

# FORESIGHT

THE INTERNATIONAL JOURNAL OF APPLIED FORECASTING

## ARTICLES *by* FORECASTERS *for* FORECASTERS: issue 65

Representativeness: A New Criterion for  
Selecting Forecasts



Join the *Foresight* readership by becoming a  
member of the International Institute of Forecasters  
[forecasters.org/foresight/](http://forecasters.org/foresight/)

*made available to you with permission from the publisher*

# SPECIAL FEATURE: SELECTING A BEST MODEL

## Representativeness: A New Criterion for Selecting Forecasts

FOTIOS PETROPOULOS AND ENNO SIEMSEN

**PREVIEW** *Statistical criteria for selecting a best forecasting method from a group of candidates have been proposed, studied, and implemented widely in forecasting software. Very well-known are information criteria, such as the AIC, which balance performance and complexity, and validation techniques, which examine forecasting performance in a holdout sample. So it's a breath of fresh air to have a distinctly new take on method selection, which is what Fotios and Enno are presenting here. They offer strong evidence that method selection can be improved by accounting for the representativeness of the forecasts.*

### INTRODUCTION: THE CHALLENGE OF SELECTING A MODEL

It's the age-old dilemma: What is the best forecasting model for you? This has always been a challenging question for researchers and practitioners alike. Being able to identify the best model could lead to substantial performance benefits. Further, performance-ranking different forecasting models supports setting better weights for forecast combinations.

A fundamental problem, however, is that the model that has performed comparatively well given the observed (historical) data will not necessarily be the model that performs best in the future. The underlying patterns in observed data are always subject to change.

While forecast model selection is important, the forecasting community has not fully coalesced around the proper ways to accomplish this, and several approaches have been developed. One aspect these approaches have in common is that they focus wholly on the observed data – not considering their forecasts of the future in the selection decision.

We believe this is a significant oversight. In this paper, after discussion of existing

approaches to model selection, we propose an enhancement that incorporates the *representativeness* of a model's forecasts as a new component in the model-selection decision. It does so by providing a reality check on the reasonableness of the generated forecasts. Testing on numerous actual time series shows this enhancement to improve forecast performance.

### JUDGMENT IN MODEL SELECTION

Recent studies have shown that individuals are able to select models such that the average forecasting performance is as good as selection based on some statistical criteria. An examination of judgment in model selection was featured in the Summer 2019 issue of *Foresight* (Issue 54), beginning with Fotios's article on the application of judgment for model selection (Petropoulos, 2019). Several Commentaries pointed out that the commendable performance of judgmental model selection offers an attractive alternative to the reliance on automatic selection in forecasting software as well as a way forward for those with algorithmic aversion.

## Key Points

- The forecasting community has not fully coalesced around the proper ways to select a forecasting method, and several approaches have been developed. None of the established model-selection rules make use of the forecasts that will eventually be used for decision making. Information criteria and validation/cross-validation approaches explicitly assume that whichever method produced the best forecasts previously will be the best model forward.
- Our new approach for selecting among forecast models is based on examination of the forecasts made for real-time future periods, the actuals for which are not yet knowable. It is based on *representativeness*, the degree to which these forecasts are a natural continuation of the observed data.
- We describe and illustrate a new criterion for method selection that considers the representative of the forecasts as well as the accuracy with which the methods fit the observed data. We call this the REP.
- Finally, we compare REP with two main existing criteria for method selection, the AIC and cross-validation (CV) using a large number and wide variety of time series from previous M competitions. We believe the results are highly promising and point to deeper exploration into the psychology of human input into forecast-method selection.

In addition, Fotios reported that combining forecasts, either by aggregating the forecasts of small groups (judgmental aggregation) or averaging the forecasts of models selected statistically with those selected judgmentally, could significantly outperform straight statistical approaches. Moreover, while he found that statistical approaches will select the best models more often (compared to humans), they will also select the worst model more frequently. In other words, humans are better able to avoid the worst outcomes.

The way we apply judgment to the task of model selection is fundamentally different from the way statistical approaches such as information criteria work. The key difference seems to lie in what information is being used. Statistical approaches work “backwards” in the sense of examining observed data values and measuring forecasting performance over past periods (both in terms of performance in tracking the in-sample fit or the out-of-sample performance on data held out from model fit).

We humans, in contrast, look “forward” and compare the pattern of forecasts produced by each forecasting model with our mental extrapolations of the data pattern. Brain imaging experiments by Weiwei Han and colleagues (2019) revealed that humans reject models for which the forecasts look unreasonable. In other words, we use a visualization of forecasts for which the actual data are not yet available to select models that produce a pattern that best matches past data: the forecasts that are *representative*.

Motivated by this realization, we devised an algorithmic approach to forecast selection based on the concept of representativeness. A full description and analysis of the algorithm is in our new article in the journal *Management Science* (Petropoulos and Siemsen, 2022). Here we will describe and illustrate how model selection works based on this new criterion of representativeness. We begin with an overview of model-selection criteria.

### MODEL-SELECTION CRITERIA

Most statistical criteria for model selection are based, directly or indirectly, on forecast-error metrics, with smaller errors over a designated period being preferable.

#### *Information Criteria*

Information criteria are based on how well the model fits the in-sample data but contains a penalty for the size/complexity of the model. The penalty is designed to avoid overfitting – introduction of additional model complexity that adds little new value.

Information criteria, such as the Akaike's Information Criterion (AIC) and Bayesian Information Criterion (BIC), are in wide use. Implementations for them are readily provided within open-source forecasting packages, such as the *forecast* and *smooth* packages for R statistical software. They are also reported in many commercial packages and used in some of these for model selection. Some versions build in corrections for small sample size. For instance, the functions `ets()` and `auto.arima()` of R's *forecast* package rely, by default, on the AIC values corrected for small sample sizes (AICc).

The in-sample fit of each model is measured either by the mean squared error (MSE) or the likelihood function (which measures how likely we would see the

In addition, information criteria may erroneously select a simpler model when the sample size is small and the sample data are highly variable. That said, the selection of a misspecified model – either one that is simpler or one that is more complex – will not necessarily degrade forecast performance, a point we return to later in this article.

Baki Billah and colleagues (2006) compared different statistical criteria for choosing among exponential-smoothing models and reported that information criteria, particularly the AIC, performed best in their simulations. Stephan Kolasza (2011) showed how the values of information criteria may be used to calculate weights for combining the forecasts of different models.

**The existing selection criteria all consider past data as well as the forecasts that correspond to observations up until the latest available time period. None of these approaches, however, make use of the forecasts for future periods.**

observed data, given the model). The penalty is based on the number of model parameters to be estimated. The premise of an information criterion follows Occam's Razor in that, among all solutions with the same performance, the simplest one should be selected. In comparing the values of an information criterion across models, lower values indicate superior performance.

One caveat is that information criteria values are not always directly comparable across different models, especially if we are comparing models from different families, such as an exponential smoothing vs. an ARIMA model. The difficulty arises from how the likelihood function is computed and how initial values have been specified; for instance, differencing is often required for ARIMA models but not exponential smoothing. In general, then, information criteria should not be used when the data are transformed in different ways, or when different orders of integration have been applied.

### **Validation and Cross-Validation**

*Validation* and *cross-validation* approaches split the available data into training and validation sets, using the former to fit a model and the latter to measure the out-of-sample performance of the forecasts. The model with the best validation performance is then taken forward to produce forecasts for future periods.

The principle of out-of-sample evaluation is that forecasting performance should be measured on data that have not been used in the training of the models.

In the validation approach, only one set of forecasts is produced for a time series, and these are tested against the data in the validation set of that series. This approach is simple and relatively fast; however, it is still slower than selection based on information criteria, as the selected model needs to be fitted twice: once for the validation step and once to produce forecasts for future periods. In addition, the forecast-error metrics that are calculated blend multiple forecast horizons

from a fixed origin, so no insights emerge as to how forecast accuracy changes as the horizon of the forecast lengthens.

A better alternative to validation is cross-validation, in which model fitting and evaluation is repeated from multiple forecast origins. While there are several forms in use, for time-series data the rolling origin evaluation is most appropriate. Here, a model is fit with the training data and forecasts made for multiple horizons in the validation set. The window of the training data is expanded to include the first time period in the validation set. The model is reestimated and forecasts made from the new origin. The process continues until all data in the validation set have been incorporated into the training set. The procedure can be time-consuming and requires the availability of long time series to support evaluations over multiple origins.

Additionally, for cross-validation in general, one needs to decide between expanding or rolling training windows, overlapping or nonoverlapping validation sets, and how often validation forecasts should be produced. For additional reading on

series in the reference set is split into a training and a test set (Talagala and colleagues, 2021). The training set of each series is used to calculate values for the time-series features as well as to produce forecasts for the test set, and this is done for each of the models in the pool.

A meta-learning model is trained to select between models by comparing how performance of the forecasts of the various models in the pool is related to the values of the features. The reference series can be publicly available data sets, such as the M4 competition data, or even synthetically generated series that possess the desired features (Kang and colleagues, 2020). The R packages *tsfeatures* and *gratis* can be used to calculate the values of time-series features and generate synthetic series based on them.

A variation of this approach is the rule-based forecasting RBF system from Armstrong, Collopy, and Adya. These are sets of “if/then” rules distilled from experience of forecasting experts to select and combine among a set of simple time-series models (Adya and colleagues, 2001).

## Representativeness of a model's forecasts as a new component in the model-selection decision provides a reality check on the reasonableness of the generated forecasts.

time-series validation and cross-validation, the reader is referred to Tashman (2000) and Bergmeir and Benítez (2012).

### ***Times-Series Features***

Apart from information criteria and validation approaches, model selection has been based on *time-series features*, such as trend, seasonality, autocorrelation, cycle, randomness, series length and variability, and interdemand interval for intermittent data (Petropoulos and colleagues, 2014). Since there is no universal method that is best for every time series, the goal is to identify the best forecasting method for the particular features of “my data.”

Given a pool of models and a set of reference series, the features are used to train *meta-learning algorithms*. Here each

### **SELECTION BY REPRESENTATIVENESS**

These various selection criteria all consider past data as well as the forecasts that correspond to observations up until the latest available time period. Information criteria use in-sample comparisons between actual and predicted – also called *fitted* – values. Validation approaches use out-of-sample forecasts; however, these forecasts still occur during past time periods (the validation sets correspond to appropriate holdouts).

None of these approaches, however, make use of the forecasts for *future* periods - that is, forecasts for which the actual observations remain unknown. More explicitly, none of the established model-selection rules make use of the forecasts

that will eventually be used for decision making. Information criteria and validation/cross-validation approaches explicitly assume that whichever method produced the best forecasts previously will be the best model forward.

This principle strikes us as, at the very least, naive. Our new approach for selecting among forecast models is based on examination of the forecasts made for real-time future periods, the actuals for which are not yet knowable. The evaluation of such forecasts, given that the actual data are not yet available, is based on the degree to which these forecasts are a natural continuation of the past observed data.

We call this *representativeness*.

For example, a flat/level forecast would not be representative of data with a strong trend, nor would a trend-only model be representative of data that exhibit strong seasonal behavior.

### **REP Defined**

We define the *representativeness gap* as the lack of representativeness of a forecast compared to past actuals, and use this concept in conjunction with the performance of the in-sample forecasts (fitted values). This approach effectively replaces the complexity penalty applied by information criteria by the representativeness gap. In essence, we propose that selecting between forecasting models should be a balance between in-sample fit and representativeness. Conceptually, this new criterion, REP, can be expressed as

$$\text{REP} = \text{performance gap} + \text{representativeness gap}$$

Similar to information criteria, model forecasts with lower REP values should be preferred to those with higher. If two sets

of forecasts offer the same in-sample performance, the one with better representativeness – a lower representativeness gap – should be selected. This is not the same as selecting the least complex of two equally performing models.

Representativeness (and its gap) can be measured in different ways, all of which are based on comparisons of the forecasts for the future with the available data up to the present. Such comparisons are asynchronous, in the sense that the forecasts and the data to be compared refer to different time periods. As such, some scaling and transformations may be required to place the “present” and the “future” on the same level.



For series with multiplicative (trend or seasonal) patterns, logarithmic transformations can be needed. If the data are periodic – with patterns that repeat at regular intervals such as monthly seasonality – the comparisons must be aligned, so that the respective periods (e.g.,

the same month) of the past data and the forecasts are compared.

Measurement of representativeness can not only be made for point forecasts but prediction intervals as well. Other choices include the error metrics used, the length of the forecast horizon (i.e., length of the data window for the comparisons), the use of a single window or multiple windows of past data, and the use of equal or unequal weights to average representativeness across multiple windows of data.

### **Calculating the REP for Point Forecasts**

The calculation of REP follows a number of steps, each of which can be adapted to

the specifics of the data at hand. **Figure 1** provides an illustrative example. The historical data are plotted in black in the top panel. These represent monthly sales of a toy product over a period of six years. An upward trend and a seasonal pattern are evident.

We produce forecasts with two methods: Holt’s linear trend exponential smoothing (that ignored the seasonal pattern) and the Holt-Winters exponential smoothing with multiplicative seasonality. The forecasts for the next year (next 12 months) are depicted by the red and blue lines in the first panel.

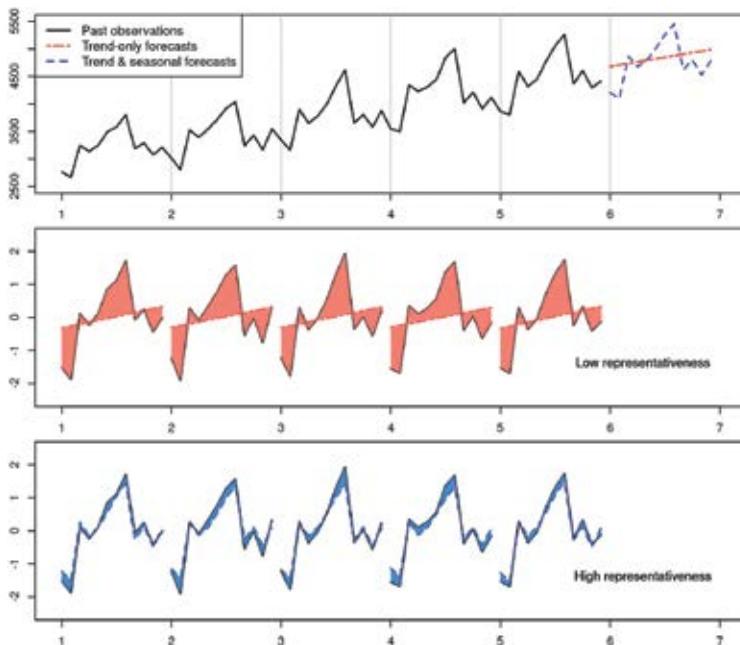
1. The first step is to produce forecasts using all available historical data, without splitting a series into training and validation sets.
2. Specify the window of forecasts over which representativeness is to be measured. It simplifies things to match the forecast horizon with the periodicity of the data, so for monthly data – a periodicity of 12 – we set the forecast horizon to be 12 months ahead.
3. Split the sample data into buckets such that (i) each bucket is at least as

long as the forecasts’ window, and (ii) the first period of each bucket corresponds to the respective first period of the forecasts. In the toy example, we split the in-sample data into five windows/buckets, each a length of 12 months to match the forecasting horizon and each bucket beginning with the same month of the year. This ensures that seasonal patterns will be aligned when comparisons are performed.

4. The fourth step is to perform any needed scaling and transformations. For each bucket in our illustrative data, we aligned the scales and did a Box-Cox transformation to stabilize the variance and transform multiplicative patterns into additive ones. This is to be done for the forecasts as well.
5. For each individual bucket, the representativeness gap is measured in terms of the closeness of the past data points to the point forecasts. Most simply, we sum the absolute distances between a (scaled and transformed) bucket of past data and the forecasts. For our toy example, the result of this step is five sums of absolute differences.
6. The final step is to take an average of the resulting bucket sums from the prior step. This could be an unweighted average, but we recommend a weighted mean to give more emphasis to the more recent years. In particular, we recommend that the weight for each bucket is one half of its more recent adjacent bucket; that is, the weights are decreasing by 50% as we move back from the most recent bucket.

In Figure 1, the middle frame highlights the representativeness gap for the trended-only forecast – the red line in the top frame. The large representativeness gap reflects mainly the seasonal departure of the data from the trend line. The bottom frame shows the forecasts from a trended/seasonal (Holt-Winters) model – the blue forecasts. It is clear that the

**Figure 1. A Toy Example on Measuring Representativeness (and Its Gap), adapted from Petropoulos and Siemsen (2022)**



representativeness gap is much lower now that we've also accounted for seasonality.

Once the representativeness gap has been calculated, then the value of the REP criterion is calculated as the sum of the representativeness gap and the in-sample performance gap. For the toy data, the in-sample performance gap is the mean of absolute differences between the actual and fitted values over the five years.

### PERFORMANCE COMPARISON OF SELECTION APPROACHES

To evaluate how our REP criterion fares against information criteria and cross-validation, we compared the three criteria on a large collection of real data from the M, M3, and M4 forecasting competitions. We used these three criteria to select and to combine models from the exponential-smoothing family. For the latter, we calculated weights based on how the individual models performed on the particular selection criterion.

**Table 1** offers a selective summary of the empirical results. (These were more fully reported in Petropoulos and Siemsen, 2022). For Table 1, we present forecast accuracy results for the yearly and monthly data frequencies. The metric reported is the mean absolute scaled error (MASE), originally proposed by Rob Hyndman (see Hyndman, 2006) as a scale-free metric suitable for measuring accuracy across multiple time series (including intermittent series).

For point forecasts, the REP selection rule outperforms both the information criterion and cross-validation, and does so for both the yearly and monthly data. For the yearly frequency, REP was more accurate than AICc (Akaike's Information Criterion corrected for small sample sizes) and cross-validation (CV) by 8.2% and 5.5% respectively. Differences in performance were statistically significant in most of the cases, especially for yearly, quarterly, and monthly data.

In addition, we found that REP selects the best (among a class of exponential smoothing) models more often than

**Table 1. Summary Accuracy Results, adapted from Petropoulos and Siemsen (2022)**

		Yearly Data	Monthly Data
Selection	AICc	3.405	0.941
	CV	3.307	0.921
	REP	3.125	0.918
Combination	AICc	3.351	0.933
	CV	3.300	0.916
	REP	3.101	0.906

information criteria, while more frequently avoiding the worst of these models. The good performance of REP was evident not only when it was used to select/combine models within the exponential-smoothing family, but also within the ARIMA family of models, or even between models of different families.

We note too that combinations of models based on REP outperformed AICc and CV in terms of estimating uncertainty as well as accuracy in the point forecasts, using the mean scaled interval score and a 95% confidence level.

We also performed sensitivity analyses to examine the performance of REP under different conditions. Two principal findings:

- First, we considered the case of using only the representativeness gap in measuring REP, leaving out the in-sample performance. We observed that even excluding the performance gap leaves the REP approach superior to selection based on the AICc. This is very important for cases in which the available forecasts are not accompanied by in-sample fits, as it is usual for purely judgmental forecasts.
- Second, we analyzed the performance of REP for different forecast horizons and found that REP performs strongly across all horizons (short, medium, and long) – and in fact, its advantage relative to the information and cross-validation criteria grows at longer horizons.



## CONCLUSION

As a means of selecting between forecast models, representativeness – with the strong empirical performance of REP – has earned its way into the criteria for model selection. Similar to information criteria, REP consists of two parts: how well the model fits the historical data and a penalty. Contrary to information criteria, the penalty is not based on model complexity but rather on the representativeness of the forecasts, the degree to which the resulting forecasts are perceived as a natural extension of the historical data.

Our study offers more evidence that the infusion of human judgment into algorithms, such as the manner in which we use visualizations of the forecasts, can improve the performance of both existing algorithms and judgment alone. That

said, there is still much to distill from how humans approach forecasting and how we can translate such insights to further improve our forecasting algorithms.

## REFERENCES

- Adya, M., Collopy, F., Armstrong, J.S. & Kennedy, M. (2001/4). Automatic Identification of Time Series Features for Rule-Based Forecasting, *International Journal of Forecasting*, 17(2), 143–157.
- Bergmeir, C. & Benítez, J.M. (2012). On the Use of Cross-Validation for Time Series Predictor Evaluation, *Information Sciences*, 191, 192–213.
- Billah, B., King, M.L., Snyder, R. & Koehler, A.B. (2006/4). Exponential Smoothing Model Selection for Forecasting, *International Journal of Forecasting*, 22(2), 239–247.
- Han, W., Wang, X., Petropoulos, F. & Wang, J. (2019). Brain Imaging and Forecasting: Insights from Judgmental Model Selection, *Omega*, 87, 1–9.
- Hyndman, R.J. (2006). Another Look at Measures of Forecast Accuracy for Intermittent Demand, *Foresight*, 4, 43–46.
- Kang, Y., Hyndman, R.J., Li, F. GRATIS: (2020). Generating Time Series with Diverse and Controllable Characteristics, *Statistical Analysis and Data Mining*, 13, 354–376.
- Kolassa, S. (2011). Combining Exponential Smoothing Forecasts Using Akaike Weights, *International Journal of Forecasting*, 27(2), 238–251.
- Petropoulos, F. (2019). Judgmental Model Selection, *Foresight*, 54, 4–10.
- Petropoulos, F., Makridakis, S., Assimakopoulos, V. & Nikolopoulos, K. (2014). “Horses for Courses” in Demand Forecasting, *European Journal of Operational Research*, 237, 152–163.
- Petropoulos, F. & Siemsen, E. (2022). Forecast Selection and Representativeness, *Management Science*, forthcoming.
- Talagala, T.S., Li, F. & Kang, Y. (2021). FFORMPP: Feature-Based Forecast Model Performance Prediction, *International Journal of Forecasting*, <https://doi.org/10.1016/j.ijforecast.2021.07.002>
- Tashman, L.J. (2000). Out-of-Sample Tests of Forecasting Accuracy: An Analysis and Review, *International Journal of Forecasting*, 16(4), 437–450.



**Fotios Petropoulos** is Professor of Management Science, University of Bath, and *Foresight* Editor for Forecasting Support Systems.

**f.petropoulos@bath.ac.uk**



**Enno Siemsen** is Professor of Operations and Information Management and Associate Dean at the University of Wisconsin School of Business.

**esiemsen@wisc.edu**

# Commentary on Representativeness

NIGEL HARVEY AND SHARI DE BAETS

## INTRODUCTION

Fotios and Enno introduce us to the *representativeness* criterion for selecting between different statistical-forecasting models. To use this criterion, we take samples from historic datasets and then compare the patterns in successive chunks of these data with the pattern in a similar-length chunk provided by the forecasts. The less similar the pattern in the forecasts is to the patterns in the historical data, the less representative the forecasts are of the data series: in other words, the bigger the representativeness gap. To obtain a value of the representativeness criterion (REP), the representativeness gap is added to the performance gap, the in-sample difference between forecasts and outcomes. The model with the lowest REP value is the one selected.

Fotios and Enno compare the success of the REP criterion with that of two established approaches, the Akaike information criterion (AIC) and cross-validation (CV), in selecting between models from the exponential smoothing family. To do this, they used forecasts produced by the models for real series drawn from those used in forecasting competitions and found that REP produced higher forecast accuracy than the other two approaches. Even using the representativeness gap alone, excluding the performance gap, produced higher forecast accuracy than using the information criterion.

that REP is more effective than AIC and CV not just in selecting between alternative exponential-smoothing models, but also in selecting between alternative ARIMA models. Furthermore, they went on to demonstrate that it is more effective than CV in selecting between heterogeneous model types (exponential smoothing, ARIMA, Theta model) for lower-frequency data (yearly, quarterly, monthly, weekly) and similar to CV for higher-frequency data (daily, hourly). Thus the REP approach generalizes well across model types.

What about generalization across series types? In their empirical work, the authors used 103,830 series; some were from those used in the M and the M3 forecasting competitions, but most came from the M4 forecasting competition. These series were drawn from a wide variety of different domains; for example, M4 series were taken from industries, services, tourism, imports and exports, demographics, education, labour and wages, government, households, bonds, stocks, insurances, loans, real estate, transportation, and natural resources and the environment. In fact, Spiliotis and colleagues (2020) provide evidence that the M4 series are statistically representative of real-world series. So we can say that the findings reported on representativeness hold, *on average*, for the types of series from which models are used to make forecasts in the real world.

**We can say that the findings reported on representativeness hold, *on average*, for the types of series from which models are used to make forecasts in the real world.**

## ARE THE RESULTS GENERALIZABLE?

These results are impressive. How generalizable are they? In their original presentation of the representativeness criterion, Petropoulos and Siemsen (2022) showed

Of course, this does not mean we can be certain that Fotios and Enno's conclusions will hold for a particular type of series. We know from the Makridakis forecasting competitions that there is no overall optimal way to make forecasts;

instead, particular approaches suit particular types of series better than others. Similarly, different ways of selecting between forecasting models may be appropriate for different types of series.

### NONLINEAR SERIES

---

Consider two types of series that might be interesting for future analysis: nonlinear series and trended series. In both these cases, patterns in past data may not be a perfect guide to the patterns we should expect in future data. Generally,

start by establishing that the series is indeed nonlinear and determining the nature of its nonlinearity by, for example, extracting its Hurst exponent. Doing so, however, requires analysis of a very long data series, which is not typically available, except in some applications, such as high-frequency trading. We can still use established forecasting models and expect them to produce adequate short-term forecasts; for example, with a Hurst exponent of 0.329, we know there is a good chance of the current trend continuing over the next few periods.

**The question would then be whether nonlinear series are better forecast by the simpler exponential models selected by AIC and CV or by the more complex ones selected by REP.**

M4 series showed a high degree of linearity (Spiliotis and colleagues, 2020), but there are specific domains (finance, meteorology) in which series are better characterized as nonlinear (Mandelbrot and Hudson, 2008).

Differences between nonlinear series can be characterized by their Hurst exponent ( $H$ ); for most financial series it varies between 0.3 and 0.7 (Sang and colleagues, 2001). An  $H$  value of 0.5 corresponds to a random walk; an  $H$  value greater than 0.5 means that the series changes direction less frequently than a random walk – it is *persistent*; an  $H$  value less than 0.5 means that it changes more frequently than a random walk – it is *anti-persistent*. For example, daily stock returns for Caterpillar can be characterised by a Hurst exponent of 0.329 (Sang and colleagues, 2001). This means that trends in the series are likely to continue for a while before reversing and continuing in the opposite direction, only to eventually reverse again. (Also, a correlogram of this series shows a long, slow decline, indicating long-range dependencies in the data.)

How should we proceed? It may look as if such series are describing a non-deterministic seasonal process, but these patterns are transient so they cannot be used to make adequate long-range forecasts. The best approach could be to

### MODEL SELECTION CRITERIA FOR NONLINEAR SERIES

---

Would REP provide the best way of selecting between models to provide forecasts from nonlinear series? Assuming (as the authors do) that the data-generation process is unknown to forecasters, the length of the forecast window would be set equal to the apparent periodicity in the data series, and the candidate models (e.g., from the exponential-smoothing family) would each provide forecasts for that window. The representativeness gap between each set of forecasts and data in each of the previous windows would then be calculated in the way that Fotios and Enno describe.

Given the apparent “seasonal” patterns in nonlinear series, the representativeness gap may be smallest for the exponential smoothing models that include seasonal component. Also, given the transience of these seasonal patterns, such models may fit better when the discount factor is higher (i.e., when greater emphasis is placed on comparing forecasts with more recent time windows).

What about the AIC approach? Petropoulos and Siemsen (2022) provide examples of where this approach fails: forecasts did not follow trends in the data series because of the penalty applied to complexity. Complex models are likely to

be needed in attempts to account for the behaviour of nonlinear series, and so the complexity penalty could result in AIC selecting a simpler model than that chosen by REP in this case as well.

How about the CV approach? It does not penalize model complexity, but it is likely to be affected by the transient nature of apparent trends and seasonalities in nonlinear series: patterns present in the training sets may not remain present in validation sets. As a result, CV may also select a simpler model than REP.

The question would then be whether nonlinear series are better forecast by the simpler exponential models selected by AIC and CV or by the more complex ones selected by REP. Here our arguments about the types of models that might be favoured by different selection criteria are, of course, speculative, but they are designed to support our point that it would be worth carrying out empirical research to compare different ways of selecting between models for particular types of series – in this case, nonlinear series.

### TRENDED SERIES

---

We now turn to trended series. The authors point out that their development of the REP criterion was at least partly triggered by findings from judgmental-forecasting research that indicate that people try to make their sequence of forecasts look like (i.e., have the same characteristics as) the data series. One major finding is that when making forecasts from trended series, people tend to place their

That knowledge tells them that nothing continues to increase or decrease forever (entropy apart). Instead, trends eventually become asymptotic or turn into cycles. Forecasters factor this knowledge into their judgments as *trend damping*. Statistical forecasters have also discovered the value of trend damping: Gardner and McKenzie (1985) showed that adding an ad hoc damping term can improve performance of exponential-smoothing models that take trends into account. In summary, the steepness of trends in past data cannot be considered to be a perfect guide to the steepness of trends in future data.

It is possible to add an ad hoc damping term to forecast sequences that contain trends. In Figure 1 of their paper, the trend-only and the trend and seasonal forecasts could be subject to a small amount of damping to improve their accuracy. However, this would have two effects. First, it would make the forecast sequence (window 6 in Figure 1) less representative of the data in the windows of the data series (windows 1-5 in Figure 1). As a result, the representativeness gap would increase. However, on the basis of Gardner and McKenzie's (1985) findings, we would expect the performance gap to decrease. If this latter effect is larger, REP would be lower when a damping term is used to improve forecasts. So our question is whether REP can select an exponential model with damping over one without damping. The AIC approach is less likely to do this because addition of the damping term will be subject to a complexity penalty.

### The steepness of trends in past data cannot be considered to be a perfect guide to the steepness of trends in future data.

forecasts below upward trend lines and above downward ones (Eggleton, 1982; Harvey and Reimers, 2013).

It appears that when forecasting from linear or exponentially increasing trends in data series, people take into account not only the data in front of them but also their knowledge of the real world.

### INFLUENCE ON JUDGMENTAL FORECASTING

---

Fotios and Enno have described how the innovative work on representativeness was at least partially triggered by research on judgmental forecasting. We can ask how their ideas might influence work on judgmental forecasting.

There has long been a debate about the relative effectiveness of purely judgmental, purely statistical, and hybrid approaches to forecasting. Lawrence and colleagues (1985) concluded from their studies that “judgmental extrapolation is on average no less accurate than statistical forecasting and, in a number of subgroups of the time series, was the most accurate” (p 25). In the M2 competition (Makridakis and colleagues, 1993), forecasts produced by purely statistical methods were compared with those produced by people who had access to statistical methods but who could judgmentally adjust the forecasts produced by those methods to take account of additional information. There was little difference between these two

types of forecasts at short horizons, but purely statistical forecasts tended to be superior at longer horizons. The authors conclude that “no judgmental revisions ought to be made to the quantitative forecasts without making sure beforehand about the need and value of such revisions” (pp. 17-18).

In conclusion, circumstances in which judgment adds quality to forecasting processes still need to be identified more precisely. However, it should be possible to use the REP criterion to select not just between purely statistical approaches but also between statistical approaches, pure judgmental forecasting, and hybrid approaches that combine statistical and judgmental contributions in various ways.

#### REFERENCES

- Eggleton, I.R.C. (1982). Intuitive Time-Series Extrapolation, *Journal of Accounting Research*, 20(1), 68–102.
- Gardner, E.S., Jr. & McKenzie, E. (1985). Forecasting Trends in Time Series, *Management Science*, 31(10), 1237–1246.
- Harvey, N. & Reimers, S. (2013). Trend Damping: Under-Adjustment, Experimental Artefact, or Adaptation to Features of the Natural Environment? *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 39(2), 598-607.
- Lawrence, M.J., Edmundson, R.H. & O'Connor, M.J. (1985). An Examination of the Accuracy of Judgmental Extrapolation of Time-Series, *International Journal of Forecasting*, 1(1), 25-35.
- Makridakis, S., Chatfield, C., Hibon, M., Lawrence, M., Mills, T., Ord, K. & Simmons, L.F. (1993). The M-2 Competition: A Real-Time Judgmentally Based Forecasting Study, *International Journal of Forecasting*, 9(1), 5-22.
- Mandelbrot, B.B. & Hudson, R.L. (2008). *The (Mis)behaviour of Markets: A Fractal View of Risk, Ruin and Reward*, London: Profile Books.
- Petropoulos, F. & Siemsen, E. (2022). Forecast Election and Representativeness, *Management Science*, In press.
- Sang H-W, Ma, T. & Wang, S-Z. (2001). Hurst Exponent Analysis of Financial Time Series, *Journal of Shanghai University (English edition)*, 5(4), 269–272.
- Spiliotis, E., Kouloumos, A., Assimakopoulos, V. & Makridakis, S. (2020). Are Forecasting Competitions Data Representative of Reality? *International Journal of Forecasting*, 36(1), 37-53.



**Nigel Harvey** is Professor of Judgment and Decision Research at University College London. His research deals with the many ways in which judgment is involved in the forecasting process.

[n.harvey@ucl.ac.uk](mailto:n.harvey@ucl.ac.uk)



**Shari De Baets** is a senior post-doctoral researcher in the Faculty of Economics and Business Administration at Ghent University. She works at the interface of operations management and behavioral science with a particular focus on judgmental forecasting and the human perception of algorithms in work life.

[Shari.DeBaets@UGent.be](mailto:Shari.DeBaets@UGent.be)