Heterogeneous Expectations Across Inflation Regimes: Evidence and Implications for Monetary Policy.

Fernando Letelier and Facundo Luna

Rutgers University

December 2025

Abstract

This paper shows that the accuracy and behavior of U.S. inflation expectations depend critically on whether inflation is driven by demand or supply shocks. Combining one-year-ahead expectations from the SPF, Michigan Survey, and Cleveland Fed with Shapiro's (2024) decomposition, we find a reversal in forecast rankings: consumers forecast CPI inflation more accurately than experts in demand-driven episodes, while professional and market-based expectations dominate in supply-driven episodes. Forecast inefficiencies and error persistence are also regime-specific. A simple New Keynesian noisy-information framework with divine coincidence in demand regimes and its breakdown in supply regimes rationalizes these patterns and their policy implications. A state-dependent Taylor rule which conditions on the prevailing demand–supply mix can reduce welfare losses by around 20 percent, highlighting the monetary policy gains from treating expectations as regime-contingent rather than uniform.

1 Introduction

Inflation expectations are a cornerstone of monetary policy design, transmission, and communication. In the United States, three main sources provide complementary yet often divergent measures of expected inflation: the *Survey of Professional Forecasters* (SPF), the *University of Michigan Survey of Consumers*, and the *Cleveland Fed inflation expectations model*. Professional forecasts and market–model–based expectations are typically viewed as more accurate and more tightly linked to monetary policy, while household expectations are central for wage bargaining, price setting, and the anchoring of long-run beliefs. Recent inflation episodes, especially the post-pandemic period, have highlighted large and persistent discrepancies across these measures precisely when inflation itself reflected a mix of demand and supply shocks.

A large literature compares the level, bias, and accuracy of survey- and market-based inflation expectations.[1] Another branch of work studies information rigidities and rational inattention, typically through revision regressions and forecast error persistence (Coibion and Gorodnichenko, 2012,1; Maćkowiak and Wiederholt, 2009; Mankiw and Reis, 2002,0). These contributions generally treat inflation as a single stochastic process and evaluate expectations over the full sample, implicitly assuming that all agents face the same type of uncertainty. In contrast, New Keynesian models emphasize that demand and supply shocks differ in their observability, persistence, and policy implications, and that "divine coincidence" holds only when inflation is mainly demand-driven. This raises a fundamental question: *whose expectations forecast better, and under which type of inflation shock?*

This paper revisits the comparison of U.S. inflation expectations by explicitly conditioning forecast performance on the nature of inflation shocks. Using the decomposition of supply and demand contributions to inflation developed by Shapiro (2024), we classify each quarter as demand-driven or supply-driven and evaluate one-year-ahead forecasts from the SPF, the Michigan Survey of Consumers, and the Cleveland Fed model for both headline CPI and PCE

[1] See, among others, Ang et al. (2007), Faust and Wright (2013), and Reis (2023).

inflation. We combine standard forecast evaluation tools—Mincer–Zarnowitz regressions, bias and efficiency tests—with state-dependent tests of information rigidity and adaptive expectations that are estimated separately in demand and supply regimes.

Three main results emerge. First, forecast accuracy is strongly state-dependent and features a *reversal* in the ranking of agents across regimes. For CPI inflation, consumers forecast better than professional forecasters during demand-driven episodes (lower RMSE and MAE), while professional and market-based expectations forecast better than consumers during supply-driven episodes. For PCE inflation—the Fed's target index—all three measures systematically under-predict realized inflation by 0.3–0.7 percentage points on average, even though CPI forecasts show no significant bias. This persistent CPI–PCE wedge suggests that, in practice, both households and experts appear to think in "CPI terms" when forming expectations, with important implications for the communication of a PCE-based target.

Second, departures from rationality are themselves state-dependent. During demand-driven inflation, professional forecasters exhibit severe forecast inefficiency: forecast errors are large, highly persistent, and strongly predictable from the forecasts themselves, with revision and persistence coefficients in the range of 0.61–0.69. Consumers, by contrast, display smaller and less persistent forecast errors in these episodes, and in some specifications their CPI expectations cannot be statistically distinguished from forecast efficiency. In supply-driven episodes, the pattern reverses: error persistence and predictability largely disappear for experts, who behave close to the rational benchmark, while consumers' forecast errors remain more backward-looking and more strongly related to past inflation.

Third, the information set used to form expectations changes across regimes. Adaptive-expectations regressions show that consumers put substantial weight on lagged inflation in all periods, whereas professional forecasters place limited weight on past inflation and respond strongly to macro and policy variables—especially the Federal Funds rate—during supply-driven inflation. In supply regimes, policy signals have clear predictive content for experts and little additional explanatory power for consumers, consistent with experts understanding the trade-off the Federal Reserve faces when divine coincidence fails. In demand regimes, however, the mapping from policy to inflation becomes harder to interpret, and experts' forward-looking models perform poorly, while consumers' simple backward-looking rules perform relatively well.

Building on these empirical findings, we embed heterogeneous and state-dependent expectations into a simple New Keynesian model with demand and supply shocks and hybrid (backward- and forward-looking) expectations. Calibrating regime-specific information rigidity to the estimated coefficients, we compare a standard Taylor rule with a rule that is allowed to respond more aggressively to inflation in demand-driven episodes than in supply-driven ones. The optimal state-dependent rule implies roughly twice as strong a response to inflation when demand shocks dominate and reduces the welfare loss associated with inflation, output, and interest rate volatility by about 20% relative to the Taylor benchmark—more than twice the gain obtainable by merely re-optimizing a constant policy rule. These results show that treating expectations as regime-contingent, rather than uniform, has quantitatively meaningful implications for monetary policy design.

The paper contributes to three strands of the literature. First, it shows that the relative performance of consumer, professional, and market-based inflation expectations is not constant but depends systematically on whether inflation is driven by demand or supply shocks, a dimension largely ignored in existing forecast comparisons. Second, it provides new evidence of *state-dependent* information frictions and forecast inefficiencies: experts are most inaccurate and rigid precisely in demand-driven episodes, when divine coincidence would suggest they should be most informative. Third, it draws out the policy implications for how central banks should interpret and use heterogeneous expectations measures: which expectations to emphasize depends critically on the underlying shock mix, and the CPI–PCE wedge implies that communicating a PCE-based target may leave the public systematically perceiving higher inflation than the target metric.

2 Literature Review

2.1 Heterogeneous Expectations and Information Rigidity

A large literature documents that inflation expectations differ systematically across types of agents and that these differences reflect heterogeneous information frictions. Coibion and Gorodnichenko (2015) show that professional forecasters, firms, and households all display substantial information rigidity, but to very different degrees: professionals update frequently and have relatively small forecast errors, while households and firms are more inattentive and exhibit more persistent mistakes. Subsequent work using rich survey data reinforces this hierarchy. Weber et al. (2022) review micro evidence from household and firm surveys and conclude that households' expectations are more volatile, biased, and dispersed than those of professional forecasters, and that firms tend to lie in between. Cornand and Hubert (2022) and Link et al. (2023) further document how the frequency of updating and the degree of disagreement differ systematically across households, firms, professional forecasters and policymakers, confirming that agents operate under very different "information regimes." In unconditional comparisons, Verbrugge and Zaman (2021) and related work typically find that professional forecasters and, to some extent, firms predict inflation more accurately on average than households.

More recent contributions dig deeper into how different agents form expectations and which signals they react to. A growing empirical literature (e.g. D'Acunto et al., 2023; Weber et al., 2022) shows that households rely heavily on salient prices and personal experience, particularly for frequently purchased items, while firms and experts rely more on aggregate macro indicators and policy communication. Singh and Mitra (2022) use high-frequency identification around macroeconomic announcements and find that household expectations respond strongly to labor-market news but much less to inflation releases per se, suggesting that households primarily map demand-side conditions into beliefs about inflation. Han (2024) shows that households and professional forecasters associate inflation with very different macro narratives: households tend to view higher expected inflation as a signal of weaker future real activity, while professionals tend to associate higher inflation with stronger demand and growth. These findings are consistent with a view in which households rely on simple heuristics centered around local demand conditions and broad narratives, whereas experts process a richer set of macro and policy signals through formal models.

A key limitation of this literature is that most comparisons of forecasting performance across agents are unconditional: they ask "who forecasts best on average?" rather than "who forecasts better under which type of shock?". Existing evidence establishes that experts generally dominate households when pooling all periods, but does not systematically condition on whether inflation is primarily demand-driven or supply-driven. We contribute to this strand by classifying quarters into demand- and supply-driven inflation episodes and comparing the forecasting performance of consumers and experts across these regimes. Our results show that households can outperform experts when inflation is mostly driven by demand shocks, whereas experts retain a clear advantage when inflation is driven by supply disturbances—a finding that reconciles households' reliance on demand-side heuristics with their surprisingly strong performance in certain periods.

2.2 Demand versus Supply Shocks and Divine Coincidence

In parallel, a growing macro literature emphasizes the importance of distinguishing between demand-driven and supply-driven inflation episodes, and how this distinction interacts with expectations and monetary policy. Recent work decomposing the post-2019 inflation surge highlights the dominant role of broad-based supply shocks and their interaction with expectations in driving inflation dynamics, in contrast to earlier demand-dominated episodes (see, for example, Beaudry et al., 2025 and related contributions). In standard New Keynesian models, "divine coincidence" holds under pure demand shocks: stabilizing inflation closely aligns with stabilizing the output gap, making the central bank's reaction function relatively transparent (Blanchard and Galí, 2007). When supply (cost-push) shocks dominate, this coincidence breaks down and the central bank faces a non-trivial trade-off between inflation and real activity, rendering policy more state-dependent and harder to infer from simple rules of thumb.

This theoretical distinction has clear implications for expectation formation that remain underexplored. If households rely on simple heuristics that work well when divine coincidence holds—mapping demand conditions directly into inflation forecasts—their forecasts should perform relatively well during demand-driven episodes. Conversely, during supply-driven episodes where the inflation-output relationship is more complex, the richer models employed by professional forecasters should confer an advantage. We provide direct empirical evidence for this mechanism by documenting that the relative forecasting performance of consumers and experts flips across demand and supply regimes, consistent with the theoretical prediction that divine coincidence episodes favor simpler forecasting approaches.

2.3 State-Dependent Monetary Policy

The theoretical literature on state-dependent monetary policy formalizes how optimal policy coefficients should vary with the nature of shocks. The seminal work of Davig and Leeper (2007) generalizes the Taylor principle to environments where policy coefficients evolve according to a Markov process, showing that unique bounded equilibria can obtain even when policy temporarily deviates from the Taylor principle. Leeper and Davig (2006) extend this framework by making regime changes endogenous, with policy switching to a more aggressive stance when inflation exceeds a threshold. Empirical applications include Rabanal (2004), who documents that Federal Reserve behavior is well characterized by a two-state model: inflation targeting during expansions and output stabilization during recessions. More recently, Chang et al. (2021) develop methods for estimating DSGE models with threshold-type switching in the policy rule, where the transition between hawkish and dovish regimes depends on the historical impacts of structural shocks. Recent policy discussions at the Bank for International Settlements highlight that the appropriate monetary policy response depends critically on identifying the nature of prevailing shocks, with demand shocks calling for aggressive stabilization and supply shocks requiring a more measured approach (Hofmann et al., 2024).

A limitation of this literature is that regime switches are typically treated as exogenous or driven by latent factors, without connecting them to observable features of expectations. We contribute by using our empirical estimates of regime-specific information rigidity to discipline a New Keynesian model with hybrid expectations. This approach provides micro-foundations for state-dependent policy: the higher information rigidity we document during demand episodes implies that expectations adjust sluggishly, warranting more aggressive policy to anchor them. During supply episodes, near-rational expectations, combined with the inflation-output trade-off, imply that moderate policy is optimal. Our calibrated model yields an optimal policy ratio of $\phi_\pi^D/\phi_\pi^S \approx 2$, providing quantitative guidance for state-dependent policy design.

2.4 Imperfect Knowledge and Learning

A related strand of the literature examines optimal policy when agents possess imperfect knowledge about the economy. Orphanides and Williams (2005) and Orphanides and Williams (2007) show that policy rules calibrated under rational expectations can perform poorly when agents learn adaptively, as the learning process interacts with policy errors to generate persistent deviations of inflation expectations from target. Their findings suggest that more aggressive policy responses to inflation may be warranted when expectations are slow to adjust. This insight is particularly relevant for our analysis, as it provides theoretical support for why high information rigidity during demand episodes calls for aggressive policy.

Our contribution to this strand is to provide direct empirical evidence on how information rigidity varies across inflation regimes. While Orphanides and Williams model learning as a constant process, our estimates reveal that information rigidity is itself state-dependent: consumers exhibit high rigidity during demand episodes but lower rigidity during supply episodes, while the pattern reverses for professional forecasters. This heterogeneity in learning across regimes and agent types has not been documented previously and suggests that the welfare gains from state-dependent policy may be larger than models with constant learning parameters would imply.

3 Conceptual Framework

This section develops the economic intuition for why forecast performance may differ across inflation regimes. We draw on the New Keynesian distinction between demand and supply shocks, and the concept of "divine coincidence," to motivate our empirical analysis and the state-dependent policy framework.

3.1 Demand vs. Supply Shocks and Divine Coincidence

Consider a standard New Keynesian environment in which inflation is driven by demand and supply disturbances. The Phillips curve takes the form

$$\pi_t = \beta \mathbb{E}_t \pi_{t+1} + \kappa x_t + u_t, \tag{1}$$

where π_t is inflation, x_t is the output gap, and u_t is a cost-push (supply) shock. Monetary policy follows a Taylor rule

$$i_t = \phi_\pi \pi_t + \phi_x x_t, \tag{2}$$

where i_t is the nominal interest rate.

When shocks are predominantly demand-driven and u_t is small, stabilizing inflation is approximately equivalent to stabilizing the output gap—a property known as "divine coincidence" (Blanchard and Galí, 2007). In this case, inflation is tightly linked to aggregate demand conditions and slack. When shocks are predominantly supply-driven, u_t is large and divine coincidence breaks down: stabilizing inflation requires accepting output losses, and the central bank faces a genuine trade-off.

This distinction has important implications for expectation formation. During demand-driven inflation, the main forces behind price changes are variables that households directly observe—job prospects, wage growth, local labor market tightness, and spending patterns. During supply-driven episodes, inflation depends on sector-specific disturbances (energy prices, supply-chain disruptions) and on the central bank's policy response to the inflation-output trade-off. These factors are arguably better understood by professional forecasters with access to detailed sectoral data and macroeconomic models.

3.2 Inflation Decomposition

To connect this intuition to our empirical framework, we decompose realized inflation into demand and supply components:

$$\pi_t = \pi_t^D + \pi_t^S, \tag{3}$$

where π_t^D reflects demand-driven price pressures and π_t^S reflects supply-driven pressures. Following Shapiro (2024), we classify each quarter by its dominant inflation source.

This decomposition allows us to examine whether forecast accuracy and information rigidity vary systematically across regimes. If consumers have a comparative advantage in observing demand conditions while professionals have a comparative advantage in tracking supply disruptions and policy responses, we would expect:

1. Consumers to forecast more accurately during demand-driven episodes;
2. Professionals to forecast more accurately during supply-driven episodes;
3. Information rigidity to vary across regimes and forecaster types.

Table 19 documents the distribution of demand and supply regimes across decades. The 1990s exhibit the highest concentration of demand-driven inflation (57.5% of quarters), corresponding to the stable growth and tight labor markets of that period. In contrast, the

Table 1: Regime Distribution by Decade

Decade	Demand	Supply	Total	% Demand	Mean Inflation
1983-1989	5	27	32	15.6%	3.96%
1990-1999	23	17	40	57.5%	3.01%
2000-2009	9	31	40	22.5%	2.57%
2010-2019	18	22	40	45.0%	1.77%
2020-2024	10	9	19	52.6%	4.28%
Total	65	106	171	38.0%	2.94%

Notes: Demand periods defined as quarters where demand factors account for $\geq$ 50% of inflation per Shapiro (2024). The 1990s and 2020s show higher shares of demand-driven inflation, while the 1980s and 2000s are predominantly supply-driven.

1980s (15.6% demand) and 2000s (22.5% demand) were predominantly supply-driven, reflecting oil price shocks and commodity price fluctuations. The 2010s show a more balanced distribution (45% demand), while the pandemic period (2020-2024) is roughly evenly split. This temporal variation in regime frequency raises a potential concern: our findings could reflect time-period effects rather than regime-specific forecaster behavior. We address this concern in the Appendix by including period fixed effects and restricting the sample to more homogeneous subperiods.

3.3 Implications for Monetary Policy

If information rigidity is state-dependent, optimal monetary policy should also be state-dependent. When information rigidity is high, expectations adjust slowly to policy actions, and aggressive policy may be warranted to anchor inflation. When information rigidity is low, expectations respond quickly, and moderate policy suffices.

We formalize this intuition in Section 7, Optimal Monetary Policy with Heterogeneous Expectations, using a New Keynesian model with hybrid expectations, where the backward-looking share ω varies across regimes. The model allows us to derive optimal Taylor rule coefficients for each regime and quantify the welfare gains from state-dependent policy.

4 Data and Empirical Framework

4.1 Data Sources

The analysis combines three distinct measures of U.S. inflation expectations with realized inflation and an ex-post decomposition of inflation shocks into supply- and demand-driven components.

Survey of Professional Forecasters (SPF). The Survey of Professional Forecasters, conducted quarterly by the Federal Reserve Bank of Philadelphia, provides professional predictions for several macroeconomic variables, including CPI inflation. We focus on the one-year-ahead expected CPI inflation rate, denoted by $\pi_{t+4}^{e,SPF}$, where t indexes the survey quarter and the horizon is four quarters ahead.[2] The SPF serves as the benchmark for expert forecasts and reflects expectations of agents closely following macroeconomic developments and monetary policy.

University of Michigan Survey of Consumers. The University of Michigan Survey of Consumers reports household expectations for inflation over the next twelve months, based on monthly interviews of U.S. consumers. We use the median one-year-ahead expected inflation, aggregate it to the quarterly frequency by taking the average of monthly values within each quarter, and denote the resulting series by $\pi_{t+4}^{e,MICH}$. These expectations capture the perceptions of non-professional agents, which are likely influenced by salient prices such as food, energy, and housing.

Cleveland Fed Inflation Expectations. The Federal Reserve Bank of Cleveland publishes model-based measures of expected inflation inferred from nominal and real Treasury yields, inflation swaps, and survey data. We use the one-year-ahead expected CPI inflation, $\pi_{t+4}^{e,CLEV}$, available at a monthly frequency and aggregated to quarters by averaging monthly

[2] In practice, we use the SPF forecast of CPI inflation over the next four quarters and align it with year-over-year CPI inflation four quarters ahead.

observations. This series reflects inflation expectations implicit in financial markets, filtered through a term-structure model that also incorporates inflation risk premia. In some specifications, we also consider the Cleveland Fed measure of one-year-ahead PCE inflation to compare expectations for the Fed's target index.

Realized Inflation. Realized inflation is computed using quarterly CPI data from the Bureau of Labor Statistics (BLS). The year-over-year CPI inflation rate is defined as:

$$\pi_t^{\text{CPI}} = 100 \times \left(\frac{\text{CPI}_t}{\text{CPI}_{t-4}} - 1 \right),$$

For robustness, we also construct year-over-year PCE inflation using data from the Bureau of Economic Analysis (BEA). Unless otherwise noted, results are shown separately for CPI and PCE inflation.

Forecast errors for expectation source $i \in \{\text{SPF}, \text{MICH}, \text{CLEV}\}$ are defined as the difference between realized and expected inflation at the one-year horizon:

$$FE_t^i = \pi_{t+4}^{\text{actual}} - \pi_t^{e,i},$$

where $\pi_{t+4}^{\text{actual}}$ is either CPI or PCE inflation, depending on the specification. Because the forecast horizon is four quarters, forecast errors are overlapping; all inference below uses HAC / GMM procedures with appropriate lag lengths to account for serial correlation.

Inflation Shock Decomposition. To distinguish periods dominated by supply versus demand shocks, we use the decomposition developed by Shapiro (2024), which attributes changes in inflation to the contributions of demand and supply factors across disaggregated categories. Shapiro provides monthly contributions of demand-driven and supply-driven components to headline PCE inflation. Following his classification, we construct a quarterly indicator variable equal to one when demand components explain a majority of the change in inflation over the quarter and zero otherwise. Formally, we define:

$$S_t = \begin{cases} 1 & \text{if demand-driven components account for more than 50\% of inflation in quarter } t \\ 0 & \text{otherwise, i.e. inflation is supply-driven.} \end{cases}$$

This binary indicator S_t is the key object that allows us to evaluate forecast performance separately in demand- and supply-driven inflation regimes.

5 Empirical Framework

This section describes the regime classification and econometric tests used to evaluate inflation forecast accuracy and rationality across demand and supply inflation regimes.

5.1 Regime Classification

We classify each quarter as demand-driven or supply-driven based on the decomposition methodology of Shapiro (2024). This approach uses sectoral price and quantity data to identify the underlying source of inflation: demand shocks generate positive comovement between prices and quantities, while supply shocks generate negative comovement.

Let s_t^D denote the share of inflation attributable to demand factors. A quarter is classified as demand-driven if $s_t^D \geq 0.50$, and supply-driven otherwise:

$$\text{Regime}_t = \begin{cases} \text{Demand} & \text{if } s_t^D \geq 0.50 \\ \text{Supply} & \text{if } s_t^D < 0.50 \end{cases} \tag{4}$$

In Appendix A, we assess a more stringent definition whether both the headline and the core inflation are explained mostly by demand or supply factors.

5.2 Forecast Accuracy Metrics

We evaluate forecast accuracy using three standard metrics. The root mean squared error (RMSE) is:

$$\text{RMSE} = \sqrt{\frac{1}{T}\sum_{t=1}^{T} FE_t^2} \tag{5}$$

The mean absolute error (MAE) is:

$$\text{MAE} = \frac{1}{T}\sum_{t=1}^{T} |FE_t| \tag{6}$$

The bias is the mean forecast error:

$$\text{Bias} = \frac{1}{T}\sum_{t=1}^{T} FE_t \tag{7}$$

We compute these metrics separately for each forecaster group, inflation measure, and regime.

5.3 Tests of Forecast Rationality

Following Mankiw et al. (2003), we implement four tests of forecast rationality.

5.3.1 Panel A: Testing for Bias

Under rational expectations, forecast errors should have zero mean unconditionally. We estimate:

$$FE_t = \alpha + \varepsilon_t \tag{8}$$

and test $H_0 : \alpha = 0$. Rejection indicates systematic over- or under-prediction.

5.3.2 Panel B: Forecast Efficiency

Under full-information rational expectations (FIRE), forecasts should be orthogonal to forecast errors. We estimate:

$$FE_t = \alpha + \beta\, E_{t-4}[\pi_t] + \varepsilon_t \tag{9}$$

and test $H_0 : \alpha = \beta = 0$. If $\beta \neq 0$, the forecast contains information about its own error, indicating inefficient use of available information.

5.3.3 Panel C: Forecast Error Persistence

Under rational expectations, forecast errors should be serially uncorrelated. We estimate:

$$FE_t = \alpha + \beta\, FE_{t-4} + \varepsilon_t \tag{10}$$

and test $H_0 : \beta = 0$. A positive β indicates that past errors predict current errors, suggesting sluggish information updating.

5.3.4 Panel D: Exploitation of Macroeconomic Information

Under FIRE, publicly available macroeconomic data should not predict forecast errors. We estimate:

$$FE_t = \alpha + \beta\, E_{t-4}[\pi_t] + \gamma\, \pi_{t-1} + \kappa\, i_{t-1} + \delta\, u_{t-1} + \varepsilon_t \tag{11}$$

where π_{t-1} is lagged inflation, i_{t-1} is the federal funds rate, and u_{t-1} is the unemployment rate. We test $H_0 : \gamma = \kappa = \delta = 0$. Rejection indicates that forecasters fail to fully incorporate observable macroeconomic conditions.

5.4 Tests of Adaptive Expectations

Finally, we test whether expectations follow a purely adaptive process. Under adaptive expectations, forecasts depend only on past realizations of inflation. We estimate:

$$E_t[\pi_{t+4}] = \sum_{j=1}^{8} \beta_j \, \pi_{t-j} + \gamma_0 \, u_t + \gamma_1 \, u_{t-1} + \kappa_0 \, i_t + \kappa_1 \, i_{t-1} + \varepsilon_t \tag{12}$$

where the regression includes eight quarterly lags of inflation along with current and lagged unemployment and interest rates.

Under pure adaptive expectations, forecasts depend only on past inflation, implying $\gamma_0 = \gamma_1 = \kappa_0 = \kappa_1 = 0$ and $\sum_{j=1}^{8} \beta_j = 1$. We report the sum of inflation coefficients $\sum \beta$ and the joint test of whether the macroeconomic variables are significant.

5.5 Estimation

All regressions are estimated by OLS with robust standard errors clustered by year to account for serial correlation and heteroskedasticity. Each specification is estimated separately for (i) all periods, (ii) demand periods only, and (iii) supply periods only, allowing us to examine whether forecast properties vary across inflation regimes.

6 Results

We study how inflation expectations are formed and how this process depends on the underlying source of inflation. We organize the analysis around three questions. First, how does forecast accuracy vary across demand- and supply-driven inflation episodes? Second, do expectations satisfy standard rationality conditions, and does this differ by regime? Third, can expectations be described as purely backward-looking, or do agents use forward-looking information in a state-dependent way?

A recurring theme will be that our empirical patterns line up naturally with the New

Keynesian distinction between demand and supply shocks and the notion of "divine coincidence." When inflation is primarily driven by demand, stabilizing inflation is approximately equivalent to stabilizing the output gap; when inflation is driven by supply shocks, this coincidence breaks down and the central bank faces a genuine trade-off between inflation and real activity. We show that this distinction helps explain why consumers forecast better than experts in demand-driven episodes, while experts forecast better in supply-driven ones.

6.1 Summary Statistics and Regime Characteristics

Figures of inflation by shock and contribution illustrate that supply-driven inflation episodes are more frequent than demand-driven ones and show the contribution of each component to quarterly inflation.

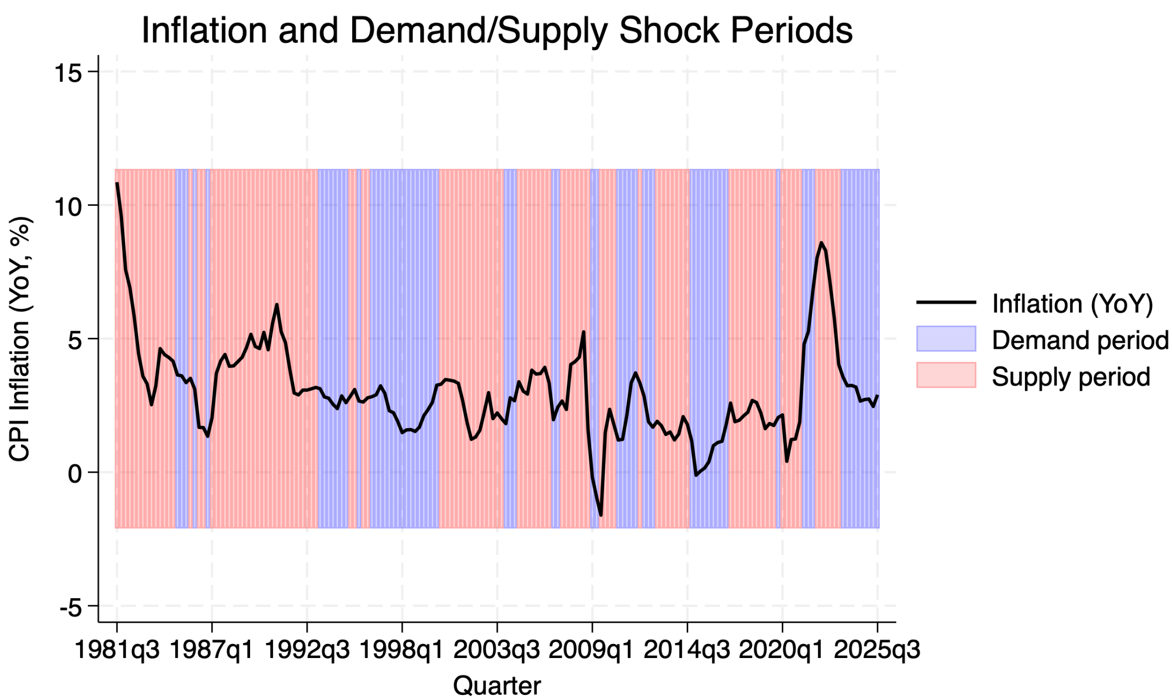

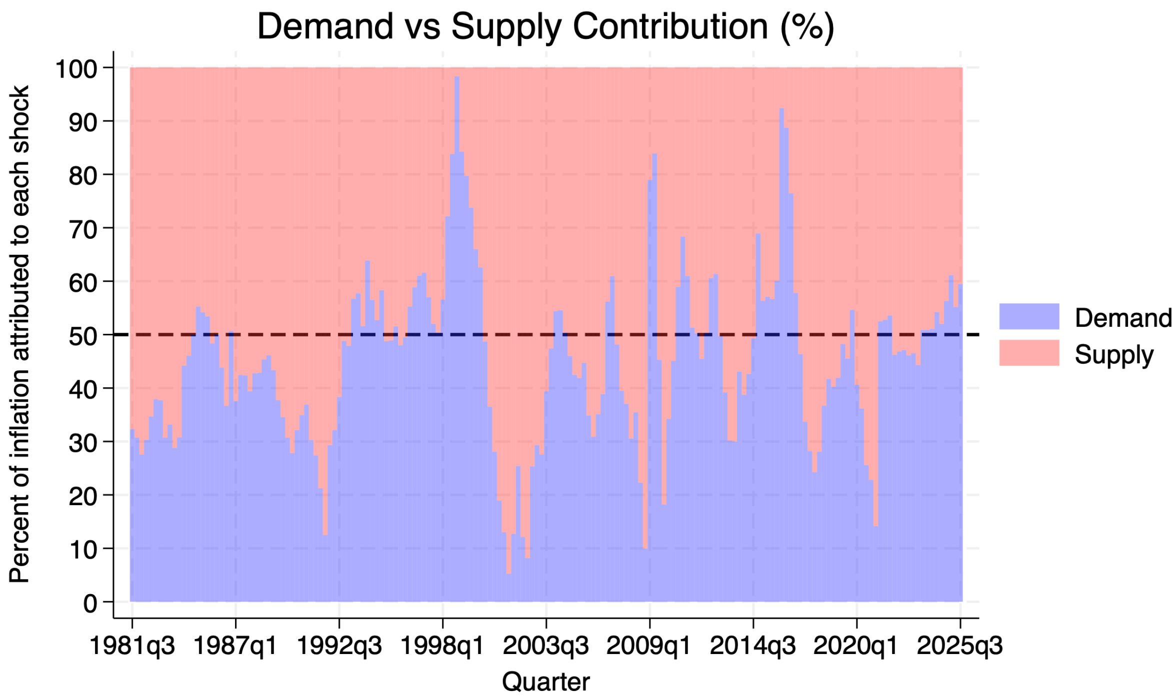

Table 2: Summary Statistics: Regime Comparison

	All Periods			Demand			Supply		
Variable	Mean	SD	N	Mean	SD	N	Mean	SD	N
CPI Inflation (YoY)	2.858	1.552	167	2.250	1.296	65	3.246	1.583	102
PCE Inflation (YoY)	2.457	1.343	167	1.887	1.152	65	2.820	1.334	102
Michigan Survey: Median Expected Inflation	3.118	0.621	167	2.978	0.505	65	3.207	0.672	102
SPF: Expected CPI Inflation (1 Year)	2.832	0.930	167	2.523	0.646	65	3.030	1.028	102
Cleveland Fed: 1-Year Expected Inflation	2.711	1.076	167	2.405	0.943	65	2.906	1.114	102
Federal Funds Effective Rate	3.735	3.014	167	3.117	2.667	65	4.129	3.165	102
Unemployment Rate	5.921	1.745	167	5.951	1.873	65	5.902	1.667	102
Output Gap	-0.662	1.920	167	-0.490	2.152	65	-0.772	1.759	102
Share of Inflation by demand	0.547	0.570	167	0.832	0.830	65	0.366	0.106	102

Notes: Demand periods: share of inflation from demand $\geq$ 50%

Table 2 reports summary statistics for the full sample and separately for demand- and supply-driven periods. A quarter is classified as demand-driven when demand factors explain at least 50% of observed inflation; this yields 65 demand quarters and 102 supply quarters over 1983q1–2024q3.

The differences across regimes are economically meaningful. Inflation is systematically higher in supply-driven periods: average CPI inflation is 3.25% in supply episodes versus 2.25% in demand episodes (a gap of 1.00 percentage points). PCE inflation shows a similar pattern (2.82% vs. 1.89%). This reflects the major inflation episodes in our sample early 1980s oil shocks and the post-pandemic supply-chain disruptions, which are largely supply-driven and large in magnitude. Demand-driven inflation episodes, in contrast, are less frequent and more moderate.

Inflation expectations move in the same direction but by much less. Consumer expectations (Michigan) average 3.21% in supply periods and 2.98% in demand periods, a difference of only 0.23 percentage points. Professional forecasts from the SPF rise from 2.52% (demand) to 3.03% (supply), and the Cleveland Fed measure from 2.41% to 2.91%. In all cases, the gap in expectations across regimes is smaller than the corresponding gap in realized inflation.

This foreshadows a central finding: agents understate the persistence and magnitude of supply-driven inflation relative to demand-driven inflation. We return to this below.

The macroeconomic environment also differs by regime. During demand-driven inflation, the Federal Funds rate is relatively low (3.12%) and unemployment moderate (5.95%). Many of these episodes occur during the Great Moderation, when inflation was low and stable and the Fed faced a relatively benign policy environment. In supply-driven episodes, the Federal Funds rate is higher (4.13%), unemployment is slightly lower (5.90%), and the output gap is more negative (-0.77% vs. -0.49%), consistent with stagflationary conditions that complicate the policy trade-off.

6.2 Forecast Accuracy: The State-Dependent Performance Puzzle

Table 3 reports three standard measures of forecast accuracy—root mean squared error (RMSE), mean absolute error (MAE), and average bias—for three expectations measures (Michigan, SPF, Cleveland) and two inflation concepts (CPI and PCE), for the full sample and separately by regime.

Table 3: Forecast Accuracy: All Periods, Demand, and Supply Inflation Periods

	Michigan		SPF		Cleveland Fed	
	CPI	PCE	CPI	PCE	CPI	PCE
Panel A: All Periods						
RMSE	1.525	1.452	1.466	1.312	1.591	1.416
MAE	1.119	1.114	0.980	0.984	1.048	1.012
Bias	-0.271	-0.702	0.015	-0.417	0.136	-0.295
N	167	167	167	167	167	167
Panel B: Demand Inflation Periods						
RMSE	1.320	1.289	1.651	1.440	1.809	1.595
MAE	1.019	1.087	1.127	1.097	1.182	1.153
Bias	-0.224	-0.726	0.232	-0.270	0.349	-0.153
N	65	65	65	65	65	65
Panel C: Supply Inflation Periods						
RMSE	1.643	1.546	1.335	1.223	1.434	1.289
MAE	1.182	1.131	0.886	0.912	0.963	0.922
Bias	-0.300	-0.687	-0.123	-0.510	0.001	-0.386
N	102	102	102	102	102	102

Notes: RMSE = Root Mean Squared Error; MAE = Mean Absolute Error; Bias = Mean Forecast Error. Forecast errors calculated as actual inflation minus 4-quarter-ahead forecast. Panel A includes all observations. Panel B restricts to demand inflation periods. Panel C restricts to supply inflation periods.

6.2.1 Overall Performance and the PCE Puzzle

Panel A (all periods) confirms standard patterns. For CPI inflation, professional forecasters slightly outperform consumers: SPF forecasts have an RMSE of 1.46%, compared with 1.53% for Michigan. Given average CPI inflation of 2.85%, this is a non-trivial difference. For PCE, SPF and Cleveland also dominate Michigan in terms of RMSE and MAE.

More striking is the pattern for PCE bias. All three measures under-predict PCE inflation on average: Michigan by 0.70 percentage points (1% significance), SPF by 0.41 points (5% significance), and Cleveland by 0.29 points (not significant). By contrast, CPI forecasts display no statistically significant bias; point estimates range from -0.27% to +0.14%, none different from zero.

This divergence between CPI and PCE forecasts is relevant for policy. The Federal Reserve targets PCE inflation, but both consumers and professional forecasters appear to think in CPI terms or some intermediate notion of "inflation." Over our sample, the average CPI–PCE wedge is about 0.4 percentage points (ranging from -1.69 in 1983q3 to 1.68

percentage points in 2022q2). This implies that when the Fed communicates about achieving a 2% PCE target, households and markets may perceive inflation as higher, potentially weakening the anchoring of expectations. This wedge is higher in supply periods (0.41) than in demand periods (0.37).

6.2.2 The State-Dependent Forecast Accuracy Reversal

Panels B and C reveal our first major result: forecast accuracy is strongly regime-dependent and behaves very differently for consumers and experts.

Consumers forecast better during demand-driven inflation. For CPI, Michigan's RMSE is 1.32% in demand periods and 1.64% in supply periods—a deterioration of 0.32 percentage points when moving from demand to supply. MAE similarly rises from 1.02% to 1.18%.

Professional forecasters show the opposite pattern. For CPI, SPF's RMSE is 1.65% in demand episodes but improves to 1.34% in supply episodes (a reduction of 0.32 percentage points). Cleveland's RMSE falls from 1.81% (demand) to 1.43% (supply), an improvement of 0.38 points.

In short, consumers forecast relatively well when inflation is demand-driven but substantially worse when inflation is supply-driven. Experts forecast relatively poorly in demand episodes but significantly better in supply episodes. This reversal—consumers outperforming experts in demand regimes and experts outperforming consumers in supply regimes—is at the heart of our results and is not predicted by standard expectation-formation models in which experts are uniformly more informed or more rational.

Bias patterns are consistent with this picture. During demand periods, professional CPI forecasts exhibit small positive biases (SPF: +0.23 pp; Cleveland: +0.35 pp, both insignificant), suggesting a tendency to overpredict inflation or to lean too heavily on long-run inflation targets and rapid mean reversion. During supply periods, professional CPI biases become slightly negative but remain small and insignificant, indicating under-prediction of very persistent supply shocks. For PCE, negative bias is present in both regimes and is especially pronounced for Michigan.

6.3 Tests of Forecast Rationality: State-Dependent Violations

We now turn to a set of standard rationality tests to assess whether inflation expectations behave consistently with the benchmark conditions from Nordhaus (1987) and Mankiw et al. (2003). The analysis focuses on four dimensions: systematic bias, forecast efficiency, error persistence, and the role of macroeconomic information. Table 4 summarizes the full-sample results.

6.3.1 Full-Sample Evidence

Panel A examines systematic bias by testing whether the average forecast error differs from zero. For CPI inflation, the estimates range from roughly -0.27 to 0.14 percentage points, and none of these estimates is statistically different from zero. For PCE inflation, however, we observe clear negative bias: forecast errors average -0.70 percentage points for Michigan (significant at the 1 percent level), -0.41 for the SPF (5 percent level), and a smaller, statistically insignificant -0.29 for Cleveland.

Panel B evaluates forecast efficiency by regressing forecast errors on the forecasts themselves. Under rational expectations, forecasts should be orthogonal to subsequent errors. The results show that professional forecasters fail this test: coefficients on the forecast are significantly negative for both SPF and Cleveland CPI, and similar patterns hold for PCE. Joint tests of $\alpha = \beta = 0$ also reject efficiency for professionals (p-values no higher than 0.078). In contrast, Michigan consumers do not show meaningful departures from efficiency for CPI, even though their RMSEs are somewhat larger.

Panel C turns to the persistence of forecast errors. Rational expectations imply that errors should not be serially correlated. Across all series, the estimated persistence coefficients fall between 0.08 and 0.32, and only one is significant at the 10 percent level. This suggests that agents do not systematically repeat the same mistakes over time.

Table 4: Tests of Forecast Rationality: All Periods

Panel A: Testing for Bias

	Michigan		SPF		Cleveland Fed	
	CPI	PCE	CPI	PCE	CPI	PCE
Constant	-0.271	-0.702***	0.015	-0.417**	0.136	-0.295
	(0.206)	(0.175)	(0.205)	(0.177)	(0.219)	(0.194)
N	167	167	167	167	167	167

Panel B: Is Information in the Forecast Fully Exploited?

	Michigan		SPF		Cleveland Fed	
	CPI	PCE	CPI	PCE	CPI	PCE
$E_{t-4}[\pi_t]$	-0.321	-0.372	-0.360**	-0.404***	-0.553***	-0.590***
	(0.444)	(0.369)	(0.171)	(0.133)	(0.175)	(0.142)
Joint test p-value	0.433	0.000	0.078	0.000	0.010	0.000
N	167	167	167	167	167	167

Panel C: Are Forecasting Errors Persistent?

	Michigan		SPF		Cleveland Fed	
	CPI	PCE	CPI	PCE	CPI	PCE
$\pi_{t-4} - E_{t-8}[\pi_{t-4}]$	0.088	0.162	0.177	0.279	0.213	0.323*
	(0.135)	(0.125)	(0.161)	(0.169)	(0.167)	(0.166)
N	167	167	167	167	167	167

Panel D: Are Macroeconomic Data Fully Exploited?

	Michigan		SPF		Cleveland Fed	
	CPI	PCE	CPI	PCE	CPI	PCE
$E_{t-4}[\pi_t]$	-0.580	-0.849*	-0.247	-0.377	-0.752**	-0.822***
	(0.631)	(0.493)	(0.323)	(0.276)	(0.330)	(0.266)
Inflation_{t-1}	0.129	0.315***	0.114	0.219*	0.217*	0.335***
	(0.124)	(0.108)	(0.111)	(0.117)	(0.116)	(0.114)
Fed Funds_{t-1}	0.092	0.041	-0.085	-0.093	0.008	-0.014
	(0.068)	(0.055)	(0.123)	(0.095)	(0.135)	(0.102)
$\text{Unemployment}_{t-1}$	0.119	0.133**	0.044	0.071	0.110	0.118*
	(0.079)	(0.057)	(0.086)	(0.061)	(0.087)	(0.065)
Joint test p-value	0.187	0.012	0.687	0.187	0.289	0.027
N	167	167	167	167	167	167

Notes: Robust standard errors clustered by year in parentheses. Dependent variable in Panels A-C: Forecast Error ($\pi_t - E_{t-4}\pi_t$). Panel B joint test: $\alpha = \beta = 0$. Panel D joint test: coefficients on inflation, Fed Funds, and unemployment = 0. Full sample. * p<0.10, ** p<0.05, *** p<0.01

Panel D investigates whether lagged macroeconomic conditions—specifically lagged inflation, the Federal Funds rate, and unemployment—help predict forecast errors once the forecast itself is included. Strong rationality implies no incremental explanatory power. The evidence is mixed: for professional CPI forecasts, we generally fail to reject the null of strong rationality, whereas for PCE and for consumer expectations these variables occasionally matter. On balance, macro indicators add little and their contribution is unstable across specifications.

Together, these tests show that departures from rationality are present but uneven: professionals fail efficiency tests despite lower RMSEs, consumers show some bias in PCE expectations, and macro fundamentals play a limited role in forecasting errors. The next subsection examines whether these departures vary systematically across inflation regimes.

6.3.2 Demand-Driven Inflation

We next examine whether forecast rationality differs during demand-driven inflation regimes. Table 5 reports the results for these episodes.

Panel A shows that CPI expectations remain broadly unbiased. SPF and Cleveland forecasts exhibit small, positive constants—roughly 0.23 and 0.35 percentage points—which are not statistically meaningful, while Michigan reports a modest negative bias of about 0.22 percentage points, also insignificant. The picture shifts for PCE inflation: all three sources display more pronounced negative bias, with Michigan's estimate of -0.73 percentage points significant at the 1 percent level and smaller, statistically insignificant values for SPF and Cleveland.

Panel B documents a sharp decline in forecast efficiency among experts. The coefficient on $E_t\pi_{t+4}$ becomes large, negative, and highly significant for professional CPI forecasts: -1.07 for SPF and -1.04 for Cleveland, both significant at the 1 percent level. These estimates are close to or beyond -1, indicating that when experts expect higher inflation during a demand regime, they tend to overpredict almost one-for-one. This is a strong violation of

Table 5: Tests of Forecast Rationality: Demand Inflation Periods

Panel A: Testing for Bias

	Michigan		SPF		Cleveland Fed	
	CPI	PCE	CPI	PCE	CPI	PCE
Constant	-0.224	-0.726***	0.232	-0.270	0.349	-0.153
	(0.258)	(0.213)	(0.339)	(0.297)	(0.363)	(0.329)
N	65	65	65	65	65	65

Panel B: Is Information in the Forecast Fully Exploited?

	Michigan		SPF		Cleveland Fed	
	CPI	PCE	CPI	PCE	CPI	PCE
$E_{t-4}[\pi_t]$	0.526	0.351	-1.068***	-1.033***	-1.039***	-1.047***
	(0.863)	(0.665)	(0.258)	(0.210)	(0.219)	(0.185)
Joint test p-value	0.254	0.001	0.001	0.000	0.000	0.000
N	65	65	65	65	65	65

Panel C: Are Forecasting Errors Persistent?

	Michigan		SPF		Cleveland Fed	
	CPI	PCE	CPI	PCE	CPI	PCE
$\pi_{t-4} - E_{t-8}[\pi_{t-4}]$	0.322	0.273	0.678**	0.691**	0.609**	0.643***
	(0.220)	(0.198)	(0.299)	(0.253)	(0.290)	(0.229)
N	65	65	65	65	65	65

Panel D: Are Macroeconomic Data Fully Exploited?

	Michigan		SPF		Cleveland Fed	
	CPI	PCE	CPI	PCE	CPI	PCE
$E_{t-4}[\pi_t]$	0.476	0.098	-1.235**	-1.294***	-0.914*	-1.066**
	(0.935)	(0.720)	(0.543)	(0.387)	(0.507)	(0.389)
Inflation$_{t-1}$	0.053	0.186	0.583*	0.643**	0.531	0.600**
	(0.243)	(0.197)	(0.341)	(0.288)	(0.324)	(0.281)
Fed Funds$_{t-1}$	-0.050	-0.060	-0.138	-0.096	-0.202	-0.128
	(0.094)	(0.077)	(0.230)	(0.172)	(0.279)	(0.212)
Unemployment$_{t-1}$	-0.077	-0.028	-0.011	0.006	-0.038	-0.009
	(0.155)	(0.121)	(0.170)	(0.138)	(0.184)	(0.153)
Joint test p-value	0.913	0.786	0.259	0.038	0.267	0.036
N	65	65	65	65	65	65

Notes: Robust standard errors clustered by year in parentheses. Dependent variable in Panels A-C: Forecast Error ($\pi_t - E_{t-4}\pi_t$). Panel B joint test: $\alpha = \beta = 0$. Panel D joint test: coefficients on inflation, Fed Funds, and unemployment = 0. Sample restricted to demand inflation periods. * p<0.10, ** p<0.05, *** p<0.01

forecast efficiency. Consumers, by contrast, show a positive but insignificant coefficient. Joint tests reinforce this pattern, with efficiency strongly rejected for both SPF and Cleveland (p-values of 0.001 and 0.000, respectively). Michigan CPI consumers, however, do not exhibit significant inefficiency ($p = 0.254$), though Michigan PCE expectations show rejection at the 1 percent level.

Panel C highlights a notable rise in error persistence for experts. Among SPF and Cleveland forecasts, coefficients fall between 0.61 and 0.69 for both CPI and PCE and are statistically significant. Consumers exhibit weaker and only marginally significant persistence. For experts, these values imply that two-thirds to three-quarters of last year's forecast error carries over into the current period during demand-driven episodes, consistent with slow adjustment or a mis-specified forecasting model.

Panel D examines the role of macroeconomic information. Lagged inflation, the Federal Funds rate, and unemployment generally offer little predictive power once the forecast itself is included. Most joint tests fail to reject the null of strong rationality, with the exceptions of SPF and Cleveland PCE forecasts, which show modest rejection at conventional levels. Taken together, the evidence suggests that the inefficiencies observed among experts stem less from a lack of information and more from a forecasting framework that misjudges the persistence of demand-driven inflation and the stance of monetary policy.

6.3.3 Supply-Driven Inflation: Near Rationality for Experts

Table 6 presents the results for periods classified as supply-driven.

Panel A shows that CPI forecast biases turn slightly negative in these episodes. SPF averages around –0.12 percentage points and Cleveland around 0.00, and both estimates are statistically insignificant. Michigan displays a somewhat larger negative bias of –0.30 percentage points, though still not significant. For PCE inflation, the familiar downward bias persists: Michigan again shows a sizeable and significant bias of –0.69 percentage points, while the SPF estimate stands at –0.51 percentage points (significant at the 1 percent level)

Table 6: Tests of Forecast Rationality: Supply Inflation Periods

Panel A: Testing for Bias

	Michigan		SPF		Cleveland Fed	
	CPI	PCE	CPI	PCE	CPI	PCE
Constant	-0.300	-0.687***	-0.123	-0.510***	0.001	-0.386**
	(0.251)	(0.220)	(0.179)	(0.154)	(0.191)	(0.169)
N	102	102	102	102	102	102

Panel B: Is Information in the Forecast Fully Exploited?

	Michigan		SPF		Cleveland Fed	
	CPI	PCE	CPI	PCE	CPI	PCE
$E_{t-4}[\pi_t]$	-0.624	-0.656*	-0.150	-0.240*	-0.321*	-0.392***
	(0.434)	(0.377)	(0.156)	(0.120)	(0.164)	(0.127)
Joint test p-value	0.334	0.014	0.321	0.000	0.149	0.000
N	102	102	102	102	102	102

Panel C: Are Forecasting Errors Persistent?

	Michigan		SPF		Cleveland Fed	
	CPI	PCE	CPI	PCE	CPI	PCE
$\pi_{t-4} - E_{t-8}[\pi_{t-4}]$	-0.033	0.099	0.004	0.097	0.058	0.171**
	(0.149)	(0.135)	(0.118)	(0.111)	(0.091)	(0.077)
N	102	102	102	102	102	102

Panel D: Are Macroeconomic Data Fully Exploited?

	Michigan		SPF		Cleveland Fed	
	CPI	PCE	CPI	PCE	CPI	PCE
$E_{t-4}[\pi_t]$	-1.163**	-1.287***	0.030	-0.033	-0.557**	-0.541**
	(0.514)	(0.423)	(0.298)	(0.276)	(0.270)	(0.227)
Inflation$_{t-1}$	0.252*	0.403***	0.009	0.060	0.130	0.204*
	(0.138)	(0.140)	(0.111)	(0.113)	(0.116)	(0.106)
Fed Funds$_{t-1}$	0.163**	0.094	-0.071	-0.104	0.045	-0.016
	(0.074)	(0.059)	(0.095)	(0.076)	(0.099)	(0.073)
Unemployment$_{t-1}$	0.142	0.142*	0.005	0.044	0.086	0.103
	(0.085)	(0.070)	(0.079)	(0.053)	(0.081)	(0.061)
Joint test p-value	0.008	0.001	0.862	0.363	0.522	0.111
N	102	102	102	102	102	102

Notes: Robust standard errors clustered by year in parentheses. Dependent variable in Panels A-C: Forecast Error ($\pi_t - E_{t-4}\pi_t$). Panel B joint test: $\alpha = \beta = 0$. Panel D joint test: coefficients on inflation, Fed Funds, and unemployment = 0. Sample restricted to supply inflation periods. * p<0.10, ** p<0.05, *** p<0.01

and Cleveland at –0.39 (significant at the 5 percent level). For experts, this pattern aligns with mild underestimation of the persistence associated with unusually large supply disturbances.

Panel B examines forecast efficiency and shows that most violations dissipate during supply episodes. The coefficients on $E_t\pi_{t+4}$ fall substantially in magnitude. For CPI, the estimates are –0.62 for Michigan, –0.15 for SPF, and –0.32 for Cleveland, with only the latter significant at conventional levels. The joint test does not reject efficiency for Michigan CPI expectations (p = 0.33), and for SPF and Cleveland the degree of inefficiency is far smaller than what we observe in demand-driven environments. PCE forecasts still display some deviations from efficiency, but these deviations are muted relative to the demand results.

Panel C shows that forecast errors exhibit very little serial dependence in supply periods. Estimated coefficients range from –0.03 to 0.17, and only one (Cleveland PCE, at 0.17) is statistically significant. This stands in marked contrast to the strong persistence documented for experts in demand episodes and suggests faster adjustment of beliefs when inflation movements originate on the supply side.

Panel D highlights a notable asymmetry in the relevance of macroeconomic information. For Michigan, lagged inflation, the unemployment rate, and the policy rate jointly help predict forecast errors during supply regimes, with joint-test p-values of 0.008 for CPI and 0.001 for PCE. Some of the individual coefficients on lagged inflation and unemployment are positive and significant. For professional forecasters, however, joint tests consistently fail to reject the null of strong rationality, with p-values between 0.11 and 0.86. Conditional on their forecasts, macro variables add no further explanatory power. This pattern is consistent with near-rational behavior by experts in supply-driven periods.

The forecast itself continues to matter for consumers: coefficients remain large and negative, around –1.16 for CPI and –1.29 for PCE. For SPF and Cleveland, in contrast, the forecast no longer predicts subsequent errors. Consumers retain some inefficiencies even when inflation is driven by supply shocks, while experts come close to the rational expectations benchmark in these periods.

6.4 Tests of Adaptive Expectations: State-Dependent Forecasting Models

We next ask whether expectations can be characterized as purely backward-looking. Table 7 reports regressions of expectations on eight lags of inflation, current and lagged unemployment, and current and lagged Fed Funds. Under pure adaptive expectations, (i) the sum of lagged inflation coefficients should equal one, and (ii) macro variables should not have predictive power conditional on lagged inflation.

We reject both implications and again find strong regime dependence in how agents combine past inflation with forward-looking information.

6.4.1 Full Sample: Forward-Looking Elements for All Agents

In the full sample (Panel A), the sum of lagged inflation coefficients ranges from 0.13 to 0.37, far below unity. Consumers place relatively high weight on past inflation (0.31 for CPI, 0.37 for PCE), while experts place much less weight (0.15–0.21), consistent with more forward-looking models.

Macroeconomic variables clearly matter. The current Fed Funds rate has positive and highly significant coefficients for SPF and Cleveland (0.35–0.38), indicating that periods with high policy rates—typically times of elevated current inflation—are also periods in which expected inflation one year ahead is high. Lagged Fed Funds has a negative coefficient for Michigan CPI (significant at 10%), consistent with the view that past tightening eventually lowers inflation, which consumers partially internalize. Unemployment also enters significantly in several specifications, typically with positive coefficients, reflecting persistence and co-movement of unemployment and inflation expectations over the sample.

Joint tests strongly reject pure adaptive expectations in almost all cases (p-values near zero).

Table 7: Tests of Adaptive Expectations Across Regimes

Panel A: All Periods

	Michigan		SPF		Cleveland Fed	
	CPI	PCE	CPI	PCE	CPI	PCE
$\sum \beta$ (inflation lags)	0.313***	0.373***	0.152**	0.214***	0.135**	0.161**
	(0.041)	(0.043)	(0.064)	(0.075)	(0.058)	(0.068)
Unemployment$_t$	0.006	0.005	0.054*	0.053*	-0.089***	-0.091***
	(0.034)	(0.028)	(0.029)	(0.028)	(0.018)	(0.019)
Unemployment$_{t-1}$	0.052**	0.034*	0.083**	0.068**	0.186***	0.173***
	(0.025)	(0.019)	(0.034)	(0.031)	(0.022)	(0.021)
Fed Funds$_t$	0.206	0.141	0.383***	0.355***	0.384***	0.361***
	(0.130)	(0.120)	(0.107)	(0.105)	(0.103)	(0.106)
Fed Funds$_{t-1}$	-0.236*	-0.176	-0.150	-0.133	-0.096	-0.075
	(0.120)	(0.111)	(0.094)	(0.096)	(0.104)	(0.109)
Reject adaptive (p-value)	0.004	0.009	0.000	0.000	0.000	0.000
N	159	159	159	159	159	159

Panel B: Demand Inflation Periods

	Michigan		SPF		Cleveland Fed	
	CPI	PCE	CPI	PCE	CPI	PCE
$\sum \beta$ (inflation lags)	0.358***	0.353***	0.134	0.152	0.097	0.116*
	(0.071)	(0.070)	(0.086)	(0.092)	(0.063)	(0.066)
Unemployment$_t$	-0.047	-0.043	0.035	0.049	-0.117**	-0.113**
	(0.041)	(0.039)	(0.049)	(0.045)	(0.044)	(0.043)
Unemployment$_{t-1}$	0.059***	0.040**	0.055	0.055	0.169***	0.175***
	(0.017)	(0.016)	(0.043)	(0.041)	(0.036)	(0.040)
Fed Funds$_t$	-0.019	-0.002	0.304	0.340	0.444	0.435
	(0.171)	(0.158)	(0.333)	(0.317)	(0.322)	(0.330)
Fed Funds$_{t-1}$	-0.085	-0.074	-0.122	-0.159	-0.149	-0.134
	(0.155)	(0.145)	(0.331)	(0.318)	(0.326)	(0.336)
Reject adaptive (p-value)	0.000	0.000	0.000	0.000	0.000	0.000
N	65	65	65	65	65	65

Panel C: Supply Inflation Periods

	Michigan		SPF		Cleveland Fed	
	CPI	PCE	CPI	PCE	CPI	PCE
$\sum \beta$ (inflation lags)	0.330***	0.441***	0.155	0.228**	0.148	0.174
	(0.070)	(0.083)	(0.092)	(0.110)	(0.097)	(0.103)
Unemployment$_t$	0.017	0.048	-0.089	-0.064	0.275	0.314*
	(0.210)	(0.200)	(0.182)	(0.168)	(0.179)	(0.165)
Unemployment$_{t-1}$	0.037	-0.027	0.265	0.218	-0.133	-0.186
	(0.208)	(0.197)	(0.178)	(0.154)	(0.151)	(0.135)
Fed Funds$_t$	0.276**	0.199*	0.390***	0.364***	0.508***	0.475***
	(0.129)	(0.111)	(0.111)	(0.111)	(0.109)	(0.113)
Fed Funds$_{t-1}$	-0.276**	-0.225**	-0.131	-0.122	-0.228**	-0.204*
	(0.118)	(0.098)	(0.099)	(0.098)	(0.109)	(0.110)
Reject adaptive (p-value)	0.161	0.046	0.000	0.000	0.000	0.000
N	94	94	94	94	94	94

Notes: Dependent variable: Expected Inflation $E_{t-4}\pi_t$. Regressions include 8 quarterly inflation lags (not shown). Robust standard errors clustered by year in parentheses. Joint test: all macro variables = 0. Panel A: Full sample. Panel B: Demand inflation periods (demand_side==1 or demand_both==1). Panel C: Supply inflation periods (demand_side==0 or demand_both==0). * p<0.10, ** p<0.05, *** p<0.01

6.4.2 Demand-Driven Inflation: Forward-Looking but Miscalibrated

Panel B shows that during demand episodes, experts rely even less on past inflation. For SPF and Cleveland, the sum of lagged inflation coefficients falls and becomes insignificant, while for Michigan it remains high and significant. Consumers thus remain predominantly backward-looking, whereas experts attempt to be more forward-looking.

However, macro variables become harder to interpret. Coefficients on current Fed Funds are large and positive for experts (0.30–0.44) but imprecisely estimated (standard errors almost as large as the coefficients), and unemployment loses significance. Agents seem to know policy matters for future inflation but face genuine uncertainty about how to map the policy stance and real activity into inflation during demand episodes. Despite these imprecise estimates, joint tests still reject pure adaptive expectations (p-values of 0.000 for all specifications).

6.4.3 Supply-Driven Inflation: Clearer Signals and Faster Learning

Panel C shows that during supply-driven inflation, consumers become even more backward-looking, especially for PCE (sum of lagged inflation coefficients around 0.42). This is intuitive: consumers observe large relative price increases (e.g. energy, goods) and extrapolate them without distinguishing temporary supply shocks from persistent inflation.

Experts, in contrast, keep low weights on past inflation and place greater emphasis on policy signals. The Fed Funds rate regains strong predictive power for SPF and Cleveland: current Fed Funds coefficients are around 0.36–0.51 and highly significant (1%). This indicates that experts interpret the policy stance as a reliable signal about the future path of inflation in supply episodes. Unemployment also helps explain SPF expectations in supply periods, with positive and sometimes significant coefficients, reflecting recognition of stagflationary dynamics.

Across all agents and regimes, adaptive expectations are rejected, but the mix of backward- and forward-looking information is clearly state dependent.

6.5 Economic Interpretation: Demand vs. Supply, Divine Coincidence, and Who Forecasts Better

Our state-dependent findings can be interpreted naturally through the New Keynesian distinction between demand and supply shocks and the notion of "divine coincidence."

In a standard New Keynesian model with predominantly demand shocks and relatively small cost-push disturbances, stabilizing inflation is approximately equivalent to stabilizing the output gap. Monetary policy can then pursue a single objective—closing the output gap—while also keeping inflation near target. In this environment, the key drivers of inflation are aggregate demand conditions and slack, which are closely related to variables that households directly observe: job-finding prospects, wage growth, the tightness of local labor markets, and the general strength of spending. Simple heuristics based on these signals ("when jobs and wages are strong, prices will keep rising") can yield reasonably accurate inflation forecasts.

This theoretical structure is consistent with our finding that consumers forecast better than experts in demand-driven inflation episodes. Our results show that during demand regimes, Michigan's CPI forecast errors are smaller than SPF's (RMSE of 1.32% vs. 1.65%), and consumer errors are less persistent. At the same time, our rationality tests reveal that experts make large, highly persistent errors in demand periods: forecast efficiency coefficients are close to or beyond -1 and forecast error persistence coefficients around 0.61–0.69. One interpretation is that professional forecasters, relying on structural models and an assumed policy reaction function, tend to overestimate the speed with which the Federal Reserve will bring inflation back to target and underappreciate the persistence of demand-driven price pressures. Households, by contrast, put more weight on recent inflation and local demand conditions, which are precisely the relevant state variables when divine coincidence holds approximately. As a result, their simpler, more backward-looking expectations can outperform sophisticated models in these episodes.

When inflation is driven by supply shocks, divine coincidence breaks down. Stabilizing inflation now implies a deterioration in the output gap, and the central bank faces a genuine trade-off between inflation and real activity. Inflation dynamics in these regimes depend not only on the size and persistence of supply disturbances (oil price movements, supply-chain bottlenecks, relative price shocks) but also on how monetary policy chooses to navigate this trade-off and on how expectations respond to both shocks and policy. Understanding inflation in this environment requires processing information about sectoral shocks, the slope of the Phillips curve, long-run inflation targets, and the credibility of policy commitments. These are precisely the dimensions on which professional forecasters have an advantage.

Our supply-regime results match this interpretation. During supply-driven episodes, experts forecast better than consumers (CPI RMSE of 1.34% for SPF vs. 1.64% for Michigan), their forecast errors exhibit little persistence, and lagged macro variables provide essentially no additional predictive power conditional on the forecast. This pattern is consistent with near-rational expectations for experts in supply regimes. Consumers, meanwhile, become more backward-looking and rely heavily on recent inflation, leading them to over-extrapolate large relative price shocks. In other words, when divine coincidence fails and inflation reflects complex policy trade-offs in response to supply disturbances, the advantage shifts from households—who see demand conditions but cannot infer the policy–inflation mapping—to experts, who are trained precisely to interpret these trade-offs.

Finally, the systematic underprediction of PCE inflation in all regimes suggests an additional layer of complexity: while the Fed targets PCE, households and many forecasters effectively think in CPI terms. This stable CPI–PCE wedge, on top of regime-dependent forecasting performance, implies that the Fed faces both a measurement/communication challenge (explaining its target) and a regime-dependent expectations challenge (different groups forecast better or worse depending on whether inflation is demand- or supply-driven). Both aspects are likely to matter for the design of monetary policy and forward guidance.

6.6 Real-Time Regime Classification

A potential concern with our analysis is that the Shapiro (2024) decomposition uses ex-post data, while forecasters must form expectations without knowledge of the contemporaneous regime. This creates an identification challenge: if regime classification is only possible with hindsight, our findings may not have clear implications for real-time forecasting or policy implementation.

We address this concern by constructing several real-time regime indicators that use only information available at the forecast date. Our preferred measure uses the one-quarter lagged demand share: a quarter is classified as demand-driven if $L1.\text{share\_demand} \geq 0.50$. We also consider more conservative alternatives: (i) the four-quarter lagged share, (ii) a four-quarter moving average of lagged shares, and (iii) a probit model predicting regime using lagged inflation, oil price changes, unemployment, and output gap.

Table 15 reports agreement rates between real-time and ex-post classification. The one-quarter lag method achieves 85.2% overall agreement, with higher accuracy for supply regimes (89.5%) than demand regimes (75.7%). The asymmetry reflects the greater persistence of supply-driven episodes, which makes them easier to identify in real time. Even the most conservative methods achieve 70–75% agreement, suggesting that regime uncertainty, while present, is manageable.

Tables 16–18 re-estimate our main specifications using real-time classification. All key findings are robust:

- *Forecast accuracy reversal*: Michigan outperforms SPF in demand regimes (RMSE: 1.28 vs. 1.56) while SPF outperforms Michigan in supply regimes (RMSE: 1.40 vs. 1.66). The ranking reversal is preserved under real-time classification.

- *Forecast efficiency*: SPF and Cleveland Fed exhibit severe inefficiency in demand regimes ($\beta \approx -0.9$ to -1.0, statistically significant) and near-efficiency in supply regimes ($\beta \approx -0.2$, often insignificant). This pattern is robust to real-time classification.

- *Error persistence*: High persistence for professional forecasters in demand regimes (0.49–0.55) and near-zero persistence in supply regimes remains evident with real-time indicators.

Figure 1 visualizes these results, showing that the efficiency coefficient pattern is nearly identical across ex-post and real-time classification methods. The robustness of our findings to real-time regime identification strengthens the case for state-dependent monetary policy, as central banks could feasibly implement such policies using lagged indicators.

Figure 1: Comparison of Regime Classification

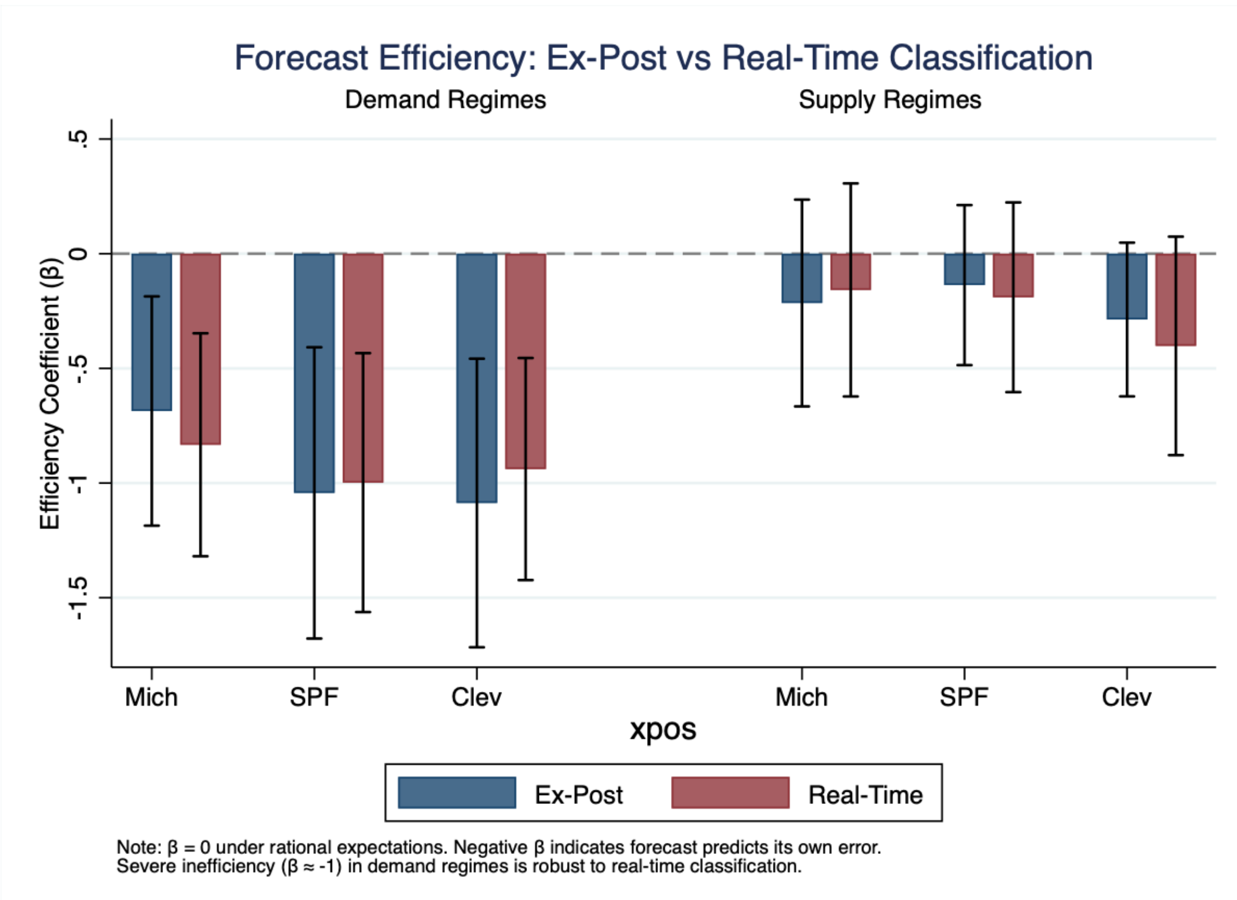

6.7 Robustness and Limitations

Several caveats apply to our analysis. First, our regime classification relies on an ex-post decomposition of inflation into demand and supply components. Forecasters in real time face uncertainty about the prevailing regime, so part of the forecast errors we document may reflect this regime uncertainty rather than conditional mistakes within a known regime.

Second, we focus on one-year-ahead forecasts, as this is the common horizon across our surveys. Expectations at other horizons may behave differently, particularly with respect to anchoring and sensitivity to policy.

Third, we use revised inflation data as the benchmark for forecast errors, whereas forecasters rely on real-time data that are subsequently revised. Some forecast errors thus reflect data revisions rather than informational or modeling failures.

Fourth, our sample spans heterogeneous historical episodes—Volcker disinflation, the Great Moderation, the Global Financial Crisis, and the post-pandemic period—that differ in credibility, financial structure, and openness. While our demand/supply classification captures one important dimension, other regime changes may also matter for expectations formation.

Finally, the three expectations measures differ in wording, respondent composition, and elicitation method. Michigan asks consumers about "prices in general," the SPF asks professional forecasters about specific indices, and the Cleveland Fed measure is inferred from financial markets. Differences across series may therefore reflect survey design as well as genuine differences in forecasting behavior.

Despite these caveats, the core patterns are robust: forecast accuracy, bias, efficiency, and the mix of backward- and forward-looking components all vary systematically across inflation regimes, and they do so in ways that differ sharply between consumers and experts. This regime dependence has been largely overlooked in the expectations literature and has important implications for macroeconomic modeling and the design of monetary policy and communication strategies.

7 Optimal Monetary Policy with Heterogeneous Expectations

Our empirical analysis reveals systematic differences in forecast accuracy and information rigidity across inflation regimes. During demand-driven episodes, consumers exhibit high information rigidity yet outperform professional forecasters. During supply-driven episodes, professionals are near-rational and forecast more accurately. This section explores the implications for optimal monetary policy using a New Keynesian model calibrated to match these empirical findings.

7.1 A New Keynesian Model with Hybrid Expectations

We extend the standard three-equation New Keynesian framework to incorporate *hybrid expectations*, where a fraction of agents form expectations adaptively while the remainder hold rational expectations. This structure captures the information rigidity documented in our empirical analysis.

Hybrid Expectations. Let $\omega \in [0, 1]$ denote the share of backward-looking agents. Aggregate expectations are formed as:

$$\pi_t^e = \omega \cdot \pi_{t-1} + (1 - \omega) \cdot \mathbb{E}_t[\pi_{t+1}] \tag{13}$$

$$y_t^e = \omega \cdot y_{t-1} + (1 - \omega) \cdot \mathbb{E}_t[y_{t+1}] \tag{14}$$

The parameter ω serves as our model analog to the empirical information rigidity coefficient β estimated in Section 6, Results. Higher ω implies slower expectation adjustment, consistent with the elevated β coefficients we document during demand-driven episodes.

Model Equations. The IS curve, derived from household optimization, relates the output gap to expected future output and the real interest rate:

$$y_t = y_t^e - \frac{1}{\sigma}\left(i_t - \pi_t^e - r_t^{nat}\right) \tag{15}$$

where y_t is the output gap, i_t is the nominal interest rate, $r_t^{nat} = \sigma(1-\rho_g)g_t$ is the natural real rate, and g_t is a demand shock.

The New Keynesian Phillips Curve, derived from Calvo pricing, links inflation to expected inflation and the output gap:

$$\pi_t = \beta\pi_t^e + \kappa y_t + u_t \tag{16}$$

where κ is the slope of the Phillips curve and u_t is a cost-push shock.

The central bank follows a Taylor-type rule:

$$i_t = \phi_\pi \pi_t + \phi_y y_t \tag{17}$$

Demand and cost-push shocks follow AR(1) processes with persistence parameters ρ_g and ρ_u, respectively.

7.2 State-Dependent Policy

Our empirical analysis suggests that both the nature of shocks and the degree of information rigidity differ systematically across regimes. We therefore consider a state-dependent Taylor rule:

$$i_t = \phi_\pi^{S_t}\pi_t + \phi_y y_t \tag{18}$$

where $S_t \in \{D, S\}$ denotes the prevailing regime (demand or supply), and the central bank may choose different inflation response coefficients ϕ_π^D and ϕ_π^S for each regime.

Two forces suggest that optimal policy should be more aggressive during demand-driven episodes ($\phi_\pi^D > \phi_\pi^S$). First, when demand shocks dominate, *divine coincidence* approximately holds: stabilizing inflation also stabilizes the output gap. As shown in the left panel of Figure 3, both Var(π) and Var(y) decline as ϕ_π increases in the demand regime. In contrast, the right panel shows that during supply-driven episodes, higher ϕ_π reduces inflation variance but *increases* output variance—the classic inflation-output trade-off induced by cost-push shocks. Second, high information rigidity during demand episodes ($\omega^D = 0.30$ vs. $\omega^S = 0.05$) means expectations adjust slowly to policy actions. More aggressive policy compensates by anchoring expectations faster, reducing the persistence of inflation deviations. This mechanism, emphasized by Orphanides and Williams (2007) in models with learning, implies that the welfare cost of passive policy is amplified when agents are slow to update their beliefs.

7.3 Calibration

Table 9 summarizes the model calibration. We discipline the regime-specific parameters using our empirical estimates.

Table 8: Model Calibration

Parameter	Description	Value	Source
Panel A: Structural Parameters			
β	Discount factor	0.99	Standard
σ	Inverse IES	1.0	Log utility
κ	NKPC slope	0.10	Galı and Gertler (1999)
ϕ_y	Output gap response	0.125	Taylor (1993)
ρ_g	Demand shock persistence	0.80	Smets and Wouters (2007)
ρ_u	Cost-push shock persistence	0.50	Smets and Wouters (2007)
Panel B: Demand Regime			
ω^D	Information rigidity	0.30	Table 5: Avg β from Consumers
σ_g^D	Demand shock std. dev.	0.010	Demand shocks dominate
σ_u^D	Cost-push shock std. dev.	0.005	Cost-push shocks present
Panel C: Supply Regime			
ω^S	Information rigidity	0.05	Table 6: Avg β from SPF
σ_g^S	Demand shock std. dev.	0.005	Demand shocks subdued
σ_u^S	Cost-push shock std. dev.	0.010	Cost-push shocks dominate
Panel D: Regime Frequencies (Table 2)			
$\Pr(S_t = D)$	Demand regime frequency	0.389	65 of 167 quarters
$\Pr(S_t = S)$	Supply regime frequency	0.611	102 of 167 quarters
Panel E: Welfare Weights			
λ_π	Weight on $\text{Var}(\pi)$	1.00	Normalized
λ_y	Weight on $\text{Var}(y)$	0.50	Woodford (2003)
λ_i	Weight on $\text{Var}(i)$	0.50	Rotemberg and Woodford (1997)

Notes: Information rigidity ω captures the share of backward-looking agents and is calibrated to match our empirical information rigidity coefficients. Shock variances are chosen so that demand shocks dominate during demand-driven episodes and cost-push shocks dominate during supply-driven episodes, consistent with our regime classification. Welfare weights on interest rate volatility follow Woodford (2003) and Orphanides and Williams (2007).

The welfare loss function, following Woodford (2003) and Rotemberg and Woodford (1997), includes interest rate volatility:

$$\mathcal{L} = \frac{1}{2}\left(\lambda_\pi \mathrm{Var}(\pi) + \lambda_y \mathrm{Var}(y) + \lambda_i \mathrm{Var}(i)\right) \tag{19}$$

The interest rate term captures welfare costs from transaction frictions and financial market disruptions associated with volatile policy rates. Including this term is particularly important for our analysis, as it penalizes excessively aggressive policy responses, incorporating financial stability concerns in the Central Bank Loss Function.

7.4 Theoretical Foundation for $\omega \approx \beta$.

Our calibration maps the empirical forecast error persistence coefficient β to the backward-looking share ω in the hybrid Phillips curve. This mapping follows from sticky-information models (Carroll, 2003; Mankiw and Reis, 2002). In these models, a fraction $(1-\lambda)$ of agents update their information set each period, while the remainder continue using outdated forecasts. Coibion and Gorodnichenko (2015) show that under sticky information, the forecast error persistence coefficient equals precisely the non-updating share: $\beta = (1 - \lambda)$.

In our hybrid expectations framework, ω represents the weight on backward-looking (adaptive) expectations, which is conceptually equivalent to the share of agents who have not updated to forward-looking, model-consistent expectations. The mapping $\omega \approx \beta$ thus has rigorous micro-foundations: both parameters capture the degree to which expectations deviate from full-information rational expectations due to information frictions.

To assess sensitivity, Table 9 reports results under four alternative calibration strategies:

1. *Baseline*: Use Michigan (consumer) average β for demand regimes and SPF average β for supply regimes, reflecting the best forecaster in each regime.

2. *Simple average*: Equal-weighted average of β across all three forecasters.

3. *Lag Regime*: We use the definition of the regime period as the previous quarter period

4. *Regime-specific best*: Same as baseline, using the most accurate forecaster in each regime.

All four calibrations preserve our central finding that $\omega^D > \omega^S$—the backward-looking share is substantially higher in demand regimes. The ratio ω^D/ω^S ranges from 5.95 to 16.32 across specifications. This robustness reflects the consistent empirical pattern: forecast error persistence is high for all forecasters in demand regimes and low in supply regimes. The qualitative policy implication—that stronger inflation responses are optimal in demand regimes—is thus insensitive to reasonable calibration choices.

Table 9: Alternative Calibrations for Information Rigidity (ω)

Calibration	ω^D	ω^S	Ratio	Rationale
1. Baseline	0.2975	0.05	5.95	Consumer (D), SPF (S)
2. Simple Average	0.536	0.033	16.32	Equal weight all forecasters
3. Lag Regime	0.244	0.024	10.16	Alternative Definition
4. Best Forecaster	0.322	0.054	5.96	Regime-specific best
Underlying Information Rigidity Coefficients (β):				
Michigan	0.322	-0.014		
SPF	0.678	0.054		
Cleveland Fed	0.609	0.058		

Notes: ω is the backward-looking share in the hybrid expectations model. β is the forecast error persistence coefficient from Table 4 Panel C. The mapping $\omega \approx \beta$ follows from sticky-information models where β represents the fraction of agents who have not updated their information set. All calibrations preserve the key finding that $\omega^D > \omega^S$.

7.5 Results

We solve the model using Dynare and compute welfare under alternative policy specifications. For each regime, we search over a grid of ϕ_π values to identify the optimal policy response. Total welfare is computed as the frequency-weighted average of regime-specific welfare losses.

Figure 2 displays welfare loss as a function of ϕ_π for each regime. During demand-driven episodes (left panel), the optimal response coefficient is $\phi_\pi^D = 3.06$, substantially above the

Taylor benchmark of 1.5. During supply-driven episodes (right panel), optimal policy is more moderate at $\phi_\pi^S = 1.26$, below the Taylor benchmark.

Figure 2: Welfare Loss by Policy Aggressiveness and Regime

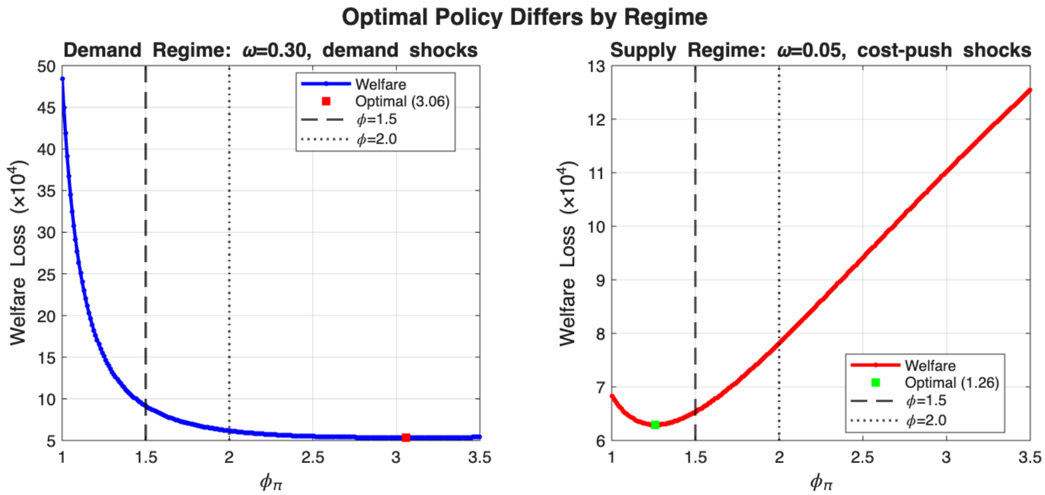

Notes: Left panel shows welfare loss in the demand regime ($\omega = 0.30$, demand shocks dominate). Right panel shows welfare loss in the supply regime ($\omega = 0.05$, cost-push shocks dominate). Dashed vertical line indicates Taylor benchmark ($\phi_\pi = 1.5$). Red/green squares mark optimal policy for each regime.

Figure 3 illustrates the mechanism underlying these results. In the demand regime (left panel), both inflation variance and output variance decline as ϕ_π increases—divine coincidence holds, making aggressive policy unambiguously beneficial. In the supply regime (right panel), higher ϕ_π reduces inflation variance but increases output variance, reflecting the policy trade-off induced by cost-push shocks.

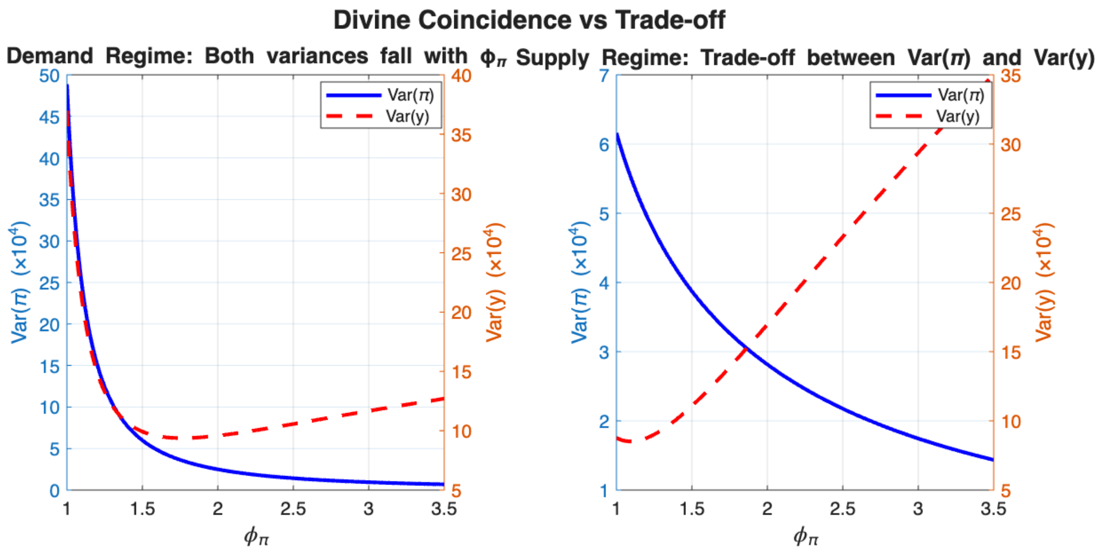

Figure 3: Divine Coincidence vs. Trade-off

Notes: Left panel shows variance of inflation (solid blue, left axis) and output gap (dashed red, right axis) as functions of ϕ_π in the demand regime. Both variances decline with ϕ_π. Right panel shows the same for the supply regime, where higher ϕ_π reduces Var(π) but increases Var(y).

In Table 10 we present the main welfare results. The optimal state-dependent rule ($\phi_\pi^S = 1.26$, $\phi_\pi^D = 3.06$) reduces welfare loss by about 20.8% relative to the constant Taylor rule. This gain is nearly four times what can be achieved by optimizing a constant rule (5.6% gain with $\phi_\pi = 1.79$).

Table 10: Welfare Comparison: Constant vs State-Dependent Taylor Rule

	Taylor ($\phi_\pi = 1.5$)	Taylor ($\phi_\pi = 2.0$)	Optimal Constant	State-Dependent Optimal
$\text{Var}(\pi) \times 10^4$	4.659	2.696	2.673	3.267
$\text{Var}(y) \times 10^4$	3.291	10.652	14.222	14.307
$\text{Var}(i) \times 10^4$	10.083	10.036	9.187	9.290
Welfare Loss $\times 10^4$	7.501	7.200	7.083	5.942
Gain vs $\phi_\pi = 1.5$	–	4.01%	5.57%	**20.78%**
Gain vs $\phi_\pi = 2.0$	–	–	1.62%	**17.47%**
Policy Coefficients:				
ϕ_π^S (Supply Regime)	1.50	2.00	1.79	1.26
ϕ_π^D (Demand Regime)	1.50	2.00	1.79	3.06
Ratio (ϕ_π^D/ϕ_π^S)	1.00	1.00	1.00	2.43
Regime Characteristics (from empirical analysis):				
Information Rigidity ω^D	0.30 (high rigidity, Table 4)			
Information Rigidity ω^S	0.05 (near-rational, Table 4)			
Regime Frequency	38.9% demand, 61.1% supply (Table 1)			

Note: Welfare loss following Woodford (2003) and Rotemberg and Woodford (1997): $\mathcal{L} = \frac{1}{2}(\lambda_\pi \text{Var}(\pi) + \lambda_y \text{Var}(y) + \lambda_i \text{Var}(i))$ with $\lambda_\pi = 1.0$, $\lambda_y = 0.50$, $\lambda_i = 0.50$. Interest rate volatility penalizes aggressive policy, following Orphanides and Williams (2007). The model features hybrid expectations where ω is the share of backward-looking agents.

The optimal ratio $\phi_\pi^D/\phi_\pi^S \approx 2.4$ implies that central banks should respond roughly two and a half times as aggressively to inflation during demand-driven episodes. The state-dependent rule achieves a 20.8% welfare gain relative to the Taylor benchmark, substantially exceeding the 5.6% gain from merely re-optimizing a constant rule.

Moving to the welfare comparison, Figure 4 summarizes our findings. The state-dependent rule delivers welfare gains that are economically meaningful: a roughly 21% reduction in welfare loss corresponds to a substantial improvement in macroeconomic stability over the business cycle.

Figure 4: Welfare Comparison Across Policy Rules

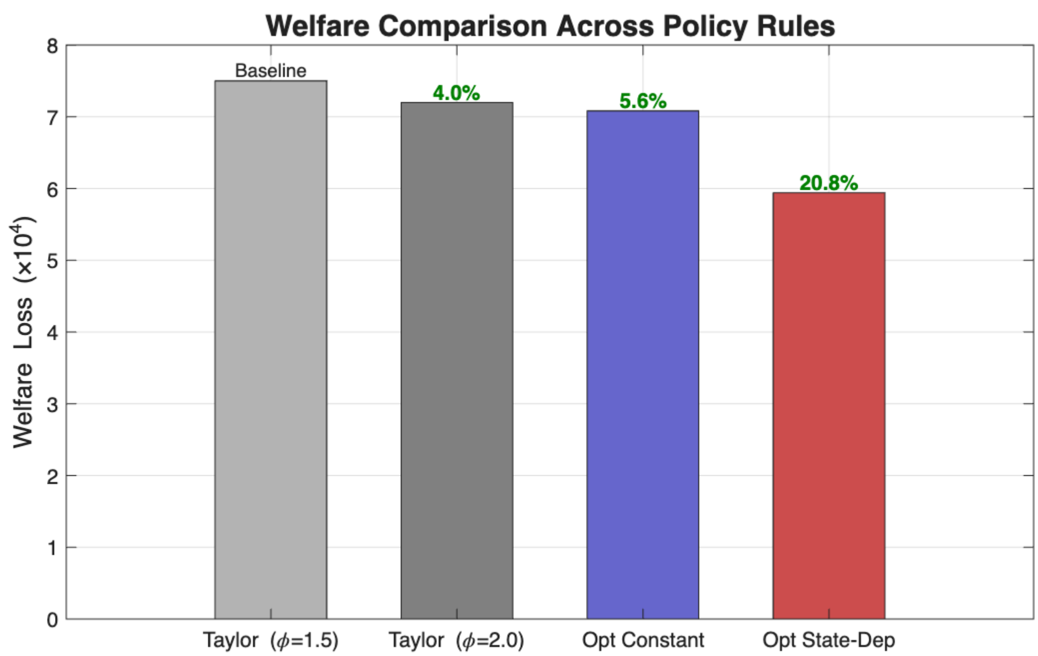

Notes: Welfare loss under policy rules: Taylor benchmark ($\phi_\pi = 1.5$), Aggressive ($\phi_\pi = 2$), optimal constant rule ($\phi_\pi = 1.83$), and optimal state-dependent rule ($\phi_\pi^S = 1.26$, $\phi_\pi^D = 3.06$). Percentages indicate welfare gain relative to Taylor benchmark.

7.6 Discussion

Our model-based analysis provides a theoretical foundation for regime-dependent monetary policy, grounded in the empirical patterns documented in this paper. Several insights emerge:

First, *information rigidity matters for policy design.* The higher backward-looking share during demand episodes ($\omega^D = 0.30$ vs. $\omega^S = 0.05$) implies that expectations respond sluggishly to policy. Aggressive policy compensates by anchoring expectations more rapidly, consistent with the findings of Orphanides and Williams (2007) in models with adaptive learning.

Second, *shock structure interacts with expectation formation.* Divine coincidence holds during demand-driven episodes, allowing aggressive policy to stabilize both inflation and output simultaneously. During supply-driven episodes, the inflation-output trade-off limits the benefits of aggressive policy. These structural differences, combined with regime-specific information rigidity, generate the optimal ratio $\phi_\pi^D/\phi_\pi^S \approx 2.4$.

Third, *regime identification is valuable.* The welfare gains from state-dependent policy (about 21%) exceed those from optimizing a constant rule (about 6%) by a factor of roughly three to four. This suggests that central banks can achieve meaningful welfare improvements by conditioning their policy stance on the nature of prevailing shocks—information that our empirical analysis shows is recoverable from disagreement patterns between consumer and professional forecasts.

These findings complement recent work on state-dependent policy rules and provide empirical discipline for calibrating such rules based on observable features of inflation expectations.

8 Conclusion

This paper has revisited the comparison of U.S. inflation expectations by explicitly conditioning forecast performance on the nature of inflation shocks. Using the supply–demand decomposition of Shapiro (2024), we classify each quarter as demand- or supply-driven and evaluate one-year-ahead expectations from the *Survey of Professional Forecasters* (SPF), the *University of Michigan Survey of Consumers*, and the *Cleveland Fed* model for both CPI and PCE inflation. Across a wide set of forecast evaluation metrics—bias, RMSE and MAE, efficiency tests, and information-rigidity regressions—a clear pattern of state dependence emerges.

The first main result is a reversal in the ranking of forecasters across regimes. During demand-driven episodes, consumer expectations forecast CPI inflation more accurately than professional forecasts and market-based measures, with smaller forecast errors and lower persistence. During supply-driven episodes, the ranking flips: professional and model-based expectations dominate, while household expectations become less accurate and more backward-looking. For PCE inflation, all three sources systematically under-predict realized inflation by 0.3–0.7 percentage points on average, even when CPI forecasts are unbiased, revealing a persistent CPI–PCE wedge that complicates communication of a PCE-based target.

The second main result is that departures from rationality are themselves regime-specific. In demand episodes, professional forecasters exhibit large, highly persistent, and predictable forecast errors, with information-rigidity coefficients around 0.61–0.69. In supply episodes, these inefficiencies largely disappear and expert forecasts come close to the rational benchmark, while household expectations remain more inertial. Adaptive-expectations regressions show that consumers consistently place substantial weight on lagged inflation, whereas experts rely more on policy and macro variables, especially the Federal Funds rate, in supply-driven regimes. The information set underlying expectations, and how it is processed, thus shifts systematically with the source of inflation.

We interpret these findings in a simple New Keynesian noisy-information framework in which divine coincidence holds approximately in demand regimes but fails in supply regimes. When inflation is mainly demand-driven, it co-moves with the output gap and demand-side conditions that households directly experience—local labor markets, income, and spending. Consumers thus have relatively precise signals about the relevant drivers of inflation, and simple backward-looking heuristics can outperform experts' model-based forecasts when the latter misjudge the persistence of demand shocks or the speed of monetary tightening. When inflation is supply-driven, divine coincidence breaks down and inflation depends on sectoral disturbances and the central bank's policy trade-off between inflation and real activity. In that environment, professional forecasters' access to sectoral data, policy communication, and structural models becomes crucial, and their expectations outperform those of households. In the noisy-information language, this shows up as state-dependent signal precisions and information rigidities that mirror the state-dependent β coefficients and error persistence estimated in the data.

Building on this interpretation, we embed heterogeneous and state-dependent expectations into a New Keynesian model with hybrid (backward- and forward-looking) expectations and regime-specific information rigidity calibrated to the empirical estimates. We compare a standard Taylor rule with a rule that is allowed to respond more aggressively to inflation in demand-driven episodes than in supply-driven ones. The optimal state-dependent rule implies roughly twice as strong a response to inflation when demand shocks dominate as when supply shocks dominate and reduces the welfare loss associated with inflation, output, and interest rate volatility by about 20% relative to the Taylor benchmark—more than twice the gain obtained by merely re-optimizing a constant rule. These results show that recognizing the regime dependence of expectations is not only of diagnostic value but also has quantitatively meaningful implications for optimal monetary policy design.

From a policy perspective, the results imply that no single expectation measure provides a uniformly reliable summary of future inflation. Instead, the informational content of each series is regime-dependent. When inflation is driven by supply disturbances, professional and market-based expectations are more informative about underlying inflation trends, while household expectations mainly reflect salient but partly transitory price changes. When inflation is demand-driven, consumer expectations contain valuable information about the public's perceptions of inflation and the credibility of monetary policy, and may even forecast realized inflation more accurately than expert forecasts. Recognizing this conditional structure can help central banks interpret heterogeneous expectations more effectively, diagnose the sources and persistence of inflation, design state-contingent policy rules, and tailor communication strategies that speak to both experts and households.

Future work could extend this framework by incorporating expectations from firm-level surveys or inflation derivatives, examining the full distribution of expectations rather than just means, and exploring whether similar state-dependent patterns arise in other advanced economies. More broadly, the evidence in this paper suggests that inflation expectations are not a single object but a collection of heterogeneous beliefs formed under different information sets, incentives, and models—and that their usefulness for policy depends crucially on the economic state in which they are formed.

References

Ang, A., Bekaert, G., and Wei, M. (2007). Do macro variables, asset markets, or surveys forecast inflation better? *Journal of Monetary Economics*, 54(4):1163–1212.

Beaudry, P., Hou, C., and Portier, F. (2025). The dominant role of expectations and broad-based supply shocks in driving inflation. *NBER Macroeconomics Annual*.

Blanchard, O. and Galí, J. (2007). Real wage rigidities and the New Keynesian model. *Journal of Money, Credit and Banking*, 39(s1):35–65.

Carroll, C. D. (2003). Macroeconomic expectations of households and professional forecasters. *Quarterly Journal of Economics*, 118(1):269–298.

Chang, Y., Maih, J., and Tan, F. (2021). Origins of monetary policy shifts: A new approach to regime switching in DSGE models. *Journal of Economic Dynamics and Control*, 133:104235.

Coibion, O. and Gorodnichenko, Y. (2012). What can survey forecasts tell us about information rigidities? *Journal of Political Economy*, 120(1):116–159.

Coibion, O. and Gorodnichenko, Y. (2015). Information rigidity and the expectations formation process: A simple framework and new facts. *American Economic Review*, 105(8):2644–2678.

Cornand, C. and Hubert, P. (2022). Information frictions across various types of inflation expectations. *European Economic Review*, 146:104167.

D'Acunto, F., Malmendier, U., and Weber, M. (2023). What do the data tell us about inflation expectations? In *Handbook of economic expectations*, pages 133–161. Elsevier.

Davig, T. and Leeper, E. M. (2007). Generalizing the Taylor principle. *American Economic Review*, 97(3):607–635.

Evans, G. W. (2001). Expectations in macroeconomics adaptive versus eductive learning. *Revue économique*, 52(3):573–582.

Faust, J. and Wright, J. H. (2013). Forecasting inflation. *Handbook of Economic Forecasting*, 2A:2–56.

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

Han, Z. (2024). Asymmetric information and misaligned inflation expectations. *Journal of Monetary Economics*, 143:99–121.

Hofmann, B., Manea, C., and Mojon, B. (2024). Targeted taylor rules: monetary policy responses to demand-and supply-driven inflation. *BIS Quarterly Review*, pages 19–35.

Leeper, E. M. and Davig, T. (2006). *Endogenous Monetary Policy Regime Change*. National Bureau of Economic Research.

Link, S., Peichl, A., Roth, C., and Wohlfart, J. (2023). Information frictions among firms and households. *Journal of Monetary Economics*, 135:99–115.

Maćkowiak, B. and Wiederholt, M. (2009). Optimal sticky prices under rational inattention. *American Economic Review*, 99(3):769–803.

Mankiw, N. G. and Reis, R. (2002). Sticky information versus sticky prices: a proposal to replace the new keynesian phillips curve. *The Quarterly Journal of Economics*, 117(4):1295–1328.

Mankiw, N. G. and Reis, R. (2007). Sticky information in general equilibrium. *Journal of the European Economic Association*, 5(2-3):603–613.

Mankiw, N. G., Reis, R., and Wolfers, J. (2003). Disagreement about inflation expectations. *NBER macroeconomics annual*, 18:209–248.

Nordhaus, W. D. (1987). Forecasting efficiency: concepts and applications. *The Review of Economics and Statistics*, pages 667–674.

Orphanides, A. and Williams, J. C. (2005). The decline of activist stabilization policy: Natural rate misperceptions, learning, and expectations. *Journal of Economic Dynamics and Control*, 29(11):1927–1950.

Orphanides, A. and Williams, J. C. (2007). Robust monetary policy with imperfect knowledge. *Journal of monetary Economics*, 54(5):1406–1435.

Rabanal, P. (2004). Monetary policy rules and the us business cycle: Evidence and implications.

Reis, R. (2023). Four mistakes in the use of measures of expected inflation. In *AEA Papers and Proceedings*, volume 113, pages 47–51. American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203.

Rotemberg, J. J. and Woodford, M. (1997). An optimization-based econometric framework for the evaluation of monetary policy. *NBER macroeconomics annual*, 12:297–346.

Shapiro, A. H. (2024). Decomposing supply-and demand-driven inflation. *Journal of Money, Credit and Banking*.

Singh, A. and Mitra, A. (2022). What determines household expectations? *Available at SSRN 4189773*.

Smets, F. and Wouters, R. (2007). Shocks and frictions in us business cycles: A bayesian dsge approach. *American economic review*, 97(3):586–606.

Taylor, J. B. (1993). Discretion versus policy rules in practice. In *Carnegie-Rochester conference series on public policy*, volume 39, pages 195–214. Elsevier.

Verbrugge, R. J. and Zaman, S. (2021). Whose inflation expectations best predict inflation? *Federal Reserve Bank of Cleveland Economic Commentary*, (2021-19).

Weber, M., D'Acunto, F., Gorodnichenko, Y., and Coibion, O. (2022). The subjective inflation expectations of households and firms: Measurement, determinants, and implications. *Journal of Economic Perspectives*, 36(3):157–184.

Woodford, M. (2003). Interest and prices: Foundations of a theory of monetary policy princeton, nj: Princeton univ.

A Robustness Checks - Alternative definitions of demand and Supply Periods

In this Appendix, we test the robustness of our results by imposing a stricter classification for demand- and supply-driven inflation periods. Using the decomposition methodology of Shapiro (2024), we classify a quarter as purely demand-driven only if the demand component accounts for at least 50% of both headline and core inflation. Conversely, a purely supply-driven quarter is defined as one where the supply component accounts for more than 50% of both headline and core inflation.

Table 11: Forecast Accuracy: All Periods,Pure Demand, and Pure Supply Inflation Periods

	Michigan		SPF		Cleveland Fed	
	CPI	PCE	CPI	PCE	CPI	PCE
Panel A: All Periods						
RMSE	1.525	1.452	1.466	1.312	1.591	1.416
MAE	1.119	1.114	0.980	0.984	1.048	1.012
Bias	-0.271	-0.702	0.015	-0.417	0.136	-0.295
N	167	167	167	167	167	167
Panel B: Demand Inflation Periods						
RMSE	1.364	1.352	1.847	1.653	1.941	1.764
MAE	0.990	1.151	1.254	1.282	1.282	1.354
Bias	-0.201	-0.807	0.085	-0.521	0.070	-0.536
N	40	40	40	40	40	40
Panel C: Supply Inflation Periods						
RMSE	1.722	1.615	1.366	1.235	1.457	1.287
MAE	1.265	1.193	0.890	0.907	0.962	0.898
Bias	-0.356	-0.735	-0.100	-0.479	0.044	-0.335
N	92	92	92	92	92	92

Notes: RMSE = Root Mean Squared Error; MAE = Mean Absolute Error; Bias = Mean Forecast Error. Forecast errors calculated as actual inflation minus 4-quarter-ahead forecast. Panel A includes all observations. Panel B restricts to demand inflation periods. Panel C restricts to supply inflation periods.

Table 12: Tests of Forecast Rationality: Pure Demand Inflation Periods

Panel A: Testing for Bias

	Michigan		SPF		Cleveland Fed	
	CPI	PCE	CPI	PCE	CPI	PCE
Constant	-0.201	-0.807**	0.085	-0.521	0.070	-0.536
	(0.361)	(0.291)	(0.510)	(0.435)	(0.534)	(0.465)
N	40	40	40	40	40	40

Panel B: Is Information in the Forecast Fully Exploited?

	Michigan		SPF		Cleveland Fed	
	CPI	PCE	CPI	PCE	CPI	PCE
$E_{t-4}[\pi_t]$	1.560**	1.257***	-1.090**	-0.929**	-1.087**	-0.994**
	(0.567)	(0.395)	(0.433)	(0.324)	(0.448)	(0.347)
Joint test p-value	0.031	0.000	0.015	0.000	0.010	0.000
N	40	40	40	40	40	40

Panel C: Are Forecasting Errors Persistent?

	Michigan		SPF		Cleveland Fed	
	CPI	PCE	CPI	PCE	CPI	PCE
$\pi_{t-4} - E_{t-8}[\pi_{t-4}]$	0.620***	0.571***	1.032***	1.003***	0.936***	0.908***
	(0.182)	(0.140)	(0.189)	(0.144)	(0.192)	(0.144)
N	40	40	40	40	40	40

Panel D: Are Macroeconomic Data Fully Exploited?

	Michigan		SPF		Cleveland Fed	
	CPI	PCE	CPI	PCE	CPI	PCE
$E_{t-4}[\pi_t]$	2.376**	1.937*	-1.356**	-1.390***	-0.731	-0.910**
	(1.082)	(0.933)	(0.502)	(0.349)	(0.438)	(0.332)
Inflation$_{t-1}$	-0.344	-0.324	0.936***	1.005***	0.806***	0.888***
	(0.405)	(0.405)	(0.289)	(0.268)	(0.246)	(0.242)
Fed Funds$_{t-1}$	0.105	0.105	-0.274	-0.177	-0.400	-0.264
	(0.145)	(0.113)	(0.210)	(0.151)	(0.259)	(0.184)
Unemployment$_{t-1}$	-0.232	-0.070	-0.029	-0.070	-0.136	-0.137
	(0.178)	(0.155)	(0.107)	(0.098)	(0.173)	(0.146)
Joint test p-value	0.511	0.676	0.030	0.003	0.035	0.008
N	40	40	40	40	40	40

Notes: Robust standard errors clustered by year in parentheses. Dependent variable in Panels A-C: Forecast Error ($\pi_t - E_{t-4}\pi_t$). Panel B joint test: $\alpha = \beta = 0$. Panel D joint test: coefficients on inflation, Fed Funds, and unemployment = 0. Sample restricted to demand inflation periods. * p<0.10, ** p<0.05, *** p<0.01

Table 13: Tests of Forecast Rationality: Pure Supply Inflation Periods

Panel A: Testing for Bias

	Michigan		SPF		Cleveland Fed	
	CPI	PCE	CPI	PCE	CPI	PCE
Constant	-0.356	-0.735***	-0.100	-0.479***	0.044	-0.335*
	(0.273)	(0.236)	(0.190)	(0.164)	(0.197)	(0.173)
N	92	92	92	92	92	92

Panel B: Is Information in the Forecast Fully Exploited?

	Michigan		SPF		Cleveland Fed	
	CPI	PCE	CPI	PCE	CPI	PCE
$E_{t-4}[\pi_t]$	-0.619	-0.667*	-0.060	-0.192	-0.230	-0.339**
	(0.444)	(0.386)	(0.147)	(0.122)	(0.165)	(0.138)
Joint test p-value	0.331	0.014	0.604	0.000	0.378	0.003
N	92	92	92	92	92	92

Panel C: Are Forecasting Errors Persistent?

	Michigan		SPF		Cleveland Fed	
	CPI	PCE	CPI	PCE	CPI	PCE
$\pi_{t-4} - E_{t-8}[\pi_{t-4}]$	-0.034	0.092	-0.006	0.075	0.045	0.145*
	(0.152)	(0.135)	(0.121)	(0.117)	(0.093)	(0.078)
N	92	92	92	92	92	92

Panel D: Are Macroeconomic Data Fully Exploited?

	Michigan		SPF		Cleveland Fed	
	CPI	PCE	CPI	PCE	CPI	PCE
$E_{t-4}[\pi_t]$	-1.192**	-1.324***	0.155	0.061	-0.467*	-0.485**
	(0.525)	(0.431)	(0.326)	(0.299)	(0.258)	(0.224)
Inflation$_{t-1}$	0.279*	0.435***	-0.015	0.033	0.125	0.198*
	(0.146)	(0.147)	(0.118)	(0.119)	(0.115)	(0.103)
Fed Funds$_{t-1}$	0.177**	0.088	-0.076	-0.113	0.044	-0.022
	(0.081)	(0.068)	(0.106)	(0.085)	(0.097)	(0.073)
Unemployment$_{t-1}$	0.149	0.139*	-0.012	0.028	0.077	0.092
	(0.093)	(0.078)	(0.079)	(0.056)	(0.090)	(0.070)
Joint test p-value	0.004	0.002	0.869	0.476	0.605	0.202
N	92	92	92	92	92	92

Notes: Robust standard errors clustered by year in parentheses. Dependent variable in Panels A-C: Forecast Error ($\pi_t - E_{t-4}\pi_t$). Panel B joint test: $\alpha = \beta = 0$. Panel D joint test: coefficients on inflation, Fed Funds, and unemployment = 0. Sample restricted to supply inflation periods. * p<0.10, ** p<0.05, *** p<0.01

Table 14: Tests of Adaptive Expectations Across Regimes

Panel A: All Periods

	Michigan		SPF		Cleveland Fed	
	CPI	PCE	CPI	PCE	CPI	PCE
$\sum \beta$ (inflation lags)	0.313***	0.373***	0.152**	0.214***	0.135**	0.161**
	(0.041)	(0.043)	(0.064)	(0.075)	(0.058)	(0.068)
Unemployment$_t$	0.006	0.005	0.054*	0.053*	-0.089***	-0.091***
	(0.034)	(0.028)	(0.029)	(0.028)	(0.018)	(0.019)
Unemployment$_{t-1}$	0.052**	0.034*	0.083**	0.068**	0.186***	0.173***
	(0.025)	(0.019)	(0.034)	(0.031)	(0.022)	(0.021)
Fed Funds$_t$	0.206	0.141	0.383***	0.355***	0.384***	0.361***
	(0.130)	(0.120)	(0.107)	(0.105)	(0.103)	(0.106)
Fed Funds$_{t-1}$	-0.236*	-0.176	-0.150	-0.133	-0.096	-0.075
	(0.120)	(0.111)	(0.094)	(0.096)	(0.104)	(0.109)
Reject adaptive (p-value)	0.004	0.009	0.000	0.000	0.000	0.000
N	159	159	159	159	159	159

Panel B: Pure Demand Inflation Periods

	Michigan		SPF		Cleveland Fed	
	CPI	PCE	CPI	PCE	CPI	PCE
$\sum \beta$ (inflation lags)	0.306***	0.335***	0.304**	0.230**	0.287***	0.208**
	(0.088)	(0.054)	(0.112)	(0.099)	(0.057)	(0.074)
Unemployment$_t$	-0.364	-0.289	-0.169	0.015	-0.232	-0.278
	(0.286)	(0.243)	(0.354)	(0.233)	(0.230)	(0.245)
Unemployment$_{t-1}$	0.365	0.244	0.432	0.311	0.361	0.464*
	(0.311)	(0.244)	(0.365)	(0.229)	(0.219)	(0.246)
Fed Funds$_t$	-0.144	-0.081	0.187	0.367***	0.447**	0.520***
	(0.093)	(0.064)	(0.186)	(0.123)	(0.193)	(0.177)
Fed Funds$_{t-1}$	0.042	-0.001	-0.053	-0.200*	-0.193	-0.215
	(0.098)	(0.065)	(0.170)	(0.106)	(0.193)	(0.182)
Reject adaptive (p-value)	0.000	0.000	0.087	0.008	0.000	0.000
N	40	40	40	40	40	40

Panel C: Pure Supply Inflation Periods

	Michigan		SPF		Cleveland Fed	
	CPI	PCE	CPI	PCE	CPI	PCE
$\sum \beta$ (inflation lags)	0.320***	0.437***	0.147	0.201*	0.147	0.172
	(0.070)	(0.086)	(0.090)	(0.102)	(0.096)	(0.105)
Unemployment$_t$	0.037	0.051	-0.050	-0.004	0.296	0.340*
	(0.219)	(0.212)	(0.166)	(0.148)	(0.180)	(0.172)
Unemployment$_{t-1}$	0.015	-0.030	0.230	0.174	-0.144	-0.198
	(0.216)	(0.207)	(0.159)	(0.136)	(0.153)	(0.144)
Fed Funds$_t$	0.295**	0.194	0.413***	0.381***	0.513***	0.472***
	(0.140)	(0.120)	(0.122)	(0.124)	(0.128)	(0.134)
Fed Funds$_{t-1}$	-0.291**	-0.219**	-0.157	-0.137	-0.236*	-0.204
	(0.129)	(0.105)	(0.109)	(0.109)	(0.122)	(0.125)
Reject adaptive (p-value)	0.221	0.081	0.000	0.000	0.000	0.000
N	87	87	87	87	87	87

Notes: Dependent variable: Expected Inflation $E_{t-4}\pi_t$. Regressions include 8 quarterly inflation lags (not shown). Robust standard errors clustered by year in parentheses. Joint test: all macro variables = 0. Panel A: Full sample. Panel B: Demand inflation periods (demand_side==1 or demand_both==1). Panel C: Supply inflation periods (demand_side==0 or demand_both==0). * p<0.10, ** p<0.05, *** p<0.01

A.1 Real-Time Regime Classification Tables

Table 15: Real-Time vs Ex-Post Regime Classification Agreement

Real-Time Method	Agreement with Ex-Post Demand	Supply	Agreement Rate (%)	N
1-Quarter Lag	75.7%	89.5%	85.2%	223
4-Quarter Lag	50.0%	79.3%	70.0%	220
Moving Average (4Q)	62.9%	83.3%	76.8%	220
Probit Prediction	44.9%	89.3%	75.2%	218

Notes: Agreement rates show percentage of quarters where real-time classification matches ex-post Shapiro (2024) decomposition. 1-Quarter Lag uses share_demand ≥ 0.50. 4-Quarter Lag uses share_demand ≥ 0.50. Moving Average uses 4-quarter average of lagged shares. Probit Prediction uses lagged inflation, oil prices, unemployment, and output gap.

Table 16: Forecast Accuracy: Ex-Post vs Real-Time Regime Classification

	Michigan		SPF		Cleveland Fed	
	CPI	PCE	CPI	PCE	CPI	PCE
Panel A: Ex-Post Classification (Baseline)						
Demand Periods						
RMSE	1.320	1.289	1.651	1.440	1.809	1.595
Supply Periods						
RMSE	1.643	1.546	1.335	1.223	1.434	1.289
Panel B: Real-Time Classification (1-Quarter Lag)						
Demand Periods						
RMSE	1.277	1.271	1.562	1.381	1.634	1.470
Supply Periods						
RMSE	1.661	1.553	1.403	1.267	1.563	1.382
Panel C: Real-Time Classification (Moving Average)						
Demand Periods						
RMSE	1.451	1.420	1.592	1.427	1.718	1.563
Supply Periods						
RMSE	1.569	1.471	1.382	1.234	1.507	1.317

Notes: RMSE = Root Mean Squared Error. Panel A uses ex-post Shapiro (2024) decomposition. Panels B and C use real-time classification based only on information available at the forecast date. The ranking reversal (Michigan outperforms in demand, SPF outperforms in supply) is robust to real-time classification.

Table 17: Forecast Efficiency Tests: Real-Time vs Ex-Post Classification

	Michigan		SPF		Cleveland Fed	
	CPI	PCE	CPI	PCE	CPI	PCE
Panel A: Ex-Post Demand Classification						
$E_{t-4}[\pi_t]$	-0.686**	-0.352	-1.043***	-0.905***	-1.087***	-0.983***
	(0.255)	(0.218)	(0.324)	(0.279)	(0.321)	(0.269)
N	65	65	65	65	65	65
Panel B: Real-Time Demand (1-Quarter Lag)						
$E_{t-4}[\pi_t]$	-0.833***	-0.468*	-0.998***	-0.889***	-0.939***	-0.883***
	(0.248)	(0.238)	(0.288)	(0.229)	(0.247)	(0.205)
N	64	64	64	64	64	64
Panel C: Ex-Post Supply Classification						
$E_{t-4}[\pi_t]$	-0.215	-0.268	-0.137	-0.238*	-0.287	-0.357***
	(0.230)	(0.184)	(0.178)	(0.130)	(0.171)	(0.129)
N	98	98	98	98	98	98
Panel D: Real-Time Supply (1-Quarter Lag)						
$E_{t-4}[\pi_t]$	-0.158	-0.222	-0.190	-0.271*	-0.402	-0.443**
	(0.237)	(0.185)	(0.211)	(0.158)	(0.243)	(0.185)
N	99	99	99	99	99	99

Notes: Dependent variable: Forecast Error ($\pi_t - E_{t-4}\pi_t$). Robust standard errors clustered by year in parentheses. Under forecast efficiency, $\beta = 0$. Panels A and C use ex-post Shapiro (2024) classification. Panels B and D use real-time classification (1-quarter lagged share). The pattern of severe inefficiency for SPF/Cleveland in demand regimes and near-efficiency in supply regimes is robust to real-time classification. * p<0.10, ** p<0.05, *** p<0.01

Table 18: Forecast Error Persistence: Real-Time vs Ex-Post Classification

	Michigan		SPF		Cleveland Fed	
	CPI	PCE	CPI	PCE	CPI	PCE
Panel A: Ex-Post Demand						
$\pi_{t-4} - E_{t-8}[\pi_{t-4}]$	0.322	0.273	0.678**	0.691**	0.609**	0.643***
	(0.220)	(0.198)	(0.299)	(0.253)	(0.290)	(0.229)
N	65	65	65	65	65	65
Panel B: Real-Time Demand (1Q Lag)						
$\pi_{t-4} - E_{t-8}[\pi_{t-4}]$	0.267**	0.221*	0.522***	0.548***	0.491***	0.534***
	(0.117)	(0.114)	(0.184)	(0.174)	(0.172)	(0.157)
N	64	64	64	64	64	64
Panel C: Ex-Post Supply						
$\pi_{t-4} - E_{t-8}[\pi_{t-4}]$	-0.033	0.099	0.004	0.097	0.058	0.171**
	(0.149)	(0.135)	(0.118)	(0.111)	(0.091)	(0.077)
N	102	102	102	102	102	102
Panel D: Real-Time Supply (1Q Lag)						
$\pi_{t-4} - E_{t-8}[\pi_{t-4}]$	-0.077	0.100	-0.037	0.085	0.028	0.166
	(0.189)	(0.191)	(0.132)	(0.147)	(0.153)	(0.168)
N	103	103	103	103	103	103

Notes: Dependent variable: Forecast Error ($\pi_t - E_{t-4}\pi_t$). Robust standard errors clustered by year in parentheses. Under rational expectations, $\beta = 0$ (no error persistence). High persistence (0.6–0.7) for SPF/Cleveland in demand regimes is robust to real-time classification. * p<0.10, ** p<0.05, *** p<0.01

A.2 Sample Composition Robustness

Table 19: Regime Distribution by Decade

Decade	Demand	Supply	Total	% Demand	Mean Inflation
1983-1989	5	27	32	15.6%	3.96%
1990-1999	23	17	40	57.5%	3.01%
2000-2009	9	31	40	22.5%	2.57%
2010-2019	18	22	40	45.0%	1.77%
2020-2024	10	9	19	52.6%	4.28%
Total	65	106	171	38.0%	2.94%

Notes: Demand periods defined as quarters where demand factors account for $\geq$ 50% of inflation per Shapiro (2024). The 1990s and 2020s show higher shares of demand-driven inflation, while the 1980s and 2000s are predominantly supply-driven.

Table 20: Forecast Efficiency: Controlling for Time Period

	Michigan		SPF		Cleveland Fed	
	CPI	PCE	CPI	PCE	CPI	PCE
Panel A: Demand Regimes - No Fixed Effects						
$E_{t-4}[\pi_t]$	-0.686**	-0.352	-1.043***	-0.905***	-1.087***	-0.983***
	(0.255)	(0.218)	(0.324)	(0.279)	(0.321)	(0.269)
Panel B: Demand Regimes - With Period Fixed Effects						
$E_{t-4}[\pi_t]$	-1.044**	-0.729**	-2.136***	-1.762***	-1.773***	-1.446***
	(0.394)	(0.346)	(0.373)	(0.286)	(0.292)	(0.241)
Panel C: Supply Regimes - No Fixed Effects						
$E_{t-4}[\pi_t]$	-0.002	-0.072	-0.117	-0.208*	-0.250*	-0.321***
	(0.264)	(0.229)	(0.137)	(0.105)	(0.138)	(0.107)
Panel D: Supply Regimes - With Period Fixed Effects						
$E_{t-4}[\pi_t]$	-0.453*	-0.486**	-0.391**	-0.498***	-0.637***	-0.676***
	(0.261)	(0.220)	(0.151)	(0.125)	(0.173)	(0.142)

Notes: Dependent variable: Forecast Error ($\pi_t - E_{t-4}\pi_t$). Period fixed effects: Pre-1990, 1990-2007, 2008-2019, 2020+. Robust standard errors clustered by year. The pattern of inefficiency in demand regimes and near-efficiency in supply regimes is robust to controlling for time period effects. * p<0.10, ** p<0.05, *** p<0.01

B Optimal Monetary Policy Model: Solution, Welfare Computation and Robustness Checks

This appendix provides technical details on the New Keynesian model with hybrid expectations and describes the numerical methods used for the welfare analysis.

B.1 Complete Model Specification

The model consists of eight equations in eight endogenous variables $(y_t, \pi_t, i_t, r_t^{nat}, \pi_t^e, y_t^e, g_t, u_t)$:

$$\pi_t^e = \omega \cdot \pi_{t-1} + (1-\omega) \cdot \mathbb{E}_t[\pi_{t+1}] \quad \text{(Hybrid } \pi \text{ expectation)}$$

$$y_t^e = \omega \cdot y_{t-1} + (1-\omega) \cdot \mathbb{E}_t[y_{t+1}] \quad \text{(Hybrid } y \text{ expectation)}$$

$$y_t = y_t^e - \frac{1}{\sigma}(i_t - \pi_t^e - r_t^{nat}) \quad \text{(IS Curve)}$$

$$\pi_t = \beta \pi_t^e + \kappa y_t + u_t \quad \text{(Phillips Curve)}$$

$$r_t^{nat} = \sigma(1-\rho_g) g_t \quad \text{(Natural Rate)}$$

$$i_t = \phi_\pi \pi_t + \phi_y y_t \quad \text{(Taylor Rule)}$$

$$g_t = \rho_g g_{t-1} + \varepsilon_t^g, \quad \varepsilon_t^g \sim N(0, \sigma_g^2) \quad \text{(Demand Shock)}$$

$$u_t = \rho_u u_{t-1} + \varepsilon_t^u, \quad \varepsilon_t^u \sim N(0, \sigma_u^2) \quad \text{(Cost-Push Shock)}$$

The hybrid expectation specification nests pure rational expectations ($\omega = 0$) and pure adaptive expectations ($\omega = 1$) as special cases. For intermediate values, the model captures gradual expectation adjustment consistent with models of sticky information (Mankiw and Reis, 2002) or adaptive learning (Evans, 2001).

B.2 Solution Method

We solve the model using Dynare 6.1. The linear rational expectations solution is computed using the method of Blanchard and Kahn (1980). For each regime $j \in \{D, S\}$, we specify regime-specific parameters $(\omega^j, \sigma_g^j, \sigma_u^j)$ and compute the unconditional variance-covariance matrix of endogenous variables.

B.3 Welfare Computation

The welfare loss function is:

$$\mathcal{L} = \frac{1}{2} (\lambda_\pi \text{Var}(\pi) + \lambda_y \text{Var}(y) + \lambda_i \text{Var}(i)) \tag{20}$$

For state-dependent policy evaluation, we compute regime-specific welfare and aggregate using empirical regime frequencies:

$$\mathcal{L}^{total} = \Pr(S_t = D) \cdot \mathcal{L}^D(\phi_\pi^D) + \Pr(S_t = S) \cdot \mathcal{L}^S(\phi_\pi^S) \tag{21}$$

The optimal state-dependent rule is found by independently minimizing $\mathcal{L}^D$ over ϕ_π^D and $\mathcal{L}^S$ over ϕ_π^S. We search over a grid $\phi_\pi \in \{1.0, \ldots, 3.5\}$.

B.4 Robustness

Table 21 examines the sensitivity of our optimal policy results to key calibration choices. Across all specifications, the central finding that optimal policy should be substantially more aggressive during demand-driven episodes ($\phi_\pi^D > \phi_\pi^S$) remains robust. The optimal ratio ranges from 1.12 to 2.61, and the welfare gains from state-dependent policy are economically meaningful in all cases, between 7.0% and 22.0%.

Increasing the weight on interest rate volatility to $\lambda_i = 0.75$ raises the optimal demand-regime coefficient to $\phi_\pi^{D*} = 3.24$ while leaving the supply coefficient nearly unchanged at

$\phi_\pi^{S*} = 1.24$, implying a ratio of 2.61. This result, while seemingly counterintuitive, reflects a mechanism emphasized by Orphanides and Williams (2007): when information rigidity is high, more aggressive policy can *reduce* the volatility of interest rates by anchoring expectations more rapidly. In the demand regime, where backward-looking behavior is prevalent, aggressive policy stabilizes inflation quickly, shortening the duration of rate adjustments. The welfare gain rises to 22.0% because a higher weight on $\text{Var}(i)$ magnifies the benefits of avoiding persistent policy swings.

Alternative calibrations of information rigidity produce similar conclusions. When information rigidity is recalibrated using the lag-structure estimates ($\omega^D = 0.244$, $\omega^S = 0.024$), the optimal ratio remains high at 2.40 and the welfare gain is 19.8%. Using best estimates of rigidity ($\omega^D = 0.322$, $\omega^S = 0.054$) yields the largest ratio, 2.46, with a welfare gain of 21.4%. These results confirm that the strength of state dependence in optimal policy is tightly linked to differences in information rigidity across regimes: the more inattentive agents are in demand episodes relative to supply episodes, the more valuable aggressive stabilization policy becomes.

The most informative robustness check equalizes shock variances across regimes, setting $\sigma_g = \sigma_u = 0.008$. This removes asymmetries in shock structure and isolates the role of expectations. In this case, the optimal ratio falls sharply to 1.12 and welfare gains decline to 7.0%. With symmetric shocks, the divine coincidence mechanism is attenuated and the inflation–output trade-off becomes more similar across regimes. The remaining difference in policy aggressiveness is therefore driven primarily by the residual asymmetry in information rigidity.

Taken together, these robustness exercises reinforce the economic interpretation of our results. Both channels matter: (i) differences in shock composition across regimes and (ii) differences in information rigidity. But even when we strip out asymmetries in shocks, meaningful welfare gains from state-dependent policy remain. The conclusion that central banks should respond more aggressively to inflation during demand-driven episodes is therefore not an artifact of a narrow calibration assumption, but a robust quantitative implication of the interaction between heterogeneous expectations and regime-dependent inflation dynamics.

Table 21: Sensitivity of Optimal Policy to Calibration

Specification	ϕ_π^{S*}	ϕ_π^{D*}	Ratio	Welfare Gain
Baseline	1.26	3.06	2.43	20.8%
Higher $\lambda_i = 0.75$	1.24	3.24	2.61	22.0%
Lag Regime $\omega^D = 0.244$ $\omega^S = 0.024$	1.24	2.98	2.40	19.8%
Best Accuracy $\omega^D = 0.322$ $\omega^S = 0.054$	1.26	3.10	2.46	21.4%
Equal shock variances	1.76	1.98	1.12	7.0%

Notes: Welfare gains in relation to Taylor Rule with $\phi_\pi = 1.5$. Baseline uses $\lambda_i = 0.50$, $\omega^D = 0.30$ $\omega^S = 0.05$. Shock variances in baseline: $\sigma_g^D = 0.010$, $\sigma_u^D = 0.005$, $\sigma_g^S = 0.005$, $\sigma_u^S = 0.010$. Equal shock variances specification uses $\sigma_g = \sigma_u = 0.008$ in both regimes.