Unforeseeable Change in Rational Participants' Inflation Expectations: Evidence from Forecast-Error Regressions*

Roman Frydman† Morten Nyboe Tabor‡

September 1, 2025

INET Center on Knightian Uncertainty Working Paper 02

Abstract

We present empirical evidence that the bias in market participants' ex post inflation forecast errors and these errors' correlation with ex ante information shift between nonrepetitive values. This finding is inconsistent with the predictions of full- and limited-information rational expectations (REH) models. Building on Knight's and Muth's pathbreaking insights, our theoretical account reconciles our empirical findings with participants' rationality by recognizing that the inflation process undergoes nonrepetitive structural change that cannot be foreseen, even in probabilistic terms. Because REH models abstract from such unforeseeable structural change, they do not represent expectations of rational participants in real-world markets.

Keywords: Inflation Expectations; Forecast Errors; Structural Shifts; Unforeseeable Change; Knightian Uncertainty; Muth's Hypothesis;

JEL Codes: D83, D84, E31, E37.

*The authors are participants in the INET Center on Knightian Uncertainty. They thank Søren Johansen and Joshua Stillwagon for providing insightful comments and constructive suggestions on the draft of this paper. The authors are grateful to the Institute for New Economic Thinking for its support of their research.

†Department of Economics, New York University. E-mail: rf3@nyu.edu.

‡INET Center on Knightian Uncertainty. E-mail: mortentabor@gmail.com

1 Introduction

Macroeconomists have relied on the rational expectations hypothesis (REH) to represent rational participants' expectations of outcomes, such as inflation, which we focus on here. Most REH models abstract from structural change in the inflation process, which they formalize by constraining the model's structure and parameters to remain constant. In such models, full-information REH (FIRE) predicts that participants' ex post forecast errors should be unbiased and uncorrelated with ex ante information.

Hajdini and Kurmann (2024) show, however, that this prediction does not hold if the inflation process is assumed to switch between two repetitive regimes governed by a Markov chain. In such a model, FIRE predicts that the bias of participants' ex post inflation forecast errors and their correlation with ex ante information should switch between repetitive positive and negative values.

Using survey data of participants' inflation forecasts, this paper uncovers novel empirical evidence that is inconsistent with FIRE's predictions. We find that the bias and correlation of participants' ex post forecast errors with ex ante information shift between nonrepetitive values.

Our findings lead us to advance a novel explanation of FIRE's empirical difficulties: REH models *do not* represent expectations of rational, profit-seeking participants in real-world markets, regardless of whether participants base their expectations on full or limited information. Building on Knight's (1921) and Muth's (1961) pathbreaking insights, the paper's main theoretical contribution is to reconcile our empirical findings with the rationality of market participants' inflation expectations.

Our empirical evidence is based on two widely used forecast-error regressions of ex post inflation-forecast errors on either ex ante forecast revisions or ex ante inflation. We construct measures of forecast errors and forecast revisions using forecast data from the Survey of Professional Forecasters covering the sample period from 1969 to 2022.

Our econometric procedure does not constrain a priori whether, when, and how the intercept and slope of these regressions shift over time. Instead, we identify the timing and estimate the magnitude of shifts in these parameters from time-series data. To do so, we use step-indicator and multiplicative step-indicator saturation to transform the identification of the timing of structural shifts into a high-dimensional model-selection problem that can be efficiently handled by the Autometrics algorithm (see Doornik, 2009, Castle, Doornik, and Hendry, 2012, Ericsson, 2012, and Castle et al., 2015).

This procedure is particularly suitable for assessing whether the parameters of the forecast-error regressions undergo shifts between repetitive or nonrepetitive values, because it neither restricts the maximum number of such shifts nor constrains the parameters to shift between a fixed number of repetitive regimes chosen prior to estimation. Moreover, it enables us to identify shifts in the intercept and slope at different points in time and throughout the entire sample period. Thus, by design, the econometric procedure nests the possibility that the parameters of the forecast-error regressions have remained constant or switched between repetitive values.

We find periods during which the intercepts and/or slopes of the forecast-error regressions have been nonzero, indicating that participants' ex post inflation-forecast errors have been biased and/or correlated with ex ante information. Importantly, these biases and correlations have undergone major shifts between nonrepetitive values with different magnitudes—including zero, positive, and negative—and statistical significances.

Most of the structural shifts, as well as the largest bias and correlation of ex post forecast errors with ex ante information, occurred during the 1970s, early 1980s, and post-2020. These were periods characterized by major structural shifts in inflation dynamics and expectations caused by the oil crises, the sharp change in monetary policy following Paul Volcker's appointment as chairman of the Fed, and the onset of the Covid-19 pandemic. In contrast, we find only one structural shift, very small estimates of the bias, and insignificant estimates of the correlation between participants' ex post forecast errors and ex ante forecast revisions during the Great Moderation from the early 1980s to 2020.

The literature has sought to account for a supposedly positive correlation between ex post forecast errors and ex ante forecast revisions, which Coibion and Gorodnichenko (2015, abstract, p. 2644) suggest is a "stylized fact." Except for Hajdini and Kurmann (2024), these accounts have presumed that a constant-parameter model characterizes inflation. Coibion and Gorodnichenko (2015) suggest the positive correlation arises because participants are rational but face limited or noisy information. Farmer, Nakamura, and Steinsson (2024), in the context of interest-rate expectations, suggest it arises from rational participants' Bayesian learning about the model's constant parameters. In contrast, Bordalo et al. (2020) attribute this positive correlation to participants' irrationality: participants' expectations deviate—in systematic, and thus predictable ways—from FIRE due to the distorting influence of Kahneman and Tversky's (1972) "representativeness heuristic," which they formalize with Gennaioli and Shleifer's (2018) diagnostic expectations.

Our empirical findings show, however, that the supposedly positive correlation between ex post forecast errors and ex ante forecast revisions is not a stylized fact, but rather an artifact of estimating this correlation as constant. Indeed, we find a significantly positive correlation only during the period from 1972 to 1974. This corroborates Angeletos, Huo, and Sastry's (2021) and Hajdini and Kurmann's (2024) findings that the estimate of this correlation changes from significantly positive to insignificantly different from zero when the estimation samples 1968:Q4-2017:Q4 and 1970:Q2-2019:Q1 are changed to start in 1984:Q1 and 1983:Q1, respectively.

Our findings show that participants' forecast errors undergo nonrepetitive structural shifts and that these shifts' timing and frequency reflect, at least in part, historical events. This suggests that accounting for these errors requires jettisoning constant-parameter models of inflation, as well as models that formalize structural change in the inflation process as repetitive. Consequently, our theoretical account of participants' ex post forecast errors recognizes that the inflation process undergoes structural change that is nonrepetitive and whose magnitude and timing therefore cannot be foreseen, even in probabilistic terms, by rational participants in real-world markets.

REH models do not represent expectations of rational participants in real-world markets because they assume that structural change in the inflation process is foreseeable: the process is assumed either to remain unchanging or to switch between repetitive regimes governed by a Markov chain. In both cases, a single conditional probability distribution—whose parameters are constant and can be consistently estimated based on historical data—constitutes the model's prediction of future inflation. Because REH represents participants' inflation expectations with the conditional expectation of this distribution, it presumes that market participants can foresee, in probabilistic terms, the magnitude and timing of all future structural changes in the inflation process. This implies that the bias and correlation of their ex post forecast errors with ex ante information should be constant, and equal to zero, or switch between repetitive non-zero values, which we found was inconsistent with the survey forecast data.

Our theoretical account rests on a novel formalization of Knight's argument that participants in real-world markets face "true uncertainty," which arises from change in the economy that cannot "by any method be reduced to an objective, quantitatively determinate probability" (Knight, 1921, pp. 231-232). Importantly, Knight argued that profit-seeking activities inherently involve attempts to exploit such unforeseeable change, and thus cannot be understood in terms of standard probabilistic "risk." As he put it: "if all changes were to take place in accordance with invariable universally known laws, [so that] they could be foreseen for an indefinite period in advance of their occurrence (...) profit or loss would not arise" (Knight, 1921, p. 198).

In Frydman and Tabor (2024), we propose an approach to representing rational participants' expectations of inflation, and other outcomes, undergoing unforeseeable change. We call our approach the Knight-Muth hypothesis (KMH).

KMH provides a tractable and empirically testable formalization of unforeseeable structural change in the inflation process by assuming that its models' parameters shift intermittently between nonrepetitive values with timings that are not characterized with a probabilistic rule. Because future parameters differ from past ones, their values are inherently unknowable in advance, even in probabilistic terms. This implies that future inflation is characterized by Knightian uncertainty, not just probabilistic risk.

KMH proposes a novel implementation of Muth's (1961) hypothesis to represent rational market participants' inflation expectations as being consistent with the model's predictions of inflation undergoing unforeseeable change. We assume that participants base their expectations on full information. In contrast to FIRE, however, KMH represents rational participants' inflation expectations in terms of subjective parameters that can differ from the parameters that characterize ex post inflation. KMH recognizes that, over time, rational market participants revise their subjective expectations in nonrepetitive ways as the inflation process shifts in unforeseeable ways.

We show that KMH predicts that the bias and correlation of rational market participants' ex post forecast errors with ex ante information intermittently shift between nonrepetitive values. During periods where the inflation process undergoes major unforeseeable change, the bias and correlation should be large and either positive or negative. In contrast, the bias and correlation should be small, or even zero, during periods where the inflation process remains stable.

By recognizing that rational participants' ex ante inflation expectations may differ from how ex post inflation unfolds, KMH explains why most of the structural shifts and the largest bias and correlation of ex post forecast errors with ex ante information occurred during the 1970s, early 1980s, and post-2020. Acknowledging that it was particularly difficult for rational market participants to foresee inflation during those periods is crucial for KMH's account of their ex post forecast errors. It is no less crucial for KMH that when the inflation process has been relatively stable, as during the Great Moderation, rational participants' forecast errors have been nearly unbiased and uncorrelated with ex ante information, which we have found.

The structure of the paper is as follows. Section 2 describes the time-series data we use for our empirical analysis. Section 3 presents our empirical estimates of structural shifts in the forecast-error regressions and shows that these estimates are inconsistent with the predictions of limited- and full-information REH models. In Section 4, we illustrate how KMH can account for the unforeseeable structural shifts in the forecast-error regressions. We compare KMH's and FIRE's predictions of those shifts, and trace FIRE's inconsistency with the survey forecast data to its premise that economists and market participants can foresee structural change in the inflation process in probabilistic terms. Concluding Section 5 argues that a "stylized fact" with which formal representations of rational expectations in macroeconomic models should be consistent is that profit-seeking participants understand that the process driving inflation and other aggregate variables undergoes unforeseeable change. We suggest that macroeconomists should not rely on models that abstract from such change and, consequently, reconsider the widespread belief that REH represents rational participants' expectations of aggregate outcomes in real-world markets.

2 Survey Data on Inflation Expectations

To construct forecast errors and revisions for the forecast-error regressions, we use quarterly survey data on U.S. inflation expectations from the Survey of Professional Forecasters (SPF), available from the Federal Reserve Bank of Philadelphia. Established in 1968, the survey has been answered each quarter by a rotating panel of professional forecasters. We consider the average forecasts of the annualized change in the price index for gross domestic product (PGDP) for four-quarter-ahead forecast horizons from 1968:Q4 to 2022:Q4.

Like Coibion and Gorodnichenko (2015) and Hajdini and Kurmann (2024), we focus on a baseline case with an annual forecast horizon and construct the average quarterly inflation forecast and forecast errors over this horizon, as described in Supplemental Appendix A. Because the professional forecasters report their forecasts in the middle of each quarter, before the quarter's inflation data has been released, they are based on the previous quarter's inflation data. To better match the theoretical specification of expectations, we follow Hajdini and Kurmann (2024) and specify the timing of the survey forecasts in terms of the information it is based on instead of when it is reported. Thus, we denote the inflation rate by π_t and the survey forecasts of inflation based on the information available at the end of period t by $F_t(\pi_{t+h})$ for the forecast horizon $h > 0$. Our baseline case with an annual horizon corresponds to $h = 4$.[1]

As is common in the literature, we construct forecast errors using the initial release of the quarterly change in the PGDP, expressed in annualized percentage points, prepared by the Bureau of Economic Research. In Supplemental Appendix C, we show that our empirical findings are robust to alternative data specifications and forecast horizons.

There are missing values for the survey forecast for the five-quarter horizon in 1968:Q4–1969:Q2, 1969:Q4, and 1974:Q2, which implies missing values for the forecast revision for the annual horizon in the subsequent observations. Thus, for the baseline case, we use a sample for the forecast-error regressions that starts in 1970:Q2, and for the regression that includes the forecast revision, we drop the observation with a missing value for the forecast revision in 1974:Q3, as is commonly done in the literature.

Figure 1 shows the time-series data for inflation, the survey forecasts of inflation, and the computed forecast errors and revisions. Panel (a) shows the inflation rate (black line), π_t, and each of the grey lines shows the sequence of survey forecasts of inflation, $F_t(\pi_{t+h})$, for forecast horizons $h = 1, 2, 3, 4, 5$, starting from the latest observed inflation rate, π_t. Panels (b) and (c) show the computed ex post forecast errors, $\pi_{t+4} - F_t(\pi_{t+4})$, and ex ante forecast revisions, $F_t(\pi_{t+4}) - F_{t-1}(\pi_{t+4})$, for the baseline case with a forecast horizon $h = 4$. The forecast errors and revisions for forecast horizons $h = 1, 2, 3, 5$ show

[1] Our baseline case, $h = 4$, corresponds to Coibion and Gorodnichenko's (2015) baseline horizon $h = 3$ because they specify the survey forecasts in terms of the period in which it is reported. This difference in the timing of the survey forecasts does not alter the empirical results.

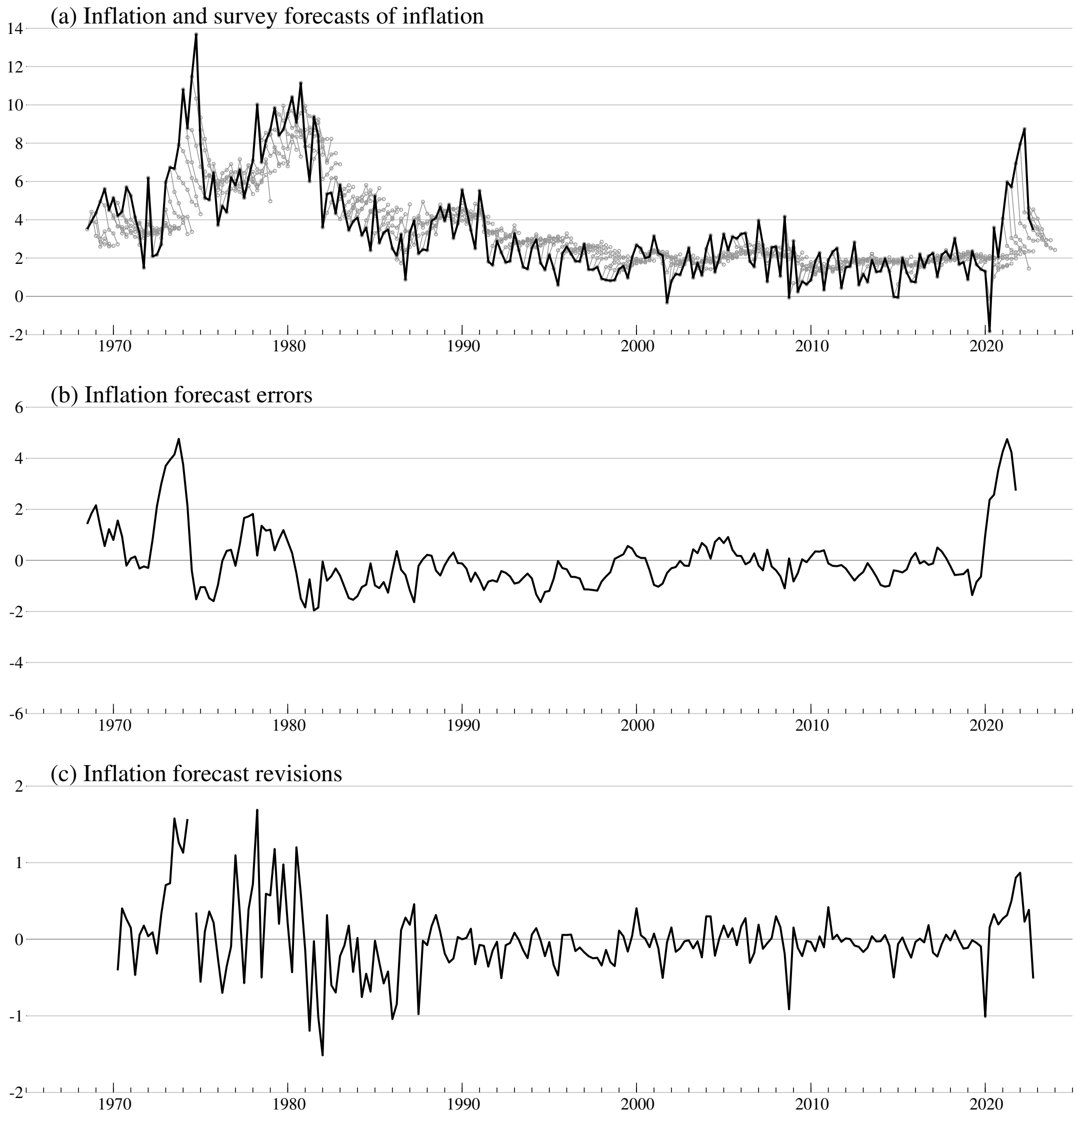

Figure 1: The figure shows inflation, survey forecasts of inflation, and forecast errors and revisions for the baseline case with an annual forecast horizon of $h = 4$. In Panel (a), the black line shows the inflation rate, π_t. The grey lines show the sequences of survey forecasts of inflation, $F_t(\pi_{t+h})$, for horizons $h = 1, 2, 3, 4, 5$ and starting from the latest observed inflation rate, π_t. Panel (b) shows the computed quarterly ex post survey forecast errors, $\pi_{t+4} - F_t(\pi_{t+4})$, and Panel (c) the computed ex ante revisions of the survey forecasts, $F_t(\pi_{t+4}) - F_{t-1}(\pi_{t+4})$, for the baseline case with a forecast horizon of $h = 4$.

similar patterns as for $h = 4$ over time, but with the magnitude of the forecast errors and revisions increasing with the forecast horizon.

3 Empirical Evidence of Unforeseeable Shifts in Forecast-Error Regressions

We consider two forecast-error regressions that have been widely used in the literature to test the predictions of alternative specifications of participants' inflation expectations.

The first is the regression of ex post forecast errors on ex ante forecast revisions, given by:

$$\pi_{t+h} - F_t(\pi_{t+h}) = \alpha^{(1)} + \beta^{(1)} \left(F_t(\pi_{t+h}) - F_{t-1}(\pi_{t+h})\right) + u_t^{(1)}, \tag{1}$$

for $t = 1, 2, \ldots, T$ and where π_t denotes inflation at time t; $\pi_{t+h} - F_t(\pi_{t+h})$ denotes the average quarterly inflation forecast error over the horizon from t to $t + h$, based on the information available at time t; $F_t(\pi_{t+h}) - F_{t-1}(\pi_{t+h})$ denotes the ex ante revision from time $t - 1$ to t of the average forecast over the horizon until $t + h$; and $u_t^{(1)}$ is an error term. This regression was proposed by Coibion and Gorodnichenko (2015) to test full-information rational expectations (FIRE) against the alternative of limited information rational expectations (LIRE) based on constant-parameter models of inflation and other macroeconomic variables.

The second forecast-error regression is that of ex post inflation forecast errors on the ex ante inflation rate, given by:

$$\pi_{t+h} - F_t(\pi_{t+h}) = \alpha^{(2)} + \beta^{(2)}\pi_t + u_t^{(2)}, \tag{2}$$

for $t = 1, 2, \ldots, T$ and where $u_t^{(2)}$ is an error term. As Hajdini and Kurmann (2024) show, this regression provides a simple and intuitive link to FIRE's representation of participants' expectations in models that assume that the inflation process undergoes repetitive regime switches governed by a two-state Markov chain.

Earlier studies testing predictions of alternative specifications of participants' expectations with survey forecast data have focused on estimating the forecast-error regressions in (1) and (2) with constant parameters. For comparison with that literature, Table 1 presents the constant-parameter estimates for the baseline case with a forecast horizon $h = 4$ based on our full sample from 1970:Q2 to 2021:Q4.

We get a positive and significant estimate of the correlation between ex post forecast errors and ex ante forecast revisions in (1), $\widehat{\beta}^{(1)} = 1.122$ with a standard error of 0.412, and an insignificant estimate of the intercept. For (2), we get a positive, but insignificant, estimate of the correlation between ex post forecast errors and ex ante inflation, $\widehat{\beta}^{(2)} = 0.090$ with a standard error of 0.069, and a negative, but insignificant, estimate of the

Table 1: Constant-Parameter Estimates of the Forecast-Error Regressions

Panel A: $\pi_{t+h} - F_t(\pi_{t+h}) = \alpha^{(1)} + \beta^{(1)} \left(F_t(\pi_{t+h}) - F_{t-1}(\pi_{t+h})\right) + u_t^{(1)}$

	Estimate	Std. error	P-value
$\alpha^{(1)}$	-0.014	0.135	0.919
$\beta^{(1)}$	1.122	0.412	0.007
Residual std. error, σ			1.122
R^2			0.163
Log-likelihood			-314.932
Observations			206

Panel B: $\pi_{t+h} - F_t(\pi_{t+h}) = \alpha^{(2)} + \beta^{(2)}\pi_t + u_t^{(2)}$

	Estimate	Std. error	P-value
$\alpha^{(2)}$	-0.329	0.183	0.073
$\beta^{(2)}$	0.090	0.069	0.192
Residual std. error, σ			1.201
R^2			0.037
Log-likelihood			-330.558
Observations			207

Notes: The table shows the constant-parameter estimates of the forecast-error regressions in (1) and (2) for the baseline case with a forecast horizon of $h = 4$. The sample period is 1970:Q2–2021:Q4. For regression (1) in Panel A, the observation with missing data in 1974:Q3 has been dropped from the estimation sample. HAC standard errors are reported and used to compute the P-values.

intercept.

Our insignificant estimate of $\beta^{(2)}$ in (2) is similar to that reported by Hajdini and Kurmann (2024). In contrast, the full-sample estimate of the slope $\beta^{(1)}$ in (1) is significantly positive. This estimate is similar to that reported by Coibion and Gorodnichenko (2015), Bordalo et al. (2020), Angeletos, Huo, and Sastry (2021), and Hajdini and Kurmann (2024). This similarity across many different studies appears to buttress Coibion and Gorodnichenko's (2015, abstract) claim that the positive correlation between the ex post forecast error and ex ante forecast revision is one of the "stylized facts" about participants' inflation expectations.

As we show next, however, the parameters of the forecast-error regressions shift during the sample, which is inconsistent with a constant-parameter characterization of the inflation process. Moreover, once the shifts in the parameters of (1) are identified from the time-series data, the estimate of $\beta^{(1)}$, and thus of the correlation between ex post forecast errors and ex ante forecast revisions, is significantly positive only in the early 1970s, and insignificantly different from zero during the rest of the sample period.

3.1 Subsample Estimates

Several studies have found that the estimates of the forecast-error regressions change when the estimation sample is changed. For example, Hajdini and Kurmann (2024) show that the estimated correlation between participants' ex post forecast errors and their ex ante forecast revisions becomes much smaller and changes from being significantly positive to insignificantly different from zero when the estimation sample 1970:Q2-2019:Q1 is changed to start in 1983:Q1. Angeletos, Huo, and Sastry (2021) present similar findings when changing their sample from 1968:Q4-2017:Q4 to start in 1984.

These studies' choice of the starting dates of their samples was based on the plausible conjecture that the inflation process and/or participants' inflation expectations underwent structural shifts around the onset of the Great Moderation. It also seems plausible that inflation and/or participants' expectations, and thus their forecast errors, underwent structural shifts around the onset of Covid-19 pandemic in early 2020.

Table 2 presents the estimates of the two forecast-error regressions for three distinct subsamples: the 1970s subsample covering the period from 1970:Q2 to 1979:Q3, the Great Moderation subsample covering the period from 1980:Q1 to 2019:Q4, and the post-2020 subsample covering the period from 2020:Q1 to 2021:Q4. The full-sample and subsamples estimates of the forecast-error regression in (1) are illustrated in Figure 2, which shows the scatterplots between ex post forecast errors and ex ante forecast revisions with regression lines corresponding to the intercept and slope of (1).[2]

[2]Marginally changing the timing of the three subsamples, for example to match the subsamples used by Hajdini and Kurmann (2024) or Angeletos, Huo, and Sastry (2021), does not alter our conclusions here.

Table 2: Subsample Estimates of the Forecast-Error Regressions

Panel A: $\pi_{t+h} - F_t(\pi_{t+h}) = \alpha^{(1)} + \beta^{(1)}\left(F_t(\pi_{t+h}) - F_{t-1}(\pi_{t+h})\right) + u_t^{(1)}$			
	Estimate	Std. error	P-value
Subsample 1970:Q2–1979:Q4 (38 obs.)			
$\alpha^{(1)}$	0.535	0.347	0.132
$\beta^{(1)}$	1.117	0.595	0.069
Subsample 1980:Q1–2019:Q4 (160 obs.)			
$\alpha^{(1)}$	-0.401	0.087	0.000
$\beta^{(1)}$	0.072	0.139	0.604
Subsample 2020:Q1–2021:Q4 (8 obs.)			
$\alpha^{(1)}$	2.896	0.414	0.000
$\beta^{(1)}$	1.482	0.509	0.027
Panel B: $\pi_{t+h} - F_t(\pi_{t+h}) = \alpha^{(2)} + \beta^{(2)}\pi_t + u_t^{(2)}$			
	Estimate	Std. error	P-value
Subsample 1970:Q2–1979:Q4 (39 obs.)			
$\alpha^{(2)}$	0.781	0.837	0.357
$\beta^{(2)}$	0.017	0.122	0.891
Subsample 1980:Q1–2019:Q4 (160 obs.)			
$\alpha^{(2)}$	-0.189	0.135	0.165
$\beta^{(2)}$	-0.084	0.047	0.072
Subsample 2020:Q1–2021:Q4 (8 obs.)			
$\alpha^{(2)}$	2.366	0.471	0.002
$\beta^{(2)}$	0.235	0.102	0.060

Notes: The table shows the subsample estimates of the forecast-error regressions in (1) and (2) for the baseline case with a forecast horizon of $h = 4$. For regression (1) in Panel A, the observation with missing data in 1974:Q3 has been dropped from the estimation sample. HAC standard errors are reported and used to compute the P-values.

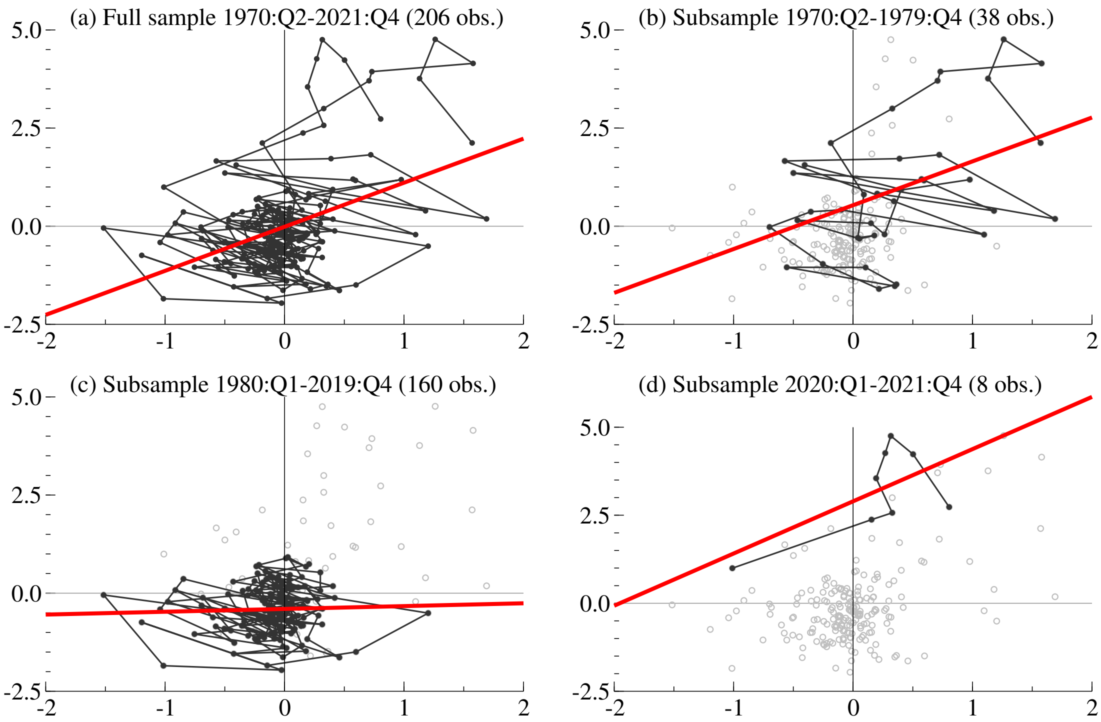

Figure 2: The figure shows scatterplots of the ex ante forecast revisions, $F_t(\pi_{t+4}) - F_{t-1}(\pi_{t+4})$, on the horizontal axes and the ex post forecast errors, $\pi_{t+4} - F_t(\pi_{t+4})$, on the vertical axes for the baseline case with a forecast horizon of $h = 4$. Panel (a) shows the scatterplot for the full sample 1970:Q2-2021:Q4 and the red line shows the regression line for this sample period. Panels (b)-(d) show the scatterplots and regression lines for the three subsamples 1970:Q2-1979:Q4, 1980:Q1-2019:Q4, and 2020:Q1-2021:Q4. The grey circles show the observations outside each of these subsamples.

The subsample estimates in Table 2 show that the correlation between ex post forecast errors and ex ante forecast revisions appears to be positive during 1970s, but turns statistically insignificant at the onset of the Great Moderation. This evidence suggests that the positive correlation, as estimated as constant by Coibion and Gorodnichenko (2015) and others, neither characterizes participants' ex post forecast errors nor their inflation expectations.

Similarly, the constant estimate of the intercept in (1) as statistically insignificantly also appears not to be a general feature throughout the sample period. The subsample estimates show that the average level of the ex post forecast errors was positive during the 1970s and post-2020 subsamples, but significantly negative during the Great Moderation. Thus, the insignificant full-sample constant-parameter estimate of the intercept might be a result of the intercept shifting between positive and negative values over the sample period. Indeed, the subsample estimates indicate that the intercepts and slopes of the forecast-error regressions shift substantially over the sample period.

When macroeconomic models recognize structural change in the process driving outcomes and participants' expectations, they typically assume that such change is repetitive and specify it as governed by a Markov chain. To investigate whether FIRE based on a model of inflation switching between repetitive regimes can account for the survey forecast data, Hajdini and Kurmann (2024) estimate a DSGE model with a monetary policy rule whose parameters switch according to a two-state Markov chain. They find that the rolling-window estimates of the correlations between ex post forecast errors and ex ante forecast revisions or ex ante inflation based on data simulated from the model "are overall quite different" from the rolling-window estimates of those correlations based on survey forecast data. Thus, as they put it, "the DSGE [model] we use as the [data generating process] to test FIRE fails to account for the dynamics of (...) inflation forecasts observed in the [survey] data" (p. 33). This, they conclude, provides "empirical motivation to consider models with alternative types of regime shifts as well as potential departures from FIRE to account for the dynamics of average expectations observed in the data" (Hajdini and Kurmann, 2024, pp. 35-36).

The estimates in Table 2 suggest that the key to accounting for the "dynamics of (...) expectations observed in the [survey] data" is to recognize that the structural shifts in the inflation process have not been repetitive and that, as a result, their magnitudes and timing could not have been foreseen, even in probabilistic terms. The estimated bias and correlation of participants' ex post forecast errors with ex ante information do not appear to switch between repetitive values. Moreover, the shifts appear to occur proximate to unprecedented historical events: the oil crises of the 1970s, the radical change in monetary policy following Paul Volcker's appointment as Federal Reserve Chair in late 1979, and the onset of the Covid-19 pandemic in the early 2020.

Although the subsample estimates in Table 2 are suggestive of which "types of regime shifts" in inflation and participants' expectations might explain structural shifts in the forecast-error regressions, they suffer from important limitations.

The subperiods for each subsample, and thus the timing of the potential structural shifts, were chosen a priori, rather than identified *from the time-series data*. Moreover, estimating both parameters of each forecast-error regression based on pre-set subsamples of the data rules out the possibility that the intercept and the slope may shift at different times. Although the estimated parameters are very different in each of the subsamples, there here are only three pre-set regimes. Moreover, the first two subsamples in Table 2 span long stretches of time. As we show next, identifying whether the parameters of the forecast-error regressions underwent shifts *within* those subperiods *from time-series data* is important for assessing whether the parameters of the forecast-error regressions underwent nonrepetitive shifts during the sample period.

3.2 Identifying Structural Shifts from the Time-Series Data

To identify the timing and estimate the magnitude of structural shifts in the parameters of the forecast-error regressions directly from the time-series data, we use step-indicator and multiplicative step-indicator saturation to transform the identification of structural shifts into a high-dimensional model-selection problem that can be efficiently handled by the Autometrics algorithm (see Doornik, 2009, Castle, Doornik, and Hendry, 2012, Ericsson, 2012, and Castle et al., 2015). We illustrate this procedure in Supplemental Appendix B.

This procedure enables us to identify nonrepetitive structural shifts, because it does not require specifying a priori the timing or maximum number of structural shifts during the sample period; nor does it require restricting the parameters to shift between a fixed set of, say, two regimes. Moreover, it enables the identification of structural shifts in the intercept and slope at different points in time and throughout the entire sample period (except for the last observation, where a shift in the slope cannot be distinguished from a shift in the intercept).

Table 3 reports the estimates of (1) and (2) with structural shifts identified from the time-series data by Autometrics for our baseline case with a forecast horizon $h = 4$. Figure 3 illustrates how the parameter estimates have evolved over time. Panels (a) and (c) plot the estimates of the two intercept terms over time, together with the inflation forecast errors, $\pi_{t+4} - F_t(\pi_{t+4})$.

Comparison of the estimates based on a priori subsamples in Table 2 with those in Table 3 shows that identifying structural shifts directly from time-series data uncovers multiple structural shifts in both the intercepts and the slopes of the two forecast-error regressions. Moreover, when the structural shifts in the parameters are identified from the time-series data, the forecast-error regressions can account for a much larger part of the variation in ex post forecast errors. As Tables 1 and 3 show, R^2 increases from 0.16

and 0.04 for the full-sample constant-parameter model to 0.76 and 0.82 for the model estimated with structural shifts.

Estimates presented in Table 3 and displayed in Figure 3 show that there are periods during which the intercepts and/or slopes of these regressions were nonzero, indicating that participants' inflation forecast errors were biased and/or correlated with ex ante information. Importantly, these biases and correlations underwent major shifts between nonrepetitive values with different magnitudes—including zero, positive, and negative signs—and statistical significances.[3]

The timing of structural shifts in the forecast-error regressions suggests that the changes in the bias of participants' inflation-forecast errors and their correlation with ex ante information reflect, at least partly, unforeseeable change in the inflation process and revisions of participants' inflation expectations in anticipation of and response to these changes. In that case, we should expect the ex post forecast errors to have the highest average level and correlation with ex ante information, and to undergo most structural shifts, during periods of major structural change in the economy and inflation dynamics.

Indeed, we find that most of the structural shifts and the largest estimates of both the intercepts and slopes occurred during the 1970s, the early 1980s, and at the end of the sample, in the post-2020 period. These were periods characterized by major changes in the economy, monetary policy, and inflation dynamics: the surge in inflation following the oil crises in the 1970s, the rapid disinflation following the appointment of Volcker as Fed Chair in 1979, and the rapid post-pandemic increase in inflation from 2020.

In contrast, we find only one structural shift during the Great Moderation period from the early 1980s to 2020 (the shift in ex post forecast errors' correlation with ex ante inflation in 1998:Q4), which is typically referred to as a stable period without major structural change in the economy, monetary policy, and inflation dynamics.

Because historical events like OPEC's decision to limit the supply of oil in 1973, the appointment of Volcker in 1979, and the Covid-19 crisis are not repetitive, at least in part, no one could have foreseen their timing and implications for the inflation process and market participants' inflation expectations. For example, prior to Volcker's appointment, no one could have foreseen, even in probabilistic terms, how and when the new monetary policy would lead to rapid disinflation. Though participants' inflation expectations adjusted gradually to the disinflation and stabilized, their forecast errors shifted from positive, on average, at the end of the 1970s to negative, on average, from the early 1980s.

[3] In Supplemental Appendix C, we show that our baseline specification's conclusions regarding the timing and magnitude of structural shifts in the parameter estimates, and these shifts' statistical significance, are robust to alternative specifications of the data used. These alternative specifications include using the final release of the inflation data instead of the initial vintage release to construct forecast errors, using the forecast error of quarter-to-quarter inflation at horizon h instead of the average forecast error over the horizon until $t + h$, and using different forecast horizons.

Table 3: Estimates of the Forecast-Error Regressions with Structural Shifts

Panel A: $\pi_{t+h} - F_t(\pi_{t+h}) = \alpha^{(1)} + \beta^{(1)} \left(F_t(\pi_{t+h}) - F_{t-1}(\pi_{t+h})\right) + u_t^{(1)}$

Parameter	Period	Estimate	Std. error	P-value
$\alpha^{(1)}$	1970:Q2–1972:Q2	0.298	0.251	0.236
	1972:Q3–1974:Q2	2.565	0.123	0.000
	1974:Q4–1976:Q1	-1.269	0.040	0.000
	1976:Q2–1980:Q2	0.871	0.235	0.000
	1980:Q3–2019:Q4	-0.384	0.089	0.000
	2020:Q1–2021:Q4	3.126	0.493	0.000
$\beta^{(1)}$	1970:Q2–1972:Q2	-0.674	0.506	0.185
	1972:Q3–1974:Q1	1.325	0.128	0.000
	1974:Q2–1980:Q4	-0.200	0.139	0.152
	1981:Q1–2021:Q4	0.299	0.209	0.154
σ				0.613
R^2				0.760
Log-likelihood				-186.456
Observations				206

Panel B: $\pi_{t+h} - F_t(\pi_{t+h}) = \alpha^{(2)} + \beta^{(2)}\pi_t + u_t^{(2)}$

Parameter	Period	Estimate	Std. error	P-value
$\alpha^{(2)}$	1970:Q2–1976:Q1	0.074	0.395	0.851
	1976:Q2–1977:Q2	0.965	0.181	0.000
	1977:Q3–1980:Q2	2.129	0.244	0.000
	1980:Q3–2019:Q4	-0.241	0.098	0.015
	2020:Q1–2020:Q3	1.952	0.239	0.000
	2020:Q4–2021:Q4	3.770	0.315	0.000
$\beta^{(2)}$	1970:Q2–1972:Q2	0.037	0.099	0.710
	1972:Q3–1973:Q4	0.621	0.060	0.000
	1974:Q1–1974:Q2	0.298	0.043	0.000
	1974:Q3–1998:Q3	-0.131	0.025	0.000
	1998:Q4–2021:Q4	0.028	0.058	0.631
σ				0.533
R^2				0.818
Log-likelihood				-157.892
Observations				207

Notes: The table shows the estimates of the forecast-error regressions in (1) and (2) with structural shifts for the baseline case with a forecast horizon of $h = 4$. The sample period is 1970:Q2-2021:Q4, where the observation with missing data for the forecast revision in 1974:Q3 has been dropped for the estimation of (1) in Panel A. HAC standard errors are reported and used to compute the P-values.

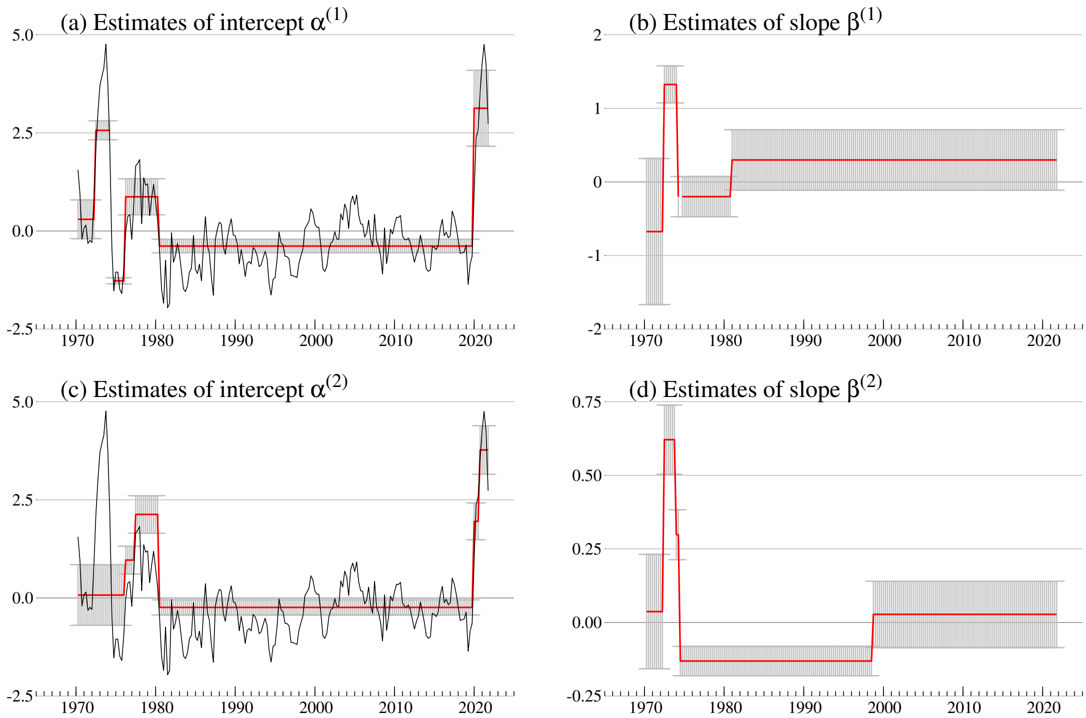

Figure 3: The figure shows the estimates of the parameters in the forecast-error regressions with structural shifts. Panels (a) and (b) show the estimates of $\alpha^{(1)}$ and $\beta^{(1)}$ (red lines) in (1) over time and their 95 percent confidence intervals (vertical grey bars) based on HAC standard errors. Panels (c) and (d) show the estimates of $\alpha^{(2)}$ and $\beta^{(2)}$ (red lines) in (2) and their 95 percent confidence intervals (vertical grey bars) based on HAC standard errors. Panels (a) and (c) also show the inflation forecast errors (black line), $\pi_{t+4} - F_t(\pi_{t+4})$.

Our empirical finding of structural shifts is inconsistent with most macroeconomic and finance models' core assumption that both the inflation process and participants' inflation expectations can be formalized with constant-parameter models. These include FIRE and LIRE, as well as behavioral models such as diagnostic expectations (Bordalo et al., 2020). In contrast to our empirical findings, a constant-parameter representation of the inflation process and participants' expectations implies that the parameters of the forecast-error regressions should have remained constant over time.

Moreover, identifying structural shifts from time-series data reveals that the constant-parameter estimates of an insignificant intercept and a significantly positive correlation between ex post forecast errors and ex ante forecast revisions are an artifact of the assumption that these parameters have remained constant.

The estimates presented in Table 3 show that the intercepts of each regression shifted between positive and negative values of substantial size during different subperiods of the sample. Moreover, the estimates are significantly different from zero in each of these subperiods, except the one at the beginning of the sample. Panels (a) and (c) in Figure 3 illustrate how these shifts in the intercept capture the large shifts between positive and negative average forecast errors during the 1970s, the negative average forecast errors during the Great Moderation period, and the large positive average forecast errors from 2020. Because the constant-parameter specifications of the forecast-error regressions ignore these shifts in the intercepts, they yield estimates that appear to be insignificant.

Moreover, the estimates of the slope in Table 3 show a positive correlation between ex post forecast errors and ex ante forecast revisions only during the short period from 1972:Q3 to 1974:Q1. The estimated correlation during this subperiod, $\widehat{\beta}^{(1)} = 1.325$ with a standard error of 0.128, captures the period's rapidly increasing and historically high inflation rates and large, positive forecast errors and revisions, as illustrated in Figure 1 and Panels (b) and (d) in Figure 3. The estimated correlation becomes negative, $\widehat{\beta} = -0.200$ with a standard error of 0.139, from 1974:Q2 to 1980:Q4, and positive, $\widehat{\beta}^{(1)} = 0.299$ with a standard error of 0.209, from 1981:Q1 until the end of the sample, but the estimates during both subperiods are insignificantly different from zero.

The finding that ex post forecast errors have been uncorrelated with ex ante forecast revisions during most subperiods of the sample suggests that participants' forecast errors do not arise from information rigidities. Coibion and Gorodnichenko (2015, pp. 2671-2675) argue that the degree of information rigidities might have changed over time and that it likely increased during the Great Moderation period. However, this would imply that the correlation between ex post forecast errors and ex ante forecast revisions, even if undergoing structural change, should have been positive throughout the sample period and should have increased during the Great Moderation period. This prediction is inconsistent with our empirical findings.

Our empirical findings are also inconsistent with Bordalo et al.'s (2020) prediction of a positive correlation between ex post forecast errors and ex ante forecast revisions. Moreover, Bordalo et al. show that diagnostic expectations implies that the intercept term in (1) should be zero. Our empirical findings show, however, that it has shifted between positive and negative estimates during different subperiods and has been significantly different from zero during all subperiods except the subperiod at the beginning of our sample, from 1970:Q2 to 1972:Q2.

In contrast, our estimates of nonrepetitive structural shifts in the forecast-error regressions may help explain Hajdini and Kurmann's (2024) findings. As discussed above, they found that FIRE-implied forecast errors simulated from a DSGE model with the structure of inflation switching according to a Markov switching process are "generally quite different from the ones we observe in the [survey forecast] data." As we show in the next section, nonrepetitive structural shifts in the inflation process render FIRE inapplicable as a representation of participants' expectations, let alone rational expectations.

4 A Theoretical Account of Unforeseeable Structural Shifts in the Forecast-Error Regressions

The Knight-Muth hypothesis (KMH), introduced in Frydman and Tabor (2024), provides a new approach to representing rational market participants' expectations of inflation, and other outcomes, that undergo unforeseeable structural change.

Here, we illustrate how formalizing unforeseeable structural change in the inflation process enables KMH to account for our empirical findings and reconcile them with the rationality of market participants' inflation expectations. We do so by implementing KMH in the simple reduced-form model for inflation Hajdini and Kurmann (2024) use to derive FIRE's implications for the forecast-error regressions when the inflation process undergoes repetitive regime switches. We trace KMH's and FIRE's different predictions for structural change in the forecast-error regressions to their respective assumptions about structural change in the inflation process.

4.1 The Reduced-Form Model for Inflation

We consider the reduced-form model for inflation given by:

$$\pi_t = \delta_t + \gamma_t x_t, \tag{3}$$

for $t = 1, 2, \ldots, T$, where $\delta_t > 0$ and $\gamma_t > 0$ are parameters at time t, which, using Muth's (1961) nomenclature, we refer to as the model's "objective" parameters, and x_t is an exogenous process that follows the mean-zero autoregressive process with constant parameters given by:

$$x_t = \rho x_{t-1} + \epsilon_t, \tag{4}$$

for $t = 1, 2, \ldots, T$, where $0 < \rho < 1$, the initial value x_0 is given, and $\epsilon_t \sim iidN(0, \sigma^2)$.

Most REH models abstract from structural change in the inflation process, which they formalize by constraining the objective parameters in (3) to remain constant over time, $\delta_t = \delta$ and $\gamma_t = \gamma$ for all t. These constant parameters can be estimated based on historical inflation data. Thus, how inflation will unfold can be foreseen in probabilistic terms, and the only uncertainty about future inflation is probabilistic risk about the future random shocks to the output gap.

Although some REH models have recognized that the inflation process undergoes structural change, they formalize this change as repetitive regime switches. Like in the constant-parameter models, formalizing structural change as repetitive rests on the premise that how future inflation will unfold can be foreseen in probabilistic terms. REH models formalize repetitive regime switches by assuming that the objective parameters in (3) depend on a latent state variable, s_t, following a Markov chain with two or more states and switching probabilities $P(s_t = j | s_{t-1} = i) = p_{ij}$ for $i, j \in \{1, 2\}$, where $p_{i1} + p_{i2} = 1$ for $i \in \{1, 2\}$, such that the objective parameters take the values $(\delta(s_t), \gamma(s_t)) = (\delta^{(i)}, \gamma^{(i)})$ when the state is $s_t = i$ for $i \in \{1, 2\}$. Because the regime switches depend only on constant parameters and switching probabilities, $(\delta^{(1)}, \delta^{(2)}, \gamma^{(1)}, \gamma^{(2)}, p_{11}, p_{22})$, that can be estimated based on historical inflation data, how inflation will unfold can be foreseen in probabilistic terms. For example, if the current state is $s_t = 1$, it can be foreseen that, with probability p_{11}, the next period's inflation will be in regime 1 with parameters $(\delta^{(1)}, \gamma^{(1)})$ and, with probability $1 - p_{11}$, in regime 2 with parameters $(\delta^{(2)}, \gamma^{(2)})$. Thus, the only uncertainty about future inflation is probabilistic risk about the future random shocks to the output gap and the future latent states, which determine which of the regimes future inflation will be in.

In contrast, building on Knight's (1921) insights, KMH formalizes structural change in the inflation process as nonrepetitive, and thus unforeseeable. Because future changes differ from past changes, neither an economist nor a rational market participant can foresee, even in probabilistic terms, how the inflation process will change in the future. Such unforeseeable change gives rise to Knightian uncertainty: the uncertainty about future inflation cannot be reduced ex ante to a single probability distribution.

4.2 Formalizing Structural Change as Unforeseeable

KMH provides a tractable and empirically testable formalization of unforeseeable change in the inflation process by assuming that the sequence of objective parameters, $\{\delta_t, \gamma_t\}_{t=1}^{T}$, shifts intermittently between nonrepetitive values restricted to lie within an interval:

$$(\delta_t, \gamma_t) = \left(\delta^{(m)}, \gamma^{(m)}\right) \in I^\delta \times I^\gamma = [\delta_L, \delta_U] \times [\gamma_L, \gamma_U]\,, \tag{5}$$

for $t = T_{m-1}, T_{m-1}+1, \ldots, T_m - 1$ and $m = 1, 2, \ldots, M$, where $T_m < T_{m+1}$ for all m, $T_0 = 1$, $T_M = T + 1$, and the bounds of the intervals I^δ and I^γ satisfy $0 < \delta_L < \delta_U$ and $0 < \gamma_U < \gamma_U$.

The first part of (5) formalizes the objective parameters to shift only intermittently at times $\{T_m\}_{m=1}^{M-1}$, which are not determined by a probabilistic rule, between nonrepetitive values $\{\delta^{(m)}, \gamma^{(m)}\}_{m=1}^{M}$. For example, the parameters are given by $(\delta_t, \gamma_t) = (\delta^{(1)}, \gamma^{(1)})$ during the first subperiod, from $t = T_0 = 1$ to $t = T_1 - 1$, but they shift to $(\delta^{(2)}, \gamma^{(2)})$ at time $t = T_1$, $(\delta^{(3)}, \gamma^{(3)})$ at time $t = T_2$, etc. The second part of (5) constrains the parameters in each of these periods to lie within the positive intervals I^δ and I^γ, such that inflation is always positively related to the output gap.

This formalization of structural change in the inflation process can be tested empirically by estimating the sequence of objective parameters $\{\delta_t, \gamma_t\}_{t=1}^{T}$ based on historical time-series data using an empirical approach that identifies the timing of these shifts from the data, such as the one we used in Section 3. This enables us to test the adequacy of a KMH model's specification.

Importantly, the specification in (5) formalizes the structural change as unforeseeable: viewed from any time t, the values of the future objective parameters and the timings of their shifts are inherently unknowable, even in probabilistic terms. As future objective parameters differ from past ones, their values cannot be estimated by any method based on historical inflation data.

4.3 Knightian Uncertainty About Future Inflation

Because the KMH model formalizes the unforeseeable change in the inflation process as unforeseeable, it provides a novel formalization of Knightian uncertainty. Whereas a single conditional distribution constitutes a REH model's prediction of future inflation, a *set* of conditional distributions constitutes the KMH model's prediction.

Formally, although the future objective parameters are unknown ex ante, they are restricted to lie within the intervals I^δ and I^γ. For each possible value of the parameters within these intervals, the model in (3) and (4) defines a unique conditional expectation of future inflation. As a result, the set of conditional distributions corresponding to all possible values of the future objective parameters within the interval constitutes the model's prediction of future inflation undergoing unforeseeable change. This set formalizes Knightian uncertainty about future inflation.

Let $\pi_{t+h}(\delta_{t+h}, \gamma_{t+h})$ denote inflation at time $t+h$, $h > 0$, as a function of the objective parameters $(\delta_{t+h}, \gamma_{t+h})$ and $\mathrm{E}(\pi_{t+h}(\delta_{t+h}, \gamma_{t+h})|x_t)$ its expectation conditional on x_t, as defined by (3) and (4). The set of conditional expectations that constitutes the model's prediction of future inflation is given by:

$$\begin{aligned}
&\left\{\mathrm{E}\left(\pi_{t+h}\left(\delta_{t+h},\gamma_{t+h}\right)|x_t\right) \mid \left(\delta_{t+h},\gamma_{t+h}\right)\in I^{\delta}\times I^{\gamma}\right\}\\
&=\left\{\delta_{t+h}+\gamma_{t+h}\rho^{h}x_t \mid \left(\delta_{t+h},\gamma_{t+h}\right)\in\left[\delta_L,\delta_U\right]\times\left[\gamma_L,\gamma_U\right]\right\}\\
&=\left[\delta_L+\mathbf{1}_{(x_t\geq 0)}\left(\gamma_L\rho^{h}x_t\right)+\mathbf{1}_{(x_t<0)}\left(\gamma_U\rho^{h}x_t\right)\right. ,\\
&\qquad\left.\delta_U+\mathbf{1}_{(x_t<0)}\left(\gamma_L\rho^{h}x_t\right)+\mathbf{1}_{(x_t\geq 0)}\left(\gamma_U\rho^{h}x_t\right)\right],
\end{aligned} \tag{6}$$

where $\mathbf{1}_{(x_t>0)}$ is an indicator variable that takes the value 1 if $x_t>0$, and 0 otherwise.

4.4 Rational Expectations of Inflation Undergoing Unforeseeable Change

Muth (1961) advanced the pathbreaking hypothesis that an economist can represent participants' expectations as consistent with his model's predictions of those outcomes. His hypothesis assumed that market participants are rational: they base their expectations of aggregate outcomes on their understanding of "the way the economy works" (p. 315). His pioneering idea was that an economist can acknowledge that participants are rational by representing their understanding of the way the economy works with his own understanding, as formalized by his model.

An REH model implements Muth's hypothesis by representing participants' inflation expectation with the conditional expectation that constitutes the model's *only* prediction of future inflation.

Although the KMH model implies that the conditional expectation of future inflation is given by $\mathrm{E}\left(\pi_{t+h}|x_t\right)=\delta_{t+h}+\gamma_{t+h}\rho^{h}x_t$ for $h>0$, there is no objective way to assess the values of the future objective parameters, $\left(\delta_{t+h},\gamma_{t+h}\right)$, ex ante. Indeed, if there was such an objective way, the structural change would be foreseeable.

Consequently, KMH represents rational participants' inflation expectations with a conditional expectation within the set in (6) based on a subjective scenario—which represents their ex ante assessment of the unknown future objective parameters—that is only restricted to lie within the interval $I^{\delta}\times I^{\gamma}$, but whose value is determined by neither the objective parameters nor the available information. This represents participants' expectations as consistent with the model's prediction of future inflation. It recognizes that although market participants are rational, there is no objective way for them to assess ex ante how the inflation process will change in the future, even if they base their expectations on full information. Moreover, it formalizes Keynes' (1936, p. 162-163) idea that *rational* participants rely on various factors, including psychological and other factors outside the model, in forming expectations. This implies that they revise their subjective scenario at times and in ways that are, at least partly, independent of the shifts in the objective parameters.

Formally, let $(\bar{\delta}_t, \bar{\gamma}_t)$ denote participants' subjective scenario at time t.[4] We assume that the sequence of subjective parameters, $\{\bar{\delta}_t, \bar{\gamma}_t\}_{t=1}^T$, shifts intermittently between nonrepetitive values restricted to lie within the objective parameters' interval:

$$\left(\bar{\delta}_t, \bar{\gamma}_t\right) = \left(\bar{\delta}^{(n)}, \bar{\gamma}^{(n)}\right) \in I^\delta \times I^\gamma, \tag{7}$$

for $t = \bar{T}_{n-1}, \bar{T}_{n-1} + 1, \ldots, \bar{T}_n - 1$ and $n = 1, 2, \ldots, N$, where $\bar{T}_n < \bar{T}_{n+1}$ for all n, $\bar{T}_0 = 1$, $\bar{T}_N = T + 1$, and the intervals I^δ and I^γ are defined in (5). Similar to the specification of the objective parameters in (5), the first part of (7) constrains participants' subjective scenario to shift intermittently between nonrepetitive values at times that are not characterized by a probabilistic rule. The second part restricts the subjective scenario to lie within the interval for the objective parameters at all times. This renders participants' subjective scenario consistent with the model's prediction that the unknown future objective parameters lie within this interval.

Although the magnitude and timing of shifts in participants' subjective scenario are not determined by the objective parameters, we assume that the subjective scenario differs most from the objective parameters around the times when the objective parameters shift, T_m for $m = 1, 2, \ldots, M - 1$. Moreover, we expect most shifts in the subjective scenario to occur proximate to those times.

Similar to FIRE, we assume that rational participants base their inflation expectations on full information. Given the subjective scenario, we represent their inflation expectation, at time t and for forecast horizon $h > 0$, as:

$$\mathrm{E}\left(\pi_{t+h}\left(\bar{\delta}_t, \bar{\gamma}_t\right) | x_t\right) = \bar{\delta}_t + \bar{\gamma}_t \rho^h x_t, \tag{8}$$

which lies within the set of conditional expectations in (6) that constitutes the model's predictions of future inflation as $(\bar{\delta}_t, \bar{\gamma}_t) \in I^\delta \times I^\gamma$.

Like FIRE, the functional form and relevant information—here x_t—are determined by the inflation process in (3) and (4). Unlike FIRE, however, KMH represents rational participants' inflation expectations in terms of subjective parameters that undergo shifts whose magnitude and timing are not determined by those in the objective parameters. As the magnitude and timing of shifts in the subjective and objective parameters differ, rational participants' inflation expectations deviate from ex post inflation in nonrepetitive, and thus unforeseeable, ways, with the largest deviations occuring during times when the inflation process is undergoing major structural shifts.

In contrast, a FIRE model with regime-switching in the inflation process implies that participants' inflation expectations deviate from ex post inflation in repetitive ways that

[4]We assume that participants' subjective scenario at time t is the same for all forecast horizons, $h > 0$. Our results do not depend on this simplifying assumption.

are *completely* determined by the regime switches in the inflation process. In such a model, given the current state, $s_t = i$ for $i \in \{1,2\}$, FIRE represents participants' inflation expectation with the state-dependent weighted average of the two regimes' conditional expectations of future inflation: $\mathrm{E}(\pi_{t+h}(s_{t+h})|x_t, s_t = i) = p_{i1}^{(h)}\delta^{(1)} + p_{i2}^{(h)}\delta^{(2)} + (p_{i1}^{(h)}\gamma^{(1)} + p_{i2}^{(h)}\gamma^{(2)})\rho^h x_t$, where $p_{ij}^{(h)} = P(s_{t+h} = j|s_t = i)$ for $i, j \in \{1,2\}$ and $h > 0$. As the inflation process switches regime, participants switch the weights attached to the two regimes, $p_{ij}^{(h)}$, in repetitive, and thus foreseeable, ways.

4.5 Implications for the Forecast-Error Regressions

Given the inflation process in (3) and the representation of participants' inflation expectations in (8), KMH implies that rational participants' ex post inflation forecast errors are given by:

$$\pi_{t+h} - \mathrm{E}\left(\pi_{t+h}\left(\bar{\delta}_t, \bar{\gamma}_t\right)|x_t\right) = \left(\delta_{t+h} - \bar{\delta}_t\right) + \left(\gamma_{t+h} - \bar{\gamma}_t\right)\rho^h x_t + \gamma_{t+h}\sum_{i=0}^{h-1}\rho^i \varepsilon_{t+h-i}. \quad (9)$$

We can directly link these ex post forecast errors to ex ante inflation, π_t, by using (3), which implies that $x_t = (\pi_t - \delta_t)/\gamma_t$:

$$\begin{aligned}\pi_{t+h} - \mathrm{E}\left(\pi_{t+h}\left(\bar{\delta}_t, \bar{\gamma}_t\right)|x_t\right) = &\left(\left(\delta_{t+h} - \bar{\delta}_t\right) - \frac{\left(\gamma_{t+h} - \bar{\gamma}_t\right)\rho^h \delta_t}{\gamma_t}\right) \\ &+ \left(\frac{\left(\gamma_{t+h} - \bar{\gamma}_t\right)\rho^h}{\gamma_t}\right)\pi_t + \gamma_{t+h}\sum_{i=0}^{h-1}\rho^i \varepsilon_{t+h-i}. \quad (10)\end{aligned}$$

This shows directly that the nonrepetitive shifts in both the objective and subjective parameters imply that the bias of participants' ex post forecast errors and their correlation with ex ante inflation, π_t, undergoes nonrepetitive, and thus unforeseeable, shifts over time. During times characterized by major structural shifts in the inflation process, we expect a large deviation between the ex ante subjective parameters and the ex post objective parameters, so that participants' ex post forecast errors have a large—either positive or negative—bias and correlation with the ex ante inflation. In contrast, we expect this deviation to be small—and thus the bias and correlation close to, or even equal to, zero—during times when the inflation process remains stable.

The following proposition states the KMH model's predictions regarding structural change in the forecast-error regressions in (2). The proof of the proposition is presented in Appendix A.[5]

Proposition 1 *Assume that inflation and the output gap, π_t and x_t, are characterized by (3)-(5), and market participants' inflation expectation is given by $F_t(\pi_{t+h}) =$*

[5] In Supplemental Appendix D, we present a similar proposition for the KMH model's predictions for the forecast-error regression in (1).

$\mathrm{E}(\pi_{t+h}(\bar{\delta}_t, \bar{\gamma}_t)|x_t)$ *for* $h > 0$*, as defined by (7)-(8). Moreover, assume that we estimate the forecast-error regression in (2) with structural shifts at known times* $\{\tau_k\}_{k=1}^{K-1}$*, as given by*

$$\pi_{t+h} - F_t(\pi_{t+h}) = \alpha^{(2,k)} + \beta^{(2,k)}\pi_t + u_t^{(2,k)}, \tag{11}$$

for $t \in [\tau_{k-1}, \tau_k - 1]$ *and* $k = 1, 2, \ldots, K$*, where*

$$\{\tau_k\}_{k=1}^{K-1} = \{T_m, T_m - h\}_{m=1}^{M-1} \cup \{\bar{T}_n\}_{n=1}^{N-1} \tag{12}$$

with $K \leq 2M + N$*,* $\tau_0 = 1$ *and* $\tau_K = T + 1$*.*

For a subperiod k*, for* $t \in [\tau_{k-1}, \tau_k - 1]$ *with* $\tau_{k-1} = \max(T_{m-1}, \bar{T}_{n-1})$ *and* $\tau_k = \min(T_m - h, \bar{T}_n)$*, where* $(\delta_t, \gamma_t) = (\delta_{t+h}, \gamma_{t+h}) = (\delta^{(m)}, \gamma^{(m)})$ *and* $(\bar{\delta}_t, \bar{\gamma}_t) = (\bar{\delta}^{(n)}, \bar{\gamma}^{(n)})$ *for some* $m, n > 0$*, the OLS estimates of (11),* $\widehat{\alpha}^{(2,k)}$ *and* $\widehat{\beta}^{(2,k)}$*, converge asymptotically towards:*

$$\widehat{\alpha}^{(2,k)} \xrightarrow{p} \alpha^{(2,k)} = \left(\delta^{(m)} - \bar{\delta}^{(n)}\right) - \frac{\left(\gamma^{(m)} - \bar{\gamma}^{(n)}\right)\rho^h \delta^{(m)}}{\gamma^{(m)}} \tag{13}$$

and

$$\widehat{\beta}^{(2,k)} \xrightarrow{p} \beta^{(2,k)} = \frac{\left(\gamma^{(m)} - \bar{\gamma}^{(n)}\right)\rho^h}{\gamma^{(m)}} \tag{14}$$

for $\tau_k \longrightarrow \infty$*. Thus, if the subperiod is sufficiently long, the estimators are approximately given by* $\widehat{\alpha}^{(2,k)} \approx \alpha^{(2,k)}$ *and* $\widehat{\beta}^{(2,k)} \approx \beta^{(2,k)}$*. If* $\gamma^{(m)} > \bar{\gamma}^{(n)}$ *then* $\beta^{(2,k)} > 0$*; if* $\gamma^{(m)} < \bar{\gamma}^{(n)}$ *then* $\beta^{(2,k)} < 0$*; and if* $\gamma^{(m)} = \bar{\gamma}^{(n)}$ *then* $\beta^{(2,k)} = 0$*. If* $\delta^{(m)} > \bar{\delta}^{(n)}$ *and* $\gamma^{(m)} \leq \bar{\gamma}^{(n)}$ *then* $\alpha^{(2,k)} > 0$*; if* $\delta^{(m)} < \bar{\delta}^{(n)}$ *and* $\gamma^{(m)} \geq \bar{\gamma}^{(n)}$ *then* $\alpha^{(2,k)} < 0$*; and if* $\delta^{(m)} = \bar{\delta}^{(n)}$ *and* $\gamma^{(m)} = \bar{\gamma}^{(n)}$ *then* $\alpha^{(2,k)} = 0$*.*

The proposition shows the KMH model predicts that the estimates of the intercept and the slope of the forecast-error regression in (2) should shift between nonrepetitive values over time. Because these shifts arise from unforeseeable shifts in the model's objective and subjective parameters, the parameters of the forecast-error regression shift in unforeseeable ways over time between nonrepetitive values that can be either positive, negative, or zero during different periods.

The KMH model's predictions differ from those of a FIRE model with repetitive regime switches. As Hajdini and Kurmann (2024, Proposition 1) show, a FIRE model implies that the correlation of participants' ex post forecast errors with ex ante inflation should switch between repetitive positive and negative values.

In contrast, the KMH model predicts that participants' ex post forecast errors are biased and/or correlated with ex ante inflation *only* during periods when their expectations are based on an ex ante subjective scenario that differs from the ex post objective parameters. For example, the estimated correlation during subperiod k should be positive if $\gamma^{(m)} > \bar{\gamma}^{(n)}$ and negative if $\gamma^{(m)} < \bar{\gamma}^{(n)}$. During periods when their ex ante subjective scenario is identical to the ex post objective parameters, the bias and/or correlation should be zero. For example, if $\gamma^{(m)} = \bar{\gamma}^{(n)}$, then the estimated correlation should be zero.

The KMH model recognizes that rational participants' ex ante inflation expectations may differ from ex post inflation undergoing unforeseeable change. Thus, it explains why most structural shifts and the largest bias and correlation of ex post forecast errors with ex ante information occurred when they did: during the 1970s, the early 1980s, and post-2020. These periods were characterized by unprecedented—and thus unforeseeable—structural changes in the economy, monetary policy, and inflation dynamics. The KMH model's account of their ex post forecast errors acknowledges that it was particularly difficult for rational market participants to forecast inflation during those periods. Likewise, it explains our empirical finding that when the inflation process has been relatively stable, as during the Great Moderation, rational participants' forecast errors have been nearly unbiased and uncorrelated with ex ante information.

5 Concluding Remarks

Relying on forecast-error regressions with parameters assumed to be constant, a number of studies have found participants' ex post inflation forecast errors to be unbiased and positively correlated with ex ante forecast revisions. This positive correlation has come to be viewed in the literature as a "stylized fact" that representations of participants' expectations of inflation, and other aggregate outcomes, should be consistent with (Coibion and Gorodnichenko, 2015, abstract, p. 2644).

However, our empirical findings reveal that the estimates of an insignificant bias and a significantly positive correlation between ex post forecast errors and ex ante forecast revisions are an artifact of the assumption that the parameters of the forecast-error regression have remained constant throughout the sample period. Once structural shifts in these parameters are identified from the time-series data, the estimate of the correlation between ex post inflation-forecast errors and ex ante forecast revisions is significantly positive only during the period from 1972 to 1974 and has remained insignificant until the end of the sample in 2022. Moreover, the estimated bias undergoes nonrepetitive shifts between significantly positive and significantly negative values during different subperiods of the sample.

Our findings suggest a different "stylized fact" that macroeconomic models of inflation should be consistent with: the inflation process and rational participants' inflation expectations undergo nonrepetitive change at times that cannot be foreseen, even in probabilistic terms. However, to establish that recognizing unforeseeable change is necessary for building empirically relevant macroeconomic and finance models requires future research to examine whether forecast-error regressions for other macroeconomic and financial variables also undergo nonrepetitive shifts.

Our findings provide econometric evidence for what seems apparent: the inflation process in real-world markets undergoes change that cannot be foreseen, even in probabilistic terms. Moreover, market participants understand that they face such change, which we interpret as a hallmark of their rationality. Building on Frank Knight's (1921) and John Muth's (1961) pathbreaking insights, Frydman and Tabor (2024) propose the Knight-Muth hypothesis (KMH) to enable building tractable and empirically testable intertemporal models that formalize unforeseeable structural change in the inflation process and in rational participants' inflation expectations. Section 4 illustrates how implementing KMH in a reduced-form model of inflation undergoing unforeseeable change can account for our empirical findings.

REH cannot represent participants' expectations of outcomes undergoing unforeseeable change because it requires that future structural change, and thus future outcomes, can be foreseen in probabilistic terms. We suggest that macroeconomists should reconsider relying on models that assume that structural change in the process driving outcomes is foreseeable, and that representing participants' expectations consistent with such models' predictions, as REH does, represents rational participants' expectations of outcomes in real-world markets.

Like REH, KMH can be applied in any macroeconomic and finance model. Doing so in models other than the simple model for inflation used here, and examining how KMH alters REH's implications for market outcomes and policy analysis, are among the main topics for future research.

References

Angeletos, G.-M., Huo, Z., & Sastry, K. A. (2021). Imperfect Macroeconomic Expectations: Evidence and Theory. *NBER Macroeconomics Annual*, *35*, 1–86. https://doi.org/10.1086/712313

Bordalo, P., Gennaioli, N., Ma, Y., & Shleifer, A. (2020). Overreaction in Macroeconomic Expectations. *American Economic Review*, *110*(9), 2748–2782. https://doi.org/10.1257/aer.20181219

Castle, J. L., Doornik, J. A., & Hendry, D. F. (2012). Model selection when there are multiple breaks. *Journal of Econometrics*, *169*(2), 239–246. https://doi.org/10.1016/j.jeconom.2012.01.026

Castle, J. L., Doornik, J. A., Hendry, D. F., & Pretis, F. (2015). Detecting Location Shifts during Model Selection by Step-Indicator Saturation. *Econometrics*, *3*(2), 240–264. https://doi.org/10.3390/econometrics3020240

Coibion, O., & Gorodnichenko, Y. (2015). Information Rigidity and the Expectations Formation Process: A Simple Framework and New Facts. *American Economic Review*, *105*(8), 2644–2678. https://doi.org/10.1257/aer.20110306

Doornik, J. A. (2009). Autometrics. In J. Castle & N. Shephard (Eds.), *The Methodology and Practice of Econometrics: A Festschrift in Honour of David F. Hendry*. Oxford University Press. https://doi.org/10.1093/acprof:oso/9780199237197.003.0004

Ericsson, N. R. (2012). Detecting crises, jumps, and changes in regime. *Draft, Board of Governors of the Federal Reserve System, Washington, DC, November.*

Farmer, L. E., Nakamura, E., & Steinsson, J. (2024). Learning about the Long Run. *Journal of Political Economy*, *132*(10), 3334–3377. https://doi.org/10.1086/730207

Frydman, R., & Tabor, M. N. (2024). *Rational Expectations Under Knightian Uncertainty Arising From Unforeseeable Change*. INET Center on Knightian Uncertainty Working Paper.

Gennaioli, N., & Shleifer, A. (2018). *A crisis of beliefs: Investor psychology and financial fragility*. Princeton University Press.

Hajdini, I., & Kurmann, A. (2024). *Predictable Forecast Errors in Full-Information Rational Expectations Models with Regime Shifts*. Federal Reserve Bank of Cleveland, Working Paper 24-08. https://doi.org/10.26509/frbc-wp-202408

Kahneman, D., & Tversky, A. (1972). Subjective probability: A judgment of representativeness. *Cognitive Psychology*, *3*(3), 430–454. https://doi.org/10.1016/0010-0285(72)90016-3

Keynes, J. M. (1936). *The General Theory of Employment, Interest and Money*. Macmillan.

Knight, F. H. (1921). *Risk, uncertainty and profit*. Houghton Mifflin.

Muth, J. F. (1961). Rational Expectations and the Theory of Price Movements. *Econometrica*, *29*(3), 315–335. https://doi.org/10.2307/1909635

A Proof of Proposition 1

Proof of Proposition 1. We assume that inflation and the output gap, π_t and x_t, are characterized by (3)-(5), and market participants' inflation expectation is given by $F_t(\pi_{t+h}) = \mathrm{E}(\pi_{t+h}(\bar{\delta}_t, \bar{\gamma}_t)|x_t)$ for $h > 0$, as defined by (7)-(8).

We consider the forecast-error regression in (11) with the timings of structural shifts, $\{\tau_k\}_{k=1}^{K-1}$, as specified in (12), assumed known. We consider a subperiod k, for $t \in [\tau_{k-1}, \tau_k - 1]$ with $\tau_{k-1} = \max(T_{m-1}, \bar{T}_{n-1})$ and $\tau_k = \min(T_m - h, \bar{T}_n)$, during which the objective and subjective parameters are given by $(\delta_t, \gamma_t) = (\delta_{t+h}, \gamma_{t+h}) = (\delta^{(m)}, \gamma^{(m)})$ and $(\bar{\delta}_t, \bar{\gamma}_t) = (\bar{\delta}^{(n)}, \bar{\gamma}^{(n)})$ for some $m \in [1, M-1]$ and $n \in [1, N-1]$.

Here, we simplify the notation by dropping the superscripts and denote these parameters by $(\delta^{(m)}, \gamma^{(m)}) = (\delta, \gamma)$ and $(\bar{\delta}^{(n)}, \bar{\gamma}^{(n)}) = (\bar{\delta}, \bar{\gamma})$. Thus, during the subperiod, inflation and the ex post forecast errors are given by:

$$\pi_t = \delta + \gamma x_t, \tag{15}$$

and

$$e_t = \pi_{t+h} - \mathrm{E}\left(\pi_{t+h}\left(\bar{\delta}, \bar{\gamma}\right)|x_t\right) = \left(\delta - \bar{\delta}\right) + \left(\gamma - \bar{\gamma}\right)\rho^h x_t + \gamma \sum_{i=0}^{h-1} \rho^i \varepsilon_{t+h-i}. \tag{16}$$

The OLS estimator of the slope in (11) for subperiod k, which we denote here by $\widehat{\beta} = \widehat{\beta}^{(2,k)}$, is given by:

$$\widehat{\beta} = \frac{\frac{1}{\tau_k - \tau_{k-1}} \sum_{t=\tau_{k-1}}^{\tau_k - 1} \left(\pi_t - \bar{\pi}\right)\left(e_t - \bar{e}\right)}{\frac{1}{\tau_k - \tau_{k-1}} \sum_{t=\tau_{k-1}}^{\tau_k - 1} \left(\pi_t - \bar{\pi}\right)^2}, \tag{17}$$

where $\bar{\pi} = (\tau_k - \tau_{k-1})^{-1} \sum_{t=\tau_{k-1}}^{\tau_k - 1} \pi_t$ and $\bar{e} = (\tau_k - \tau_{k-1})^{-1} \sum_{t=\tau_{k-1}}^{\tau_k - 1} e_t$.

Using (15) and defining $\bar{x} = (\tau_k - \tau_{k-1})^{-1} \sum_{t=\tau_{k-1}}^{\tau_k - 1} x_t$, the term $\pi_t - \bar{\pi}$ in the nominator and denominator of (17) is given by:

$$\pi_t - \bar{\pi} = \delta + \gamma x_t - \frac{1}{\tau_{k-1} - \tau_k} \sum_{t=\tau_{k-1}}^{\tau_k - 1} \left(\delta + \gamma x_t\right) = \gamma \left(x_t - \bar{x}\right). \tag{18}$$

Using (16) and defining $\bar{\varepsilon}^{(h-i)} = (\tau_k - \tau_{k-1})^{-1} \sum_{t=\tau_{k-1}}^{\tau_k - 1} \varepsilon_{t+h-i}$ for $h > 0$ and $i = 0, 1, \ldots, h-1$, the term $e_t - \bar{e}$ in the nominator of (17) is given by:

$$\begin{aligned} e_t - \bar{e} &= \left(\delta - \bar{\delta}\right) + \left(\gamma - \bar{\gamma}\right)\rho^h x_t + \gamma \sum_{i=0}^{h-1} \rho^i \varepsilon_{t+h-i} \\ &\quad - \frac{1}{\tau_{k-1} - \tau_k} \sum_{t=\tau_{k-1}}^{\tau_k - 1} \left(\left(\delta - \bar{\delta}\right) + \left(\gamma - \bar{\gamma}\right)\rho^h x_t + \gamma \sum_{i=0}^{h-1} \rho^i \varepsilon_{t+h-i} \right) \\ &= \left(\gamma - \bar{\gamma}\right)\rho^h \left(x_t - \bar{x}\right) + \gamma \sum_{i=0}^{h-1} \rho^i \left(\varepsilon_{t+h-i} - \bar{\varepsilon}^{(h-i)}\right). \end{aligned} \tag{19}$$

Inserting these expressions, the denominator of (17) converges asymptotically towards:

$$
\begin{aligned}
&\frac{1}{\tau_k-\tau_{k-1}}\sum_{t=\tau_{k-1}}^{\tau_k-1}\left(\pi_t-\bar{\pi}\right)\left(e_t-\bar{e}\right)\\
&\quad=\frac{1}{\tau_k-\tau_{k-1}}\sum_{t=\tau_{k-1}}^{\tau_k-1}\gamma\left(x_t-\bar{x}\right)\left(\left(\gamma-\bar{\gamma}\right)\rho^h\left(x_t-\bar{x}\right)+\gamma\sum_{i=0}^{h-1}\rho^i\left(\varepsilon_{t+h-i}-\bar{\varepsilon}^{(h-i)}\right)\right)\\
&\quad=\gamma\left(\gamma-\bar{\gamma}\right)\rho^h\left(\frac{1}{\tau_k-\tau_{k-1}}\sum_{t=\tau_{k-1}}^{\tau_k-1}\left(x_t-\bar{x}\right)^2\right)\\
&\qquad+\gamma^2\sum_{i=0}^{h-1}\rho^i\left(\frac{1}{\tau_k-\tau_{k-1}}\sum_{t=\tau_{k-1}}^{\tau_k-1}\left(x_t-\bar{x}\right)\left(\varepsilon_{t+h-i}-\bar{\varepsilon}^{(h-i)}\right)\right)\\
&\quad\xrightarrow{p}\gamma\left(\gamma-\bar{\gamma}\right)\rho^h\operatorname{Var}(x_t)+\gamma^2\sum_{i=0}^{h-1}\rho^i\operatorname{Cov}(x_t,\varepsilon_{t+h-i})=\gamma\left(\gamma-\bar{\gamma}\right)\rho^h\frac{\sigma^2}{1-\rho^2},
\end{aligned}
\tag{20}
$$

for $\tau_k \longrightarrow \infty$, as $0<\rho<1$ implies that $\operatorname{Var}(x_t)=\sigma^2/(1-\rho^2)$, and $\operatorname{Cov}(x_t,\varepsilon_{t+h-i})=0$ for all $h>0$ and $i=0,1,\ldots,h-1$. The denominator in (17) converges asymptotically towards:

$$
\begin{aligned}
\frac{1}{\tau_k-\tau_{k-1}}\sum_{t=\tau_{k-1}}^{\tau_k-1}\left(\pi_t-\bar{\pi}\right)^2&=\frac{1}{\tau_k-\tau_{k-1}}\sum_{t=\tau_{k-1}}^{\tau_k-1}\left(\gamma\left(x_t-\bar{x}\right)\right)^2\\
&=\gamma^2\left(\frac{1}{\tau_k-\tau_{k-1}}\sum_{t=\tau_{k-1}}^{\tau_k-1}\left(x_t-\bar{x}\right)^2\right)\\
&\xrightarrow{p}\gamma^2\operatorname{Var}(x_t)=\gamma^2\frac{\sigma^2}{1-\rho^2},
\end{aligned}
\tag{21}
$$

for $\tau_k \longrightarrow \infty$. It follows that the slope estimator in (17) converges asymptotically towards the following:

$$
\widehat{\beta}=\frac{\frac{1}{\tau_k-\tau_{k-1}}\sum_{t=\tau_{k-1}}^{\tau_k-1}\left(\pi_t-\bar{\pi}\right)\left(e_t-\bar{e}\right)}{\frac{1}{\tau_k-\tau_{k-1}}\sum_{t=\tau_{k-1}}^{\tau_k-1}\left(\pi_t-\bar{\pi}\right)^2}\xrightarrow{p}\frac{\gamma\left(\gamma-\bar{\gamma}\right)\rho^h\frac{\sigma^2}{1-\rho^2}}{\gamma^2\frac{\sigma^2}{1-\rho^2}}=\frac{\left(\gamma-\bar{\gamma}\right)\rho^h}{\gamma}=\beta \tag{22}
$$

for $\tau_k \longrightarrow \infty$. Thus, if the subperiod is sufficiently long, the slope estimator is approximately given by $\widehat{\beta}\approx\beta$. As $\gamma>0$ and $0<\rho<1$, it follows that if $\gamma>\bar{\gamma}$ then $\beta>0$, if $\gamma<\bar{\gamma}$ then $\beta<0$, and if $\gamma=\bar{\gamma}$ then $\beta=0$.

The OLS estimator of the intercept in (11) for subperiod k, which we denote here by $\widehat{\alpha}=\widehat{\alpha}^{(2,k)}$, is given by:

$$
\widehat{\alpha}=\frac{1}{\tau_k-\tau_{k-1}}\sum_{t=\tau_{k-1}}^{\tau_k-1}e_t-\widehat{\beta}\frac{1}{\tau_k-\tau_{k-1}}\sum_{t=\tau_{k-1}}^{\tau_k-1}\pi_t. \tag{23}
$$

Using the expressions for π_t and e_t in (15) and (16), the intercept estimator converges asymptotically towards the following:

$$
\begin{aligned}
\widehat{\alpha} &= \frac{1}{\tau_k - \tau_{k-1}} \sum_{t=\tau_{k-1}}^{\tau_k - 1} e_t - \widehat{\beta} \left(\frac{1}{\tau_k - \tau_{k-1}} \sum_{t=\tau_{k-1}}^{\tau_k - 1} \pi_t \right) \\
&= \frac{1}{\tau_k - \tau_{k-1}} \sum_{t=\tau_{k-1}}^{\tau_k - 1} \left(\left(\delta - \bar{\delta}\right) + \left(\gamma - \bar{\gamma}\right) \rho^h x_t + \gamma \sum_{i=0}^{h-1} \rho^i \varepsilon_{t+h-i} \right) \\
&\quad - \widehat{\beta} \left(\frac{1}{\tau_k - \tau_{k-1}} \sum_{t=\tau_{k-1}}^{\tau_k - 1} (\delta + \gamma x_t) \right) \\
&= \left(\delta - \bar{\delta}\right) + \left(\gamma - \bar{\gamma}\right) \rho^h \left(\frac{1}{\tau_k - \tau_{k-1}} \sum_{t=\tau_{k-1}}^{\tau_k - 1} x_t \right) \\
&\quad + \gamma \sum_{i=0}^{h-1} \rho^i \left(\frac{1}{\tau_k - \tau_{k-1}} \sum_{t=\tau_{k-1}}^{\tau_k - 1} \varepsilon_{t+h-i} \right) - \widehat{\beta}\delta - \widehat{\beta}\gamma \left(\frac{1}{\tau_k - \tau_{k-1}} \sum_{t=\tau_{k-1}}^{\tau_k - 1} x_t \right) \\
&\xrightarrow{p} \left(\delta - \bar{\delta}\right) + \left(\gamma - \bar{\gamma}\right) \rho^h \,\mathrm{E}(x_t) + \gamma \sum_{i=0}^{h-1} \rho^i \,\mathrm{E}(\varepsilon_{t+h-i}) - \beta\delta - \beta \,\mathrm{E}(x_t) \\
&= \left(\delta - \bar{\delta}\right) - \beta\delta = \left(\delta - \bar{\delta}\right) - \frac{(\gamma - \bar{\gamma})\rho^h \delta}{\gamma} = \alpha,
\end{aligned}
\tag{24}
$$

for $\tau_k \longrightarrow \infty$, where we have used that $\mathrm{E}(x_t) = 0$, as $0 < \rho < 1$, $\mathrm{E}(\varepsilon_{t+h-i}) = 0$ for $h > 0$ and $i = 0, 1, \ldots, h-1$, and that $\widehat{\beta} \longrightarrow \beta$, as given in (22). Thus, if the subperiod is sufficiently long, the intercept estimator is approximately given by $\widehat{\alpha} \approx \alpha$. As $\gamma > 0$, $\delta > 0$, and $0 < \rho < 1$, it follows that if $\delta > \bar{\delta}$ and $\gamma \leq \bar{\gamma}$ then $\alpha > 0$; if $\delta < \bar{\delta}$ and $\gamma \geq \bar{\gamma}$ then $\alpha < 0$; and if $\delta = \bar{\delta}$ and $\gamma = \bar{\gamma}$ then $\alpha = 0$. ■