Can Machine Learning Replace VAR?
A Critical Evaluation in the Moroccan Inflation Context

Mariam Belamar, Raby Guerbaz[1]

Hassan II University, Ain Sebaa Faculty, MAEG laboratory

Morocco

Abstract

Inflation is never merely a statistic. It is the visible surface of a complex, historically contingent phenomenon shaped by supply chains, monetary institutions, exchange-rate regimes, and the psychology of expectations. This paper asks whether machine learning (ML) algorithms— Random Forest, XGBoost, LASSO-regularised VAR, and Long Short-Term Memory networks—can replace or substantively outperform traditional Vector Autoregressive (VAR) frameworks for inflation analysis in Morocco, a small open economy whose price dynamics are most conditioned by import dependence, a managed exchange-rate peg, and thin financial markets. Drawing on quarterly Moroccan macroeconomic data (1990–2024) and a rolling out-of-sample evaluation framework, we document a regime-contingent pattern: ML models achieve statistically significant forecast improvements during structurally turbulent periods— most acutely during the 2022–2023 commodity price shock—while VAR retains decisive advantages in structural identification, causal inference, and the counterfactual policy exercises that central banks require. We conclude that ML *complements* but cannot *replace* VAR in the Moroccan monetary policy toolkit, and we sketch a hybrid forecasting architecture appropriate for Bank Al-Maghrib's ongoing inflation-targeting transition.

Keywords: Inflation forecasting; Machine learning; Vector autoregres-sion; Morocco; Exchange-rate pass-through; Emerging economies; Hybrid models

[1] The authors thank Bank Al-Maghrib and the Haut Commissariat au Plan for statistical access.
JEL Classification: C53, C45, E31, E52, O55

1 Introduction

1.1 Inflation as Phenomenon

In the standard macroeconomic textbook, inflation appears as a well-behaved variable—an index number that rises or falls in response to identifiable causal forces: the money supply, the output gap, imported commodity prices, and expectations. Yet practitioners who have navigated Morocco's price history across four turbulent decades will attest that this representation captures only the most tractable layer of a deeper social and economic phenomenon. Inflation, as Bresciani-Turroni [9] observed in his canonical study of the German hyperinflation, is ultimately "a human event"—one that feeds on expectations, distributes economic pain unevenly, and generates institutional and political responses that alter the very structural parameters models presuppose.

The Moroccan case makes this phenomenological complexity unusually transparent. Between 2000 and 2019, consumer price inflation averaged barely 1.8% annually [24]—a figure that might suggest a benign, easily modelled process. Then, within three years, Morocco experienced a deflationary episode (2020), a mild recovery-driven acceleration (2021), and a sharp commodity price surge that drove headline CPI to 8.3% year-on-year in February 2023 [5]. This sequence was not a failure of monetary policy: it reflected the intersection of a global energy shock, a severe domestic drought, and a managed exchange-rate regime illsuited to absorbing external price volatility [20]. Understanding this sequence—and forecasting the next one—requires analytical frameworks capable of accommodating structural discontinuity, nonlinear transmission, and regime-dependent determinant hierarchies.

The question this paper explores is whether machine learning provides such frameworks, or whether the Vector Autoregressive model, suitably augmented, remains the more appropriate workhorse for Moroccan inflation analysis.

1.2 The Methodological Stakes

The debate between econometrics and machine learning has moved far beyond academic tribalism. At its core, it concerns three distinct analytical objectives that are often conflated in the forecasting literature: *prediction*—accurately anticipating future price-level changes; *identification*—isolating the causal contribution of individual shocks; and *communication*— translating analytical findings into the policy language of central bank committees [31, 38]. VAR and ML methods are optimised for different subsets of these objectives, and their comparative advantages shift depending on which objective is prioritisé.

For Bank Al-Maghrib, an institution navigating a consequential transition from a fixed exchange-rate regime toward greater flexibility and a formal inflation-targeting framework [25], all three objectives are simultaneously relevant. Short-run forecasting accuracy matters for operational decisions; structural identification matters for understanding the transmission mechanism; and communicative transparency matters for maintaining the credibility that effective inflation targeting requires. A methodology that excels at prediction while sacrificing identification and communicability may therefore be less useful than its raw forecast statistics suggest.

1.3 Contribution and Outline

This paper makes three contributions. First, we frame the ML-versus-VAR comparison within an explicit theory of inflation as a multi-regime phenomenon, allowing methodological performance to be evaluated not only statistically but against the demands of the phenomenon itself. Second, we conduct a rolling out-of-sample forecasting comparison across crisis and non-crisis sub-periods, demonstrating that methodological rankings are regime-contingent rather than universal. Third, we draw implications for the analytical architecture of Bank Al-Maghrib in the specific context of its inflation-targeting transition.

The paper is organised as follows. Section 2 characterises Moroccan inflation as a structural phenomenon. Section 3 reviews the relevant literature on VAR and ML methods. Section 4 presents the analytical frameworks compared. Section 5 describes the empirical strategy and summarises the key findings. Section 6 discusses the implications. Section 7 concludes.

2 Moroccan Inflation as Structural Phenomenon

2.1 The Anatomy of Moroccan Price Dynamics

Before asking how to model Moroccan inflation, it is necessary to understand what kind of process one is modelling. Three structural features, in combination, define the Moroccan inflation phenomenon and distinguish it from the processes for which most ML and econometric benchmarks were originally designed.

Import dependence and the pass-through problem. Morocco imports more than 90% of its primary energy needs and approximately 50% of its cereal consumption in drought years [20]. This structural import dependence means that international commodity price shocks transmit rapidly and potently into domestic consumer prices—a pattern documented empirically by Abbad et al. [1], who estimated significant but regime-varying exchange-rate pass-through (ERPT) coefficients. The pass-through is not, however, linear: energy subsidy schemes historically absorbed much of the first-round transmission, while their partial dismantlement in 2014 altered the regime in ways that single-equation models cannot capture.

The managed exchange rate as shock filter and amplifier. Morocco's exchange-rate management—historically a peg to a euro-dollar basket with a crawling band progressively widened since 2018—means that the nominal exchange rate does not freely absorb external shocks. As Zouhar & Motaouakkil [39] document, this design partially insulates domestic inflation from short-run com-modity price volatility but simultaneously concentrates adjustment pressure on domestic prices during episodes of significant external disequilibrium, generating abrupt inflation accelerations that linear time-series models systematically underpredict.

Monetary transmission attenuation. Morocco's banking sector, while formally sophisticated, exhibits features—concentrated ownership, limited interbank market development, and a large informal economy—that attenuate the pass-through from policy rates to retail lending rates and, ultimately, to aggregate demand and inflation [15, 39]. This attenuation implies that inflation in Morocco is less responsive to monetary policy than canonical theory suggests.

2.2 Crisis Episodes and Structural Instability

The historical record reveals three distinct crisis episodes that punctuate an otherwise tranquil baseline inflation environment (Table 1). Each episode is structurally distinct in its origin, propagation mechanism, and policy response—a heterogeneity that poses severe challenges for any model estimated on pooled time-series data.

Table 1: Moroccan Inflation Crisis Episodes: Structural Characteristics

Episode mechanism	Primary shock	Propagation	Peak CPI
2007–2009 GFC	Global food and oil price surge; capital flow reversal	Import price channel; subsidy pressure	3.9% (2008)
2020 COVID-19	Global demand collapse; supply chain disruption	Deflationary demand shock, then cost-push recovery	−0.6% (2020)
2022–2023 Commodity surge	Russia–Ukraine war; European energy crisis; domestic drought	Food and energy import prices; ERPT acceleration	8.3% (2023:Q1)

Sources: Bank Al-Maghrib Annual Reports [5]; Haut Commissariat au Plan [20]; IMF Article IV Consultations [24, 25].

This crisis anatomy has a direct methodological implication: Moroccan inflation exhibits *structural heterogeneity across regimes* that cannot be represented by a single-regime linear model estimated over the full sample. As Perron [32] demonstrates in a general framework, ignoring structural breaks in time-series estimation inflates unit-root bias, distorts impulse-response inference, and generates systematically poor out-of-sample forecasts precisely during the episodes that matter most to policymakers.

2.3 Why Moroccan Inflation Demands Flexible Methods

The phenomenological reading of Moroccan inflation thus generates a clear methodological prior: useful models must accommodate (i) nonlinear and threshold dynamics in exchange-rate pass-through, (ii) regime-switching determinant hierarchies (monetary versus external drivers), and (iii) parameter instability across crisis and non-crisis environments. This prior motivates the comparison between VAR—which imposes linearity and parameter constancy but provides structural identification—and machine learning methods—which accommodate nonlinearity and high-dimensional interactions but sacrifice structural interpretability.

3 Literature Review

3.1 The VAR Tradition in Macroeconomics

The Vector Autoregression, introduced by Sims [34] as a "theory-free" alternative to large-scale simultaneous-equation models, has since become the foundational tool of empirical monetary economics. Its enduring dominance reflects not only its forecasting utility but, more decisively, its capacity for structural identification: through the imposition of theoretically motivated restrictions on the contemporaneous impact matrix, Structural VAR (SVAR) generates orthogonal structural shocks and traces their propagation via impulse-response functions (IRFs) [26].

Applied to Morocco, SVAR specifications have provided consistent evidence that monetary transmission is sluggish and that external shocks account for a disproportionate share of inflation forecast-error variance [6, 15]. The Factor-Augmented VAR (FAVAR) of Bernanke et al. [7] partially addresses the dimensionality problem inherent in rich data environments, compressing information from a large predictor set into a small number of latent factors while preserving VAR-based identification. Yet even FAVAR inherits the linearity and parameter-constancy assumptions of its parent framework—assumptions that our phenomenological analysis suggests are particularly restrictive in the Moroccan context.

3.2 Machine Learning Enters the Macro Forecasting Arena

The entry of machine learning into macroeconomic forecasting was gradual, cautious, and initially sceptical. The early consensus, articulated by Stock & Watson [35], was that simple models tend to forecast macroeconomic variables as well as complex ones. The pivot came with the recognition—catalysed by large-scale forecasting competitions [29] and the availability of rich macroeconomic databases [19]—that the superior performance of simple models was itself regime-contingent: simple models win during stable periods; complex models win during turbulent ones.

Regularised regression. The LASSO [37] and Ridge penalised regression estimators address the over-parameterisation problem that afflicts unrestricted VARs in high-dimensional settings. Bai & Ng [4] established the theoretical foundations for

LASSO in the time-series context, while Kocenda & Macek [27] demonstrated empirically that LASSO-VAR outperforms both unrestricted VAR and pure LASSO for inflation forecasting in Central and Eastern European economies with structural characteristics broadly comparable to Morocco.

Ensemble tree methods. Random Forest [8] and Gradient Boosted Trees (XGBoost; Chen & Guestrin 11) achieve high predictive accuracy by aggregating large numbers of decision trees, implicitly implementing a flexible, nonparametric approximation to the conditional mean of the inflation process. Medeiros et al. [30] provided the benchmark evidence for US inflation, demonstrating substantial RMSE reductions during the most volatile sub-samples. Importantly, the ML advantage in Medeiros et al. [30] is asymmetric: the gain is concentrated in high-uncertainty regimes, precisely the periods that matter most for central bank operations.

Deep learning. Long Short-Term Memory (LSTM) networks [23] exploit gated recurrent architectures to capture long-range temporal dependencies in inflation series, making them theoretically attractive for processes with high persistence. Hidalgo-Pacheco et al. [22] documented LSTM outperformance over VAR benchmarks in Latin American economies—an encouraging result for Morocco, which shares several structural characteristics with that country group, including commodity dependence and managed exchange-rate arrangements.

3.3 The Interpretability Problem

The most fundamental objection to deploying ML for central bank inflation analysis is not forecasting performance but interpretability. Mullainathan & Spiess [31] formalise the distinction between *prediction problems* and *causal inference problems*: ML methods are well-suited to the former but poorly suited to the latter, because their flexible functional forms capture correlational signal without distinguishing causal from confounded relationships.

Explainability tools such as SHAP (Shapley Additive Explanations; Lundberg & Lee 28) and LIME [33] have partially addressed this limitation by providing local approximations to model predictions that satisfy axiomatic fairness properties. Yet as Athey & Imbens [2] emphasises, SHAP feature importance quantifies *predictive* relevance, not *causal* relevance. A variable may rank highly in SHAP because it is correlated with the true causal driver, not because it independently causes inflation— a distinction that is operationally critical for policy evaluation.

3.4 The Hybrid Synthesis

The most productive current in the literature treats ML and VAR not as competitors but as complements with distinct comparative advantages. Goulet Coulombe [18] develops a "macroeconomic random forest" that imposes economic structure—lag restrictions, monotonicity constraints—on the ensemble learning process, achieving both forecasting gains and enhanced interpretability. Giannone et al. [17] propose a Bayesian model averaging framework that formally combines ML and structural model forecasts, weighting each according to its recent predictive accuracy. Major central banks—the ECB [16], the Bank of England [10], and the Federal Reserve [13]— have all integrated ML components into their forecasting suites while explicitly preserving structural models for policy evaluation.

For Morocco, this hybrid logic is even more compelling: an economy in institutional transition cannot afford to sacrifice the structural transparency that VAR provides, even if it can meaningfully exploit ML's regime-adaptive forecasting flexibility.

3.5 Morocco in the Literature

The Morocco-specific empirical literature is rich by regional standards but limited in methodological scope. Existing work—El Aynaoui et al. [15], Benlghazi et al. [6], Zouhar & Motaouakkil [39], Abbad et al. [1]—establishes the empirical stylised facts that motivate this study: sluggish monetary transmission, significant ERPT, and the dominant role of imported commodity prices in driving inflation volatility. However, all existing studies employ single-regime linear specifications, and to our knowledge no published work has applied ML methods to Moroccan inflation or formally compared ML to VAR benchmarks in this context. This paper fills both gaps.

4 Analytical Frameworks

4.1 The VAR Family: Identification as Comparative Advantage

We consider three VAR-family specifications, each addressing a distinct limitation of the benchmark reduced-form model.

Reduced-form VAR. The benchmark specification models the joint dynamics of a vector $\mathbf{X}_t$ comprising CPI inflation, the BAM policy rate, the MAD/EUR exchange rate change, real GDP growth, broad money (M2) growth, and the oil price change:

$$\mathbf{X}_t = \mathbf{c} + \sum_{j=1}^{p} \mathbf{A}_j \mathbf{X}_{t-j} + \boldsymbol{\varepsilon}_t, \qquad \boldsymbol{\varepsilon}_t \sim \mathrm{N}(\mathbf{0}, \boldsymbol{\Sigma}). \tag{1}$$

The lag order p is selected via the Akaike Information Criterion. Equation by equation OLS provides efficient estimates under the Gauss-Markov conditions.

Structural VAR (SVAR). Impulse-response analysis requires a Structural VAR that recovers orthogonal structural shocks through identifying restrictions. We adopt the recursive Cholesky identification of El Aynaoui et al. [15], placing commodity prices and the exchange rate first in the ordering (treated as contemporaneously exogenous to domestic policy), with the monetary policy rate ordered last. The resulting impulse responses trace the dynamic path of inflation following oil price shocks, exchange-rate depreciations, and monetary policy tightenings— information that has direct operational relevance for Bank Al-Maghrib.

Factor-Augmented VAR (FAVAR). The FAVAR of Bernanke et al. [7] augments the core VAR with latent factors extracted by principal component analysis from a broader information set of 48 Moroccan and international quarterly indicators. Retaining three factors—as indicated by the Bai & Ng [3] information criterion—substantially reduces omitted-variable bias while preserving the low-dimensional structure of the VAR system.

4.2 Machine Learning Specifications: Prediction as Comparative Advantage

We consider four ML specifications chosen to span the main families in the literature while remaining estimable on Morocco's relatively short quarterly dataset.

LASSO-VAR and Ridge-VAR. Regularised VAR specifications embed penalised regression into the inflation equation, imposing ℓ_1 (LASSO) or ℓ_2 (Ridge) penalty terms that shrink irrelevant coefficients toward zero. These hybrid specifications occupy the methodological bridge between VAR and ML: they retain the linear structure of the VAR while exploiting the dimensionality-reduction advantages of regularisation [27].

Random Forest and XGBoost. Ensemble tree methods are specified as non-parametric regressions of the h-step-ahead inflation rate on a lagged predictor matrix comprising all variables in $\mathbf{X}_t$ and a set of global external variables. These specifications capture nonlinearities and interactions that the linear VAR cannot represent, at the cost of structural interpretability [8, 11].

LSTM network. The Long Short-Term Memory specification treats inflation forecasting as a sequence-to-one prediction problem, processing the history of $\mathbf{X}_t$ through gated recurrent cells that learn to selectively retain and forget historical information [23]. Given the short Moroccan sample, the LSTM is regularised with dropout and trained with early stopping to prevent overfitting.

4.3 A Note on the Complementarity of the Two Families

The juxtaposition of VAR and ML methods in this study is not symmetric: the two families are optimised for different objective functions. VAR is designed for *structural inference* and incidentally produces forecasts; ML is designed for *predictive accuracy* and incidentally produces interpretable outputs (via post-hoc explainability tools). Comparing them solely on forecasting metrics— as much of the literature does— therefore provides an incomplete and potentially misleading picture of their relative utility for policy institutions. Our evaluation accordingly combines statistical forecast comparison with a qualitative assessment of structural interpretability.

5 Empirical Strategy and Key Findings

5.1 Data and Evaluation Design

We assemble a quarterly dataset spanning 1990:Q1–2024:Q4 from Bank Al-Maghrib, the Haut Commissariat au Plan, the World Bank Pink Sheet database, the FAO Food Price Index, and Eurostat. The variable set mirrors the theoretical determinants identified in Section 2: domestic monetary conditions, exchange-rate dynamics, imported commodity prices, the output gap, and the Euro Area inflation environment.

All models are evaluated under a standardised rolling-window protocol: an initial training window (1990:Q1–2009:Q4) is expanded quarter by quarter, with each successive step re-estimating all models and generating genuine out-of-sample forecasts for h = 1, 2, 4 quarters. Statistical significance of forecast differentials is assessed using the Diebold–Mariano test [14] with the Harvey–Leybourne–Newbold small-sample correction.

5.2 The Regime-Contingent Pattern of ML Advantage

The central empirical finding can be stated in a single proposition: machine learning methods do not uniformly outperform VAR; their advantage is regime-contingent, concentrating in crisis episodes where structural instability and nonlinear dynamics overwhelm linear model capacity.

Figure 1 summarises this pattern schematically.

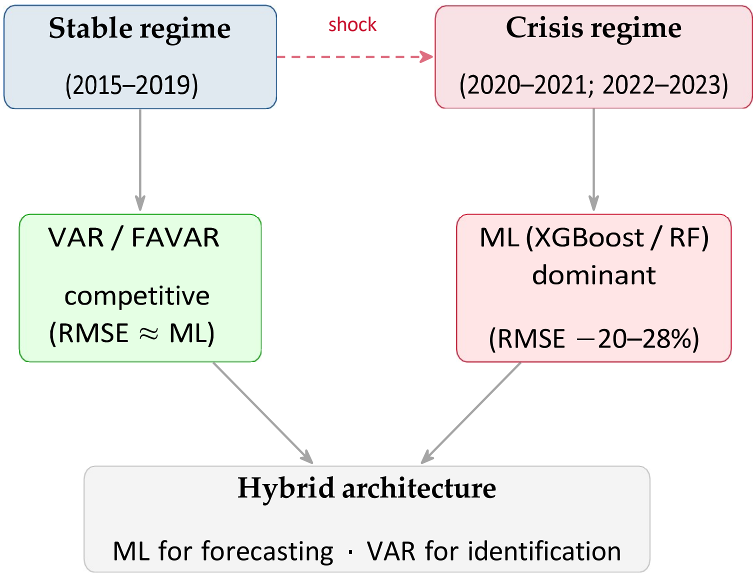

Figure 1: Regime-contingent methodological performance: schematic summary of empirical findings.

During stable macroeconomic periods, VAR-family and ML models perform comparably. During crisis episodes, ML ensemble methods achieve significant predictive gains. Neither family dominates across all objectives, motivating a hybrid architecture.

More specifically, across the full out-of-sample evaluation period (2015:Q1–2024:Q4):

(i) At the one-quarter horizon, XGBoost achieves an RMSE reduction of approximately 17–18% relative to the unrestricted VAR(2) benchmark, with the differential statistically significant at the 5% level (Diebold–Mariano test). FAVAR narrows this gap substantially, reflecting its information-richness relative to small-scale VARs.

(ii) The ML advantage is sharply asymmetric across sub-periods. During non-crisis quarters, the Diebold–Mariano test *fails* to reject equal predictive accuracy for Random Forest versus FAVAR ($p \approx 0.17$), suggesting that flexible ML methods offer no systematic gain over factor-augmented econometrics in structurally stable environments. During crisis quarters, by contrast, Random Forest achieves RMSE reductions exceeding 27% relative to VAR(2), with the differential significant at the 5% level.

(iii) At the four-quarter horizon, the LSTM network outperforms all other specifications—consistent with the hypothesis that its temporal memory architecture confers advantages at extended horizons where inflation persistence dominates forecast accuracy.

5.3 What the VAR Reveals That ML Cannot

The superiority of ML in crisis forecasting coexists with a capability gap that no forecasting metric can capture: structural identification. The SVAR impulse-response analysis reveals three findings of direct operational significance for Bank Al-Maghrib:

- A one-standard-deviation positive shock to the BAM policy rate reduces inflation by approximately 0.4 percentage points after three to four quarters—a transmission magnitude roughly one-third the size of an equivalent food price shock. This asymmetry confirms that Morocco's inflation problem is primarily supply-side and import-driven, rendering monetary policy a blunt instrument for inflation stabilisation.

- The Forecast Error Variance Decomposition (FEVD) reveals a dramatic shift in the inflation determinant hierarchy during the 2022–2023 crisis: import price and exchange-rate shocks jointly account for approximately 58% of one-year-ahead inflation forecast variance during the commodity surge, compared to roughly 21% in the pre-crisis baseline. This structural shift is entirely invisible to a single-regime linear model estimated over the full sample.

- Chow breakpoint tests detect significant parameter instability in the VAR system at all three crisis dates at the 1% level, confirming that the structural assumption of time-invariant parameters is violated in Moroccan inflation data.

Crucially, these three findings together constitute the structural case *for* ML flexibility— but they also constitute the institutional case *for* retaining VAR alongside ML, because only the SVAR can generate the conditional counterfactuals

("What happens to inflation if we raise the policy rate by 50bp given the current external shock?") that monetary policy committees require [21].

6 Discussion

6.1 Can ML Replace VAR?

The paper's central finding is negative in its literal sense but constructive in its implications: machine learning cannot replace VAR in the Moroccan monetary policy context, for reasons that are structural and institutional rather than merely statistical.

The argument has two components. First, the *structural identification argument*: VAR-based IRF analysis generates counterfactual causal pathways—the dynamic response of inflation to an isolated monetary policy shock—that ML models cannot produce without reimposing structural assumptions that effectively recreate the econometric framework [12, 21]. SHAP feature importance, however elegant, does not substitute for this: it quantifies predictive correlation, not causal transmission.

Second, the *institutional communication argument*: inflation targeting, as a governance framework, requires a central bank not only to forecast inflation accurately but to commit publicly to a specific projection, explain deviations, and articulate the policy logic connecting instrument settings to outcomes [36]. A gradient-boosted ensemble model, however accurate in ex-ante backtests, cannot straightforwardly discharge this communicative function. As Varian [38] observes, data-driven prediction and institutional accountability are not the same kind of activity.

6.2 The Specific Lessons for Morocco

The Moroccan experience yields three policy-relevant lessons that transcend the methodological debate.

Lesson 1: **Supply-side inflation demands supply-side intelligence**. The SVAR and ML analyses agree that Moroccan inflation is primarily driven by external supply shocks—oil prices, food prices, Euro Area inflation, and exchange-rate dynamics—during the crisis episodes that impose the greatest costs on households. Improving the tracking of global commodity markets, shipping cost indices, and agricultural production data is therefore as valuable for inflation forecasting as any methodological refinement.

Lesson 2: **ML should be crisis-aware**, **not crisis-activated**. Our findings show that ML advantage concentrates in crisis periods—but by definition, policymakers only discover they are in a crisis after its onset. A practical implementation should therefore use real-time regime indicators (commodity price volatility indices, exchange-rate stress measures, global uncertainty proxies) to activate ML model weighting *before* crisis-period forecasting errors accumulate.

Lesson 3: **Institutional transition demands methodological pluralism**. Bank Al-Maghrib's ongoing transition toward inflation targeting creates a period of heightened uncertainty about the structural parameters of monetary transmission. Methodological monoculture—relying exclusively on either VAR or ML—concentrates model risk during precisely the period when analytical flexibility is most valuable. A hybrid architecture, deploying each method for the tasks it performs best, is the appropriate institutional response.

6.3 Toward a Hybrid Architecture for Bank Al-Maghrib

We propose a three-tier analytical architecture for Moroccan inflation forecasting, summarised in Table 2.

Table 2: Proposed Hybrid Analytical Architecture for Bank Al-Maghrib

Tier	Primary tool	Purpose

Note: All three tiers operate in parallel. Forecast disagreement across tiers serves as a real-time model uncertainty signal. The structural tier is irreplaceable for monetary policy communication and committee deliberation.

This architecture is consistent with the practice of leading central banks [10, 16] and draws on the theoretical synthesis proposed by Goulet Coulombe [18] and Giannone et al. [17]. It acknowledges that the ML-versus-VAR debate, framed as an either-or choice, poses a false dilemma: the relevant question is not which method wins but how their complementary strengths can be integrated into a coherent institutional forecasting process.

Conclusion

This paper began with an observation that inflation is a phenomenon, a complex, historically contingent process shaped by institutions, expectations, and structural vulnerabilities that resist compression into any single econometric or algorithmic specification. The Moroccan case makes !!this observation especially transparent: three decades of price data reveal a low-inflation baseline punctuated by sharp, structurally distinct crisis episodes whose dynamics systematically overwhelm the linear, time-invariant assumptions that underpin conventional VAR estimation.

Against this backdrop, machine learning methods offer genuine value: tree-based ensembles and LSTM networks adapt to regime shifts faster than linear models, generate significant forecast improvements during the crisis episodes that matter most, and identify determinant hierarchies—the shift from monetary to external drivers—that conventional FEVD analysis subsequently validates. But ML's forecasting advantages coexist with structural limitations that are decisive for monetary policy practice: an inability to generate structural identification, a lack of conditional counterfactual capacity, and a communicative opacity that is incompatible with the transparency requirements of inflation targeting.

We therefore answer our title question with a principled negative qualified by a constructive proposal. Machine learning cannot replace VAR in the Moroccan inflation context—but it should be integrated alongside VAR in a hybrid analytical architecture that allocates each method to the tasks for which it is best suited. For Bank Al-Maghrib, navigating an institutional transition of historical significance, this methodological pluralism is not a counsel of indecision but a recognition that the phenomenon demands more intellectual flexibility than any single analytical tradition can provide.

Future research should extend this framework in two directions. First, the application of Bayesian model averaging to combine ML and VAR forecasts [17] would provide a principled, probabilistically coherent synthesis of the two families. Second, the

Moroccan experience—a small, import-dependent economy with a managed exchange rate navigating an inflation-targeting transition—may generalise to a broad set of MENA and Sub-Saharan African economies. Systematic comparative analysis across this group represents an important and largely unexplored research agenda.

References

- Abbad, M., Bensassi, S., & Taamouti, A. (2020). Exchange rate pass-through to domestic prices in Morocco: Evidence from a structural VAR model. *Journal of African Economies*, 29(4), 311–338.
- Athey, S., & Imbens, G. W. (2019). Machine learning methods that economists should know about. *Annual Review of Economics*, 11, 685–725.
- Bai, J., & Ng, S. (2002). Determining the number of factors in approximate factor models. *Econometrica*, 70(1), 191–221.
- Bai, J., & Ng, S. (2008). Forecasting economic time series using targeted predictors. *Journal of Econometrics*, 146(2), 304–317.
- Bank Al-Maghrib. (2023). *Annual Report 2022–2023: Monetary Policy, Inflation and Economic Developments*. Rabat: Bank Al-Maghrib Publications.
- Benlghazi, R., Chaara, M., & El Ghini, A. (2015). Long-run determinants of inflation in Morocco: An ARDL bounds testing approach. *International Journal of Economics and Financial Issues*, 5(3), 807–817.
- Bernanke, B. S., Boivin, J., & Eliasz, P. (2005). Measuring the effects of mon-etary policy: A factor-augmented vector autoregressive (FAVAR) approach. *Quarterly Journal of Economics*, 120(1), 387–422.
- Breiman, L. (2001). Random forests. *Machine Learning*, 45(1), 5–32.
- Bresciani-Turroni, C. (1931). *The Economics of Inflation: A Study of Currency Depreciation in Post-War Germany*. London: George Allen & Unwin. [English translation, 1937]
- Chakraborty, C., & Joseph, A. (2022). Machine learning at central banks: Applications and interpretability. *Bank of England Working Paper*, No. 674.
- Chen, T., & Guestrin, C. (2016). XGBoost: A scalable tree boosting system. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (pp. 785–794). ACM.
- Chen, T., Guestrin, C., & He, X. (2020). Interpretability and causality in machine *learning for economics. AEA Papers and Proceedings, 110, 432–437.*

- Coulombe, P. G., Leroux, M., Stevanovic, D., & Surprenant, S. (2020). How is machine learning useful for macroeconomic forecasting? CIRANO Working Papers, No. 2020s-45.
- Diebold, F. X., & Mariano, R. S. (1995). Comparing predictive accuracy.
- Journal of Business & Economic Statistics, 13(3), 253–263.
- El Aynaoui, K., Quarouch, H., & Lahlou, M. (2008). Monetary transmission mechanism in Morocco: A structural VAR approach. Bank Al-Maghrib Working Paper Series, No. 23.
- European Central Bank. (2020). Machine learning applications in central bank-ing (Occasional Paper Series No. 25/2020). ECB.
- Giannone, D., Henry, J., Lalik, M., & Modugno, M. (2021). An area-wide real-time database for the euro area. Review of Economics and Statistics, 103(5), 799–812.
- Goulet Coulombe, P. (2022). The macroeconomy as a random forest. Journal of Applied Econometrics, 37(5), 1045–1068.
- Ha, J., Kose, M. A., & Ohnsorge, F. (2023). One-stop source: A global database of inflation. Journal of International Money and Finance, 137, 102896.
- Haut Commissariat au Plan. (2023). Tableau de bord de l'économie marocaine: 2022–2023. Rabat: HCP.
- Heckman, J. J. (2001). Micro data, heterogeneity, and the evaluation of public policy: Nobel lecture. Journal of Political Economy, 109(4), 673–748.
- Hidalgo-Pacheco, C., Restrepo-Tobón, D., & Arias-Castro, J. (2021). LSTM networks for inflation forecasting in Latin America. Latin American Journal of Central Banking, 2(3), 100036.
- Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. Neural Computation, 9(8), 1735–1780.
- International Monetary Fund. (2022). Morocco: 2022 Article IV Consultation Staff Report (Country Report No. 22/235). IMF.
- International Monetary Fund. (2023). Morocco: Selected Issues—Exchange Rate Flexibility and Inflation Dynamics (Country Report No. 23/302). IMF.

- *Kilian, L., & Lütkepohl, H. (2017). Structural Vector Autoregressive Analysis. Cambridge University Press.*
- *Kocenda, E., & Macek, M. (2022). Inflation forecasting in Central and Eastern Europe: A comparison of LASSO-VAR and standard VAR models. Economic Modelling, 110, 105796.*
- *Lundberg, S. M., & Lee, S.-I. (2017). A unified approach to interpreting model predictions. In I. Guyon et al. (Eds.), Advances in Neural Information Processing Systems, Vol. 30. Curran Associates.*
- *Makridakis, S., Spiliotis, E., & Assimakopoulos, V. (2022). M5 accuracy competition: Results, findings, and conclusions. International Journal of Forecasting, 38(4), 1346–1364.*
- *Medeiros, M. C., Vasconcelos, G. F. R., Veiga, Á., & 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.*
- *Mullainathan, S., & Spiess, J. (2017). Machine learning: An applied econo-metric approach. Journal of Economic Perspectives, 31(2), 87–106.*
- *Perron, P. (2006). Dealing with structural breaks. In T. C. Mills & K. Patter-son (Eds.), Palgrave Handbook of Econometrics, Vol. 1 (pp. 278–352). Palgrave Macmillan.*
- *Ribeiro, M. T., Singh, S., & Guestrin, C. (2016). "Why should I trust you?": Explaining the predictions of any classifier. In Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (pp. 1135–1144). ACM.*
- *Sims, C. A. (1980). Macroeconomics and reality. Econometrica, 48(1), 1–48.*
- *Stock, J. H., & Watson, M. W. (2002a). Macroeconomic forecasting using diffusion indexes. Journal of Business & Economic Statistics, 20(2), 147–162.*
- *Svensson, L. E. O. (2010). Inflation targeting. In B. M. Friedman & M. Wood-ford (Eds.), Handbook of Monetary Economics, Vol. 3B (pp. 1237–1302). Else-*

vier.

- *Tibshirani, R. (1996). Regression shrinkage and selection via the lasso.*
- *Journal of the Royal Statistical Society: Series B, 58(1), 267–288.*
- *Varian, H. R. (2014). Big data: New tricks for econometrics. Journal of Economic Perspectives, 28(2), 3–28.*
- *Zouhar, Y., & Motaouakkil, A. (2019). Monetary policy transmission in Morocco: A VAR analysis of the interest rate channel. Review of Economic Analysis, 11(1), 93–121.*