Pontifícia Universidade Católica
do Rio de Janeiro

Felipe Gomes de Vasconcelos Musa

Inflation Risk Forecasting in Brazil

Dissertação de Mestrado

Masters dissertation presented to the Programa de Pós–graduação em Economia, do Departamento de Economia da PUC-Rio in partial fulfillment of the requirements for the degree of Mestre em Economia.

Advisor : Prof. Márcio Gomes Pinto Garcia
Co-advisor: Prof. Lucas Lima

Rio de Janeiro
April 2025

Pontifícia Universidade Católica
do Rio de Janeiro

Felipe Gomes de Vasconcelos Musa

Inflation Risk Forecasting in Brazil

Masters dissertation presented to the Programa de Pós–graduação em Economia da PUC-Rio in partial fulfillment of the requirements for the degree of Mestre em Economia. Approved by the Examination Committee:

Prof. Márcio Gomes Pinto Garcia
Advisor
Departamento de Economia – PUC-Rio

Prof. Lucas Lima
Co-advisor
Departamento de Economia – PUC-Rio

Prof. Marcelo Cunha Medeiros
Departamento de Economia - PUC-Rio

Prof. João Victor Issler
Escola Brasileira de Economia e Finanças - FGV EPGE

Rio de Janeiro, April 24th, 2025

Felipe Gomes de Vasconcelos Musa

B.A. in Economics, Pontifícia Universidade Católica do Rio de Janeiro (PUC-Rio), 2021.

Bibliographic data

Musa, Felipe Gomes de Vasconcelos

Inflation Risk Forecasting in Brazil / Felipe Gomes de Vasconcelos Musa; advisor: Márcio Gomes Pinto Garcia; co-advisor: Lucas Lima. – 2025.

51 f: il. color. ; 30 cm

Dissertação (mestrado) - Pontifícia Universidade Católica do Rio de Janeiro, Departamento de Economia, 2025.

Inclui bibliografia

1. Economia – Teses. 2. Previsão. 3. Risco de Inflação. 4. Densidades Preditivas. 5. Probabilidades. 6. Meta de Inflação. I. Garcia, Márcio Gomes Pinto. II. Lima, Lucas. III. Pontifícia Universidade Católica do Rio de Janeiro. Departamento de Economia. IV. Título.

CDD: 004

Acknowledgments

First, I would like to sincerely thank my family and lifelong friends. This journey would not have been possible without their constant support. Especially, I am deeply thankful to my parents, Roberta and Paulo, for their love, their dedication throughout my life, and all the sacrifices they made to invest in my education. Their encouragement and belief in me at every step made it possible for me to pursue and complete my master's degree. I also extend my heartfelt gratitude to my sister, Isabel, for her companionship, and to my grandmother, Marilia, for her constant love and care.

I am grateful to my classmates and friends at PUC-Rio, who made the challenging moments of the program lighter and more enjoyable. I am also immensely thankful to the professors and staff of the Department of Economics at PUC-Rio for providing a world-class education during both my undergraduate and master's years.

I would like to express my sincere appreciation to my advisors, Márcio Garcia and Lucas Lima, for their guidance and support throughout this project. I am also grateful to the committee members, Marcelo Medeiros and João Victor Issler, for their time and insightful feedback.

Finally, I would like to acknowledge CAPES and PUC-Rio for the financial support that made this work possible.

This study was financed in part by the Coordenação de Aperfeiçoamento de Pessoal de Nível Superior - Brasil (CAPES) - Finance Code 001.

Abstract

Musa, Felipe Gomes de Vasconcelos; Garcia, Márcio Gomes Pinto (Advisor); Lima, Lucas (Co-Advisor). **Inflation Risk Forecasting in Brazil**. Rio de Janeiro, 2025. 51p. Dissertação de Mestrado – Departamento de Economia, Pontifícia Universidade Católica do Rio de Janeiro.

Inflation forecasting is central to monetary policy, and effective decision-making under uncertainty requires a risk management approach that considers the full distribution of outcomes rather than relying solely on point estimates. This paper evaluates the accuracy and policy relevance of probabilistic inflation forecasts under Brazil's inflation-targeting regime. We begin by analyzing the Central Bank of Brazil's (BCB) predictive distributions and document a systematic underestimation of inflation uncertainty. As an alternative, we implement a Quantile Phillips Curve (QPC) model, which shows improved calibration—particularly in upper-intermediate quantiles, a critical area given Brazil's historically upward-skewed inflation environment. The models' ability to assess inflation target breach scenarios is then compared. While BCB forecasts provide only short-term alerts, the QPC model delivers informative signals up to 12 months ahead. The paper's main original contribution is the introduction of a multi-period target risk measure that combines persistence and timing dimensions, estimated via a flexible copula-based approach. Applied to the QPC model, this framework identifies persistent inflation breaches up to six months in advance, offering policymakers a valuable tool for more nuanced, forward-looking risk management.

Keywords

Forecasting; Inflation Risk; Predictive Densities; Probabilities; Inflation Target.

Resumo

Musa, Felipe Gomes de Vasconcelos; Garcia, Márcio Gomes Pinto; Lima, Lucas. **Previsão de Risco de Inflação no Brasil**. Rio de Janeiro, 2025. 51p. Dissertação de Mestrado – Departamento de Economia, Pontifícia Universidade Católica do Rio de Janeiro.

A previsão de inflação é central para a política monetária, e a tomada de decisão eficaz em contextos de incerteza exige uma abordagem de gestão de riscos que considere toda a distribuição de possíveis cenários, em vez de se basear apenas em estimativas pontuais. Este trabalho avalia a acurácia e a relevância para a política econômica das previsões probabilísticas de inflação no regime de metas de inflação do Brasil. Iniciamos analisando as distribuições preditivas do Banco Central do Brasil (BCB) e documentamos uma subestimação sistemática da incerteza inflacionária. Como alternativa, implementamos um modelo de Curva de Phillips Quantílica (QPC), que apresenta melhor calibração — especialmente nos quantis intermediários superiores, uma faixa crítica diante do histórico de assimetria à alta inflação no Brasil. Em seguida, comparamos a capacidade dos modelos em avaliar riscos de descumprimento da meta de inflação. Enquanto as previsões do BCB oferecem apenas alertas de curto prazo, o modelo QPC fornece sinais informativos com até 12 meses de antecedência. A principal contribuição original do trabalho é a introdução de uma medida de risco multiperíodo de descumprimento da meta, que combina as dimensões de persistência e momento temporal, estimada por meio de uma abordagem flexível baseada em cópulas. Aplicada ao modelo QPC, essa estrutura identifica descumprimento persistentes da meta de inflação com até seis meses de antecedência, oferecendo aos formuladores de política uma ferramenta valiosa para uma gestão de riscos mais refinada e prospectiva.

Palavras-chave

Previsão; Risco de Inflação; Densidades Preditivas; Probabilidades; Meta de Inflação.

Table of contents

List of figures

List of tables

List of Abbreviations

BCB – Brazilian Central Bank
QPC – Quantile Phillips Curve
PIT – Probability Integral Transforms
CDF – Cumulative Distribution Function
FCI – Financial Conditions Index
IT – Inflation Target
MAR – Mixed Causal-Noncausal Autoregressive
ROC – Receiver Operating Characteristic
AUC – Area Under the ROC Curve

1
Introduction

Inflation forecasting plays a critical role in monetary policymaking. While point estimates have traditionally received primary attention, the inherent uncertainty of economic projections requires policymakers to also consider the full distribution of potential outcomes:

> "[A] central bank seeking to maximize its probability of achieving its goals is driven, I believe, to a risk-management approach to policy. By this I mean that policymakers need to consider not only the most likely future path for the economy but also the distribution of possible outcomes about that path."
>
> *Federal Reserve Board Chairman Alan Greenspan*, August 2003.

In this context, central banks increasingly complement point forecasts with probabilistic assessments to better communicate and quantify the uncertainty surrounding inflation outlook. A well-specified distribution allows policymakers to evaluate the likelihood, costs, and benefits of alternative inflation scenarios, thereby supporting monetary policy decision-making.

In emerging economies—where inflation dynamics tend to be more volatile and frequently influenced by structural shocks—the development of reliable probabilistic forecasts is particularly important for effective monetary policymaking.

This study focuses on Brazil, a major emerging country that has adopted an inflation-targeting regime in 1999. Since then, the Central Bank of Brazil has regularly published its Inflation Fan Chart (*Leque de Inflação*), a probabilistic representation of the uncertainty regarding expected inflation paths. Given its importance as a primary input into monetary policy decisions[1], it is paramount to assess the accuracy of the Central Bank of Brazil's (henceforth BCB) distribution forecasts.

To this end, we conduct a systematic assessment of these forecasts across multiple horizons. Findings indicate that BCB's predictive densities demonstrate persistent underestimation of inflation outcomes, with poor calibration[2]

[1] As stated in the September 1999 Monetary Policy Report: "In this sense, the inflation fan chart is an extremely useful and important tool since its objective is to announce the conclusions drawn by the Monetary Policy Committee regarding the expected inflation trajectory, (...), and is used both as the basis for discussions and as an aid in the decision-making process."

[2] Calibration measures how well forecasted probabilities correspond to actual outcome frequencies.

at key monetary policy horizons (six and eight quarters ahead). In contrast, we introduce a Quantile Phillips Curve (QPC) model, inspired by Lopez-Salido and Loria (2024), which exhibits improved calibration properties.

Considering the uncertainty in inflation forecasting, policy should adapt to emerging risks that could threaten monetary goals. Hence, the ability to monitor and quantify inflation risks is essential for central banks to fulfill their mandates effectively. This is particularly relevant for central banks operating under an Inflation Targeting (IT) regime, as their credibility is often based on their capacity to maintain inflation within predefined targets.

Brazil's experience with inflation targeting illustrates the institutional diversity that can exist across monetary policy frameworks. From 1999 to 2024, the country followed a fixed event approach, evaluating inflation at year-end (December) against pre-announced targets and tolerance bounds.

In January 2025, Brazil transitioned to a "continuous target" regime: inflation is now assessed monthly, with a 12-month target of 3.00%, and a tolerance range of ± 1.50 percentage points. Non-compliance requires that inflation stays outside this band for six consecutive months, thus underscoring the need for a continuous risk assessment for efficient policy-making.

This paper also analyzes how effectively BCB captures inflation target risk, by evaluating the likelihood of inflation target breaches scenarios based on its probabilistic forecast tool. Results indicate that BCB's forecasts are only informative at shorter horizons, whereas our proposed QPC model provides more reliable assessments of inflation target risk at medium-term horizons.

Furthermore, a key original contribution of this paper is the introduction of an inflation target *duration* risk measure, which also accounts for the persistence of deviations from the target. We propose estimating this measure using a flexible Copula-based approach that requires only individual horizon-specific forecast densities as input, making it applicable across a wide range of forecasting frameworks.

We assess this novel metric using forecasts from the suggested QPC model, finding it delivers informative early warnings of sustained inflation target breaches scenarios up to six months ahead. Therefore, offering policymakers a valuable tool for forward-looking risk management.

The document is structured as follows. Chapter 2 summarizes related works and situates our contributions. Chapter 3 presents the density forecasting methodology, the dataset, and the Quantile Phillips Curve model. Chapter 4 evaluates the calibration of predictive densities, while Chapter 5 analyzes their performance in gauging inflation target risk. Chapter 6 concludes with a synthesis of the key findings.

2
Related Literature

Evaluations of central bank density forecasts have largely centered on the Bank of England (e.g., Mitchell and Hall (2005)). For Brazil, Galvão (2005) and Knuppel and Schultefrankenfeld (2019) assess the BCB Inflation Fan Charts under the assumption of an underlying Gaussian distribution, whereas the present study permits a more flexible distributional format.

Galvão (2005), based on data up to 2004, evaluate forecast uncertainty across the distribution and reports that forecasts are well-calibrated at the h=0 (nowcast) median, with uncertainty overestimation for other forecast horizons and quantiles. In comparison, this work consider a substantially larger sample and longer horizons. While we also find accurate median forecasts at $h = 0$, our results indicate uncertainty underestimation at longer horizons. This divergence is likely due to their sample period, marked by high inflation volatility, which prompted the BCB to produce wider forecast intervals.

Knuppel and Schultefrankenfeld (2019) examine forecast uncertainty from 1999 to 2016, focusing on the average distribution calibration. They find that BCB tends to underestimate uncertainty at most horizons. This study differs in both scope and focus: we use an extended sample, broader horizons and calibration across quantiles. Our findings similarly indicate systematic underestimation, with significant miscalibration in the upper quantiles at 6 and 8 quarters ahead.

Empirical density forecasting of macroeconomic variables can be estimated using parametric methods, assuming specific forms, or via non-parametric methods, which are flexible but prone to overfitting. Seminal work of Adrian, Boyarchenko and Giannone (2019) proposes a semi-parametric two-step approach: first, estimate quantiles via Quantile Regression; second, fit them with a skewed t-distribution to yield a full distribution forecast

Originally applied to GDP, this approach has been used to model future inflation distribution in works such as Lopez-Salido and Loria (2024), Tagliabracci (2020), and Banerjee et al. (2022). All grounded in Phillips Curve-based quantile models for the U.S., Euro Area and a panel of advanced economies. This study adds to the empirical distribution estimation literature by applying the Quantile Phillips Curve method to Brazil, a major emerging economy.

Despite its policy importance, research on quantifying inflation target risk remains scarce. Hecq, Issler and Voisin (2024) proposes a risk measure based on the conditional probability that inflation will be within the bounds at h period ahead. Applied to Brazilian data, the authors use a Mixed Causal-Noncausal Autoregressive (MAR) model to derive the probabilities for 1,3 and 6-month ahead. Results indicate that the model generates informative risk alerts only up to 3 months ahead.

In this context, we assess the capacity of both BCB and QPC models to deliver informative early-warning signals of target breaches. Using the predictive densities to derive conditional breach probabilities, we find that both models outperform the MAR model: the BCB provides accurate signals up to 6 months ahead, while the QPC remains informative up to 12 months.

Hecq, Issler and Voisin (2024) introduced a timing-based target risk measure, evaluating compliance only at horizon h. This paper propose a new risk measure that integrates timing and duration dimensions, offering a more robust inflation target risk assessment.

The new approach considers the conditional probability that inflation breaches bounds at horizon h and persists outside for k consecutive periods. Applied to our QPC model,using a copula-based estimation inspired by Mogliani and Odendahl (2025), we produce informative early-warning predictions for windows up to 6 months ahead.

3
Predictive Density Estimation

3.1
Methodology

In our setting, each model provides a finite set of conditional inflation quantile forecasts.[1] Following a standard approach in the macroeconomic risk literature, as in Adrian, Boyarchenko and Giannone (2019), we recover the underlying predictive inflation density by mapping these discrete quantile estimates into a continuous distribution.

Specifically, the skew t-distribution proposed by Azzalini and Capitanio (2003) is fitted, which provides a flexible yet parsimonious specification capable of capturing asymmetric and heavy-tailed behavior. The distribution is characterized by four parameters: location (μ), scale (σ), shape (α), and kurtosis (ν).

For each forecast origin T and horizon h, parameters are estimated by minimizing the ℓ_2 norm between the empirical conditional inflation forecast quantiles and theoretical quantiles implied by the skew-t distribution:

$$\underset{\mu_{T+h},\sigma_{T+h},\alpha_{T+h},\nu_{T+h}}{\arg\min} \sum_{\tau} \left(\hat{Q}_{\pi_{T+h}|\mathcal{I}_T}(\tau) - \hat{F}^{-1}(\tau;\, \mu_{T+h},\, \sigma_{T+h},\, \alpha_{T+h},\, \nu_{T+h})\right)^2$$

where $\mathcal{I}_T$ is the information set available at T and F is the cumulative distribution function (CDF) of the skewed-t probability density function (PDF) f, given by:

$$f(\pi;\mu,\sigma,\alpha,\nu) = \frac{2}{\sigma}\, s\left(\frac{\pi-\mu}{\sigma};\nu\right)\, S\left(\alpha\frac{\pi-\mu}{\sigma}\sqrt{\frac{\nu+1}{\nu+\left(\frac{\pi-\mu}{\sigma}\right)^2}};\nu+1\right),$$

where $s(\cdot)$ and $S(\cdot)$ denote the PDF and CDF of the Student-t distribution, respectively. When $\alpha = 0$, the distribution reduces to the standard Student-t. In the limiting case as $\nu \to \infty$, becomes a Gaussian density with mean μ and standard deviation σ. Similarly to Adrian, Boyarchenko and Giannone (2019), we focus on the exactly identified case, matching the $\tau = 0.05$,

[1] If the estimated quantiles are not monotonic, the uncrossing procedure of Chernozhukov, Fernández-Val and Galichon (2010) is applied to enforce monotonicity.

0.25, 0.75, and 0.95 quantiles.[2]

3.2
Brazilian Central Bank Model

The construction of Inflation Fan Chart relies on a predetermined parametric probability distribution, whose functional form is not publicly disclosed. This distribution is characterized by parameters[3] that are specified using a combination of quantitative methods and the subjective judgment of the Monetary Policy Committee regarding potential inflationary risks.

Published fan charts display the forecast median and quantiles[4] associated with the probability distributions at various points along the projection horizon. Figure 3.1 provides an example of a BCB Inflation Fan Chart, as published in the Monetary Policy Report of December 2024.

Figure 3.1: BCB Inflation Projections and Fan Chart

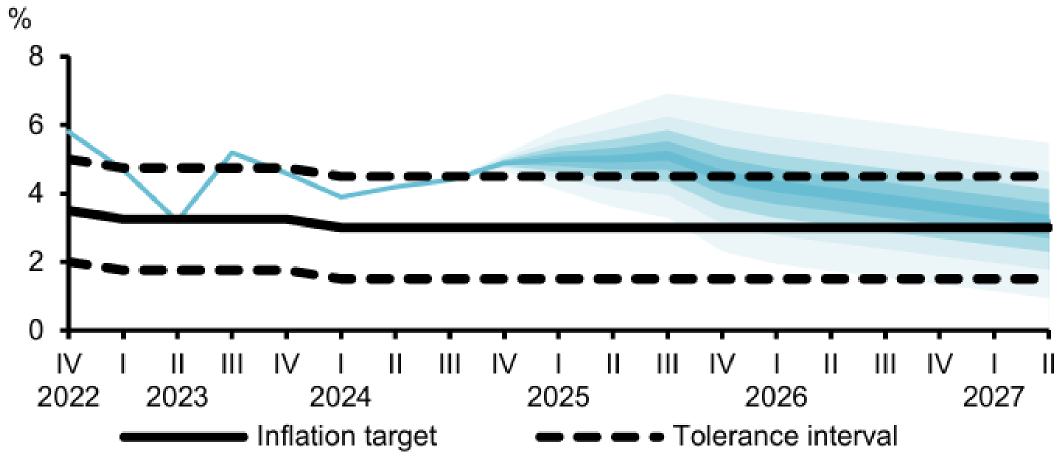

Figure 3.1: Displays the BCB's outlook, made in 2024Q4, for 12-month IPCA inflation rate. Shaded areas represent projections intervals associated with the following probabilities (from the inner to the outer interval): 10%, 30%, 50%, 70% and 90%.

[2]An alternative approach, beyond the scope of this paper, would be to empirically explore the use of additional quantiles (e.g., include also $\tau = 0.50$) to pin down the parameters of f, allowing cases of over-identification.

[3]These include a measure of central tendency, dispersion, and asymmetry.

[4]Quantiles typically displayed are 5, 15, 25, 35, 45, 55, 65, 75, 85, and 95.

3.2.1 Data

Inflation Fan Charts are published quarterly in the Monetary Policy Report. The forecast for the current quarter ($h = 0$) corresponds to the 12-month inflation projection for the final month of the respective quarter.[5] Subsequent h-quarter-ahead forecasts correspond to $3h$ month-ahead projections.

Table 3.1 below summarizes the available data for each h-quarter-ahead horizon considered.[6] Last forecast origin is selected such that the final out-of-sample projection corresponds to 2024Q4.

Table 3.1: BCB Model - Inflation Target Sample

Horizon	**Forecast Origins**	**Sample Size**
$h = 0$	1999Q4 – 2024Q4	101
$h = 1$	1999Q2 – 2024Q3	102
$h = 2$	1999Q2 – 2024Q2	101
$h = 3$	1999Q2 – 2024Q1	100
$h = 4$	1999Q2 – 2023Q4	99
$h = 5$	1999Q2 – 2023Q3	98
$h = 6$	1999Q2 – 2023Q2	90
$h = 7$	1999Q2 – 2023Q1	82
$h = 8$	1999Q2 – 2022Q4	73

Note: For horizons $h = 0, 6, 7, 8$, predictions were unavailable for certain periods within the forecast origin range.

For each forecast origin, we collect inflation quantiles at horizons $h = 0, \ldots, 8$. In particular, these predictive quantiles assume that the future nominal policy interest rate coincides with market expectations and the exchange rate follows a predefined path.

3.2.2 Estimation

BCB's predictive densities, obtained from the quantile forecasts as detailed in Section 3.1, are visualized in Figure 3.2. Each density is aligned such that the probability function for period t corresponds to the forecast issued at time t-h.

[5] Forecasts are released during the last month of the respective quarter, prior to the publication of the IPCA for that month.

[6] Although the maximum horizon extends up to 13 quarters ahead, sample sizes for horizons beyond 8 quarters are limited.

The figures reveal a structural shift in the format of the predictive distributions around 2008. Prior to this period, the distributions were relatively flexible, exhibiting asymmetry and time-varying dispersion. After 2008, however, the distributions resemble a Gaussian shape, with variation occurring primarily in the central tendency.

Figure 3.2: BCB Predictive Densities

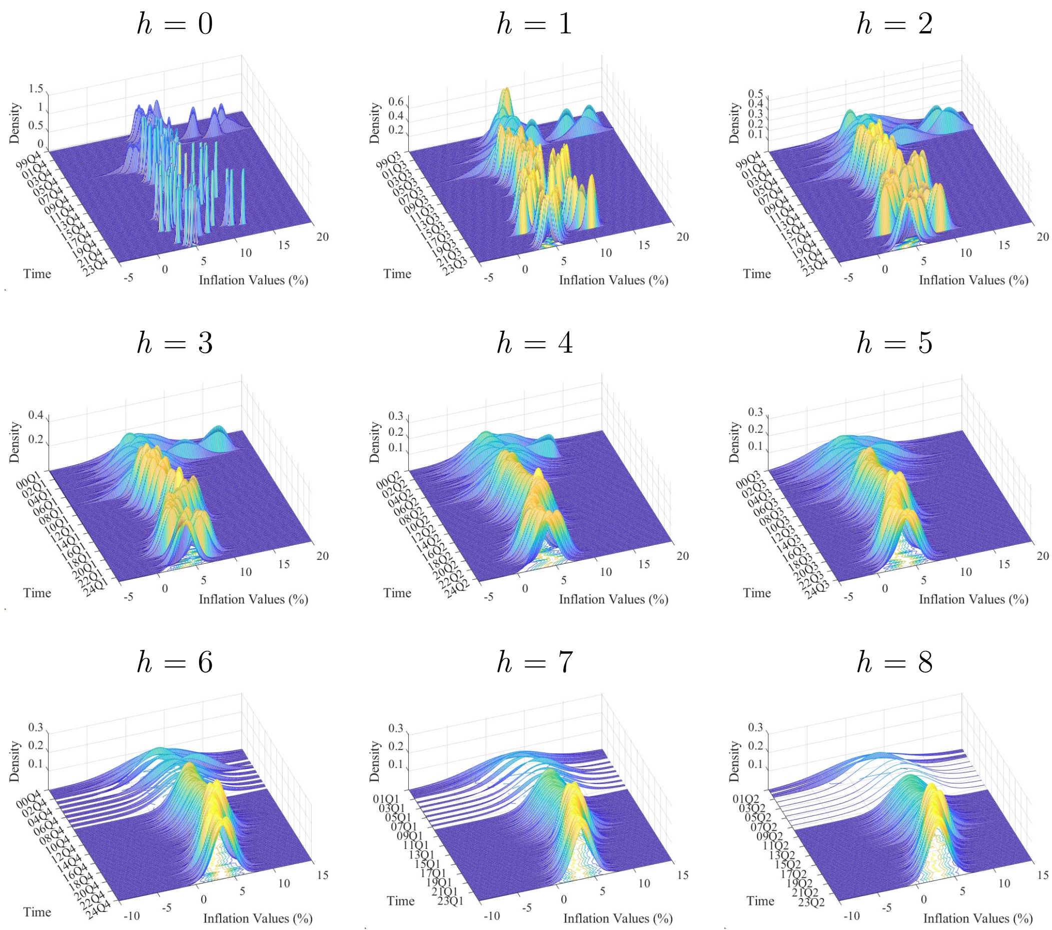

Figure 3.2: Each panel presents BCB's predictive density for period t, based on forecasts issued at time t-h. The sample encompasses all available projections during the Inflation Targeting period, with the final out-of-sample forecast corresponding to 2024Q4. Predictions for horizons $h = 6, 7, 8$ were unavailable for certain periods.

3.3 Quantile Phillips Curve Model

Our econometric framework for conditional inflation quantile forecasting adopts a distinct approach from the BCB's methodology. While the Central Bank assumes an ex-ante parametric distribution and derives quantiles accordingly, we directly estimate the quantiles and subsequently recover the implied predictive density.[7]

The model specification builds upon the Phillips-curve framework of Blanchard, Cerutti and Summers (2015), estimating the conditional τ-quantile of h-step-ahead inflation (π_{t+h}) as:[8]

$$\begin{aligned} Q_\tau\left(\pi_{t+h} \mid X_t\right) = \alpha_\tau + (1-\lambda_\tau)\pi_t + \lambda_\tau \pi_t^{\text{exp}} + \beta_{1,\tau}\text{Output\_Gap}_t \\ +\beta_{2,\tau}\pi_t^{\text{imp}} + \beta_{3,\tau}\text{FCI}_t + \beta_{4,\tau}\text{Deficit}_t \end{aligned} \tag{3-1}$$

Contemporaneous inflation (π_t) captures the role of inertia in the price-setting process, while long-term inflation expectations (π_t^{exp}) reflect the forward-looking price-setting behavior of firms. The dominant behavior is determined by the value of the parameter λ_τ. Other common factors in the literature include economic slack measure (Output_Gap_t), relative prices shocks (π_t^{imp}) and financial conditions (FCI_t).

We also incorporate the government's fiscal primary result (Deficit_t). This inclusion is motivated by Banerjee et al. (2022), which estimates a quantile-augmented Phillips curve that includes fiscal balance, using data from 21 economies over four decades. Their findings reveal that higher deficits significantly affect not only average inflation but also the tails of the distribution of future inflation outcomes in "fiscally led" regime economies.[9] The rationale is that larger deficits are associated with tail movements due to risks of monetization and fiscal dominance. Given Brazil's structural challenges regarding fiscal sustainability and recurring political pressures on its central bank, the inclusion of fiscal balance as an explanatory variable can be justified.[10]

[7] Quantiles are estimated independently for each forecast horizon, allowing the predictive distribution to vary without imposing a specific functional form.

[8] Formally, the dependence between x_t and the τ-quantile of π_{t+h} is measured by:

$$\hat{\beta}_\tau = \underset{\beta_\tau \in \mathbb{R}^k}{\arg\min} \sum_{t=1}^{T-h} \left(\tau \cdot \mathbf{1}_{(\pi_{t+h} \geq X_t\beta)} |\pi_{t+h} - X_t\beta_\tau| + (1-\tau) \cdot \mathbf{1}_{(\pi_{t+h} < X_t\beta)} |\pi_{t+h} - X_t\beta_\tau|\right),$$

where $\mathbf{1}_{(\cdot)}$ denotes the indicator function.

[9] Defined as a regime in which the government does not adjust the primary balance to stabilize debt, and the central bank exhibits reduced independence.

[10] Appendix A.1 presents an inference exercise assessing the relationship between explanatory variables and movements in different parts of the inflation distribution. Results show

3.3.1 Data

The model utilizes monthly data from January 2004 to December 2024. Parameters are re-estimated each period using an expanding training window of 120 months (10 years, approximately 50% of the sample), which sets the first forecast origin at December 2013.

Horizons range from $h = 0$ (nowcast) to $h = 4$ quarters (12 months) ahead, as the direct forecasting approach—estimating parameters for each horizon h to predict h-period-ahead values from current data—is not advisable for longer horizons due to declining efficiency.[11]

The last forecast origin is chosen such that, for each horizon, the final out-of-sample prediction corresponds to December 2024. This setup yields an initial sample of 133 monthly quantile forecast observations for $h = 0$, 130 for $h = 1$ and decreasing linearly up to 121 for $h = 4$.

To align with the BCB's predictions, made at the end of each quarter, initial sample is restricted to include only predictions issued at the last month of each quarter. This adjustment results in 45 quantile prediction observations for $h = 0$, 44 for $h = 1$ and decreasing linearly up to 41 for $h = 4$.

Table 3.2 summarizes the variables used in the analysis, along with their typical release lags. These lags are incorporated to construct data vintages that approximate the information set available at each forecast origin.[12]

statistically significant effects of fiscal deficits, especially on the upper tail.

[11]Empirical evidence (e.g., Marcellino, Stock and Watson (2006)) indicates that direct forecasts remain efficient relative to iterated forecasts up to 12 steps ahead, consistent with the maximum horizon in Mogliani and Odendahl (2025), that also employs a QPC model.

[12]BCB specifies the exact day within the month on which the Fan Charts are produced. As forecasts are issued during month t, and data for that month are incomplete, the minimum feasible release lag is one month. Descriptive statistics are available in Appendix A.2.

Table 3.2: Variables Description

Variable	Definition	Release Lag
π_t	Headline IPCA (12-month % change)	1 Month
π_t^{exp}	FOCUS Survey Median 4 Years Ahead Expected Inflation	1 Month
Output_Gap_t	Cycle decomposition of IBC-Br activity indicator using HP Filter	2–3 Months
π_t^{imp}	Monthly % changes nominal BRL/USD + % changes Brazil Commodity Index	1 Month
FCI_t	BCB Financial Conditions Index	1 Month
Deficit_t	12-Month Accumulated Primary Fiscal Deficit (% GDP)	2 Months

3.3.2 Estimation

At each forecast origin, we estimate Equation 3-1 to generate, almost[13] real time, out-of-sample h-step-ahead quantile forecasts for inflation. Using the methodology described in Section 3.1, we derive predictive densities for horizons $h = 0, 1, 2, 3, 4$ quarters ahead, which are visualized in Figure 3.3.

[13] Data correspond to revised (final) values.

Figure 3.3: QPC Predictive Densities

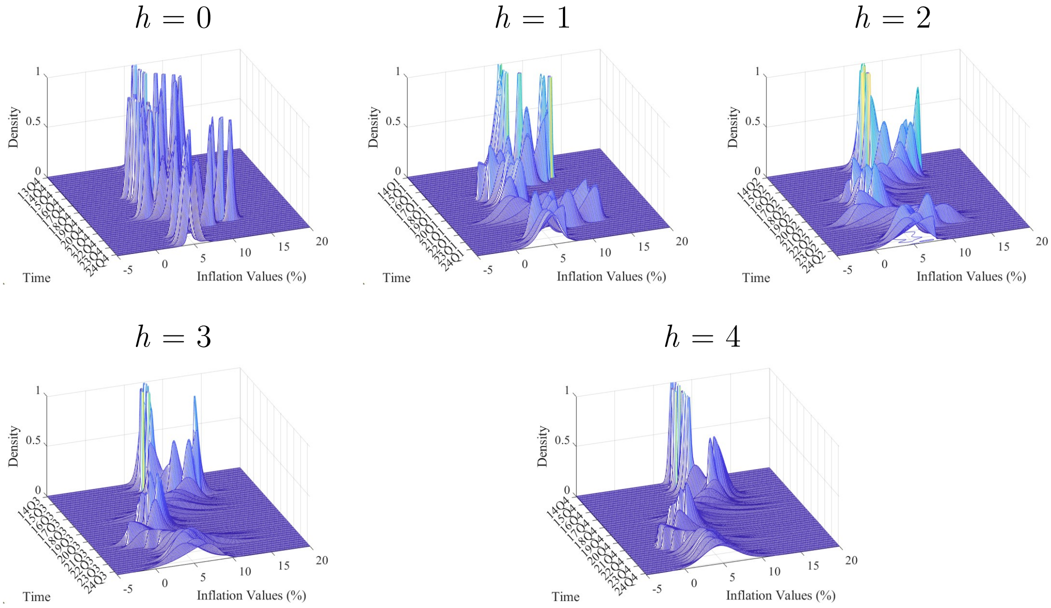

Figure 3.3: Each panel presents QPC predictive density for period t, based on forecasts issued at time t-h. The sample encompasses all projections starting in 2013Q4, with the final out-of-sample forecast corresponding to 2024Q4.

The panels display flexible distributional shapes with time-varying dispersion, skewness, and kurtosis. This contrasts with the BCB's constant Gaussian format across most horizons during the same sub-sample period, as illustrated in Figure 3.4 for the $h = 4$ forecasts.

Figure 3.4: BCB Four Quarters Ahead Predictive Densities

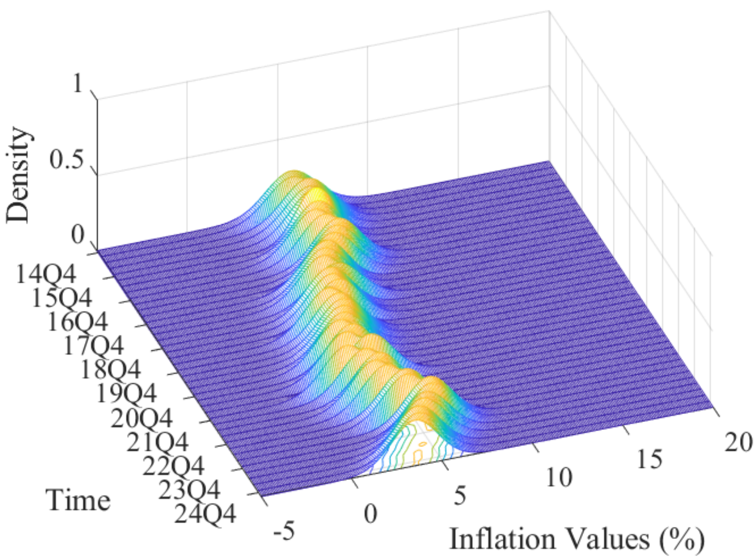

Figure 3.4: Presents BCB predictive densities for period t, based on forecasts issued at time t-4. The sub-sample encompasses all projections starting in 2013Q4, with the final out-of-sample forecast corresponding to 2024Q4.

4
Predictive Density Calibration

This Chapter outlines the framework for evaluating the calibration of predictive densities presented in Chapter 3. It assesses the calibration of the BCB model across the entire Inflation Target period and also compares it with the QPC model over the 2013:Q4 to 2024:Q4 period.

4.1
Density Evaluation Test

We formally assess the forecast distributions of both models by analyzing their calibration—that is, the extent to which the predictive distributions align with the true, unobserved data-generating process.

To do so, we employ probability integral transforms (PITs). For each forecast, the PIT is obtained by computing the conditional predictive cumulative distribution function (CDF)

$$\widehat{F}_{\pi_{t+h}|\mathcal{I}_t}(\cdot) = \mathbb{P}(\pi_{t+h} \leq \cdot \mid \mathcal{I}_t)$$

at the realized inflation π^*_{t+h}.

$$PIT_{\widehat{F}_{\pi_{t+h}}|\mathcal{I}_t}(\pi^*_{t+h}) \equiv \widehat{F}_{\pi_{t+h}}|\mathcal{I}_t(\pi^*_{t+h})$$

In a perfectly calibrated model, these PITs are independently and identically distributed as $U(0,1)$ (Diebold et al. (1998)); that is, if the model assigns a *ex-ante* probability τ to an event, then approximately a fraction τ of the *ex-post* observations should fall below the corresponding forecast quantile.

4.2
BCB's Density Calibration over Inflation Target Period

First, we evaluate BCB's quarterly density forecasts over the Inflation Target period. Figure 4.1 shows, for every forecast horizon, the empirical CDF of the PITs (solid blue line) alongside the 45-degree line of an ideal uniform distribution. Deviations from this line signal potential miscalibration of the predictive distribution. To address sampling uncertainty, 95% confidence bands (dashed blue lines) are added, following Rossi and Sekhposyan (2019).[1] A curve

[1]Bands are based on a block weighted bootstrap from Inoue (2001), assuming uniformity. As the BCB's estimation process is undisclosed, these bands, derived from rolling window estimates, serve as indicative benchmarks.

falling outside the bands indicates rejection of the null hypothesis of correct calibration.

Figure 4.1: BCB Empirical Distribution of PITs

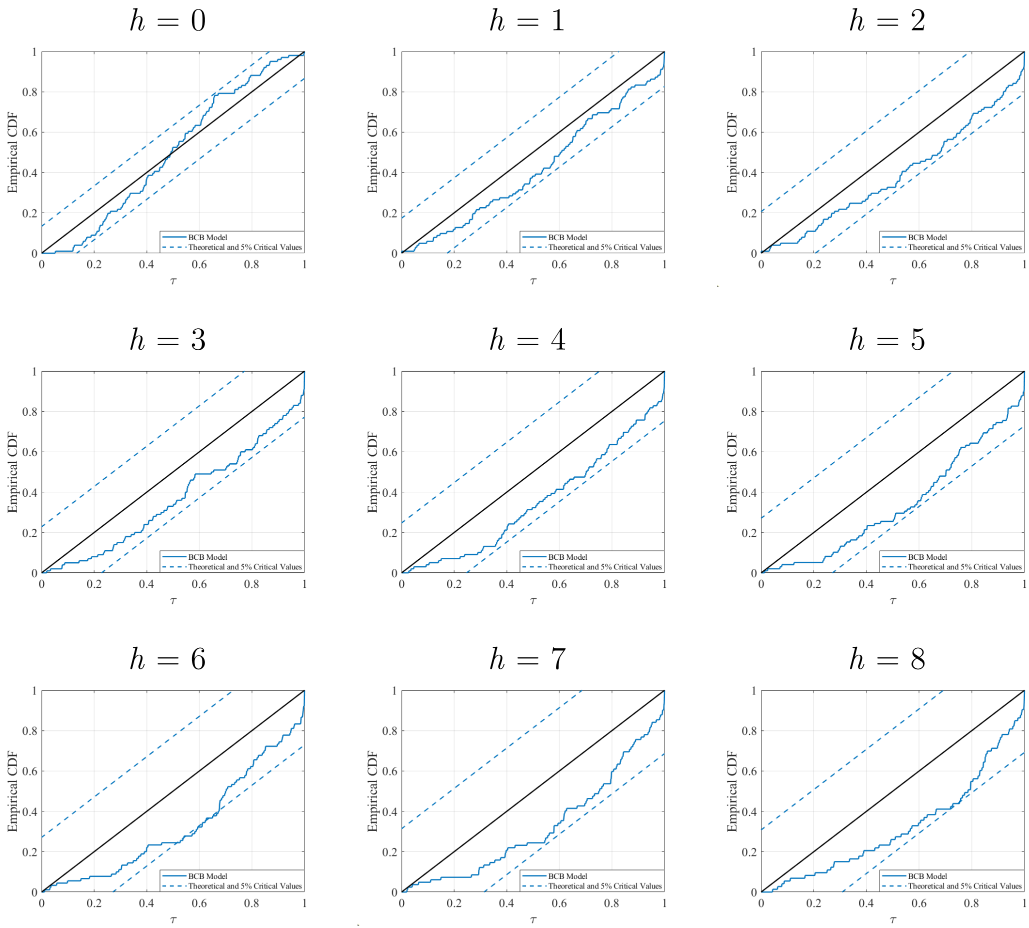

Figure 4.1: Each figure reports, for a specific forecast horizon, the empirical cumulative distribution of the probability integral transform (PITs) for BCB's quarterly forecasts from 1999:Q2 to 2024:Q4. The general guidance 95% confidence bands (dashed lines) are obtained as in Rossi and Sekhposyan (2019) and plotted parallel to the 45-degree line, representing the theoretical ideal calibration.

The PIT plot analysis reveals that forecast calibration deteriorates as the forecast horizon increases. At $h = 0$, the predictive distribution is well-calibrated at the median but overly wide, with realized inflation outcomes clustering around the center of the distribution. This produces an distinct S-shaped PIT CDF deviation, indicating underconfidence, though the curve remains narrowly within the general guidance confidence bounds for calibrated forecasts.

Across forecast horizons ranging from one quarter to two years ahead, the PIT CDF systematically lies below the theoretical 45-degree line, indicating a systematic underestimation of realized inflation. This suggests that inflation outcomes frequently fall in the upper tail—or even exceed—the predicted distributions, leading to lower than anticipated PIT CDF values across nearly all quantiles.

Despite the empirical distribution remaining within general guidance confidence bounds for most horizons, calibration is rejected at the current relevant monetary policy horizon ($h = 6$)[2] and for long-term projections ($h = 8$), where the plots reveal a systematic underprediction notably in the upper intermediate quantiles of the inflation distribution.

4.3 Models Density Calibration over 2013:Q4 to 2024:Q4

In order to compare both models, we employ the PIT CDF distribution test considering the 2013:Q4 to 2024:Q4 sample, constrained by the QPC model's densities forecasts availability.

Figure 4.2 presents the empirical CDF of the PITs for each forecast horizon, with the BCB model depicted by a solid blue line and the QPC model by a solid green line.

Figure 4.2: BCB and QPC Empirical Distribution of PITs

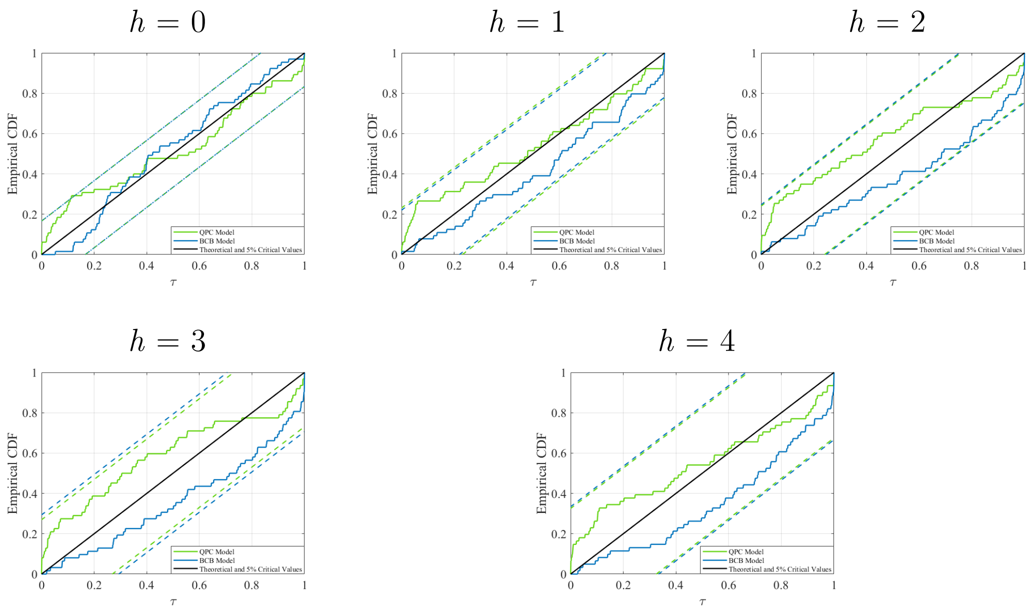

Figure 4.2: Each figure reports, for a specific forecast horizon, the empirical cumulative distribution of the probability integral transform (PITs) for BCB's and QPC's quarterly forecasts from 2013:Q4 to 2024:Q4. The general guidance 95% confidence bands (dashed lines) are obtained as in Rossi and Sekhposyan (2019) and plotted parallel to the 45-degree line, representing the theoretical ideal calibration.

[2]Considered as the timeframe that the Central Bank views it can steer inflation in order to achieve its target.

Analysis of the PIT plots indicates that, across most horizons, both models remain within their respective 95% confidence bands. The plots also suggest that the QPC model generally exhibits superior forecast calibration compared to the BCB model. Across all forecast horizons, the QPC model's empirical CDF intersects the 45-degree reference line at least once and remains closer to it, particularly in the central and upper intermediate regions of the inflation distribution.

This indicates a more accurate assessment of moderate to high inflation scenarios—arguably more critical in the Brazilian context, given the historical upward asymmetry of inflationary pressures.

For $h = 0$, the predictive densities of both models appear well-calibrated, as their empirical CDFs align closely with the ideal 45-degree line across most of the inflation distribution. However, the QPC model shows signs of miscalibration in the lower tail of the distribution.

Across forecast horizons ranging from one quarter to one years ahead, the BCB densities predominantly fall below the 45-degree reference line, suggesting a systematic underestimation of realized inflation during the analyzed subperiod. This pattern also aligns with the findings for the entire inflation target period.

In contrast, the QPC model's predictive densities do not exhibit systematic directional bias. However, they display signs of overconfidence, as the predictive distribution is too narrow relative to the actual dispersion of inflation outcomes. Specifically, the observed inflation are more frequently extreme than the model predicts. This is indicated by inverse S-shaped deviations from the 45-degree reference line for certain horizons.

5
Risk Assessment on Inflation-Targeting Regime

As outlined in Chapter 1, assessing the risks associated with future inflation developments is a critical component of decision-making for economic agents.

Hecq, Issler and Voisin (2024) proposes a statistical warning mechanism for Central Banks under an IT regime. This risk measure considers if it is sustainable for inflation to stay within the actual tolerance bounds at any given horizon. In particular, it considers the conditional probabilities that inflation will be within the bounds at h period ahead.

Among the countries that adopt the widely used continuous-time inflation target framework, most of them define their target breach as a persistent deviation from the announced target or bands. In order, to evaluate the duration risk of deviation episodes, a multi-period consideration is required, extending beyond single-h period-ahead compliance assessment.

The primary contribution of this paper lies in proposing a more robust target risk measure for Central Banks by incorporating timing and duration dimensions. In the next sections, we assess models performance in gauging inflation target risk.

5.1
Single-period Inflation Target Risk Assessment

In this context, our analysis focuses on the conditional probability that inflation will fall outside bounds at a given horizon h:

$$\mathbb{P}\left(\pi_{T+h} > ub_{T+h} \vee \pi_{T+h} < lb_{T+h} \mid \mathcal{I}_T\right) = \tag{5-1}$$

$$\mathbb{P}\left(\pi_{T+h} > ub_{T+h} \mid \mathcal{I}_T\right) + \mathbb{P}\left(\pi_{T+h} < lb_{T+h} \mid \mathcal{I}_T\right)$$

Where lb_t and ub_t are the lower and upper bound target for time t and $\mathcal{I}_T$ is the information set in the period T which the forecast is computed.

5.1.1
Evaluation of BCB's forecasts for Inflation Target Period

A key question is whether the BCB's probabilistic forecasts can reliably assess inflation target breaches scenarios. Figure 5.1 illustrates the inflation dynamics within Inflation Target period.

In Figure 5.2, we present the corresponding conditional probabilities computed from the BCB's quarterly density forecasts. The solid line corresponds to a forecast made at time *t-h* of time-*t* inflation, with grey-shaded areas indicating periods when realized inflation missed the target.

Figure 5.1: 12-month Inflation Rate in Brazil and Target Bound

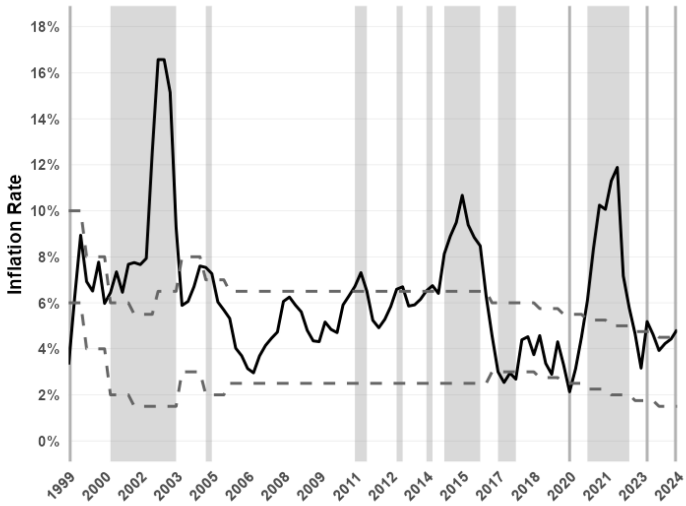

Figure 5.1: Displays the 12-month IPCA inflation rate in Brazil since the adoption of the Inflation Targeting regime. The dashed lines represent the upper and lower target bounds, while the grey-shaded areas highlight periods when realized inflation fell outside the target range.

Figure 5.2: BCB Probabilities of Inflation Target Breach at t

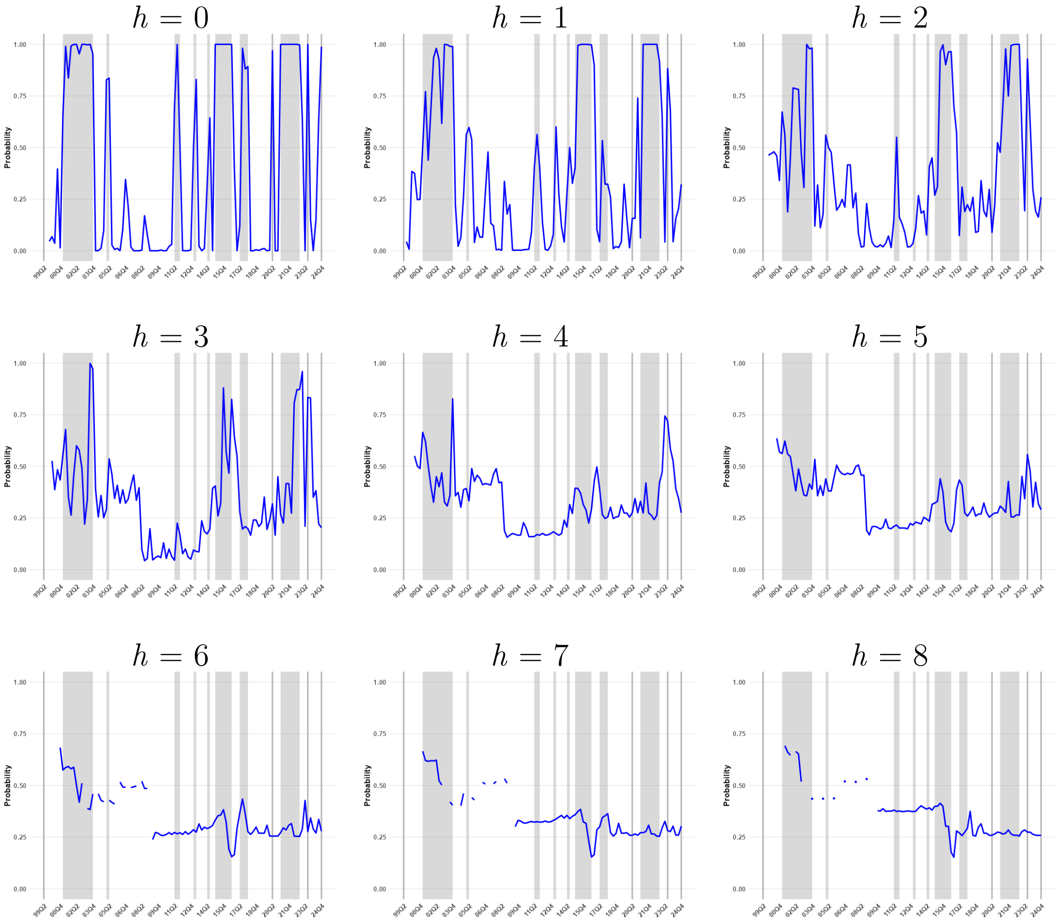

Figure 5.2: Each figure reports the conditional probability, computed from the BCB's predictive distributions (from 1999 to 2024), that inflation will be outside the bounds at a given horizon h. Solid line corresponds to a forecast made at time t-h of time-t inflation, with grey-shaded areas indicating periods when realized inflation missed the target.

An effective Inflation Target Risk monitoring tool would accurately assign, at time $t-h$, a higher probability of missing the target at time t in cases where realized inflation indeed deviates from the target. Specifically, the model should yield higher probabilities for grey-shaded areas periods (when realized inflation deviates from the target) and lower probabilities during non-shaded areas periods (when inflation remains within the target range).

Preliminary visual inspection indicates that, for forecast horizons up to three quarters ahead, the probabilities associated with the BCB's density forecasts exhibit a pattern reasonably aligned with the expected behavior of an Inflation Target Risk warning tool. Notably, the blue line tends to be higher during grey-shaded periods and lower during non-shaded periods.

For forecast horizons ranging from one year ahead up to two years ahead, the pattern shifts significantly. Several instances emerge where lower probabilities are associated with grey-shaded areas while higher probabilities appear in non-shaded periods. For instance, during the post-COVID inflation target breach period (grey-shaded area between 2021Q1 and 2022Q4), the associated probabilities for all forecast horizons beyond $h = 4$ were relatively low. Conversely, during much of the non-shaded period between 2003Q4 and 2008Q2, the associated probabilities were relatively higher.

Additionally, while probabilities exhibit significant variation for short-term forecast, from 2009Q4 onward, the inflation target breach risk appears to converge around 0.25 for longer horizons. This suggests that BCB's density forecasts often assign a 25% probability of an inflation target breach occurring at horizons ranging from one to two years ahead, regardless of the prevailing economic conditions and monetary policy stance at the time of the forecast.[1]

Overall, the analysis suggests that BCB's density forecasts may be useful for assessing an inflation target breach risk, particularly at short horizons. Nonetheless, such visual assessments are not appropriate for comparing models or conducting rigorous quantitative evaluations. To address this, we rely on a widely adopted method for evaluating the accuracy of probabilistic models: the *Receiver Operating Characteristic* (ROC) curve.

[1]In comparison, over the Inflation Targeting period, inflation remained outside the target band approximately 43% of the time. This unconditional pattern may reflect a statistical modeling artifact, a deliberate strategy to anchor long-term expectations around a constant risk level, or suggest overconfidence in its ability to meet the target.

5.1.1.1
Performance Metric for Probabilistic Forecast

In order to evaluate the accuracy of probability forecasts, the primary approach is to frame the problem as a binary classification task.

An *event* occurs when realized inflation breaches the target at time $t+h$, while a *nonevent* occurs if inflation remains within bounds:[2]

$$\begin{aligned} \textit{Event:} \quad & \pi^*_{t+h} < lb_{t+h} \quad \text{or} \quad \pi^*_{t+h} > ub_{t+h}, \\ \textit{Nonevent:} \quad & lb_{t+h} \leq \pi^*_{t+h} \leq ub_{t+h}. \end{aligned}$$

Define the indicator variable I^{Out}_{T+h} as:

$$I^{\text{Out}}_{T+h} = \begin{cases} 1, & \text{if an } \textit{Event} \text{ occurs,} \\ 0, & \text{if a } \textit{Nonevent} \text{ occurs.} \end{cases}$$

The estimated probability of a breach at $t+h$, given information at forecast origin T, is:

$$P^{\text{Out}}_{T+h} \equiv \hat{\mathbb{P}} \left(\pi_{T+h} > ub_{T+h} \vee \pi_{T+h} < lb_{T+h} \mid \mathcal{I}_T \right).$$

A *signal* occurs when P^{Out}_{T+h} exceeds a threshold $c \in (0,1)$, while a *nonsignal* occurs otherwise:

$$\textit{Signal}: \quad P^{\text{Out}}_{T+h} \geq c, \quad \textit{Nonsignal}: \quad P^{\text{Out}}_{T+h} < c.$$

The signal indicator function is:

$$\widehat{I}^{\text{Out}}_{T+h}(c) = \begin{cases} 1, & \text{if a } \textit{Signal} \text{ occurs,} \\ 0, & \text{if a } \textit{Nonsignal} \text{ occurs.} \end{cases}$$

The *True Positive Rate* (TPR) and *False Positive Rate* (FPR) are

[2]Hecq, Issler and Voisin (2024) defines *event* as inflation remaining within the target bounds at h. We opt to use the complementary formulation, as it aligns more directly with the concept of target risk and is also formulated in this manner in other studies, such as Galvão (2005). Naturally, accuracy results are invariant regarding definition of binary events.

defined as follows:

$$TPR(c) = \frac{\#(\widehat{I}^{\text{Out}}_{T+h}(c) = 1 \cap I^{\text{Out}}_{T+h} = 1)}{\#(I^{\text{Out}}_{T+h} = 1)},$$
$$FPR(c) = \frac{\#(\widehat{I}^{\text{Out}}_{T+h}(c) = 1 \cap I^{\text{Out}}_{T+h} = 0)}{\#(I^{\text{Out}}_{T+h} = 0)}.$$

The TPR measures the proportion of true signals—cases where the model correctly predicts a target breach—relative to all actual breach events. It reflects the model's accuracy in detecting actual target breach episodes. The FPR, in contrast, captures the proportion of false alarms—instances where the model incorrectly predicts a breach—relative to all non-breach events, indicating the rate of erroneous breach predictions.

The *Receiver Operating Characteristic (ROC)* curve plots $TPR(c)$ against $FPR(c)$ for all $c \in (0, 1)$, illustrating the trade-off between true and false positives. As $c \to 1$, no signals are emitted, yielding $TPR(c) = 0$ and $FPR(c) = 0$; as $c \to 0$, $TPR(c) = 1$ and $FPR(c) = 1$. For an uninformative model, for all $c \in (0, 1)$, $TPR(c) = FPR(c)$, producing a 45° ROC line; a perfect classifier traces the upper-left boundary of the unit square ($TPR(c) = 1$ and $FPR(c) = 0$).

Selecting a threshold c is critical—yet inherently subjective—when assessing inflation breach risk. Jorda and Berge (2011) links rational decision-making to a binary classification problem, demonstrating that the optimal choice of c depends on the agent's assumed utility function.

Alternatively, the *Area Under the ROC Curve* (AUC), defined as:

$$AUC = \int_0^1 ROC(r)\, dr, \tag{5-2}$$

serves as a summary statistic for classification models. An $AUC = 1$ indicates a perfect model, while $AUC = 0.5$ suggests performance no better than random guessing. The metric aggregates model performance across all classification thresholds c, eliminating the need to justify a specific cutoff. As shown by Yang and Pagan (2024), the use of a higher AUC model can be associated with greater expected utility from a rational decision-making perspective.

The Figure 5.3 displays the AUC values for the BCB's conditional probabilities of inflation falling outside target bounds at horizon h.[3]

Results are consistent with the preliminary visual analysis in Section 5.1.2, ratifying that probabilities derived from BCB's density forecasts

[3]The associated ROC curves are displayed in Appendix A.3.

Figure 5.3: AUC statistic of BCB's Single-Period Target Risk Probabilities for the Inflation Target Period

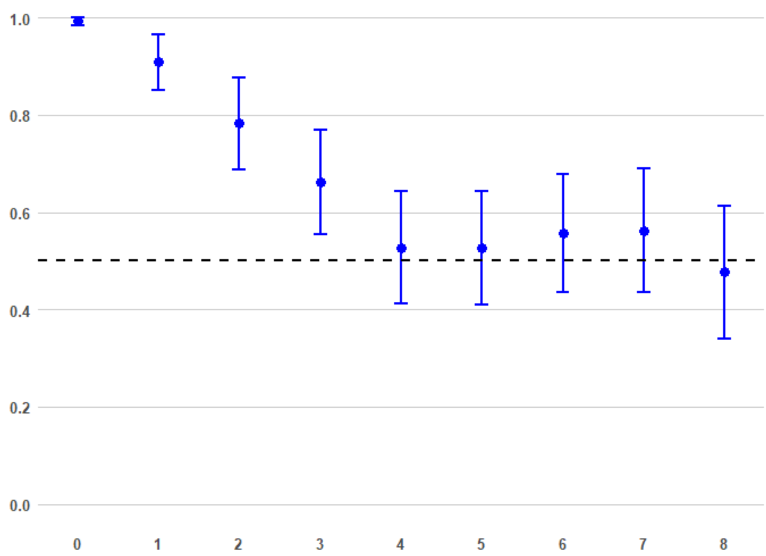

Figure 5.3: Displays the AUC values for BCB's conditional probability that inflation will be outside the bounds at a given horizon h. The plot includes 95% confidence intervals computed via DeLong method, and a dashed gray line at 0.5, representing the threshold for a random classifier. Probabilities are derived from BCB's predictive distributions from 1999 to 2024

are effective for short-term inflation target risk assessment. At the nowcasting for the end-month of the quarter ($h = 0$), the AUC is approximately 0.99, indicating that the model correctly distinguishes between breach and non-breach periods 99% of the time.

As the forecast horizon extends, predictive accuracy declines, with AUC values beyond one to two years ahead becoming statistically indistinguishable from a random classifier performance. This underscores the reliability of the BCB's predictive densities for short-term inflation risk (up to three quarters) but reveals limitations for medium and longer term risk assessment, including the six-quarter horizon currently emphasized for monetary policy decisions.

5.1.2
Comparison of QPC and BCB forecasts

For forecast horizons extending up to one year ahead, conditional probabilities 5-1 are derived from the QPC and BCB's quarterly density forecasts, over the 2013–2024 sample period.

Figure 5.4 presents the corresponding probabilities for this evaluation period and Figure 5.5 displays the AUC values.

Figure 5.4: QPC and BCB's Probabilities of Inflation Target Breach at t

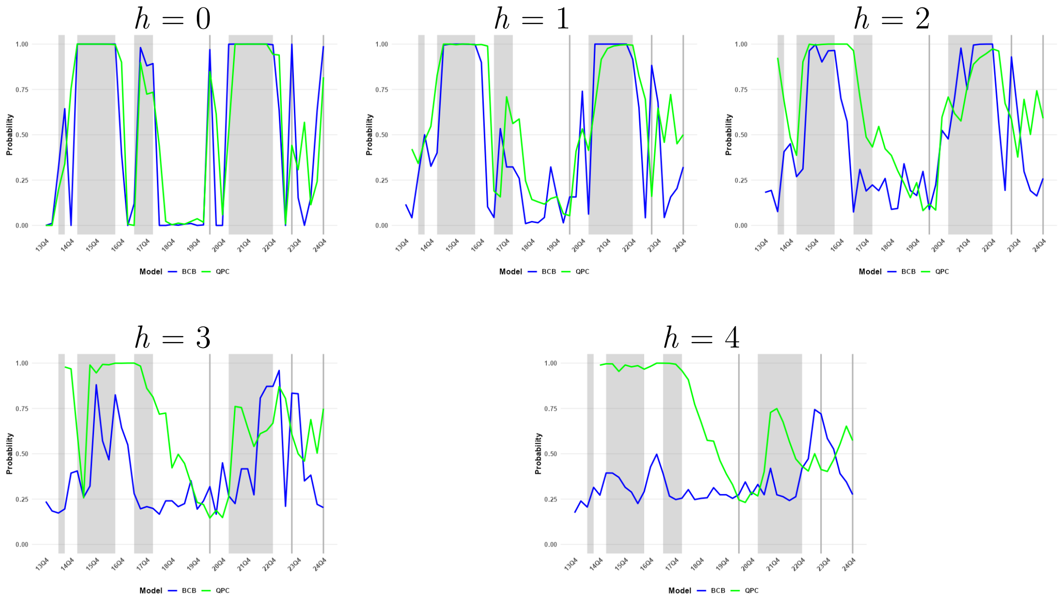

Figure 5.4: Each panel presents the conditional probability, computed from the QPC and BCB's predictive distributions (from 2013:Q4 to 2024:Q4), that inflation will be outside the bounds at a given horizon h. Solid line corresponds to a forecast made at time t-h of time-t inflation, with grey-shaded areas indicating periods when realized inflation missed the target.

Figure 5.5: *AUC* statistic of QPC and BCB's Single-Period Target Risk Probabilities

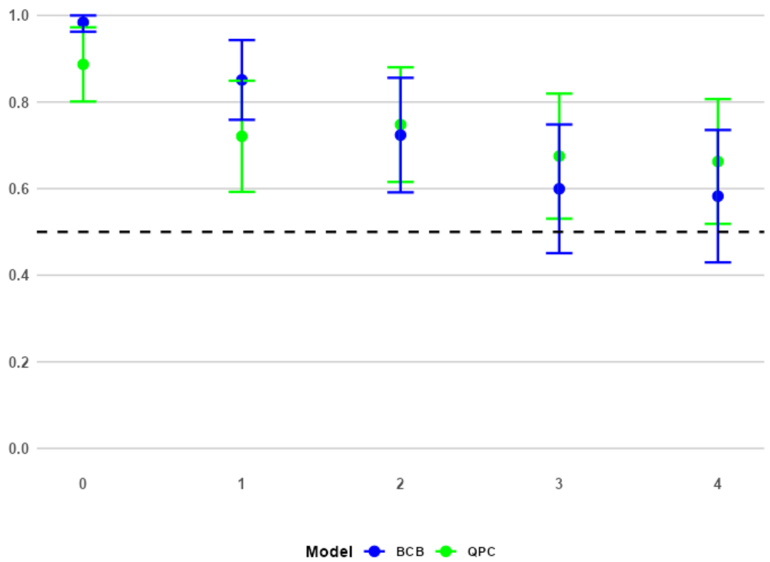

Figure 5.5: Displays the *AUC* values for QPC and BCB's conditional probability that inflation will be outside the bounds at a given horizon *h*. Plot includes 95% confidence intervals computed via DeLong method, and a dashed gray line at 0.5, representing the threshold for a random classifier. Probabilities are derived from predictive distributions from 2013:Q4 to 2024:Q4

Both models probabilities for short-term predictions in 5.4 exhibit close alignment and are consistent with the expected performance of an Inflation Target risk warning tool. Beyond two quarters, models signals seem to diverge.

During the 2017 disinflation period, BCB forecasts made in 2016 (h=3,4 quarters in advance) underestimated the inflation target risk, while the QPC model indicated elevated risk.

This discrepancy is even more pronounced during the post-COVID inflation overshoot, when BCB forecasts, made in 2020, failed to detect target risk. This misassessment suggests that the BCB underestimated upside risks when aggressively lowered the *Selic* rate to 2%, requiring sharp corrective rate hikes in the following quarters. Conversely, the QPC model accurately signaled increasing risks post-COVID at these horizons, though it overestimated risks following the 2017 disinflation period.

Results based on the *AUC* statistic, presented in 5.5,[4] indicate that both models effectively distinguish between breach and non-breach periods for horizons up to two quarters ahead. Specifically, the BCB's estimates exhibit higher accuracy for the nowcast and one-quarter-ahead horizon. However, for horizons of three quarters to one year ahead, the BCB's probabilities performance are

[4]The associated ROC curves are displayed in Appendix A.4.

statistically indistinguishable from a random model, highlighting the advantages of incorporating quantile Phillips Curve relationships for medium-term forecasting.

5.2 Multi-period Inflation Target Risk Assessment

In the previous section, we examined the conditional risk of an Inflation Target miss at a specific horizon h. As most countries define a breach as a sustained deviation from the announced target, a multi-period consideration is required, extending beyond single-h period-ahead compliance assessment.

We define the multi-period risk as the conditional joint probability that inflation falls outside the target bounds at horizon h and remains outside for at least k consecutive periods. This statistic captures not only the timing but also the persistence of a deviation event, providing a more robust target risk measure.

$$\mathbb{P}\left(\bigcap_{j=0}^{k-1}\left(\pi_{T+h+j} > ub_{T+h+j} \vee \pi_{T+h+j} < lb_{T+h+j}\right) \mid \mathcal{I}_T, k \leq H-h+1\right) = \tag{5-3}$$

$$\mathbb{P}\left(\pi_{T+h} \notin [lb_{T+h}, ub_{T+h}], \pi_{T+h+1} \notin [lb_{T+h+1}, ub_{T+h+1}], \ldots,\right.$$

$$\left.\pi_{T+h+k-1} \notin [lb_{T+h+k-1}, ub_{T+h+k-1}] \mid \mathcal{I}_T, k \leq H-h+1\right).$$

Note that the number of consecutive periods k is constrained by the maximum horizon H, which denotes the farthest projection being considered. The definition is flexible and ensures the framework can adapt to diverse IT regime requirements and economic conditions[5].

To compute the probability in (5-3), the conditional joint probability distribution function can be used:

$$F_{\boldsymbol{\pi}_{T+h:T+h+k-1}}\left(\boldsymbol{B} \mid \mathcal{I}_T\right) = \mathbb{P}\left(\pi_{T+h} \leq b_{T+h}, \ldots, \pi_{T+h+k-1} \leq b_{T+h+k-1} \mid \mathcal{I}_T\right), \tag{5-4}$$

where $\boldsymbol{\pi}_{T+h:T+h+k-1} = (\pi_{T+h}, \pi_{T+h+1}, \ldots, \pi_{T+h+k-1})'$ denotes a vector of future inflation rates and $\boldsymbol{B} = [lb, ub]^k$ is the set of bounds over k periods.

One may assume a predefined function (ie., Multivariate Normal Distribution) or use some other estimation method. This paper proposes to combine individual h-step-ahead predictive distributions into a joint forecast using a copula function. The method is fully flexible and requires only the horizon-

[5] It can incorporate endogenous $k(\cdot)$ and $H(\cdot)$, enabling, for instance, k to vary during periods of external shocks or adaptive projection horizons H. For a summary on different forecast horizons H within IT frameworks globally, refer to the "*Governance for the Communication of the Inflation Projections Horizon*" section Box in the BCB's September 2024 Inflation Report (available at:https://www.bcb.gov.br/content/ri/inflationreport/202409/ri202409b5i.pdf).

specific forecast densities as input, regardless of assumptions about the distributional form or the models used to generate them. This makes it broadly applicable for practitioners aiming to construct joint predictive distributions for inflation without imposing rigid parametric specifications.

5.2.1 Estimating Inflation Joint Probability Distribution via Copulas

For the single-period inflation target risk assessment, we used the estimated distributions from the BCB and QPC models to derive the conditional probabilities that inflation will be outside the bounds at a given h period ahead. In other words, this involves using the marginal predictive densities functions for each horizon to obtain the associated probabilities.

Our setup follows a *direct* forecasting scheme, that is individual predictions do not contain information on cross-horizon dependence. A simple approach would be to assume independence between the different marginal predictive densities (i.e., no correlation between the direct h-step-ahead predictions at different horizons) and compute the multi-period joint probability in (5-3)[6] by multiplying the single-period probabilities for each horizon:

$$\mathbb{P}\left(\bigcap_{j=0}^{k-1} \left(\pi_{T+h+j} > ub_{T+h+j} \vee \pi_{T+h+j} < lb_{T+h+j}\right)\right) =$$

$$\prod_{j=0}^{k-1} \left(\mathbb{P}\left(\pi_{T+h+j} > ub_{T+h+j}\right) + \mathbb{P}\left(\pi_{T+h+j} < lb_{T+h+j}\right)\right).$$

However, this procedure neglects the serial dependence common in macroeconomic indicators, notably pronounced in inflation dynamics.

To address this issue, we build on the method presented by Mogliani and Odendahl (2025). Their framework focuses on constructing predictive objects that depend on several horizons using direct density forecasts, while accounting for cross-horizon dependence.[7]

[6] Conditioning in $\mathcal{I}_T$ and $k \leq H - h + 1$ is dropped for better readability

[7] This method has also been applied in related contexts, such as in Charemza, Makarova and Wu (2018), where the authors use a copula-based approach to forecast the duration of short-term deflation episodes.

5.2.2 Methodology

The forecaster aims to estimate $F(\pi_1, \ldots, \pi_H)$, the multivariate distribution of $(\Pi_1, \ldots, \Pi_H)$, using copula functions. Sklar's theorem states that for any $F(\pi_1, \ldots, \pi_H)$, there exists a copula function $C(\cdot|R)$, such that:

$$F(\pi_1, \ldots, \pi_H) = C(G_{\Pi_1}(\pi_1), \ldots, G_{\Pi_H}(\pi_H)|R), \tag{5-5}$$

where $G_{\Pi_1}, \ldots, G_{\Pi_H}$ are the marginal CDFs and R parameters governs the dependence structure. Inversely, a copula function C, combined with marginal CDFs, reconstructs the full multivariate distribution.

We employ a Gaussian copula, C_{Ga}, with correlation matrix R, to estimate the joint predictive CDF, $F_T(\pi_{T+1}, \ldots, \pi_{T+H})$, from the marginal predictive CDFs $G_{\Pi_{T+1}}, \ldots, G_{\Pi_{T+H}}$ obtained in Chapter 3:[8]

$$F_T(\pi_{T+1}, \ldots, \pi_{T+H}) = C_{\text{Ga}}(G_{\Pi_{T+1}}(\pi_{T+1}), \ldots, G_{\Pi_{T+H}}(\pi_{T+H})|R) \tag{5-6}$$

The following Monte Carlo algorithm is applied to estimate joint probabilities:

[8]The copula approach is highly adaptable with respect to the marginals, irrespective of their distributional form or underlying model.

Algorithm 1: Estimating the Joint Predictive Distribution

1. Compute realized PITs:

$$\{\{\mathrm{PIT}_{t+h}\}_{h=1}^{H}\}_{t=1}^{T-H},$$

using estimated predictive CDFs $\left\{\left\{\widehat{G}_{\Pi_{t+h}}\right\}_{h=1}^{H}\right\}_{t=1}^{T-H}$ and realized inflation $\left\{\left\{\pi^{*}_{t+h}\right\}_{h=1}^{H}\right\}_{t=1}^{T-H}$.

2. Estimate the rank correlation matrix $\widehat{\mathbf{R}}$ from $\{\{\mathrm{PIT}_{t+h}\}_{h=1}^{H}\}_{t=1}^{T-H}$..

3. Compute the lower Cholesky decomposition of $\widehat{\mathbf{R}}$, denoted by $\mathbf{P}$.

4. For each $s = 1, \ldots, S$, where s represents a simulation round:

   (a) Draw a $H \times 1$ vector $\mathbf{X} \sim_{iid} \mathcal{N}(\mathbf{0}, \mathbf{I}_H)$.

   (b) Compute:

   $$[Z_1, ..., Z_H]' = \mathbf{Z} = \mathbf{PX}.$$

   Let $\mathbf{U} = [U_1, ..., U_H] = [\Phi(Z_1), ..., \Phi(Z_H)]$, where $\Phi(\cdot)$ is the CDF of a standard normal distribution.

   (c) Evaluate inverse CDFs:

   $$\widehat{G}^{-1}_{\Pi_{T+1}}(U_1), \ldots, \widehat{G}^{-1}_{\Pi_{T+H}}(U_H),$$

   to obtain simulated draws $[\pi^{s}_{T+1}, ..., \pi^{s}_{T+H}]'$, that maintains the dependence structure.

5. Use the simulated draws:

$$\{[\pi^{s}_{T+1}, ..., \pi^{s}_{T+H}]'\}_{s=1}^{S}$$

to approximate $\widehat{F}_T(\pi_{T+1}, \ldots, \pi_{T+H})$.

5.2.3
Aplication: Brazil Inflation Target Framework

The proposed multi-period target risk considers the probability (5-3) of inflation breaching the target at horizon h and remaining outside for k consecutive periods. While the duration of such deviations is typically not predefined, the new Brazilian Inflation Targeting framework defines a target as unmet if the 12-month inflation rate deviates beyond the tolerance interval for six consecutive months ($k = 6$). This explicit criterion positions Brazil as an ideal candidate for applying the introduced duration risk measure.

The continuous monthly target assessment prevents the evaluation of the BCB's predictive densities, as these are available only at quarter-end and defined over quarterly horizons. Accordingly, the analysis is based solely on probabilities derived from the QPC model, which can be updated at a monthly frequency.

Computation of (5-3) follows the algorithm previously described. The probability is estimated by drawing 10,000 joint simulated paths, using a maximum forecast horizon H of 12 months. Although the initial monthly forecast origin is December 2013, the use of a 32-month rolling window to estimate the correlation matrix $\widehat{\mathbf{R}}$ from empirical PITs shifts the effective starting point to August 2017.

To better interpret the results, Table 5.1 summarizes the output of a real out of sample forecast conducted in June 2021 using our QPC model. For each possible 6-month window within the 12-month forecast horizon, we display the joint probability of inflation remaining outside the target bounds for every month in that window.

Table 5.1: Example of Forecasted Probability of Inflation Breach Duration (%)

Forecast Date	**Forecast Window (Months Ahead)**							
	(0 – 5)	(1 – 6)	(2 – 7)	(3 – 8)	(4 – 9)	(5 – 10)	(6 – 11)	(7 – 12)
June 2021	68%	60%	47%	34%	24%	16%	12%	10%

The results indicate a moderate to high probability (68%) of consecutive breaches in the near term (June–November 2021). However, as the forecast horizon extends, the probability of a sustained deviation declines significantly, reaching only 10% for the last six-month window (January–June 2022). This suggests that, while inflation was likely to remain outside the target range in the short run, the medium-term outlook—arguably more relevant given monetary policy transmission lags—indicates a reversion toward the target. These findings imply that, as of June 2021, prevailing monetary policy may have been insufficient to bring inflation back within bounds immediately but remained consistent with stabilization in longer horizons.

The heatmap in Figure 5.6, extends the example of Table 5.1 to all forecast origins.

Figure 5.6: Heatmap of 6 Consecutive Breach Probabilities

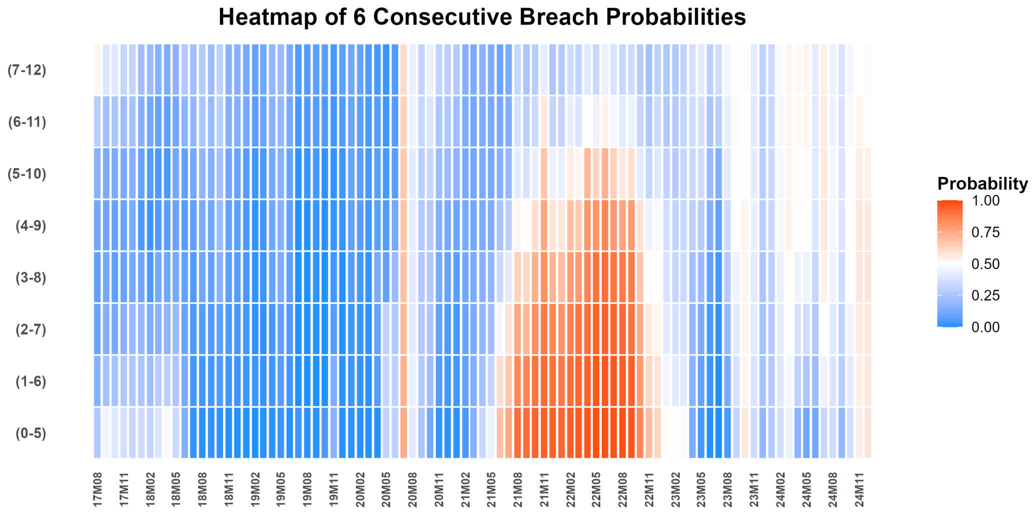

Figure 5.6: Display the conditional probability, computed from the QPC's predictive distributions, that inflation will be outside the target bound during every period of the considered 6-month windows.

Figure 5.7 reports the AUC values for the QPC's joint conditional probabilities of inflation falling outside the target bounds for $k = 6$ consecutive periods.[9]

The model's joint probabilities accurately distinguish between persistent and non-persistent breach periods in 91% of cases within the immediate six-month window (0 to 5 months ahead).

Predictive efficiency decreases with longer horizons, yet the AUC values indicate that the model remains informative for risk assessment over windows up to six months ahead. These findings highlight the reliability of the Quantile Phillips Curve in anticipating short-term persistent inflation events, offering valuable insights for continuous risk monitoring and effective policymaking.

[9] In this case, an *event* is defined as realized inflation breaching the target at $t + h$ and remaining outside bounds for five subsequent periods. Corresponding ROC curves are shown in Appendix A.5.

Figure 5.7: AUC statistic of QPC's Multi-Period Target Risk Probabilities

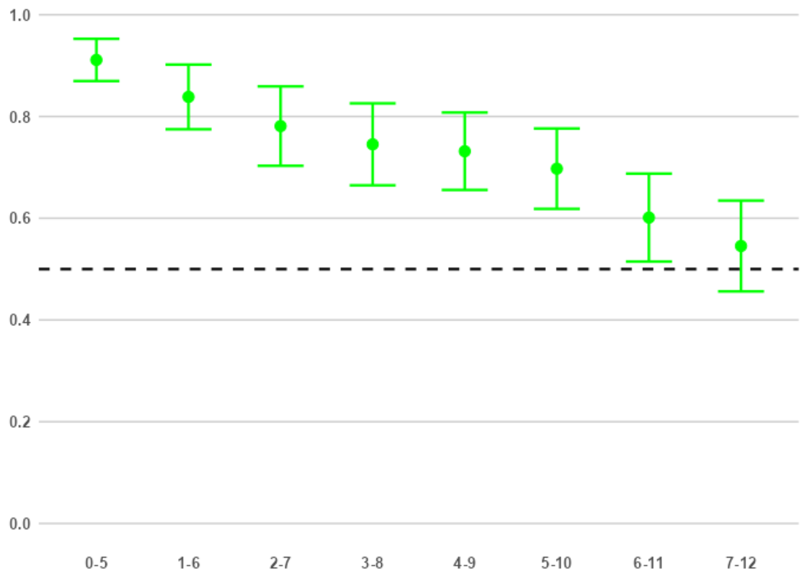

Figure 5.7: Displays the AUC values for QPC's conditional joint probability that inflation will be outside the bounds for each 6-months consecutive windows. Plot includes 95% confidence intervals computed via DeLong method, and a dashed gray line at 0.5, representing the threshold for a random classifier. Probabilities are derived from predictive distributions from 2017:M8 to 2024:M12

6
Conclusion

In complex and shock-prone economic environments, central banks must consider not only the most likely future path for the economy but also the distribution of possible outcomes about that path. This paper assessed the accuracy and informativeness of the Central Bank of Brazil's probabilistic inflation forecasts, a key tool in the country's inflation targeting framework. The analysis spans Brazil's full Inflation Targeting period from 1999 to 2024, while also examining the 2013:Q4–2024:Q4 subsample for comparative evaluation.

Findings show that BCB's probabilistic tool consistently underestimate inflation uncertainty, exhibiting significant miscalibration at key monetary policy horizons ($h = 6$ and $h = 8$ quarters). In contrast, the QPC model shows improved calibration properties, especially in the central and upper ranges of the inflation distribution—particularly relevant for Brazil's historically upward-skewed inflation dynamics.

We also evaluate both models capacity to generate informative inflation target risk assessment. Findings show that BCB's forecasts are only efficient at shorter horizons (up to 6-months ahead), whereas our proposed QPC model provides reliable assessments of inflation target risk even at medium-term horizons (up to 12-months ahead).

An original contribution of this study is a novel inflation target duration risk measure that integrates persistence and timing risk by assessing the probability of consecutive target breaches across different periods along the forecast horizon. Applied to the QPC model via a Copula-based approach, adaptable to diverse forecasting setups, it delivers informative early warnings of persistent inflation episodes up to six months ahead. This advancement enhances forward-looking risk management by signaling prolonged deviations from the inflation target, offering nuanced insights for more effective policymaking.

In conclusion, while the BCB's Inflation Fan Chart remains a valuable communication tool, its forecasting accuracy and risk assessment capabilities are limited at short-term horizons. The QPC model suggested in this study addresses these shortcomings by providing policymakers with improved tools to navigate Brazil's volatile inflation dynamics. Future research could extend this framework to other economies, incorporate additional macroeconomic variables, or explore alternative models to further refine probabilistic forecasting under inflation-targeting regimes.

7
Bibliography

ADRIAN, T.; BOYARCHENKO, N.; GIANNONE, D. Vulnerable growth. **The American Economic Review**, American Economic Association, v. 109, n. 4, p. 1263–1289, 2019. ISSN 00028282, 19447981.

AZZALINI, A.; CAPITANIO, A. Distributions generated by perturbation of symmetry with emphasis on a multivariate skew t-distribution. **Journal of the Royal Statistical Society. Series B (Statistical Methodology)**, [Royal Statistical Society, Wiley], v. 65, n. 2, p. 367–389, 2003. ISSN 13697412, 14679868.

BANERJEE, R. N. et al. Fiscal deficits and inflation risks: the role of fiscal and monetary policy regimes. Bank for International Settlements, 2022.

BLANCHARD, O.; CERUTTI, E.; SUMMERS, L. Inflation and activity – two explorations and their monetary policy implications. n. 21726 National Bureau of Economic Research Working Paper Series, 2015.

CHAREMZA, W.; MAKAROVA, S.; WU, Y. Forecasting the duration of short-term deflation episodes. **Journal of Forecasting**, v. 37, n. 4, p. 475–488, 2018.

CHERNOZHUKOV, V.; FERNáNDEZ-VAL, I.; GALICHON, A. Quantile and probability curves without crossing. **Econometrica**, v. 78, n. 3, p. 1093–1125, 2010.

DIEBOLD, F. X. et al. Evaluating density forecasts with applications to financial risk management. **International Economic Review**, [Economics Department of the University of Pennsylvania, Wiley, Institute of Social and Economic Research, Osaka University], v. 39, n. 4, p. 863–883, 1998.

GALVÃO, A. B. C. Uma avaliação do leque das previsões de inflação do banco central do brasil. Instituto de Pesquisa Econômica Aplicada (Ipea), 2005.

HECQ, A.; ISSLER, J. V.; VOISIN, E. A short term credibility index for central banks under inflation targeting: An application to brazil. **Journal of International Money and Finance**, v. 143, p. 103057, 2024. ISSN 0261-5606.

INOUE, A. Testing for distributional change in time series. **Econometric Theory**, v. 17, n. 1, p. 156–187, 2001.

JORDA, O.; BERGE, T. J. Evaluating the classification of economic activity into recessions and expansions. **American Economic Journal: Macroeconomics**, v. 3, n. 2, p. 246–77, April 2011.

KNUPPEL, M.; SCHULTEFRANKENFELD, G. Assessing the uncertainty in central banks' inflation outlooks. **International Journal of Forecasting**, v. 35, n. 4, p. 1748–1769, 2019. ISSN 0169-2070.

LOPEZ-SALIDO, D.; LORIA, F. Inflation at risk. **Journal of Monetary Economics**, v. 145, p. 103570, 2024. ISSN 0304-3932. Inflation: Expectations Dynamics October 14-15, 2022.

MARCELLINO, M.; STOCK, J. H.; WATSON, M. W. A comparison of direct and iterated multistep ar methods for forecasting macroeconomic time series. **Journal of Econometrics**, v. 135, n. 1, p. 499–526, 2006. ISSN 0304-4076.

MITCHELL, J.; HALL, S. G. Evaluating, comparing and combining density forecasts using the klic with an application to the bank of england and niesr 'fan' charts of inflation. **Oxford Bulletin of Economics and Statistics**, v. 67, n. s1, p. 995–1033, 2005.

MOGLIANI, M.; ODENDAHL, F. Density forecast transformations. n. 2511 Documentos de Trabajo Banco de España, 2025.

ROSSI, B.; SEKHPOSYAN, T. Alternative tests for correct specification of conditional predictive densities. **Journal of Econometrics**, v. 208, n. 2, p. 638–657, 2019.

TAGLIABRACCI, A. Asymmetry in the conditional distribution of euro-area inflation. Bank of Italy Temi di Discussione (Working Paper), v. 1270, 2020.

YANG, K. L. L.; PAGAN, A. Getting the roc into sync. **Journal of Business & Economic Statistics**, ASA Website, v. 42, n. 1, p. 109–121, 2024.

A
Appendix

A.1
Inference Exercise

In Table A.1, we report the coefficients of each explanatory variable from an in-sample estimation of the Quantile Phillips Curve model (Equation 3-1), using data from January 2004 to December 2024.

Expected inflation significantly influences inflation forecasts across all quantiles (5th, 50th, and 95th), with its positive effect diminishing at higher quantiles, underscoring the importance of forward-looking expectations across inflation regimes. Inflation inertia, however, gains prominence in the upper quantile (95th), indicating a stronger role in extreme inflation scenarios. The output gap significantly affects the lower (5th) and median (50th) quantiles but shows no notable impact in the upper (95th) quantile. In contrast, fiscal deficits exhibit a significant influence across all parts of the distribution, highlighting fiscal policy's relevance across inflation regimes.

In summary, upside inflation risk (i.e., movements in the right tail of the distribution) is primarily associated with higher fiscal deficits, elevated inflation expectations, and stronger inflation inertia. Conversely, downside inflation risk (i.e., movements in the left tail) is linked to fiscal surpluses, lower inflation expectations, and a weaker output gap.

Table A.1: Dependent Variable: 12-Months Ahead Inflation

	Quantile Estimates		
Variable	**5th**	**50th**	**95th**
Expected Inflation	**0.91**	**0.85**	**0.73**
	[0.77–1.05]	[0.72–0.97]	[0.52–0.93]
Current Inflation	0.09	**0.15**	**0.27**
	[-0.05–0.23]	[0.03–0.28]	[-0.48–0.48]
Output Gap	**0.84**	**0.99**	0.48
	[0.45–1.22]	[0.56–1.41]	[-0.32–1.28]
Financial Conditions	-0.14	**-0.61**	-0.52
	[-0.28–0.56]	[-0.14––1.09]	[-0.37–1.41]
Imported Inflation	0.02	**0.11**	0.06
	[-0.02–0.07]	[0.07–0.16]	[-0.03–0.15]
Fiscal Deficit	**0.32**	**0.80**	**0.82**
	[0.03–0.61]	[0.52–1.08]	[0.29–1.35]

Note: Coefficients in bold are statistically significant. Brackets show 95% confidence intervals based on standard errors estimated using 5,000 block-bootstrap replications.

A.2 Descriptive Statistics

Table A.2: Descriptive Statistics

Variable	**Mean**	**Std. Dev.**	**Min**	**Max**
Current Inflation	5.74	2.11	1.88	12.13
Expected Inflation	4.26	0.65	3.00	5.50
Output Gap	-0.07	2.41	-7.91	3.27
Imported Inflation	0.70	3.47	-7.14	11.84
Financial Conditions	-0.06	1.19	-2.50	3.06
Fiscal Deficit	-0.91	2.52	-4.08	9.24

Note: Data span monthly observations from January 2004 to December 2024 (N = 203).

A.3
BCB *ROC* Curves

Figure A.1: BCB Single-period Inflation Target Risk Assessment ROC Curve

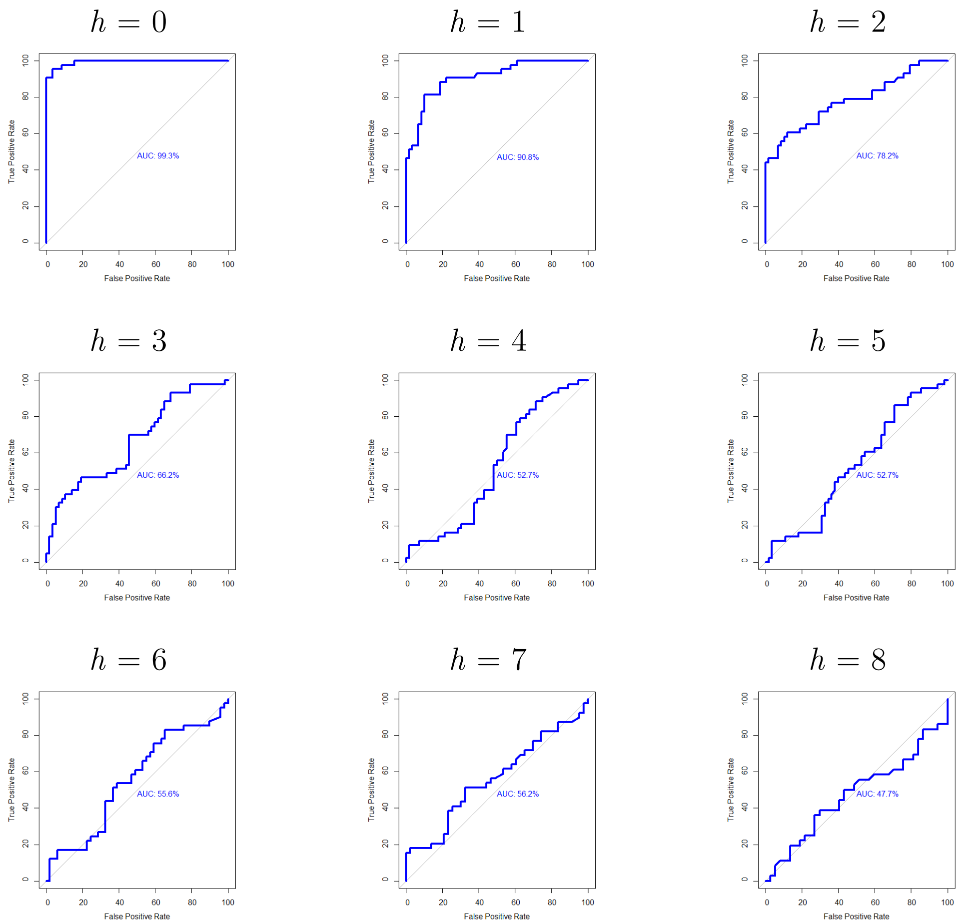

Figure A.1: Each panel displays the ROC curve for BCB's conditional probability that inflation will be outside the bounds at a given horizon. The curve plots $TPR(c)$ against $FPR(c)$ for all $c \in (0, 1)$. Sample encompasses all available projections during from 1999:Q2 to 2024:Q4.

A.4
QPC and BCB *ROC* Curves

Figure A.2: BCB and QPC Single-period Inflation Target Risk Assessment *ROC* Curves

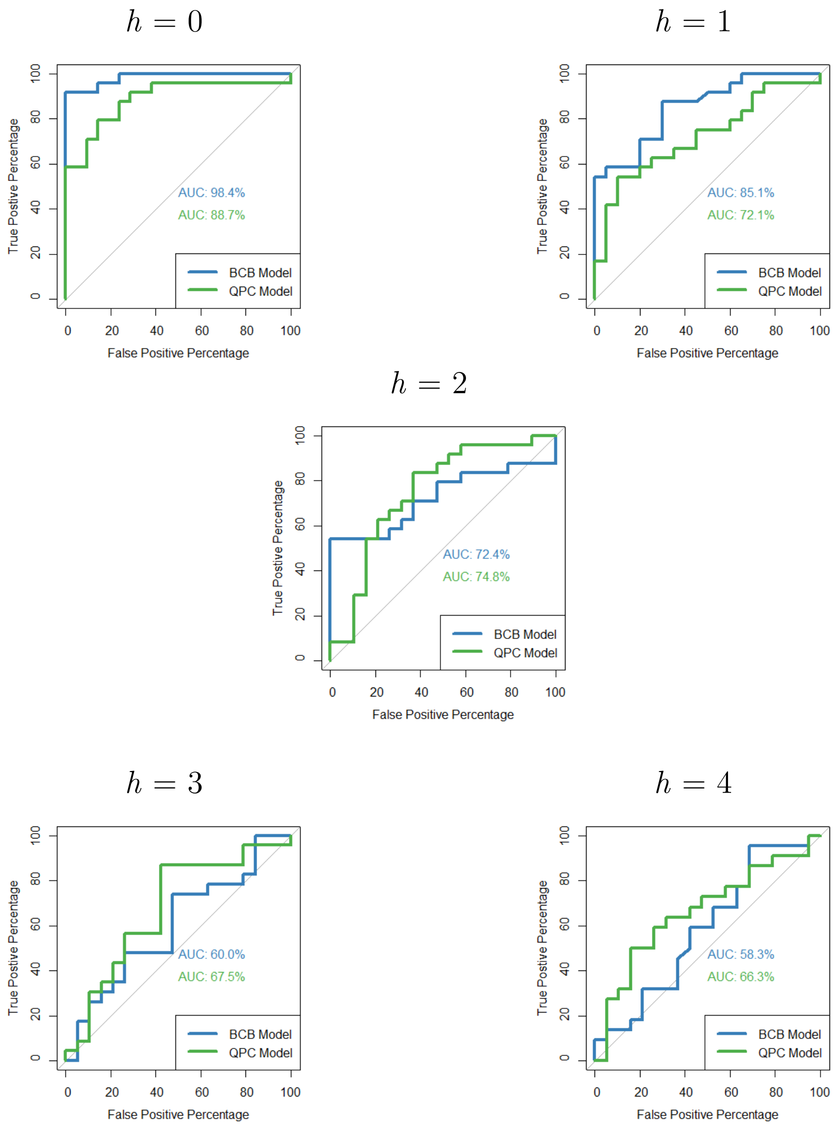

Figure A.2: Each panel displays the *ROC* curve for QPC and BCB's conditional probability that inflation will be outside the bounds at a given horizon. The curve plots $TPR(c)$ against $FPR(c)$ for all $c \in (0, 1)$. Sample encompasses all available projections from 2013:Q4 to 2024:Q4.

A.5
QPC Multi-Period *ROC* Curves

Figure A.3: QPC Multi-Period Inflation Target Risk Assessment *ROC* Curves

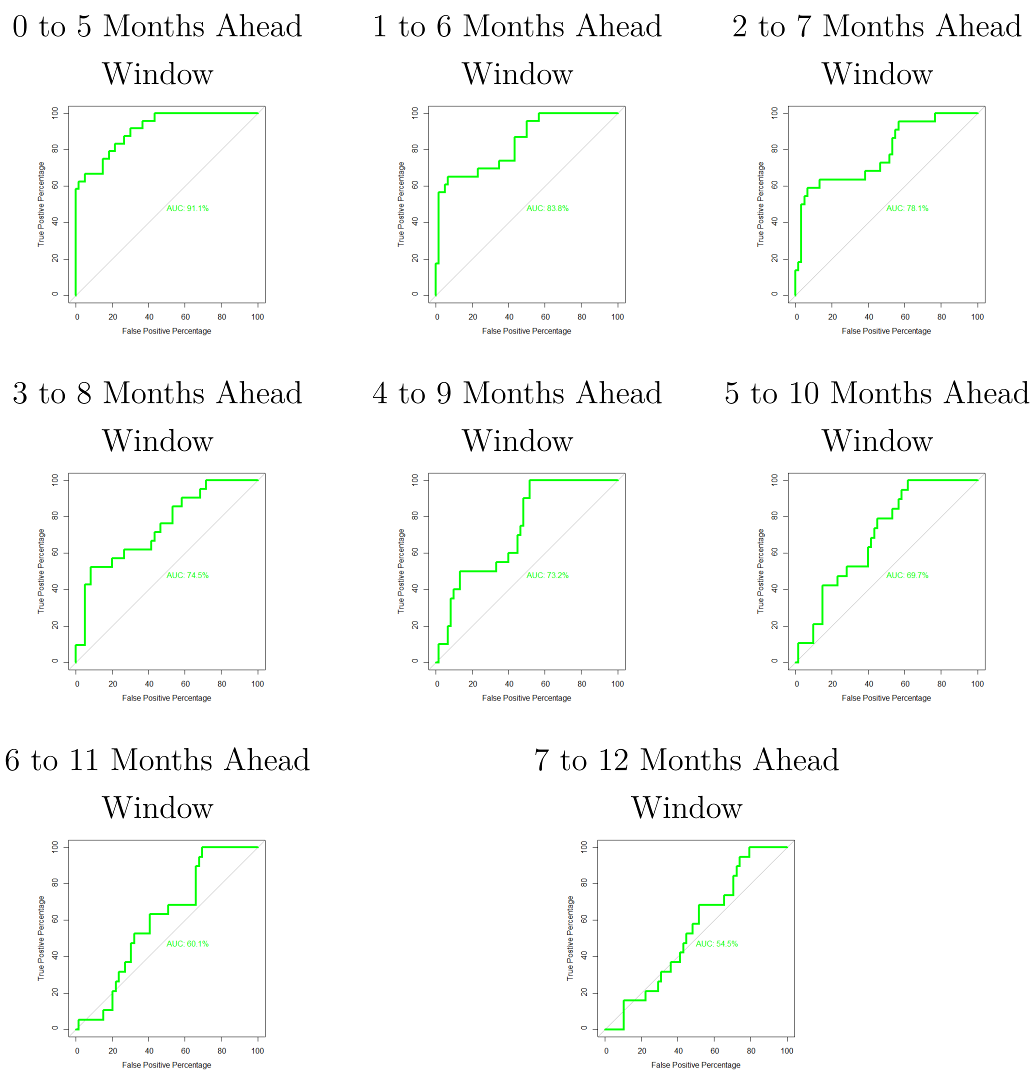

Figure A.3: Each panel displays the *ROC* curve for the QPC's joint conditional probabilities of inflation falling outside the target bounds for $k = 6$ consecutive periods. The curve plots $TPR(c)$ against $FPR(c)$ for all $c \in (0, 1)$. Sample encompasses all available projections from 2017:M8 to 2024:M12.