NCEA Level 3 · AS 91580: Time Series Investigations

Name: _____ Date: _____

Hi! This standard is all about spotting patterns in data collected over time, trend, seasonal ups and downs, and using them to forecast. Work through the sheet in order, and always link your maths back to the real-world context of the data.

Revision summary

- Time series data = measurements recorded at regular time intervals (e.g. quarterly sales). Look for: trend (long-term direction), seasonality (regular repeating pattern), and irregular/random variation.
- Additive model: Data = Trend + Seasonal + Irregular ($y = T + S + I$). Use when seasonal swings stay roughly constant in size over time.
- Centred Moving Average (CMA) smooths out seasonality to reveal the trend. For quarterly data (period 4) use a 4-point moving average then centre it (average two consecutive 4-point means).
- Seasonal effect for a season = average of (actual value − CMA) for that season. Seasonal effects should sum to ≈ 0 across the cycle; if not, adjust each by subtracting the mean.
- Forecasting: extend the trend line (from the CMA or a fitted trend), then add the relevant seasonal effect for that season: Forecast = Trend estimate + Seasonal effect.
- Interpret in context: describe trend (increasing/decreasing, by how much), seasonal pattern (which seasons are high/low and why), and comment on the reliability of forecasts.
- Statistical enquiry cycle (PPDAC): Problem, Plan, Data, Analysis, Conclusion, state a purpose, explore the data, and write conclusions in context.

Key skills

1. A café in Devonport records quarterly hot-chocolate sales (in hundreds of cups): Q1 = 8, Q2 = 5, Q3 = 4, Q4 = 9. Calculate the 4-point moving average for these four quarters.

2. State what the 'seasonal effect' of a particular quarter tells you in an additive time series model.

3. The additive model gives a trend estimate of 42 units for the next winter quarter, and winter has a seasonal effect of −6.5. Forecast the value for that quarter.

4. Explain why we 'centre' a moving average when working with quarterly data.

Achieved

5. A garden centre in Henderson records quarterly plant sales ($000): 2022 Q1 22, Q2 30, Q3 18, Q4 14; 2023 Q1 26, Q2 34, Q3 22, Q4 18. Describe the trend and the seasonal pattern you can see.

6. Using the garden-centre data above, calculate the first two 4-point moving averages and the first centred moving average.

7. A power company's data shows the CMA (trend) for electricity use rising by about 5 units per quarter, reaching 210 units at Q4 2023. The Q1 seasonal effect is +18. Forecast Q1 2024.

8. The seasonal effects for a quarterly series are found to be Q1 = +6, Q2 = +11, Q3 = −4, Q4 = −11. Check whether they need adjusting, and adjust if required.

Merit

9. A Ponsonby ice-cream shop's quarterly revenue shows a clear seasonal pattern with Q2 (summer months Jan–Mar) highest. The CMA trend is increasing steadily. Explain how you would produce a forecast for next summer (Q2), and why the additive model is suitable here.

10. A tourism operator sees quarterly visitor numbers with a strong seasonal pattern, but the summer peaks are getting larger and larger each year as the business grows. Explain why forecasts from a simple additive model with constant seasonal effects could become unreliable, and what this means for the operator.

11. For a quarterly dataset, a student calculated seasonal effects but forgot to centre their 4-point moving average. Explain what problem this causes and how it would show up when they compare their trend to the actual data.

Excellence

12. Two analysts model the same quarterly rainfall series for Whangārei. One uses an additive model, the other a multiplicative model (Data = Trend × Seasonal × Irregular). The trend is roughly flat over the period. Argue what you would expect about the two sets of forecasts, and under what condition they would give nearly identical results.

13. A student claims: 'If I add up all the CMA (centred moving average) values and compare with the sum of the original data over the same complete years, they should be about equal.' Evaluate whether this claim is justified.

14. A local council forecasts quarterly water demand using an additive model built on 3 years of data. The forecasts extend 8 quarters (2 years) ahead. Discuss, with justification, the factors that limit how much confidence should be placed in the 8th forecast compared with the 1st.

Answer key (for parents and tutors)

1. A café in Devonport records quarterly hot-chocolate sales (in hundreds of cups): Q1 = 8, Q2 = 5, Q3 = 4, Q4 = 9. Calculate the 4-point moving average for these four quarters.
   4-point moving average = (8 + 5 + 4 + 9) ÷ 4 = 26 ÷ 4 = 6.5 (hundred cups).
   Achieved: applies the moving-average calculation.

2. State what the 'seasonal effect' of a particular quarter tells you in an additive time series model.
   The seasonal effect is the average amount by which that quarter's actual value sits above (+) or below (−) the underlying trend. For example, a seasonal effect of +2.3 means that quarter is typically 2.3 units higher than the trend line for that period.
   Achieved: defines seasonal effect within the additive model.

3. The additive model gives a trend estimate of 42 units for the next winter quarter, and winter has a seasonal effect of −6.5. Forecast the value for that quarter.
   Forecast = Trend + Seasonal effect = 42 + (−6.5) = 35.5 units.
   Achieved: applies Forecast = Trend + Seasonal effect.

4. Explain why we 'centre' a moving average when working with quarterly data.
   A 4-point moving average of quarterly data falls between two quarters (it has no single time point). Centring, averaging two consecutive 4-point means, realigns the smoothed value with an actual quarter, so the trend estimate lines up with the original data points.
   Key skill: understands purpose of the centred moving average.

5. A garden centre in Henderson records quarterly plant sales ($000): 2022 Q1 22, Q2 30, Q3 18, Q4 14; 2023 Q1 26, Q2 34, Q3 22, Q4 18. Describe the trend and the seasonal pattern you can see.
   Trend: sales are generally increasing over time, each 2023 quarter is about $4000 higher than the matching 2022 quarter. Seasonal pattern: Q2 is consistently the highest (30, 34) and Q4 the lowest (14, 18), with Q1 and Q3 in between. This makes sense as spring/summer (Q2) drives higher plant sales.
   Achieved: identifies trend and seasonal features from raw data.

6. Using the garden-centre data above, calculate the first two 4-point moving averages and the first centred moving average.
   First 4-point MA (Q1–Q4 2022) = (22 + 30 + 18 + 14) ÷ 4 = 84 ÷ 4 = 21. Second 4-point MA (Q2 2022 – Q1 2023) = (30 + 18 + 14 + 26) ÷ 4 = 88 ÷ 4 = 22. First CMA (aligned with Q3 2022) = (21 + 22) ÷ 2 = 21.5.
   Achieved: computes 4-point and centred moving averages.

7. A power company's data shows the CMA (trend) for electricity use rising by about 5 units per quarter, reaching 210 units at Q4 2023. The Q1 seasonal effect is +18. Forecast Q1 2024.
   Q1 2024 trend estimate = 210 + 5 = 215 units. Forecast = Trend + Seasonal effect = 215 + 18 = 233 units.
   Achieved: extends trend one step and applies seasonal effect.

8. The seasonal effects for a quarterly series are found to be Q1 = +6, Q2 = +11, Q3 = −4, Q4 = −11. Check whether they need adjusting, and adjust if required.
   Sum = 6 + 11 − 4 − 11 = +2. For an additive model the seasonal effects should sum to 0, so they need adjusting. Mean adjustment = 2 ÷ 4 = 0.5. Subtract 0.5 from each: Q1 = +5.5, Q2 = +10.5, Q3 = −4.5, Q4 = −11.5. These now sum to 0.
   Achieved: applies the seasonal-effect adjustment method.

9. A Ponsonby ice-cream shop's quarterly revenue shows a clear seasonal pattern with Q2 (summer months Jan–Mar) highest. The CMA trend is increasing steadily. Explain how you would produce a forecast for next summer (Q2), and why the additive model is suitable here.
   Method: (1) Fit/extend the CMA trend line to next Q2's time position to get the trend estimate. (2) Take the Q2 seasonal effect (the average of actual − CMA for past Q2s, after adjusting to sum to 0). (3) Forecast = trend estimate + Q2 seasonal effect. The additive model is suitable because the seasonal swings appear to stay roughly constant in absolute size over the years (the summer peak is about the same number of dollars above trend each

year), rather than growing in proportion to the level, which would need a multiplicative model.
Merit: links forecasting steps and justifies model choice from data behaviour.

10. A tourism operator sees quarterly visitor numbers with a strong seasonal pattern, but the summer peaks are getting larger and larger each year as the business grows. Explain why forecasts from a simple additive model with constant seasonal effects could become unreliable, and what this means for the operator.
An additive model assumes each season sits a fixed amount above/below the trend. If summer peaks are growing (the gap above trend widens each year), a constant seasonal effect will underestimate future summers and the error will grow over time. The seasonal variation is proportional to the level, so a multiplicative model would fit better. For the operator, additive forecasts would understate demand in peak seasons, risking under-staffing or under-stocking, so they should treat far-ahead additive forecasts with caution.
Merit: connects model assumptions to reliability and real consequences.

11. For a quarterly dataset, a student calculated seasonal effects but forgot to centre their 4-point moving average. Explain what problem this causes and how it would show up when they compare their trend to the actual data.
An uncentred 4-point moving average lands between two quarters, so it is time-shifted by half a quarter relative to the actual data points. When the student computes actual − trend to find seasonal effects, they'd be subtracting a trend value that doesn't align with the correct time point. This distorts the seasonal effects (systematically over/under-estimating them) and, when plotted, the smoothed trend line would appear offset from where it should sit against the raw data. Centring fixes the alignment so residuals reflect true seasonal deviations.
Merit: links a procedural step to its effect on results across representations.

12. Two analysts model the same quarterly rainfall series for Whangārei. One uses an additive model, the other a multiplicative model (Data = Trend × Seasonal × Irregular). The trend is roughly flat over the period. Argue what you would expect about the two sets of forecasts, and under what condition they would give nearly identical results.
When the trend is flat (roughly constant level), the additive and multiplicative models will give very similar forecasts. In an additive model, seasonal effect = actual − trend; in a multiplicative model, seasonal index = actual ÷ trend. If the trend level L is constant, then a multiplicative index of, say, 1.2 corresponds to an additive effect of 0.2 × L, a fixed amount, so both describe the same fixed seasonal deviation. The models diverge only when the trend changes over time: then the multiplicative model scales the seasonal swing with the level while the additive keeps it constant. So the near-flat trend here means the two forecasts should be almost identical; the choice matters far more for strongly trending series.
Excellence: generalises the relationship between additive and multiplicative models and justifies the condition for equivalence.

13. A student claims: 'If I add up all the CMA (centred moving average) values and compare with the sum of the original data over the same complete years, they should be about equal.' Evaluate whether this claim is justified.
The claim is broadly justified for complete cycles. A centred moving average is a weighted average of surrounding data that smooths out seasonal highs and lows; over one full cycle the positive seasonal deviations cancel the negative ones (they sum to ≈ 0 in an additive model). So summing the CMA over a whole number of years gives approximately the same total as summing the raw data, because the trend represents the average level and the seasonal parts net out. It is 'about' rather than exactly equal because the CMA cannot be computed for the first and last few data points (end effects), so those points are missing from the CMA sum. If we restrict both sums to the same set of time points where the CMA exists, they will be very close, differing only by the residual irregular variation which also tends to cancel.
Excellence: reasons abstractly about the averaging property and identifies the end-effect limitation.

14. A local council forecasts quarterly water demand using an additive model built on 3 years of data. The forecasts extend 8 quarters (2 years) ahead. Discuss, with justification, the factors that limit how much confidence should be placed in the 8th forecast compared with the 1st.
Several factors reduce confidence as forecasts extend further: (1) Trend extrapolation risk, the trend line is fitted to only 3 years; assuming it continues unchanged 2 years ahead ignores possible turning points (e.g. new water restrictions, population change). Small errors in the trend slope compound over 8 quarters, so the 8th forecast could be far off while the 1st is close. (2) Constant-seasonal assumption, the additive model uses fixed seasonal effects from limited data; if climate or usage patterns shift, later forecasts drift further from reality. (3) Irregular/unpredictable events (droughts, policy changes) accumulate uncertainty over time. (4) Only 3 cycles

gives seasonal effects based on just 3 values each, a small, potentially unrepresentative sample. Overall, the 1st forecast relies mostly on the recent, well-supported trend and is more reliable; the 8th depends on the trend and seasonal pattern holding for two more years, which is far less certain, so its confidence interval would be much wider. The council should update the model as new data arrives rather than commit to the far forecasts.
Excellence: justifies how forecast uncertainty grows and evaluates model limitations in context.