1999

Level Adjusted Exponential Smoothing: A Method For Judgmentally Adjusting Exponential Smoothing Models for Planned Discontinuities

Dan Williams
CUNY Bernard M Baruch College

Don Miller
Virginia Commonwealth University

LEVEL ADJUSTED EXPONENTIAL SMOOTHING: A METHOD FOR JUDGMENTALLY ADJUSTING EXPONENTIAL SMOOTHING MODELS FOR PLANNED DISCONTINUITIES

Dan W. Williams, Baruch College, City University of New York, 17 Lexington Ave, New York, NY, 10010 US
Don M. Miller, Virginia Commonwealth University, School of Business, Box 844000, Richmond, VA, 23284 US

ABSTRACT

Forecasters often make judgmental adjustments to exponential smoothing forecasts to account for the effects of a future planned change. While this approach may produce sound initial forecasts, it can result in diminished accuracy for forecast updates. A proposed technique lets the forecaster include policy change adjustments *within* an exponential smoothing model. For 20 real data series representing Virginia Medicaid expenses, initial forecasts and forecast updates are developed using the proposed technique and several alternatives, and they are updated through various simulated level shifts. The proposed technique was more accurate than the alternatives in updating forecasts when a shift in level occurs approximately as planned.

INTRODUCTION

Forecasters frequently have contextual information about anticipated events such as policy changes that suggest a future shift in the level of a time series. When statistical extrapolation models are used, such planned discontinuities often confound the forecasting process. Since they are based on historical data patterns, their forecasts cannot incorporate anticipated changes for which there is no historical precedent.

In our experience, a two-stage procedure is commonly used. A statistical time series model is used to project the current pattern. Then the forecaster judgmentally adjusts the forecast to account for the planned shift in level. If the contextual information is good, this approach can be expected to produce reasonable initial forecasts; but it creates difficulties in tracking and updating. During and immediately following the period of expected change, the forecaster may wish to confirm that the change has occurred as planned and to update the forecast. But he or she is likely to experience difficulties in updating due to lack of fit in the time series model caused by the discontinuity in level within the now-historical data. Once the shift has occurred, it is not clear how to integrate the judgmental component with the forecast of a time series model: The change is now in the data, but the model is in the process of adapting itself to a new level.

This paper proposes a new approach to forecasting planned shifts in level that mitigates this problem by providing a more effective way to integrate contextual knowledge with the statistical forecast. Twenty real data series representing Virginia Medicaid cost components are used to examine the forecasting accuracy of the proposed method and several alternative approaches when applied to real data. In practice, actual shifts of level may differ from planned shifts in a variety of ways, such as magnitude, length of the change period, and timing. These 20 series are used in a simulation experiment to compare these methods under various forms of planning error.

A PROPOSED TECHNIQUE

A good model should require minimum intervention when a forecast remains accurate and be easy to adjust when adjustments are needed. These criteria are better achieved by making the adjustment for a planned intervention *within* the time series model. In this paper, we describe such an adjustment within the T. M. Williams variant [1] of Holt exponential smoothing (called "Holt-independent" in this paper), which is used because of its simplicity when expressed in its error-correction form.) Level-adjusted exponential smoothing (LAES) is defined as follows:

$F_{t,t+m}$ = Forecast at t for t + m
$= S_t + mB_t + (P_{t+m} - P_t)$ (1)
F_{t+1} = One-step-ahead forecast at t for t+1
$= S_t + B_t + A_{t+1}$ (2)
where
X_t = An observation from the series X at t (3)
e_t = Error at time t $= X_t - F_t$ (4)
S_t = Level at time t $= F_t + \quad e_t$ (5)
B_t = Trend at time t $= B_{t-1} + \quad e_t$ (6)
A_t = Adjustment factor at time t $= P_t - P_{t-1}$ (7)
P_t = The cumulative effect of policy through t (8)
m = Any arbitrary number of periods after period t

This model differs from the Holt-independent model only in the inclusion of the term $P_{t+m} - P_t$ within the m-step-ahead forecast function (or the term A_{t+1} within the one-step-ahead forecast function F_{t+1}).

Suppose that a policy change is expected to raise the level of a series over 4 periods, starting at 20 units and progressing cumulatively to 50, 65, and 70 units, after which the policy is fully in effect. In this case, the P series is (..., 0, 0, 0, 20, 50, 65, 70, 70, 70, ...) and the A series is (..., 0, 0, 0, 20, 30, 15, 5, 0, 0, 0, ...).

When a shift in level occurs as expected, large errors do not develop for LAES, so neither change in the model nor adjustment to the forecast is suggested. In contrast, an exponential smoothing model produces a large error in this case, suggesting a need for model intervention. When the anticipated effect of an intervention does not occur, the magnitude of errors within the LAES model increases, generating a signal that alerts the forecaster to an implementation failure or delay.

A SIMULATION STUDY

LAES is compared with the Holt-independent model, where in both cases the parameter values were chosen to optimize fit. For the Holt-independent model, the expected intervention effect is incorporated outside the model. LAES is also compared to two logical alternatives that might be used to facilitate "catching up" as rapidly as possible when a shift in level occurs: the Holt-independent model using a large parameter value ($\alpha = .9$) for smoothing the level component, and T. M. Williams' adaptive-parameter variation [1] of the Holt-independent model in which only the level parameter is adaptive.

The simulation uses 20 non-seasonal monthly Medicaid-related series for which forecasts are regularly prepared and updated by the Virginia Department of Medical Assistance Services. These 20 series have been found to be approximately independent based on a correlation matrix of their first differences.

Each series contains 62 observations. The first 24 observations are used for model initialization. The first through 39th observations are used for model fitting; however, errors for the initialization period are not included in the loss function. The 40th through 44th observations are adjusted to simulate the various level-shift scenarios to be examined. The remaining observations ($t = 45$ through 62) are also adjusted to simulated levels following a level shift and are then used in ex-ante evaluation. The models are fit using minimum root mean squared error as the criterion.

The simulation uses a factorial design involving four factors that commonly contribute to forecasting error: the expected shift pattern (number of periods over which the shift is expected to occur); the actual shift pattern; the magnitude of the shift; and the timing of the actual shift. For all scenarios, the expected magnitude of a shift is +25% of the level at period 39.

Scenarios are generated by using the following factor levels:

1. The planned pattern of the discontinuity is set alternately at a 1-period step and a 3-period ramp.
2. Simulated outcomes are as follows:
   a. Five level shifts: 0%, 20%, 25%, 30% and 50%.
   b. Two shift patterns: a 1-period step and a 3-period ramp. This creates four possible scenarios with respect to the expected and actual patterns: step-step; step-ramp; ramp-step; ramp-ramp.
   c. Two timing conditions: on time and 2-period delay.

There are 34 scenarios, with 2 expectation scenarios matched with 17 unique outcome scenarios. Summarizing:

1. Planned shifts of 25% of the original series in a 1-period step or a 3-period ramp.
2. Simulated shifts of 50%, 30%, 25%, 20%, or 0% of the original series,
3. Simulated shifts in a 1-period step or 3-period ramp.
4. Simulated shifts occurring when planned or after a two-period delay.

FORECASTS AND ERRORS

To evaluate forecast accuracy for the other techniques, within-model forecasts are adjusted outside the model by adding the change in level that is expected but has not yet been attained. Thus, with respect to forecasting accuracy, all models are compared on the basis that the forecaster used his knowledge of expected future disruptions in the series due to intervention.

Beginning with the last period prior to the simulated updating ($t = 39$) and for five successive updating periods, forecasts are projected through an 18 month horizon. Forecasting errors are summarized by the symmetrical absolute percent errors (SAPE), defined as SAPE = 100 |actual - forecast|/((actual + forecast)/2). SAPE is the individual error used in calculating symmetrical mean absolute percent error (S-MAPE).

To evaluate forecasting accuracy, SAPEs for the initial forecast and all updates (including extra-model adjustments for the exponential smoothing models) are calculated for horizons 1, 6, 12, and 18, and these errors are retained for evaluation. SAPE values for horizons 6, 12, and 18 use cumulative values for forecasts and simulated actuals.

SIMULATION RESULTS

We begin with a comparison of the forecasting accuracy of LAES and its direct counterpart, the augmented Holt-independent model (i.e., with the expected level shift added to the forecast outside the model). The initial forecasts are always the same over all horizons. For subsequent updates, however, their relative accuracy depends on the scenario. For scenarios in which the magnitude of change is approximately as expected (20%, 25%, or 30%), LAES was more accurate. For scenarios in which the expected change fails to materialize, magnitudes LAES was less accurate. When the level shift is twice the expected size (i.e. 50%), LAES is even more accurate in comparison to augmented Holt-independent. The explanation is that the greater the actual shift, the greater the internal disruption of the Holt-independent model's level and trend components upon updating. The disruption is substantially less for LAES because it has anticipated half of the shift internally. This general pattern persists through five successive updates, although the differences between the two methods diminish somewhat for succeeding updates.

Forecasting Horizon Effects

The differentiation between the models' S-MAPEs increases with increasing forecasting horizons. That is, when the shift size is approximately as expected or greater than expected, LAES's advantage in accuracy is most pronounced for longer horizons. When the expected shift fails to occur at all, LAES's deficit in accuracy is also most pronounced for longer horizons.

Pattern and Timing Effects

For the first update, the relative accuracy of the models depends substantially on the shift pattern and timing. When the expected and actual patterns are step/step, the results are as discussed previously. For a step/ramp pattern and a step/step pattern with a 2-period delay, the results are similar but less favorable to LAES. When the expected pattern is a ramp, whether the actual pattern is a step or ramp, there are only minor differences in accuracy between the two models regardless of shift size. Upon further updates, however, these pattern and timing effects disappear. By the third update, the entire shift has occurred and is in the data. Now the results are essentially the same for all 5 cases. Unless the shift fails to occur, LAES is more accurate than augmented Holt-independent, and its relative advantage increases with the size of the shift.

Summarizing, LAES-updated forecasts are clearly more accurate than those of the augmented Holt-independent model as long as the expected level shift occurs. The extent of its advantage in accuracy depends on the size of the shift, the expected and actual patterns of shift, and the timeliness of the shift. When the expected level shift fails to occur or is delayed, LAES forecasts are less accurate. However, these relatively large LAES errors are desirable because they signal failed expectations about policy implementation. The Holt-independent model does not provide such a signal.

Alternative Approaches

Now we compare LAES to Holt-independent with a large value of $\alpha = .9$ and T. M. Williams' adaptive exponential smoothing [1], both augmented outside the model with the expected shift in level. For the first update, the $\alpha = .9$ model was more accurate for 14 of 20 conditions studied in which a shift actually occurred. By the third update, the $\alpha = .9$ model was always more accurate than the adaptive-α model unless the expected shift failed to occur. Henceforth, the high-α variation is considered the primary competitor of LAES.

LAES was compared to Holt-independent, $\alpha = .9$ for the first three updates following the shift in level. The effect of expected and actual shift pattern on the comparative accuracy of the models is similar to the effect previously observed for augmented Holt-independent with optimal-fit α. After the full shift has occurred (i.e., step/step and ramp/step patterns for update 1 and all patterns for update 3), LAES is more accurate than augmented Holt-independent with $\alpha = .9$ unless the shift size is zero or the simulated shift is much larger than the expected shift. By the third update, the entire shift has occurred and is in the data for all shift patterns. Now the results are essentially the same for all pattern and timing combinations. If the shift is approximately the expected size, LAES is somewhat more accurate than augmented Holt-independent with $\alpha = .9$. If the shift fails to occur, then LAES is less accurate, although this means that it generates a much stronger signal of a failed implementation. If the shift is twice the expected size, the models are equally accurate or LAES is slightly less accurate depending on the scenario. As the length of the horizon increases, the differences between the two models are accentuated. Overall, LAES's advantage over the augmented Holt-independent $\alpha = .9$ model is less than for the augmented Holt-independent optimal fit model. The practitioner who is concerned with longer horizons obtains the greatest gain through the use of LAES.

CONCLUSIONS

In comparison to several exponential smoothing-based alternatives in which adjustment for planned level shifts are made outside the model, it was found that LAES usually provided more accurate forecast updates than the alternative models after adjusting them (outside the model) for the expected level shift. Generally, LAES forecasts were more accurate than the alternatives except when the shift failed to occur or was delayed. When adjustments are required, they are simpler for LAES than for the alternatives. The forecaster needs only to adjust the timing, size, or pattern of the shift component of the model, but not the level and trend components.

REFERENCES

[1] Williams, T. M. (1987), Adaptive Holt-Winters forecasting, *Journal of the Operational Research Society*, 38, 553-560.