}_{T+k}$.

Proof. Without loss of generality, assume $\hat{y}_{T+k,1} < \hat{\hat{y}}_{T+k} < \hat{y}_{T+k,n}$. Consider the set $S = \{x \in [1, n] : \hat{y}_{T+k,x} < \hat{\hat{y}}_{T+k}\}$. Since $\hat{y}_{T+k,1} < \hat{\hat{y}}_{T+k}$, $1 \in S$, making S non-empty. Also, S is bounded above by n, so by the completeness property of the real numbers, S has a least upper bound, say i^*, where $i^* \in [1, n]$.

To show $\hat{y}_{T+k,i^*} = \hat{\hat{y}}_{T+k}$, we proceed by contradiction:

- If $\hat{y}_{T+k,i^*} < \hat{\hat{y}}_{T+k}$, then by the continuity of $\hat{y}_{T+k,i}$, there exists a $\delta > 0$ such that for all x in $(i^*, i^* + \delta)$, $\hat{y}_{T+k,x} < \hat{\hat{y}}_{T+k}$. This contradicts i^* being the least upper bound of S.
- If $\hat{y}_{T+k,i^*} > \hat{\hat{y}}_{T+k}$, then there exists a $\delta > 0$ such that for all x in $(i^* - \delta, i^*)$, $\hat{y}_{T+k,x} > \hat{\hat{y}}_{T+k}$. This implies there exists some $x \in S$ with $x > i^*$, contradicting i^* being the least upper bound of S.

The only remaining possibility is $\hat{y}_{T+k,i^*} = \hat{\hat{y}}_{T+k}$

□

A1.2 Proof of Corollary 1

Proof. Assume without loss of generality that $\hat{y}_{T+k,0}$ and $\hat{y}_{T+k,n}$ are not both the maximum or both the minimum of $\hat{y}_{T+k,i}$. By assumption $\hat{y}_{T+k,0} \leq \hat{\hat{y}}_{T+k} \leq \hat{y}_{T+k,n}$. There exists at least one maximum value $\dot{y}_{T+k,i^{max}}$ and one minimum value $\dot{y}_{T+k,i^{min}}$ of the linear piece wise interpolation, and by definition $\dot{y}_{T+k,i^{min}} \leq \hat{\hat{y}}_{T+k} \leq \dot{y}_{T+k,i^{max}}$. By the Intermediate Value Theorem (IVT) if $\dot{y}_{T+k,i}$ is a continuous function on an interval $i \in [0, n]$, then there exists at least one point $i^{\#}$ between $\hat{y}_{T+k,0}$ and $\hat{y}_{T+k,n}$ for which $\dot{y}_{T+k,i^{\#}} = \hat{\hat{y}}_{T+k}$. □

A1.3 Proof of Proposition 1

Proof. Under bottom-up, $\mathbf{E}_{T,n}\left(y_{T+k,i}\right) = \rho^{(k-1)n+i} y_{T,n}$. Suppose $k=1$, then

$$\begin{aligned}\sum_{i=1}^{n}\left[y_{T+1,i}-\mathbf{E}_{T,n}\left(y_{T+1,i}\right)\right] = y_{T+1,1}-\mathbf{E}_{T,n}\left(y_{T+1,1}\right)+y_{T+1,2}-\mathbf{E}_{T,n}\left(y_{T+1,2}\right) \\ +\cdots+y_{T+1,n}-\mathbf{E}_{T,n}\left(y_{T+1,n}\right).\end{aligned} \tag{30}$$

Note that

$$\begin{aligned}y_{T+1,1}-\mathbf{E}_{T,n}\left(y_{T+1,1}\right) &= \rho y_{T,n}+\varepsilon_{T+1,1}+\rho y_{T,n}=\varepsilon_{T+1,1} \\ y_{T+1,2}-\mathbf{E}_{T,n}\left(y_{T+1,2}\right) &= \rho y_{T+1,1}+\varepsilon_{T+1,2}+\rho^{2} y_{T,n}=\rho\varepsilon_{T+1,1}+\varepsilon_{T+1,2}=\sum_{j=1}^{2}\rho^{2-i}\varepsilon_{T+1,i} \\ \vdots &= \vdots \\ y_{T+1,n}-\mathbf{E}_{T,n}\left(y_{T+1,n}\right) &= \sum_{j=1}^{n}\rho^{n-i}\varepsilon_{T+1,i}.\end{aligned}$$

Therefore,

$$\sum_{i=1}^{n}\left[y_{T+1,i}-\mathbf{E}_{T,n}\left(y_{T+1,i}\right)\right]=\varepsilon_{T+1,1}+\sum_{j=1}^{2}\rho^{2-i}\varepsilon_{T+1,i}+\cdots+\sum_{j=1}^{n}\rho^{n-i}\varepsilon_{T+1,i}=\sum_{i=1}^{n}\sum_{j=1}^{i}\rho^{i-j}\varepsilon_{T+1,j}.$$

Hence, the relation in Claim 1 holds. Assume the relation in Claim 1 holds for $k-1$, then

$$\begin{aligned}\sum_{i=1}^{n}\left[y_{T+(k-1),i}-\mathbf{E}_{T,n}\left(y_{T+(k-1),i}\right)\right] &= y_{T+(k-1),1}-\mathbf{E}_{T,n}\left(y_{T+(k-1),1}\right)+y_{T+(k-1),2}-\mathbf{E}_{T,n}\left(y_{T+(k-1),2}\right) \\ &\quad+\cdots+y_{T+(k-1),n}-\mathbf{E}_{T,n}\left(y_{T+(k-1),n}\right) \\ &= \sum_{i=1}^{n}\sum_{j=1}^{i}\sum_{l=1}^{(k-1)}\rho^{((k-1)-l)n-j+i}\varepsilon_{T+l,j}\end{aligned} \tag{31}$$

In particular, we have that

$$y_{T+(k-1),n}-\mathbf{E}_{T,n}\left(y_{T+(k-1),n}\right)=\rho^{(k-1)n}y_{T,n}+\sum_{j=1}^{n}\sum_{l=1}^{k-1}\rho^{(k-l)n-j}\varepsilon_{T+l,j}-\rho^{(k-1)n}y_{T,n}. \tag{32}$$

We show now that the relation holds for any k. Consider

$$
\begin{aligned}
y_{T+k,1} - \mathbf{E}_{T,n}\left(y_{T+k,1}\right) &= \rho y_{T+k-1,n} + \varepsilon_{T+k,1} - \rho^{(k-1)n+1} y_{T,n} \\
&= \rho \left[\rho^{(k-1)n} y_{T,n} + \sum_{j=1}^{n} \sum_{l=1}^{k-1} \rho^{(k-l)n-j} \varepsilon_{T+l,j} \right] + \varepsilon_{T+k,1} - \rho^{(k-1)n+1} y_{T,n} \\
&= \sum_{j=1}^{n} \sum_{l=1}^{k-1} \rho^{(k-l)n-j+1} \varepsilon_{T+l,j} + \varepsilon_{T+k,1} \\
y_{T+k,2} - \mathbf{E}_{T,n}\left(y_{T+k,2}\right) &= \rho y_{T+k,1} + \varepsilon_{T+k,2} - \rho^{(k-1)n+2} y_{T,n} \\
&= \rho^2 y_{T+k-1,n} + \rho \varepsilon_{T+k,1} + \varepsilon_{T+k,2} - \rho^{(k-1)n+2} y_{T,n} \\
&= \rho^2 \left[\rho^{(k-1)n} y_{T,n} + \sum_{j=1}^{n} \sum_{l=1}^{k-1} \rho^{(k-l)n-j} \varepsilon_{T+l,j} \right] + \rho \varepsilon_{T+k,1} + \varepsilon_{T+k,2} - \rho^{(k-1)n+2} y_{T,n} \\
&= \sum_{j=1}^{n} \sum_{l=1}^{k-1} \rho^{(k-l)n-j+2} \varepsilon_{T+l,j} + \sum_{j=1}^{2} \rho^{2-j} \varepsilon_{T+k,j} \\
\vdots \quad &= \quad \vdots \\
y_{T+k,n} - \mathbf{E}_{T,n}\left(y_{T+k,n}\right) &= \sum_{j=1}^{n} \sum_{l=1}^{k-1} \rho^{(k-l)n-j+n} \varepsilon_{T+l,j} + \sum_{j=1}^{n} \rho^{n-j} \varepsilon_{T+k,j} \\
&= \sum_{j=1}^{n} \sum_{l=1}^{k-1} \rho^{(k-l)n-j+n} \varepsilon_{T+l,j} + \sum_{j=1}^{n} \rho^{n-j} \varepsilon_{T+k,j} \\
&= \sum_{j=1}^{n} \sum_{l=1}^{k-1} \rho^{(k-l)n-j+n} \varepsilon_{T+l,j} + \sum_{j=1}^{n} \rho^{n-j} \varepsilon_{T+k,j} \,,
\end{aligned}
$$

using (32) in the first equality above. Hence, we have that

$$
\begin{aligned}
\sum_{i=1}^{n} \left[y_{T+k,i} - \mathbf{E}_{T,n}\left(y_{T+k,i}\right) \right] = \sum_{i=1}^{n} \Bigg[& \sum_{j=1}^{i} \sum_{l=1}^{k-1} \rho^{(k-l)n-j+1} \varepsilon_{T+l,j} + \varepsilon_{T+k,1} \\
& + \sum_{j=1}^{i} \sum_{l=1}^{k-1} \rho^{(k-l)n-j+2} \varepsilon_{T+l,j} + \sum_{j=1}^{2} \rho^{2-j} \varepsilon_{T+k,j} \\
& + \cdots \qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad (33) \\
& + \sum_{j=1}^{i} \sum_{l=1}^{k-1} \rho^{(k-l)n-j+n} \varepsilon_{T+l,j} + \sum_{j=1}^{n} \rho^{n-j} \varepsilon_{T+k,j} \Bigg] \\
= & \sum_{i=1}^{n} \sum_{j=1}^{i} \sum_{l=1}^{k} \rho^{(k-l)n-j+i} \varepsilon_{T+l,j} \,.
\end{aligned}
$$

$\square$

A1.4 Proof of Proposition 2

Proof. First, for $k = 1$, we have

$$\sum_{i=1}^{n} \left[y_{T+1,i} - \mathbf{E}_{T,n}\left[y_{T+1,n}\right]\right] = y_{T+1,1} - \mathbf{E}_{T,n}\left[y_{T+1,n}\right] + y_{T+1,2} - \mathbf{E}_{T,n}\left[y_{T+1,n}\right] + \cdots + y_{T+1,n} - \mathbf{E}_{T,n}\left[y_{T+1,n}\right]$$

where

$$\begin{aligned}
y_{T+1,1} - \mathbf{E}_{T,n}\left[y_{T+1,n}\right] &= \rho y_{T,n} + \epsilon_{T+1,1} - \rho^n y_{T,n} = \rho y_{T,n}(1-\rho^{n-1}) + \epsilon_{T+1,1} \\
y_{T+1,2} - \mathbf{E}_{T,n}\left[y_{T+1,n}\right] &= \rho y_{T+1,1} + \epsilon_{T+1,2} - \rho^n y_{T,n} = \rho^2 y_{T,n} + \rho\epsilon_{T+1,1} + \epsilon_{T+1,2} - \rho^n y_{T,n} \\
&= \rho^2 y_{T,n}(1-\rho^{n-2}) + \sum_{i=1}^{2} \rho^{2-i}\epsilon_{T+1,i} \\
\vdots \quad &= \quad \vdots \\
y_{T+1,n} - \mathbf{E}_{T,n}\left[y_{T+1,n}\right] &= \sum_{i=1}^{n} \rho^{n-i}\epsilon_{T+1,i}\,.
\end{aligned}$$

Therefore,

$$\begin{aligned}
\sum_{i=1}^{n} \left[y_{T+1,i} - \mathbf{E}_{T,n}\left[y_{T+1,n}\right]\right] = {}& \rho y_{T,n}(1-\rho^{n-1}) + \epsilon_{T+1,1} \\
&+ \rho^2 y_{T,n}(1-\rho^{n-2}) + \sum_{i=1}^{2} \rho^{2-i}\epsilon_{T+1,i} \\
&+ \cdots \\
&+ \sum_{i=1}^{n} \rho^{n-i}\epsilon_{T+1,i} \\
&= \sum_{i=1}^{n-1} \rho^{i} y_{T,n}(1-\rho^{n-i}) + \sum_{i=1}^{n}\sum_{j=1}^{i} \rho^{i-j}\epsilon_{T+1,j}.
\end{aligned}$$

The relation in Claim 2 holds. Assume it holds for $k-1$. As in Claim 1, we have that

$$y_{T+(k-1),n} - \mathbf{E}_{T,n}\left[y_{T+(k-1),n}\right] = y_{T,n}\left(\rho^{(k-1-1)n+n} - \rho^{(k-1)n}\right) + \sum_{j=1}^{n}\sum_{l=1}^{k-1}\rho^{(k-l)n-j}\varepsilon_{T+l,j} \tag{34}$$

We show now that the relation holds for any k. Consider

$$\begin{aligned}
y_{T+k,1} - \mathbf{E}_{T,n}\left[y_{T+k,n}\right] &= \rho y_{T+k-1,n} + \varepsilon_{T+k,1} - \rho^{kn} y_{T,n} \\
&= y_{T,n}\left(\rho^{(k-1-1)n+n+1} - \rho^{(k-1)n+1}\right) + \sum_{j=1}^{n}\sum_{l=1}^{k-1}\rho^{(k-l)n-j+1}\varepsilon_{T+l,j} + \varepsilon_{T+k,1} \\
y_{T+k,2} - \mathbf{E}_{T,n}\left[y_{T+k,n}\right] &= y_{T,n}\left(\rho^{(k-1-1)n+n+2} - \rho^{(k-1)n+2}\right) + \sum_{j=1}^{n}\sum_{l=1}^{k-1}\rho^{(k-l)n-j+2}\varepsilon_{T+l,j} + \sum_{j=1}^{2}\rho^{2-j}\varepsilon_{T+k,j} \\
\vdots &= \vdots \\
y_{T+1,n} - \mathbf{E}_{T,n}\left[y_{T+1,n}\right] &= y_{T,n}\left(\rho^{(k-1-1)n+n+n} - \rho^{(k-1)n+n}\right) + \sum_{j=1}^{n}\sum_{l=1}^{k-1}\rho^{(k-l)n-j+n}\varepsilon_{T+l,j} + \sum_{j=1}^{n}\rho^{n-j}\varepsilon_{T+k,j} \\
&= y_{T,n}\left(\rho^{(k-1)n+n} - \rho^{kn}\right) + \sum_{j=1}^{n}\sum_{l=1}^{k-1}\rho^{(k-l)n-j+n}\varepsilon_{T+l,j} + \sum_{j=1}^{n}\rho^{n-j}\varepsilon_{T+k,j}
\end{aligned}$$

using (34) in the first equality above. Therefore,

$$\sum_{i=1}^{n}\left[y_{T+k,i} - \mathbf{E}_{T,n}\left(y_{T+k,n}\right)\right] = \sum_{i=1}^{n} y_{T,n}\left(\rho^{(k-1)n+i} - \rho^{kn}\right) + \sum_{i=1}^{n}\sum_{j=1}^{i}\sum_{l=1}^{k}\rho^{(k-l)n-j+i}\varepsilon_{T+l,j}\,.$$

□

A2 Data Appendix

This appendix describes the construction of the data series used in the empirical exercises.

Copper spot prices Daily closing spot prices of grade A copper at the London Metal Exchange were obtained from Bloomberg (LMCADY). Monthly average data used for backcasting is obtained from the World Bank Pink sheets. https://www.worldbank.org/en/research/commodity-markets

Interest rate Daily data on interest rates were obtained from Board of Governors of the Federal Reserve System. The long-term rate is the Market Yield on U.S. Treasury Securities at 10-Year Constant Maturity, Quoted on an Investment Basis (DGS10), retrieved from FRED, Federal Reserve Bank of St. Louis.

OECD CLI The Composite Leading Indicator (CLI) for the G20 countries is provided by the Organisation for Economic Co-operation and Development, https://doi.org/10.1787/4a174487-en. The CLI provides early signals of turning points in business cycles.

Consumer price index Real-time vintages of the seasonally adjusted U.S. consumer price index are obtained from the real-time database of the Philadelphia Federal Reserve.

Nowcasts of CPI for real forecasts Missing real-time observations for the CPIs are nowcasted using the average historical growth rate from 1973M1. All vintages of U.S. CPI are available and are observed with a one-month publication delay.

A3 Robustness Results

A3.1 Comparison of No-Change Forecasts

Table A1. End-of-Month Versus Monthly Average No-change Forecasts

Series	Simulated Ran. Walk	10 Year Bonds	Copper
Horizon		MSFE Ratio	
1	0.54	0.56 (0.000)	0.55 (0.000)
3	0.89	0.84 (0.000)	0.89 (0.007)
6	0.95	0.92 (0.000)	0.96 (0.027)
12	0.97	0.95 (0.000)	0.99 (0.321)
		Success Ratio	
1	0.74	0.75 (0.000)	0.71 (0.000)
3	0.61	0.65 (0.000)	0.60 (0.000)
6	0.58	0.66 (0.000)	0.55 (0.020)
12	0.55	0.64 (0.000)	0.56 (0.004)

Note: Real-time, out-of-sample forecasts in levels, 1992M1–2021M1. End-of-month no-change forecasts relative to the monthly average no-change forecast, with p-values reported in parentheses.

A3.2 Alternative Averaging

Table A2. Alternative End-of-Month No-changes Versus Monthly Average No-change Forecasts

Series	Simulated Ran. Walk	10 Year Bonds	Copper
Average		MSFE Ratio	
EoM	0.54	0.56 (0.000)	0.55 (0.000)
2 days	0.55	0.58 (0.000)	0.56 (0.000)
1 week	0.62	0.65 (0.000)	0.62 (0.000)
2 weeks	0.74	0.76 (0.000)	0.73 (0.000)
		Success Ratio	
EoM	0.74	0.75 (0.000)	0.71 (0.000)
2 days	0.73	0.73 (0.000)	0.73 (0.000)
1 week	0.71	0.72 (0.000)	0.72 (0.000)
2 weeks	0.68	0.67 (0.000)	0.68 (0.000)

Note: Real-time, out-of-sample forecasts in levels, 1992M1–2021M1. End-of-month no-change forecasts relative to the monthly average no-change forecast, with p-values reported in parentheses.

A3.3 Real Price of Copper Backcasts

Table A3. Real-Time Model-Based Forecasts of the Real Price of Copper, no Monthly Backcast

Method	Aggregate	Bottom-up	PEPSI	PEPSI(est)	EoP
Data	Average	Daily	EoM	EoM	EoM
Horizon			MSFE Ratio		
1	1.76 (0.998)	**0.99 (0.299)**	**0.98 (0.307)**	1.01 (0.636)	1.05 (0.731)
3	1.14 (0.960)	1.03 (0.872)	1.01 (0.603)	1.07 (0.852)	1.06 (0.755)
6	1.20 (0.980)	1.06 (0.906)	1.14 (0.959)	1.16 (0.976)	1.18 (0.968)
12	1.26 (0.990)	1.14 (0.958)	1.25 (0.984)	1.21 (0.969)	1.29 (0.985)
24	1.47 (1.000)	1.29 (0.998)	1.51 (1.000)	1.19 (0.926)	1.55 (1.000)
			Success Ratio		
1	**0.51 (0.352)**	**0.54 (0.081)**	**0.52 (0.191)**	0.50 (0.439)	**0.52 (0.191)**
3	0.48 (0.759)	0.50 (0.726)	**0.52 (0.305)**	0.48 (0.775)	**0.54 (0.241)**
6	0.45 (0.980)	0.50 (0.720)	0.47 (0.880)	0.45 (0.944)	0.47 (0.923)
12	0.44 (0.893)	0.42 (0.977)	0.41 (1.000)	0.46 (0.853)	0.41 (1.000)
24	0.38 (1.000)	0.38 (1.000)	0.41 (1.000)	0.45 (0.942)	0.41 (1.000)

Note: Real-time, out-of-sample forecasts of the monthly average real price of copper in levels, 2000M1–2021M1. End-of-period (EoP) uses end-of-period forecasts as the forecast of the average. AR(12) selected using AIC. Monthly value begin in 1986 and are not backcast. Forecast criteria expressed relative to the end-of-month no-change forecast, with p-values reported in parentheses. Bold values indicate forecast gains relative to the no-change forecast for the MSFE ratio and gain in directional accuracy for the success ratio. "Backcast" data refers to backcasting monthly series using monthly average data to 1973M1.

A3.4 Alternative Parameterizations

Table A4. Real-Time ARMA(1,1) Forecasts of the Real Price of Copper

Method	Aggregate	Bottom-up	PEPSI	PEPSI(est)	EoP
Data	Average	Daily	EoM	EoM	EoM
Horizon			MSFE Ratio		
1	1.66 (1.000)	1.01 (0.919)	**0.96 (0.084)**	**0.98 (0.356)**	**0.96 (0.248)**
3	1.06 (0.862)	1.03 (0.896)	**0.96 (0.156)**	1.00 (0.451)	**0.96 (0.218)**
6	**0.98 (0.408)**	1.06 (0.908)	**0.94 (0.197)**	1.02 (0.673)	**0.94 (0.234)**
12	**0.93 (0.249)**	1.13 (0.958)	**0.91 (0.189)**	1.01 (0.566)	**0.92 (0.215)**
24	**0.85 (0.025)**	1.29 (0.998)	**0.84 (0.016)**	**0.88 (0.015)**	**0.84 (0.019)**
			Success Ratio		
1	**0.52 (0.249)**	0.47 (0.837)	0.47 (0.842)	0.47 (0.842)	0.47 (0.842)
3	**0.52 (0.213)**	0.49 (0.755)	**0.54 (0.141)**	**0.52 (0.213)**	**0.54 (0.100)**
6	**0.55 (0.089)**	0.49 (0.819)	**0.55 (0.083)**	**0.54 (0.021)**	**0.55 (0.072)**
12	**0.61 (0.031)**	0.42 (0.985)	**0.64 (0.008)**	**0.56 (0.099)**	**0.64 (0.006)**
24	**0.66 (0.004)**	0.38 (1.000)	**0.68 (0.001)**	**0.62 (0.000)**	**0.68 (0.001)**

Note: Real-time, out-of-sample forecasts of the monthly average real price of copper in levels, 2000M1–2021M1, using an ARMA(1,1). End-of-period (EoP) uses end-of-period forecasts as the forecast of the average. Monthly value begin in 1973 and are backcast. Forecast criteria expressed relative to the end-of-month no-change forecast, with p-values reported in parentheses. Bold values indicate forecast gains relative to the no-change forecast for the MSFE ratio and gain in directional accuracy for the success ratio. "Backcast" data refers to backcasting monthly series using monthly average data to 1973M1.

Table A5. Real-Time ARMA(1,1) Forecasts of the Nominal 10-Year Treasury Bonds

Method	Aggregate	Bottom-up	PEPSI	PEPSI(est)	EoP
Data	Average	Daily	EoM	EoM	EoM
Horizon			MSFE Ratio		
1	1.77 (1.000)	1.01 (0.886)	1.00 (0.546)	**0.99 (0.295)**	1.03 (0.840)
3	1.08 (0.986)	1.00 (0.716)	1.01 (0.854)	1.05 (0.900)	1.01 (0.861)
6	1.04 (0.958)	1.00 (0.427)	1.00 (0.619)	1.17 (0.990)	1.01 (0.610)
12	1.03 (0.872)	**0.98 (0.205)**	**0.98 (0.252)**	1.43 (0.999)	**0.98 (0.259)**
24	**0.98 (0.293)**	**0.96 (0.178)**	**0.96 (0.192)**	1.06 (0.865)	**0.96 (0.198)**
			Success Ratio		
1	0.48 (0.650)	0.44 (0.997)	**0.52 (0.115)**	**0.52 (0.115)**	**0.52 (0.115)**
3	**0.53 (0.082)**	0.49 (1.000)	0.47 (0.816)	0.47 (0.865)	0.48 (0.745)
6	0.50 (0.552)	**0.53 (1.000)**	**0.53 (0.577)**	0.46 (1.000)	**0.53 (0.608)**
12	**0.51 (0.841)**	**0.63 (1.000)**	**0.61 (0.984)**	0.37 (1.000)	**0.61 (0.984)**
24	0.50 (0.993)	**0.58 (1.000)**	**0.57 (1.000)**	0.49 (0.565)	**0.57 (1.000)**

Note: Real-time, out-of-sample forecasts of the nominal monthly average 10-year Treasury bonds in levels, 2000M1–2021M1, using an ARMA(1,1). End-of-period (EoP) uses end-of-period forecasts as the forecast of the average. Forecast criteria expressed relative to the end-of-month no-change forecast, with p-values reported in parentheses. Bold values indicate forecast gains relative to the no-change forecast for the MSFE ratio and gain in directional accuracy for the success ratio.