Adaptive Seasonal Demand Forecasting Model

Portfolio methodology: how the model was designed, tested, rejected, rebuilt, and validated

Objective: estimate expected weekly New Customers while separating recurring seasonality, the current demand level, and the changing rate of demand growth or decline.

Project component	Portfolio summary
Primary target	Weekly New Customers
Development period	2023–2024
Strict holdout	2025
Later out-of-time check	2026 YTD, with calendar-alignment sensitivity
Final model type	Adaptive seasonal forecast with Fourier seasonality + locally trended normalized demand level
Final 2025 status	Provisionally validated

This document is intended to accompany a portfolio case study. It focuses on analytical decisions rather than only the final formula: what was tested, why certain approaches were rejected, how leakage was controlled, and what evidence was required before calling the model useful.

1. Problem framing

The project began with a narrow forecasting question: can observable historical conditions support a defensible weekly Demand Score for expected customer demand? The dependent variable was fixed from the start as New Customers. Leads, revenue, conversions, and normalized customer value were not allowed to replace the target.

What the model needed to accomplish

- Preserve the strong annual demand cycle without assuming the same absolute level every year.
- Adapt to the current year's demand level as new weeks are completed.
- React differently when demand is accelerating, peaking, or declining.
- Avoid giving weather or economic variables credit for simply sharing the same seasonal pattern.
- Operate prospectively: a forecast for week t can only use information available before week t.
- Remain simple enough to explain, reproduce, audit, and deploy.

Core modeling principle: a complex model only earns a place if it improves genuine forward forecasts against simple adaptive alternatives.

2. Data and validation design

Item	Design
Historical modeling data	156 weekly observations across 2023–2025
Development sample	104 weeks from 2023–2024
Final holdout	52 weeks from 2025
Target	New Customers
Chronology	Never shuffled
Internal validation	Rolling-origin / expanding-window forecasting
Scoring scale	Defined only from development-period predictions

The 2025 target was intentionally withheld during model development. Variable selection, lag selection, seasonal specification, transformations, and score construction were frozen before the 2025 outcome was evaluated. This avoided repeatedly tuning the model until the holdout looked favorable.

Target behavior

Development statistic	Result
Mean New Customers	8.10
Variance	68.69
Variance / mean	8.48
Zero-demand weeks	23
Zero rate	22.1%
Lag-1 autocorrelation	0.693

The target is a nonnegative, overdispersed count series with strong persistence and seasonality. This motivated testing count-regression families as well as ordinary regression and adaptive time-series baselines rather than treating the series as independent continuous observations.

3. Analytical workflow

Stage	What was done	Why it mattered
1. Audit	Checked grain, missingness, duplicates, target distribution, leakage risks, low-variation fields.	Prevents modeling artifacts.
2. Baselines	Historical mean, seasonal mean, prior-year period, rolling means, recent-level seasonal methods.	Sets a minimum performance bar.
3. Seasonality	Month controls and Fourier harmonics.	Separates annual structure from external signals.
4. Feature screen	Current values, lags, rolling windows, effect sizes, correlations, incremental forecast tests.	Tests each signal beyond seasonality.
5. Model families	OLS/GLM, count models, regularization, robust regression, shallow nonlinear candidates.	Avoids assuming one structure is correct.
6. Rolling backtest	Chronological one-step-ahead development forecasts.	Measures forward accuracy, not training fit.
7. Ablation/diagnostics	Removed features, checked residuals, multicollinearity, stability, influential periods.	Confirms complexity is earning its place.
8. Freeze	Locked model and score scale using 2023–2024 only.	Protects holdout integrity.
9. Holdout	Applied once to 2025.	Tests generalization.
10. 2026 check	Out-of-time backtest with calendar-alignment sensitivity.	Tests behavior under a new operating year.

4. Why the first model was rejected

The first frozen specification combined annual seasonality, prior-week demand, and precipitation hours lagged three weeks. It looked statistically reasonable and ranked 2025 demand well, but it failed the forecasting standard.

2025 holdout	First frozen model	Recent 4-week mean
MAE	2.90	2.43
RMSE	3.89	3.70
Actual vs predicted correlation	0.855	0.846

The issue was structural. A recent rolling average adapts to the current level, but it does not know the expected seasonal slope. Conversely, a fixed seasonal model knows the annual shape but may remain anchored to an outdated level. The model needed both.

Rejected conclusion: high correlation was not enough. Because a simpler recent-level baseline had lower error, the first Demand Score was classified as not validated.

5. Revised architecture: seasonality + adaptive level + local trend

The second architecture reframed the problem. Instead of forecasting raw demand directly from a flat rolling average or a fixed seasonal curve, each completed week's demand is first normalized by its expected seasonal level. A short local trend is then estimated on those normalized ratios.

Step 1 — Estimate the recurring seasonal shape

$$\eta(w) = 1.66650229 + 0.32560289 \sin(2\pi w/52) - 1.35700057 \cos(2\pi w/52) + 0.07409782 \sin(4\pi w/52) - 0.32840420 \cos(4\pi w/52)$$

$$S(w) = \exp(\eta(w))$$

Two Fourier harmonics were retained because they provided a parsimonious representation of the annual demand curve without creating 12 month coefficients or a large set of week-of-year dummy variables.

Step 2 — Normalize recent actual demand by seasonality

$$r_j = (\text{NewCustomers}_j + 1) / S(w_j)$$

The ratio r_j answers a more useful operational question than raw demand: is the current year running above or below what would normally be expected at this point in the seasonal cycle?

Step 3 — Estimate the current six-week local level and slope

$$L_t = (-4r[t-6] - r[t-5] + 2r[t-4] + 5r[t-3] + 8r[t-2] + 11r[t-1]) / 21$$

The weights are the endpoint estimate from a local linear trend across the previous six normalized observations. More recent weeks receive progressively larger weight. This allows the model to respond differently during a spring ramp-up than during a late-season decline.

Step 4 — Reapply the seasonal shape to the current level

$$\text{Predicted New Customers}_t = \max(0, S(w_t) \times \max(0.01, L_t) - 1)$$

Interpretation: seasonality determines the expected direction and rate of movement; the adaptive component determines how high or low the current year is running relative to that seasonal curve.

6. Feature testing and why external variables were excluded

Weather and economic variables were tested rather than assumed. The screen included raw correlations, lagged versions, rolling windows, seasonality-adjusted relationships, and incremental rolling-origin performance.

- Temperature and dew point showed strong raw correlation with demand, but most of that relationship disappeared after removing the annual seasonal cycle.
- Precipitation variables showed some lagged signal in earlier specifications, but the gain weakened once the adaptive seasonal level was modeled correctly.

- Economic variables were constrained by low effective sample size at weekly frequency, repeated monthly values, and uncertainty about publication timing.
- No external variable improved the revised adaptive baseline enough during development to justify the extra complexity in the final frozen model.

This was a deliberate negative result. The final model is smaller because the evidence suggested that much of the apparent external-variable signal was actually shared seasonality or short-term demand momentum.

7. Model families and statistical tools tested

Category	Examples tested	Role
Baseline	Historical mean, seasonal mean, same-period prior year, rolling means, adaptive seasonal baselines	Benchmark forecasting value
Regression	OLS and seasonal regressions	Interpretable expected-demand models
Count models	Poisson and Negative Binomial	Target distribution check
Regularization	Ridge, Lasso, Elastic Net	Control redundancy / overfit
Robust regression	Huber-style regression	Check sensitivity to spikes
Nonlinear	Shallow boosting / constrained low-complexity alternatives	Check whether modest nonlinearity added value

Diagnostics used

- MAE, RMSE, mean error/bias, actual-predicted correlation, predictive R^2.
- AIC/BIC where appropriate for likelihood-based models.
- Correlation matrices and VIF for multicollinearity.
- HC3 robust standard errors for coefficient inference.
- Ljung–Box residual autocorrelation checks.
- Breusch–Pagan heteroskedasticity check.
- CUSUM stability diagnostics.
- Influential-observation review and ablation testing.

8. Development backtest results

2024 rolling-origin development	MAE	RMSE	Bias	Correlation
Adaptive seasonal K2 + 6-week trend	5.22	7.65	-0.51	0.598
Seasonal month mean	5.44	8.48	-1.33	0.478
Recent 4-week mean	5.51	7.82	+0.29	0.559
Recent 6-week mean	6.38	8.69	+0.33	0.420
Same week prior year	6.98	10.01	-1.78	0.345

The revised architecture won development selection before 2025 was opened. That matters because the later holdout result was confirmatory rather than a basis for choosing the structure.

9. Strict 2025 holdout validation

2025 model	MAE	RMSE	Bias	Corr.	Predictive R^2
Adaptive seasonal	1.97	3.23	+0.13	0.891	0.778
Recent 4-week mean	2.43	3.70	0.00	0.846	0.709
Recent 6-week mean	2.89	4.32	0.00	0.786	0.603
Seasonal month mean	3.52	4.90	+2.32	0.788	0.490

Holdout improvement vs the recent 4-week baseline: 19.2% lower MAE and 12.7% lower RMSE.

The 2025 result supports provisional validation rather than full validation. Only one untouched annual cycle was available, so the model demonstrated useful generalization but still required a later out-of-time check.

10. Converting expected demand into a 0–100 Demand Score

The Demand Score is not a manually weighted index. It is a percentile transformation of the model's predicted New Customers. The reference distribution was defined only from usable 2023–2024 development predictions.

Score band	Approx. predicted New Customers	Interpretation
0–20	< 0.86	Unusually weak expected demand
20–40	0.86–3.32	Below-normal expected demand
40–60	3.32–11.70	Typical expected demand
60–80	11.70–15.67	Above-normal expected demand
80–100	> 15.67	Unusually strong expected demand

In 2025, the score's rank relationship with actual demand was strong (Spearman correlation ≈ 0.905), and observed demand rose monotonically across score quintiles. That supports using the score as an interpretable demand-intensity layer on top of the point forecast.

11. 2026 out-of-time check and alignment lesson

A second-year check used 2026 Customer Since counts. This surfaced an important data-engineering issue: the 2026 pivot was Monday–Sunday, while the frozen historical model used fixed Jan 1–7, Jan 8–14, etc. buckets. A direct Week 1-to-Week 1 match therefore shifts the target by three days.

To avoid presenting an invalid comparison as definitive, the portfolio analysis reports an overlap-weighted approximation and explicitly labels it as such. Exact 2026 validation requires daily Customer Since dates or daily counts.

2026 overlap-weighted approximation (W01–W36)	Adaptive seasonal	Recent 4-week
MAE	1.560	1.621
RMSE	2.304	2.324

2026 overlap-weighted approximation (W01–W36)	Adaptive seasonal	Recent 4-week
Bias	-0.032	-0.335
Predictive R^2	0.591	0.584

2026 phase	Adaptive MAE	Recent4 MAE	What it suggests
Rising season (W01–W20)	1.27	1.47	Seasonal slope adjustment adds value
Declining season (W21–W36)	1.92	1.80	Recent-level baseline adapts slightly faster after peak

Portfolio lesson: alignment sensitivity is itself a modeling finding. Calendar definitions are part of the model specification, not a formatting detail.

12. Tooling and implementation stack

Tool / method	How it was used
Python	Primary analysis and reproducible modeling environment.
pandas	Time-series cleaning, joins, lag/rolling features, calendar alignment, metric tables.
NumPy	Numerical transforms, Fourier features, vectorized calculations.
statsmodels	Regression/count-model estimation and statistical diagnostics.
scikit-learn	Regularized and robust candidate models; standardized model comparison.
Excel / XLSX	Input review, reproducible output tables, audit-ready model workbooks.
Rolling-origin validation	Chronological model selection without random cross-validation.
Fourier seasonality	Compact annual seasonal shape using sine/cosine terms.
Empirical percentile scoring	0–100 score derived from development predictions rather than subjective weights.

Reproducibility controls

- Chronological ordering preserved throughout.
- No random K-fold cross-validation.
- 2025 kept out of feature and model selection.
- Forecast features restricted to information available at or before the forecast origin.
- Negative and failed tests retained in the audit trail.
- Model frozen before holdout evaluation.
- Score scale defined from development predictions only.
- Calendar-bucket definition treated as part of the model contract.

13. What I would show on the portfolio page

Section	Recommended content
Hero	Adaptive Seasonal Demand Forecasting — forecasting expected weekly demand while separating seasonality, current level, and local momentum.
Problem	Fixed seasonal models miss changing annual levels; flat rolling averages lag during rapid seasonal acceleration or decline.
Method	Fourier seasonal curve → seasonally normalized recent demand → six-week local trend → expected New Customers → percentile Demand Score.
Evidence	2025 holdout: 19.2% lower MAE and 12.7% lower RMSE vs recent 4-week baseline.
Validation discipline	2023–2024 development, 2025 untouched holdout, 2026 out-of-time alignment sensitivity.
Technical credibility	Rolling-origin backtesting, ablation, robust/count models, multicollinearity checks, residual diagnostics.
Honest limitation	2026 exact validation requires daily re-bucketing because weekly calendar definitions changed.

Suggested portfolio headline

Built and validated an adaptive weekly demand model that combines recurring seasonality with the current demand level and local rate of change, improving 2025 holdout MAE by 19.2% versus a recent-level baseline.

Suggested one-sentence methodology description

I used chronological rolling-origin validation to compare seasonal, adaptive, count, regularized, robust, and nonlinear candidates; the final model uses Fourier seasonality and a locally trended six-week seasonally normalized demand level, with a percentile-based 0–100 score defined entirely from development data.

14. Limitations and next validation step

- Only two complete annual cycles were available for development, which limits the certainty of seasonal and external-signal estimates.
- One strict holdout year is strong evidence for a small dataset, but not enough to call the model permanently stable.
- Economic features repeated monthly observations across weekly rows, reducing their effective independent sample size.
- Current-week realized weather is explanatory but not automatically forecast-safe unless historical forecast vintages are available.
- The 2026 weekly target uses a different week definition; daily re-bucketing is required for a definitive 2026 comparison.
- The model performed relatively better during the rising seasonal ramp than after the peak, which is a candidate area for future refinement.

Next step: reconstruct 2026 New Customers at the original daily-to-week bucket definition, run the frozen model without retraining, and use that exact out-of-time result as the next validation checkpoint.

Appendix — model-development decision log

Decision	Evidence	Outcome
Use New Customers as target	Defined business objective and count-series behavior	Kept throughout
Random CV	Would leak time ordering	Rejected
Poisson as sole final model	Strong overdispersion	Rejected
Fixed seasonal model	Did not adapt sufficiently to changing annual level	Rejected
Flat recent rolling mean as final model	Adapts level but not seasonal rate of change	Retained only as benchmark
Weather/economic signals in final model	No material incremental gain after adaptive seasonality	Rejected
Adaptive seasonal + local trend	Won development and 2025 holdout	Selected
Full validation label	Only one strict annual holdout	Deferred; status = provisionally validated
Naive 2026 week-number match	3-day calendar shift	Rejected as formal evidence

All performance figures in this document come from the model-development and backtesting workflow described above.