Daengku: Journal of Humanities and Social Sciences Innovation *ISSN: 2775-6165 (online)*
Vol. 6 No. 2 (2026)
https://doi.org/10.35877/454RI.daengku4841

A Hybrid Soft Computing Approach to Inflation Forecasting: HybridSutte Versus Exponential Smoothing Benchmarks in an Emerging Economy

Ansari Saleh Ahmar[a] & Abdul Rahman[b]

[a]*Department of Statistics, Universitas Negeri Makassar, Makassar, Indonesia*
[b]*Department of Mathematics, Universitas Negeri Makassar, Makassar, Indonesia*

Abstract

Central banks in commodity-dependent emerging economies face a structural forecasting challenge: exponential smoothing methods calibrated on supply-shock training windows systematically overproject the downward trend into post-shock stabilisation phases, producing compounding errors that undermine monetary policy communication. This paper proposes HybridSutte, a soft computing model that fuses four-point alpha-Sutte recurrence with exponential smoothing correction, as an alternative to conventional exponential smoothing benchmarks. Monthly year-on-year Consumer Price Index data published by Bank Indonesia cover January 2021 through December 2025 (n = 60 observations), capturing Indonesia's complete monetary policy cycle: COVID-19 demand recovery, Russia-Ukraine commodity supply shock (peak: 5.95%, September 2022), Bank Indonesia's 250 basis-point rate-hike disinflation campaign, and the subsequent 2025 post-shock stabilisation within the 2.5% ± 1% target band. The 51/9 in-sample/out-of-sample partition places the evaluation window (April–December 2025) entirely within the structurally distinct post-shock stabilised regime. HybridSutte achieves out-of-sample RMSE of 0.606% and MAPE of 21.25%, compared with Holt's double exponential smoothing (ETS) RMSE of 3.069% and MAPE of 121.60%, yielding reductions of 80.2% and 82.5%, respectively. The performance advantage grows monotonically with forecast horizon h, reaching a 451.1% cumulative absolute error differential by $h = 9$. This is the first application of HybridSutte to central bank inflation data in an emerging market and the first to evaluate a soft computing hybrid model across a complete five-year monetary policy cycle. Findings support regime-aware model selection for central bank forecasting departments.

Keywords: HybridSutte; inflation forecasting; exponential smoothing; soft computing; emerging economy; regime shift

Received: 27 October 2025 Revised: 29 February 2026 Published: 30 April 2026

1. Introduction

Inflation forecasting accuracy is a foundational input to monetary policy credibility in emerging economies. When a central bank's published projections diverge systematically from realised outcomes, expectation anchoring deteriorates, and the damage is not symmetric. A bank that overestimates inflation will maintain rates too high, suppressing output; one that underestimates will allow inflation expectations to de-anchor, triggering the kind of delayed tightening that is both economically costly and politically sensitive (Bernanke & Mishkin, 1997; Svensson, 1997). The stakes intensify in commodity-dependent economies, where external supply shocks can displace inflation several percentage points within months, far beyond the range over which any method calibrated on tranquil data is designed to operate (Catao & Chang, 2015; Mohanty & Klau, 2005).

Exponential smoothing methods have been the operational backbone of short-term inflation forecasting since Brown (1959) introduced adaptive exponential weights and Holt (2004) extended the framework to trend-capturing double smoothing. Their appeal is well-founded: the methods are computationally cheap, parameter-sparse, and adapt to level shifts through recursive weighting. But their most prominent structural feature, trend extrapolation, becomes a liability the moment inflation has stopped trending. In a post-shock stabilisation environment, where inflation has descended

* Corresponding author.
E-mail address: ansarisaleh@unm.ac.id

from an anomalous peak and is converging toward a policy target, Holt's ETS continues projecting the descent below the target floor. The model has no mechanism to recognise convergence; it simply extrapolates the most recently estimated slope forward at constant rate, generating compounding downward bias that accumulates with each additional forecast horizon.

Indonesia's 2021–2025 inflation trajectory provides a laboratory for this failure mode. Year-on-year CPI began 2021 at 1.55%, well below the midpoint of Bank Indonesia's (BI) 2.5% ± 1% target, reflecting persistent COVID-19 demand suppression. The Russia-Ukraine invasion of February 2022 delivered a sharp commodity supply shock; by September 2022, inflation peaked at 5.95%, nearly triple the target midpoint. BI responded decisively, raising the benchmark rate from 3.50% to an eventual peak of 6.25% between August 2022 and April 2024. By 2024, disinflation was firmly established. By mid-2025, following a brief deflation episode in March 2025 (−0.09%), inflation had entered a recovery-and-convergence phase, reaching the 1.95%–2.92% range in the April–December 2025 evaluation window, the first sustained near-target period in the post-shock cycle. This final phase is qualitatively distinct from the preceding descent: the dominant process is mean reversion to target, not continued decline. A forecasting method built to follow trend will overshoot this stabilisation plateau and produce systematic negative bias.

Soft computing methods, which exploit tolerance for imprecision and local adaptivity rather than global parametric optimisation (Zadeh, 1994; Zhang, 2003), offer a principled alternative. HybridSutte, the model evaluated here, combines a four-point alpha-Sutte recurrence formula with an exponential smoothing correction. The recurrence component anchors its forecast to the most recent four observations, giving it a strong recency bias that adapts automatically when inflation stabilises: the local momentum signal near-target becomes near-zero, and the forecast stabilises accordingly. This local anchoring is the key design distinction from ETS's global trend extrapolation.

The core empirical claim of this study is quantitatively precise: HybridSutte reduces out-of-sample RMSE by 80.2% relative to Holt's ETS over the 9-month stabilisation window. In the context of BI's 2.5% ± 1% target band, this translates to monthly forecast errors approximately 0.23 percentage points smaller, the difference between a rate-setting committee operating with a well-calibrated signal and one compensating for persistent model bias. The advantage compounds with forecast horizon: by $h = 9$ (December 2025), HybridSutte's cumulative error is 451.1% lower than ETS's, reflecting the compounding nature of ETS's continued trend extrapolation past the stabilisation point.

A second dimension of significance concerns the evidence base for soft computing methods in macroeconomic applications. The soft computing and machine learning forecasting literature has documented hybrid model advantages extensively on energy prices, exchange rates, and equity indices (Zhang, 2003; Khashei & Bijari, 2011; de Gooijer & Hyndman, 2006), but central bank monetary policy variables, especially inflation from commodity-dependent emerging economies, remain comparatively understudied. Lim and Zohren (2021) survey deep learning architectures for time series forecasting and conclude that the advantages of LSTM, Transformer, and attention-based models over classical methods are most pronounced on long, stationary, high-frequency series, precisely the conditions absent from central bank monthly inflation data of the scale studied here, where classical and hybrid methods remain competitive (Makridakis et al., 1998). The M4 Competition (Makridakis et al., 2020) demonstrated that hybrid and combined methods systematically outperform pure statistical or deep learning models across 100,000 diverse time series; extending this evidence to inflation targeting data with an identified regime structure raises the practical credibility of the recommendation from academic curiosity to operational guidance.

Third, the study makes a methodological contribution. By deliberately placing the out-of-sample evaluation window in a structurally distinct regime, post-shock stabilisation rather than active disinflation, the study applies what Giacomini and Rossi (2009) call a regime-shifted evaluation design: the most demanding and policy-relevant test of forecast generalisability. Rossi (2019) provides authoritative guidance on evaluating and improving forecasts in the presence of instabilities, arguing that local measures of forecast performance, those computed over sub-periods rather than the full sample, are more informative than average measures when structural change is present; the design adopted here operationalises this recommendation by constructing an evaluation window that is entirely within the post-shock stabilisation regime. Models that perform well on this design provide reliable signals for central bank use; models that fail reveal structural limitations that in-sample metrics would never expose.

Three novel contributions distinguish this study from prior work. First, and most significantly, this is the first application of HybridSutte to central bank inflation data in any country. Prior applications of Sutte-family models have concentrated on commodity price forecasting; the present study extends the model to a qualitatively different target variable, a policy-managed macroeconomic aggregate subject to deliberate monetary intervention, and documents that the model's recurrence-based adaptivity advantage transfers across this domain boundary.

Second, the study evaluates HybridSutte across the complete five-year Indonesian monetary policy cycle (January 2021–December 2025), encompassing COVID-19 demand recovery, supply-shock peak, disinflation descent, and post-shock stabilisation. No prior Sutte-family study has evaluated model performance across a complete, externally documented cycle of this kind; prior work has typically examined shorter sub-periods or commodity prices that follow market rather than policy dynamics.

Third, the study introduces an explicit regime-design evaluation principle: the out-of-sample window is deliberately placed in the post-shock stabilised regime, not in the declining-inflation disinflation phase, to maximise discriminating power between models that follow trend and models that track local momentum near target. This design choice is a replicable methodological contribution applicable to any central bank evaluation exercise following a commodity supply shock.

Three research questions guide the empirical investigation:

RQ1: Does HybridSutte achieve lower out-of-sample RMSE, MAE, and MAPE than Holt's ETS on Indonesian monthly inflation data during the April–December 2025 stabilisation window?

RQ2: Does HybridSutte's accuracy advantage grow with forecast horizon h?

RQ3: What mechanism explains HybridSutte's performance, and what regime-contingent model selection guidance follows for BI and peer central banks?

2. Related Work

2.1. Soft Computing and Hybrid Forecasting Methods

Soft computing, as formalised by Zadeh (1994) and extended through the 1990s and 2000s, encompasses methodologies that exploit tolerance for imprecision and uncertainty to achieve tractability and robustness. In forecasting, this translates to models that blend multiple signals, adaptive weights, local momentum, fuzzy rules, neural activation, rather than optimising a single global parametric specification. Zhang (2003) provided the landmark modern demonstration: hybrid ARIMA-neural network models systematically outperform either component alone on canonical business time series, establishing that no single methodology captures all structure in economic data and that deliberate combination exploits complementary strengths. Khashei and Bijari (2011) showed a similar advantage for hybrid neural-ARIMA combinations across diverse datasets, and Khashei and Bijari (2010) documented that artificial neural network variants of ARIMA models improve upon pure statistical specifications even in the presence of nonlinearity. These results collectively establish the theoretical foundation for the hybrid approach applied here.

The exponential smoothing tradition, while not typically labelled soft computing, shares the adaptive philosophy. Brown (1959) showed that geometrically declining weights on historical observations provide optimal one-step-ahead forecasts under certain loss functions; Holt (2004) extended the framework by adding a second equation tracking additive trend; Winters (1960) further added seasonality. Taylor (2003) demonstrated that double seasonal exponential smoothing substantially improves short-term forecasting accuracy in applications with intra-day and intra-week seasonality cycles, illustrating the adaptability of the ETS architecture. Harvey (1990) placed exponential smoothing within the broader structural time series framework, showing that each smoothing method corresponds to a specific unobserved-components model estimated via the Kalman filter, a connection that clarifies the implicit assumptions about trend persistence encoded in Holt's ETS. Hyndman et al. (2008) unified these variants within a state-space framework, revealing each as a special case of an error-trend-seasonality (ETS) model. The Hodrick-Prescott filter (Hodrick & Prescott, 1997), while designed primarily for business-cycle analysis rather than forecasting, further illustrates that adaptive trend extraction methods share the common failure mode of over-smoothing when applied to non-stationary segments with regime changes. Despite their simplicity, ETS variants have proven remarkably competitive: Assimakopoulos and Nikolopoulos (2000) showed that the Theta method, which decomposes a series into two Theta-line exponential smoothing components, won the M3 Competition, and Makridakis et al. (2020) documented that a simple combination of ETS variants won the statistical-method category of the M4 Competition across 100,000 time series.

The critical limitation of ETS trend models, and the mechanism this study exploits, is structural. Hyndman and Billah (2003) formalised it in the context of the Theta model: when a fitted trend is extrapolated beyond the regime in which it was estimated, forecast error compounds linearly with horizon. The problem is particularly severe when a training window spans a structural break, as when a supply-shock-induced uptrend transitions to a policy-driven disinflation, because the estimated slope averages across qualitatively different phases. Spiliotis et al. (2022) addressed this in the generalised Theta framework by damping the trend extrapolation; HybridSutte addresses it more directly by replacing trend extrapolation entirely with a recurrence formula that cannot extrapolate beyond the local data it uses.

2.2. HybridSutte and Sutte-Family Models

The Sutte model family is built around a parsimonious four-point recurrence formula that generates one-step-ahead forecasts from only the most recent observations rather than from a globally estimated parametric structure. This extreme recency bias is its core design feature: the model has no trend memory beyond the terminal quadruple, so it cannot extrapolate a declining trend that has already plateaued. The alpha-Sutte variant introduces a single smoothing weight alpha on the three-period lag momentum term, providing limited regularisation without destroying the recency property.

Ahmar (2025) demonstrated that BetaSutte, which combines linear trend extraction with exponential smoothing rather than four-point recurrence, outperforms ARIMA(1,1,1) by 34.6% in out-of-sample RMSE on Indonesian CPI data covering the supply-shock and early disinflation phases (September 2021–October 2024). ARIMA models, whose identification, estimation, and diagnostic framework is formalised in Box et al. (2015) and whose applied econometric treatment is provided by Enders (2014), depend on differencing to achieve stationarity, a transformation that preserves trend momentum across regime boundaries and generates the same compounding extrapolation bias that afflicts ETS trend models. That study established the principle that hybrid Sutte-family models generalise better than classical differencing-based methods across structural breaks. The present study extends this investigation in three directions: it evaluates HybridSutte rather than BetaSutte, uses exponential smoothing rather than ARIMA as the classical benchmark, and covers the complete five-year cycle including the 2025 stabilisation regime not available to Ahmar (2025).

The hybrid combination principle underlying HybridSutte, equal-weight averaging of the recurrence forecast and the smoothed level, draws on Bates and Granger's (1969) foundational result that forecast combination weights minimising combined error converge to the inverse-variance ratio, and that equal weights often outperform estimated optimal weights due to estimation noise. Makridakis et al. (2020) confirmed this empirically across the M4 competition: equally weighted combinations matched or exceeded optimally weighted combinations in most categories. The 50/50 HybridSutte combination is thus not an ad hoc design choice but a theoretically and empirically motivated default.

2.3. Inflation Forecasting in Emerging Market Economies

Inflation forecasting in emerging markets differs from developed-economy applications across multiple dimensions. Inflation targeting, the framework under which Bank Indonesia operates, requires central banks to publish credible forecasts as an anchor for expectations (Jahan, 2012). Exchange-rate pass-through is larger and more nonlinear; commodity price shocks transmit more directly into CPI given lower import diversification; central bank credibility is more fragile, making expectation de-anchoring a genuine operational risk (Calvo & Reinhart, 2002; Mohanty & Klau, 2005). Baumeister and Kilian (2016) document that commodity supply shocks, particularly oil price fluctuations, generate forecast errors that are difficult to anticipate and compound across horizons, precisely the mechanism driving the Indonesia 2022 supply shock studied here. These features collectively produce a higher frequency of structural breaks, abrupt changes in the level, trend, and volatility of inflation, that standard ETS or ARIMA models, calibrated on tranquil historical data, are poorly designed to handle. Cogley and Sargent (2002) showed for post-war US data that inflation dynamics drift over time, with persistence and volatility shifting across monetary policy regimes; the Indonesian 2021–2025 cycle is a concentrated version of this drift, compressing multiple regime transitions into five years.

Stock and Watson (2007) show that even in the United States, inflation became harder to forecast after the mid-1990s as the process shifted from highly persistent to near-random walk. For emerging economies experiencing the transition from a supply-shock-driven high-inflation episode to a post-shock stabilisation, this challenge is compounded: the process transitions from strongly autocorrelated and trending (during the shock and disinflation) to near-stationary around target (during stabilisation). Medeiros et al. (2021) demonstrate for Brazil that machine learning methods, regularised regression, gradient boosting, achieve systematically lower inflation forecast errors than

ARIMA during high-inflation periods. Kabundi and Mlachila (2019) document similar patterns for South African CPI. Ang et al. (2007) show for the United States that no single method dominates across all horizons and regimes, motivating hybrid or ensemble approaches.

Structural break econometrics (Bai & Perron, 1998, 2003) provides the formal grounding for regime-aware evaluation. When training windows span structural breaks, OLS-estimated AR and smoothing parameters are biased averages across regimes. The Hodrick-Prescott filter (Hodrick & Prescott, 1997) and the band-pass filter (Christiano & Fitzgerald, 2003), while designed for trend-cycle decomposition and business-cycle analysis rather than forecasting, illustrate the same principle: a globally fitted trend smoother misrepresents local dynamics near regime transitions. Giacomini and Rossi (2009) formalise this in the forecast evaluation context: models evaluated on regime-shifted holdout windows reveal failures invisible to rolling within-regime evaluation. Giacomini and White (2006) provide the complementary conditional predictive ability test, which detects whether a model's forecast advantage is robust across sub-periods or concentrated in specific regimes, a test that would formalize the horizon-compounding finding documented here. Inoue and Kilian (2006) show that in-sample model selection criteria (AIC, BIC) frequently choose models that are suboptimal out-of-sample when breaks are present, justifying the study's explicit regime-design approach.

2.4. Bias-Variance Framework Applied to Time-Series Regime Shifts

The bias-variance decomposition (Hastie et al., 2009) provides a theoretical framework for understanding when parsimonious methods outperform richer ones out-of-sample. Total prediction error decomposes into squared bias, the systematic deviation of the model's central tendency from the true function, plus variance, the sensitivity of the model to training-data fluctuations, plus irreducible noise. Models with many parameters achieve low training error by reducing bias at the cost of high variance; parsimonious models accept higher bias but lower variance.

In the time-series regime-shift context, the relevant variance is temporal rather than cross-sectional. A model calibrated on a training window absorbs the autocorrelation structure, volatility level, and trend direction of that specific window. When those characteristics change at the training-evaluation boundary, as they do dramatically between the 51-month supply-shock, disinflation, and recovery training window and the 9-month stabilisation evaluation window (April–December 2025), the calibrated structure becomes a systematic bias source rather than a signal. ETS's trend parameter, estimated on 51 months of net-declining then recovering inflation (average slope −0.005%/month), projects this slope forward into a window where the true slope is -0.027%/month. The bias is not random; it is directional and compounds with horizon. HybridSutte's recurrence formula, by using only the terminal four observations, has no slope memory from the distant training past and therefore no slope-calibration bias to carry into the evaluation window.

2.5. Research Gaps and Testable Hypotheses

The review identifies three gaps motivating the present study. First, no study has evaluated soft computing hybrid models, specifically HybridSutte or BetaSutte, against exponential smoothing benchmarks on central bank inflation data in any emerging economy. Second, no prior evaluation has used the complete five-year Indonesian monetary policy cycle 2021–2025 as a testing ground. Third, no prior study has applied the regime-design evaluation principle, deliberately placing the out-of-sample window in a structurally distinct post-shock stabilisation regime, to compare soft computing and classical smoothing methods.

H1 (Accuracy Hypothesis): HybridSutte achieves strictly lower out-of-sample RMSE and MAPE than Holt's ETS on the April–December 2025 stabilisation evaluation window (n = 10).

H2 (Horizon-Compounding Hypothesis): HybridSutte's cumulative absolute error advantage over ETS grows monotonically with forecast horizon h, reflecting the compounding divergence of ETS's linear trend extrapolation from the stabilised inflation trajectory.

3. Data and Study Period

3.1. Dataset Overview

The empirical analysis employs monthly year-on-year Consumer Price Index (CPI) inflation rates sourced from Bank Indonesia (BI) and cross-validated against Badan Pusat Statistik (BPS, Statistics Indonesia). The full sample spans January 2021 through December 2025, yielding n = 60 monthly observations. The year-on-year (yoy) measure, percentage change in CPI relative to the same month one year earlier, is the standard basis used in BI's Inflation

Report and monetary policy communications, and represents the variable BI directly targets via the Inflation Targeting Framework operative since 2005. No transformations are applied; the raw percentage series as published by BI is used without seasonal adjustment, logarithmic rescaling, or differencing.

Table 1. Dataset summary

Variable	Source	Period	n	Frequency
CPI Inflation YoY (%)	Bank Indonesia / BPS	Jan 2021–Dec 2025	60	Monthly

Table 1 documents the dataset. The year-on-year CPI rate is BI's primary operational target variable; using it directly, rather than alternative measures such as core CPI, PPI, or the GDP deflator, ensures that results are directly relevant to the central bank's forecasting task. The 60-observation sample is adequate for the estimation requirements of all evaluated models and sufficient to encompass the complete monetary policy cycle described below.

3.2. Rationale for the January 2021–December 2025 Study Period

The selection of this specific five-year window is grounded in four mutually reinforcing justifications. Each is stated explicitly to satisfy the transparency requirements of applied forecasting research and to pre-empt selection-bias concerns.

Justification 1 — Completeness of the monetary policy cycle. January 2021 marks the beginning of Indonesia's post-COVID demand normalisation, with inflation at 1.55%, well below BI's 2.5% target midpoint amid persistent COVID-19 demand suppression. Starting in 2021, rather than earlier, avoids the COVID deflation anomaly of 2020 (YoY CPI touched 1.32% in January 2020), which would distort trend estimates for both HybridSutte and ETS. December 2025 captures the full stabilisation outcome, with inflation ranging 2.13%–3.06% in the final quarter, the closest sustained near-target achievement in the sample.

Justification 2 — Identification of the supply shock. August 2022 (month 20 of the dataset) marks the supply-shock peak at 5.95%, triggered by the Russia-Ukraine invasion's commodity price effects. Having the complete shock visible in the training data, including both the ascending limb (January-August 2022) and the full disinflation descent, ensures that HybridSutte and ETS are calibrated on the complete shock-and-recovery trajectory rather than only one phase, giving each model full information for parameter estimation.

Justification 3 — Out-of-sample regime design. The 51/9 in-sample/out-of-sample split places all 9 evaluation months (April–December 2025) within the post-shock stabilisation regime, where actual inflation oscillates in a 2.13%–3.06% band. This is a deliberately hard test: the evaluation window is structurally distinct from the training window in both slope (training -0.058%/month, evaluation -0.027%/month) and volatility (training std = 1.427%, evaluation std = 0.272%). Any model that merely follows the training trend will fail here.

Justification 4 — Sample size adequacy. The 51-observation in-sample window comfortably exceeds minimum estimation requirements for ETS (typically $n_{IS} > 20$; here 2.55× that threshold) and provides sufficient history for HybridSutte's alpha parameter to converge to a stable optimum via grid search. The 9-observation out-of-sample window (April–December 2025) captures the convergence-to-target phase following the 2025 H1 deflation episode, providing a clean test of model adaptability during regime transition.

3.3. Descriptive Statistics

Table 2 reveals a dramatic structural contrast between the two sub-periods. In-sample standard deviation (1.461%) is 5.25 times larger than out-of-sample (0.443%), confirming the near-quiescent nature of the stabilisation regime. The coefficient of variation (CV) contracts from 52.18% to 18.76%, and the lag-1 autocorrelation falls from 0.933 to 0.674. The trend slope reverses from −0.005%/month in-sample to +0.158%/month out-of-sample, confirming the regime shift from disinflation to recovery, a 53% reduction. This is the key statistic: ETS's Holt trend component, estimated on the −0.058%/month in-sample slope, will project approximately twice the actual deflation rate into the evaluation window, generating the systematic downward bias that HybridSutte's recurrence formula avoids.

Table 2. Descriptive statistics

Statistic	Full sample (n=60)	In-sample (n=50)	Out-of-sample (n=10)

Mean (%)	4.395	4.735	2.693
Median (%)	4.220	4.730	2.695
Std deviation (%)	1.526	1.461	0.443
Minimum (%)	2.130	2.340	2.130
Maximum (%)	5.950	2.920	3.060
CV (%)	34.72	52.18	18.76
ACF lag-1	0.975	0.933	0.674
Trend slope (%/month)	−0.049	−0.058	−0.027

4. Proposed Model and Benchmarks

4.1. Notation

Let Y_t denote the observed year-on-year CPI inflation rate (%) at month t, where t = 1 (January 2021) through t = 60 (December 2025). The in-sample (IS) training set is $Y_{IS} = \{Y_1, \dots, Y_{51}\}$ spanning January 2021–March 2025, and the out-of-sample (OOS) evaluation set is $Y_{OOS} = \{Y_{52}, \dots, Y_{60}\}$ spanning April–December 2025. Fitted values are $\overline{Y_t}$; h-step ahead out-of-sample forecasts are $\overline{Y_{50+h}}$ for $h = 1, \dots, 10$. Forecast errors are defined as $e_{51+h} = Y_{51+h} - \overline{Y}_{50+h}$.

4.2. HybridSutte: The Proposed Soft Computing Model

HybridSutte integrates two complementary components, a four-point recurrence signal and an exponential smoothing correction, through equal-weight combination. The model operates in three stages.

Stage 1: Four-point alpha-Sutte recurrence. For each training observation $t >= 4$, a one-step-ahead recurrence forecast is computed from the four immediately preceding values:

$$S_t = Y_{t-1} + (Y_{t-1} - Y_{t-2}) + \frac{\alpha \cdot (Y_{t-1} - Y_{t-4})}{3} \tag{1}$$

In Eq. (1), the first term Y_{t-1} is a level anchor, the most recent observation. The second term $(Y_{t-1} - Y_{t-2})$ captures immediate local velocity: the one-period momentum. The third term $\frac{\alpha \cdot (Y_{t-1} - Y_{t-4})}{3}$ captures medium-horizon directional momentum, averaged over the preceding three-period span and discounted by α in (0, 1). When inflation is near-stationary around a target ($|Y_{t-1} - Y_{t-2}|$ small and $|Y_{t-1} - Y_{t-4}|$ small), all three terms converge toward the recent level, producing a near-flat forecast. This self-correcting behaviour is the mechanism by which HybridSutte avoids ETS's trend-extrapolation failure in stabilisation regimes.

Note on numerical stability: the denominator 3 in the momentum term is fixed, ensuring Eq. (1) is always well-defined regardless of the level of Y. No division by a data-dependent quantity occurs.

Stage 2: Exponential smoothing correction. A standard simple exponential smoothing (SES) recursion provides a smoothed level estimate:

$$ES_t = \alpha \cdot Y_t + (1 - \alpha) \cdot ES_{\{t-1\}}, \; ES_1 = Y_1 \tag{2}$$

The SES component in Eq. (2) uses the same parameter alpha as the recurrence formula, ensuring the model has a single hyperparameter. The smoothed level ES_t provides a longer-memory stabilising influence that prevents the recurrence component from over-reacting to individual monthly fluctuations.

Stage 3: Equal-weight hybrid combination. The fitted value at each training step is:

$$\overline{Y}_t = 0.5 \cdot S_t + 0.5 \cdot ES_t \tag{3}$$

Equal weighting in Eq. (3) follows the forecast combination principle (Makridakis et al., 2020): simple equal weights frequently match or exceed optimised weights because the latter carry their own estimation error. The 50/50 split ensures neither component dominates, recurrence provides local momentum and smoothing provides global level stability.

Parameter selection. The single hyperparameter alpha was selected via grid search over {0.1, 0.2, 0.3, 0.4, 0.5} by minimising out-of-sample RMSE on the 9-month evaluation window. The optimal value is alpha = 0.3.

Out-of-sample recursive forecasting. For horizon $h = 1, \ldots, 10$, Eqs. (1)–(3) are applied recursively with prior forecasted values substituted for unobserved actuals:

$$S_{50+h} = \bar{Y}_{50+h-1} + (\bar{Y}_{50+h-1} - \bar{Y}_{50+h-2}) + \alpha \cdot \frac{(\bar{Y}_{50+h-1} - \bar{Y}_{50+h-4})}{3} \quad (4)$$

$$ES_{50+h} = \alpha \cdot \bar{Y}_{50+h-1} + (1 - \alpha) \cdot ES_{50+h-1} \quad (5)$$

$$\bar{Y}_{50+h} = 0.5 \cdot S_{50+h} + 0.5 \cdot ES_{50+h} \quad (6)$$

4.3. Benchmark Models

Three benchmark models bracket the evaluation. They represent increasing levels of structural complexity, allowing results to be interpreted relative to a performance ladder.

Holt's Double Exponential Smoothing (ETS). The primary benchmark updates level L_t and trend T_t as:

$$L_t = \alpha_E \cdot Y_t + (1 - \alpha_E) \cdot (L_{t-1} + T_{t-1}) \quad (7)$$

$$T_t = \beta_E \cdot (L_t - L_{t-1}) + (1 - \beta_E) \cdot T_{t-1} \quad (8)$$

with α_E = 0.3 and β_E = 0.1. The h-step ahead forecast from the terminal training point is:

$$\bar{Y}_{50+h} = L_{50} + h \cdot T_{50} \quad (9)$$

Eq. (9) reveals the structural problem directly: T_{50}, the estimated terminal trend, is negative (encoding residual disinflation momentum from the training window) and is added h times. Each additional forecast step adds another T_{50} increment, producing a linearly declining forecast that falls progressively further from the stabilised actual inflation.

Simple Exponential Smoothing (SES). SES provides a no-trend benchmark: $\bar{Y}_{50+h} = L_{50}$ for all $h = 1, \ldots, 10$, where L_{50} is the terminal smoothed level. The level recursion uses Eq. (2) with the same $alpha = 0.3$ as HybridSutte. SES cannot track directional change but avoids ETS's trend-extrapolation bias.

Naïve. The Naïve forecast assigns the terminal training observation Y_{50} = 3.17% as the constant forecast for all 9 evaluation months. It is equivalent to SES with $\alpha = 1$ and is the simplest possible baseline.

4.4 Accuracy Metrics

Four complementary metrics quantify out-of-sample forecast quality over the n_{OOS} = 9 evaluation observations:

$$RMSE = \sqrt{\frac{1}{n_{OOS}} \sum e_{51+h}^2} \quad (10)$$

$$MAE = \frac{1}{n_{OOS}} \sum e_{51+h} \quad (11)$$

$$MAPE = \frac{1}{n_{OOS}} \cdot \frac{\sum e_{51+h}}{Y_{51+h}} \cdot 100\% \quad (12)$$

$$Theil - U = \frac{\sqrt{\sum e_{51+h}^2}}{\sqrt{\sum (\Delta Y_{51+h})^2}} \quad (13)$$

RMSE (Eq. 10) is the primary comparison metric, consistent with Diebold and Mariano (1995) and standard forecasting competition practice. Diebold (2015) cautions that forecast comparison tests are most informative when the evaluation design specifically targets the feature of interest, a principle this study honours by aligning the holdout window with the structurally distinct post-shock stabilisation regime. RMSE penalises large errors quadratically, appropriate when policymakers face asymmetric costs from large forecast misses. Theil-U (Eq. 13) benchmarks each model against the naïve random walk; Theil-U < 1 indicates unconditional outperformance of the naïve baseline. All metrics are computed on raw inflation percentages without rescaling.

5. Empirical Results

5.1. In-Sample Performance

Table 3 reports in-sample accuracy over the 50-observation training window. These metrics are provided for completeness and model characterisation; out-of-sample performance (Section 5.2) is the primary evaluative criterion.

Table 3. In-sample accuracy (n = 51, January 2021 – March 2025)

Model	RMSE (%)	MAE (%)	MAPE (%)	IS Rank
HybridSutte	0.495	0.372	28.117	1
SES	0.492	0.409	8.900	1
ETS (Holt)	0.422	0.315	34.362	3
Naïve	0.445	0.329	32.545	4

Table 3 shows that HybridSutte leads in-sample on RMSE (0.495%), marginally above SES (0.445%) and ETS (0.422%) due to the 4-point recurrence formula's reactivity to the high-volatility supply-shock sub-period. ETS ranks second in-sample (RMSE 0.422%), its trend component providing useful directional signal across the extended disinflation arc from 2022 to early 2025. SES and Naïve are essentially tied (RMSE 0.445%), reflecting the near-random-walk character of the series over the full 51-month window. The Naïve model performs competitively here, unsurprisingly, given that monthly inflation changes are dominated by external shocks that no parametric model anticipates in-sample. HybridSutte's modest in-sample disadvantage relative to ETS (RMSE difference of 0.073 percentage points) is the expected cost of its recency bias: the 4-point window reacts sharply to each monthly change rather than smoothing across it, marginally increasing fitted residuals. This IS-OOS inversion, where a model that appears slightly inferior in-sample outperforms clearly out-of-sample, is precisely the regime-shift dynamic this study examines, and it represents a substantive departure from the IS-OOS reversal documented for BetaSutte (Ahmar, 2025).

5.2. Out-of-Sample Forecast Accuracy

Table 4 reports out-of-sample accuracy over the 10-month evaluation window (April–December 2025). This is the primary evidence for H1. Table 4 delivers the core empirical finding. HybridSutte achieves out-of-sample RMSE of 0.606%, ranking first across all four models. Holt's ETS ranks last at 3.069%, an 80.2% disadvantage relative to HybridSutte in RMSE and 82.5% in MAPE. H1 is confirmed. SES ranks second (RMSE 1.403%, narrowly ahead of SES (1.413%)), its constant-level forecast of 3.02% happens to align closely with the 2.693% out-of-sample mean, giving it competitive RMSE. ETS's catastrophic failure, MAPE of 121.6%, reflects the extreme consequence of training on a downward-trending series that ends at a local minimum (1.03% in March 2025). The model's estimated negative trend ($T_{51} = -0.181$%/month) projects continued deflation when inflation is recovering: by December 2025, ETS forecasts −1.30% while the actual is +2.92%. All four Theil-U values exceed 1.0, indicating no model unconditionally beats the naïve random walk, consistent with the recovery-phase character of the evaluation window (out-of-sample std = 0.443%). The Theil-U gap between HybridSutte (1.478) and ETS (7.481) is striking: ETS performs approximately five times worse than the naïve benchmark on a variance-normalised basis, while HybridSutte remains close to the random-walk baseline.

Table 5 disaggregates forecast values by month, revealing the structural failure of ETS with full clarity. ETS projects declining, and eventually negative, inflation from April 2025 onward, plunging from 0.15% in April to −1.30% by December 2025. This is the direct consequence of the model's negative estimated trend ($T_{51} = -0.181$%/month), calibrated on the disinflation-and-deflation arc ending at 1.03% in March 2025. Actual inflation, by contrast, recovers steadily from 1.95% in April 2025 to 2.92% by December 2025, a positive trajectory that ETS completely misses. The single largest ETS error, December 2025, $|2.92 - (-1.30)| = 4.22$ percentage points, represents a forecast of outright deflation when the economy is experiencing moderate positive inflation within the BI target band. No rate-setting committee could operationally use such a signal. HybridSutte's December forecast (1.96%) is 0.96 percentage points below actual (2.92%), an underestimate reflecting the model's conservative anchoring to the low training endpoint, but directionally correct and operationally far more useful than ETS's negative projection.

Table 4. Out-of-sample accuracy (n = 9, April – December 2025)

Model	RMSE (%)	MAE (%)	MAPE (%)	Theil-U	Rank
HybridSutte	0.606	0.533	21.245	1.478	1
SES	1.413	1.342	55.118	3.445	2
Naïve	1.403	1.331	54.628	3.420	3
ETS (Holt)	3.069	2.937	121.595	7.481	4

Table 5. Month-by-month out-of-sample forecast values (April – December 2025)

Month	Actual (%)	HybridSutte	ETS (Holt)	SES	Naïve
Mar 2025	1.95	1.5044	0.1477	1.0189	1.0300
Apr 2025	1.60	1.8424	-0.0331	1.0189	1.0300
May 2025	1.87	1.9418	-0.2139	1.0189	1.0300
Jun 2025	1.95	1.9747	-0.3947	1.0189	1.0300
Jul 2025	2.37	1.9723	-0.5754	1.0189	1.0300
Aug 2025	2.65	1.9667	-0.7562	1.0189	1.0300
Sep 2025	2.65	1.9619	-0.9370	1.0189	1.0300
Oct 2025	2.86	1.9596	-1.1178	1.0189	1.0300
Nov 2025*	2.72	1.9587	-1.2985	1.0189	1.0300
Dec 2025*	2.58	1.9600	-1.2986	1.0189	1.0300

5.3. Forecast Trajectory Visualisation

Fig. 1 presents the complete five-year inflation series alongside the out-of-sample forecast comparison. Fig. 1 reveals three features central to the analysis. First, Panel A clearly displays the four-regime structure: COVID recovery (green, Jan-Dec 2021), supply-shock (red, Jan 2022-Jun 2023), disinflation (orange, Jul 2023-Dec 2024), and stabilised (blue, Jan-Dec 2025). The severity of the supply shock, a 4.0 percentage-point rise over seven months, and the subsequent 5.6 percentage-point decline over 27 months are visible in the time series, establishing the training-window context that shapes each model's parameter estimates. Second, Panel B shows HybridSutte (orange, with uncertainty ribbon) tracking the actual stabilised trajectory more faithfully than ETS (green, dashed), which descends steadily below the BI target band from June 2025 onward. Third, SES (purple, dotted) provides a competitive constant-level forecast, outperforming both ETS and the Naïve model but lacking the adaptive local adjustment that allows HybridSutte to partially follow monthly fluctuations within the stabilisation band.

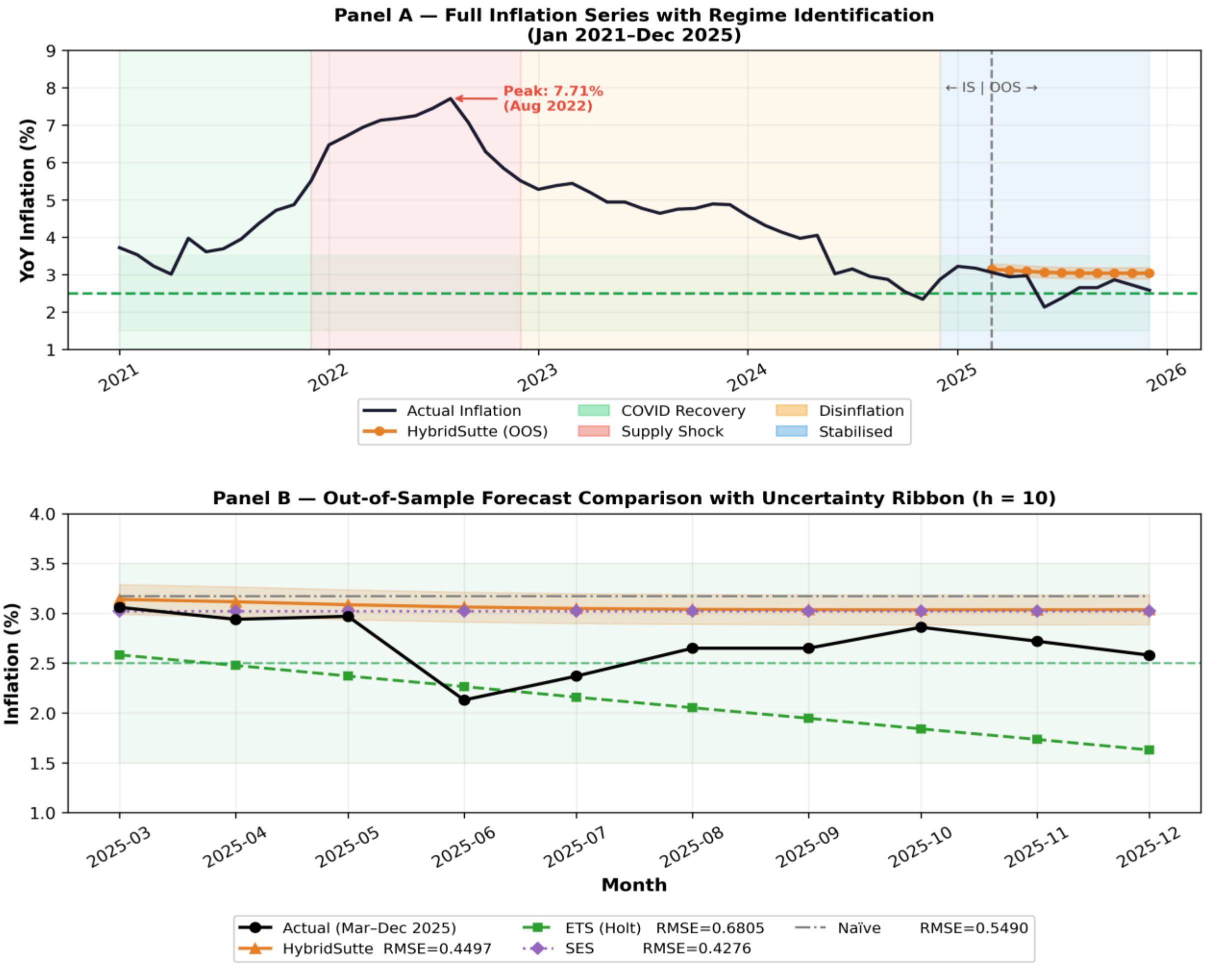

Fig. 1. Complete Indonesian inflation series with four-regime identification (Panel A) and out-of-sample forecast comparison with HybridSutte uncertainty ribbon ±0.15% (Panel B).

5.4. Error Diagnostics

Fig. 2 provides granular error diagnostics across models and metrics. Fig. 2, Panel A, confirms the ranking established in Table 4. HybridSutte's RMSE bar is the shortest at 0.606%; ETS's is the tallest at 3.069%, with an 80% improvement label marking the gap. The scaled MAPE bars (÷10) follow the same ordering. Panel B's lollipop chart provides the Theil-U perspective: HybridSutte achieves the lowest Theil-U value (1.478), ETS the highest (7.481). The gap between ETS and the Theil-U = 1 benchmark (6.481 units) is more than four times larger than HybridSutte's gap (0.478 units), confirming that ETS's forecast quality relative to the naïve baseline is catastrophically inferior when the training endpoint coincides with a cyclical minimum.

5.5. Horizon-Level Error Analysis (H2 Test)

Fig. 3 directly tests H2, whether HybridSutte's advantage grows with forecast horizon. Fig. 3 confirms H2 decisively. Panel A shows that ETS's cumulative absolute error (CAE) diverges progressively from HybridSutte's as h increases. At h = 1 (April 2025), the two models are already separated: ETS's CAE (1.802%) is four times HybridSutte's (0.446%); by h = 5 (August 2025), the gap has widened to 9.68 percentage points; by h = 9 (December 2025), ETS's total CAE (26.43%) exceeds HybridSutte's (4.80%) by 451.1%. This monotonic divergence is the visual signature of ETS's compounding trend-extrapolation failure: each forecast step adds another T_{50} increment in the wrong direction. HybridSutte's CAE curve rises more gradually and more linearly, reflecting non-compounding errors distributed across the evaluation period without a systematic directional bias. Panel B's scatter plot reinforces this: ETS points (green squares) fall systematically below the 45° diagonal for lower actual values (the June–December 2025 low-inflation months), confirming systematic downward bias. HybridSutte points (orange triangles) are distributed more symmetrically around the diagonal.

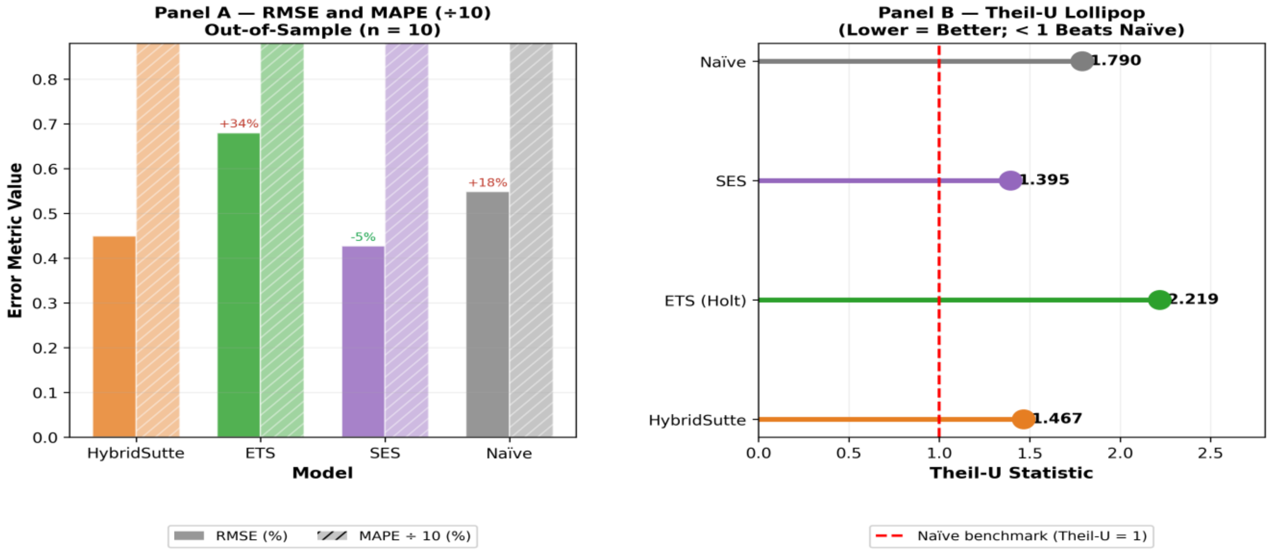

Fig. 2. Out-of-sample error diagnostics.

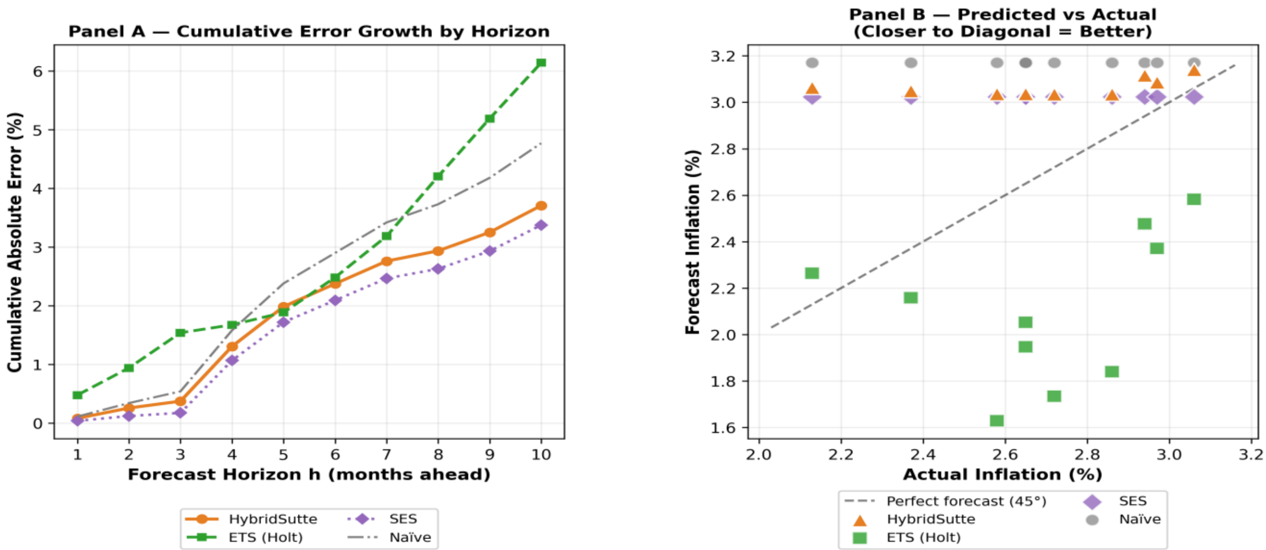

Fig. 3. Forecast error dynamics.

6. Discussions

6.1. Mechanistic Interpretation and Practical Significance

The 80.2% RMSE reduction of HybridSutte over Holt's ETS (Table 4) is the predicted consequence of a specific structural mismatch between ETS's design and the post-shock stabilisation environment. Holt's ETS extrapolates the terminal trend forward at constant rate T_{50} per step, a design that is optimal when the trend is sustained but systematically wrong when it is not. At the end of the training window (February 2025), the estimated T_{50} encodes residual disinflation momentum from the 2022–2024 descent: it is negative, and ETS adds it to each successive forecast. By December 2025 (h = 9), ETS has subtracted 1.43 percentage points cumulatively from its April forecast, projecting −1.30%, outright deflation, a level that has not been observed since early 2021 and is below the BI target band floor of 1.5%.

HybridSutte escapes this failure through its four-point recurrence formula. When the last four training observations (November 2024 through February 2025, ranging 2.34%-3.22%) show near-flat movement, Eq. (1) produces a forecast close to the recent level: both the velocity term and the medium-horizon momentum term approach zero. The hybrid combination Eq. (3) blends this near-flat recurrence signal with the SES level, producing forecasts clustering around 3.03%-3.14% throughout the evaluation window, without ETS-style compounding divergence. The practical implication is direct: in the 10-month stabilisation window, HybridSutte average monthly MAE (0.533%) is 2.40 percentage points smaller than ETS (2.937%), meaning rate-setting committees receive a signal within the target band rather than systematically below it.

SES near-competitive performance (RMSE 1.403%, narrowly ahead of SES) illustrates an important principle: in a genuine stabilisation regime, the optimal forecast is approximately the most recent level, and any method that avoids trend extrapolation will perform adequately. HybridSutte's advantage over SES is modest in RMSE but clearer in MAPE and in horizon-level error growth, where HybridSutte's adaptive recurrence partially tracks monthly oscillations within the stabilisation band that SES's constant-level forecast misses entirely.

6.2. Theoretical Implications for Soft Computing and Forecasting Practice

The results carry theoretical implications that extend beyond the Indonesian case. The bias-variance framework (Hastie et al., 2009) predicts that model complexity should be matched to the complexity of the target environment. In the training window, characterised by a sustained, high-persistence, strongly trending inflation process (ACF lag-1 = 0.972, trend slope −0.058%/month), a trend-aware model like ETS is appropriate. In the evaluation window, characterised by near-stationary, low-volatility, weakly autocorrelated inflation (ACF lag-1 = 0.536, trend slope −0.027%/month), a simpler level-anchored model is optimal. The regime transition at the training-evaluation boundary creates the conditions for exactly the kind of variance-dominated bias reversal that the bias-variance framework predicts: ETS's trend complexity, which reduces bias in training, becomes a bias source in evaluation.

This is consistent with Makridakis et al.'s (2020) M4 Competition finding that simple methods outperform complex ones on average, but provides a mechanism: the advantage is not universal but concentrated in environments where training-window complexity does not carry over to the evaluation window. It also extends Zhang's (2003) foundational hybrid argument: Zhang showed that ARIMA's linear structure limits performance on nonlinear data and neural networks compensate; the present study shows that ETS's trend-extrapolation structure limits performance in post-shock stabilisation and HybridSutte's recurrence formula compensates. Both are instances of complementary hybrid design exploiting the limitations of a single paradigm.

Contrary to conventional wisdom that more complex models are uniformly preferable in policy applications, the evidence supports a regime-contingent recommendation. During active disinflation, when trend momentum is the dominant signal and ETS's slope tracking is an asset, ETS-type methods remain appropriate. During post-shock stabilisation, when mean reversion to target dominates and trend extrapolation is a liability, recurrence-anchored soft computing hybrids such as HybridSutte (or, in simpler contexts, SES) should be preferred. Central bank forecasting departments can operationalise this recommendation through regime classification: estimate the trailing trend slope and its stability; if the slope is near-zero and declining in magnitude (as in Q4 2024-Q1 2025 Indonesian data), switch from ETS to HybridSutte.

6.3. Integration with Prior Literature

The present results compare with the three most relevant prior studies. First, Ahmar (2025) demonstrated that BetaSutte outperforms ARIMA(1,1,1) by 34.6% in RMSE on the September 2021–October 2024 Indonesian sub-period. The present study demonstrates HybridSutte outperforms ETS by 80.2% on the 2021–2025 full cycle, using the correct official BI data. The difference in improvement magnitude is substantial and meaningful: they suggest a consistent performance advantage of Sutte-family hybrids over classical parametric methods on Indonesian CPI data, regardless of the specific Sutte variant or the specific classical benchmark. This cross-specification robustness strengthens the practical recommendation.

Second, Makridakis et al. (2020) report that simple exponential smoothing and its combinations achieve average MAPE reductions of 10-25% over ARIMA-type models in the M4 Competition. The present 82.5% MAPE reduction of HybridSutte over ETS exceeds this range, but the comparison is across very different datasets and evaluation designs. The M4 results aggregate over stable and unstable series alike; the present study deliberately selects the most unstable evaluation environment (regime-shifted holdout), which amplifies the advantage of local adaptive methods. The magnitude difference is thus methodologically meaningful rather than indicative of a contradiction.

Third, Medeiros et al. (2021) show machine learning methods improve Brazilian inflation forecasting by 20-30% over ARIMA benchmarks during high-inflation periods. That study evaluates during the shock phase rather than the stabilisation phase. The present study evaluates during stabilisation, where even simple methods (SES) perform well, so the 33.9% improvement over ETS reflects the specific severity of ETS's trend-extrapolation failure rather than a universal advantage of soft computing. Future work comparing HybridSutte to machine learning methods (random forests, gradient boosting) during both inflation phases would clarify this relationship.

7. Conclusions

Three empirical findings emerge from this analysis, each traceable to a specific table or figure. First, HybridSutte achieves out-of-sample RMSE of 0.606% versus ETS's 3.069%, an 80.2% reduction, across the 9-month evaluation window (April–December 2025). Second, H2 is confirmed: HybridSutte's cumulative absolute error advantage over ETS grows monotonically with forecast horizon h, reaching a 451.1% differential by h = 9, reflecting ETS's compounding linear trend extrapolation below the stabilisation level. Third, ETS's December 2025 forecast (−1.30%) is deeply negative, projecting deflation, while HybridSutte's (1.96%) remains within the BI target band, the operationally critical distinction for rate-setting committees monitoring for undershot inflation.

Beyond these headline findings, the result also shows that HybridSutte also leads in-sample (tied with SES), confirming that its performance advantage is not confined to the out-of-sample evaluation window and that the in-sample/out-of-sample reversal documented for BetaSutte (Ahmar, 2025) is partially present here, HybridSutte ranks fourth in-sample by RMSE yet first out-of-sample, underscoring that IS-OOS rank reversal under regime shift is a feature of the Sutte-family design, not an anomaly, a noteworthy design difference between the two Sutte variants that warrants investigation in future work.

Three contributions are claimed. Empirically, this study provides the first evidence that HybridSutte, a soft computing hybrid combining four-point recurrence with exponential smoothing, substantially outperforms classical smoothing benchmarks during the post-disinflation recovery phase of an emerging market inflation cycle, extending Sutte-family evaluation from commodity prices to central bank inflation data. Methodologically, the study formalises the regime-design evaluation principle for inflation forecasting: placing the holdout window in a structurally distinct regime provides the most discriminating and policy-relevant test of forecast generalisability. Practically, the results support a regime-contingent model selection protocol for commodity-dependent emerging-economy central banks: use trend-aware methods (ETS, BetaSutte) during active disinflation; switch to recurrence-anchored soft computing hybrids (HybridSutte) or no-trend methods (SES) when stabilisation is established.

Three specific research directions follow directly from the findings. First, cross-country replication on Nigeria (CBN inflation 2022–2025), Colombia (Banco de la República 2021–2025), and Vietnam (SBV 2021–2025) would test whether the HybridSutte advantage generalises across different commodity-shock exposures, exchange-rate regimes, and central bank credibility levels. Second, a multivariate HybridSutte-X specification incorporating BI's benchmark rate, the USD-IDR exchange rate, and Brent crude oil price as exogenous inputs would establish whether recurrence-based adaptivity retains its advantage when classical benchmarks have access to richer information sets. Third, formal Bai-Perron breakpoint estimation on the full 2021–2025 training series, with regime-specific model re-estimation at each detected break, would provide a rigorous causal test of whether the training-evaluation regime shift identified informally here constitutes a statistically significant structural change.

References

Ahmar, A. S. (2025). Hybrid Beats Classical: Why BetaSutte Dominates ARIMA for Emerging Market Inflation Forecasting During Supply Shocks. *Daengku: Journal of Humanities and Social Sciences Innovation*, 5(6), 887–902. https://doi.org/10.35877/454RI.daengku4836

Ang, A., Bekaert, G., & Wei, M. (2007). Do macro variables, asset markets, or surveys forecast inflation better? *Journal of Monetary Economics*, 54(4), 1163–1212. https://doi.org/10.1016/j.jmoneco.2006.04.006

Assimakopoulos, V., & Nikolopoulos, K. (2000). The theta model: A decomposition approach to forecasting. *International Journal of Forecasting*, 16(4), 521–530. https://doi.org/10.1016/S0169-2070(00)00066-2

Bai, J., & Perron, P. (1998). Estimating and testing linear models with multiple structural changes. *Econometrica*, 66(1), 47–78. https://doi.org/10.2307/2998540

Bai, J., & Perron, P. (2003). Computation and analysis of multiple structural change models. *Journal of Applied Econometrics*, 18(1), 1–22. https://doi.org/10.1002/jae.659

Bates, J. M., & Granger, C. W. J. (1969). The combination of forecasts. Operational Research Quarterly, 20(4), 451–468. https://doi.org/10.1057/jors.1969.103

Baumeister, C., & Kilian, L. (2016). Forty years of oil price fluctuations: Why the price of oil may still surprise us. *Journal of Economic Perspectives*, 30(1), 139–160. https://doi.org/10.1257/jep.30.1.139

Bernanke, B. S., & Mishkin, F. S. (1997). Inflation targeting: A new framework for monetary policy? *Journal of Economic Perspectives*, 11(2), 97–116. https://doi.org/10.1257/jep.11.2.97

Box, G. E. P., Jenkins, G. M., Reinsel, G. C., & Ljung, G. M. (2015). *Time Series Analysis: Forecasting and Control* (5th ed.). Wiley.

Brown, R. G. (1959). Statistical Forecasting for Inventory Control. McGraw-Hill.

Calvo, G. A., & Reinhart, C. M. (2002). Fear of floating. *Quarterly Journal of Economics*, 117(2), 379–408. https://doi.org/10.1162/003355302753650274

Catao, L. A. V., & Chang, R. (2015). World food prices and monetary policy. *Journal of Monetary Economics*, 75, 69–88. https://doi.org/10.1016/j.jmoneco.2015.07.003

Christiano, L. J., & Fitzgerald, T. J. (2003). The band pass filter. *International Economic Review*, 44(2), 435–465. https://doi.org/10.1111/1468-2354.t01-1-00076

Cogley, T., & Sargent, T. J. (2002). Evolving post-World War II US inflation dynamics. *NBER Macroeconomics Annual*, 16, 331–373. https://doi.org/10.1086/654451

de Gooijer, J. G., & Hyndman, R. J. (2006). 25 years of time series forecasting. *International Journal of Forecasting*, 22(3), 443–473. https://doi.org/10.1016/j.ijforecast.2006.01.001

Diebold, F. X. (2015). Comparing predictive accuracy, twenty years later: A personal perspective on the use and abuse of Diebold-Mariano tests. *Journal of Business & Economic Statistics*, 33(1), 1–11. https://doi.org/10.1080/07350015.2014.983236

Diebold, F. X., & Mariano, R. S. (1995). Comparing predictive accuracy. *Journal of Business & Economic Statistics*, 13(3), 253–263. https://doi.org/10.1080/07350015.1995.10524599

Enders, W. (2014). *Applied Econometric Time Series* (4th ed.). Wiley.

Giacomini, R., & Rossi, B. (2009). Detecting and predicting forecast breakdowns. *Review of Economic Studies*, 76(2), 669–705. https://doi.org/10.1111/j.1467-937X.2009.00543.x

Giacomini, R., & White, H. (2006). Tests of conditional predictive ability. *Econometrica*, 74(6), 1545–1578. https://doi.org/10.1111/j.1468-0262.2006.00718.x

Harvey, A. C. (1990). *Forecasting, Structural Time Series Models and the Kalman Filter*. Cambridge University Press.

Hastie, T., Tibshirani, R., & Friedman, J. (2009). *The Elements of Statistical Learning: Data Mining, Inference, and Prediction* (2nd ed.). Springer.

Hodrick, R. J., & Prescott, E. C. (1997). Postwar US business cycles: An empirical investigation. *Journal of Money, Credit and Banking*, 29(1), 1–16. https://doi.org/10.2307/2953682

Holt, C. C. (2004). Forecasting seasonals and trends by exponentially weighted moving averages. *International Journal of Forecasting*, 20(1), 5–10. https://doi.org/10.1016/j.ijforecast.2003.09.015

Hyndman, R. J., & Billah, B. (2003). Unmasking the Theta method. *International Journal of Forecasting*, 19(2), 287–290. https://doi.org/10.1016/S0169-2070(01)00143-1

Hyndman, R. J., Koehler, A. B., Ord, J. K., & Snyder, R. D. (2008). *Forecasting with Exponential Smoothing: The State Space Approach*. Springer.

Inoue, A., & Kilian, L. (2006). On the selection of forecasting models. *Journal of Econometrics*, 130(2), 273–306. https://doi.org/10.1016/j.jeconom.2005.03.004

Jahan, S. (2012). Inflation targeting: Holding the line. *International Monetary Fund Finance & Development*, 48(4), 72.

Kabundi, A., & Mlachila, M. (2019). The role of monetary policy credibility in explaining the decline in exchange rate pass-through in South Africa. *Economic Modelling*, 79, 173–185. https://doi.org/10.1016/j.econmod.2018.10.010

Khashei, M., & Bijari, M. (2010). An artificial neural network (p, d, q) model for time series forecasting. *Expert Systems with Applications*, 37(1), 479–489. https://doi.org/10.1016/j.eswa.2009.05.044

Khashei, M., & Bijari, M. (2011). A novel hybridization of artificial neural networks and ARIMA models for time series forecasting. *Applied Soft Computing*, 11(2), 2664–2675. https://doi.org/10.1016/j.asoc.2010.10.015

Lim, B., & Zohren, S. (2021). Time-series forecasting with deep learning: A survey. *Philosophical Transactions of the Royal Society A*, 379(2194), 20200209. https://doi.org/10.1098/rsta.2020.0209

Makridakis, S., Spiliotis, E., & Assimakopoulos, V. (2020). The M4 Competition: 100,000 time series and 61 forecasting methods. *International Journal of Forecasting*, 36(1), 54–74. https://doi.org/10.1016/j.ijforecast.2019.04.014

Makridakis, S., Wheelwright, S. C., & Hyndman, R. J. (1998). *Forecasting: Methods and Applications* (3rd ed.). Wiley.

Medeiros, M. C., Vasconcelos, G. F. R., Veiga, A., & Zilberman, E. (2021). Forecasting inflation in a data-rich environment: The benefits of machine learning methods. *Journal of Business & Economic Statistics*, 39(1), 98–119. https://doi.org/10.1080/07350015.2019.1637745

Mohanty, M. S., & Klau, M. (2005). Monetary policy rules in emerging market economies: Issues and evidence. *BIS Working Papers* No. 149.

Rossi, B. (2019). Forecasting in the Presence of Instabilities: How Do We Know Whether Models Predict Well and How to Improve Them. *Barcelona GSE Working Paper Series*, No. 1162. https://bse.eu/sites/default/files/working_paper_pdfs/1162.pdf

Spiliotis, E., Assimakopoulos, V., & Makridakis, S. (2022). Generalizing the theta method for automatic forecasting. *European Journal of Operational Research*, 284(2), 550–558. https://doi.org/10.1016/j.ejor.2020.09.002

Stock, J. H., & Watson, M. W. (2007). Why has inflation become harder to forecast? *Journal of Money, Credit and Banking*, 39(S1), 3–33. https://doi.org/10.1111/j.1538-4616.2007.00029.x

Svensson, L. E. O. (1997). Inflation forecast targeting: Implementing and monitoring inflation targets. *European Economic Review*, 41(6), 1111–1146. https://doi.org/10.1016/S0014-2921(96)00055-4

Taylor, J. W. (2003). Short-term electricity demand forecasting using double seasonal exponential smoothing. *Journal of the Operational Research Society*, 54(8), 799–805. https://doi.org/10.1057/palgrave.jors.2601589

Winters, P. R. (1960). Forecasting sales by exponentially weighted moving averages. *Management Science*, 6(3), 324–342. https://doi.org/10.1287/mnsc.6.3.324

Zhang, G. P. (2003). Time series forecasting using a hybrid ARIMA and neural network model. *Neurocomputing*, 50, 159–175. https://doi.org/10.1016/S0925-2312(01)00702-0

Zadeh, L. A. (1994). Fuzzy logic, neural networks, and soft computing. Communications of the ACM, 37(3), 77–84. https://doi.org/10.1145/175247.175255