The Phillips Curve Above and Below 2%

Engin Kara*

Abstract

I estimate the Phillips curve's cost and expectations channels using UK micro-data matched with upstream cost indices. Below 1.9% sectoral inflation, sticky-price firms filter noisy signals; above, filtering stops and pass-through rises. The expectations coefficient doubles (0.52 to 0.99); the cost coefficient quadruples (0.03 to 0.13). Identification exploits an asymmetry: above threshold, the gap between sticky and flexible firms closes on expectations (98%) but not costs (15%). Confirming signal extraction, competitor prices are ambiguous, costs are not. Cognitive discounting emerges as the low-inflation limit of rational signal extraction. Credible central banks can "look through" supply shocks by anchoring expectations below threshold.

Keywords: Inflation targeting, signal extraction, Phillips curve, behavioural macroeconomics, price-setting.

JEL classification numbers: E31, E52, E58, D21, L13.

*Corresponding author: Cardiff Business School, University of Cardiff, Cardiff, CF10 3EU, UK. E-mail: karae1@cardiff.ac.uk

1. Introduction

Monetary policy operates through the Phillips curve. Understanding when and why the slope changes is essential for policy design. Yet the empirical literature remains contested. Some find it flat—cost shocks do not pass through to prices. Others find it steep—firms respond strongly to marginal costs. This paper shows both are correct. The slope depends on signal-to-noise: flat below 1.9% inflation, steep above. The mechanism is signal extraction.

I estimate both channels of the Phillips curve—expectations and costs—directly from micro data. Using six million UK price observations matched with upstream cost indices constructed from input-output tables and producer prices, I show that both channels are state-dependent. Below 1.9% sectoral inflation, sticky-price firms filter price signals; above threshold, filtering stops and pass-through rises sharply. The expectations coefficient nearly doubles (0.52 to 0.99) and the cost coefficient quadruples (0.03 to 0.13). Adjustment frequency is identical across regimes, ruling out menu cost explanations.

The identification strategy exploits the divergence between the two channels. The fundamental asymmetry is that competitor price changes are ambiguous signals, whereas cost shocks are not. Competitor price movements conflate persistent aggregate inflation with transitory idiosyncratic shocks. Below the threshold, idiosyncratic volatility dominates. A sticky firm observes a competitor's price rise and faces a signal extraction problem: is this a permanent shift or a temporary overshooting? If the competitor is flexible, they may be reacting to a temporary shock and will likely reverse price in the next period. A sticky firm, anticipating this reversal but unable to adjust quickly itself, optimally *leans against* the movement to maintain its average competitive position. Above the threshold, the signal clarifies. Large movements are attributed to common aggregate shocks, which are persistent. The expectation of competitor reversal vanishes, so sticky firms stop leaning against and catch up. This generates the observed **convergence** in the expectations channel.

Costs behave differently. Input cost indices are verifiable and persistent; they do not suffer from the same signal ambiguity. When the threshold is

crossed, both firm types pay more attention to costs, but because neither expects costs to revert, they increase pass-through by similar amounts. This generates the **parallel shift** in the cost channel. The data confirm this asymmetry: 98% convergence on expectations, 15% on costs. A formal test rejects equality ($F = 26.89$, $p < 0.001$). The threshold is discrete: I test against smooth state-dependence and reject it. The discontinuity is consistent with signal extraction—below threshold, noise dominates and firms filter; above, signals clarify and pass-through rises discretely.

Signal extraction alone predicts pass-through between zero and one (discounting). The data show something stronger: below threshold, sticky-price firms exhibit negative pass-through (-0.72). This requires a second ingredient: **competitor dynamics**. Sticky-price firms observe flexible-price competitors who can quickly correct pricing errors. The data confirm that below threshold, flexible prices actively revert ($AR(1) = -0.26$); above threshold, they do not ($AR(1) = 0.00$). A sticky firm, locked into its price for multiple periods, anticipates that flexible competitors will reverse. Optimal forward-looking pricing therefore requires leaning against current movements. Above threshold, even flexible prices persist because shocks are recognised as aggregate—no reversal is expected, so no leaning against occurs. Consistent with this, the negative pass-through is strongest precisely for sticky firms facing flexible competitors (-1.45), and it vanishes above threshold (-0.20).

Calibrating the aggregate Phillips curve using the micro estimates yields dramatic state-dependence. Below threshold, the expectations coefficient is 0.52—firms heavily discount future inflation. The cost coefficient is 0.03—the curve is nearly flat. Above threshold, the expectations coefficient rises to 0.99—the rational expectations benchmark. The cost coefficient rises to 0.13—standard Calvo. The Phillips curve steepens by a factor of four. This has immediate policy implications. Below threshold, cost shocks are absorbed, expectations are anchored by filtering, and the output-inflation trade-off is favourable. Above threshold, both channels activate: cost shocks pass through rapidly, expectations become forward-looking, and the trade-off deteriorates. The Taylor principle becomes harder to satisfy. Optimal policy shifts from ag-

gressive (responding to current conditions) to gradualist (smoothing the path to avoid expectational disruption). The 2% inflation target approximately coincides with the signal extraction threshold—marking the boundary where price-setting behaviour fundamentally changes.

These findings reconcile conflicting evidence on the Phillips curve. Hazell et al. (2022) estimate a flat Phillips curve using US data from 1978–2020. Gagliardone et al. (2025) argue this reflects their use of output gaps rather than marginal costs; using Belgian firm-level cost data, they estimate substantial pass-through. My findings are consistent with both and add another dimension: cost pass-through quadruples above threshold, so the Phillips curve is flat below and steep above. Gabaix (2020) proposes cognitive discounting—agents place reduced weight on future variables, flattening the Phillips curve. He calibrates an expectations coefficient of 0.76 and treats it as a behavioural primitive. My weighted average is 0.73—strikingly close. But this is not a primitive. It is the average of two rational regimes: deep filtering below threshold (0.52) and full-information pricing above (0.99). What appears as "bounded rationality" may be a rational response to a genuine information problem. Relatedly, what appears as "muddled thinking" in survey evidence on firm expectations (Coibion and Gorodnichenko, 2015) may be rational filtering in low-inflation environments. When inflation rises, filtering stops—not because agents become more sophisticated, but because the signal-to-noise ratio improves.

The findings also speak to current policy debates. Nakamura et al. (2025) document that the Federal Reserve "looked through" post-Covid inflation, achieving disinflation without a severe recession. My results identify why this was possible: a credible central bank anchors long-run expectations in the low-inflation regime, where firms filter. Credibility allows the central bank to borrow the stability properties of the filtering regime even when inflation temporarily exceeds threshold.

The findings also address the markup shock puzzle. Chari et al. (2009) criticise New Keynesian models for requiring large, unexplained markup shocks to fit the data. But in a one-sector model that ignores filtering, the econometri-

cian must attribute the dampened response of sticky-price sectors to exogenous shocks. The shocks are not exogenous—they are the endogenous response of sticky-price firms to noisy signals from flexible-price sectors. What appears as an inexplicable markup shock is partly the systematic filtering of sector-specific volatility. The puzzle is a misspecification: one-sector models miss the filtering mechanism that multi-sector economies generate automatically.

The findings connect to the theory of sectoral heterogeneity and monetary policy. In standard one-sector New Keynesian models, there is no tradeoff between output and price stability—the "divine coincidence" of Blanchard and Gali (2007). As established by Kara (2010), this coincidence breaks in the presence of sectoral heterogeneity. Aoki (2001) shows that optimal policy places greater weight on sticky-price sectors. Yet Kara (2010) finds that aggregate inflation targeting performs nearly as well as optimal sectoral weighting. The current paper suggests an explanation. Below threshold, sticky-price firms filter flexible-price noise, so aggregate inflation already reflects informative signals—and sticky sectors deserve greater weight precisely because they filter. Above threshold, sticky and flexible firms converge; behavioural heterogeneity disappears even though structural heterogeneity remains, and aggregate inflation contains as much information as sticky prices. In both regimes, the gains from sectoral targeting are limited.

The paper proceeds as follows. Section 2 presents the theoretical framework, deriving the signal extraction threshold and horizon filter. Section 3 describes the data, including the construction of upstream cost indices from input-output tables. Section 4 presents the empirical results: the joint test of both channels, the competitor flexibility mechanism, and the threshold versus smooth comparison. Section 5 calibrates the aggregate Phillips curve and compares with Gabaix (2020). Section 6 derives implications for optimal monetary policy. Section 7 concludes.

2. Theoretical Framework

This section develops a signal extraction model that generates state-dependent pass-through on both channels of the Phillips curve. I build on

Lucas (1972), extending his framework in two ways: inference weights depend on realised signals rather than only unconditional variances, and the mechanism applies to both cost shocks and competitor price signals.

2.1. Environment

A continuum of firms indexed by i produce differentiated goods. Firm i's demand depends on its price relative to competitors:

$$y_i = Y \left(\frac{P_i}{P_c} \right)^{-\eta} \tag{1}$$

where P_c is the sectoral price index. Firms face marginal cost MC_i and set prices to maximise profits subject to pricing frictions.

Two forces shape pricing decisions: (1) **Marginal costs**: The firm's static optimum is a markup over marginal cost: $P_i^* = \mu \cdot MC_i$. When costs rise, the optimal price rises. (2) **Strategic complementarities**: Because demand depends on relative prices, firms wish to align with competitors. A firm priced above competitors loses market share; a firm priced below leaves surplus on the table. This creates incentives to track competitor prices when forming expectations.

The New Keynesian Phillips Curve captures both channels:

$$\pi_t = \gamma \mathbb{E}_t[\pi_{t+1}] + \kappa \cdot mc_t \tag{2}$$

where γ governs the expectations channel (through strategic complementarities) and κ governs the cost channel. The literature typically treats both as constants. I show both are state-dependent.

2.2. The Lucas Problem

Lucas (1972) showed that agents observing local prices face a fundamental inference problem: does a price change reflect aggregate conditions affecting all markets, or factors specific to their own?

In Lucas's framework, an agent observes a signal z that combines aggregate and idiosyncratic components:

$$z = a + \varepsilon \tag{3}$$

where $a \sim N(0, \sigma_a^2)$ is the aggregate shock and $\varepsilon \sim N(0, \sigma_\varepsilon^2)$ is idiosyncratic noise. The agent cannot observe a and ε separately.

Lucas derived the optimal inference weight on the aggregate component:

$$\mathbb{E}[a|z] = \frac{\sigma_a^2}{\sigma_a^2 + \sigma_\varepsilon^2} \cdot z \tag{4}$$

A key feature of Lucas's solution is that the weight $\frac{\sigma_a^2}{\sigma_a^2+\sigma_\varepsilon^2}$ depends only on unconditional variances—it is constant regardless of the realised signal z. An agent observing a large price movement faces the same inference problem as one observing a small movement.

2.3. Extending Lucas: Signal Extraction with Unknown Signal Variance

A firm observes sectoral inflation π_c, which combines a signal component s and noise ε:

$$\pi_c = s + \varepsilon \tag{5}$$

where $\varepsilon \sim N(0, \sigma_c^2)$ is idiosyncratic noise. The firm knows σ_c^2 from historical data—it is the sector's typical volatility around aggregate conditions.

The firm's problem is to extract the signal from the noisy observation. In standard signal extraction, the optimal weight on the signal is given by the Kalman filter:

$$\lambda = \frac{\sigma_s^2}{\sigma_s^2 + \sigma_c^2} \tag{6}$$

where σ_s^2 is the variance of the signal.

The challenge is that σ_s^2 is not constant and varies over time. Some periods feature large aggregate disturbances (oil shocks, financial crises, pandemics); others are quiescent. The firm does not know *ex ante* which type of period it currently faces. It must estimate the signal environment from the realisation.

2.3.1. Estimating Signal Variance

Since s and ε are independent, the variance of the observed signal is:

$$\text{Var}(\pi_c) = \sigma_s^2 + \sigma_c^2 \tag{7}$$

For a mean-zero variable, the squared realisation is an unbiased estimator of variance. The firm can therefore estimate total variance as π_c^2 and back out signal variance:

$$\hat{\sigma}_s^2 = \pi_c^2 - \sigma_c^2 \tag{8}$$

This is the unbiased method-of-moments estimator. However, it has a fundamental problem: for small realisations where $\pi_c^2 < \sigma_c^2$, the estimator yields a negative variance—an impossibility.

This is not a technical curiosity. It reflects the core identification problem: when the observed movement is smaller than typical noise, the signal *cannot be statistically distinguished from noise*. The estimator fails precisely because identification fails.

2.3.2. The Threshold

To ensure a well-defined variance, the firm uses the squared realisation directly:

$$\hat{\sigma}_s^2 = \pi_c^2 \tag{9}$$

This is intuitive: a firm observing large price movements infers it is in a high-variance environment; a firm observing small movements cannot distinguish signal from noise.

Substituting into (6):

$$\lambda = \frac{\pi_c^2}{\pi_c^2 + \sigma_c^2} \tag{10}$$

The complementary weight is:

$$\theta(\pi_c) = 1 - \lambda = \frac{\sigma_c^2}{\sigma_c^2 + \pi_c^2} \tag{11}$$

Defining the signal-to-noise ratio as $SNR = |\pi_c|/\sigma_c$:

$$\theta = \frac{1}{1 + SNR^2} \tag{12}$$

Maximum ambiguity ($\theta = 0.5$) occurs when $SNR = 1$—that is, when $|\pi_c| = \sigma_c$. This defines the **threshold**: the boundary where inference shifts from "probably noise" to "probably signal."

The inference weight λ is continuous in $|\pi_c|$. But the threshold at $SNR = 1$ marks a qualitative shift in behaviour. Below threshold, $\lambda < 0.5$: the firm places more weight on noise than signal and optimally filters. Above threshold, $\lambda > 0.5$: the firm places more weight on signal than noise and optimally responds. The threshold is not a discontinuity in the function—it is the point where the balance of evidence shifts from "probably noise" to "probably signal"—and with it, a discrete shift in optimal pricing.

In UK data, the expenditure-weighted average of sectoral volatility is $\bar{\sigma}_c = 1.9\%$. The threshold is therefore 1.9%—derived from the data, not imposed.

2.3.3. State-Dependent Inference

Figure 1 illustrates the mechanism. The solid curve shows the weight on the signal as a function of sectoral inflation. Below the threshold (shaded region), the weight is low—firms discount the signal because it is indistinguishable from noise. Above the threshold, the weight rises toward unity—firms respond because the signal dominates.

The dashed line shows the benchmark from Lucas (1972). In Lucas's framework, the signal-to-noise ratio is a structural constant: all agents face the same inference problem regardless of the signal they observe. A sector experiencing 5% inflation faces the same uncertainty as one experiencing 0.5%.

This extension makes the effective signal-to-noise ratio endogenous to the shock size. Large realisations reveal high-signal-variance environments; small realisations reveal low-signal-variance environments. This is Bayesian learning about a time-varying state—not bounded rationality.

The key departure from Lucas is substantive: the inference weight varies with the magnitude of the signal. This generates state-dependent pass-through

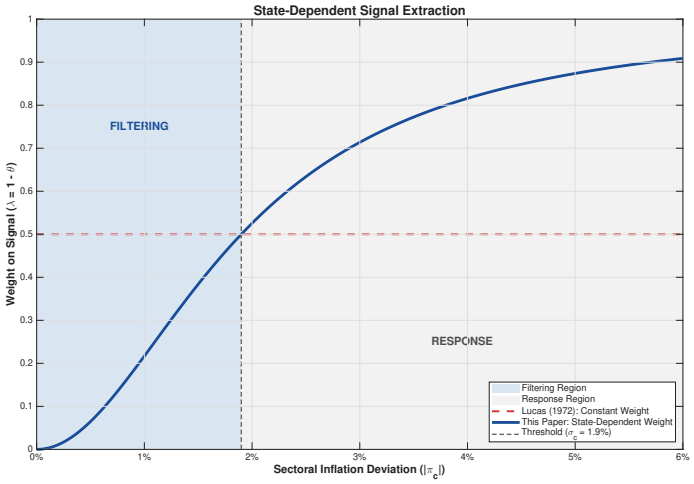


Figure 1: State-Dependent Signal Extraction

Notes: The solid curve shows the weight on the signal, $\lambda = \pi_c^2/(\pi_c^2 + \sigma_c^2)$, as a function of sectoral inflation. The dashed line shows the constant weight implied by Lucas (1972). Below the threshold ($|\pi_c| < \sigma_c = 1.9\%$), signals are indistinguishable from noise and firms filter. Above threshold, signals dominate and firms respond.

and, after aggregation, a state-dependent Phillips curve.

2.4. The Horizon Filter

Inference alone does not determine pricing behaviour. A firm that believes a shock is idiosyncratic must still decide how to respond. The answer depends on the firm's pricing horizon.

definition 1 (Horizon Filter)**.** *Let* $\Lambda(\Phi)$ *denote the weight a firm with pricing horizon* Φ *places on current idiosyncratic movements, accounting for expected dynamics over the pricing spell.*

For a flexible firm that re-optimises every period ($\Phi = 1$), $\Lambda \approx 1$. Current movements matter fully because the firm can correct any error next period. For a sticky firm with horizon $\Phi > 1$, Λ depends on the expected path of competitor prices over the pricing spell.

To build intuition, consider idiosyncratic shocks that follow $\varepsilon_{t+1} = \rho_\varepsilon \varepsilon_t + u_t$

with $0 < \rho_\varepsilon < 1$. The firm cares about the average level over its pricing spell:

$$\mathbb{E}_t \left[\frac{1}{\Phi} \sum_{k=0}^{\Phi-1} \varepsilon_{t+k} \right] = \frac{1 - \rho_\varepsilon^\Phi}{\Phi(1 - \rho_\varepsilon)} \cdot \varepsilon_t \tag{13}$$

With positive autocorrelation, the horizon filter $\Lambda = \frac{1-\rho_\varepsilon^\Phi}{\Phi(1-\rho_\varepsilon)}$ lies between zero and one. Transitory movements partially wash out over the pricing horizon, so responding fully today creates mispricing. The filter is decreasing in Φ—stickier firms place less weight on current movements because more of the shock decays before their next adjustment.

But the horizon filter can also be negative. This occurs when $\rho_\varepsilon < 0$—when shocks are expected to *overshoot* rather than merely revert. The mechanism operates through observation of flexible-price competitors. Flexible firms can re-optimise quickly; when they discover a shock was idiosyncratic, they reverse—and may overshoot to recover lost market share. A sticky firm, locked into its price for multiple periods, observes these flexible competitors and anticipates reversal. If flexible prices are expected to overshoot, the sticky firm optimally *leans against* current movements to maintain competitive position over its pricing spell. This implies $\Lambda < 0$.

The model remains agnostic about the sign of Λ and identifies it empirically. This generates a testable prediction: negative pass-through should be strongest for sticky firms facing flexible competitors, and should vanish above threshold when signals clarify. Section 3.6 tests this directly.

2.5. Optimal Pass-Through

A firm observing a signal—whether competitor prices or upstream costs—faces inference uncertainty: is this persistent or transitory? The optimal response combines the probability the shock is noise (θ) with the horizon filter (Λ):

$$\rho(\theta, \Lambda) = \theta \cdot \Lambda + (1 - \theta) \cdot 1 = 1 - \theta(1 - \Lambda) \tag{14}$$

Below threshold ($\theta \to 1$), pass-through approaches Λ. Above threshold ($\theta \to 0$), pass-through approaches unity. The regime shift is common to

both channels. But the two channels differ in Λ—and this generates distinct predictions.

Expectations channel. Competitor prices are endogenous. Idiosyncratic movements revert as firms correct errors—and may overshoot as they recover market share. Sticky firms, anticipating reversal, filter aggressively (Λ small or negative). Flexible firms can wait and see ($\Lambda \approx 1$). The gap is large below threshold and closes above as signals clarify.

Cost channel. Upstream costs are exogenous. Once identified as real, cost shocks persist—commodity prices follow near-random walks, supply chain disruptions do not self-correct. Both sticky and flexible firms expect persistence conditional on a real shock ($\Lambda_{mc} \approx 1$ for both). The horizon wedge collapses. State-dependence operates through θ alone, shifting all firms equally.

Figure 2 summarises. The left panel shows the expectations channel: sticky firms (dashed) filter below threshold and converge to flexible firms (solid) above—the gap closes. The right panel shows the cost channel: both firm types increase pass-through above threshold, but the gap remains constant—a parallel shift.

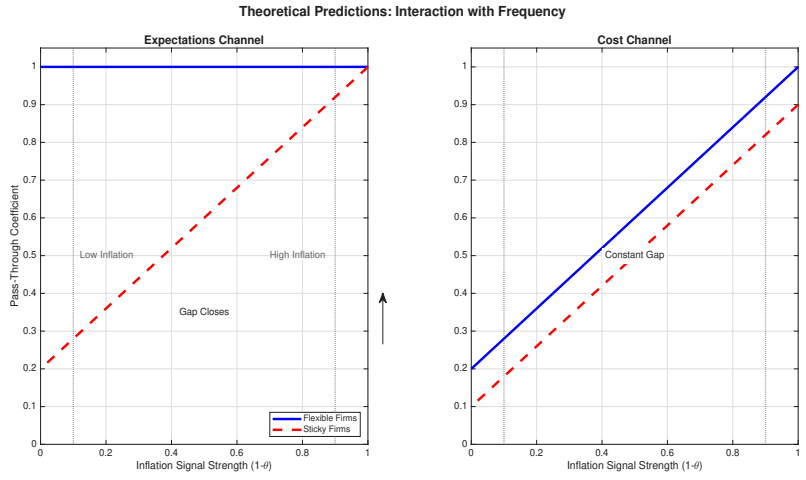


Figure 2: Theoretical Predictions: Convergence vs Parallel Shift
Notes: Left panel: expectations channel. Sticky firms (dashed) converge to flexible firms (solid) above threshold. Right panel: cost channel. Both firm types shift upward equally—the gap remains constant.

The two channels thus provide a joint test of the mechanism. If state-

dependence reflects signal extraction, we should observe convergence on expectations and parallel shift on costs. If it reflects menu costs or attention, both channels should show the same pattern.

2.6. The Aggregate Phillips Curve

Aggregating with standard Calvo pricing yields a state-dependent Phillips curve:

$$\pi_t = \gamma(\theta)\mathbb{E}_t[\pi_{t+1}] + \kappa(\theta) \cdot mc_t \tag{15}$$

where $\gamma(\theta) = \beta[1-(1-\omega)(1-\rho(\theta))]$, $\kappa(\theta) = \kappa_0 \cdot \rho_{mc}(\theta)$ and $\kappa_0 = {}^{(1-\omega)(1-\beta\omega)}/_{\omega}$, where $1-\omega$ is the probability of price adjustment.

Below threshold, both coefficients are attenuated—firms discount expectations and absorb costs. Above threshold, both approach the standard Calvo benchmark.

This provides a micro-foundation for Gabaix (2020)'s cognitive discounting. His constant $\bar{M} \approx 0.85$ reflects the below-threshold regime. In my framework, the discount is state-dependent: it binds when inflation is low and vanishes when inflation rises. What appears as a behavioural primitive may be the low-inflation limit of a rational model.

3. Empirical Evidence

3.1. CPI Micro-Data

I use UK CPI micro-price data compiled by Davies (2021), covering 2003–2021. The dataset encompasses approximately 60% of the CPI basket, with over 6 million price observations across 335,000 unique products at quarterly frequency. I exclude temporary sales to focus on fundamental pricing decisions and winsorise extreme price changes at the 1st and 99th percentiles.[1]

[1]The correlation between micro-data implied inflation and official data is 0.60 excluding sales and 0.39 including sales.

The dependent variable is the non-sale price change ($\Delta p_{i,t}$). I construct aggregate and sectoral inflation as:

$$\pi_t = \sum_i \left(\pi_{i,c,t} \times \frac{w_{i,t}}{\sum_j w_{j,t}} \right) \tag{16}$$

$$\pi_{c,t} = \sum_{i \in c} \left(\pi_{i,c,t} \times \frac{w_{i,t}}{\sum_{j \in c} w_{j,t}} \right) \tag{17}$$

where $w_{i,t}$ is the expenditure weight of product i and c denotes the COICOP category. Aggregate inflation π_t captures economy-wide price movements; sectoral inflation $\pi_{c,t}$ isolates category-specific dynamics.

3.1.1. Measuring Competitor Price Movements

Each firm's competitor price change is the expenditure-weighted mean of all other products in the same COICOP category:

$$\Delta \bar{p}_{-i,c,t} = \frac{\sum_{j \in c, j \neq i} w_{j,t} \cdot \Delta p_{j,c,t}}{\sum_{j \in c, j \neq i} w_{j,t}} \tag{18}$$

The subscript $-i$ indicates exclusion of product i, avoiding mechanical correlation with its own price. I use lagged values ($\Delta \bar{p}_{-i,c,t-1}$) to mitigate simultaneity.

3.2. Upstream Cost Index

I construct upstream cost indices by combining Producer Price Indices (PPI) with Input-Output coefficients to capture exogenous supply-side variation in firms' marginal costs. The PPI data are obtained from the ONS MM22 dataset, providing quarterly price indices for 24 manufacturing industries at the 2-digit SIC level (C10–C33). The Input-Output coefficients come from the 2023 UK Supply and Use Tables, which map inter-industry purchases across the economy.

For each retail product category j and quarter t, the upstream cost index

is:

$$C_{jt} = \sum_{i \in \text{Manufacturing}} \omega_{ij} \cdot \Delta \ln PPI_{it} \tag{19}$$

where ω_{ij} is the share of industry j's intermediate inputs sourced from upstream industry i, and $\Delta \ln PPI_{it}$ is the quarterly log change in producer prices for industry i. The weights are fixed from the 2023 I-O tables; all time-series variation comes from PPI changes.

A COICOP-SIC concordance maps each consumption category to its primary supplier industry (e.g., Food → C10 Food manufacturing, Clothing → C14 Wearing apparel). This construction provides plausibly exogenous variation: upstream PPI changes reflect supply conditions in manufacturing—input costs, productivity, capacity utilisation—rather than downstream retail demand. Details on data sources and construction are in Appendix Appendix A.

3.3. The 1.9% Threshold

Figure 3 plots the signal-to-noise ratio against sectoral inflation. The ratio crosses unity at approximately $\pm 1.9\%$, reflecting average sectoral volatility ($\bar{\sigma}_c = 1.9\%$). Below this threshold, idiosyncratic noise dominates; above it, common shocks become distinguishable. The symmetry at $\pm 1.9\%$ confirms this is an information boundary, not an artefact of inflation direction.

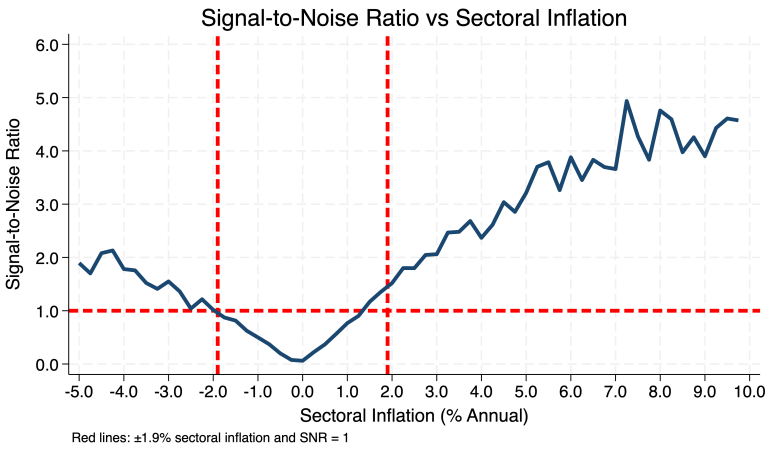


Figure 3: Signal-to-Noise Ratio and the Information Threshold

Notes: Signal-to-noise ratio (absolute sectoral inflation divided by sector-specific volatility σ_c) averaged across sectors. SNR crosses unity at approximately ±1.9% (vertical lines).

This threshold applies broadly because low-volatility sectors—where $\sigma_c \approx 1.9\%$—account for 90% of expenditure. High-volatility sectors have higher thresholds but adjust frequently regardless.

3.4. Descriptive Statistics

Table 1 reports summary statistics by inflation regime. The sample splits roughly evenly: 56% below the 1.9% threshold, 44% above.

Table 1: Summary Statistics by Inflation Regime

	Below 1.9%	Above 1.9%	Overall
Panel A: Price Adjustment			
Adjustment frequency	0.295	0.294	0.295
Mean price change (%)	0.29	0.66	0.45
Panel B: Signals			
Competitor price change (%)	−0.04	0.13	0.03
Sectoral inflation (%)	0.25	0.57	0.39
Aggregate inflation (%)	0.42	0.49	0.45
Panel C: Costs			
Cost index (%)	0.12	0.18	0.14
Panel D: Sample			
Observations (millions)	3.58	2.82	6.40
Share of sample (%)	55.9	44.1	100

Notes: All rates are quarterly. The 1.9% threshold is based on annualised sectoral inflation. Adjustment frequency is the fraction of observations with non-zero price changes.

Three patterns emerge. First, adjustment frequency is identical across regimes: 29.5% below threshold versus 29.4% above. This rules out menu cost explanations, which predict higher frequency when inflation rises.

Second, conditional on adjusting, price changes are larger above threshold: 0.66% versus 0.29%. Pass-through operates through the intensive margin—firms do not adjust more often, they adjust by more.

Third, the informational environment differs. Competitor price changes average −0.04% below threshold but +0.13% above. Price signals are more informative when sectoral inflation is high.

3.5. Regression Results

The model predicts convergence for the expectations channel (sticky firms catch up to flexible firms above threshold) but parallel shifts for costs (all firms respond more, gap remains constant). I test both predictions jointly.

3.5.1. The Joint Test

As I have discussed, signal extraction generates opposite predictions for the two Phillips curve channels (see Figure 2). The expectations channel should show **convergence**: sticky firms filter competitor prices below the threshold because they anticipate reversion; above the threshold, signals clarify and the gap closes. The cost channel should show **parallel shift**: upstream costs are exogenous and persistent, so both firm types respond equally when the threshold is crossed.

I test both predictions in a single regression:

$$\begin{aligned}\Delta p_{i,t} = \alpha_i &+ \beta_1 \Delta \bar{p}_{-i,t-1} + \beta_2 \Delta \bar{p}_{-i,t-1} \times T + \beta_3 \Delta \bar{p}_{-i,t-1} \times F + \beta_4 \Delta \bar{p}_{-i,t-1} \times T \times F \\ &+ \gamma_1 C_{t-1} + \gamma_2 C_{t-1} \times T + \gamma_3 C_{t-1} \times F + \gamma_4 C_{t-1} \times T \times F \\ &+ \delta_1 \pi_{c,t-1} + \delta_2 \pi_t + \delta_3 \Delta p_{i,t-1} + \varepsilon_{i,t}\end{aligned} \tag{20}$$

where $\Delta p_{i,t}$ is the quarterly price change for product i, $\Delta \bar{p}_{-i,t-1}$ is the lagged average price change of competitors in the same category, C_{t-1} is the upstream cost index constructed from input-output tables and producer prices, T is an indicator equal to one when annualised sectoral inflation exceeds 1.9%, F is the product's historical adjustment frequency, $\pi_{c,t-1}$ is lagged sectoral inflation, π_t is aggregate inflation, and α_i are product fixed effects. Standard errors are clustered by product.

The key coefficients are β_4 (expectations channel) and γ_4 (cost channel). Signal extraction predicts $\beta_4 < 0$ (convergence) and $\gamma_4 = 0$ (parallel shift).

Table 2 reports the results.

Table 2: Joint Test: Expectations and Cost Channels

	Expectations (Competitor Price)	Costs (Cost Index)
Base (β_1, γ_1)	-1.002^{***}	-0.203^{***}
	(0.050)	(0.043)
× Threshold (β_2, γ_2)	0.979^{***}	0.226^{***}
	(0.061)	(0.056)
× Frequency (β_3, γ_3)	1.413^{***}	1.449^{***}
	(0.066)	(0.159)
× Threshold × Frequency (β_4, γ_4)	-1.398^{***}	-0.219
	(0.094)	(0.201)
Lagged Price Change	-0.270^{***}	
	(0.001)	
Aggregate Inflation	0.953^{***}	
	(0.014)	
Sectoral Inflation	0.312^{***}	
	(0.006)	
Asymmetry test: $\beta_4 = \gamma_4$		
F-statistic	26.89	
p-value	< 0.001	
Observations	5,680,862	
R-squared (within)	0.073	

Notes: The dependent variable is the quarterly price change (percent). Both channels estimated jointly. All specifications include product fixed effects. Standard errors clustered by product in parentheses. Threshold equals one when annualised sectoral inflation exceeds 1.9%. Frequency is the product's historical adjustment frequency. Cost Index is the upstream PPI-weighted cost shock. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

The results confirm both predictions. The expectations triple interaction is large and significant ($\beta_4 = -1.40$, $p < 0.001$): sticky firms converge toward flexible firms above threshold. The cost triple interaction is small and insignificant ($\gamma_4 = -0.22$, $p = 0.28$): no convergence. A formal test rejects equality

($F = 26.89$, $p < 0.001$). Only signal extraction generates this asymmetry.[2]

3.5.2. Interpreting the Results

Table 3 translates these coefficients into pass-through rates for sticky firms (frequency = 0.20) and flexible firms (frequency = 0.60).

Table 3: Pass-Through by Firm Type and Regime

	Below 1.9%	Above 1.9%	Difference
Panel A: Expectations Channel			
Sticky firms (freq = 0.20)	−0.72	−0.02	+0.70
	(0.039)	(0.023)	
Flexible firms (freq = 0.60)	−0.15	−0.01	+0.14
	(0.028)	(0.021)	
Gap (Flexible − Sticky)	0.57	0.01	
Convergence			**98.2%**
Panel B: Cost Channel			
Sticky firms (freq = 0.20)	0.09	0.27	+0.18
	(0.021)	(0.022)	
Flexible firms (freq = 0.60)	0.67	0.76	+0.09
	(0.062)	(0.055)	
Gap (Flexible − Sticky)	0.58	0.49	
Convergence			15.5%

Notes: Pass-through rates computed from Table 2. Standard errors in parentheses.

Expectations channel. Below threshold, sticky firms exhibit negative pass-through (−0.72)—they lean against competitor movements. The gap between firm types is 0.57. Above threshold, the gap collapses to 0.01: 98% convergence. Figure 4 illustrates this pattern.

[2]Restricting to products observed throughout the sample period strengthens the asymmetry: the expectations triple interaction is −1.76 (s.e. = 0.12) while the cost triple interaction is 0.02 (s.e. = 0.28). The formal test strongly rejects equality ($F = 32.30$, $p < 0.001$).

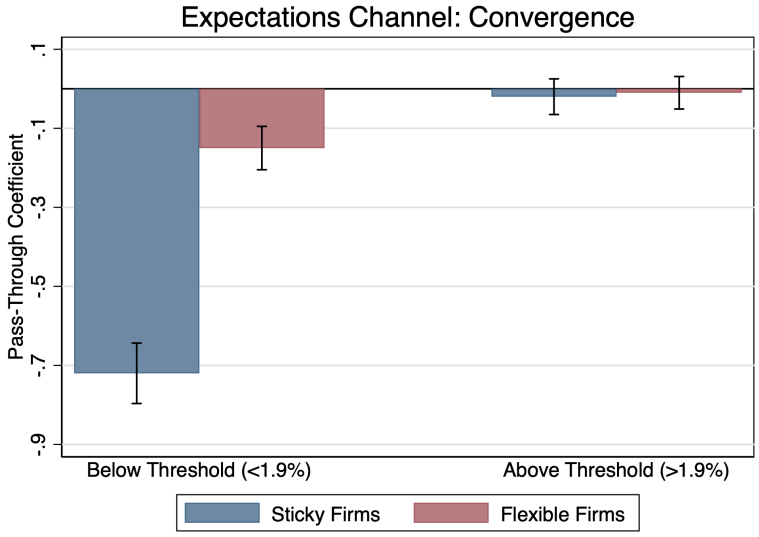


Figure 4: Expectations Channel: Convergence

Notes: Pass-through coefficients for sticky firms (frequency = 0.20) and flexible firms (frequency = 0.60) computed from Table 2. Error bars show 95% confidence intervals. Below threshold, sticky firms lean against competitor prices (-0.72); flexible firms show weaker negative pass-through (-0.15). Above threshold, both converge to near zero. Gap closes from 0.57 to 0.01: 98% convergence.

Cost channel. Both firm types increase pass-through by similar amounts (+0.18 for sticky, +0.09 for flexible). The gap remains nearly constant (0.58 below, 0.49 above): only 15% convergence, statistically indistinguishable from zero given that γ_4 is insignificant. Figure 5 shows the parallel shift.

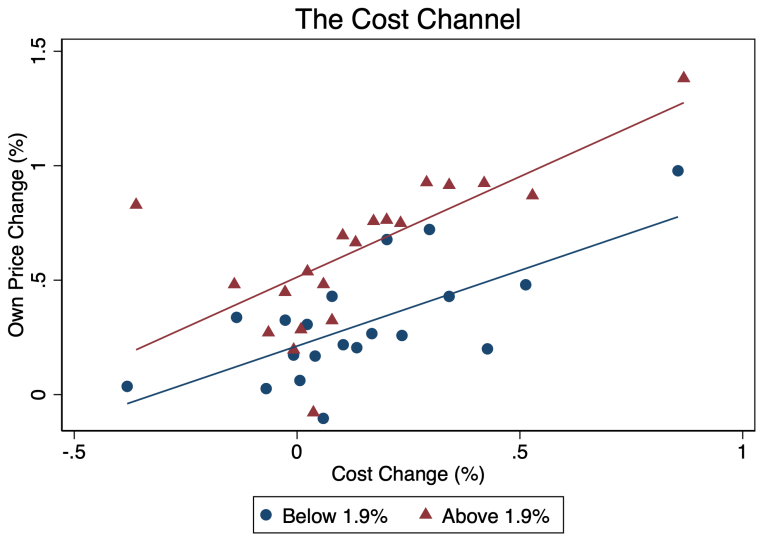


Figure 5: Cost Channel: Parallel Shift

Notes: Binned scatter plot of own price change against upstream cost change by regime. Each point represents the mean within a ventile. Lines show linear fit. Both regimes exhibit similar slopes—pass-through increases above threshold but the gap between firm types remains constant.

This parallel shift contrasts with the convergence observed in the expectations channel.

3.6. The Mechanism: Competitor Flexibility

The negative pass-through coefficient below threshold requires explanation. Standard Bayesian signal extraction generates weights in $[0, 1]$—firms discount noisy signals, but do not lean against them. Yet the estimates show $\rho < 0$ for sticky firms below threshold. What generates negative pass-through?

The answer combines two mechanisms: signal extraction determines *beliefs*; competitor flexibility determines the *response* to those beliefs.

Signal extraction (belief formation). Below threshold, signal-to-noise is low. Firms attribute competitor price movements to idiosyncratic factors and expect mean reversion. Above threshold, signals are clear. Firms recognise movements as reflecting common shocks—no reversion expected.

Competitor flexibility (response to beliefs). If competitors are

flexible-price firms, they can quickly correct transitory movements—and may overshoot to recover lost market share. A sticky firm, anticipating reversal, optimally leans against current movements. If competitors are also sticky, no quick reversal is possible, so no leaning against.

The key test is the **four-way interaction**: does the flexible-competitor effect disappear above threshold? If signal extraction drives the regime switch, then below threshold flexible competitors lead to expected reversal and leaning against; above threshold the common shock is recognised and even flexible competitors won't reverse, so no leaning. Without signal extraction, the flexible-competitor effect would be constant across regimes.

I construct the average adjustment frequency of competitors in each product category:

$$\text{CompetitorFreq}_{i,c,t} = \frac{\sum_{j \in c, j \neq i} w_{j,t} \cdot \text{Freq}_j}{\sum_{j \in c, j \neq i} w_{j,t}} \tag{21}$$

where Freq_j is the adjustment frequency of product j and weights are expenditure shares.

Table 4 reports results from interacting competitor prices with own frequency, competitor frequency, and the threshold.

Table 4: Pass-Through and Competitor Flexibility

	Coefficient
Base effects	
Competitor price	0.38***
	(0.11)
× Own frequency	−0.68**
	(0.30)
× Competitor frequency	−3.74***
	(0.34)
× Own freq × Comp freq	4.61***
	(0.40)
Threshold interactions	
× Threshold	−0.17
	(0.13)
× Threshold × Own freq	0.51
	(0.39)
× Threshold × Comp freq	2.97***
	(0.42)
× Threshold × Own freq × Comp freq	−3.87***
	(0.49)
Observations	5,798,850

Notes: Fixed effects panel regression with product fixed effects. Includes lagged dependent variable, aggregate inflation, and sectoral inflation. Standard errors clustered by product in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

The four-way interaction is large and highly significant (-3.87, s.e. $= 0.49$, p<0.001). This confirms that the flexible-competitor effect weakens above threshold—exactly as predicted by signal extraction.

Tables 5 and 6 translate the coefficients into pass-through rates across the full distribution of own and competitor flexibility.

Table 5: Pass-Through Below Threshold

	Competitor Frequency		
Own Frequency	Sticky (0.2)	Medium (0.4)	Flexible (0.6)
Sticky (0.2)	−0.32	−0.89	−1.45
Medium (0.4)	−0.27	−0.65	−1.03
Flexible (0.6)	−0.23	−0.42	−0.62

Notes: Pass-through computed from Table 4. Below threshold (annualised sectoral inflation $< 1.9\%$).

Table 6: Pass-Through Above Threshold

	Competitor Frequency		
Own Frequency	Sticky (0.2)	Medium (0.4)	Flexible (0.6)
Sticky (0.2)	+0.05	−0.08	−0.20
Medium (0.4)	+0.05	−0.05	−0.15
Flexible (0.6)	+0.04	−0.03	−0.09

Notes: Pass-through computed from Table 4. Above threshold (annualised sectoral inflation $\geq 1.9\%$).

Three patterns emerge. First, negative pass-through concentrates where expected: below threshold, the most negative pass-through (−1.45) occurs for sticky firms facing flexible competitors—the upper-right cell of Table 5. These firms cannot quickly re-optimise but face competitors who can. They anticipate reversal and lean against. Second, convergence above threshold: all cells converge toward zero. The range compresses from [−1.45, −0.23] below threshold to [−0.20, +0.05] above. Firms stop leaning against competitors because they recognise movements as persistent. Third, the largest convergence occurs for the most affected: the sticky-vs-flexible cell shows the largest change, from −1.45 to −0.20, an improvement of 1.25.

Do flexible prices actually revert? The negative pass-through assumes sticky firms correctly anticipate that flexible competitors will reverse. Table 7 tests this directly by estimating AR(1) persistence of sectoral price changes by frequency quartile and regime.

Table 7: Price Persistence by Firm Type and Regime

	Q1 Sticky	Q2 Med-Low	Q3 Med-High	Q4 Flexible
Below 1.9%	0.37***	0.26*	0.32***	−0.26
	(0.12)	(0.14)	(0.10)	(0.20)
Above 1.9%	0.46***	0.32**	0.30**	0.00
	(0.08)	(0.15)	(0.13)	(0.08)
Change	+0.09	+0.06	−0.02	**+0.26**

Notes: AR(1) coefficients from regressing sectoral price changes on lagged values by frequency quartile. Standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Flexible-price sectors (Q4) exhibit mean reversion below threshold (AR(1) $= -0.26$) but not above (AR(1) $= 0.00$). Sticky firms are right to lean against flexible competitors below threshold—those prices actually reverse. Above threshold, flexible prices persist, so leaning against is no longer optimal.

The four-way interaction confirms that signal extraction and competitor flexibility work together. Signal extraction determines whether firms expect reversal; competitor flexibility determines the magnitude of expected reversal; the interaction generates the regime switch. Below threshold, sticky firms facing flexible competitors show the most negative pass-through (-1.45). Above threshold, the signal clears—even flexible competitors won't reverse a persistent common shock—and pass-through converges toward zero (-0.20).

4. Threshold vs Smooth Nonlinearity

A natural concern is whether the 1.9% threshold reflects a genuine discrete shift or merely smooth nonlinearity. The critical test: what happens to sticky firms (frequency $= 0.2$) when sectoral inflation moves from 1.8% to 2.0%? The threshold model predicts a discrete jump; smooth alternatives predict gradual adjustment spread across the entire inflation range.

Table 8: Threshold vs Smooth: Behavioural Change at 1.9%

Channel	Threshold Model	Linear Model	% Discrete
Expectations	0.70	0.001	99.8%
Costs	0.18	0.007	96.4%

Notes: Predicted change in pass-through for sticky firms (frequency = 0.2) when sectoral inflation crosses 1.9%. Threshold model computed from Table 2: $\beta_2 + \beta_4 \times 0.20$ for expectations, $\gamma_2 + \gamma_4 \times 0.20$ for costs. Linear model captures change from 1.8% to 2.0%—a 0.2pp step on a smooth gradient.

The threshold model captures nearly all the behavioural change. Smooth models miss it entirely. Smooth models fail because they spread behavioural change across the entire inflation range. The threshold model concentrates all change at 1.9%, capturing the discrete shift from filtering to responding. This is precisely what signal extraction predicts: at $|\pi_c| = \sigma_c$, inference flips from "probably noise" to "probably signal."

5. From Micro Pass-Through to Aggregate Dynamics

Standard Calvo aggregation with incomplete pass-through yields state-dependent Phillips curve coefficients:

$$\gamma(\theta) = \beta \left[1 - (1 - \omega)(1 - \rho(\theta))\right] \tag{22}$$

$$\kappa(\theta) = \kappa_0 \cdot \rho_{mc}(\theta) \tag{23}$$

where $\rho(\theta)$ is pass-through from competitor prices, $\rho_{mc}(\theta)$ is pass-through from costs, and $\kappa_0 = \frac{(1-\omega)(1-\beta\omega)}{\omega}$ is the standard Calvo slope.

When $\rho = 1$, firms fully incorporate sectoral conditions and $\gamma \to \beta$. When $\rho < 1$, firms discount sectoral information and $\gamma < \beta$. Negative pass-through ($\rho < 0$) implies firms lean against sectoral movements, generating even larger discounting.

A key distinction arises between empirical estimation and calibration. Above threshold, the coefficient on competitor prices converges to zero in the regressions—not because firms ignore aggregate conditions, but because they respond through the sectoral inflation control. The information friction has

vanished. Therefore, I use empirical estimates below threshold and the rational benchmark ($\rho = 1$, $\rho_{mc} = 1$) above threshold.

I calibrate using $\beta = 0.99$, $\omega = 0.70$, yielding $\kappa_0 = 0.13$. Pass-through is evaluated at mean adjustment frequency (0.30). From Table 2, below threshold expectations pass-through is -0.58 and cost pass-through is 0.23.

Expectations channel. Below threshold ($\rho = -0.58$):

$$\gamma_{below} = 0.99 \times [1 - 0.30 \times (1 - (-0.58))] = 0.52 \tag{24}$$

Above threshold ($\rho \to 1$): $\gamma_{above} = 0.99$.

Cost channel. Below threshold ($\rho_{mc} = 0.23$):

$$\kappa_{below} = 0.13 \times 0.23 = 0.03 \tag{25}$$

Above threshold ($\rho_{mc} \to 1$): $\kappa_{above} = 0.13$.

Table 9: Calibrated Phillips Curve Coefficients

	Expectations (γ)		Costs (κ)	
	Pass-through	Coefficient	Pass-through	Coefficient
Below 1.9%	-0.58	0.52	0.23	0.03
Above 1.9%	1.00	0.99	1.00	0.13
Ratio		1.9×		4.3×

Notes: Calibration uses $\beta = 0.99$, $\omega = 0.70$, $\kappa_0 = 0.13$. Pass-through evaluated at mean frequency (0.30) from Table 2. Below-threshold values from empirical estimates; above-threshold uses rational benchmark.

The state-dependence is substantial. The expectations coefficient nearly doubles; the cost coefficient quadruples.

5.1. Comparison with Gabaix (2020)

Gabaix (2020) proposes cognitive discounting with constant $\bar{M} \approx 0.85$, yielding $\gamma \approx 0.76$. My framework provides a micro-foundation: Gabaix's $\bar{M}$ corresponds to $1-(1-\omega)(1-\rho)$. The "cognitive discount" arises from rational signal extraction when signal-to-noise is low.

Crucially, my discount is state-dependent. Gabaix calibrates on post-1984 U.S. data—predominantly below 2% inflation. His estimate reflects a weighted average. In my sample (56% below threshold):

$$\gamma_{avg} = 0.56 \times 0.52 + 0.44 \times 0.99 = 0.73 \tag{26}$$

Table 10: Comparison with Gabaix (2020)

	Standard NK	Gabaix (2020)	This Paper
Below 2%	0.99	0.76	0.52
Above 2%	0.99	0.76	0.99
Weighted average	0.99	0.76	0.73
State-dependent?	No	No	Yes

Notes: Coefficient on expected future inflation (γ). Weighted average uses sample distribution: 56% below threshold, 44% above.

What appears as a fixed behavioural parameter is state-dependent. Below 2%, the economy is in the filtering regime ($\gamma = 0.52$). Above 2%, it returns to rationality ($\gamma = 0.99$). Gabaix's cognitive discount is not a primitive—it is the low-inflation regime of a rational model.

6. Optimal Monetary Policy

The state-dependent Phillips curve has direct implications for optimal policy. I derive the optimal targeting rule under commitment and the determinacy condition.

6.1. Optimal Targeting Rule

The central bank minimises:

$$\mathcal{L} = \frac{1}{2}\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left[\pi_t^2 + \alpha x_t^2\right] \tag{27}$$

subject to the state-dependent Phillips curve. Solving from the timeless perspective (Woodford, 2003) yields:

$$\pi_t = -\frac{\alpha}{\kappa}\left(x_t - \frac{\gamma}{\beta}x_{t-1}\right) \tag{28}$$

Two coefficients govern optimal policy. The ratio γ/β determines policy inertia—how much current targets depend on past conditions. The ratio α/κ determines how aggressively the central bank trades off output gaps against inflation.

Table 11: Optimal Policy Under Commitment

	Below Threshold	Above Threshold	Ratio
Expectations coefficient (γ)	0.52	0.99	1.9×
Cost coefficient (κ)	0.03	0.13	4.3×
Policy inertia (γ/β)	0.53	1.00	1.9×
Inflation weight (α/κ)	33α	7.7α	4.3×

Notes: Optimal targeting rule: $\pi_t = -(\alpha/\kappa)(x_t - (\gamma/\beta)x_{t-1})$. Parameters from Table 9.

Below threshold, optimal policy is aggressive: the central bank responds primarily to current conditions ($\gamma/\beta = 0.53$). Above threshold, policy must be gradualist: current targets depend one-for-one on past output gaps ($\gamma/\beta = 1.00$). When γ is high, firms incorporate expected future policy immediately, so deviations from a smooth path create expectational disruption.

6.2. The State-Dependent Taylor Principle

Consider a Taylor rule $i_t = r^n + \phi_\pi \pi_t + \phi_x x_t$. For a unique bounded equilibrium:

$$\phi_\pi + \frac{1-\gamma}{\kappa}\phi_x > 1 \tag{29}$$

The condition is state-dependent through both γ and κ:

Table 12: State-Dependent Taylor Principle

	Below Threshold	Above Threshold
Determinacy condition	$\phi_\pi + 16.0\phi_x > 1$	$\phi_\pi + 0.08\phi_x > 1$
Minimum ϕ_π if $\phi_x = 0$	1.0	1.0
Minimum ϕ_π if $\phi_x = 0.5$	-7.0	0.96

Notes: Below threshold: $(1 - 0.52)/0.03 = 16.0$. Above threshold: $(1 - 0.99)/0.13 = 0.08$.

Below threshold, stabilising output substantially contributes to inflation stabilisation. Even modest output gap response ($\phi_x = 0.5$) allows the central bank to *decrease* rates in response to inflation while maintaining determinacy. Above threshold, the coefficient is tiny (0.08): the central bank must respond nearly one-for-one to inflation itself.

6.3. Determinacy and the "Looking Through" Strategy

Nakamura et al. (2025) document that the Federal Reserve partially "looked through" post-Covid inflation, achieving an "immaculate disinflation" without a severe recession. They argue that this was only possible due to established credibility. My results identify the precise mechanism behind this success.

The argument connects closely to Galí (2025), who proposes monetary policy rules expressed in terms of the long-term real rate. The link between these views lies in the stability conditions of the economy. As emphasized by Woodford (2003), the condition for determinate equilibrium depends on the long-run properties of the system—specifically, the Phillips curve parameters in the steady state to which the economy is expected to return:

$$\phi_\pi + \frac{1 - \gamma_{\text{LR}}}{\kappa_{\text{LR}}}\phi_x > 1 \tag{30}$$

Crucially, the determinacy condition reflects the *expected future regime*, not necessarily the current one. This distinction is vital in a state-dependent economy.

Consider a central bank facing inflation above the threshold today. In the current high-inflation regime, the Phillips curve is steep ($\gamma_H = 0.99$, $\kappa_H = 0.13$), and the determinacy condition is tight: $\phi_\pi + 0.08\phi_x > 1$. A standard short-rate Taylor rule would face a stringent constraint to ensure stability. However, the long-term real rate depends on the entire expected path of future short rates and inflation dynamics.

If the central bank is credible, agents expect inflation to eventually return below the threshold. Consequently, long-run expectations are anchored in the low-inflation regime ($\gamma_L = 0.52$, $\kappa_L = 0.03$). The effective determinacy condition becomes $\phi_\pi + 16.0\phi_x > 1$—dramatically relaxed. Credibility allows the central bank to borrow the stability properties of the low-inflation regime. Because agents anticipate a return to filtering, they do not unmoor long-run expectations.

This provides the micro-foundation for the "immaculate disinflation" of Nakamura et al. (2025). A credible central bank can "look through" a temporary spike above the threshold because the private sector anticipates the restoration of the filtering regime. An incredible central bank faces the opposite dynamic: agents expect inflation to persist, locking expectations into the high-γ regime where the Phillips curve is steep and immediate aggressive action is required.

The state-dependent Phillips curve thus clarifies the mechanics of credibility. Credibility determines which regime governs long-run expectations—and therefore which γ and κ enter the effective determinacy condition. The 2% threshold is not arbitrary; it is the boundary where this mechanism switches on, and credibility is the force that ensures the switch is temporary.

7. Conclusions

This paper shows that the Phillips curve is state-dependent: flat when inflation is low, steep when inflation is high. The mechanism is signal extraction. When sectoral inflation is below 1.9%, price movements are indistinguishable from idiosyncratic noise—firms filter. When inflation exceeds 1.9%, signals

clarify—firms respond. The threshold emerges from the data as the point where signal-to-noise equals unity.

The empirical evidence confirms two distinct channels. On the expectations channel, sticky firms converge to flexible firms above threshold: the gap in pass-through closes by 98%. On the cost channel, all firms increase pass-through equally: the gap remains constant. Same threshold, different mechanisms—exactly as signal extraction predicts.

The findings provide a micro-foundation for Gabaix (2020)'s cognitive discounting. What appears as a constant behavioural parameter is in fact state-dependent. The "cognitive discount" is not a primitive—it is the low-inflation regime of a rational model. Below 2%, firms filter and the Phillips curve is flat ($\gamma = 0.52$, $\kappa = 0.03$). Above 2%, filtering stops and the curve steepens ($\gamma = 0.99$, $\kappa = 0.13$).

The implications for monetary policy are substantial. Optimal policy is state-dependent: aggressive below threshold, gradualist above. The Taylor principle is state-dependent: output stabilisation contributes to inflation stabilisation below threshold but not above. And the costs of crossing the threshold are asymmetric: below threshold, policy errors self-correct through filtering; above threshold, errors compound through expectations.

The 2% target is not merely a number to stabilise around—it is the boundary where price-setting behaviour changes, where credibility is tested, and where the feasibility of 'looking through' is determined.

Appendix A. Upstream Cost Index Construction

This appendix provides detailed documentation of the upstream cost index construction. The index combines three data sources: (i) Producer Price Indices from the ONS, (ii) Input-Output tables from the UK Supply and Use Tables, and (iii) a concordance mapping retail product categories to manufacturing industries.

Appendix A.1. Producer Price Index Data

I obtain monthly Producer Price Indices (PPI) from the ONS MM22 dataset. The PPI measures prices received by UK producers for goods sold on the domestic market, providing a direct measure of upstream cost pressures facing downstream retailers.

The data cover 24 manufacturing industries at the 2-digit SIC level (C10–C33), listed in Table A.13. I use the "Output Domestic, excluding Duty" series to capture producer costs net of excise taxes. Where duty-inclusive and duty-exclusive versions exist (e.g., beverages, tobacco, petroleum), I use the excluding-duty version.

Table A.13: PPI Industry Coverage

SIC	Industry	SIC	Industry
C10	Food products	C22	Rubber and plastics
C11	Beverages	C23	Non-metallic minerals
C12	Tobacco	C24	Basic metals
C13	Textiles	C25	Fabricated metals
C14	Wearing apparel	C26	Electronics
C15	Leather products	C27	Electrical equipment
C16	Wood products	C28	Machinery
C17	Paper products	C29	Motor vehicles
C18	Printing	C30	Other transport
C19	Petroleum products	C31	Furniture
C20	Chemicals	C32	Other manufacturing
C21	Pharmaceuticals	C33	Repair services

Monthly indices (base year 2015 = 100) are collapsed to quarterly frequency by taking the average of the three monthly values within each quarter. Quarterly log changes are computed as:

$$\Delta \ln PPI_{it} = \ln(PPI_{it}) - \ln(PPI_{i,t-1}) \tag{A.1}$$

where i indexes the 2-digit SIC industry and t indexes the quarter. The sample covers 1996Q1–2024Q4.

Appendix A.2. Input-Output Tables

I obtain Input-Output coefficients from the ONS Supply and Use Tables (2023 edition). Sheet A of the published tables contains technical coefficients showing, for each industry pair (i, j), the value of inputs from industry i required to produce one unit of output in industry j.

The raw tables are disaggregated to approximately 100 industries using a hybrid SIC/product classification. I aggregate to 2-digit SIC level by summing coefficients across detailed industries within each 2-digit category. For example, detailed food manufacturing industries (C101 Meat processing, C102 Fish processing, C103 Fruit and vegetables, etc.) are aggregated to C10 Food products.

For downstream industry j, the input share from upstream industry i is:

$$\omega_{ij} = \frac{U_{ij}}{\sum_k U_{kj}} \tag{A.2}$$

where U_{ij} is the value of purchases by industry j from industry i, and the sum is over all upstream industries k. By construction, $\sum_i \omega_{ij} = 1$ for each downstream industry. Weights are renormalised to sum to unity over manufacturing industries for which PPI data are available.

Own-sector purchases are excluded (i.e., $\omega_{jj} = 0$) to capture purely upstream cost pressures. When a food retailer buys from food manufacturers, this is included; when a food manufacturer buys from other food manufacturers, this is excluded.

Table A.14 reports input shares for selected downstream industries, showing the concentration of upstream purchases.

Appendix A.3. COICOP-SIC Concordance

The micro-price data are classified by COICOP (Classification of Individual Consumption According to Purpose), while the PPI and Input-Output tables use SIC (Standard Industrial Classification). I construct a concordance mapping each COICOP division to its primary upstream manufacturing industry based on the dominant source of goods in each consumption category.

Table A.14: Input-Output Structure: Selected Industries

Downstream Industry	Top Upstream Suppliers	Share
G47 Retail trade	C10 Food products	0.31
	C20 Chemicals	0.08
	C22 Rubber/plastics	0.06
I56 Food service	C10 Food products	0.42
	C11 Beverages	0.18
	C17 Paper products	0.04

Table A.15: COICOP to SIC Industry Mapping

COICOP	Category	SIC	Upstream Industry
01	Food & non-alcoholic beverages	C10	Food products
02	Alcoholic beverages & tobacco	C11	Beverages
03	Clothing & footwear	C14	Wearing apparel
04	Housing & utilities	C23	Non-metallic minerals
05	Furnishings & household	C31	Furniture
06	Health	C21	Pharmaceuticals
07	Transport	C29	Motor vehicles
08	Communications	C26	Electronics
09	Recreation & culture	C32	Other manufacturing
10	Education	C18	Printing
11	Restaurants & hotels	C10	Food products
12	Miscellaneous goods & services	C32	Other manufacturing

The mapping is necessarily approximate—most consumption categories source from multiple manufacturing industries. I assign each COICOP division to the single SIC industry that represents its dominant upstream supplier. For service-dominated categories (e.g., COICOP 04 Housing), the manufacturing PPI captures only the goods component of upstream costs. Robustness checks using alternative mappings (e.g., splitting COICOP 03 Clothing between C14 Wearing apparel and C15 Leather) yield similar results.

Appendix A.4. Cost Index Construction

The upstream cost index for downstream industry j in quarter t is:

$$C_{jt} = \sum_{i \in \text{Manufacturing}} \omega_{ij} \cdot \Delta \ln PPI_{it} \tag{A.3}$$

The construction proceeds as follows:

1. Load 2-digit I-O coefficients and reshape to wide format (one row per purchasing industry, columns for each supplying industry's coefficient)
2. Expand to panel by crossing with all quarters in the PPI data
3. Merge PPI log changes by quarter
4. Compute weighted sum: for each purchasing industry j and quarter t, multiply each supplying industry's coefficient by its PPI change and sum across suppliers
5. Lag the cost index by one quarter for use in regressions

The cost index is expressed in log points (approximately percentage changes). A value of 0.01 indicates that upstream costs increased by approximately 1%.

Appendix A.5. Merging to Micro-Price Data

The cost index is merged to the micro-price panel using:

1. Extract COICOP division from the product identifier: coicop_div = $\lfloor$coicop_id/10000$\rfloor$
2. Map COICOP division to SIC using Table A.15

3. Merge on SIC × quarter

The merge rate is approximately 91%. Non-matches occur primarily for service categories (e.g., education, health services) where manufacturing PPI is not applicable.

Appendix A.6. Summary Statistics

Table A.16 reports summary statistics for the upstream cost index.

Table A.16: Upstream Cost Index Summary Statistics

Variable	N	Mean	SD	Min	Max
Cost index (%)	6,260,000	0.14	0.26	−1.82	2.41
By inflation regime:					
Below 1.9%	3,460,000	0.12	0.24	−1.82	2.41
Above 1.9%	2,800,000	0.18	0.28	−1.54	2.38

Cost shocks are somewhat larger above threshold (mean 0.18% vs 0.12%), but the standard deviations are similar (0.28% vs 0.24%), implying substantial overlap in cost shock magnitudes across regimes.

References

Aoki, K., 2001. Optimal monetary policy response to relative price changes. Journal of Monetary Economics 48, 55–80.

Blanchard, O., Gali, J., 2007. Real wage rigidities and the new keynesian model. Journal of Money, Credit and Banking 39, 35–65.

Chari, V.V., Kehoe, P.J., McGrattan, E.R., 2009. New Keynesian models: Not yet useful for policy analysis. American Economic Journal: Macroeconomics 1, 242–66.

Coibion, O., Gorodnichenko, Y., 2015. Information rigidity and the expectations formation process: A simple framework and new facts. American Economic Review 105, 2644–2678. doi:10.1257/aer.20110306.

Davies, R., 2021. Prices and inflation in a pandemic: A micro data approach. URL: http://cep.lse.ac.uk/pubs/download/dp017.pdf.

Gabaix, X., 2020. A behavioral New Keynesian model. American Economic Review 110, 2271–2327. doi:10.1257/aer.20162005.

Gagliardone, L., Gertler, M., Lenzu, S., Tielens, J., 2025. Anatomy of the phillips curve: Micro evidence and macro implications. American Economic Review Forthcoming.

Galí, J., 2025. Rethinking the New Keynesian model. Keynes Lecture, University of Cambridge. Revised October 2025.

Hazell, J., Herreño, J., Nakamura, E., Steinsson, J., 2022. The slope of the Phillips curve: Evidence from U.S. states. Quarterly Journal of Economics 137, 1299–1344. doi:10.1093/qje/qjac010.

Kara, E., 2010. Optimal monetary policy in the generalized Taylor economy. Journal of Economic Dynamics and Control 34, 2023–2037. doi:10.1016/j.jedc.2010.04.006.

Lucas, R.E., 1972. Expectations and the neutrality of money. Journal of Economic Theory 4, 103–124. doi:https://doi.org/10.1016/0022-0531(72)90142-1.

Nakamura, E., Riblier, V., Steinsson, J., 2025. Beyond the Taylor Rule. Working Paper 34200. doi:10.3386/w34200. presented at the Jackson Hole Economic Policy Symposium.

Woodford, M., 2003. Interest and prices: Foundations of a theory of monetary policy. Princeton Univeristy Press, Princeton, NJ .