CARMA 2024

6th Int. Conf. on Advanced Research Methods and Analytics
Universitat Politècnica de València, València, 2024
DOI: https://doi.org/10.4995/CARMA2024.2024.17441

Read between the headlines: Can news data predict inflation?

Alan Chester Arcin, Ma. Ellysah Joy Guliman, Genna Paola Centeno, Jacqueline Margaux Herbo, Sanjeev Parmanand, Cherrie Mapa

Department of Economic Research, Bangko Sentral ng Pilipinas, Philippines.

How to cite: Arcin, A.; Guliman M.; Centeno G.; Herbo J.; Parmanand S.; Mapa C. 2024. Read between the headlines: Can news data predict inflation? In: 6th International Conference on Advanced Research Methods and Analytics (CARMA 2024). Valencia, 26-28 June 2024. https://doi.org/10.4995/CARMA2024.2024.17441

Abstract

Big data and machine learning applications are increasingly gaining traction in central bank operations. Among others, central banks have tapped big data in their nowcasting exercises. In this study, we construct inflation news indices and examine if such indices can help predict inflation. The indices are developed using lexicon-based sentiment analysis refined using reinforcement learning and supervised machine learning methods, particularly artificial neural networks and long short-term memory models. These indices are then used as additional feature variables in time series models and machine learning models for nowcasting regional and nationwide inflation in the Philippines. We find empirical evidence that our constructed inflation news indices can improve the predictive capability of these forecasting models.

Keywords: *Inflation; news; nowcasting; sentiment analysis; machine learning*

1. Introduction

Big data and machine learning (ML) based applications are increasingly gaining traction in central bank operations. The 2020 survey conducted by the Irving Fisher Committee on Central Bank Statistics of the Bank for International Settlements indicates that 80 percent of the central banks surveyed utilize big data for their operations, up from just 30 percent in 2015. Among various types of applications, the survey showed that central banks have used big data in their nowcasting exercises. This study adds to this existing body of empirical studies by tapping novel data sources for nowcasting regional and nationwide inflation in the Philippines. Inflation nowcasting and forecasting is important for central banks like the Bangko Sentral ng Pilipinas (BSP) given that it operates an inflation targeting framework for monetary policy formulation.

While Beck et al (2023) and Macias and Stelmasiak (2019) used online price data to nowcast inflation, our research examines if textual data such as online news articles could be useful for inflation nowcasting. This study also builds upon an earlier work by Gabriel et al (2020) for regional inflation nowcasting in the Philippines which employed machine learning techniques but did not incorporate big data. We opt to explore online news data as an indicator for inflation trends as major developments that are relevant for inflation are likely captured by news reports.

To answer our research question, we employ a two-step approach. First, we build inflation news indices which provide information on whether there are more reports of increasing/higher inflation vis-à-vis decreasing/lower inflation in the news media. We take a comprehensive approach and explore two methods of text analytics: (a) lexicon-based sentiment analysis refined by reinforcement learning and (b) ML-based approach using neural networks and long short-term memory models. Second, we employ these indices as additional feature variables in time series models and in ML models for inflation nowcasting.

Overall, we find empirical evidence that our indices are useful for nowcasting regional and nationwide inflation in the Philippines. The news-based indices have information content that help improve the forecasting accuracy of models that utilize the said indices as additional feature variables relative to models that did not.

2. Constructing the Inflation News Indices

Our proposed inflation news indices (INIs) harness the power of online news media to capture real-time information on price trends of goods and services and even other factors that could affect overall inflation. This section outlines the steps in the construction of our own lexicon- and ML-based INIs.

2.1. Data sources and annotation

News articles for this study are sourced from media outlets in the Philippines that allowed us to web scrape the data from their websites' business, finance, or economy sections.[1] Our news database contains articles from different time periods but the common period when we have data from all media outlets is from January 2018 to present. We further filtered our data to news articles containing any of the words: "inflation", "price" and "prices". A total of 3,000 randomly selected articles are annotated. These are manually labeled as 1, -1, or 0, indicating increase, decrease or no change in inflation, respectively. The annotation of sample sentences is a crucial first step to generate and to implement further enhancements of the lexicon- and ML-based INIs.

[1] These media outlets include Business Mirror, Business World, Manila Bulletin, Manila Standard and Philippine Daily Inquirer.

2.2. Lexicon with reinforcement learning method

This method involves the creation of a set of words and usage of predefined set of rules. In this study, the words are grouped to indicate *increase* or *decrease* in inflation. The negation rule is used but instead tweaked to indicate *no change*.[2]

Following Church and Hanks (1990), pointwise mutual information (PMI) was leveraged to identify which words are classified as *increase*/*decrease*. The overall score $Score_{PMI}$ for a given word w is the difference between PMI score of the word with respect to *increase* and *decrease.* Words having $Score_{PMI}$> 0 and $Score_{PMI}$< 0 are therefore assigned under *increase* and *decreas*e wordlists, respectively.

The resulting wordlists is manually evaluated for suitability and reinforcement learning is used to further refine the wordlists. In particular, we will use Q-learning (a type of reinforcement learning (Watkins and Dayan, 1992)) to retrieve the optimal set of words under *increase* and *decrease* that maximizes the Macro-F1 score of the test set within a reasonable amount of time.

2.3. ML-based method using neural networks and long short-term memory models

This method relies on supervised ML models trained using our manually annotated sentences, The sentences are transformed using Word2Vec and are fed into artificial neural networks (ANNs) such as multilayer perceptron (MLP) and Bidirectional Long Short-Term Memory (BiLSTM) models. We test several model configurations and explore the use of Synthetic Minority Oversampling Technique (SMOTE) (Chawla et al, 2002) to create synthetic entries and remedy data scarcity concerns. For the purposes of this study, all ANNs are designed to classify three output classes: *increase*, *decrease*, and *no change.* We then compare the Macro-F1 scores of the different ANNs.

2.4. Index construction and evaluation

To construct the INIs, we evaluate the lexicon and trained models based on their accuracy and Macro-F1 scores. Table 1 presents a summary of evaluation of different lexicons and ML models. The results suggest that INIs derived using lexicon refined by reinforcement learning (hereafter referred to as INI-RL) and using BiLSTM (32) (hereafter referred to as INI-BiLSTM) have the best accuracy and Macro-F1 scores are therefore preferred methods for index construction. These are then used to score the sentences in the news articles with scores aggregated per article and indexed for the month. Afterwards, the INIs are normalized to a mean of 100. Thus, any estimate above 100 suggests that relative to average, news about increasing

[2] Negation rule applies when a word is preceded by a negation word (e.g., “not” or “never”)

inflation outnumber reports of declining inflation while any estimate below 100 suggests otherwise.

Table 1. INI methods: Evaluation metrics for test set. Note: Number in parenthesis for ANNs refers to number of hidden units. Source: Authors' estimates

	Accuracy (In percent)	**Macro-F1 score (In percent)**
Initial lexicon	64.3	41.8
Lexicon with reinforcement learning	89.3	69.3
MLP	68.9	45.7
BiLSTM (16)	78.4	56.2
BiLSTM (32)	79.8	57.9

We find that the INI-RL and INI-BiLSTM tend to move together with correlation coefficient of 0.93. This provides assurance in our measurement of inflation data contained in news as two different methodologies tend to yield similar results.

3. Forecasting inflation using news data

The key research question in our research is to determine the usefulness of news data in nowcasting inflation in the Philippines. Following Gabriel et al (2020), we first estimate time series (TS) model and ML models without the INIs. We compare the forecast accuracy of these models based on their mean absolute error (MAE). We then augment these models with the INIs and perform another round of forecast evaluation to determine if the inclusion of INIs leads to improved forecasting capability.

In this study, we employ a TS and ML models for regression tasks, namely: Support Vector Machines (SVM), Gradient Boosted Machines (GBM) and Extremely Randomized Trees (EXT).

SVM (Cortes and Vapnik, 1995) finds a hyperplane that separates various classes while maximizing the distance between these classes. It uses kernel functions to map data into higher dimensions to expedite separation of classes. Well-known for solving complex problems, SVM performs well in small datasets but scales poorly.

GBM (Friedman, 2001) is an algorithm known for its speed and accuracy, with large datasets in particular. Popular in machine learning competitions, it is an ensemble model, initially starting and combining multiple weak decision trees. Sequentially, subsequent models improve on the previous ones and are evaluated on the loss function of the ensemble.

EXT (Geurts et al, 2006), also known as extra-trees, are similar to random forests (RF) where it fits a number of random decision trees. The main difference is ERT uses the entire dataset to generate trees while RF selects from different variations of data via bagging. This reduces variance and improves computational speed relative to RF.

To match the sample period for our INIs, we restrict all estimations from January 2018 to December 2023 with a train-test split of 80:20. This corresponds to January 2018 to September 2022 and October 2022 to December 2023 for the train and test sets, respectively. For all these models, the target variable is the year-on-year inflation rate.[3] We build individual models for each of the 16 regions and a separate model for nationwide inflation using the different modeling techniques. For our baseline models, we use the following as feature variables: autoregressive component (lag order of 2), previous month's month-on-month inflation and percentage change in year-on-year inflation rate per month to retrieve information embedded in the historical time series of inflation itself. Following Gabriel et al (2020), a shock variable *S* defined as the difference of the year-on-year inflation rate between the current month and the previous month is introduced as an additional variable in the baseline models. This is defined in Equation 1 below:

$$S_t = \begin{cases} 1, if\ YoY\pi_t > YoY\pi_{t-1} \\ 0, \qquad otherwise \end{cases} \tag{1}$$

For the models with INIs, we add another shock variable S^{INI} defined in Equation 2:

$$S_t^{INI} = \begin{cases} 1,\ if\ INI_t > INI_{t-1} \\ -1, if\ INI_t < INI_{t-1} \\ 0, \qquad otherwise \end{cases} \tag{2}$$

where INI_t refers to either INI-RL or INI-BiLSTM as a shock variable at month *t*.

4. Analysis of Results

This section presents the results of our inflation nowcasting exercises. As we are concerned mainly with out of sample forecasting performance, we only present the MAE from the test set. Table 2 displays the results for the baseline models. Similar to the findings of Gabriel et al (2020), we find evidence that the SVM outperforms the TS as it registers lower MAE across all regions. However, the GBM and EXT fail in this aspect. Given the weaker forecasting performance of GBT and EXT, we opt to no longer include these models in the assessment of models that will be augmented with the INIs.

[3] Data on regional and nationwide inflation for the Philippines are sourced from the Philippines Statistics Authority.

Table 2. Forecast Evaluation: MAE for test set (October 2022 – December 2023). Source: Authors' estimates.

	TS	**SVM**	**GBT**	**EXT**
NCR	0.49	0.47	0.85	0.82
CAR	0.62	0.40	0.55	0.62
R1	0.62	0.35	0.88	0.61
R2	0.60	0.46	0.61	0.68
R3	0.66	0.40	1.83	1.14
R4A	0.57	0.41	0.74	0.89
R4B	0.57	0.40	0.78	0.75
R5	0.64	0.43	0.50	0.45
R6	0.78	0.59	1.35	1.25
R7	0.44	0.48	0.47	0.52
R8	0.58	0.37	0.87	0.63
R9	0.98	0.72	0.76	0.55
R10	0.69	0.37	0.62	0.61
R11	0.59	0.40	0.69	0.68
R12	0.55	0.34	0.51	0.34
BARMM	0.68	0.37	0.90	0.44
R13	0.48	0.36	0.44	0.38
Philippines	0.44	0.34	0.85	0.65
Average	0.61	0.43	0.79	0.67

Table 3 shows the forecast evaluation of the TS and SVM models that are augmented with INI-RL and INI-BiLSTM. The results show that the models augmented with new-based indices have registered lower forecasting errors compared to baseline models. This holds true for either the INI-RL or the INI-BiLSTM. This provides empirical evidence that the INIs contain information that can help nowcast regional and nationwide inflation in the Philippines.

Table 3. Forecast evaluation: MAE for test set (October 2022 – September 2023). Source: Authors' estimates.

	Base SVM	SVM+INI-RL	SVM+INI-BiLSTM	TS	TS+INI-RL	TS+INI-BiLSTM
NCR	0.31	0.23	0.31	0.25	0.25	0.27
CAR	0.29	0.03	0.22	0.26	0.26	0.24
R1	0.35	0.31	0.13	0.31	0.26	0.28
R2	0.32	0.28	0.28	0.33	0.30	0.29
R3	1.24	0.13	0.10	0.34	0.32	0.31
R4A	0.33	0.23	0.26	0.34	0.27	0.28
R4B	0.40	0.07	0.45	0.34	0.27	0.25
R5	0.32	0.20	0.23	0.30	0.27	0.30
R6	0.44	0.37	0.36	0.35	0.36	0.34
R7	0.41	0.38	0.32	0.38	0.41	0.38
R8	0.34	0.29	0.31	0.27	0.24	0.24
R9	0.56	0.47	0.52	0.53	0.47	0.47
R10	0.22	0.25	0.21	0.24	0.23	0.23
R11	0.30	0.30	0.29	0.36	0.34	0.33
R12	0.35	0.29	0.36	0.36	0.38	0.38
BARMM	0.43	0.20	0.16	0.34	0.21	0.22
R13	0.19	0.34	0.22	0.20	0.24	0.22
Philippines	0.34	0.22	0.05	0.25	0.25	0.24
Average	0.40	0.26	0.27	0.32	0.29	0.29

5. Conclusion and Future Directions

Can news data help predict inflation? The short answer to this research question is yes. In this study, we provide empirical evidence that the inclusion of INIs can help improve the predictive capability of time series and SVM models for regional and nationwide inflation nowcasting in the Philippines. Our news-based indices have information content that are relevant for monitoring inflation trends. The models that we have built in this study can serve as complementary models for the suite of forecasting models maintained by the BSP.

Our findings contribute to the existing literature on the relevance of big data for nowcasting macroeconomic variables. Future plans on this exercise include the addition of other ML models for generating regional forecasts and combining multiple ML models as a ensemble of its own as another model for prediction.

References

Beck, G. W., Carstensen, K., Menz, J. O., Schnorrenberger, R., & Wieland, E. (2023). Nowcasting consumer price inflation using high-frequency scanner data: Evidence from Germany (No. 34/2023). Deutsche Bundesbank Discussion Paper.

Chawla, N. V., Bowyer, K. W., Hall, L. O., & Kegelmeyer, W. P. (2002). SMOTE: synthetic minority over-sampling technique. Journal of artificial intelligence research, 16, 321-357.

Church, K., & Hanks, P. (1990). Word association norms, mutual information, and lexicography. Computational linguistics, 16(1), 22-29.

Cortes, C., & Vapnik, V. (1995). Support-vector networks. Machine learning, 20, 273-297.

Friedman, J. H. (2001). Greedy function approximation: a gradient boosting machine. Annals of statistics, 1189-1232.

Gabriel M., Bautista D., & Mapa C. (2020). Forecasting regional inflation in the Philippines using machine learning techniques: A new approach. Bangko Sentral ng Pilipinas Working Paper Series No. 2020-10.

Geurts, P., Ernst, D., & Wehenkel, L. (2006). Extremely randomized trees. Machine learning, 63, 3-42.

Macias, P., & Stelmasiak, D. (2019). Food inflation nowcasting with web scraped data (p. 302). Warsaw: Narodowy Bank Polski, Education & Publishing Department.

Mahadevaswamy, U. B., & Swathi, P. (2023). Sentiment analysis using bidirectional LSTM network. Procedia Computer Science, 218, 45-56.

Mikolov, T., Sutskever, I., Chen, K., Corrado, G. S., & Dean, J. (2013). Distributed representations of words and phrases and their compositionality. Advances in neural information processing systems, 26.

Serena, J. M., Tissot, B., Doerr, S., & Gambacorta, L. (2021). Use of big data sources and applications at central banks (No. 13). Bank for International Settlements.

Watkins, C. J., & Dayan, P. (1992). Q-learning. Machine learning, 8, 279-292.