Adaptive Demand Level

Measuring whether the current year is running above or below its expected seasonal baseline

STUDY QUESTION

How can a demand forecast preserve the annual seasonal shape while adapting when the current year is systematically stronger or weaker than historical norms?

Executive Summary

A fixed seasonal curve answers when demand normally rises and falls, but not whether the current year is running at the same absolute level as prior years. To solve that problem, completed weekly demand was normalized by its expected seasonal level.

The resulting ratio converts raw customer counts into a relative demand measure: a value above the seasonal baseline means the current year is running stronger than expected for that point in the cycle; a lower value means it is running weaker. This adaptive level became the bridge between the recurring seasonal curve and the current operating year.

Why Fixed Seasonality Is Not Enough

A seasonal forecast can be directionally correct and still be systematically too high or too low if the business, market, or acquisition environment changes between years. A flat recent average has the opposite weakness: it adapts to the current level but does not know the expected seasonal slope.

The model therefore separates shape from level instead of forcing one mechanism to solve both problems.

Normalization Step

$$r_j = (NewCustomers_j + 1) / S(w_j)$$

r_j measures realized demand relative to the seasonal demand level expected for week j.

Interpretation

The normalization changes the analytical question from “How many customers did we acquire?” to “How is the current year performing relative to what is normally expected at this point in the season?”

That distinction is operationally useful because the same raw count can have very different meaning in a naturally weak week versus a peak-demand week.

How the Adaptive Level Fits the Forecast

Once recent weekly observations are expressed relative to seasonality, the model can estimate the current local level without being distorted by the normal spring ramp-up or late-season decline. The adaptive estimate is then multiplied back through the seasonal curve to return to the original New Customers scale.

$$Predicted\ New\ Customers_t = max(0, S(w_t) * max(0.01, L_t) - 1)$$

Seasonality supplies the expected shape; L_t supplies the current relative level and local trend.

Evidence for an Adaptive Architecture

The source material does not report a holdout result for a level-only model in isolation, so this component should not be assigned a standalone lift. The available evidence supports the broader architectural decision: fixed seasonal benchmarks were materially weaker than the complete adaptive seasonal forecast on the 2025 holdout.

2025 holdout model	MAE	RMSE	Bias	Correlation
Adaptive seasonal	1.97	3.23	+0.13	0.891
Seasonal month mean	3.52	4.90	+2.32	0.788

Business Applications

- Benchmark actual acquisition against the demand level expected for the current point in the season.
- Distinguish a weak business result from a naturally weak market-demand week.
- Create a dynamic baseline for marketing pacing, staffing, and performance reviews.
- Reduce dependence on prior-year absolute levels when the current year is structurally different.

Key Takeaways

Problem solved	Fixed seasonal curves can preserve shape while missing the current year's level.
Core method	Normalize realized customers by the seasonal factor before estimating current demand strength.
Interpretation	The normalized ratio tells whether the year is running above or below seasonal expectation.
Role in system	The adaptive level allows the seasonal curve to move up or down as new weeks are completed.

Limitations

The adaptive level is descriptive of realized acquisition and should not automatically be interpreted as a pure market-demand measure. Changes in marketing execution, conversion, capacity, or tracking can also affect New Customers. The broader forecasting project treats New Customers as the fixed target, so these operating influences remain embedded in the observed series.

Portfolio-Ready Summary

SUGGESTED HEADLINE

Separating seasonal shape from the current year's demand level

I converted weekly customer counts into seasonally normalized demand ratios so the model could distinguish "where we are in the year" from "how strong this year is running." That adaptive level lets the forecast preserve the annual demand curve while moving up or down as new information arrives.