e-ISSN 2549-7472
Edumatic: Jurnal Pendidikan Informatika
Volume 10, Issue 3, **520-529.**
https://e-journal.hamzanwadi.ac.id/index.php/edumatic
DOI: 10.29408/edumatic.v10i3.36932

Beyond Historical Prices: Weather and Eid Features in Weekly Curly Red Chili Price Forecasting

Andin Sabilla Janna [1], Dedy Kurniawan [1,*], Ken Ditha Tania [1], Mira Afrina [1], Apriansyah Putra [1]

[1] Universitas Sriwijaya, Indonesia
[*] Corresponding author: Dedy Kurniawan, Universitas Sriwijaya, Indonesia
✉ dedykurniawan@unsri.ac.id

Copyright: © 2026 by the authors

Received: 28 September 2026 | Revised: 16 October 2026 | Accepted: 24 October 2026 | Published: 25 November 2026

Abstract
Chili prices fluctuate sharply; prior forecasts used weather and holidays, but their incremental value across algorithms and periods remains uncertain. This study assessed how consistently they forecast Padang's median curly red chili prices for the next week. A Knowledge Discovery in Databases workflow integrated PIHPS Nasional prices, NASA POWER weather, and official Eid al-Fitr and Eid al-Adha dates for 2017–2026. We compared Ridge Regression, Random Forest, and Gradient Boosting across historical-temporal features (S1), added weather (S2), added Eid indicators (S3), and both (S4). Five-fold TimeSeriesSplit with a one-observation gap selected S1 Ridge. Evaluation used 76 final-test weeks, Naive Forecast, and 2022–2024 backtests. S1 reduced test RMSE by 2.94% against Naive Forecast. S2 lowered test RMSE only for Ridge; S3 reduced it by 0.73–1.01% across algorithms. All nine 95% block-bootstrap confidence intervals for RMSE reductions relative to S1 included zero. For Ridge, S3 reduced holiday-week RMSE by 29.78%, falling to 4.01% after excluding one of four Eid events. Pre-holiday RMSE increased by 5.95%. Historical backtests showed inconsistent gains. These findings support evaluating external features through incremental accuracy, uncertainty, and temporal consistency; positive average gains alone do not demonstrate reliable predictive value. Transferability beyond Padang remains untested.

Keywords: agricultural commodity price prediction; eid calendar indicators; exogenous features; feature evaluation; time-series forecasting

To cite this article: Janna, A. S., Kurniawan, D., Tania, K. D., Afrina, M., & Putra, A. (2026). Beyond Historical Prices: Weather and Eid Features in Weekly Curly Red Chili Price Forecasting. *Edumatic: Jurnal Pendidikan Informatika*, *10*(3), 520–529. https://doi.org/10.29408/edumatic.v10i3.36932

INTRODUCTION

Forecasting agricultural commodity prices can support market planning, decision-making, and price monitoring, although accuracy varies across commodities and datasets (Manogna et al., 2025; Sari et al., 2024). Red chili is a volatile strategic commodity in Indonesia (Matondang et al., 2024; Utari et al., 2025). In West Sumatra, curly red chili is a leading horticultural commodity facing production challenges, including yield losses from anthracnose (Warman et al., 2025). These findings underscore price volatility and local production vulnerability but do not establish forecasting performance in Padang. This study forecasts next week's median price to provide an advance market-monitoring reference and to assess pre-Eid, Eid, and post-Eid weeks separately. Daily prices preserve within-week fluctuations, whereas

Edumatic: Jurnal Pendidikan Informatika is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License

monthly aggregation may combine these phases. The weekly horizon thus fits the research objective without implying greater accuracy than other horizons.

Advances in artificial intelligence have expanded agricultural price forecasting through machine learning alongside statistical and deep learning approaches. LSTM performed well in short-term vegetable price forecasting (Zhang et al., 2024), while an optimized GA-ELM model outperformed GA-LSTM and Auto-ARIMA for the commodities examined (Sari et al., 2024). Across 23 commodities, forecasting errors varied with commodity and data characteristics (Manogna et al., 2025). Thus, model complexity alone does not guarantee accuracy; selection requires relevant data and appropriate evaluation. Predictor choice is equally critical: historical prices may capture persistent patterns (Devianto et al., 2024) earlier weather may indirectly indicate crop conditions (Siaga et al., 2024) and Eid dates provide known-ahead calendar information (Salsabila et al., 2025). However, Padang's weather may not represent producing-area conditions, while historical prices may already reflect external changes. The incremental predictive value of weather and Eid indicators beyond historical-temporal features must therefore be empirically tested, not assumed.

Feature selection is central to forecasting accuracy. Nayak et al. (2024) examined rainfall and temperature in weekly tomato, onion, and potato price forecasts, whereas Vairo et al. (2026) found no systematic benefit from adding weather to monthly wheat forecasts. Pardede et al. (2025) incrementally added price lags, moving averages, holiday indicators, and seasonal features to XGBoost using daily cayenne-pepper prices and weather data. This supported feature engineering in that setting but did not establish gains across algorithms or separate Eid phases. Salsabila et al. (2025) reported lower MAPE for ARIMAX with Eid al-Adha and Christmas indicators than for ARIMA (12.04% versus 17.50%) in daily East Java curly chili forecasts. However, the models had different ARIMA orders, the test covered only December 25–31, 2024, and each indicator combined pre-holiday, holiday, and post-holiday days. Thus, this comparison cannot isolate calendar-feature gains or establish their consistency across phases and individual Eid events. Mutiara et al. (2025) examined food prices and inflation before, during, and after Ramadan, whereas Murtani et al. (2026) forecast multiple commodities ahead of Eid al-Fitr using monthly data, seasonal indicators, and walk-forward validation. Neither assessed phase- and event-specific calendar gains in weekly curly red chili forecasts. The unresolved question is whether weather and Eid indicators add predictive value beyond a common historical-temporal baseline for one-week-ahead prices in Padang, and whether gains persist across algorithms, holiday phases, and individual events. Aggregate error reductions alone cannot establish such consistency.

To evaluate the incremental predictive value of external features, this study applies Knowledge Discovery in Databases (KDD), linking data selection, preprocessing, feature transformation, modeling, and evaluation to pattern interpretation (Bahri & Tania, 2025). The data-mining stage compares Ridge Regression, Random Forest, and Gradient Boosting in forecasting Padang's next-week median curly red chili price, with Naive Forecast as a persistence benchmark. All three algorithms use the same feature scenarios: historical-temporal features (S1), S1 plus weather (S2), S1 plus Eid al-Fitr and Eid al-Adha indicators (S3), and S1 plus both (S4).

TimeSeriesSplit preserved chronological order during validation, while RMSE, MAE, MAPE, and R^2 assessed predictive performance. Three research questions guided the study: (1) Do weather and Eid features, separately or jointly, improve forecasting accuracy relative to S1? (2) How does their incremental predictive value differ across Ridge Regression, Random Forest, and Gradient Boosting? (3) For the validation-selected algorithm, how does adding Eid indicators affect prediction errors across pre-holiday, holiday, and post-holiday weeks, and how sensitive is the aggregate holiday-week result to individual Eid events?

This study contributes empirical evidence on the incremental value of weather and Eid indicators beyond historical-temporal features across three algorithms and evaluation periods, including Ridge's errors by Eid phase and individual event. These comparisons reveal gains and setbacks obscured by aggregate accuracy, demonstrating why feature value must be assessed against existing predictors, model structure, and holiday context. The findings may inform feature selection for Padang's weekly price monitoring, but transfer beyond the city, 76 final-test weeks, and four observed Eid events remains untested.

METHOD

KDD organized selection (price, weather, and Eid sources), preprocessing (date alignment, weekly aggregation, and filtering), transformation (S1: historical-temporal; S2: S1 + weather; S3: S1 + Eid; S4: S1 + both), data mining (three algorithms per scenario), and evaluation (chronological model selection and within-algorithm S2–S4 RMSE comparisons against S1 across test periods) (Yandar et al., 2026) (Figure 1).

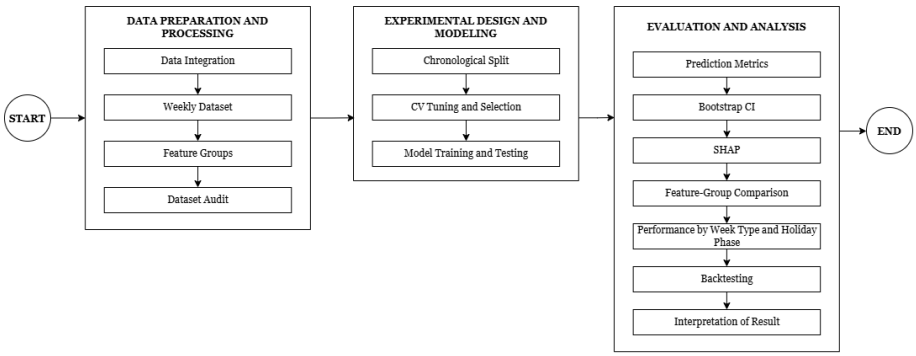

Figure 1. Research workflow

PIHPS recorded 2,477 weekdays (2 Jan 2017–30 Jun 2026; 181 missing prices). Removing 60 initially unpriced dates left 2,417 (121 missing) from 27 Mar 2017. Prices were aligned with 3,383 complete NASA POWER weather days, retaining weekends for Monday–Sunday aggregation. Weekly prices were the median of at least two observed weekdays; rainfall was summed, and temperature and humidity were averaged over seven days. From 484 candidate weeks, exclusions for incomplete weather (1), insufficient input and target prices (2 each), incomplete target week (1), and invalid four-week histories (9) left 469 next-week pairs without forward filling. Weekly aggregation may obscure daily spikes around Eid.

We represented input-week price and one- and two-week lags; a four-week rolling median and standard deviation summarized price level and dispersion. The median limits isolated spikes (Chaitra & Meena, 2026), while Li et al. (2026) used a four-week moving-average window. Input month represented seasonality; rainfall (one- and two-week totals), temperature, and humidity represented weather. Six binary target-week indicators encoded two pre-Eid weeks, Eid week, and two post-Eid weeks for Eid al-Fitr and Eid al-Adha. Window lengths were design choices, not established optima; one- and three-week Eid spans were tested for sensitivity (Table 1).

Table 1. Feature group scenarios

Scenario	Feature Group	Number of Features
S1	Historical Temporal	6
S2	Historical Temporal Features + Weather	10
S3	Historical Temporal Features + Holidays	12
S4	Historical Temporal Features + Weather + Holidays	16

Ridge Regression uses L2 regularization to address overlapping price-history predictors (Faisa & Salam, 2026); Random Forest averages trees to capture nonlinear thresholds and interactions; and Gradient Boosting models residuals sequentially (Zhu et al., 2025). These complementary linear, bagging, and boosting approaches allow S1–S4 comparisons of weather and Eid feature values under identical chronological evaluation. Chronological splitting assigned 392 pairs (2017–2024) to training and 76 (2025–2026) to testing, excluding one pair crossing the year boundary. All 12 combinations were tuned on training data using five-fold expanding-window TimeSeriesSplit with a one-observation gap, median imputation fitted within each training fold, Ridge-only scaling, and the grids in Table 2. Each model was refitted on the full training set; final selection used mean validation RMSE alone.

All 12 combinations were tested on 76 weeks using RMSE, MAE, MAPE, R^2, and Naive Forecast (current-week median); S2–S4 were compared with S1 within each algorithm. Original S1/S3 predictions supported holiday-phase, event-level, and leave-one-event-out RMSE analyses. Backtests (2022–2024) retuned and refitted all combinations on expanding training sets with the original grids, validation scheme, and boundary exclusions to prevent leakage. S3 was retuned for one-, two-, and three-week pre/post-Eid spans (Eid week fixed) and evaluated on the final test and Ridge backtests; S1 Ridge was revalidated without four-week price standard deviation. These analyses did not change final model selection. Paired circular block bootstrap used 20,000 paired resamples of 76 weeks (four-week blocks) to estimate percentile 95% CIs for RMSE reductions versus S1; two- and eight-week blocks assessed sensitivity without refitting (Table 5). SHAP evaluated fitted S1/S3 Ridge models on 76 test weeks using a training-data reference (Figure 2).

RESULTS AND DISCUSSION

Results

Data preparation yielded 469 eligible weekly pairs; excluding one 2024–2025 boundary pair left 392 training and 76 final-test pairs. Table 2 lists the grids used to tune all 12 algorithm–scenario combinations through five-fold TimeSeriesSplit on training data. Random Forest and Gradient Boosting used random_state = 42. We selected each combination's hyperparameters based on the lowest mean validation RMSE and, across combinations, the final model before testing. Table 3 reports these RMSEs and their across-fold standard deviations.

Table 2. Hyperparameter tuning configuration

Model	Parameters evaluated
Ridge Regression	α = 1, 10, 50
Random Forest	max_depth = 6, 10; min_samples_leaf = 2, 4; n_estimators = 120; max_features = "sqrt"
Gradient Boosting	learning_rate = 0.03, 0.05; max_depth = 2, 3; n_estimators = 150; min_samples_leaf = 3; subsample = 0.8

Across S1–S4, Ridge's validation RMSE was 22.7–30.8% below tree models (Table 3). Its S1 advantage was largest for 272 validation weeks with ≤20% absolute week-to-week price changes (Rp4,416/kg versus Rp7,909–8,271/kg); for 53 weeks above 20%, all models yielded Rp13,908–14,677/kg. Strong price persistence (r = 0.92; 392 training pairs) and overlapping price-history features plausibly favored Ridge regularization. Gradient Boosting outperformed Random Forest in every scenario, with rankings unchanged as features increased from six to 16 (Tables 1 and 3). These comparisons remain descriptive and grid-dependent. S1 Ridge (α = 1) was selected by the lowest mean validation RMSE (Rp6,714.46/kg; SD Rp1,667.92/kg).

Table 4 reports results for all 12 combinations and Naive Forecast on the 76 test weeks. Ridge achieved the lowest RMSE within each scenario. S4 Ridge had the lowest test RMSE

(Rp7,399.25/kg; MAE Rp5,461.39/kg; MAPE 11.15%; R^2 = 0.7056), 5.69% below Naive Forecast. The validation-selected S1 Ridge yielded RMSE Rp7,614.60/kg and MAE Rp5,731.31/kg, reducing Naive Forecast errors by Rp230.97/kg (2.94%) and Rp58.16/kg, respectively. S4's lower test error did not change the validation-based selection of S1 Ridge.

Table 3. Validation results and model selection

Scenario	Algorithm	RMSE ± SD (Rp/kg)
S1	Ridge Regression	6,714.46 ± 1,667.92
	Random Forest	9,372.44 ± 2,130.73
	Gradient Boosting	8,985.14 ± 2,178.25
S2	Ridge Regression	6,819.06 ± 1,563.56
	Random Forest	9,597.63 ± 2,224.47
	Gradient Boosting	9,055.51 ± 2,044.35
S3	Ridge Regression	6,789.61 ± 1,529.17
	Random Forest	9,604.52 ± 2,217.84
	Gradient Boosting	8,985.10 ± 2,102.04
S4	Ridge Regression	6,948.79 ± 1,501.18
	Random Forest	10,036.16 ± 2,285.18
	Gradient Boosting	8,993.73 ± 2,037.51

Table 4. Model performance comparison on the test set

Scenario	Algorithm	RMSE	MAE	MAPE	R^2
S1	Ridge Regression	7,614.60	5,731.31	11.80	0.6882
	Random Forest	7,737.27	5,802.40	11.74	0.6781
	Gradient Boosting	7,619.35	5,540.64	11.35	0.6879
S2	Ridge Regression	7,452.42	5,630.27	11.66	0.7014
	Random Forest	7,911.33	5,926.28	11.91	0.6635
	Gradient Boosting	7,718.08	5,557.38	11.40	0.6797
S3	Ridge Regression	7,553.87	5,542.49	11.28	0.6932
	Random Forest	7,658.99	5,806.29	11.71	0.6846
	Gradient Boosting	7,563.78	5,462.20	11.17	0.6924
S4	Ridge Regression	7,399.25	5,461.39	11.15	0.7056
	Random Forest	7,946.12	6,052.79	12.10	0.6605
	Gradient Boosting	7,623.70	5,478.10	11.19	0.6875
Baseline	Naive Forecast	7,845.57	5,789.47	11.98	0.6690

In S1, Ridge's test RMSE was only Rp4.75/kg below Gradient Boosting (Table 4). Relative to S1 Ridge, S4 Ridge had 3.49% higher validation RMSE, 2.83% lower final-test RMSE, and 0.25% higher RMSE in the 2024 backtest (Tables 3, 4, and 7). These period-dependent reversals and S4's 95% block-bootstrap confidence interval including zero (Table 5) do not establish a reliable gain from added features. Operational value remains untested because we did not evaluate monitoring decisions and intervention thresholds.

We compared S2, S3, and S4 with S1 within the same algorithm on the same 76 test weeks to assess the added feature groups. Table 5 reports each scenario's percentage RMSE reduction relative to S1 and its paired block-bootstrap 95% CI. Positive values indicate lower RMSE after adding features; negative values indicate higher RMSE. Using the same algorithm and test weeks makes the comparisons consistent, while the intervals quantify uncertainty in the observed differences.

On the 76-week test, S3 reduced RMSE by 0.73–1.01% (Rp55.57–78.28/kg) relative to S1 across all three algorithms, whereas S2 and S4 reduced RMSE only for Ridge. All nine 95% CIs included zero, including in the block-length sensitivity checks. Thus, the observed differences did not establish a statistically reliable gain over S1. Operational benefit remains untested because we did not evaluate monitoring thresholds or intervention outcomes.

Table 5. RMSE reduction relative to S1 (%, 95% CI)

Scenario	Ridge Regression	Gradient Boosting	Random Forest
S2	+2.13 [−1.70, 6.68]	−1.30 [−8.14, 5.02]	−2.25 [−9.13, 3.46]
S3	+0.80 [−4.01, 7.21]	+0.73 [−0.61, 2.74]	+1.01 [−2.48, 4.49]
S4	+2.83 [−3.46, 10.48]	−0.06 [−7.47, 6.94]	−2.70 [−11.24, 5.01]

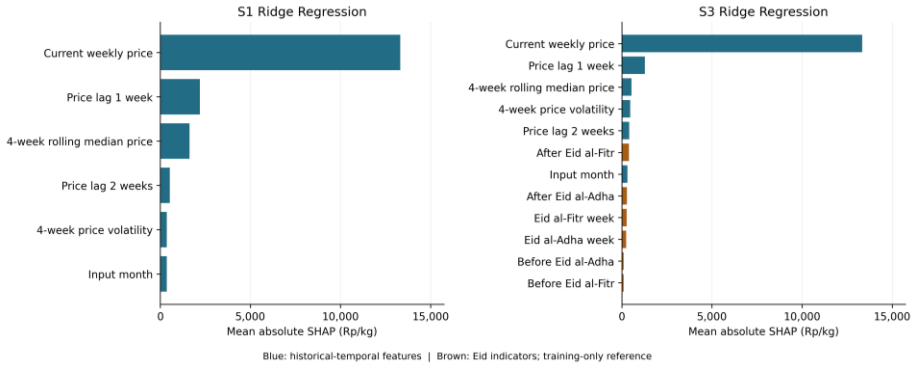

Figure 2. Mean absolute SHAP values for S1 and S3.

Figure 2 shows that the current-week price had the largest mean absolute SHAP value in S1 and S3 Ridge (Rp13,318 and Rp13,339/kg). Across four Eid weeks, the indicators contributed an average +Rp5,876/kg to S3 predictions relative to its training reference. This contribution did not consistently improve accuracy. S3 lowered error in three of four events (Table 6), and its backtest gains varied by year and algorithm (Table 7). Weather improved only Ridge (Table 5), but Ridge-only SHAP cannot explain why the tree models responded differently.

We examined model performance on the test set across four phases: normal weeks, before the holidays, during the holidays, and after the holidays. This analysis compared S1 and S3 to assess how forecasting error changed after adding holiday calendar features. Figure 3 presents the RMSE comparison for Ridge Regression across these phases.

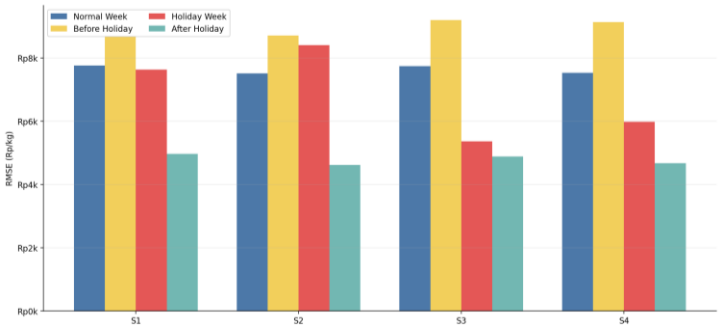

Figure 3. Comparison of Ridge Regression RMSE across holiday phases

Figure 3 shows that forecasting errors were highest before the holidays. Compared with S1, S3 reduced RMSE by 29.78% during holiday weeks but increased it by 5.95% before the

holidays. Thus, the reduction observed during holiday weeks did not extend to the preceding weeks. These phase-level figures summarize errors across multiple weeks and do not show whether S3 improved forecasts consistently across individual Eid events. The event-level analysis therefore examines whether the holiday-week reduction was shared across events or concentrated in particular ones.

Table 6. Event-level comparison of S1 and S3 errors during holiday weeks

Holiday/target week	Actual price	S1 absolute error	S3 absolute error	Error difference
Eid al-Fitr / 31 Mar 2025	63,500.00	1,265.58	7,126.69	−5,861.10
Eid al-Adha / 2 Jun 2025	34,750.00	3,396.82	1,960.51	+1,436.30
Eid al-Fitr / 16 Mar 2026	43,500.00	6,824.16	614.69	+6,209.47
Eid al-Adha / 25 May 2026	63,750.00	13,169.03	7,742.62	+5,426.40

Only four Eid weeks were tested; Table 5's 95% confidence intervals cover all 76 weeks, so event-level findings remain descriptive and exploratory. S3 Ridge reduced absolute error in three events but worsened Eid al-Fitr 2025 (Table 6). Excluding Eid al-Adha 2026 shrank the holiday-week RMSE gain from 29.78% to 4.01%, revealing event sensitivity rather than consistent calendar-feature benefits across Eid events.

Table 7. RMSE reduction relative to S1 across historical test blocks

Test year (weeks)	Algorithm	S2 (%)	S3 (%)	S4(%)
2022 (46)	Ridge Regression	+1.87	+0.95	+2.56
	Random Forest	−5.33	−4.19	−9.64
	Gradient Boosting	−1.18	−0.05	−1.52
2023 (51)	Ridge Regression	−0.94	+9.34	+7.85
	Random Forest	−4.74	−12.90	−13.40
	Gradient Boosting	+0.42	+4.10	+2.79
2024 (52)	Ridge Regression	+3.44	−2.33	−0.25
	Random Forest	+4.31	−0.12	+1.51
	Gradient Boosting	−0.06	−0.31	−0.72

Table 7's two-week pre/post-Eid spans show S3 worsened RMSE relative to S1 for Ridge (2024) and Random Forest (2022–2024); one- and three-week spans also worsened Ridge (2024). Across spans, final-test S3 gains persisted for Ridge and Gradient Boosting (0.73–1.54%) but reversed for Random Forest (+1.01% for two weeks; −1.72% for one week). Historical backtests included only five Eid weeks. Excluding the four-week price standard deviation reduced S1 Ridge validation RMSE from Rp6,714.46 to Rp6,622.39/kg. Calendar gains therefore varied by algorithm, period, and Eid span; the dispersion feature provided no benefit to Ridge validation.

Discussion

Ridge achieved the lowest RMSE in each feature scenario, but S1 led validation while S4 led the final test; S4 backtests were inconsistent (Tables 3, 4, and 7). S1 remained the final model under the validation-based selection rule, while these reversals did not establish stable external-feature gains. Current-week price dominated Ridge SHAP (Figure 2), consistent with short-term price persistence (Mahajan et al., 2026). Overlapping price-history predictors and 392 training observations may favor Ridge regularization over tree ensembles (Zhu et al., 2025), although this mechanism remains unverified.

S2 reduced Ridge test RMSE by 2.13% but worsened both tree models; its Ridge gain was absent in validation, and its 95% confidence interval included zero (Tables 3-5). Ridge may capture a weak additive weather signal, while limited training data may destabilize tree splits, and Padang weather may not represent producing areas. Ridge-only SHAP cannot explain the tree models' different responses. S3 reduced test RMSE by 0.73–1.01% across algorithms, but corresponding confidence intervals included zero, and some backtests reversed the gains. Random Forest's S3 gain also reversed under the one-week Eid span, while Ridge's pre-holiday RMSE increased. Excluding Eid al-Adha 2026 reduced Ridge's holiday-week gain from 29.78% to 4.01% (Tables 5-7; Figure 3). These results do not establish statistically reliable or temporally consistent benefits from external features.

Pardede et al. (2025) reported a small RMSE reduction between daily cayenne-pepper forecasting configurations with and without general holiday features, whereas Vairo et al. (2026) found no systematic weather benefit for monthly wheat. Differences may reflect commodity production cycles, supply networks, temporal aggregation, forecast horizons, and model design. Eid indicators provide known-ahead timing but not demand magnitude (Devianto et al., 2024). Holiday demand pressure may intensify under constrained harvests or deliveries, but sufficient market arrivals can buffer it (Devianto et al., 2024; Muflikh et al., 2024). S3 Ridge's fixed additive calendar coefficients cannot represent these unmeasured conditions, plausibly explaining lower errors in three Eid events but higher error in one (Table 6). These mechanisms, along with commodity, regional, and horizon effects, remain untested. The findings support context-specific feature evaluation against a common historical-temporal baseline: positive aggregate gains alone cannot establish reliable predictive value across algorithms, periods, and events.

Missing market-arrival, production, distribution, and realized-demand data do not invalidate observed test errors. Across 76 Padang test weeks, validation-selected S1 Ridge achieved RMSE Rp7,614.60/kg, 2.94% below Naive Forecast (Table 4). This modest, period-specific gain supports provisional price monitoring with predictions checked against observed prices, especially around Eid and abrupt changes. Pre-holiday deterioration warrants further chronological testing; the study did not evaluate intervention thresholds, decisions, or practical benefits.

Evidence is limited to one city, 76 final-test weeks, three earlier annual backtests, and four Eid events, including one influential event. Unknown supply origins and unmeasured market conditions prevent direct mechanism testing. Tuning covered only three Ridge and four tree configurations per scenario, with fixed ensemble sizes that may constrain performance, particularly boosting learning-rate interactions. Thus, rankings and feature gains depend on the tested grids, while excluded models remain unassessed. Future research should integrate forecast-available production, market arrivals, distribution, observed demand, and supply-origin weather; test producing and consuming markets; extend SHAP to both tree models; and apply repeated rolling-origin evaluation with retuning and refitting across more years, Eid events, and window lengths.

CONCLUSION

This study provides evidence that external-feature value in commodity forecasting requires jointly assessing incremental accuracy and stability. Within the tested configurations, gains varied across algorithms and periods, while Ridge's Eid gains were sensitive to individual events. Final-test confidence intervals did not establish reliable improvements over S1. The methodological implication is to compare added features against a common historical-temporal baseline on matched weeks, quantify uncertainty, and test period and event sensitivity before adoption. For Padang's price-monitoring staff, the validation-selected S1 Ridge model provides a provisional next-week reference whose errors should be reviewed as new prices

arrive; operational benefits remain untested. Evidence is limited to one market and four final-test Eid weeks. Future research should extend SHAP to tree models, integrate supply, distribution, and observed demand available at forecast time, and apply repeated rolling-origin evaluation across producing and consuming markets to test stability under different market conditions.

REFERENCES

Bahri, C. A., & Tania, K. D. (2025). Perbandingan Kinerja LSTM, Random Forest, dan SVR Berbasis Knowledge Discovery untuk Prediksi Harga Beras Sumatera Selatan. *JURIKOM (Jurnal Riset Komputer)*, *12*(5), 721–732. https://doi.org/10.30865/jurikom.v12i5.9140

Chaitra, B., & Meena, K. (2026). Fuel-driven hybrid deep learning model for forecasting finger millet prices. *Scientific Reports*, *16*(1). https://doi.org/10.1038/s41598-025-34947-8

Devianto, D., Wahyuni, E., Maiyastri, M., & Yollanda, M. (2024). The seasonal model of chili price movement with the effect of long memory and exogenous variables for improving time series model accuracy. *Frontiers in Applied Mathematics and Statistics*, *10*. https://doi.org/10.3389/fams.2024.1408381

Faisa, D. K. K., & Salam, A. (2026). Beyond Predictive Accuracy: Enhancing Parameter Stability in Multicollinear Time Series Forecasting via Regularisation. *Edumatic: Jurnal Pendidikan Informatika*, *10*(1), 70–79. https://doi.org/10.29408/edumatic.v10i1.33925

Li, Q., Wang, X., Qin, X., Yu, F., Luan, R., Ping, Y., Yang, W., & Zhang, L. (2026). A Weekly E-Commerce Seed Price Index and Short-Term Forecasting System: Evidence from Tomato, Pepper, and Lettuce. *Agriculture (Switzerland)*, *16*(15). https://doi.org/10.3390/agriculture16151620

Mahajan, P., Muley, A. A., & Fegade, M. R. (2026). Forecasting agricultural commodity prices using machine learning models. *Discover Informatics*, *1*(1), 15. https://doi.org/10.1007/s44564-026-00015-0

Manogna, R. L., Dharmaji, V., & Sarang, S. (2025). Enhancing agricultural commodity price forecasting with deep learning. *Scientific Reports*, *15*(1). https://doi.org/10.1038/s41598-025-05103-z

Matondang, M. R., Krisnamurthi, B., & Herawati, H. (2024). Price fluctuations and volatility of national strategic food commodities in Indonesia. *Agrisocionomics: Jurnal Sosial Ekonomi Pertanian,* *8*(1), 134–146. https://doi.org/10.14710/agrisocionomics.v8i1.17753

Muflikh, Y. N., Smith, C., Brown, C., Kusnadi, N., Kiloes, A. M., & Aziz, A. A. (2024). Integrating system dynamics to value chain analysis to address price volatility in the Indonesian chilli value chain. *Food Policy*, *128*. https://doi.org/10.1016/j.foodpol.2024.102713

Murtani, M., Kobi, D., & Imelda, I. (2026). Multi-commodity food price forecasting ahead of Eid al-Fitr in Indonesia (2026–2030): A comparative study of machine learning algorithms using time-series data. *MALCOM: Indonesian Journal of Machine Learning and Computer Science*, 6(1), 404–413. https://doi.org/10.57152/malcom.v6i1.2572

Mutiara, T., Al Arif, M. N. R., & Anwar, M. (2025). The Rising Food Prices in Ramadan and Its Impact on Inflation. *Signifikan: Jurnal Ilmu Ekonomi*, *14*(2). https://doi.org/10.15408/sjie.v14i2.46417

Nayak, G. H. H., Alam, M. W., Singh, K. N., Avinash, G., Kumar, R. R., Ray, M., & Deb, C. K. (2024). Exogenous variable driven deep learning models for improved price forecasting of TOP crops in India. *Scientific Reports*, *14*(1). https://doi.org/10.1038/s41598-024-68040-3

Pardede, J., Setyaningrum, A. P., Premitasari, M., & Al-Fadhlih, M. I. (2025). Impact of feature

engineering on XGBoost model for forecasting cayenne pepper prices. *Scientific Journal of Informatics, 12*(4), 651–664. https://doi.org/10.15294/sji.v12i4.32157

Salsabila, N., Damaliana, A. T., & Saputra, W. S. J. (2025). ARIMAX dengan variasi kalender Idul Adha dan Natal untuk prediksi harga komoditas cabai keriting di Jawa Timur. *Jurnal INSTEK (Informatika Sains dan Teknologi), 10*(2), 389–401. https://doi.org/10.24252/instek.v10i2.59520

Sari, M., Duran, S., Kutlu, H., Guloglu, B., & Atik, Z. (2024). Various optimized machine learning techniques to predict agricultural commodity prices. *Neural Computing and Applications*, *36*(19), 11439–11459. https://doi.org/10.1007/s00521-024-09679-x

Siaga, E., Meihana, M., Lumbantoruan, S. M., Sakagami, J.-I., & Lakitan, B. (2024). Karakter Morfo-agronomi Tanaman Cabai Merah (Capsicum annum L.) Fase Awal Vegetatif pada Kondisi Stres Jenuh Air. *Jurnal Ilmu Pertanian Indonesia*, *29*(2), 236–243. https://doi.org/10.18343/jipi.29.2.236

Utari, C. T., Arya, M. T., Sembiring, P., Rizq, M. H., & Siregar, Z. (2025). Analisis komparasi algoritma ARIMA dan LSTM pada prediksi harga cabai merah keriting harian. *Building of Informatics, Technology and Science (BITS)*, *7*(3), 1932-1942. https://doi.org/10.47065/bits.v7i3.8784

Vairo, A., Romano, F., & Grilli, L. (2026). Assessing the role of weather in commodity price forecasting: evidence from fine durum wheat. *Computational Management Science*, *23*(2). https://doi.org/10.1007/s10287-026-00576-3

Warman, R., Rahma, H., Darnetty, & Noveriza, R. (2025). Fungal diversity associated with anthracnose in chili peppers of West Sumatra, Indonesia, using conventional and NGS approaches. *Biodiversitas*, *26*(9), 4257–4266. https://doi.org/10.13057/biodiv/d260901

Yandar, K. L., Revelo Sánchez, O., & Bolaños Gonzales, M. (2026). Joint Forecasting of Energy Consumption and Generation in P2P Networks Using LSTM–CNN and Transformers. *Energies*, *19*(3). https://doi.org/10.3390/en19030760

Zhang, Q., Yang, W., Zhao, A., Wang, X., Wang, Z., & Zhang, L. (2024). Short-term forecasting of vegetable prices based on LSTM model—Evidence from Beijing's vegetable data. *Plos one*, *19*(7), e0304881. https://doi.org/10.1371/journal.pone.0304881

Zhu, Y., Chuluunsaikhan, T., Choi, J. H., & Nasridinov, A. (2025). Integrating structured and unstructured data for livestock price forecasting: a sustainability study from South Korea. *Frontiers in Sustainable Food Systems*, *9*. https://doi.org/10.3389/fsufs.2025.1613616