White Paper

News Sentiment as a Forward-Looking Driver of U.S. CPI-U Data

September 2025

Executive Summary

We build monthly news-sentiment indices and test their ability to predict U.S. CPI-U (2015–2025). Our findings indicate that there is a statistically significant causal relationship between changes in CPI and trailing news sentiments. Adding this signal to a basic inflation model improves accuracy and cuts forecast errors. The economic interpretation of our findings fit a simple expectations-to-prices story where optimism lifts spending and pricing power and the lagged impacts can be explained by infrequent price adjustments. Sentiment complements fundamentals; use it as an exogenous input with robust backtesting and governance.

Written by

Muhammad Daniyal Khan
YUKKA Working Student – Quantitative Research

Alexander Pieper
YUKKA AI & Quant Engineer

Contact Us

+49-30-610 818 69-0

product@yukkalab.com

www.yukka.ai

YUKKA AG, Hardenbergstrasse 32, 10623 Berlin

Background

Recent work shows there is meaningful room to improve macro forecasts by adding sentiment-based variables derived from news. For example, lexicon ("bag-of-words") indices built from 16 major U.S. newspapers improve the tracking and forecasting of macro conditions (Shapiro et al., 2022). Moreover, a UK study using three newspapers finds that term-count features combined with supervised learning materially boost forecast accuracy—especially in stressed periods (Kalamara et al., 2022). These studies rely largely on bag-of-words approaches; with today's NLP, the attainable gains are likely larger. Yukka's coverage spans 200K+ sources (including 10K+ licensed editorial outlets) and applies NLP layers such as Factuality and Temporality, making sentiment signals more robust and context aware.

Economic conditions are continuously reflected in news narratives. Converting text into structured sentiment features provides:

- Timeliness: near real-time read between official releases.
- Breadth: coverage across sectors and geographies that single surveys or markets can't match.
- Expectations channel: media tone shapes—and reflects—beliefs that feed into pricing and wage decisions.

This paper evaluates whether Yukka's proprietary sentiment variables improve forecasts of U.S. CPI-U month-over-month (MoM) changes. We quantify: (i) how far ahead sentiment movements affect CPI (lead time), (ii) how much variance in MoM CPI changes sentiment explains, and (iii) marginal gains in predicting magnitude, direction, and turning points. A simple plot of CPI changes and leading sentiment variables provides forward looking guidance in addition to existing nowcasts. We therefore recommend integrating sentiment as an exogenous input within the full model-risk lifecycle (governance, monitoring, and periodic re-estimation).

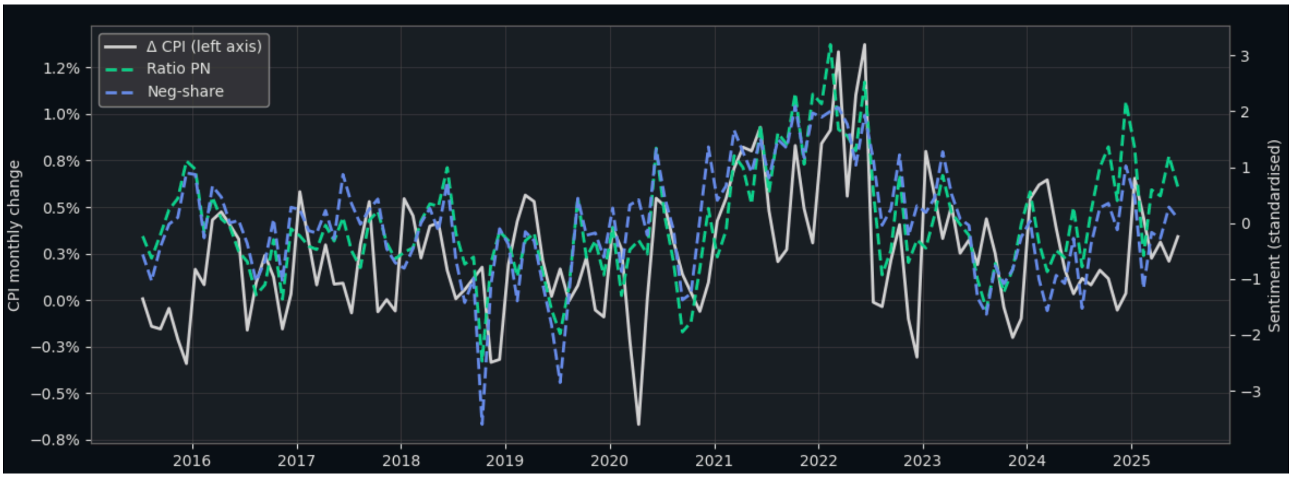

Figure 1 ΔCPI vs Forward Looking Sentiment Measures

Inputs

This study analyzes the US Bureau of Labor Statistics CPI-U since 2015 (U.S. Bureau of Labor Statistics, n.d.). The modelling target is:

$$\Delta CPI_t = \frac{CPI_t}{CPI_{t-1}} - 1$$

The CPI data is mapped against Yukka's 'Global News Based Sentiment Detection' dataset aggregated to monthly frequency. We create multiple sentiment variables based on the count of monthly Positive (Pos), Neutral (Neut) and Negative (Neg) news against US Broad market index S&P500. From these, we construct sentiment variables (Appendix 1).

The final dataset is bifurcated into an in-sample (2015 to mid 2022) and out-of-time sample (mid 2022 - mid 2025) giving us an approximately 70-30% split.

Methodology

Macroeconomic monthly changes are persistent. Therefore, we benchmark our model against a AR model of order *p*:

$$\Delta CPI_t = \alpha + \sum_{i=1}^{p} \phi_t \Delta CPI_{t-i} + \varepsilon_t,$$

Where,

ΔCPI_t is the % change in CPI for the month,

ΔCPI_{t-i} is the % change in CPI for the $t - i$ month.

ARX equation:

$$\Delta CPI_t = \alpha + \sum_{i=1}^{p} \phi_t \Delta CPI_{t-i} + \sum_{k=1}^{L} \gamma_k^{(R)} R_{t-k} + \sum_{k=1}^{L} \gamma_k^{(N)} N_{t-k} + \varepsilon_t,$$

Where,

R_{t-k} is the ratio of Positive to Negative news at $t - k$ month,

N_{t-k} is the Negative news as a % of total news at $t - k$ month,

The lag selection utlizes Granger Causality test where multiple iterations of lagged sentiments are checked for statistical significance.

Results

We tested up to 6 month lags for each sentiment construct for causal relation (Granger, 1969). The Granger grid shows that explanatory power strengthens at longer lags, peaking at L=6. At a 5% threshold we reject the null for Pos/Neg ratio and Negative-share at L=6.

Granger p-values						
Lags	1	2	3	4	5	6
Sentiment	0.52	0.62	0.77	0.78	0.30	0.16
chg_sent	0.71	0.48	0.70	0.30	0.27	0.39
neg_share	0.18	0.20	0.46	0.24	0.14	0.01
net	0.52	0.62	0.77	0.78	0.30	0.16
pos_share	0.86	0.41	0.18	0.24	0.07	0.15
ratio_pn	0.14	0.17	0.54	0.39	0.10	0.00

Next, we regressed ΔCPI on the L=6 versions of both variables. The result indicates that individual t-statistics weaken when both variables enter simultaneously—evidence of collinearity (Appendix 2). Simple correlations over the lookback window confirms this:

- $corr\ (\Delta CPI_t\ , R_{t-6}) = +0.46$
- $corr\ (\Delta CPI_t\ , N_{t-6}) = -0.41$

Therefore, we retain Pos/Neg ratio (R) as the primary predictor and exclude Negative-share to avoid redundancy and our final model becomes:

$$\Delta CPI_t = \alpha + \phi_1 \Delta CPI_{t-1} + \gamma_R R_{t-6} + \varepsilon_t$$

Comparing the AR baseline with the ARX model (with Pos/Neg ratio), the adjusted R squared increases by 0.15. We then map the predicted ΔCPI from each model into CPI-U levels to visualize fit.

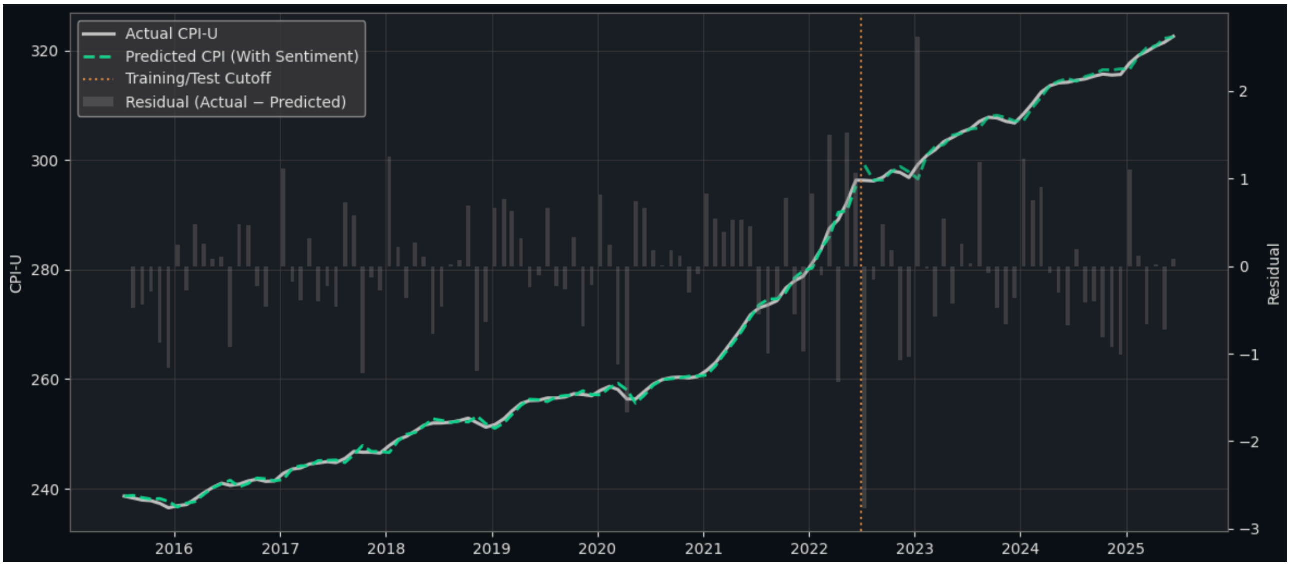

Figure 2 Actual vs Predicted CPI-U (with Sentiment) and Residuals

Finally, on the evaluation window, ARX the absolute error relative to the baseline as seen in the graph below.

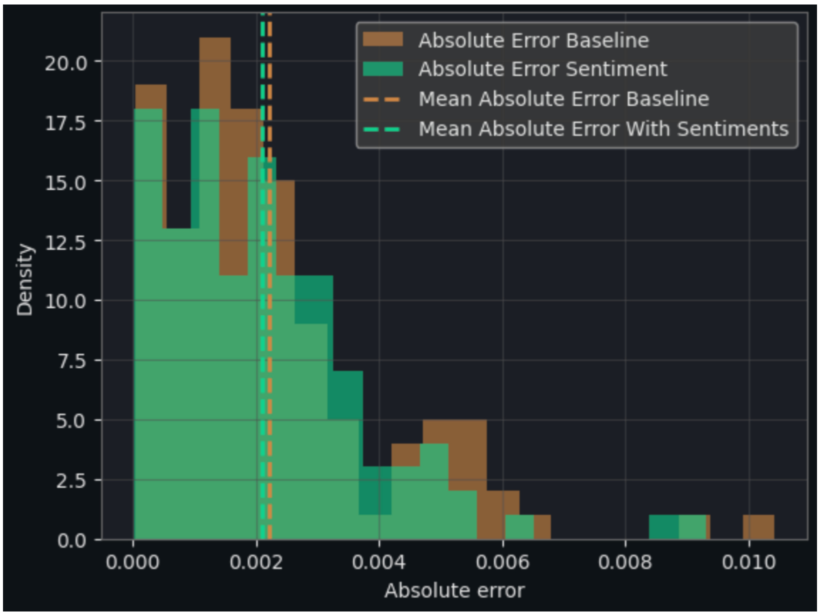

Figure 3 Accuracy Comparison - Baseline vs Model with Sentiment

Economic Interpretation

Our results fit a simple expectations-to-prices story. Media tone moves households' and firms' beliefs first: a high Pos/Neg ratio signals optimism, which lifts planned spending and raises firms' pricing power. Because prices and wages are reset infrequently—via contracts, catalogue/menu updates, procurement cycles—the changes in sentiment translates to changes in CPI with a lag of 6 months. The opposite holds for pessimistic tone where more negative coverage tightens demand and pricing power, dampening subsequent CPI changes.

This mechanism is strongest around turning points, when traditional indicators are stale and surveys under-react. Sentiment doesn't replace fundamentals (energy, rents, FX) but provides a forward looking signal that improves short-run inflation forecasts precisely because it captures shifts in expectations before they pass through to prices.

Conclusion

News sentiments provides early, interpretable signals of macro dynamics. In the CPI-U application, the Pos/Neg ratio and Negative-share proxy the economy's tone balance and pass Granger causality tests, with signs and simple correlations consistent with the expectations channel. Incorporating sentiment as exogenous inputs in AR/ARX models improves explanatory power and forecast accuracy, while remaining simple to operationalize.

We recommend integrating these features alongside standard controls (e.g., energy prices, labor indicators) and validating in rolling out-of-sample backtests. Extensions are straightforward: combine sentiment with additional macro controls in VARX, monitor distributional drift as part of the model-risk lifecycle.

Disclaimer: This material is provided solely for informational and educational purposes by Yukka.ai. It does not constitute investment advice, a recommendation, or an offer or solicitation to transact in any security, financial instrument, or strategy. Analyses and outputs generated by Yukka.ai may incorporate assumptions, algorithms, and third-party data believed to be reliable. They are provided "as is," without warranty as to accuracy, completeness, or suitability, and may change without notice. Past performance, simulated results, or model outputs are not indicative of future outcomes. All investments involve risk, including the potential loss of principal. Any forward-looking statements are inherently uncertain and may differ materially from actual results. Yukka.ai is not acting as a fiduciary and does not provide legal, regulatory, accounting, or tax advice. Users remain solely responsible for their investment decisions. Yukka.ai assumes no obligation to update this information or liability for its use.

References

Granger, C. W. J. (1969). Investigating Causal Relations by Econometric Models and Cross-spectral Methods. *Econometrica*, *37*(3), 424. https://doi.org/10.2307/1912791

Kalamara, E., Turrell, A., Redl, C., Kapetanios, G., & Kapadia, S. (2022). Making text count: Economic forecasting using newspaper text. *Journal of Applied Econometrics*, *37*(5), 896–919. https://doi.org/10.1002/jae.2907

Shapiro, A. H., Sudhof, M., & Wilson, D. J. (2022). Measuring news sentiment. *Journal of Econometrics*, *228*(2), 221–243. https://doi.org/10.1016/j.jeconom.2020.07.053

U.S. Bureau of Labor Statistics. (n.d.). *Consumer Price Index for All Urban Consumers (CPI-U)* [Data set]. https://www.bls.gov/regions/mid-atlantic/data/consumerpriceindexhistorical_us_table.htm

Appendices

Appendix 1: Sentiment variables used

Feature	Formula (math)	Pandas expression	Notes
tot	Pos + Neut + Neg	df["tot"] = df[["Pos","Neut","Neg"]].sum(axis=1)	Monthly total article count (all polarities).
net	(Pos − Neg) / tot	df["net"] = (df["Pos"] - df["Neg"]) / df["tot"]	Net sentiment in [-1, 1]; 0 = balanced
pos_share	Pos / tot	df["pos_share"] = df["Pos"] / df["tot"]	Share of positive stories.
neg_share	Neg / tot	df["neg_share"] = df["Neg"] / df["tot"]	Share of negative stories.
ratio_pn	(Pos + 1) / (Neg + 1)	df["ratio_pn"] = (df["Pos"] + 1) / (df["Neg"] + 1)	Pos-to-Neg ratio with +1 add-one smoothing to avoid division by zero.
Sentiment	net	df["Sentiment"] = df["net"]	Alias used in plots and tests.
chg_sent	Sentiment_t − Sentiment_{t−1}	df["chg_sent"] = df["Sentiment"].diff(	Month-on-month change; first row is NaN by construction.

Appendix 2: Regression Output

```
Results: Ordinary least squares
==================================================================
Model:              OLS              Adj. R-squared:     0.493
Dependent Variable: CPI_delta        AIC:                -756.8831
Date:               2025-08-06 15:05 BIC:                -747.1598
No. Observations:   84               Log-Likelihood:     382.44
Df Model:           3                F-statistic:        27.93
Df Residuals:       80               Prob (F-statistic): 1.85e-12
R-squared:          0.512            Scale:              6.8255e-06
-------------------------------------------------------------------
                  Coef.   Std.Err.     t     P>|t|    [0.025  0.975]
-------------------------------------------------------------------
const            -0.0065    0.0085  -0.7622  0.4482  -0.0234  0.0105
ratio_pn_L6       0.0030    0.0016   1.8148  0.0733  -0.0003  0.0062
neg_share_L6      0.0014    0.0324   0.0439  0.9651  -0.0631  0.0659
DeltaCPIL1        0.5270    0.0923   5.7098  0.0000   0.3433  0.7106
------------------------------------------------------------------
Omnibus:               0.887         Durbin-Watson:            1.896
Prob(Omnibus):         0.642         Jarque-Bera (JB):         0.970
Skew:                  -0.159        Prob(JB):                 0.616
Kurtosis:              2.581         Condition No.:            898
==================================================================
Notes:
[1] Standard Errors assume that the covariance matrix of the errors
is correctly specified.
```