Mixture Models for Predictive Densities: A Bayesian Synthesis of Survey-Based Inflation Expectations in Colombia*

Received: January 15, 2026 – Accepted: February 04, 2026

Doi: https://doi.org/10.12804/revistas.urosario.edu.co/economia/a.16206

Héctor Manuel Zárate-Solano†

Abstract

Understanding inflation expectations is central to monetary policy analysis, particularly in emerging economies, where uncertainty and structural change can amplify forecasting challenges. Surveys of professional forecasters increasingly rely on histogram-based density forecasts to capture not only point predictions but also the full distribution of expected inflation. Aggregating these density forecasts provides a richer characterization of the shape, dispersion, and uncertainty embedded in expert expectations. This study analyzes quarterly inflation expectations reported by Colombian professional forecasters from January 2016 to October 2023. We construct measures of disagreement and uncertainty at both the individual and aggregate levels, documenting their evolution through periods of macroeconomic stability, policy shifts, and the inflationary surge following the COVID-19 pandemic. We synthesize individual predictive densities using the Bayesian predictive synthesis (BPS) framework with time-varying mixture weights that adapt to changes in relative forecast performance over time.

* The views expressed herein are those of the author and not necessarily those of the Banco de la República (Central Bank of Colombia), nor its Board of Directors.

† Senior econometrician at Banco de la República de Colombia (Central Bank of Colombia). Email: hzaratso@banrep.gov.co

Para citar este artículo: Zárate-Solano, H. (2026). Mixture Models for Predictive Densities: A Bayesian Synthesis of Survey-Based Inflation Expectations in Colombia. *Revista de Economía del Rosario, 29*(1), 1-23. https://doi.org/10.12804/revistas.urosario.edu.co/economia/a.16206

Keywords: forecast density combination; Bayesian predictive synthesis; surveys of professional forecasters; symbolic data.
JEL Classification: C11, C32.

Modelos de mezcla para densidades predictivas: una síntesis bayesiana de las expectativas de inflación basadas en encuestas en Colombia

Resumen

Comprender las expectativas de inflación es fundamental para el análisis de la política monetaria, especialmente en las economías emergentes, donde la incertidumbre y los cambios estructurales pueden agravar los retos que plantea la elaboración de pronósticos. Las encuestas a pronosticadores profesionales recurren cada vez más a pronósticos de densidad basados en histogramas para captar no solo las predicciones puntuales, sino también la distribución completa de la inflación esperada. La agregación de estas previsiones de densidad permite una caracterización más detallada de la forma, la dispersión y la incertidumbre implícitas en las expectativas de los expertos. Este estudio analiza las expectativas de inflación trimestrales reportadas por pronosticadores profesionales colombianos desde enero de 2016 hasta octubre de 2023. Construimos medidas de desacuerdo e incertidumbre tanto a nivel individual como agregado, documentando su evolución a lo largo de periodos de estabilidad macroeconómica, cambios en las políticas y el repunte inflacionario tras la pandemia de COVID-19. Sintetizamos las densidades predictivas individuales utilizando el marco de síntesis predictiva bayesiana (BPS) con pesos de mezcla que varían en el tiempo y se adaptan a los cambios en el desempeño relativo de los pronósticos a lo largo del tiempo.

Palabras clave: combinación de densidades de pronóstico; síntesis predictiva bayesiana; encuestas a pronosticadores profesionales; datos simbólicos.
Clasificación JEL: C11, C32.

Modelos de mistura para densidades preditivas: uma síntese bayesiana das expectativas de inflação obtidas em pesquisas com especialistas em previsão na Colômbia

Resumo

Compreender as expectativas de inflação é fundamental para a análise da política monetária, especialmente em economias emergentes, nas quais a incerteza e as mudanças estruturais podem ampliar os desafios associados à elaboração de previsões. Os levantamentos realizados com especialistas em previsão recorrem cada vez mais a previsões de densidade baseadas em histogramas para captar não apenas previsões pontuais, mas também toda a distribuição da inflação esperada. A agregação dessas previsões de densidade permite caracterizar de forma mais abrangente a forma, a dispersão e a incerteza implícitas nas expectativas dos especialistas. Este estudo analisa as expectativas trimestrais de inflação informadas por especialistas em previsão na Colômbia entre janeiro de 2016 e outubro de 2023. Foram construídas medidas de divergência e de incerteza, tanto no nível individual quanto no agregado, documentando sua evolução ao longo de períodos de estabilidade macroeconômica, mudanças de política econômica e da alta inflacionária observada após a pandemia da COVID-19. As densidades preditivas individuais são sintetizadas por meio do arcabouço da síntese preditiva bayesiana (BPS), com pesos de mistura variantes no tempo, que se adaptam às mudanças no desempenho relativo das previsões ao longo do período analisado.

Palavras-chave: combinação de densidades preditivas; síntese preditiva bayesiana; pesquisas com especialistas em previsão; dados simbólicos.
Classificação JEL: C11, C32.

Introduction

The inflation-forecast-targeting framework for monetary policy is anchored in a long-term inflation objective. Achieving this objective requires the central bank to produce and continuously evaluate inflation forecasts while managing the stabilization trade-off between deviations of inflation from its target and fluctuations in real activity (often summarized by the output gap).

In Colombia, the central bank targets 3 % inflation and adjusts the policy interest rate in response to the projected path of inflation and economic slack. These decisions are inherently made under uncertainty and are influenced by measures of expected inflation reported by professional forecasters. In empirical work, forecast uncertainty is typically quantified using statistical and econometric methods that summarize both the dispersion of beliefs and the uncertainty surrounding the predictive distribution.

A widely used approach to evaluating *ex ante* uncertainty relies on survey-based expectations that elicit both point and density forecasts. These surveys capture professional forecasters' subjective assessments by collecting probability distributions of inflation across multiple horizons. By design, survey instruments harness informed judgment as a direct measure of perceived uncertainty about inflation and related macroeconomic conditions. The reported distributions therefore reflect each expert's stated beliefs, conditional on their information set.

In this paper, the recent availability of detailed forecaster beliefs allows us to construct histogram time series (HTS) from survey-based density forecasts. These distributions provide the basis for measuring uncertainty and risk balance at the individual level and then aggregating across respondents. We quantify individual uncertainty using the width (dispersion) of each reported histogram, and we infer risk balance from the distribution's asymmetry, which indicates whether perceived risks are skewed toward higher or lower future inflation. Aggregating the individual histograms yields summary measures that capture the collective information embedded in the survey panel.

Combining individual forecast densities also enables the construction of a unified consensus measure. Two considerations are central to selecting an appropriate combination scheme: (i) the weighting structure that determines how individual densities contribute to the aggregate, and (ii) the extent of model incompleteness, which determines how many and which densities should be incorporated to mitigate misspecification and omitted-information risks.

The general density-combination framework offers a flexible way to merge forecasts from both frequentist and Bayesian perspectives, allowing for time-varying weights, adaptive learning, and explicit adjustments for model incompleteness. Bayesian predictive synthesis (BPS), developed by McAlinn and West (2019) and rooted in the agent-opinion analysis of West and Crosse (1992), combines multiple predictive distributions by treating each agent's forecast draws as latent states within a dynamic linear regression model.

This structure enables decision-makers to update and synthesize expert predictive densities into a posterior distribution for inflation, thereby improving robustness, calibration, and forecast accuracy. This paper makes two main contributions. First, we conduct a static, exploratory comparative analysis using symbolic histogram data, summarizing group-level behavior with aggregate histograms that characterize uncertainty and the risk balance. Second, we combine multiple predictive distributions of inflation expectations using the dynamic BPS framework, which accommodates time-varying relationships, corrects biases and miscalibration, and accounts for cross-forecaster dependence, thereby addressing model incompleteness and improving predictive performance.

Forecast uncertainty is commonly quantified using statistical measures and econometric models. A complementary approach to evaluating ex-ante uncertainty relies on survey expectations that elicit both point and density forecasts. In these surveys, professional forecasters use their judgment to report probability distributions for future inflation across multiple horizons. The primary objective is to extract expectations about inflation and economic activity by leveraging informed human judgment as a direct measure of perceived uncertainty. The resulting densities can be interpreted as subjective probabilities—each forecaster's degree of belief conditional on their information set—and provide a natural basis for models that reconcile disagreement, dependence, and potential incoherence across individual forecasts.

In this paper, the recent availability of detailed forecaster beliefs allows us to construct HTS as symbolic data. These density forecasts are used to compute individual-level measures of uncertainty and risk balance, which are then aggregated across respondents to characterize group-level expectations. By providing a full probability distribution over future inflation outcomes, these uncertainty measures enrich forecast assessment and can be incorporated to improve economic forecasting and risk evaluation.

Individual perceived uncertainty is measured by the width (dispersion) of each reported histogram. The balance of risks is assessed using the distribution's asymmetry and spread, indicating whether probability mass is skewed toward higher- or lower-than-expected inflation outcomes. Aggregating individual histograms yields group-level measures that summarize the "wisdom of the crowd."

Combining experts' forecast densities also facilitates the construction of a unified consensus measure. Designing an effective combination scheme hinges on two issues: (i) the weighting structure that governs how individual densities contribute to the aggregate, and (ii) model incompleteness, which concerns how many and which densities to include to mitigate misspecification and omitted-information risks.

The *fundamental density combination equation* is a general framework for combining forecasts from both frequentist and Bayesian methods. It includes optimized time-varying weights, learning features, and mechanisms to address model incompleteness. The BPS approach combines multiple predictive distributions using agent/expert opinion analysis, as discussed in McAlinn and West (2019), providing a comprehensive approach to density forecasting methods that include longer-term inflation expectations.

The BPS framework provides a method for combining forecast densities, building on earlier research by West and Crosse (1992) in agent opinion analysis and formalized by McAlinn and West (2019). The analysis unfolds in two steps. First, decision-makers gather forecasted inflation densities from economic experts through surveys that capture these experts' individual assessments. Next, the decision-maker uses this information to predict inflation by merging individual densities into a posterior distribution. The weighting scheme is determined by specifying a functional form known as the synthesis function.

The synthesis function assumes a dynamic linear regression model, treating the draws from the J competing densities as latent variables, denoted as $X_{t+h|t}$. This framework yields gains in robustness and accuracy.

This paper contributes by conducting a static, exploratory comparative analysis using *symbolic histogram data* to summarize group levels through aggregate histograms for measuring uncertainty and risk balance. It also combines multiple predictive distributions of inflation expectations based on the dynamic BPS theory, which adjusts for biases, miscalibrations, and dependencies among forecast models to address model incompleteness.

The remainder of the paper is organized as follows. Section 2 provides background on inflation-forecast targeting. Section 3 describes the survey of professional forecasters. Section 4 presents a static comparative analysis of uncertainty and the risk balance. Section 5 reports the forecast evaluation and BPS results. Section 6 offers concluding remarks.

In summary, Figure 1 provides a visual representation of the flowchart that outlines the structure of this paper.

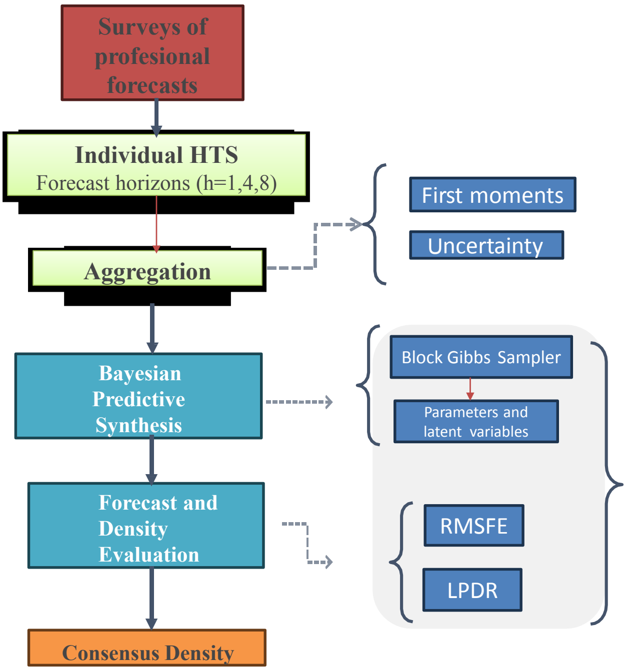

Figure 1. Flowchart of the Paper

Note: The process starts with surveys of professional forecasters (SPF), followed by histogram analysis. Upon completion, Bayesian predictive synthesis is analyzed. Finally, a predictive density is presented. HTS: histogram time series; RMSFE: root mean squared forecast error; LPDR: log predictive density ratio.

Literature Chronology

This section provides an overview of predictive density implementation, drawing on both academic research and practical applications.

First, Lindley et al. (1979) presented a framework for reconciling incoherent probabilities that conflict with the laws of probability. According to Gerrard (2023), the subjective interpretation of probability captures the numerical representation of a respondent's personal confidence level. This personal subjective probability, known as the logic of partial belief, adheres to the laws of probability and is linked to decision-making and action. In this context, West and Crosse (1992) discussed Bayesian updating as a method for reconciling probability assessments and refining an agent's opinion.

On the other hand, Billio et al. (2011) developed a Bayesian multivariate combination framework that incorporates time-varying weights by using a distributional state space representation of predictive densities from various models. The dynamics of the weights are influenced by the historical performance of these predictive densities through learning mechanisms.

This approach is evaluated using both statistical measures and utility-based performance metrics, with a focus on density forecasts for U.S. macroeconomic time series and on surveys of stock market prices.

McAlinn and West (2019) enhanced a Bayesian model and forecast combination by incorporating agent opinion analysis theory. The framework introduced dynamic latent factor models for time-series forecasting with forecast synthesis, using a simulation-based approach for computation. These models are designed to adapt to time-varying biases, miscalibration, and interdependencies among multiple models or forecasters, while also including various methods for pooling forecasts.

Furthermore, Chernis et al. (2023) improved BPS by nonparametrically modeling the synthesis function with regression trees. This innovation led to better forecast accuracy and interoperability in two macroeconomic forecasting studies: Chernis et al. (2025) focused on gross domestic product (GDP) growth forecasts from the euro area's Survey of Professional Forecasters, and Tallman et al. (2024) combined density forecasts of U.S. inflation using various regression models with different predictors.

Diebold et al. (2023) introduced methods for calculating regularized mixtures of density forecasts for Eurozone inflation and real interest rates. The study indicated that specific combinations, such as the simplex and best-average methods, outperformed other approaches in terms of forecast accuracy. Additionally, McAlinn and West (2019) presented tools for analyzing monetary policy transmission channels, allowing an examination of the degree to which inflation expectations are anchored.

Clements and Tracy (2022) provide a detailed analysis of surveys applied to professional forecasters, addressing issues with their structure and data collection methods. The authors discuss the computation of measures such as disagreement and uncertainty at both aggregate and individual levels. They examine how point forecasts align with central tendency measures obtained from density forecasts, evaluate density forecast coverage, and address issues such as forecast rounding. Additionally, they compare forecast accuracy, density forecast coverage, and forecast rounding across respondents to investigate heterogeneity and persistence in forecasting behavior.

Tallman and West (2024) recently extended BPS by proposing Bayesian predictive decision synthesis (BPDS), a framework that goes beyond traditional predictive modeling to explicitly integrate decision-analytic objectives. BPDS links predictive accuracy to decision relevance by evaluating and combining models based on their implications for downstream actions, rather than solely on their fit to realized outcomes. This perspective is particularly useful in policy settings, where synthesized forecasts of key economic indicators inform intervention design, timing, and resource allocation. When applied to survey-based predictive densities from professional forecasters within a Bayesian updating framework, BPDS provides a coherent basis for model comparison and combination in settings where decision quality is as important as predictive performance.

In the Colombian context, Iregui (2021) analyzes survey data from economic experts on point forecasts of inflation expectations. The study argues that the simple average of individual forecasts is a weak predictor of future inflation. In particular, in a changing macroeconomic environment, forecasters may adjust their expectations sluggishly to incoming information, so pooled averages can inherit biases and persistence that render them stale or systematically inaccurate.

In sum, this section develops a Bayesian framework for combining predictive densities, drawing on key concepts from the time-series forecasting literature. The following section describes the professional forecaster survey data used in the empirical application.

Survey of Professional Forecasters

In Colombia, the Banco de la República (Central Bank of Colombia) conducts a monthly survey to gather assessments from professional forecasters on key macroeconomic and financial variables, as well as on monetary policy. The survey sample includes respondents from banks, labor unions, pension funds, and academic institutions.[1] Since the first quarter of 2016, the Banco de la República has asked respondents to provide point forecasts for inflation

1 The Survey of Professional Forecasters (SPF) has been administered by the Banco de la República since 2001 and targets experts from both financial and non-financial institutions. The questionnaire elicits expectations for inflation, real GDP growth, and the exchange rate across multiple horizons.

and GDP and to assign probabilities to different outcome ranges.[2] This process aims to quantify the uncertainty surrounding inflation and GDP by collecting probabilities across various horizons each quarter. The information gathered through this survey helps policymakers understand market participants' expectations and sentiments while enhancing the transparency and communication of the central bank's monetary policy stance. By analyzing the collected data, the Banco de la República can better assess potential risks to its inflation targets and growth forecasts, enabling more informed decision-making.

Including probabilities for different scenarios helps provide a clearer understanding of potential future economic conditions. This probabilistic approach encourages respondents to think critically about their forecasts and recognize the uncertainties inherent in the macroeconomic environment.

Over the years, the insights from these surveys have been invaluable. They have fostered a more nuanced dialogue between the central bank and the market, leading to better alignment between expectations and actual economic developments. As a result, this initiative not only aids in the formulation of effective monetary policy but also supports the broader goal of financial stability in Colombia.

The survey not only considers inflation and GDP but also captures expectations for other key financial indicators, such as interest rates and the exchange rate, providing a comprehensive view of economic prospects. The Banco de la República continually refines its methodologies and expands the survey's scope to ensure it remains an essential tool for understanding economic dynamics in an ever-changing global landscape.

The Colombian Survey of Professional Forecasters (Col-SPF) requests that experts provide density forecasts in histogram form, using the specific class intervals outlined in the survey instrument. These density forecasts are essential for creating measures of uncertainty and risk balance at the individual level, which can then be extended to the entire sample of respondents. Although survey respondents are given fixed ranges, the bins have changed over time due to unforeseen developments.[3]

2 Survey participants report expectations for inflation, the Banco de la República's intervention rate, the representative market exchange rate (TRM, for its Spanish initials), and real GDP growth at specified horizons. The survey has been conducted monthly since 2002 and has collected histogram-based density forecasts quarterly from 2016.I through 2023.IV.

3 Various events, such as the COVID-19 pandemic and the war in Ukraine, have influenced the macroeconomic environment through different channels. It is evident that the placement of the bins shifts over time in response to the changing dynamics of the economy.

The central bank administers the survey on the same day the Colombian statistical authority releases the Consumer Price Index (CPI). Respondents have about five days to submit their forecasts. After data cleaning, the sample comprises 18 density forecasters who report distributions using nine histogram bins at three forecast horizons. The density-forecast questionnaire format is shown in Figure 2.

Section 3. Probabilities of CPI Inflation

Please indicate what probabilities you would attach to the percentage changes. The probabilities of these alternatives should up to 1.

	Probability of annual percent change CPI		
	Next quarter	Next year	Two years
0-1			
1-1.5			
1.5-2.0			
.			
.			
.			
14.5-15			
Sum	1		

Figure 2. COL-SPF density forecasts for the next quarter of the current year, the following year, and two years later.

Source: https://suameca.banrep.gov.co/estadisticas-economicas/encuestas

On average, the survey includes about 35 respondents per quarter from a panel of roughly 100 participants. Respondents report histogram probabilities over pre-specified intervals, with bin endpoints that may vary over time with the target variable. To ensure comparability across vintages, we convert each histogram into a continuous probability density function (PDF) evaluated on a common fine grid of 720 points (Conflitti et al., 2015). Because forecasters may enter or exit the panel at any time, the dataset is unbalanced and contains missing observations. We therefore assemble a harmonized dataset to construct the longest feasible consistent panel, excluding forecasters who fail to respond for five consecutive rounds. For the remaining sample, we identify nine forecasters with minimal missingness; the few remaining gaps are imputed using draws from an unconditional normal distribution.

Uncertainty reflects respondents' degree of confidence in their predictions and is computed from the dispersion of the probability distribution over potential inflation outcomes at each time horizon. A wider histogram—i.e., a greater spread of probabilities across bins—indicates higher perceived uncertainty, whereas a more concentrated histogram indicates lower uncertainty.

In summary, this task focuses on developing a strategy to update the probability distribution of expected inflation each quarter, accurately capturing the underlying density process. Unlike point estimates, density forecasts offer a significant advantage by avoiding errors associated with normality assumptions. They provide richer information about uncertainty and enable decision-makers to incorporate the full distribution into their strategies, rather than relying solely on a single expected value. This approach promotes more informed and robust decisions in dynamic economic environments.

Static Comparative Analysis

In this section, we conduct an exploratory analysis using symbolic histogram data to summarize group levels through aggregate histograms. This representation enables the construction of informative measures, including uncertainty and risk balance. Billiard and Diday (2020) define symbolic data as histogram-valued observations and develop basic descriptive statistics—such as sample means, variances, covariances, and higher-order moments—under the implicit assumption that observations are uniformly distributed within each interval (and, for histogram data, within each sub-interval).

Statistics from histograms

A histogram time series (HTS) is a sequence of distributions observed at times t =1, …, T, for respondents j = 1, …, J. Each distribution is represented as a histogram h_{xjt}. Thus, the density forecasts provided by experts in the SPF are valuable for identifying a representative mixture of these forecasts. They also help construct measures of uncertainty and bias, which are especially useful during periods of structural change.

There are complications, such as the number of bins used to define the histogram, which has changed over time. We have fixed the number of bins. In accordance with Poncela and Senra (2018), we assume that the probability mass is located at the midpoint of each bin. Therefore, the mean of forecasters' histograms is given by:

$$\mu_{i,t+h|t} = \sum_{k=1}^{K} p_{i,t+h|t}(k) \times mid(k) \qquad (1)$$

Where $p_{i,t+h|t}$ (k) is partial rational belief forecaster i attached to the k - th interval corresponding to an h step-ahead prediction from a survey applied in quarter t. Furthermore, $mid(k)$ is the midpoint of the kth bin.

The variance of every forecaster *h step-ahead* is:

$$\sigma^2_{i,t+h|t} = \sum_{k=1}^{K} p_{i,t+h|t}(k) \times [mid(k) - \mu_{i,t+h|t}]^2 \qquad (2)$$

Uncertainty reflects the confidence associated with a prediction and is determined by the outcomes across different percentiles of the forecast density. An increase in the spread of the predictive probabilities indicates higher uncertainty. In particular, the interquartile range is an indicator of uncertainty for each forecaster,

$$p^{IQR}_{i,t+h|t} = p^{0.75}_{i,t+h|t} - p^{0.25}_{i,t+h|t} \qquad (3)$$

Where $p^{0.75}_{i,t+h|t}$ and $p^{0.75}_{i,t+h|t}$ are the 0.75th and 0.25th percentiles of the density forecast for each respondent in the survey. The uncertainty measure based on histograms is *ex-ante*, prior to observing inflation.

Additionally, the moment skewness is a vital measure of distribution shape and is determined by calculating the third moment of the mean.

$$a = \frac{1}{2}\gamma = \frac{\mu_3}{2\sigma^2} \qquad (4)$$

Where γ represents the skewness used to measure the symmetry of the expected inflation distribution.

Comparative Analysis

This section presents a comparative analysis to assess changes in uncertainty and potential bias across three time periods. Accordingly, the analysis of economic expectations is organized into the following three periods:

1. Pre-COVID-19 period (2015.I to 2019.IV), marked by migration pressures and exchange rate fluctuations.
2. The COVID period (2020.I to 2021.II).
3. Post-COVID-19 period (2021.III to 2023.IV), characterized by large shocks to supply and demand and by changes in firms' pricing and profit dynamics.

The overlapping histograms in Figure 3 show a significant increase in the expected annual inflation over a four-step-ahead forecast horizon. Before the COVID-19 pandemic, respondents anticipated 4 % inflation, but this figure rose to 10 % afterward.

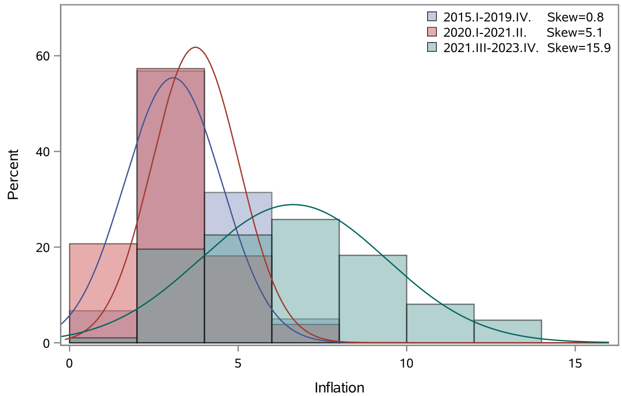

Figure 3. Histograms for four-step-ahead expected inflation forecasts

Note: Expected annual inflation projections are shown as overlapping aggregated histograms across three time periods. The histograms indicate a shift toward higher values for anxious respondents during the COVID-19 period.

Source: COL-SPF, Banco de la República.

Similarly, Figure 4 indicates that inflation uncertainty has been rising over time, as pessimistic forecasts have shifted to the right.

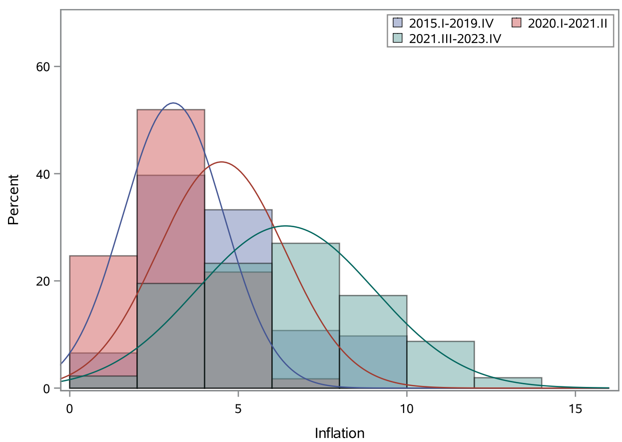

Figure 4. Histogram for four-step-ahead expected inflation forecasts for pessimists.

Note: Expected annual inflation for the forecasting horizons, shown as overlapping aggregated histograms across the three study periods.

Source: COL-SPF, Banco de la República.

Figure 4 shows a marked shift in the distribution of inflation expectations over time, particularly among respondents who express greater concern about economic conditions. The distribution widens, indicating greater dispersion and, consequently, higher perceived uncertainty. At the same time, the distribution shifts toward higher values, suggesting that these respondents increasingly assign greater probability to elevated inflation outcomes in the near term.

The evolution of inflation expectations among more anxious respondents reflects a combination of recent macroeconomic signals, media narratives, and personal financial experiences. These psychological and informational channels are important because they can shape expectations, influence household and firm behavior, and affect broader economic sentiment and inflation dynamics.

The patterns in Figure 5 underscore the value of monitoring shifts in the distribution of expectations, which may signal evolving perceptions of inflation risk and emerging inflationary pressures. For policymakers and economists, understanding the beliefs of more anxious respondents can be informative about expectation dynamics and potential implications for future decisions.

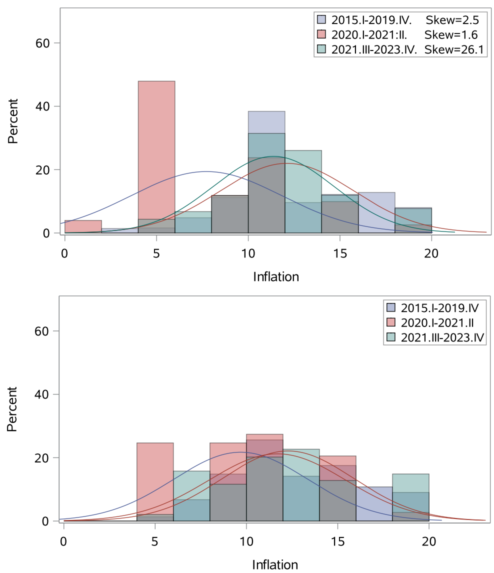

A.

B.

Figure 5. Comparison of expectation distributions between respondents and pessimists. A. Histograms for one-step-ahead expected inflation forecasts. B. Histograms for one-step-ahead expected inflation forecasts for pessimists.

Note: A pessimistic expected annual inflation for forecasting horizons, shown as overlapping aggregated histograms across the three study periods.

Source: COL-SPF, Banco de la República.

Figure 6 reports the interquartile range (IQR) across the three analyzed subperiods, indicating that uncertainty, measured by dispersion, peaks during the COVID-19 period for four-step-ahead forecasts for inflation expectations.

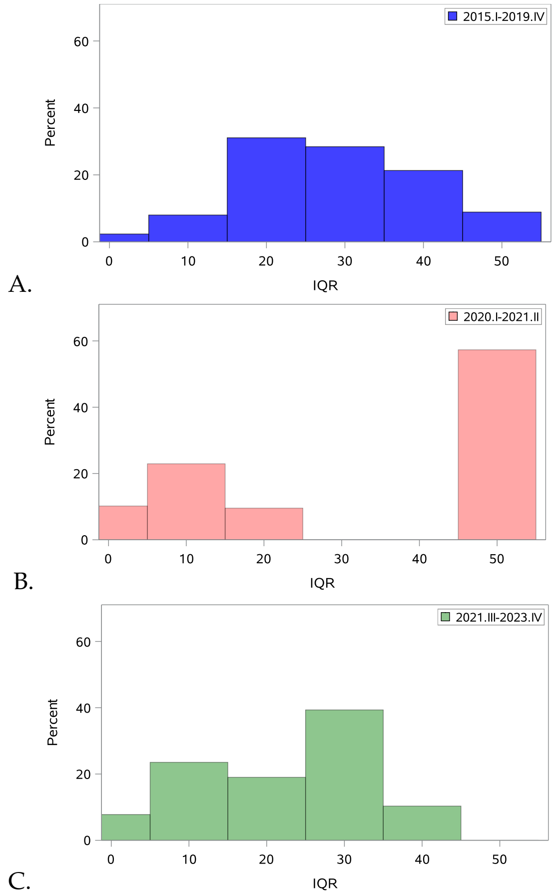

Figure 6. Interquartile ranges for four-step-ahead expected inflation forecasts across the three study periods. A. 2015.I–2019.IV. B. 2020.I–2021.II. C. 2022.III–2023.IV.

Source: Col-SPF, Banco de la República.

Bayesian Predictive Synthesis

This section combines multiple inflation-expectation predictive distributions via dynamic BPS, allowing time-varying relations and correcting bias, miscalibrations, and cross-model dependencies under model incompleteness.

The fundamental density combination equation is a mixture density, where the mixing density is the joint density of predictive densities from individual models, addressing the issues of weighting and incompleteness.

This combination equation is given by:

$$p(y_t \big| I) = \int (p(y_t, \tilde{y}_t) \,|I)\; d\,\tilde{y}_t = \int (p(y_t|\tilde{y}_t)p(\tilde{y}_t|I)d\tilde{y}_t \qquad (5)$$

Where $\tilde{y}_{1t}, \ldots \tilde{y}_{jt}$ are the predicted values from $J = 1, \ldots J$ economic experts. These represent latent variables that relate to inflation, derived from individual density forecasts, $p(y_t|I_i)$.

The fundamental combination density is given by $p(y_t, \tilde{y}_{1t}|I)$.

The BPS framework provides a Bayesian method for combining forecast densities based on analysis of expert opinions (for more details, see West & Crosse [1992]). A decision maker D receives N expert inflation densities π_t from experts A_i, $i = 1 \ldots, N$. These densities reflect each expert's assessment and form the information set $I = p(\pi_t \big| I_1), \ldots p(\pi_t \big| I_N)$. Bayesian analysis assures that D will use the information set I to predict π_t via the posterior $p(\pi_t \big| I_t))$ given by:

$$p(\pi_t|I_t) = \int \alpha(\pi_t|\tilde{\pi}_t) \prod_{i=1}^{N} p(\tilde{\pi}_{it}|I_t)\; d\tilde{\pi}_t \qquad (6)$$

Where $\tilde{\pi}_t$ is a latent vector comprising the latent π_{it}; and $\alpha((\pi_t|\tilde{\pi}_t)$ is a conditional PDF for π_t given $\tilde{\pi}_t$.

A predictive distribution is formed by combining expert forecasts via a multivariate regression that links observed inflation to these predictors. The synthesis process involves two steps using Markov Chain Monte Carlo (MCMC): first, we generate samples from the individual forecasts; second, we use these samples as explanatory variables in a time series model, with annual inflation as the dependent variable, to obtain the posterior conditional distribution. Assuming D is a decision maker with access to individual histograms, the dynamic linear regression framework, using the J experts' draws as generated regressors $X_{t+h|t}$, frames the synthesis function in the following form:

$$\alpha(\pi_{t+h}|X_{t+h|t}, \Psi_{t+h}) \approx \mathrm{N}(\pi_{t+h}|X_{t+h|t}, \sigma^2_{t+h}) \qquad (7)$$

Where Ψ_{t+h} is the latent states formed by the dynamic regression coefficients and variances; $X_{t+h|t}$ is the latent agent states or draws generated from the distributions. A dynamical setting structures sequential learning in a state-space representation, allowing the decision maker at time t to update biases and interdependencies among experts based on the received information. A Bayesian decision maker, D, predicts annual inflation sequentially, receiving predictive densities J from economic experts: A_j, $j = 1, \ldots J$. The information set for prediction is $I_{t\text{-}1} = \{\pi_{1:t\text{-}1}, H_{1:t\text{-}1}\}$, which is used to compute $p(\pi_t | I_t)$.

As data accumulates, D learns about relations among experts, their predictions, and dependency characteristics, combining forecast distributions in a Bayesian fashion with the following equation:

$$P(\pi_t \middle| \Phi_t, H_t) = \int \alpha(\pi_t \middle| \tilde{\pi}_t, \Phi_t) \prod_{j=1}^{J} h_{tj}(\tilde{\pi}_{jt}) \, d\, \tilde{\pi}_{tj} \qquad (8)$$

Where $\alpha(\pi_t | \tilde{\pi}_t, \Phi_t) \sim \mathrm{N}(\pi_t | F_t \theta_t, \upsilon_t)$ means that the calibration function follows a normal distribution for π_t. with $F_t = (1, \tilde{\pi}')'$ and $\theta_t = (\theta_{t0}, \ldots, \theta_{tj})$. On the other side, the latent variables are computed as $P(\tilde{\pi}_t | \Phi_t, \pi_t, H_t) = \prod_{j=1:J} h_{tj}(\pi_{tj})$.

Following the work of McAlinn and West (2019), we implemented an MCMC simulation algorithm for the BPS approach. This involves using a two-component block Gibbs sampler to handle the latent agent states and the parameter dynamics, denoted as θ.

In the initial step, the MCMC algorithm iterates repeatedly to resample conditional posteriors, generating samples from the target posterior distribution: $p(\tilde{\pi}_{1:T}, \Phi_{1:T} | \pi_{1:T}, H_{1:T}$, where the sampling parameters are conformed by: $\Phi_{1:T} = [\theta_{1:T}, \nu_{1:T}]$.

The BPS specification is based on the dynamic linear model (DLM):

$$\pi_t = F'_t \theta_t + \upsilon_t \sim \mathrm{N}(0, \upsilon_t) \qquad (9)$$

$$\theta_t = G\theta_{t\text{-}1} + \omega_{\mathrm{t}} \sim \mathrm{N}(0, \upsilon_t W_t) \qquad (10)$$

In the second step, we sample the latent agent states, denoted as $\tilde{\pi}_{1:T}$, from the following distribution: $p(\tilde{\pi}_t | \Phi_t, \pi_t, H_t) \propto T_{(gl)} (\pi | F_t, \theta_t, \upsilon_t) \prod_{j=1:J} h_{tj} (\pi_{tj})$.

BPS offers a formula for the distribution of inflation based on agents' histogram forecasts from various experts, while also accounting for model uncertainty. The modeler uses predictive densities provided by experts to forecast inflation at time t, employing a calibrated and synthesized collection of agents.

Forecast Density Evaluation

This section outlines the out-of-sample evaluation of expected inflation for economic experts and their BPS combination. The forecasting evaluation uses two main diagnostics. First, we assess forecast accuracy by computing the root mean squared forecast error (RMSFE) over the relevant h-step ahead forecast horizons. We compare the point forecasts provided by respondents to those generated by the BPS. Second, we evaluate the log predictive density ratio.

The log predictive density ratios (LPDR) at horizon h for time indices t is defined as:

$$LPDR_{1:t}(h) = \sum_{i=1:t} \log \left[(p_s(y_{t+h}|y_{1:t})/p_{PBS}(y_{t+h}|y_{1:t})\right) \qquad (11)$$

Where $p_s(.)$ and $p_B(.)$ are the competing predictive densities conditional on real-time information.

Figure 7 presents the forecasting errors of the survey-based predictive density models. Early in the evaluation period, forecasters tended to overpredict one-quarter-ahead inflation; over time, this bias reversed, and forecasts shifted toward underprediction. By the end of the sample, errors are smaller and more stable, consistent with a more predictable forecasting environment.

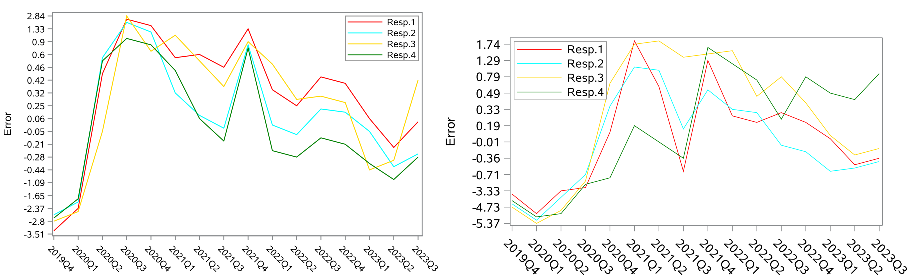

Figure 7. Forecasting errors from predictive density models. A. BPS model-based posterior median trajectories of the error in latent agent states. B. BPS model-based posterior median trajectories of the error in latent agent states.

Source: Author's computations based on col-spf, Banco de la República.

Table 1 reports RMSFE results for expected CPI inflation across forecast horizons, with the *h*-step-ahead point forecasts computed using posterior median parameter estimates. The results suggest that a parsimonious specification—placing weight on a smaller subset of consistently reliable experts—outperforms simple equal-weight ("follow-the-crowd") aggregation. Overall, BPS improves forecast accuracy relative to individual experts by about 10 %.

Table 1. RMSFE evaluation for one- and four-quarter-ahead forecasts

Model	h=1	%	h=4	%
Respondent 1	0.9326	-37.17	1.5326	-43.49
Respondent 2	0.7179	-24.96	1.2179	-26.58
Respondent 3	0.9253	-33.49	1.4253	-43.04
Respondent 4	0.6356	-20.93	1.1356	-17.08

Source: Authors' computations based on the best forecasting models.

Figure 8 presents LPDR values for one- and four-step-ahead forecasts relative to BPS. The results indicate that BPS outperforms individual experts in density forecast accuracy, provides a better characterization of predictive uncertainty, and adaptively adjusts the forecast location over time.

Bayesian Predictive Synthesis **00.5270** **1.2141**·

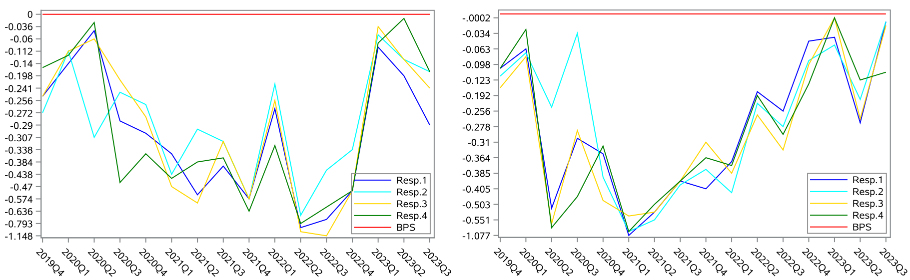

Figure 8. Log predictive density ratios (LPDR). A. One-step ahead LPDR, sequentially computed at each of the t=1:20 quarters. B. Four-step ahead LPDR, sequentially computed at each of the *t*=1:20 quarters.

Note: The horizontal baseline at zero for all *t* corresponds to the BPS benchmark. The plotted series report log density ratios at each forecast horizon, computed relative to the BPS predictive density.

Source: COL-SPF, Banco de la República.

BPS Results

Figure 9 shows that the coefficients evolve over time, with substantial changes in the implied weights. This dynamic specification accommodates shifts in agents' relative forecast performance and captures periods when previously reliable forecasters lose (or regain) predictive content. Overall, the results indicate that synthesizing a subset of consistently accurate experts is more effective than relying on a simple majority or equal-weight aggregation.

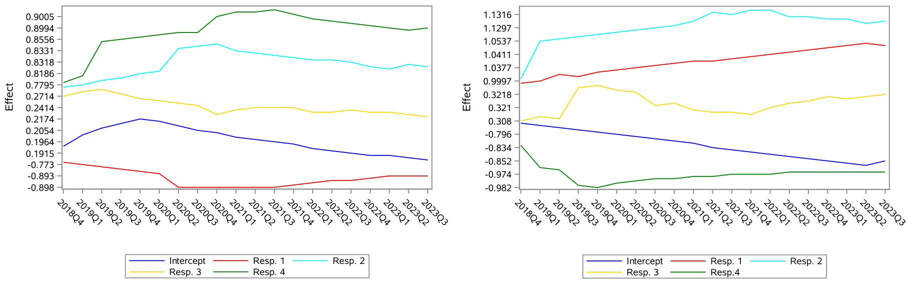

Figure 9. Evolution of the coefficients to synthesize agent expectations. A. On-line posterior means of BPS model coefficients for a one-quarter horizon, sequentially computed at each of the t=1:20 quarters. B. On-line posterior means of BPS model coefficients for a one-year horizon, sequentially computed at each of the *t*=1:20 quarters.

Source: COL-SPF, Banco de la República.

Finally, Figure 10 presents the resulting BPS-based consensus predictive densities, along with selected percentiles that summarize time variation in central tendency, uncertainty, and tail risks.

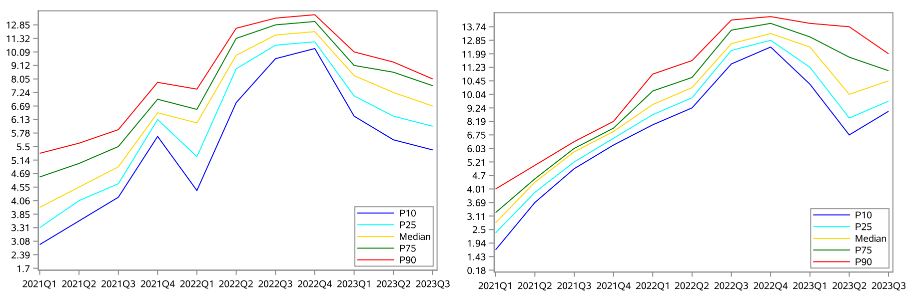

Figure 10. Consensus predictive densities based on Bayesian predictive synthesis (BPS). A. Consensus BPS for one-step-ahead predictive density. B. Consensus BPS for four-step-ahead predictive density.

Note: Forecasting horizons based on BPS predictive density show that inflation is generally well predicted for both time horizons.

Source: COL-SPF, Banco de la República.

Conclusions

We exploit a unique dataset of survey-based density forecasts of inflation expectations from economic experts. We document salient features of these distributions and implement an alternative density-combination approach to synthesize individual forecasts into a coherent measure of expected inflation. The proposed framework improves the characterization of forecast uncertainty and allows the location of predictive distributions to adapt over time, as reflected in our evaluation metrics.

Several extensions remain for future research. First, the synthesis function could be modeled using factor structures or nonparametric specifications as alternatives to shrinkage-based Bayesian regression. Second, because richer synthesis models can be computationally demanding, variational Bayes methods offer a promising approach to scalable posterior approximation and faster real-time implementation.

Authors' contribution

The author carried out the conceptualization, methodological design, formal analysis, and writing of the manuscript.

Funding

This work has not received funding.

Conflict of interest

The author declares no conflicts of interest with respect to this article.

Ethical considerations

The research adhered to international ethical standards for the social sciences. Data were either publicly available or collected with permission. There were no conflicts of interest or result manipulation.

Data availability

When they do not contain sensitive information, the methods, codes, and databases used will be available upon request from the authors.

Use of AI

Grammarly was used to check spelling and grammar in the text before submission.

References

Billio, M., Cesarini, R., Ravazzolo, F., & van Dijk, H. (2011). Combining Predictive Densities Using Bayesian Filtering With Applications to US Economics Data. *Tinbergen Institute Discussion Papers* 11-003/4. https://doi.org/10.2139/ssrn.1967435

Billiard, L., & Diday, E. (2020). *Clustering Methodology for Symbolic Data*. Jhon Wiley Sons Ltd. https://doi.org/10.1002/9781119010401

Chernis, T., Hauzenberger, N., Huber, F., Koop, G., & Mitchell, J. (2023). Predictive Density Combination Using a Tree-Based Synthesis Function. *Working Paper*, N.° 23-30. Federal Reserve Bank of Cleveland. https://doi.org/10.26509/frbc-wp-202330

Chernis, T., Hauzenberger, N., Huber, F., Koop, G., & Mitchell, J. (2025). Predictive Density Combination Using Bayesian Machine Learning. *International Economic Review*, *66*(3), 1289–1315. https://doi.org/10.1016/j.jeconom.2018.11.010

Clements, R., Rich, M., & Tracy, J. (2022). Surveys of Professionals. Federal Reserve Bank of Cleveland. *Working Paper*, N.° 13-22. https://doi.org/10.26509/frbc-wp-202213

Conflitti, C., De Mol, C., & Giannone, D. (2015). Optimal Combination of Survey Forecasts. *International Journal of Forecasting*, *31*(4), 1096–1103. https://doi.org/10.1016/j.ijforecast.2015.03.009

Diebold, M., Shin, F., & Hhang, B. (2023). On the Aggregation of Probability Assessments: Regularized Mixtures of Predictive Densities for Eurozone Inflation and Real Interest Rates. *Journal of Econometrics*, *237*(2-part C), 1–43. https://doi.org/10.1016/j.jeconom.2022.06.008

Gerrard, B. (2023). Ramsey and Keynes Revisited. *Cambridge Journal of Economics*, *47*(1), 195–213. https://doi.org/10.1093/cje/beac068

Iregui, A. (2021). ¿Qué nos dicen las encuestas sobre la formación de expectativas de inflación? *Revista Ensayos sobre Política Económica (ESPE)*, *39*(92), 30–44. https://doi.org/10.32468/espe.100

Lindley, D. V., Tversky, A., & Brown, R. V. (1979). On the Reconciliation of Probability Assessments. *Journal of the Royal Statistical Society. Series A*, *142*(2), 146–62. https://doi.org/10.2307/2345078

McAlinn, K., & West, M. (2019). Dynamic Bayesian Predictive Synthesis in Time Series Forecasting. *Journal of Econometrics*, *210*(1), 1555–1569. https://doi.org/10.1016/j.jeconom.2018.11.010

Poncela, P., & Senra, E. (2018). Measuring Uncertainty and Assessing Predictive Power in the Euro-Area. *Empirical Economics*, *53*(1), 165–182. https://doi.org/10.1007/s00181-016-1181-6

Tallman, E., & West, J. (2024). Bayesian Predictive Decision Synthesis. *Journal of the Royal Statistical Society Series B: Statistical Methodology, 86*(2), 340–63. https://doi.org/10.1093/jrsssb/qkad109

West, M., & Crosse, J. (1992). Modelling Probabilistic Agent Opinion. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, *54*(1), 285–99. https://doi.org/10.1111/j.2517-6161.1992.tb01882.x