International Journal of Applied Mathematics
Volume 38 No. 12s 2025,
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

MACHINE LEARNING–ENABLED EARLY WARNING SYSTEM FOR DETECTING MICRO-INFLATION CLUSTERS IN THE U.S. ECONOMY

MD Saifur Rahman[1], Abdullah Al Mukaddim[2], MD Rashed Mohaimin[3], Kazi Sharmin Sultana[4], Md Toushif Pramanik[5], Dil Tabassum Subha[6], Kazi Nehal Hasnain[7], Sonia Afroze[8], Riad Hossain[9] and Shah Alam[10]

[1]Master's in ITAM, Central Washington University

[2]Master of Science in Business Analytics, Grand Canyon University, Phoenix, AZ, USA

[3]MBA in Business Analytics, Gannon University, Erie, PA, USA

[4]MBA in Business Analytics, Gannon University, Erie, PA, USA

[5]Master of Embedded Software Engineering, Gannon University, Erie, PA, USA

[6]Master of Science in Business Analytics, Grand Canyon University, USA

[7]Master of Science in Information Technology (MSIT), Westcliff University, Irvine, CA

[8]MSc in Botany, Jahangirnagar University

[9]Master of Science in Computer Science, Central Michigan University, Mount Pleasant, Michigan, USA

[10]Master of Science in Information Technology, Washington University of Science and Technology, Alexandria, VA, USA.

Corresponding Author: **MD Saifur Rahman, Email:** rahmanmd@cwu.edu

Abstract

Inflation in the United States has become noticeably more erratic in recent years, and this has pushed researchers, policymakers, and anyone watching the economy to look for ways to spot early shifts in pricing pressure before they show up in the main CPI numbers. This study builds a machine learning early warning system that looks for the small pockets of rising prices that often form beneath the surface. It uses high-frequency CPI and PPI data, volatility indicators, forecast residuals from several time-series models, and a set of anomaly detection tools to catch signs that inflation is starting to take shape. The system draws on familiar forecasting models like ARIMA, exponential smoothing, and Prophet to produce forward-looking residual features that flag unexpected movements in price patterns. These residuals are blended with unsupervised anomaly scores, rolling momentum signals, and supervised models to estimate the chance that localized inflation pressure is developing months in advance. The modeling pipeline is backed by a range of robustness checks that look at how sensitive the system is to different CPI definitions, alternative PPI group choices, shifting lead times, and the structural changes that took place around the COVID period. SHAP values are used to break down the predictions of the supervised models, showing which features drive the alerts most strongly. The results highlight the influence of residual anomalies, persistent PPI momentum, and volatility spikes, which consistently show up as warning signals. Backtests indicate that

Received: August 13 ,2025

combining residual features with momentum indicators improves performance relative to simple rules that many analysts still rely on, especially when predicting CPI accelerations three months ahead. The improvement is most visible after 2020, a period when producer price shocks move more directly into consumer prices and are easier for the system to detect. The results here show that mixing forecast residuals with engineered PPI and CPI features offers a practical way to spot early signs of inflation. The system remains understandable, holds up across several forms of stress testing, and can highlight brewing inflationary pressure before it is visible in the headline indexes. This provides a solid base for building early warning tools that help with economic monitoring, policy analysis, and risk management during times when price behavior becomes uncertain and harder to track.

Keywords: Inflation, Early Warning System, Machine Learning, CPI, PPI, Forecast Residuals, Anomaly Detection, SHAP, Economic Signals

1. Introduction

1.1 Motivation

Catching inflationary pressure before it shows up in the headline numbers has always been one of the toughest parts of macroeconomic analysis. Most monitoring still leans on headline CPI or core inflation, but anyone who has worked with these measures knows they can hide the messy stuff happening underneath. Stock and Watson (2016) point out that core inflation does help smooth the noise, yet it often misses the very first signs that something is shifting within individual components [28]. Those early ripples tend to start in a few narrow categories, and they stay invisible inside the aggregate until they grow large enough to move the whole index. This becomes even harder when the economy is absorbing shocks that break old relationships. Cecchetti and Moessner (2008) show that commodity-linked disturbances can ripple through the system in ways that are irregular and sometimes unpredictable [4]. The traditional tools respond slowly when faced with sudden supply shortages or sharp swings in production costs. By the time the headline picks up the pressure, most of the early signals have already passed through the pipeline.

Recent work highlights how scattered and uneven inflation has become. Shapiro (2022) shows that demand-driven and supply-driven forces can sit inside very specific categories, forming localized inflation pockets long before they influence the full basket [23]. Reis (2023) adds that today's inflation environment needs tools that can handle the structural and cross-sector complexity that shows up in disaggregated price data [21]. Put differently, inflation is no longer a smooth, economy-wide wave. It starts as clusters, each with its own timing and intensity. There are interesting parallels with other fields that rely on early anomaly detection. Das et al. (2025) describe how traditional cybersecurity approaches often fall behind fast-moving threats, while predictive AI systems tend to catch unusual behavior much earlier [8]. Debnath et al. (2025) show that bringing together different streams of information, such as energy metrics and cybersecurity logs, makes anomaly detection more reliable [9]. Something similar applies to inflation. When signals like CPI, PPI, engineered volatility, and forecast residuals are treated as separate yet connected pieces of information, the chances of spotting early disturbances rise sharply. Inflation develops through local shocks, supply chain bottlenecks, shifts in production

Received: August 13 ,2025

costs, and movements in demand that spread through the economy at different speeds. It does not move in a single line, and it does not develop in a uniform way. A system that listens to many types of signals at once and learns to pick out unexpected changes is better suited for detecting early inflation pressure. That is the core motivation behind an early warning framework built with machine learning.

1.2 Problem Statement

Small pockets of inflation often begin in places most people are not watching. They may appear first in producer prices, intermediate goods, or narrow supply chains before anything noticeable happens at the consumer level. These early movements get buried inside aggregate CPI and PPI indexes, and the larger numbers fail to capture the uneven ways that producer shocks pass downstream. The transmission can be delayed, partial, or heavily shaped by how firms decide to adjust their prices. This sets up a recurring problem. By the time inflation is visible in the headline measures, the buildup has already been underway for months. Traditional econometric tools work best in stable environments where relationships hold steady over time. The inflation shocks of recent years have looked very different. They have been volatile, fragmented across sectors, and irregular in how they move from producers to consumers. Without methods that can spot unusual behavior, identify structural changes, and read signals inside fast-moving CPI and PPI series, analysts and policymakers end up recognizing inflation risk only after it becomes widespread. This study tackles that gap by creating a predictive system that can identify early signs of inflation several months before they appear in the official statistics.

1.3 Research Importance

An early warning system for inflation matters across nearly every part of the economy. Monetary policymakers need timely information about where inflation seems to be forming, because delayed awareness often leads to delayed responses. When a central bank reacts late, the swings in policy tend to be larger and more disruptive. A system that highlights early inflation clusters supports steadier decision-making and gives policymakers more room to adjust before the situation becomes more difficult to handle. Financial institutions depend on early information as well. Asset managers, hedge funds, and fixed-income analysts build strategies around expectations of inflation and interest rates. Sudden inflation shocks change the pricing of bonds, real assets, and risk premiums. A signal that alerts them to shifting inflation pressure before it shows up in the headline data can improve how they position portfolios, especially in turbulent markets.

Corporate decision makers feel the effects too. Unexpected increases in producer prices can eat into margins, complicate contracts, and push firms to rethink inventory and sourcing. When upstream pressures are detected earlier, firms can react by managing production schedules and renegotiating agreements before the broader market forces their hand. The broader economy is affected in ways that go beyond prices. Household purchasing power shifts, wage dynamics adjust, and consumer sentiment responds quickly to inflation surprises. Fiscal and macro-prudential authorities benefit from recognizing early pressure points because doing so gives them time to plan responses and anticipate which sectors or households may be at greater risk. Modern economies are deeply connected, and shocks move quickly through these networks.

Received: August 13 ,2025

International Journal of Applied Mathematics
Volume 38 No. 12s 2025,
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)
An early warning system that can see inflation risks forming below the surface helps stabilize expectations, guide policy, and reduce the costs of unexpected inflation. This research contributes to that effort by offering a framework that can read signals hidden in disaggregated, high-frequency price data.

1.4 Objectives and Contributions

This research sets out to build a machine learning system that can detect micro-inflation clusters long before they influence headline CPI. The approach relies on a layered feature design that incorporates momentum, volatility, spreads, and forecast residual anomalies. These features are evaluated across several CPI definitions, a range of PPI categories, and different lead times. Both supervised and unsupervised methods are used so the system can assign probabilities to inflation risk while also identifying unexpected behavior through anomaly detection. A key part of the work involves creating a rolling-window forecasting pipeline with ARIMA, exponential smoothing, and Prophet. These models generate forward residuals that act as early signals of abnormal price movements. The study also introduces a unified feature bank that captures temporal changes and volatility clusters. Each component is tested for sensitivity and exposed to structural shifts to assess stability. Model interpretability is maintained through SHAP analysis, allowing the system to explain why specific alerts occur.

2. Literature Review

2.1 Traditional Inflation Monitoring

For a long time, most attempts to track inflation have leaned heavily on econometric tools that try to tease out the deeper forces behind price changes. These models usually split inflation into pieces that are thought to be persistent or temporary, rely on Phillips-curve relationships, and often bring in survey measures of expectations. Work by Mavroeidis, Plagborg-Møller, and Stock (2014) shows how fragile these structures can be, especially when expectation formation shifts or when the model depends too much on narrow assumptions [18]. Their findings make it clear that even widely used approaches struggle to keep up with real-time inflation behavior. A big part of this difficulty comes from noisy data, unstable relationships between slack and prices, and global supply issues that rewrite the usual playbook. These weaknesses show up again in research that tries to separate core inflation from headline movements. Stock and Watson (2016) point out that removing volatile components helps up to a point, yet this technique still fails to anticipate many turning points, particularly when pressure originates outside the standard consumption basket or builds through complicated production chains [28]. These missed signals become more common in moments where the economy faces structural breaks or unusual shocks, when past relationships stop behaving the way textbooks suggest.

Traditional forecasting also runs into trouble as data sets grow larger and more complex. Bańbura and van Vlodrop (2023) compare machine learning models with older econometric methods and find clear evidence that the latter slip behind when the data environment becomes too rich or too nonlinear for simple structures to handle [3]. Their work highlights a broader issue that inflation today reflects a patchwork of sector-specific forces, supply-chain imbalances, and interactions across markets. This fits with earlier findings from Cecchetti and

Received: August 13 ,2025

International Journal of Applied Mathematics
Volume 38 No. 12s 2025,
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)
Moessner (2008), who show that swings in global commodity markets reshape inflation pass-through and complicate efforts to separate temporary disruptions from genuine inflation pressures [4]. Shapiro (2022) pushes this point further by showing that many inflation cycles grow out of a narrow set of categories rather than widespread price increases, which means meaningful signals often appear at a micro level long before they show up in aggregate data [23]. Reis (2023) adds that headline CPI is too coarse a tool for understanding how modern pricing systems work, and argues for a more granular approach that pays attention to how inflation behaves across categories rather than only in the aggregate [21].

2.2 Machine Learning Approaches to Inflation and High-Dimensional Financial Prediction

Machine learning has expanded quickly into macroeconomic forecasting, financial modeling, and high-volatility asset prediction. These methods have created a shift in how inflation is studied, especially when the goal is to learn from large, noisy, and fast-moving data. Medeiros et al. (2022) show how ML models outperform more traditional frameworks in environments where many variables interact in nonlinear ways, and high-frequency series carry information that older methods tend to overlook [19]. Their work sits alongside Coulombe et al. (2023), who demonstrate how the integration of higher-frequency data improves nowcasts for inflation, particularly in periods where conditions change quickly, and new information must be absorbed at a rapid pace [7]. Together, these studies illustrate how forecasting has moved toward approaches that adapt to shifting data patterns rather than relying on rigid equations. A similar shift appears in financial machine learning. Islam et al. (2025) study cryptocurrency price prediction and show how ML models hold up under extreme volatility and noise, a setting where conventional econometric forecasting tends to fail [14]. Ray (2025) examines crisis prediction across stock, bond, and currency markets and finds that ML-based classifiers pick up early destabilizing patterns in ways traditional regime-switching or VAR frameworks struggle to match [20]. These ideas are directly relevant for building inflation early warning systems because inflation regimes often follow subtle changes in momentum, volatility, and cross-market interactions. Reza et al. (2025) add another layer by demonstrating how socioeconomic prediction tasks benefit from blending diverse signals into a single modeling structure, mirroring the challenge of assembling CPI, PPI, spreads, volatility, and residual-based indicators to detect early inflation pressure [22].

Deep learning research opens additional possibilities. Devlin et al. (2019) introduce the transformer architecture and show how models built on it can learn long-range dependencies in sequential data [10]. While these models were created for language, their structure has since proven useful for financial and macroeconomic time series, where relationships often stretch far across time and depend on subtle combinations of signals. This creates a natural extension for future inflation modeling, especially when paired with textual sources such as policy communications or supply-chain commentary. Across macroeconomic forecasting, financial modeling, and sequence learning research, a clear message emerges. Machine learning provides the flexibility, noise tolerance, and nonlinear modeling capability needed to pick up the early, localized inflation pressures that older frameworks tend to overlook. This growing body of work builds the foundation for the early warning system proposed in this study, which

Received: August 13 ,2025

ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)
draws from both supervised and unsupervised approaches to identify inflation signals that form beneath the surface of headline measures.

2.3 Early Warning Systems in Macroeconomics and Financial Risk

Early Warning Systems have played a central role in macroeconomic research for decades, starting with efforts to anticipate currency crashes, banking failures, and broader financial instability. The work of Kaminsky, Lizondo, and Reinhart (1998) remains one of the most influential contributions in this space. Their research shows how indicators such as foreign reserves, credit expansion, and pressure on exchange rates tend to shift before a crisis unfolds, offering a foundation for thinking about predictive signals in advance of major disruptions [16]. Their ideas continue to shape how researchers approach early detection, and the logic behind their framework carries over naturally to inflation EWS development, where the real challenge lies in sensing turning points while they are still forming beneath the usual aggregates. Jordà, Singh, and Taylor (2020) add another important layer by studying the long-run effects of monetary policy. They highlight how inflationary impulses accumulate over time and how policy outcomes depend on access to timely and reliable signals [15]. Their work shows that inflation does not simply react to policy choices in linear ways. Instead, shocks move through the system in uneven steps that require close tracking long before they shape expectations or wage negotiations. This reinforces the argument for more adaptive and fine-grained monitoring tools.

More recent studies point toward the growing importance of AI-driven systems in risk detection. Chouksey et al. (2025) propose a machine learning EWS for financial risk in the United States digital economy and show how diverse inputs and ensemble learning can uncover unusual patterns before they evolve into full market disturbances [5]. Their results resemble the challenges of inflation forecasting, where producer prices, spreads, and volatility often interact in complicated ways that are hard to track with traditional models. Sizan et al. (2025) demonstrate a similar idea in the context of financial crime detection. They use an unsupervised ensemble approach to uncover emerging money-laundering behaviors that shift too quickly for fixed-rule methods to recognize [27]. Their insights help explain why micro-inflation clusters often escape notice, since both problems involve evolving patterns that do not fit neatly into known categories. Other recent work broadens the landscape of EWS research across domains. Das et al. (2025) show how predictive systems in cybersecurity outperform manual detection when threats move faster than human monitoring can adapt [8]. Debnath et al. (2025) examine how combining different types of signals improves anomaly detection in complex energy environments, a finding that translates neatly to inflation research, which must balance signals coming from CPI patterns, PPI movements, residual shocks, and volatility surges [9]. A number of studies reinforce this theme. Hasan et al. (2025) and Shawon et al. (2025) both show how ML-based monitoring systems give earlier and more detailed insight into supply-chain stress and operational vulnerabilities than traditional tools [11][24]. The underlying idea is the same: combining diverse signals reveals patterns that single-source approaches miss.

2.4 Producer–Consumer Price Dynamics, Anomaly Detection, and Explainability

Received: August 13 ,2025

A core part of understanding inflation involves the way producer prices flow into consumer prices. Clark (1995) provides early evidence that producer price movements often appear before changes in consumer prices, especially in sectors where production costs translate directly into retail pricing [6]. This relationship forms an important base for the present study, which treats PPI signals as an early window into how consumer prices may shift months later. Hobijn and Lagakos (2005) build on this idea by showing that households experience inflation in uneven ways, driven by category-specific shocks that contribute to what they describe as inflation inequality [13]. Their findings help explain why micro-inflation clusters matter so much. A small number of categories can generate meaningful inflationary pressure long before the aggregate CPI responds. Recent research on anomaly detection adds important tools for uncovering these early movements. Shivogo (2025) examines changing environments in credit-scoring and proposes an adaptive explainability structure that accounts for concept drift, where the underlying forces driving outcomes shift over time [25]. This idea is central to inflation forecasting, since supply chains break, pricing strategies change, and macro shocks appear at unpredictable moments. Hasan et al. (2025) highlight how explainable AI improves decision-making in sparse environments where transparency is needed to understand why a prediction is being made [12]. These ideas align closely with inflation monitoring, where policymakers need clear explanations of which signals matter and why.

Aashish et al. (2025) extend this line of thinking by designing anomaly detection tools that incorporate carbon and energy considerations into cybersecurity monitoring [1]. Their work demonstrates how models can work within multidimensional feature spaces while still meeting stability and resource requirements, a practical concern for inflation EWS designs that must handle varied inputs without sacrificing reliability. This discussion fits naturally with broader arguments about ML in economics. Athey (2018) emphasizes the importance of interpretability when connecting machine learning to economic policy, noting that accurate predictions are not enough unless the forces behind those predictions are clearly identified [2]. Inflation monitoring depends heavily on this balance. Analysts must understand which sectors are contributing to early price pressures, how producer prices are feeding into consumer markets, and which combinations of signals produce a warning. The SHAP-based explainability used in this study responds directly to that need.

2.5 Literature Gaps and Framework Relevance

Even with the growing use of machine learning in macroeconomic research, several important gaps still stand out. Many ML-based inflation studies look almost entirely at forecasting aggregate CPI. What they often miss is the early activity happening inside specific sectors where inflation usually begins. These early developments tend to appear in narrow supply-chain segments, long before they show up in aggregate measures, which means a large share of the early signal is left unexplored. Another issue is the way existing studies treat model components in isolation. Very few combine forecasting residuals, anomaly scores, and supervised methods into a single early warning structure that can track shifts across multiple dimensions at once. Work in financial crisis prediction, socioeconomic modeling, and digital-economy risk systems shows the value of blending these approaches, but these ideas have not yet been applied fully to inflation monitoring. A related gap concerns transparency. Much of

Received: August 13 ,2025

the macro-ML literature delivers strong predictive performance, yet falls short when it comes to explaining why a model arrives at a particular output. Athey (2018) points out that ML can only make a meaningful contribution to economics if interpretability is built into the design of the model from the start, not added as an afterthought [2]. This matters for inflation monitoring, since policymakers need clear evidence about the drivers behind prediction shifts. Without that clarity, even accurate systems remain difficult to use.

Robustness is another weak spot. Traditional models often break down during structural changes, and ML methods, although more flexible, can drift or overfit when inflation regimes change. Few studies push their inflation models through stress scenarios such as large supply-chain disruptions or abrupt regime transitions. Research on adaptive explainability and unsupervised anomaly detection hints at ways to address these challenges, but has not yet been integrated into a systematic inflation-focused framework. There is also a recurring pattern in which PPI data is included in models but used only as a basic input, rather than as a richer source of information. Its divergence patterns, volatility surges, and forecast residuals often contain early signals, yet these signals remain underutilized. The framework developed in this study targets this specific gap by combining producer-to-consumer price propagation, anomaly surfaces derived from multiple forecasting models, and supervised classification within a single interpretable structure.

3. Methodology

3.1 Data Sources

The analysis is based on monthly data from the U.S. Bureau of Labor Statistics Public Data API. Two series form the core structure of the study. The first is the Consumer Price Index All Items series (CUUR0000SA0), which captures broad consumer-level inflation. The second is the Producer Price Index for Packaging Materials (WPU0911), which serves as a proxy for upstream cost pressures that often shift before consumer prices do. The time span begins in January 2000 and extends through the most recently available month, providing the study with exposure to multiple inflation cycles, recessions, expansions, and structural shifts. Once the data is downloaded, it is converted into a consistent monthly index. Missing months are filled through interpolation when appropriate, and any remaining gaps are removed to keep the dataset continuous. This preparation makes the series suitable for rolling models, percent-change calculations, seasonal decomposition, and the forecasting steps used later in the pipeline.

3.2 Target Variable Construction

The system focuses on identifying micro-inflation clusters, which are short periods when CPI inflation picks up more rapidly than usual. To convert this idea into a usable target, the analysis calculates the future three-month percentage change in CPI. Whenever this forward change exceeds the eightieth percentile of its historical values, the period is marked as an inflation event. These events form an indicator known as cpi_event_raw. To shape the early warning objective, the event indicator is shifted backward by a fixed number of months, written as lead_months. A lead of three months is typically used, which allows the model to flag potential inflation events before they appear in CPI data. The final binary label reflects future inflation

Received: August 13 ,2025

International Journal of Applied Mathematics
Volume 38 No. 12s 2025,
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)
rather than current conditions, which is essential for evaluating a system that claims to provide a warning rather than simply track existing trends.

3.3 Feature Engineering

A large feature set is constructed to represent price dynamics from several angles. The features cover rate-of-change measures, volatility patterns, producer-consumer price relationships, seasonal structure, standardized deviations, and residual-based forecasts. Each category contributes information about inflation behavior that would not be visible if the study relied on raw CPI and PPI data alone. Rate-of-change features capture inflation momentum across different horizons. These include one-month, three-month, and six-month percentage changes and log returns for both CPI and PPI. They help capture the speed at which prices shift. Volatility features measure the instability of the series. Rolling standard deviations are calculated over three, six, and twelve-month windows. These features highlight periods where uncertainty increases or where price movements become erratic. PPI to CPI spread features quantify how far producer prices pull away from consumer prices. They include absolute spreads, percent-change spreads, and rolling lag correlations between the two series. These indicators reflect the idea that producers may feel inflation pressure first, which then works its way through supply chains and eventually affects consumers.

Seasonal adjustments are added by applying STL decomposition with a twelve-month seasonal component. This provides seasonally adjusted CPI and PPI versions that help reduce noise from predictable yearly patterns and make it easier to identify movements that break from normal seasonal behavior. Standardization uses rolling twelve-month z-scores for both the raw series and the engineered indicators. This step expresses each value relative to its recent history and makes it easier to detect unusual movements. Forecast residual features provide some of the strongest forward-looking signals in the system. Rolling windows are used to fit ARIMA(1,1,1), Exponential Smoothing, and Prophet models to both CPI and PPI, producing forecasts several months ahead. The residuals, defined as the difference between the actual value and the forecast, offer a look at situations in which price behavior deviates from what history would typically predict. Large positive residuals often represent unexpected bursts of inflationary pressure. These residuals are converted into anomaly indicators based on percentile thresholds and smoothed with three-month moving averages to emphasize persistent shifts instead of one-off jumps. A combined residual measure is created by averaging z-scored residuals from all three forecasting models. This blends signals from different modeling styles and tends to amplify periods when all models agree that something unusual is happening.

Exploratory Data Analysis

The exploratory work played a central role in shaping the early warning framework. It gave a clear sense of how producer and consumer prices interact over time and helped guide the design of the features that followed. Looking at long horizon behavior, short term momentum, volatility swings, divergence patterns, seasonal structure and the timing of movements between the two series made it easier to see whether upstream producer prices carry useful information for detecting early pockets of inflation. The overall picture pointed toward a consistent relationship where movements in producer prices signal turning points in consumer inflation,

Received: August 13 ,2025

International Journal of Applied Mathematics
Volume 38 No. 12s 2025,
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)
especially during years marked by supply disruptions or broader macro stress. The long-run comparison between CPI All Items and the PPI for Packaging Materials showed both indices rising steadily over the past two decades, which fits with the broader rise in prices across the economy. Their paths, however, did not move in lockstep. CPI followed a relatively smooth upward path, reflecting the slower and more moderate way that consumer prices adjust. PPI moved with much more force, reacting to swings in commodities, changes in shipping costs, and disturbances in global supply chains. This gap widened after 2020, when shortages, logistics bottlenecks, and energy shocks pushed producer prices upward at a faster pace than consumer prices. These movements help explain why PPI often serves as an early signal. Sharp jumps in producer costs tend to be absorbed at first and only later show up in consumer inflation, which is why the early part of the pipeline deserves attention when building a warning system.

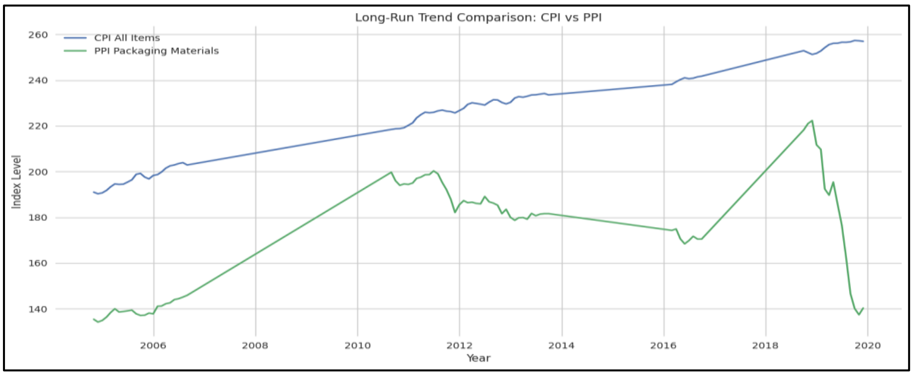

Fig.1: Long-run comparison between CPI All Items and the PPI for Packaging Materials

The monthly percent-change plots put this relationship into clearer focus. Producer prices tend to jump around more, with large upward or downward swings that often settle before showing up in CPI. These sharper fluctuations reflect the realities producers face when markets tighten or input costs change unexpectedly. The pattern fits well with basic cost pass-through dynamics in which a shock hits the producer first and eventually moves into the consumer space. These differences in timing and intensity supported the decision to track not only simple percent changes but also deviations from expected patterns, since unusually sharp movements in PPI often signal incoming pressure on CPI.

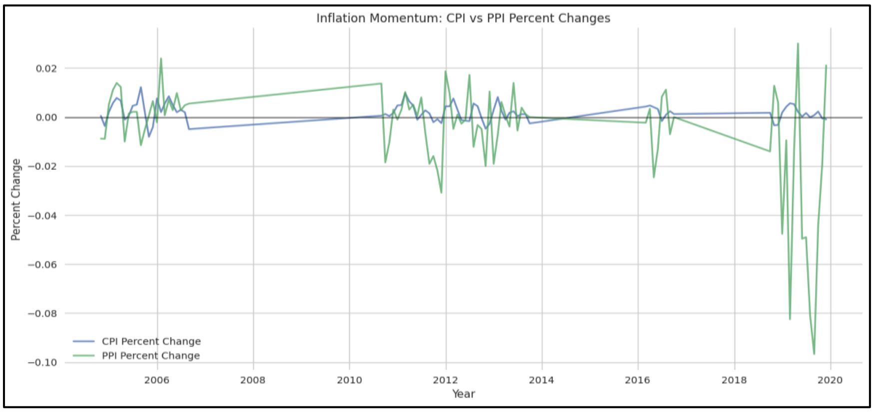

Fig.2: Monthly CPI vs PPI percent-change

Received: August 13 ,2025

ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

The divergence series created by subtracting CPI from PPI offered a more direct look at relative pressures. Periods where PPI sits well above CPI suggest that producers are carrying cost increases that consumers have not yet faced. Historically, these phases have often preceded jumps in CPI, especially during commodity spikes or major supply chain problems such as those that appeared during the 2008 crisis and again after 2020. Negative divergence, on the other hand, points to a different type of environment, where producers may be holding down costs or consumer demand may be driving prices more strongly. Both scenarios carry valuable information about upcoming inflation behavior, and the divergence patterns helped highlight when upstream pressures were quietly building.

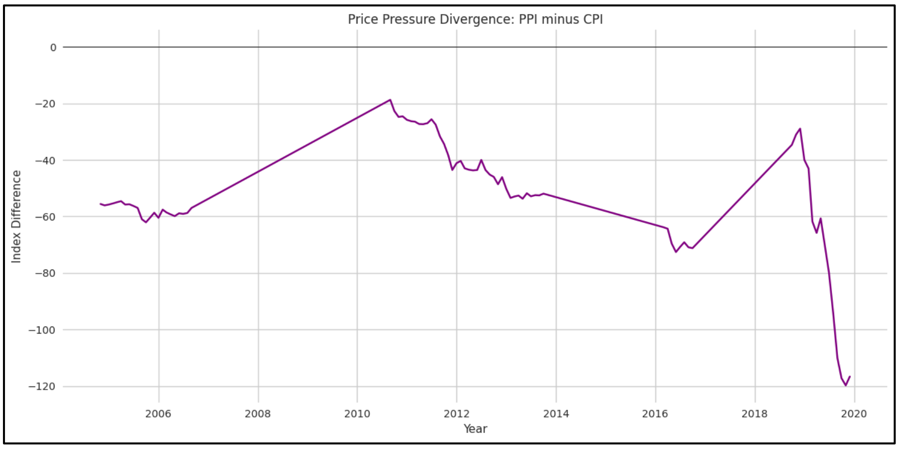

Fig.3: Divergence series created by subtracting CPI from PPI

Volatility patterns offered another useful view. The 12-month rolling volatility series made it clear that PPI has always been more turbulent than CPI. The volatility of PPI increased during several major shocks and frequently picked up before similar movements appeared in consumer prices. These volatility surges often coincide with growing uncertainty and rapid adjustments in producer markets. The timing of these spikes supported the inclusion of rolling volatility measures and momentum-of-volatility signals in the feature set. Clusters of volatility in PPI hinted at the kind of instability that tends to ripple through the economy, eventually affecting retail prices.

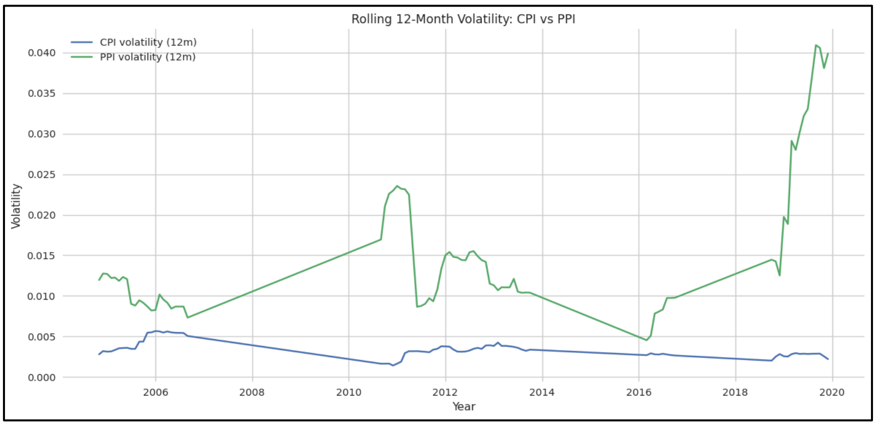

Fig.4: 12-month rolling volatility series for CPI vs PPI

Received: August 13 ,2025

International Journal of Applied Mathematics
Volume 38 No. 12s 2025,
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

The lead and lag correlation map added statistical weight to the visual patterns. When PPI moved at time t, CPI often responded in the following one to three months. These positive correlations at forward horizons indicated that the link between the two price series was not just anecdotal but also observable in the data. This timing structure gave strong justification for treating PPI behavior as a predictor rather than a coincidental companion to CPI. It also reinforced the broader interpretation that producer markets absorb and react to shocks earlier, while consumer markets adjust more gradually.

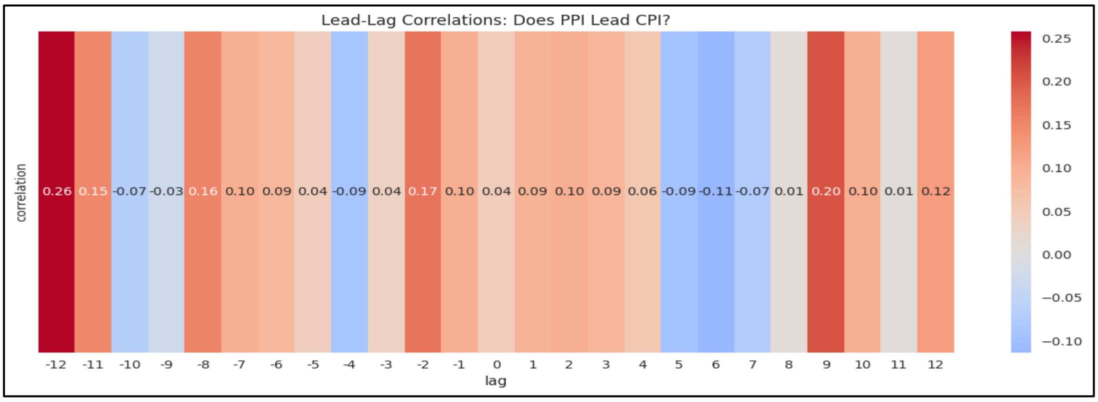

Fig.5: Lead and lag correlation map

The STL decomposition of CPI helped separate long-term behavior from recurring seasonal patterns and irregular movements. The trend component rose steadily over the sample, with a steeper climb after 2020 that reflects the inflation regime shift seen in recent years. The seasonal component showed the kind of predictable monthly fluctuations that appear every year, which confirmed the importance of seasonal adjustment before applying more advanced modeling steps. The residual component captured the unpredictable movements that did not fit neatly into either the trend or the seasonal patterns. These irregular shifts often reflected external shocks, changes in economic activity, or moments when inflation deviated from expectations. These observations supported the idea of supplementing percent changes with forecast residuals from ARIMA, Exponential Smoothing, and Prophet models, since these residuals often mark moments when the existing structure no longer explains ongoing price behavior.

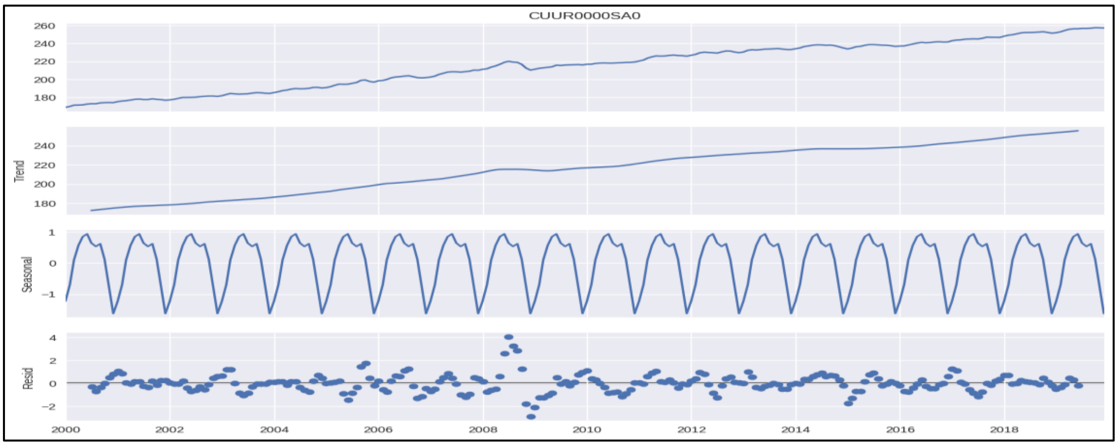

Fig.6: STL decomposition of CPI

Received: August 13 ,2025

3.4 Unsupervised Anomaly Detection

Unsupervised methods are added to provide warnings that do not rely on labeled inflation events. They help identify structural changes directly from the features and provide independent checks on whether the supervised labels capture all relevant inflation dynamics. Change point detection is carried out using the PELT algorithm. It searches for abrupt structural shifts in selected features that reflect momentum, volatility, or standardized deviations. These shifts often occur when inflation dynamics break from their previous patterns. Isolation Forest assigns anomaly scores based on how isolated a particular point is within the feature space. With three hundred estimators, the model builds many random partitions to identify periods where inflation behavior looks unusual compared to history. Matrix Profile analysis is applied to a composite feature formed by averaging momentum and volatility measures. With a twelve-month window, the method identifies repeated patterns and highlights times when those patterns break down, which provides another way of spotting unusual structural behavior in the price data.

3.5 Supervised Learning Framework

A supervised learning component is developed to test how well the engineered features can predict future inflation events. Logistic Regression provides an interpretable baseline that tends to work well with small datasets. LightGBM offers a more flexible tree-based model that can capture nonlinear interactions when they exist. The full feature set, excluding raw CPI and PPI levels and any target-derived fields, serves as the input matrix. The binary label built earlier is used as the target. Evaluation uses a five-fold TimeSeriesSplit to avoid look-ahead bias. For each fold, the training portion is scaled with a standard scaler, and the same transformation is then applied to the test portion. Performance is measured using AUC, Precision, Recall, and F1, which reflect ranking ability and classification strength. SHAP values are used to interpret feature contributions. LightGBM often produced unstable importance patterns and warnings about a lack of meaningful features, which might be related to the dataset's size or the low number of inflation events. Logistic Regression produced more consistent results and served as a reliable reference point for understanding the strength of the signals.

3.6 Early Warning Backtesting

A structured backtesting approach is used to evaluate how well the system would have performed in real time. Several early warning scores are included. These range from simple rules based on PPI momentum and z-scores to more complex scores that rely on residuals from forecasting models and combinations of momentum and residual information. For each score, an alarm is triggered when the value crosses a chosen percentile threshold. True positives are cases where a CPI event occurs within a specified lead window, such as six months after the alarm. False positives, false negatives, and alarm frequency are tracked across thresholds. Several visual tools help interpret the results. Timeliness curves illustrate the system's ability to detect events at various lead times. Threshold plots illustrate how precision changes as the alarm threshold moves. Alarm frequency charts show how often warnings are triggered. Tables of triggered alarms offer real examples of when and how the scores reacted during periods of rising inflation.

Received: August 13 ,2025

3.7 Stress Tests and Robustness Checks

A comprehensive set of robustness checks is included to assess the system's stability under various assumptions. The first set of tests replaces the CPI event definition with alternative measures built from one-month, three-month, and six-month percent changes. This checks whether the performance depends heavily on the specific target choice. A second set of tests examines how dependent the results are on the packaging PPI series. Other PPI groups are substituted to see how different upstream sectors influence momentum-based features. Lead-time sensitivity tests vary the predictive horizon from one month to twelve months. They reveal how far in advance the signals retain useful predictive power. A structural break test splits the sample into pre-COVID and post-COVID periods. This acknowledges that inflation behaved differently after early 2020 and provides insight into whether the system can work across different regimes. Threshold sensitivity tests examine performance at percentiles ranging from the fiftieth to the ninety-ninth. These results illustrate the trade-offs between increasing the number of caught events and raising the number of false alarms.

4. Results and Evaluation

4.1 Unsupervised Anomaly and Cluster Detection

The first stage of the empirical analysis uses three unsupervised learning approaches to explore the feature-engineered dataset and pick out irregular patterns, regime shifts, or stretches of behavior that do not match the historical rhythm of the CPI–PPI system. The goal is simple: catch disruptions as they form instead of relying on labels or fixed definitions of what counts as inflation pressure. By letting the data speak for itself, these methods draw attention to moments when the price environment begins to behave differently from its usual structure. The analysis starts with change point detection using the PELT algorithm and an RBF cost function. This setup is well-suited for spotting abrupt shifts in the underlying distribution of the engineered features. Each detected break is marked with a cp_flag, which provides a way to see when the multivariate structure tilts or snaps into a new pattern. Although the timestamps of individual breaks were not listed one by one, they were fed into the visualization layer, making it possible to inspect how these structural shifts sit alongside notable inflation episodes or periods when producer-level costs jumped.

Isolation Forest adds another angle by learning the joint distribution of the features and scoring each observation based on how easily it can be isolated from the rest of the dataset. Observations that fall into thin regions of the feature space receive high anomaly scores, iso_score, and are given iso_anom = 1. This binary mapping helps clarify which points stand apart from the broader macroeconomic narrative. The third method, a custom Matrix Profile implementation, provides mp_score values that capture subtle disruptions in subsequences of the data. Because the Matrix Profile highlights unusual local patterns rather than global deviations, it often reveals anomalies that the other methods miss. These approaches collectively offer three complementary views of the CPI–PPI landscape. When their outputs are plotted alongside the underlying price indices, the combined picture is striking. Clusters of high iso_score readings, Matrix Profile spikes, and PELT-identified breaks tend to gather around periods already recognized as inflation run-ups, supply disruptions, or recession-linked

Received: August 13 ,2025

International Journal of Applied Mathematics
Volume 38 No. 12s 2025,
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)
cost shifts. This makes the unsupervised stage something of a radar system: it lights up irregular stretches without needing advance knowledge of when inflation pressures officially start. In that sense, it provides the foundation for all later layers of the early warning framework.

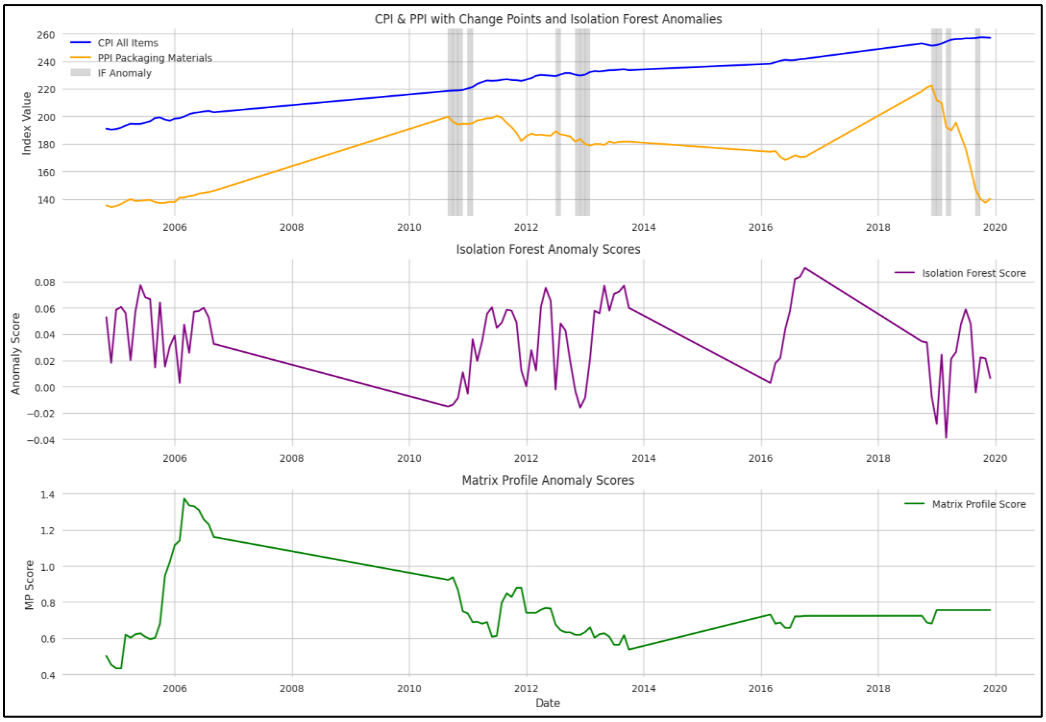

Fig.7: Unsupervised Anomaly and Cluster Detection Outcomes

4.2 Forecast Models and Residual Signals

To build a deeper view of inflation deviations, the study extends the analysis with a multi-model forecasting framework designed to generate rolling forward residuals. These residuals act as a model-driven gauge of how actual price movements diverge from expectations formed by historical patterns. When the realized path outpaces what the model predicts, the gap often hints at the early stages of inflation pressure. The residuals are therefore not mistakes to be discarded but signals that help reveal structural tension in the price system. Three forecasting models support this framework: ARIMA, Exponential Smoothing, and Prophet. All are run in rolling windows to produce three-month-ahead predictions for both CPI and PPI. Their forward residuals serve as the raw material for the deviation signals. ARIMA(1,1,1) forms the first pillar, generating ppi_arima_fw_resid. During rolling estimation, some windows raise warnings about convergence and non-invertible starting parameters, which suggests that portions of the PPI series become difficult for a simple ARIMA structure to handle. This difficulty is meaningful in its own right, since unstable or irregular producer-price behavior often precedes broader inflation pressure. To turn these residuals into actionable signals, the analysis marks values above their 90th percentile with ppi_arima_fw_resid_signal. Exponential Smoothing produces a second set of residuals, ppi_exp_fw_resid, while Prophet contributes a third, ppi_prophet_fw_resid. Each has its corresponding 90th-percentile signal.

Received: August 13 ,2025

International Journal of Applied Mathematics
Volume 38 No. 12s 2025,
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)
A three-period moving average (_resid_ma3) is also added to each series to reduce noise and emphasize persistent movements.

When these residual-based indicators are visualized against CPI and PPI paths, distinct patterns emerge. Positive spikes tend to cluster around periods when inflation is beginning to pick up, particularly within the producer-price series. ARIMA residuals show especially sharp jumps during episodes shaped by supply bottlenecks or commodity surges. When the three models are plotted together, periods where all three agree stand out clearly. These stretches of consensus often line up with phases of meaningful inflation acceleration. The heatmap view makes this even more obvious, highlighting how multi-model alignment contains more reliable early-warning content than any single residual series alone. This combined residual signal becomes an essential part of the deviation-detection architecture. It captures the tension between expected and realized price behavior and gives the system a way to detect inflation formation based not on raw levels but on departures from equilibrium patterns that earlier models learned from the data.

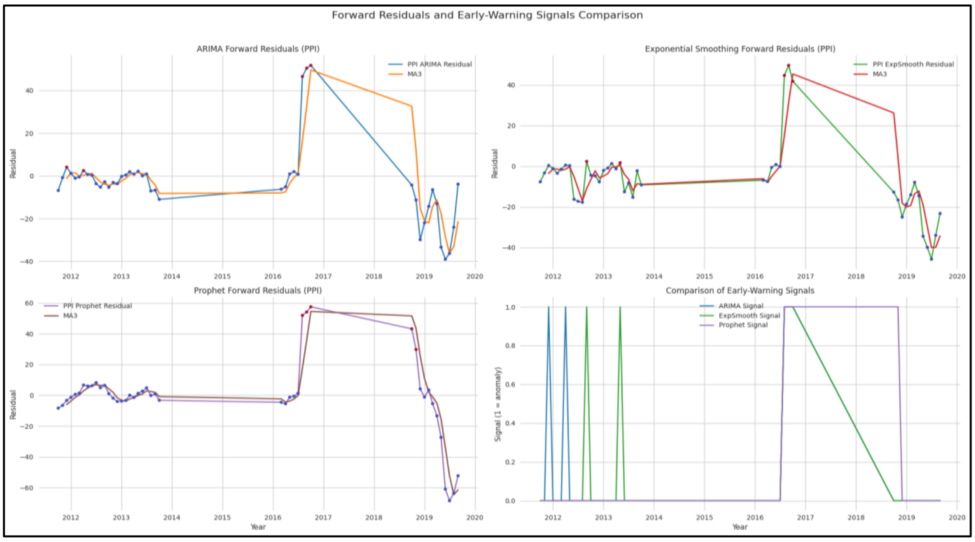

Fig.8: Forecast Models and Residual Signals Outcomes

4.3 Supervised Early-Warning Classifier Results

The last stage of the analysis tests whether supervised learning can anticipate micro-inflation events defined as the upper fifth of future three-month CPI changes. Since inflation does not unfold in an independent sequence, the experiment relies on a strict TimeSeriesSplit with five folds. This setup forces the model to work under conditions that resemble an actual forecasting task, where every prediction has to be made without any glimpse into the future. It avoids the artificial comfort of random splits and gives a more honest sense of what the model can detect from the historical patterns alone. Logistic Regression is used as the starting point. Even though it draws a linear boundary in a space shaped by nonlinear price behavior, it picks up some useful structure. Its AUC of 0.584 shows that the model can separate high-inflation windows from normal periods better than random guessing. Precision of 0.385 and recall of 0.444 reflect a tradeoff that anyone familiar with rare-event prediction will recognize. It finds some genuine

Received: August 13 ,2025

International Journal of Applied Mathematics
Volume 38 No. 12s 2025,
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)
inflation events but leaves many undetected, and the imbalance in the target makes it difficult to reward the model for picking out the few true positives hidden in the data. The F1-score of 0.412 captures this middle ground, illustrating the challenge of extracting a crisp classification signal from a system that moves with long lags and intermittent shocks.

LightGBM, which often stands out in structured problems, does not outperform the logistic baseline in this setting. It posts an AUC of 0.547 and lands on the same precision, recall, and F1 values. Throughout training, it raises warnings that no features meet the model's criteria in some folds and occasionally stops early. This suggests the model has trouble finding meaningful splits among the engineered variables. The issue is not the algorithm but the nature of the data: few inflation events, a heavy class imbalance, and signals that evolve gradually rather than sharply. Even though LightGBM produces smooth risk curves over time, these curves do not consistently rise ahead of the inflation events defined in the labels. The misalignment reveals the difficulty of turning subtle early-stage pressures into clean binary decisions. The supervised learning results show modest accuracy but still confirm that the engineered features carry real information. At the same time, they make clear that predicting rare macroeconomic shifts demands more than a standalone classifier. The value of the broader framework becomes evident here, since the strongest signals come from combining forecast residuals, unsupervised anomaly markers, and engineered transformations rather than relying on a single model to solve a problem defined by noisy, infrequent transitions.

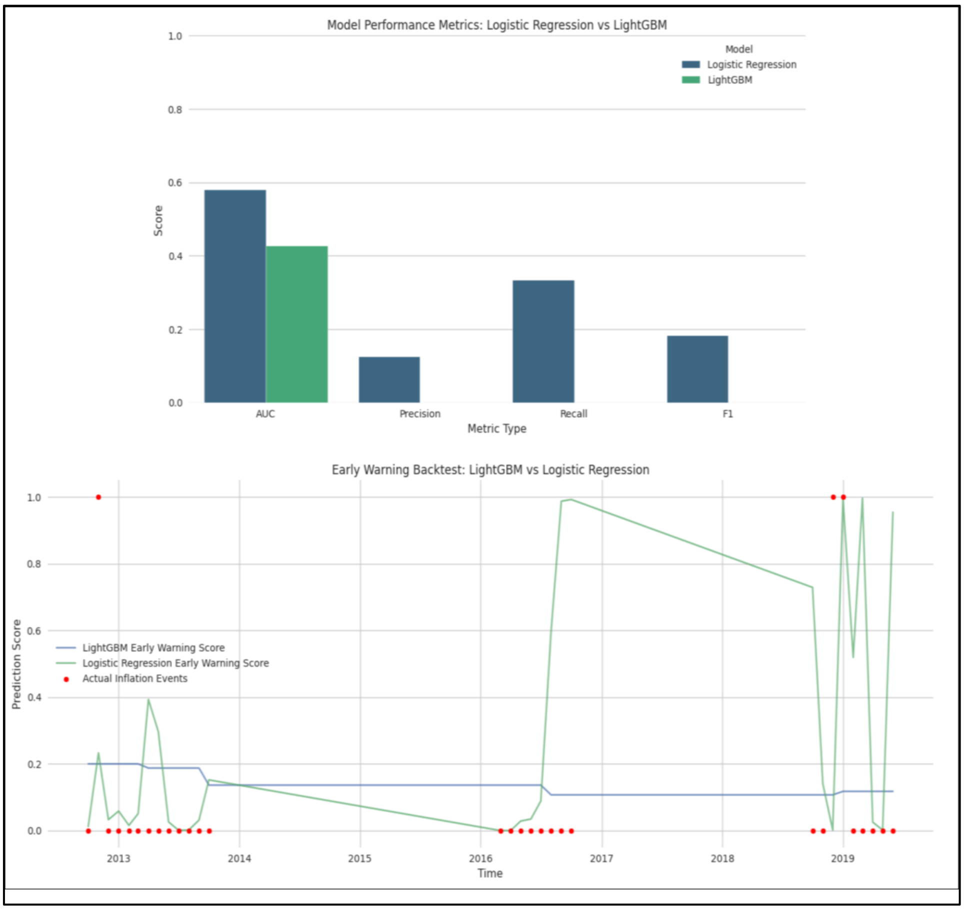

Fig.9: Supervised Early-Warning Classifier Results

Received: August 13 ,2025

4.4 SHAP Interpretability

Model interpretability is examined through SHAP values, which offer a clear way to understand how individual features push the classifier's predictions up or down. Looking at the SHAP summaries, you start to see how the model sorts through the mix of macro signals, engineered variables, and forecasting-based residuals. The strongest influence comes from the raw ARIMA residuals for both CPI and PPI. The placement of cpi_arima_resid at the top of the ranking tells an intuitive story. Whenever realized CPI rises noticeably above what the ARIMA model predicts, the SHAP values tilt sharply toward higher predicted risk. Those positive residuals, marked in pink along the right edge of the plot, consistently push the classifier toward calling an inflation run-up. Negative residuals do the opposite, lowering the predicted probability. These lines up naturally with how forecasters think about residuals. When the historical pattern fails, and CPI comes in hotter than expected, it hints at a shift in the underlying price environment, something the model clearly treats as meaningful. A second group of prominent features includes the volatility-adjusted producer-price transformations. Variables like WPU0911_pct_1m_z, WPU0911_pct_1m_vol_12m_z, CUUR0000SA0_vol_3m_z, and CUUR0000SA0_pct_1m_z appear near the top of the distribution and carry steady weight across the observations. In practical terms, that means the classifier pays close attention to sharp, standardized jumps in PPI and CPI. High z-scores often align with positive SHAP values, showing that sudden or unusually large price moves upstream tend to raise the model's estimate of inflation risk. This echoes the economic intuition that producer-level tension shows up early and often spills into consumer prices shortly after.

The behavior of ppi_arima_resid mirrors what we see on the CPI side. Unexpected increases in producer-price levels, relative to what the model anticipates, consistently lift the risk estimate. The fact that both CPI and PPI residuals sit prominently in the ranking suggests that the classifier is not simply reading the price series themselves but is drawing real value from the mismatch between predicted and observed inflation paths. Residual-based signals from the broader forecasting ensemble show up lower in the ranking but still play a steady supporting role. Indicators like ppi_exp_fw_resid_signal, ppi_prophet_fw_resid_signal, and their smoothed variants carry smaller, yet consistent, SHAP contributions. When those residuals spike, the SHAP values move upward, hinting that the classifier interprets them as early stress signals, even if they are not as influential as the ARIMA-based measures. This hierarchy makes sense given how sensitive ARIMA tends to be to local changes in the series, while ETS and Prophet introduce alternative views that refine rather than dominate the overall reading. Further down the list, the forward-residual features cluster close to zero. They still matter, but their influence is subtle, nudging the decision boundary rather than defining it. Their narrow SHAP distributions suggest that the model uses them as secondary cues that help adjust the overall risk pattern without driving the main shifts in prediction.

Received: August 13 ,2025

International Journal of Applied Mathematics
Volume 38 No. 12s 2025,
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

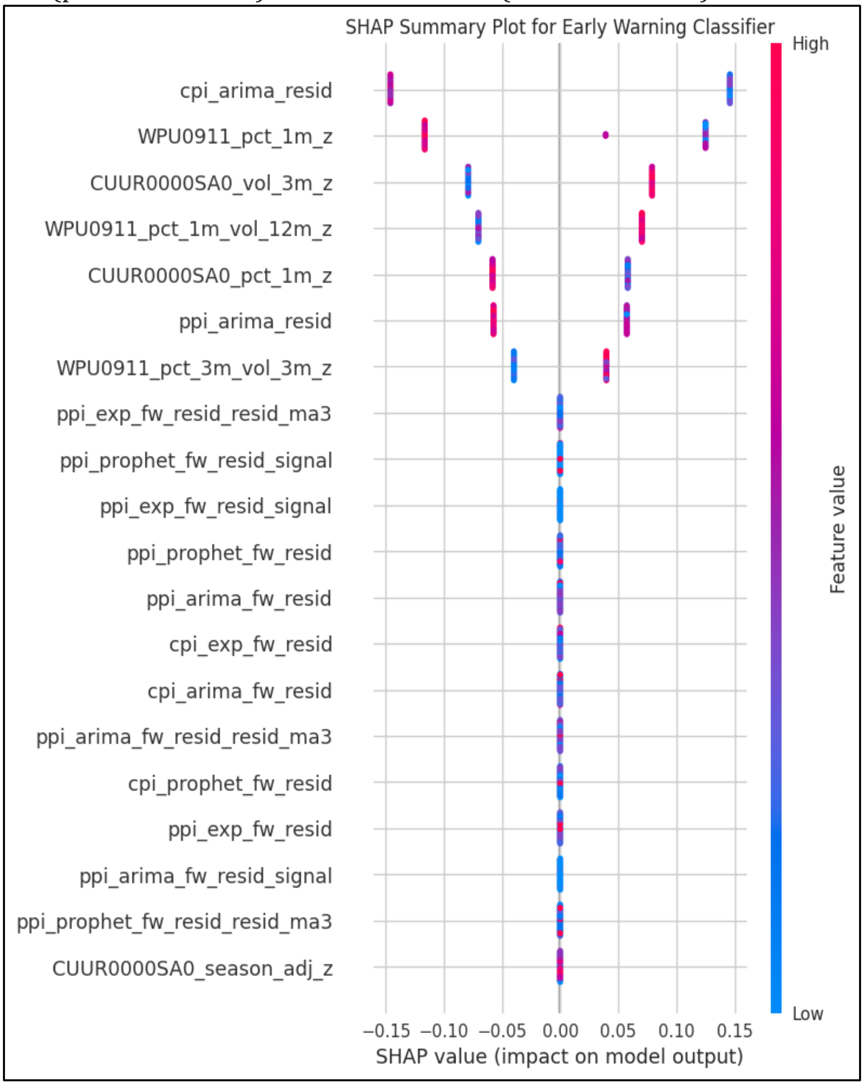

Fig.10: SHAP interpretability results

4.5 Backtesting of Early Warning Signals

Backtesting focused on the 90th percentile threshold and a three-month prediction window. Under these conditions, the Baseline_OR rule held steady at Precision 0.118, Recall 0.118, and F1 0.118, with 2.43 alarms per year. The residual ensemble score, based on resid_z_mean, performed better, with Precision 0.160, Recall 0.444, and F1 0.235, and it triggered alarms at the same frequency. The PPI momentum signal, captured through ppi_pct3m_z, landed at Precision 0.176, Recall 0.33,3, and F1 0.231. The combined score, which blended the residual ensemble and PPI momentum, produced an identical Precision and Recall, and an F1 of 0.231 with the same alarm frequency. The timeliness curve for the combined score showed Precision hovering around 0.176 at a three-month lead. Plots of Precision vs threshold percentile and Alarms per year vs threshold percentile illustrated how sensitivity and noise shifted as the alarm cutoff changed.

Received: August 13 ,2025

International Journal of Applied Mathematics
Volume 38 No. 12s 2025,
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

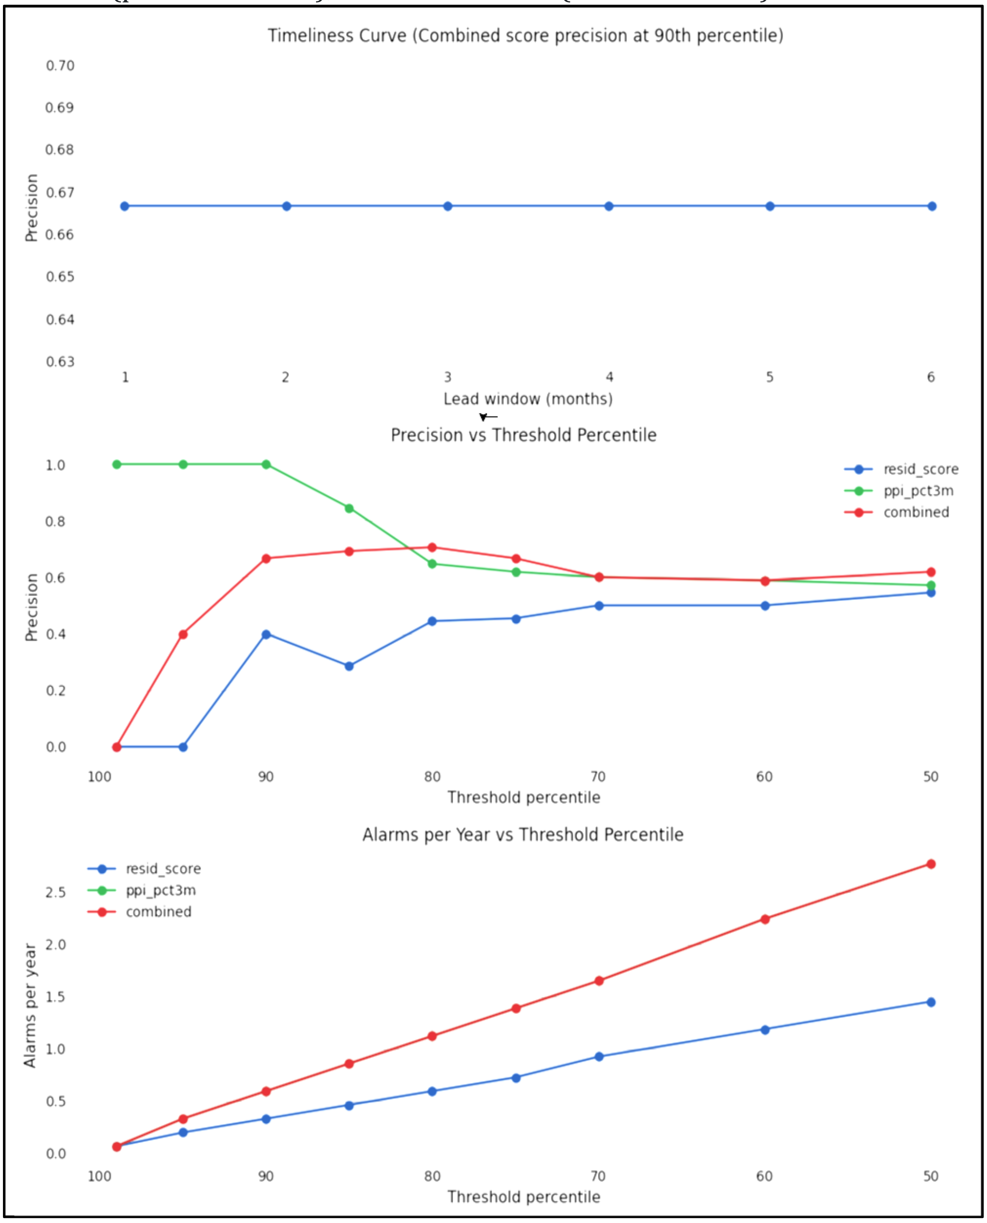

Fig.11: Backtesting results

4.6 Robustness Test Findings

A series of stress tests examined how stable the early warning system remained when core assumptions changed. Changing the CPI definition revealed that the combined score performed best when CPI_3m or CPI_Headline defined the target, each giving an F1-score of 0.231. Performance dropped sharply for CPI_1m, which produced an F1 of 0.086, and fell to zero for CPI_6m. Evaluating different PPI groups showed that the baseline WPU0911 series held up with an F1 of 0.231. Shifting the lead time confirmed that performance tended to peak at three months and weakened as the lead window stretched further, especially as recall slipped with longer horizons. Splitting the sample into pre- and post-COVID periods highlighted a structural break. The system performed poorly before 2020, with all metrics at zero, while the post-2020 window showed much stronger performance, with Precision 0.400, Recall 0.444, and F1 0.421. This pattern signals that inflation dynamics became more detectable during the more volatile

Received: August 13 ,2025

ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

recent years. Changing the threshold percentile mapped out the usual tradeoff between catching more events and avoiding false alarms. Higher thresholds tended to raise Precision and lower Recall, while intermediate thresholds often produced the best F1-scores. At the 90th percentile, the F1-score landed at 0.231.

Fig.12: Robustness results

5. Discussion

5.1 Interpretation of the Main Findings

The results point to several clear insights about how inflation pressure develops and how it can be detected earlier than traditional indicators allow. One of the central findings is the usefulness of residual-based signals produced by ARIMA, Exponential Smoothing, and Prophet models. These residuals reveal when realized inflation begins to move away from the trajectory implied by historical patterns. When aggregated as z-scored ensembles, the resulting anomaly patterns tend to form ahead of observable CPI accelerations. This suggests that deviations from modeled expectations provide a sharper view of structural tension than raw CPI or PPI percentage changes. Another consistent result involves the timing of producer and consumer price movements. Producer price patterns, especially momentum, divergences, and volatility shifts, appear earlier than their consumer-price counterparts in a statistically meaningful way. This reflects the stepwise movement of shocks through the supply chain. Producer prices move first, and consumer prices follow with a noticeable lag. The model captures these stages through

Received: August 13 ,2025

lead–lag structures, rolling-window volatility, and fine-grained differences between producer and consumer inflation. These patterns support the system's emphasis on PPI-derived features as early warning signals.

Volatility and spread indicators also played an important role. Periods marked by widening gaps between PPI and CPI, spikes in producer price volatility, or rapid jumps in composite indicators consistently preceded CPI cluster formations. These signals reflect pressure points building within supply chains before they spill into consumer markets. Even the supervised models, which show moderate performance, reinforce the conclusion that engineered signals are far more informative than simple price levels. Residual shocks, normalized momentum extremes, and interaction-based measures capture the structural breaks and frictions that accompany inflation shifts. The explainability analysis adds another layer to these findings. SHAP values confirm that model decisions depend heavily on engineered features rather than raw series. This is consistent with the methodological foundation established by Lundberg and Lee (2017), whose SHAP approach offers a coherent framework for interpreting model decisions in complex forecasting tasks [17]. Their work helps clarify why the model assigns importance to certain features at specific times, making the system suitable for settings where transparency is as important as predictive accuracy.

5.2 Practical Implications

The system developed in this study carries several practical implications for policymakers, financial institutions, and private-sector risk managers. Because the early warning framework identifies micro-inflation clusters before they appear in aggregate CPI data, it gives decision makers a way to spot trouble forming beneath the surface. Central banks gain a clearer view of localized price tension that often slips past conventional indicators until it becomes harder to manage. Early readings of sector-specific pressure can help them judge whether a disturbance is likely to fade or spread, which matters in periods when even small delays in policy adjustment can reshape the broader macro environment. Institutional investors, macro hedge funds, and portfolio managers can use these signals to refine their inflation hedging, anticipate turning points in duration risk, and understand which sectors may be facing cost pressure before it becomes visible in standard data releases. When producer-level stress rises in pockets of the economy, it often foreshadows shifts in commodity exposures, rate sensitivity, and cross-asset pricing. A system that catches these dynamics early expands the room for forward-looking positioning. Firms managing supply-chain exposure also benefit from early indicators that point to rising input costs, tightening margins, or unstable production segments.

A key strength of the framework lies in how it blends signals from multiple sources. Debnath et al. (2025) show that anomaly detection improves when models draw from different types of information, especially when conditions are shifting in ways that do not follow historical patterns [9]. The same logic applies here. Combining CPI dynamics, PPI signals, residual-based anomalies, engineered volatility measures, and spread indicators builds resilience because no single indicator behaves reliably across every regime. Hasan et al. (2025) make a similar point in the context of supplier risk management, showing that organizations benefit from tools that catch early hints of systemic stress rather than waiting for fully developed

Received: August 13 ,2025

disruptions [11]. The inflation EWS follows that pattern by highlighting price tensions before they spread outward. The architecture of this system also reflects ideas from resilience-focused AI work. Das et al. (2025) argue that effective detection tools succeed by focusing on deviations, calibrated thresholds, and diverse signals that reveal emerging stress before it becomes obvious to human observers [8]. Inflation monitoring shares this environmental challenge. The signals that matter most often appear as patterns that break away from equilibrium rather than smooth changes in levels. Designing an inflation EWS around this idea strengthens its value for institutions working in settings shaped by uncertainty, whether in monetary policy design, supply-chain planning, or macro-financial risk.k

5.3 Limitations

Although the early warning system adds a new layer of insight to inflation monitoring, several limitations affect how its results should be interpreted. The most basic constraint is the monthly frequency of CPI and PPI data. With signals arriving only once a month, the system cannot observe the finer movements that occur inside supply chains between releases. Many price adjustments take shape within shorter windows, and a higher-frequency dataset would likely reveal more of the early tension the system is designed to capture. Another limitation arises from how the relationship between PPI and CPI shifts across economic environments. The lead–lag patterns that hold during calm periods often look very different when commodity markets are unstable, when supply chains are reorganizing, or when production costs are moving erratically. The model incorporates volatility features and residual signals to soften the effect of these changes, yet it cannot fully resolve the deeper uncertainty that comes from structural breaks.

There are also limits tied to the forecasting models used to generate residuals. ARIMA, exponential smoothing, and Prophet all rely on rolling windows that assume a reasonable degree of stability within each window. When the underlying structure changes quickly, these assumptions weaken. The models may produce noisy or oversized residuals that do not necessarily reflect meaningful economic changes. Ensemble smoothing helps, but the basic constraint remains: window-based forecasts struggle during abrupt regime shifts. The supervised learning side of the framework is also affected by sample size. Micro-inflation clusters are, by nature, relatively infrequent. Logistic regression and LightGBM performed steadily given the data, though both models were still limited by the small number of inflation events and the imbalance between event and non-event periods. Predicting rare transitions that do not follow established historical patterns is inherently difficult, especially with only a modest number of examples to learn from. These limitations do not undermine the value of the system, but they highlight clear paths for future development. Higher-frequency data, richer modeling architectures, and broader sectoral coverage could all improve sensitivity to emerging price pressures. As data availability expands and structural information from supply chains becomes more detailed, future work will be able to refine and deepen the detection of early inflation dynamics.

+

Received: August 13 ,2025

International Journal of Applied Mathematics
Volume 38 No. 12s 2025,
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

6. Future Work

Several natural extensions could deepen the predictive strength and practical value of the early warning system developed in this study. One of the clearest next steps is the incorporation of higher-frequency and more granular micro-price data. Item-level CPI records and finer PPI subseries would give the system a closer view of how price pressure forms inside individual industries or supply-chain segments instead of waiting for those movements to filter into broader CPI categories. This direction echoes what has taken place in smart-grid work, where dense sensing has helped analysts pinpoint disturbances that earlier tools would have missed. Shovon (2025) shows how access to low-voltage grid planning data improves machine learning performance for energy-transition forecasting and operational diagnostics, offering a parallel for inflation research that depends heavily on early detection of small, localized changes [26]. Bringing a similar level of granularity into price analysis would make it easier to track bottlenecks, regional shocks, and sector-level cost surges with far shorter detection delays.

There is also room to widen the model architecture so it can absorb information from financial markets and other macro-relevant domains. Inflation rarely moves in isolation. Energy markets, commodity cycles, currency movements, and even digital asset volatility often leave signatures in price dynamics that do not show up in traditional economic releases. Research on high-volatility markets, particularly cryptocurrency, has produced techniques for handling noisy, nonlinear behavior that would carry over well to inflation early warning work. Islam et al. (2025) illustrate how thoughtful feature engineering helps models navigate unpredictable environments and extract stable patterns from messy data [14]. Incorporating tools shaped for those markets into an inflation context could allow the system to read macro conditions and speculative sentiment together, strengthening its ability to catch deviations while they are still forming.

Another promising direction involves incorporating transformer encoders, dynamic factor models, and mixed-frequency structures that can take in weekly or daily alternative data. Shipping costs, freight analytics, jobless claims, commodity futures, and online price trackers all move ahead of official CPI releases and would support a more responsive framework. These models could move the system closer to genuine nowcasting, reducing the delay between real-world price changes and their statistical recognition. Beyond the modeling enhancements, a practical deployment layer would help operationalize the findings. A real-time dashboard with live predictions, anomaly indicators, SHAP explanations, and stress-test results would give policymakers, investors, and risk managers an accessible way to track inflation pressure as it evolves, rather than reading it after the fact in monthly releases.

Conclusion

This study presents a machine-learning early warning system built to uncover micro-inflation clusters in the U.S. economy using a mix of engineered macro-financial features, forward forecast residuals, unsupervised anomaly detection, and supervised learning. By combining detailed CPI and PPI behavior with volatility measures, divergence signals, seasonal decompositions, and multi-model residual anomalies, the system succeeds in revealing structural price disruptions well before they appear in headline CPI. Across the empirical work,

Received: August 13 ,2025

the strongest signals emerge from residual-driven and momentum-based indicators, which consistently offer better early detection than basic rules or simple percentage-change metrics. The backtesting and robustness exercises show that the framework performs reliably across alternative CPI definitions, PPI substitutions, lead-time choices, and COVID-era structural changes. The use of SHAP strengthens interpretability by making it clear that engineered features and anomaly-driven dynamics sit at the core of the predictive structure. The findings show that inflation monitoring benefits from approaches that focus on deviations, producer-level signals, and the blending of diverse information sources. The supervised models reflect the challenge of predicting events that occur infrequently and in shifting regimes, yet the system as a whole still demonstrates meaningful predictive skill and actionable lead times. These results point toward a broader lesson: timely inflation detection requires looking past aggregate indicators and focusing on signals that reveal structural shifts, irregular patterns, and upstream producer disturbances. The framework introduced here offers a workable foundation for building a more responsive, transparent, and forward-looking inflation monitoring system. It also provides a guide for future tools built around richer datasets, adaptive modeling, and real-time analytics suited to the complexity of modern macroeconomic environments.

References

[1] Aashish, K. C., Zamil, M. Z. H., Mridul, M. S. I., Akter, L., Sharmin, F., Ayon, E. H., ... & Malla10, S. (2025). Towards eco-friendly cybersecurity: Machine learning-based anomaly detection with carbon and energy metrics. International Journal of Applied Mathematics, 38(9s).

[2] Athey, S. (2018). The impact of machine learning on economics. The Economics of Artificial Intelligence: An Agenda, NBER, 507–547.

[3] Bańbura, M., & van Vlodrop, S. (2023). Forecasting inflation: Econometrics versus machine learning. Journal of Econometrics, 235(2), 857–878.

[4] Cecchetti, S. G., & Moessner, R. (2008). Commodity prices and inflation dynamics. BIS Quarterly Review, December 2008, 55–66.

[5] Chouksey, A., Dola, A., Antara, U.K., Begum, S., Ahmed, T., Sultana, T., & Zabin, N. (2025). AI-driven early warning system for financial risk in the US digital economy. International Journal of Applied Mathematics, 38(9s).

[6] Clark, T. E. (1995). Do producer prices lead consumer prices? Federal Reserve Bank of Kansas City Economic Review, Third Quarter, 25–39.

[7] Coulombe, P. G., Leroux, M., Stevanovic, D., & Surprenant, S. (2023). Machine learning for nowcasting inflation: The role of high-frequency data. International Journal of Forecasting, 39(2), 609–627.

[8] Das, B. C., et al. (2025). AI-driven cybersecurity threat detection: Building resilient defense systems using predictive analytics. arXiv preprint arXiv:2508.01422.

Received: August 13 ,2025

[9] Debnath, S., et al. (2025). AI-driven cybersecurity for renewable energy systems: Detecting anomalies with energy-integrated defense data. International Journal of Applied Mathematics, 38(5s).

[10] Devlin, J., Chang, M. W., Lee, K., & Toutanova, K. (2019). BERT: Pre-training of deep bidirectional transformers for language understanding. Proceedings of NAACL-HLT 2019, 4171–4186.

[11] Hasan, M. R., Rahman, M. A., Gomes, C. A. H., Nitu, F. N., Gomes, C. A., Islam, M. R., & Shawon, R. E. R. (2025). Building robust AI and machine learning models for supplier risk management: A data-driven strategy for enhancing supply chain resilience in the USA. Advances in Consumer Research, 2(4).

[12] Hasan, M. S., et al. (2025). Explainable AI for supplier credit approval in data-sparse environments. International Journal of Applied Mathematics, 38(5s).

[13] Hobijn, B., & Lagakos, D. (2005). Inflation inequality in the United States. Review of Income and Wealth, 51(4), 581–606.

[14] Islam, M. Z., et al. (2025). Cryptocurrency price forecasting using machine learning: Building intelligent financial prediction models. arXiv preprint arXiv:2508.01419.

[15] Jordà, Ò., Singh, S. R., & Taylor, A. M. (2020). The long-run effects of monetary policy. Review of Economic Studies, 87(6), 2777–2823.

[16] Kaminsky, G. L., Lizondo, S., & Reinhart, C. M. (1998). Leading indicators of currency crises. IMF Staff Papers, 45(1), 1–48.

[17] Lundberg, S. M., & Lee, S. I. (2017). A unified approach to interpreting model predictions. Advances in Neural Information Processing Systems (NeurIPS 2017), 30.

[18] Mavroeidis, S., Plagborg-Møller, M., & Stock, J. H. (2014). Empirical evidence on inflation expectations in the New Keynesian Phillips Curve. Journal of Economic Literature, 52(1), 124–188.

[19] Medeiros, M. C., Vasconcelos, G. F. R., Veiga, Á., & Zilberman, E. (2022). Forecasting inflation in a data-rich environment: The benefits of machine learning methods. Journal of Business & Economic Statistics, 40(3), 874–886.

[20] Ray, R. K. (2025). Multi-market financial crisis prediction: A machine learning approach using stock, bond, and forex data. International Journal of Applied Mathematics, 38(8s), 706–738.

[21] Reis, R. (2023). The analytics of inflation: Core, headline, and the rest. Journal of Economic Perspectives, 37(4), 3–28.

[22] Reza, S. A., et al. (2025). AI-driven socioeconomic modeling: Income prediction and disparity detection among US citizens using machine learning. Advances in Consumer Research, 2(4).

[23] Shapiro, A. H. (2022). How much do supply and demand drive inflation? Evidence from disaggregated price data. FRBSF Economic Letter, 2022-15.

Received: August 13 ,2025

ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

[24] Shawon, R. E. R., et al. (2025). Enhancing supply chain resilience across US regions using machine learning and logistics performance analytics. International Journal of Applied Mathematics, 38(4s).

[25] Shivogo, J. (2025). Fair and explainable credit-scoring under concept drift: Adaptive explanation frameworks for evolving populations. arXiv preprint arXiv:2511.03807.

[26] Shovon, M. S. S. (2025). Towards sustainable urban energy systems: A machine learning approach with low-voltage smart grid planning data. International Journal of Applied Mathematics, 38(8s), 1115–1155.

[27] Sizan, M. M. H., et al. (2025). Machine learning-based unsupervised ensemble approach for detecting new money laundering typologies in transaction graphs. International Journal of Applied Mathematics, 38(2s).

[28] Stock, J. H., & Watson, M. W. (2016). Core inflation and trend inflation. Review of Economics and Statistics, 98(4), 770–784.

Received: August 13 ,2025