参赛学生姓名：吴子晋

中学：华南师范大学附属中学

省份：广东省

国家/地区：中华人民共和国

指导老师姓名：黄华林

指导老师单位：华南师范大学附属中学

论文题目：Understanding the Flattening Phillips Curve in China: The Role of Central Bank Communication and Inflation Expectations

Understanding the Flattening Phillips Curve in China: The Role of Central Bank Communication and Inflation Expectations

Zijin Wu

The Affiliated High School of SCNU

Abstract

This paper investigates the flattening slope of the New Keynesian Phillips Curve (NKPC) in China, focusing on the role of inflation expectations and central bank communication. Using quarterly data from 1993 to 2025, we construct a latent expectation index from forecasts and surveys, and develop novel text-based indices of the People's Bank of China's Monetary Policy Reports with large language models. IV-GMM estimates show that inflation dynamics are dominated by expectations, while the Phillips curve slope becomes insignificant or negative after 2013. Subsample and rolling analyses reveal a structural shift: before 2013 both expectations and demand pressures shaped inflation, whereas afterward demand effects vanish and inflation becomes expectation-driven. Communication indices further explain expectations: aggregate stance and forward-looking guidance reduce their level, financial stability concerns raise them, and forward guidance strength and structural policy focus enhance anchoring. Overall, the evidence suggests that conventional demand-based mechanisms have weakened, with expectations and communication now central to inflation dynamics and policy transmission in China.

Keywords: New Keynesian Phillips Curve; Inflation expectations; Central bank communication; Expectation anchoring; Monetary policy in China

1 Introduction

Recent developments in China's inflation dynamics present an unusually puzzling picture. On the production side, producer prices (PPI) have been in persistent deflation, reflecting weak industrial pricing power and global supply pressures. Yet at the macro level, real GDP growth has remained relatively robust, and measures of aggregate slack do not suggest severe underutilization. At the same time, households often perceive elevated inflation in daily consumption goods, with concerns over monetary expansion feeding into higher inflation expectations. The result is a striking divergence: deflation at the factory gate, stable but modest CPI inflation in the aggregate, and heterogeneous perceptions among firms, households, and policymakers. Such a paradox—falling costs in production, stable aggregates, but rising inflation concerns in perception—is rarely observed in other major economies. It makes China an especially compelling case for re-examining the Phillips Curve and, more broadly, for reassessing the central role of expectations in monetary policy transmission.

While salient in China, this challenge echoes a broader global puzzle: the flattening of the New Keynesian Phillips Curve (NKPC). A large literature has proposed several, not mutually exclusive, explanations. *(i) Anchored expectations.* With explicit inflation targets and systematic communication, expectations have become more firmly anchored, weakening the pass-through from real activity to prices. Evidence of declining slopes is documented by Blanchard (2016), while studies of forward guidance and central bank communication further underscore the importance of expectations management (Campbell et al., 2012; De Haan & Sturm, 2019; Swanson, 2021; Coibion et al., 2022). *(ii) Supply-side mechanisms.* Another strand emphasizes that marginal-cost pressures translate less into prices due to structural changes in price setting, producing unusually quiescent inflation dynamics (Del Negro et al., 2015, 2020). Microfounded models of price rigidity provide theoretical foundations for such flattening (Rotemberg, 1983; Sbordone, 2002; Galı & Gertler, 1999; Gali et al., 2005). *(iii) Expectation mismeasurement.* Coibion & Gorodnichenko (2015) show that professional forecasts fail to capture the expectations relevant for price setters; household and firm surveys, more sensitive to salient price movements, reconcile "missing disinflation" episodes. This implies that empirical slopes are partly an artifact of poor expectation proxies. Taken together, these explanations suggest that reduced-form NKPC slopes can appear small due to anchored expectations, supply-side shifts, noisy measurement of slack and expectations, and policy-regime non-linearities.

Turning to China, the Phillips Curve puzzle is closely tied to the formation and measurement of inflation expectations. Unlike advanced economies where professional forecasts or market-based expectations are often reliable proxies, China lacks a long history of consistent, high-quality household and firm surveys. As a result, empirical work has relied on indirect or noisy proxies, such as consumer sentiment indices or implicit measures extracted from bond yields, which may not reflect the expectations relevant for price and wage setters. This limitation partly explains why estimates of China's NKPC often yield unstable or insignificant slope coefficients (Zhang & Murasawa, 2011; Zhang, 2013).

Moreover, expectation formation in China is complicated by its transitional and state-influenced institutional context. Households frequently perceive high inflation in daily consumption items even when aggregate CPI remains subdued, while producer expectations are shaped by persistent PPI deflation and administrative price guidance. Such heterogeneity across agents creates a gap between aggregate measures of slack and the expectations channel that drives price-setting behavior. Indeed, recent evidence highlights that inflation dynamics in China are more closely tied to expectation shifts than to output gap fluctuations, underscoring the importance of better capturing how expectations are formed and managed by policy communication (Fernald et al., 2014; Girardin et al., 2017).

In addressing the difficulty of measuring expected inflation in China, we adopt a two-pronged approach. First, we draw on available survey proxies, including professional forecasts and diffusion indices of expected product and consumption prices. Second, we propose that the People's Bank of China's (*Monetary Policy Reports*) provide a systematic and forward-looking source of signals, which we quantify using large language models. By integrating these expectation measures into an extended NKPC framework, we test the role of communication in anchoring expectations and explaining China's inflation dynamics. This approach not only provides a new lens on the Phillips Curve in China but also contributes to the broader literature on the role of communication in monetary policy transmission.

2 Literature Review

2.1 The Phillips Curve and the Flattening Puzzle

The empirical debate on the NKPC has evolved substantially since its micro-founded formulation in the 1990s (Galı & Gertler, 1999; Sbordone, 2002). Early evidence suggested that inflation dynamics could be well explained by a hybrid model with forward- and backward-looking components (Gali et al., 2005). However, subsequent studies across the US, Euro Area, and Japan documented a striking decline in the slope of the Phillips Curve.

Blanchard (2016) characterizes this decline as a shift "back to the 1960s": expectations are better anchored, but inflation has become less sensitive to slack. Del Negro, Giannoni, and Schorfheide (2015) and Del Negro et al. (2020) show, using DSGE and VAR methods, that inflation appears unusually unresponsive to cyclical fluctuations, suggesting structural changes in price-setting behavior or global supply conditions. Other explanations include globalization, which disciplines domestic prices through import competition, and declining frequency of price adjustment, consistent with higher nominal rigidity.

Coibion and Gorodnichenko (2015) provide a complementary perspective. They argue that the missing disinflation after 2008 largely disappears once one uses household survey expectations rather than professional forecasts. This critique highlights expectation heterogeneity: while financial markets and forecasters may be well anchored, households and firms—the actual price- and wage-setters—remain sensitive to salient price shocks such as energy or food. Recent studies reinforce this view, showing that central bank credibility may bifurcate across agents, with implications for the stability of the NKPC.

For China, empirical evidence has been mixed. Zhang and Murasawa (2011) and Zhang (2013) attempt to estimate extended NKPCs including open-economy factors, but find that the output gap has limited explanatory power for inflation. Fernald, Spiegel, and Swanson (2014) argue that the effectiveness of monetary policy transmission itself is weaker than in advanced economies. He and Pauwels (2008) and Girardin et al. (2017) show that the PBoC's policy reaction function differs systematically from standard Taylor-rule type rules. These studies converge on a central point: standard NKPC formulations fail to capture China's inflation dynamics, not least because expectations and policy stance are poorly proxied by observable variables.

The remainder of the paper is organized as follows. Section 3 details the data and the construction of the communication indices. Section 4 empirically estimates Chinas NKPC slope using hybrid specifications (IV-GMM), with subsample, rolling-window, and robustness analyses. Section 5 analyzes the determinants of inflation expectations using the communication indices and evaluates their anchoring effects. Section 6 concludes with implications for monetary policy design in China and broader lessons for central bank communication.

2.2 Central Bank Communication and Textual Analysis

Parallel to the NKPC debate, a rich literature has emerged on central bank communication as a policy tool in its own right. Modern monetary policy operates not only through actions but also through words: forward guidance, policy reports, and press conferences shape expectations and financial conditions (Campbell et al., 2012; Swanson, 2021). Empirical studies show that communication shocks can move yields, exchange rates, and asset prices independently of policy rate changes.

Early research measured communication using simple event studies or dictionary-based sentiment indicators. Hansen and McMahon (2016) use computational linguistics to quantify the effects of Federal Reserve statements. Hansen, McMahon, and Prat (2018) examine FOMC deliberations, showing how transparency interacts with decision-making. De Haan and Sturm (2019) provide a comprehensive survey of central bank communication, emphasizing its evolution into a "fourth policy instrument."

Recent advances in natural language processing (NLP) and large language models (LLMs) have transformed this field. Gambacorta et al. (2024) introduce CB-LMs, specialized language models for central banking text, while Geiger et al. (2025) develop MILA, a monetary-intelligent language agent designed to classify central bank communication in a transparent manner. These innovations demonstrate the potential of LLMs to extract multi-dimensional policy signals, including tone, forward-looking orientation, and structural focus.

Despite this progress, applications to China remain scarce. Existing studies typically rely on aggregate indices or qualitative assessments of PBoC speeches and reports. To our knowledge, no study has systematically leveraged LLM-based methods to quantify the multidimensional stance of PBoC communication. This gap is particularly important given the PBoC's unique institutional architecture, where aggregate and structural objectives coexist and where transparency is more limited than in advanced economies.

Our study bridges these literatures. By introducing novel text-based indices of PBoC communication into NKPC-type models, we provide new evidence on the anchoring role of communication in shaping inflation expectations. More broadly, we contribute to comparative studies of monetary policy by examining how communication functions in a large transition economy with a dual-mandate framework.

Building on these insights, the next section introduces our data sources and the construction of the communication indices.

3 Data and Methodology

This section introduces the dataset and outlines our methodological framework. We first describe the macroeconomic variables used in the empirical analysis, including inflation, output gap, monetary aggregates, and policy rates. We then document the textual data, namely the *Monetary Policy Implementation Reports* published quarterly by the People's Bank of China since 2001, and explain how we preprocess these texts. Finally, we describe the construction of expectation proxies and the large language model (LLM)-based classification framework that produces four indices.

3.1 Data Sources

Our quarterly dataset spans 1993Q1–2025Q2 and combines both international and domestic sources. Inflation is measured using the Consumer Price Index (CPI, 2010=100) from CEIC (original source: IMF). The raw CPI series is seasonally adjusted using the U.S. Census Bureau's X-13 procedure, and quarterly inflation is defined as the log difference of the seasonally adjusted CPI (annualized).

Output is taken from CEIC's real GDP series (constant 2020 prices). We apply X-13 seasonal adjustment to the level series and then use the Hodrick–Prescott filter ($\lambda = 1600$) to separate trend and cycle. The output gap is defined as the percentage deviation of actual GDP from the HP-filtered trend.

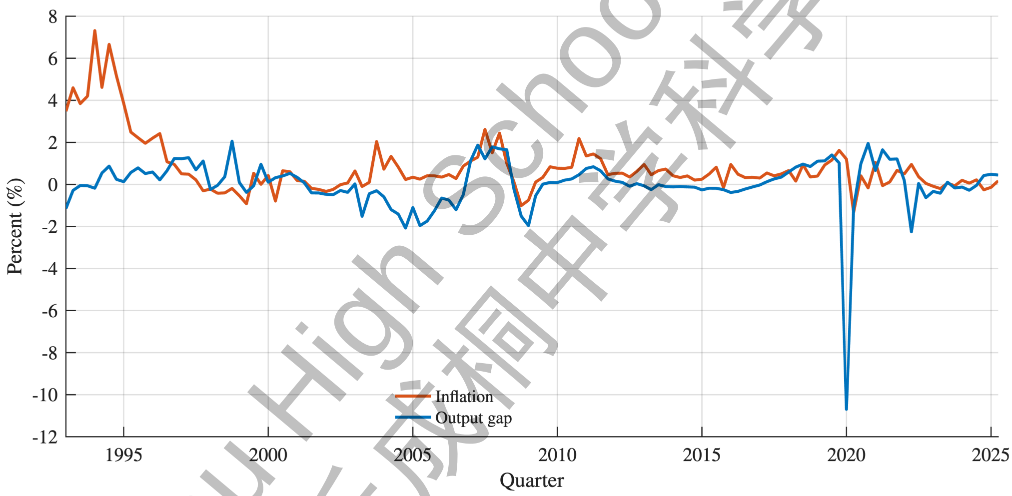

Figure 1: Stylized fact: Inflation-output gap disconnect in China.

Figure 1 illustrates the weak and unstable relationship between inflation and the output gap in China. While inflation shows pronounced swings in the 1990s and early 2000s, the output gap remains relatively muted, and the two series often diverge in recent decades. This stylized fact motivates our re-examination of the Phillips Curve mechanism.

In addition to these core variables, we construct two external controls. First, oil price shocks are measured using Brent crude oil prices (FRED), deflated by U.S. CPI to obtain a real oil price series. We aggregate the monthly data to quarterly averages and define the shock as the log quarterly change. Second, exchange rate shocks are based on the USD/CNY nominal exchange rate (FRED, daily), averaged to quarterly frequency, with the shock defined as the log quarterly change.

Finally, inflation expectations are proxied using three complementary sources: professional forecasts from the IMF World Economic Outlook (WEO), the diffusion index of expected product prices from the PBoC's enterprise survey, and the diffusion index of expected consumer prices from the PBoC's household survey. In the next subsection, we detail how these series are combined into a latent inflation expectation index.

3.2 Inflation Expectation Proxy: Construction and Stylized Patterns

Method. To measure expected inflation in China, we combine three complementary sources: (i) professional forecasts from the IMF World Economic Outlook (WEO), originally available at annual frequency but interpolated to quarterly using a state-space smoothing procedure, (ii) the household diffusion index of expected consumer prices, and (iii) the enterprise diffusion index of expected product prices. All series are mapped to quarterly frequency and standardized (z-score). We then estimate a one-factor Dynamic Factor Model (DFM) with scalar idiosyncratic variances to extract the common component. Finally, we linearly rescale the factor to inflation units by matching the factor's standardized average to the level of WEO forecasts over the overlap sample. The resulting series is our *latent inflation expectation*.

Main patterns. Figure 2 shows that the latent index (i) smooths the high-frequency noise in individual series while preserving turning points, (ii) tracks professional forecasts more closely in the 1990s but aligns increasingly with survey perceptions thereafter, and (iii) sits between the household and enterprise measures when they diverge, providing an aggregated, theoretically grounded proxy for expected inflation. We use this latent series as the benchmark expectation term in the extended NKPC estimations below.

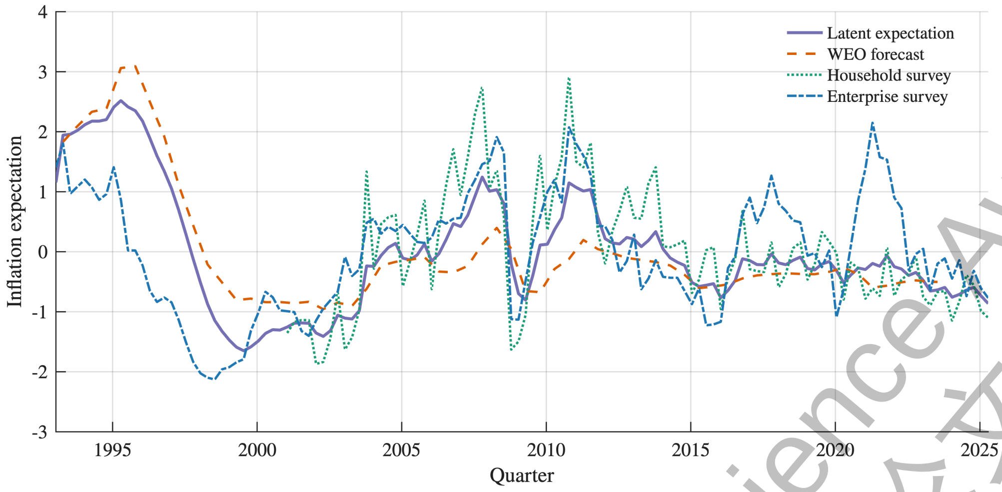

Figure 2: Latent inflation expectation (rescaled DFM factor) versus components

3.3 Text Feature

We construct a set of quarterly text-based indices from the *PBoC Monetary Policy Reports* to quantify the central bank's communication along different dimensions. Our approach is motivated by the growing literature that employs textual analysis to measure monetary policy stance and guidance (e.g., Hansen & McMahon (2016); Shapiro et al. (2022); Bholat et al. (2015)). In particular, we focus on four dimensions that capture distinct facets of central bank communication:

- **Aggregate Policy Stance Index (APSI):** captures the overall hawkish or dovish tone when the report discusses aggregate monetary policy, following the idea that the policy stance embedded in narrative text can shift public expectations.
- **Forward-Looking Guidance Indices (FLGI):** consist of both the *strength* of forward guidance (the share of sentences with explicit forward-looking content) and the *direction* of such guidance (the weighted stance of forward-looking statements). These indices are motivated by the literature emphasizing the role of forward guidance in anchoring expectations (Campbell et al. (2012); Del Negro et al. (2023)).
- **Structural Policy Focus Index (SPFI):** measures the relative attention given to structural or sectoral policies (e.g., credit support to SMEs, green finance, or real estate). This reflects the increasing importance of targeted policies in China's monetary framework.

- **Financial Stability Concern Index (FSCI):** averages the reported level of concern over systemic risk, asset bubbles, or macroprudential issues. This dimension draws on recent work linking central bank communication about risks with market perceptions of uncertainty (Born et al. (2014)).

A key methodological innovation of this paper is the use of a large language model (LLM, implemented here with `gpt-5-mini`) for large-scale text mining. Our corpus spans 2001Q1 to 2025Q2 and contains 67,437 sentences from the *PBoC Monetary Policy Reports*, a scale that renders manual coding practically impossible. The LLM-based approach allows us to handle this volume of text consistently and reproducibly. All LLM calls were executed in a deterministic setup (temperature = 0), ensuring consistency and reproducibility of the sentence-level classifications across the entire corpus. Moreover, our design feeds information at the sentence level, which eliminates the risk of contextual leakage: the model classifies each sentence without access to surrounding passages or contemporaneous data, and thus cannot implicitly condition on knowledge of the broader economic environment. This feature reduces the potential biases that might arise if human experts, aware of the historical context, let prior beliefs or hindsight influence their classification. Finally, compared with traditional bag-of-words or topic models, which face well-known difficulties in processing Chinese text due to segmentation ambiguity and weaker sentiment lexicons, the LLM approach can capture semantic nuance and institutional jargon more effectively. In this way, the LLM offers both scalability and impartiality that are difficult to achieve through conventional text-mining techniques.

Our contribution relative to the existing literature is therefore twofold. On the substantive side, we introduce a new set of indices that jointly cover stance, forward guidance, structural orientation, and financial stability concerns from the PBoC's official reports. On the methodological side, we demonstrate the feasibility and advantages of applying LLM-based labeling to central bank text, highlighting how this approach can extend beyond standard sentiment analysis toward multidimensional policy indicators.

Figure 3 shows that the five communication indices exhibit heterogeneous dynamics. The **APSI** displays a gradual upward drift, with notable reversals around the 2007–2008 global financial crisis and the 2010–2011 inflationary episode, consistent with shifts between hawkish and dovish aggregate policy tones. The **FLGI_strength** rises steadily across the sample, highlighting the increasing use of forward-looking language in PBoC reports. By contrast, the **FLGI_direction** is highly volatile, alternating between dovish and hawkish signals in line with cyclical macroeconomic conditions, even though the smoothed trend tilts upward. The **SPFI** shows a clear and persistent upward trend, underscoring the growing emphasis on structural and targeted policy objectives since the early 2010s. Finally, the **FSCI** exhibits episodic peaks—notably during 2007–2009, 2015–2017, and 2020–2022—corresponding to periods of elevated systemic risk. Since 2022, however, the index has declined, in line with easing inflation expectations and a reduced emphasis on financial stability concerns.

Robustness Check: Consistency of LLM-based Coding

A potential concern with LLM-based text mining is whether the resulting indices are stable across repeated runs of the same model, given that large language models may exhibit stochastic variation. To address this, we re-processed the *PBoC Monetary Policy Reports* for the year 2001 using identical prompts and parameter settings with `gpt-5-mini`, and compared the newly generated indices with those from our main run.

At the *quarterly index level*, the agreement is very high. The correlations across the two runs are 0.90 for APSI, 0.97 for FLGI_strength, 0.81 for SPFI, and 0.95 for FSCI. The only dimension with somewhat lower correlation is FLGI_direction (0.74), which is expected because this index is a weighted stance measure and therefore more sensitive to small differences at the sentence level. On average, the mean absolute deviations across indices remain below 0.20.

At the *sentence level*, where 320 overlapping sentences were double-coded, consistency is also high. The classification matches 87% of the time for `time_orientation`, 89% for `policy_scope`, and 92% for `policy_stance`. The correlation in stance values is 0.72, with a mean absolute difference of 0.08. For financial stability concerns, the correlation is 0.74 with a mean absolute difference of 0.34 on a 0–3 scale.

Overall, these results suggest that the indices constructed with `gpt-5-mini` are highly reproducible. While some variation exists at the margin—particularly for forward guidance direction—the aggregate indices are remarkably robust across repeated runs. This reassures us that the findings reported below are not artifacts of model randomness but reflect stable features of the underlying text.

4 The New Keynesian Phillips Curve in China

We now turn to estimating the slope of the New Keynesian Phillips Curve (NKPC) for China. Existing studies often document weak or unstable Phillips curve coefficients when using simple output gap measures, suggesting that the forward-looking NKPC alone provides an incomplete description of China's inflation dynamics. Consistent with this evidence, our baseline regressions indicate that a purely forward-looking specification yields imprecise and time-unstable coefficients. We therefore adopt a hybrid (indexation) NKPC, which nests both forward- and backward-looking terms, and evaluate the role of expectations using survey-based and latent proxies.

4.1 Empirical Specification and Estimation

Our benchmark specification is the hybrid NKPC:

$$\pi_t = \alpha\,\mathbb{E}_t \pi_{t+1} + \gamma\,\pi_{t-1} + \kappa\,y_t + \mathbf{z}_t'\delta + \varepsilon_t, \tag{1}$$

where π_t denotes quarterly CPI inflation, y_t is the output gap (HP-filtered real GDP), and $\mathbf{z}_t$ collects cost-push controls (real oil price growth and the USD/CNY depreciation rate). Expectations $\mathbb{E}_t \pi_{t+1}$ are proxied by a latent factor extracted from three sources (professional forecasts, household diffusion index, and enterprise diffusion index), as discussed in Section 3.2. The forward- and backward-looking weights, α and γ, capture the importance of expectations versus inertia in driving inflation dynamics, while κ measures the slope of the Phillips curve.

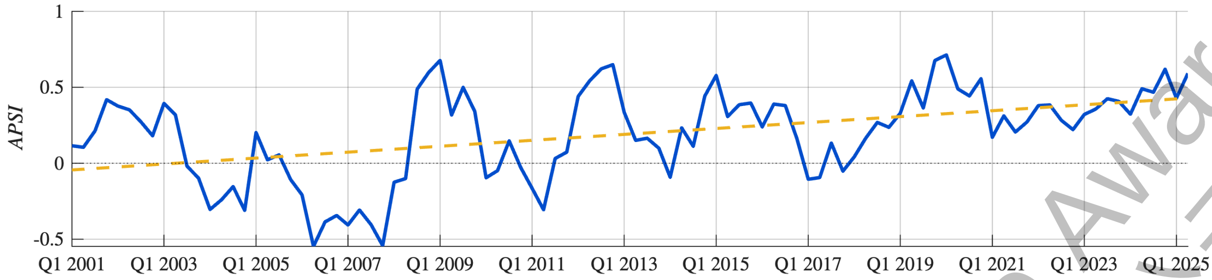

(a) APSI

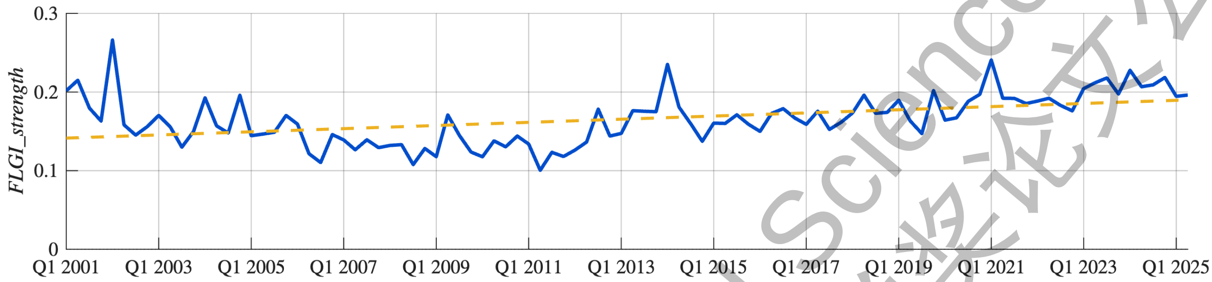

(b) FLGI strength

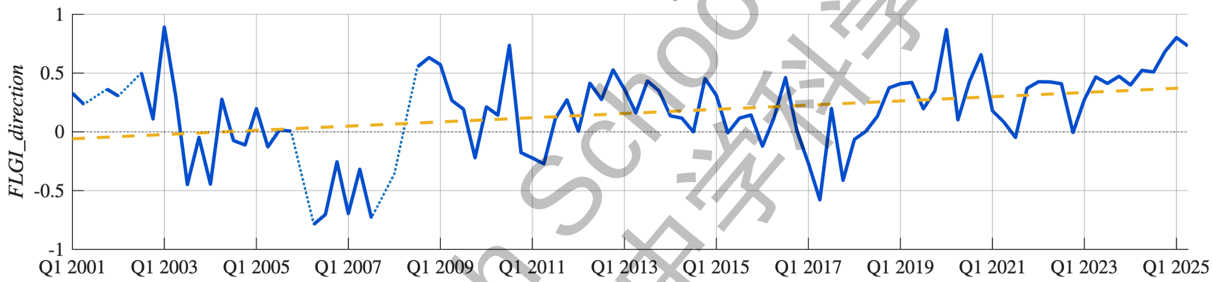

(c) FLGI direction

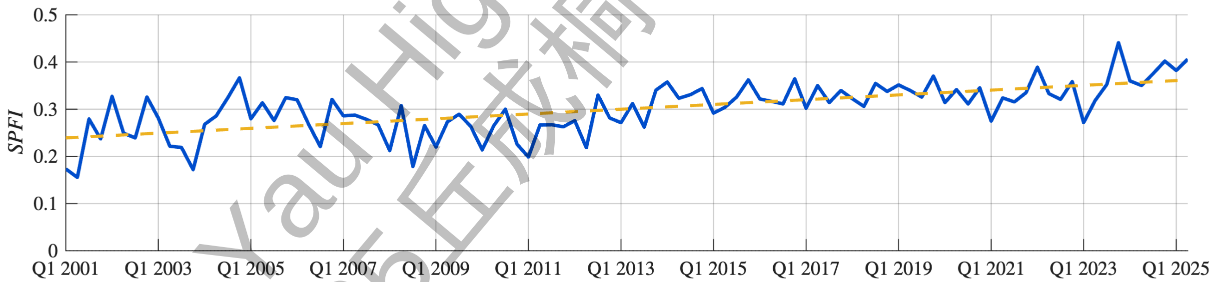

(d) SPFI

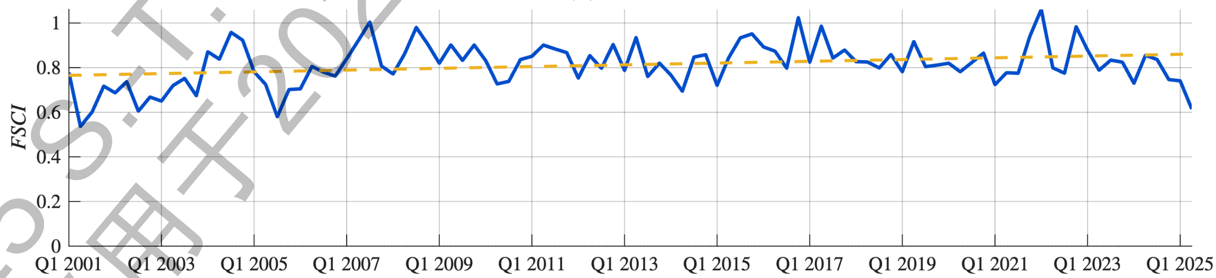

(e) FSCI

Figure 3: Time series of communication indices (2001Q1-2025Q2)

Estimation. Equation (1) is estimated by GMM to account for endogeneity of expectations and the output gap. We use lagged values as instruments:

$$\mathscr{I}_{t-1} = \{\pi_{t-1}, \pi_{t-2}, y_{t-1}, y_{t-2}, \text{oil}_{t-1}, \text{fx}_{t-1}, 1\},$$

and report heteroskedasticity- and autocorrelation-robust (HAC) standard errors. As robustness checks, we: (i) replace the latent factor with individual expectation proxies (WEO forecasts, household DI, enterprise DI), (ii) set $\gamma = 0$ to estimate a purely forward-looking NKPC, and (iii) vary the instrument set and lag length. In addition, we present subsample (pre-2013 vs. post-2013) and rolling-window estimates to assess structural change and time variation in κ and α.

In this framework, α is expected to be close to one if the proxy captures rational expectations accurately; $\gamma > 0$ would indicate the presence of indexation or inflation inertia; and κ reflects the sensitivity of inflation to economic slack. A flatter Phillips curve corresponds to a smaller κ, while stronger expectation anchoring is reflected in a larger and more stable α, together with reduced volatility of the expectation-formation residuals.

4.2 Main Estimation Results

Table 1 reports the IV–GMM estimates of the hybrid NKPC. In the full sample (2001Q1–2024Q2), the expectation term enters positively and is highly significant ($\hat{\alpha} = 0.136$, s.e. $= 0.028$), underscoring the central role of expectations in shaping Chinese inflation dynamics. Lagged inflation has a positive sign (0.220, s.e. $=$ 0.134) but is imprecisely estimated, suggesting limited inertia once expectations are explicitly controlled for. The Phillips curve slope is close to zero and statistically insignificant ($\hat{\kappa} = 0.021$, s.e. $= 0.058$), consistent with earlier findings that conventional output-gap measures have little predictive power for CPI inflation in China. Among the controls, oil price shocks are robustly significant (0.011, s.e. $=$ 0.003), while exchange rate effects remain small and indistinguishable from zero.

Splitting the sample at 2013 reveals a sharper contrast. In the earlier period, both expectations and the output gap exerted strong influence: $\hat{\alpha} = 0.219$ (s.e. $= 0.079$) and $\hat{\kappa} =$

0.295 (s.e. $=$ 0.099), indicating that inflation was sensitive to both forward-looking beliefs and demand pressures. In the later period, by contrast, the slope coefficient turns negative ($\hat{\kappa} = -0.073$, s.e. $=$ 0.026), while expectations remain precisely estimated ($\hat{\alpha} = 0.154$, s.e. $=$ 0.024). A simple Wald test rejects equality of the two κ estimates ($t \approx 3.6$), whereas the difference in α across subsamples is small and not statistically significant. This "flattening of the Phillips curve alongside stable expectations" suggests that demand-driven inflation dynamics have weakened, while expectations continue to anchor price setting.

Taken together, the subsample results point to a structural change: before 2013, inflation displayed a "dual engine" structure—driven by both expectations and demand—whereas afterwards, demand effects vanish and inflation dynamics become expectation-dominated. We next examine rolling window estimates to trace this transition more finely over time and to assess whether the subsample contrast reflects a gradual flattening or a sharper structural break.

Table 1: Hybrid NKPC Estimates for China (IV-GMM, robust SEs)

	Full sample (2001Q1–2021Q4)	Pre-2013	Post-2013
Output gap (κ)	0.0206	0.2945***	-0.0728***
	(0.0579)	(0.0994)	(0.0263)
Expected inflation (α)	0.1356***	0.2195***	0.1538***
	(0.0282)	(0.0786)	(0.0240)
Lagged inflation (γ)	0.2200	-0.1602	0.0650
	(0.1342)	(0.2992)	(0.1360)
Oil shock	0.0107***	0.0144***	0.0060
	(0.0034)	(0.0033)	(0.0038)
FX shock	0.0085	-0.0525	0.0174
	(0.0281)	(0.1476)	(0.0235)
Constant	–	–	–
Observations	98	48	50
R^2	0.658	0.784	0.516
Estimator	IV-GMM	IV-GMM	IV-GMM
Covariance	Robust (HAC)	Robust (HAC)	Robust (HAC)

Notes: Standard errors in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$. The dependent variable is quarterly CPI inflation. Instruments: $\{\pi_{t-2}, \pi_{t-3}, y_{t-1}, y_{t-2}\}$.

4.3 Extention: Rolling Estimates

To further examine the stability of the NKPC, we estimate the hybrid specification in rolling 60-quarter windows (Figures 4a–4c). The rolling trajectories provide a complementary, more continuous view of the dynamics already suggested by the subsample regressions.

The slope coefficient κ exhibits a clear time profile. In early windows, confidence intervals lie well above zero, implying that demand conditions exerted a positive and significant effect on inflation. After 2012, however, the rolling estimates collapse and oscillate around zero, with intervals frequently dipping into negative territory. This mirrors the post-2013 subsample evidence of a flattened or even inverted Phillips curve. By contrast, the forward-looking weight λ_f remains positive and stable throughout, with confidence bands rarely overlapping zero. The backward-looking weight λ_b, while sizeable in the early 2000s, steadily declines towards zero and loses significance, consistent with the idea that nominal inertia has weakened over time.

Why has this structural shift occurred? We highlight three complementary explanations. First, measurement error in the output gap may have become more pronounced over time: slowing potential growth, structural rebalancing, and capacity adjustment after 2013 mean that HP-filtered gaps increasingly misrepresent "effective slack," possibly misattributing cost–price dynamics to demand, and thereby yielding negative κ in linear regressions. Second, cost-push forces have gained importance: oil price shocks are consistently significant, while the CPI basket has become more sensitive to food, energy, and regulated prices; if cost shocks dominate while demand fluctuations converge, the estimated Phillips slope naturally flattens. Third, expectation management and policy frameworks have strengthened: both subsample and rolling results show a more stable λ_f, alongside an insignificant λ_b, consistent with stronger anchoring of expectations and diminished nominal inertia. Together these factors reduce the persistent effect of demand shocks, making the Phillips curve appear flatter or even countercyclical at observable business-cycle frequencies.

From both identification and policy perspectives, these findings are notable. On the identification side, our first-stage diagnostics indicate strong instruments across windows and subsamples, ruling out weak-IV concerns. On the policy side, a natural interpretation is that in the later period, conventional demand management has lost traction in steering CPI. Instead, stabilizing and guiding inflation expectations (via communication frameworks and credible targets) and managing cost-push pressures (energy, food, key imports) have become more direct and effective levers for influencing the short-run inflation path.

4.4 Robustness Analysis

We conduct a series of robustness checks to verify that the main findings are not artifacts of a particular sample, specification, or instrument set. Table 2 summarizes the preferred results, focusing on minimal IV sets with kernel-HAC standard errors.

Table 2: Robustness checks of NKPC estimates (preferred specifications)

Spec	Sample	$\hat{\alpha}$	$\hat{\gamma}$	$\hat{\kappa}$	N	$p(J)$
FL	Full (2001Q1–2024Q2)	0.413 (0.059)	—	-0.024 (0.049)	98	0.800
HY	Full (2001Q1–2024Q2)	-0.356 (0.538)	-0.599 (1.265)	0.064 (0.078)	98	0.331
FL	Pre-2013	0.248 (0.108)	—	0.224 (0.050)	48	0.905
HY	Pre-2013	0.656 (0.162)	-0.529 (0.332)	-0.005 (0.074)	48	0.365
FL	Post-2013	0.363 (0.108)	—	-0.071 (0.021)	50	0.070
HY	Post-2013	0.689 (0.309)	-0.804 (0.352)	-0.165 (0.043)	50	0.372

Notes: FL = forward-looking NKPC ($\gamma = 0$); HY = hybrid NKPC. Reported coefficients with HAC (Bartlett, bw=4) standard errors in parentheses. N is the number of observations. $p(J)$ reports Hansen's J-test p-value.

Three key observations emerge.

First, **expectations remain the dominant driver**. Across all subsamples and specifications, the forward-looking coefficient α is consistently positive and typically significant (ranging between 0.25 and 0.70). This confirms that expectations are a robust determinant of Chinese inflation dynamics, regardless of instrument choice or covariance estimation.

Second, **output-gap sensitivity is sample-dependent**. Before 2013, κ is positive and significant in the forward-looking specification (0.22, s.e. $=$ 0.05), indicating that demand pressures played a material role in inflation. After 2013, however, κ turns negative and statistically significant in both FL and HY regressions, consistent with a flattened or even inverted Phillips

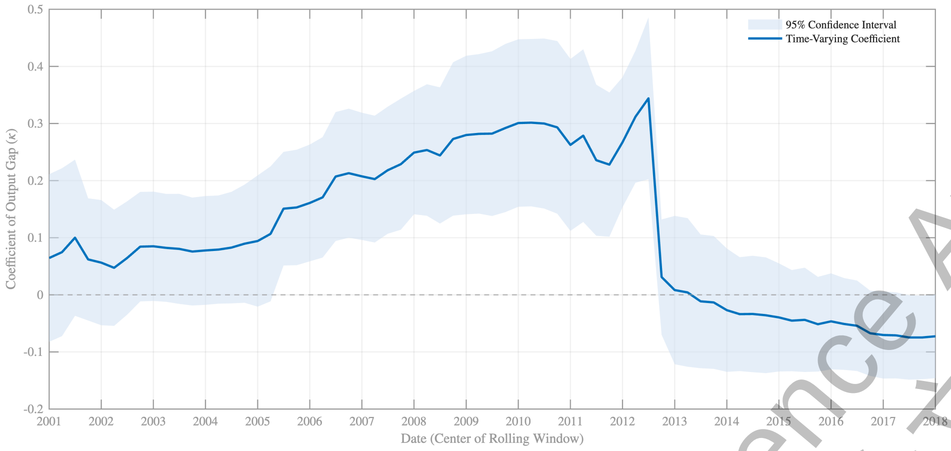

(a) Time-varying Phillips curve slope κ.

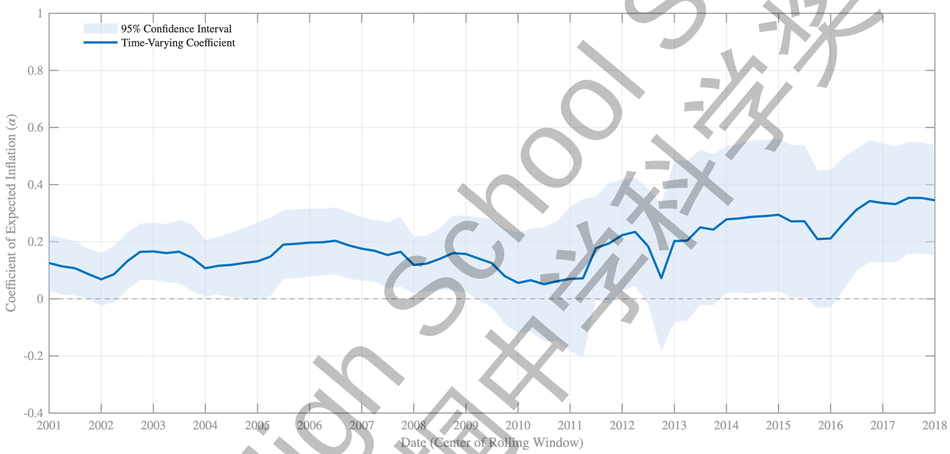

(b) Time-varying forward-looking weight λ_f.

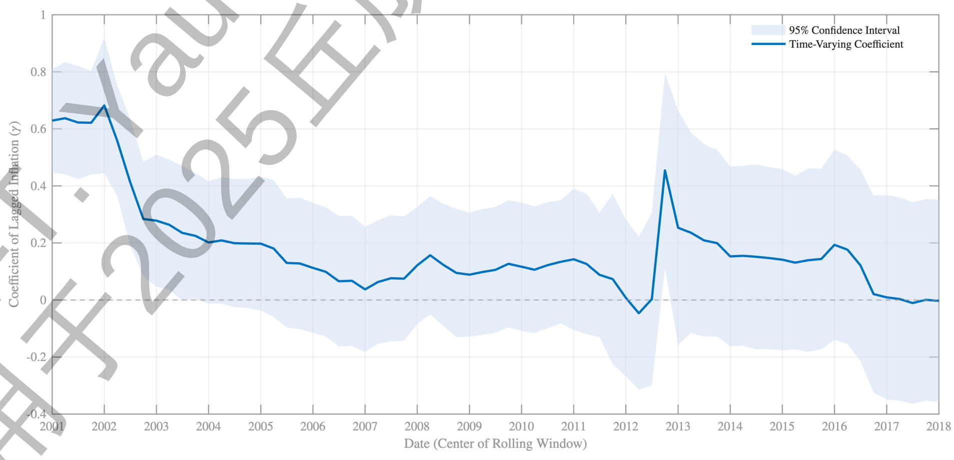

(c) Time-varying backward-looking weight λ_b.

Figure 4: Rolling 60-quarter estimates of the hybrid NKPC coefficients. Each panel reports the coefficient path with 95% confidence intervals.

curve.

Third, **hybrid specifications are less stable**. In full-sample estimates, HY regressions yield imprecise coefficients for both α and γ, reflecting sensitivity to instrument sets and possible multicollinearity. Nonetheless, the forward-looking component α remains positive and comparable in magnitude to FL results, reinforcing the conclusion that expectations dominate inertia in shaping inflation dynamics.

Overall, these robustness exercises strengthen our baseline conclusion: inflation in China was jointly driven by expectations and demand pressures before 2013, but since then the Phillips curve has flattened, leaving expectations and cost-push shocks as the primary channels of inflation dynamics.

5 Inflation Expectation Driven Factors

We next turn to the drivers of inflation expectations. Using the latent index derived from professional forecasts, household surveys, and enterprise surveys, we examine how central bank communication helps shape and anchor expectations.

5.1 Communication Indices and Inflation Expectations

We begin by examining whether central bank communication helps to explain movements in inflation expectations. Let Exp_t denote the latent index of inflation expectations, constructed from professional forecasts, household surveys, and enterprise surveys. We consider the following baseline regression specification:

$$Exp_t = \alpha + \beta_1 APSI_t + \beta_2 FLGI_t^{\text{strength}} + \beta_3 FLGI_t^{\text{direction}} + \beta_4 SPFI_t + \beta_5 FSCI_t + \varepsilon_t, \quad (2)$$

where the coefficients β_j capture the contemporaneous association between each communication dimension and inflation expectations.

We first estimate equation (2) in levels using OLS with heteroskedasticity- and autocorrelation-robust (HAC) standard errors. To address potential non-stationarity, we also conduct unit-root and cointegration tests. When cointegration is present, we employ the Engle-Granger two-step approach to estimate long-run relationships and corresponding error-correction models (ECM). Finally, we extend the analysis with dynamic specifications that include lags of both Exp_t and the communication indices, allowing us to assess predictive content.

Table 3: Pairwise correlation matrix (levels)

	Exp	APSI	FLGI_strength	FLGI_direction	SPFI	FSCI
Exp	1.000	**-0.498**	**-0.522**	**-0.419**	**-0.149**	**0.389**
APSI	-0.498	1.000	0.321	0.803	0.271	-0.050
FLGI_strength	-0.522	0.321	1.000	0.353	0.495	-0.195
FLGI_direction	-0.419	0.803	0.353	1.000	0.216	-0.189
SPFI	-0.149	0.271	0.495	0.216	1.000	0.212
FSCI	0.389	-0.050	-0.195	-0.189	0.212	1.000

Notes: Pearson correlations in levels over 2001Q1–2025Q2. Exp denotes the latent inflation-expectation index (`exp_filtered`). **Bold** highlights Corr(Exp, ·).

Table 3 shows that APSI and FLGI_strength are strongly negatively correlated with inflation expectations (both around -0.5), while FSCI is moderately positively correlated (0.39). FLGI_direction also displays a negative link (-0.42), whereas SPFI is weak. These patterns suggest that communication tone and stability concerns are systematically related to expectations. To assess the independent contribution of each index beyond simple correlations, we next turn to multivariate regressions with HAC standard errors.

Table 4 summarizes the baseline regressions. In the single-index specifications (columns 1–5), APSI, FLGI_strength, and FLGI_direction all enter with significant negative coefficients, while FSCI enters significantly positive; SPFI is weak and not statistically different from zero. When all indices are included simultaneously (column 6), APSI and FLGI_strength remain robustly negative and FSCI remains positive, whereas FLGI_direction and SPFI lose significance due to collinearity. Overall, the communication indices together explain about half of the variation in expectations ($R^2 \approx 0.50$).

While informative, these OLS regressions may be sensitive to persistence in the underlying series. To examine whether the relationships reflect a stable long-run equilibrium rather than spurious correlation, we next test for cointegration and estimate an error-correction specification (Table 5).

Table 5 presents the cointegration and ECM results. The long-run regression indicates that APSI and FLGI_strength exert significant negative effects, while FSCI enters positively, consistent with the baseline OLS. Importantly, the residuals are stationary, confirming a stable cointegrating relation between expectations and communication indices.

Table 4: Inflation expectations regressed on communication indices (OLS with HAC)

	(1) APSI	(2) FLGI_str	(3) FLGI_dir	(4) FSCI	(5) SPFI	(6) All indices
APSI	-0.498***					-0.521***
	(0.127)					(0.154)
FLGI_strength		-0.522***				-0.405***
		(0.128)				(0.101)
FLGI_direction			-0.419***			0.179
			(0.117)			(0.123)
FSCI				0.389***		0.298***
				(0.143)		(0.107)
SPFI					-0.149	0.091
					(0.190)	(0.104)
Observations	98	98	98	98	98	98
R^2	0.248	0.272	0.176	0.151	0.022	0.498

Notes: Dependent variable is standardized latent inflation expectations (`exp_filtered_z`). Columns (1)–(5) report single-index OLS with HAC robust SE (maxlags=4); with standardized variables, coefficients coincide with pairwise correlations, so $R^2 \approx \text{Corr}^2$. Column (6) includes all indices simultaneously (HAC SE). Stars denote 1% (***), 5% (**), and 10% (*).

Table 5: Long-run cointegration and short-run adjustment (Engle–Granger + ECM)

	Long-run OLS (levels)	Error-correction model (ECM)
APSI	-1.918*** (0.475)	–
FLGI_strength	-12.603*** (2.821)	dL2: 4.320** (1.906); dL3: 3.547** (1.783)
FLGI_direction	0.549 (0.398)	not significant
FSCI	3.761*** (0.888)	dL1: -1.224** (0.477)
ECM(-1)	–	-0.214*** (0.038)
Observations	98	93
R^2	0.49	0.45

Notes: Left panel reports the long-run cointegration regression; residuals reject unit root (ADF $p < 0.001$). Right panel reports the short-run ECM with 4 lags of differences and the lagged error-correction term. HAC robust SE in parentheses. Stars denote 1% (***), 5% (**), and 10% (*).

The ECM shows that short-run adjustments are also meaningful. Changes in FLGI_strength (lags 2–3) raise expectations, while higher FSCI (lag 1) temporarily lowers them. The error-correction term is negative and significant (-0.21), implying that deviations from the long-run equilibrium are corrected at about 20% per quarter. This supports the view that communication indices shape both the level and the adjustment dynamics of expectations.

Table 6: Granger causality tests: do communication indices predict inflation expectations?

	Lag 1	Lag 2	Lag 3	Lag 4
APSI	0.0005***	0.008***	0.003***	0.002***
FLGI_strength	0.130	0.206	0.178	0.010**
FLGI_direction	0.0002***	0.007***	0.008***	0.021**
SPFI	0.192	0.314	0.172	0.108
FSCI	0.227	0.565	0.606	0.563

Notes: Entries are p-values from Granger causality tests of whether lagged values of each index jointly predict inflation expectations (up to the specified horizon). Stars denote 1% (***), 5% (**), and 10% (*).

Table 6 investigates whether communication indices help forecast expectations. APSI and FLGI_direction consistently Granger-cause expectations across multiple lags, underscoring the predictive content of aggregate stance and forward-looking direction. By contrast, FSCI and SPFI do not show predictive power once expectations' own persistence is accounted for. Taken together, these results highlight that communication matters not only contemporaneously but also in shaping agents' forward-looking beliefs.

5.2 Anchoring of Expectations

To examine whether central bank communication contributes to the anchoring of inflation expectations, we construct a deviation measure relative to the long-run trend. Specifically, we extract the smooth trend component of the latent expectation index (exp_t) using the Hodrick–Prescott filter with $\lambda = 1600$ for quarterly data. The anchoring gap is then defined as

$$gap_t = exp_t - trend_t,$$

and we consider both the absolute deviation ($|gap_t|$) and its square (gap_t^2) as dependent variables. Smaller values indicate that expectations lie closer to their long-run trend and are thus better anchored.

We regress these measures on the four communication indices—APSI, FLGI_strength, FLGI_direction, SPFI, and FSCI—using OLS with HAC-robust standard errors:

$$gap_t^{(m)} = \alpha + \beta_1 APSI_t + \beta_2 FLGI\_strength_t + \beta_3 FLGI\_direction_t + \beta_4 SPFI_t + \beta_5 FSCI_t + \varepsilon_t,$$

where $gap_t^{(m)}$ denotes either $|gap_t|$ or gap_t^2.

Table 7: Anchoring regressions: deviations from HP-trend

	$\lvert gap \rvert$	gap^2
APSI	0.041 (0.291)	0.281 (0.544)
FLGI_strength	-2.996** (1.364)	-4.345** (2.163)
FLGI_direction	-0.022 (0.144)	-0.011 (0.177)
SPFI	-2.712*** (0.967)	-4.487** (2.194)
FSCI	0.980** (0.438)	1.677** (0.827)
Observations	98	98
R^2	0.255	0.168

Notes: Dependent variables are the absolute ($|gap|$) and squared (gap^2) deviations of the expectation index from its HP-filtered trend ($\lambda = 1600$). OLS with HAC-robust standard errors (maxlags=4). Standard errors in parentheses. ** $p < 0.05$, *** $p < 0.01$.

Table 7 reports the results. Two indices emerge as significant stabilizers: forward-looking guidance strength (FLGI_strength) and structural policy focus (SPFI) both enter with negative and significant coefficients, indicating that stronger forward-looking communication and emphasis on structural objectives help anchor expectations. By contrast, financial stability concerns (FSCI) exhibit positive and significant coefficients, suggesting that heightened stability discussions are associated with larger deviations from the trend, i.e. weaker anchoring. APSI and FLGI_direction do not significantly affect anchoring in these regressions.

5.3 Discussion: Level vs. Anchoring Effects

The empirical results can be organized into a simple framework distinguishing between *level effects* and *anchoring effects* of central bank communication.

Level effects refer to shifts in the *mean level* of inflation expectations. Our regressions on the expectation index itself show that the Aggregate Policy Stance (APSI) and Forward-Looking Guidance strength (FLGI_strength) both enter with significantly negative coefficients, while Financial Stability Concerns (FSCI) enters positively. This pattern suggests that more dovish aggregate stances or stronger forward guidance tend to reduce inflation expectations, whereas heightened stability discussions are associated with higher expectations. In other words, communication shapes the direction of expectations in the short run.

Anchoring effects, by contrast, capture whether expectations remain close to their long-run trend. Using the HP-filter-based deviation measures, we find that FLGI_strength and the Structural Policy Focus (SPFI) significantly reduce deviations, thereby reinforcing anchoring. By contrast, FSCI again works in the opposite direction, with stability-related concerns being associated with larger gaps and weaker anchoring. The aggregate stance (APSI) and the directional component of guidance (FLGI_direction) do not show robust anchoring effects.

Taken together, this framework highlights two distinct channels. Some communication dimensions primarily influence the level of expectations, while others are more relevant for anchoring around the long-run path. The dual perspective provides a richer interpretation of the role of central bank communication: it not only repositions expectations in the short run but also determines whether those expectations remain stably anchored over time.

6 Conclusion

This paper documents a structural transformation in China's Phillips curve dynamics. Whereas inflation was once responsive to both expectations and demand pressures, since 2013 the sensitivity to demand has largely disappeared, and expectations have emerged as the principal driver of price-setting behavior. This flattening, and in some instances inversion, of the Phillips curve highlights the limitations of conventional output-gap-based models for understanding Chinese inflation.

Our analysis shows that central bank communication plays a systematic role in shaping expectations. Different dimensions of communication affect expectations in distinct ways: aggregate stance and forward-looking guidance shift the level of expectations, while forward-looking guidance strength and structural focus contribute to their anchoring. In contrast, concerns about financial stability, while important in their own right, tend to raise expectations and weaken anchoring.

From a policy perspective, these findings suggest that managing expectations through transparent and credible communication may be more effective than relying on traditional demand management tools to stabilize inflation. The People's Bank of China's growing use of forward-looking and structural policy language appears to have enhanced the anchoring of expectations, even as the Phillips curve has flattened. At the same time, heightened emphasis on financial stability can introduce volatility in expectations, underscoring the trade-offs in communication design.

Overall, the results call for a reevaluation of monetary policy frameworks in economies facing similar structural shifts. In particular, strengthening the credibility and consistency of communication can provide a powerful complement to interest rate policy, ensuring that expectations remain well anchored and inflation dynamics are stable in the face of evolving macroeconomic challenges.

References

Bholat, D., Hansen, S., Santos, P., & Schonhardt-Bailey, C. (2015). Text mining for central banks: handbook. *Centre for Central Banking Studies Handbook*, (33), 1–19.

Blanchard, O. (2016). The phillips curve: back to the'60s? *American Economic Review*, 106(5), 31–34.

Born, B., Ehrmann, M., & Fratzscher, M. (2014). Central bank communication on financial stability. *The Economic Journal*, 124(577), 701–734.

Campbell, J. R., Evans, C. L., Fisher, J. D., Justiniano, A., Calomiris, C. W., & Woodford, M. (2012). Macroeconomic effects of federal reserve forward guidance [with comments and discussion]. *Brookings papers on economic activity*, (pp. 1–80).

Coibion, O. & Gorodnichenko, Y. (2015). Is the phillips curve alive and well after all? inflation expectations and the missing disinflation. *American Economic Journal: Macroeconomics*, 7(1), 197–232.

Coibion, O., Gorodnichenko, Y., & Weber, M. (2022). Monetary policy communications and their effects on household inflation expectations. *Journal of Political Economy*, 130(6), 1537–1584.

De Haan, J. & Sturm, J.-E. (2019). Central bank communication. *The Oxford handbook of the economics of central banking*, (pp. 231).

Del Negro, M., Giannoni, M. P., & Patterson, C. (2023). The forward guidance puzzle. *Journal of Political Economy Macroeconomics*, 1(1), 43–79.

Del Negro, M., Giannoni, M. P., & Schorfheide, F. (2015). Inflation in the great recession and new keynesian models. *American Economic Journal: Macroeconomics*, 7(1), 168–196.

Del Negro, M., Lenza, M., Primiceri, G. E., & Tambalotti, A. (2020). *Whats up with the Phillips Curve?* Technical report, National Bureau of Economic Research.

Fernald, J. G., Spiegel, M. M., & Swanson, E. T. (2014). Monetary policy effectiveness in china: Evidence from a favar model. *Journal of International Money and Finance*, 49, 83–103.

Galı, J. & Gertler, M. (1999). Inflation dynamics: A structural econometric analysis. *Journal of monetary Economics*, 44(2), 195–222.

Gali, J., Gertler, M., & Lopez-Salido, J. D. (2005). Robustness of the estimates of the hybrid new keynesian phillips curve. *Journal of Monetary Economics*, 52(6), 1107–1118.

Gambacorta, L., Kwon, B., Park, T., Patelli, P., & Zhu, S. (2024). *CB-LMs: language models for central banking*. Bank for International Settlements, Monetary and Economic Department Basel .

Geiger, F., Kanelis, D., Lieberknecht, P., & Sola, D. (2025). *Monetary-Intelligent Language Agent (MILA)*. Technical report, Technical Paper.

Girardin, E., Lunven, S., & Ma, G. (2017). China's evolving monetary policy rule: from inflation-accommodating to anti-inflation policy.

Hansen, S. & McMahon, M. (2016). Shocking language: Understanding the macroeconomic effects of central bank communication. *Journal of International Economics*, 99, S114–S133.

Hansen, S., McMahon, M., & Prat, A. (2018). Transparency and deliberation within the fomc: A computational linguistics approach. *The Quarterly Journal of Economics*, 133(2), 801–870.

He, D. & Pauwels, L. L. (2008). What prompts the people's bank of china to change its monetary policy stance? evidence from a discrete choice model. *China & World Economy*, 16(6), 1–21.

Malmendier, U. & Nagel, S. (2016). Learning from inflation experiences. *The Quarterly Journal of Economics*, 131(1), 53–87.

Rotemberg, J. J. (1983). Aggregate consequences of fixed costs of price adjustment. *The American Economic Review*, 73(3), 433–436.

Sbordone, A. M. (2002). Prices and unit labor costs: a new test of price stickiness. *Journal of Monetary economics*, 49(2), 265–292.

Shapiro, A. H., Sudhof, M., & Wilson, D. J. (2022). Measuring news sentiment. *Journal of econometrics*, 228(2), 221–243.

Swanson, E. T. (2021). Measuring the effects of federal reserve forward guidance and asset purchases on financial markets. *Journal of Monetary Economics*, 118, 32–53.

Zhang, C. (2013). Inflation dynamics and an extended new keynesian phillips curve for china. *Emerging Markets Finance and Trade*, 49(5), 82–98.

Zhang, C. & Murasawa, Y. (2011). Output gap measurement and the new keynesian phillips curve for china. *Economic Modelling*, 28(6), 2462–2468.

Acknowledgment

The past we have personally experienced is the lens through which we view the economic future (*Malmendier & Nagel* (*2016*)). This idea deeply informs my own intellectual journey, as my formative years have unfolded during an era of profound transformation for the Chinese economy. From the US-China trade tensions starting around 2016 to the global pandemic in 2020, and the subsequent new rounds of tariff negotiations, I have grown up amidst constant change and uncertainty.

This macroeconomic turbulence was not just news headlines; it was the daily conversation at my family's dinner table. As business people in foreign trade, my parents' participation in the Canton Fair each year served as my personal window into the global economy. Since I was a child, I have always wanted to understand: Why were they sometimes full of worry, and other times able to close deals with ease? What invisible forces were driving the business cycle? These childhood questions, which had fascinated me for so long, finally began to find clues and answers in my high school studies and extracurricular explorations. I organized weekly economics seminars at school where we debated everything from inflation to interest rate policy, and through this process, I discovered a central question: how should the Chinese government manage inflation in our constantly evolving domestic context? Driven by this question and my past experience in Informatics Olympiads, I began to self-study machine learning and its research applications.

The mentorship of my high school advisor, Mr. Hualin Huang, was instrumental in transforming this burgeoning interest into a rigorous research project. Our collaboration began when he suggested using the IMF's WEO professional forecasts as a proxy for inflation expectations. However, I soon began to question if this single data series, dominated by international experts, could truly capture the nuances of China's unique, transition-era economy. Through my own literature review, I discovered that scholars studying China often utilized the PBoC's household and firm survey diffusion indices. This led me to an idea: since each of these proxies only reflects the views of a single group-forecasters, consumers, or firms-a more robust measure could be created by synthesizing them. When I proposed this to Mr. Huang, he supported this line of inquiry and suggested a sophisticated econometric technique: using a Kalman filter to esti-

mate the unobserved, latent expectation variable. I then embarked on a journey of self-study, learning the principles of state-space models and the Kalman filter through online courses, and successfully implemented it to construct my initial expectation index.

This index was a significant step forward, but this first stage of growth only led me to a deeper challenge. I grew increasingly aware of the inherent flaws of any survey-based data: they are noisy, subject to behavioral biases, and too infrequent to capture rapid economic shifts. This led me to a crucial insight, inspired by my background in the Informatics Olympiads: in an era of big data, the richest information often lies not in structured numbers, but in unstructured text. I hypothesized that the narrative of the PBOC's Monetary Policy Reports-a key instrument of policy guidance in China-must contain valuable signals. I therefore proposed and developed the novel solution that became the cornerstone of my paper: applying Large Language Models to systematically quantify this narrative. The execution was demanding, involving the meticulous preprocessing of hundreds of pages of text and carefully designing prompts to classify thousands of sentences without contextual leakage. This successful pivot, from a conventional method to an innovative one, marked the second and most creative step in my growth as a young explorer in this field.

To summarize the division of labor: the research topic originated from my personal background and academic interests. I was independently responsible for all stages of the project's execution, including data acquisition, processing, the implementation of both the Kalman filter and the Large Language Model analyses, and the writing of the initial manuscript. Mr. Huangs invaluable role was that of a mentor and critical evaluator. He provided the initial direction and suggested foundational techniques. Crucially, once the draft was complete, his rigorous feedback and critical evaluation challenged my arguments, pushed me to refine my logic, and helped me clarify the paper's core contributions.

Finally, I would like to thank the Yau Mathematical Sciences Center and the Yau Awards for providing this incredible opportunity to engage in scientific research and a platform to present my work. This experience has solidified my determination to become an outstanding researcher for my country and to contribute to scientific development.