

Kara, Engin

**Working Paper**

## The Natural Rate of Inflation

CESifo Working Paper, No. 12306

**Provided in Cooperation with:**

Ifo Institute – Leibniz Institute for Economic Research at the University of Munich

*Suggested Citation:* Kara, Engin (2025) : The Natural Rate of Inflation, CESifo Working Paper, No. 12306, Munich Society for the Promotion of Economic Research - CESifo GmbH, Munich

This Version is available at:

<https://hdl.handle.net/10419/336008>

**Standard-Nutzungsbedingungen:**

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

**Terms of use:**

*Documents in EconStor may be saved and copied for your personal and scholarly purposes.*

*You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.*

*If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.*

**CES ifo**

**12306  
2025**

December 2025

# Working Papers

## The Natural Rate of Inflation

Engin Kara

CES ifo

Imprint:

**CESifo Working Papers**

ISSN 2364-1428 (digital)

Publisher and distributor: Munich Society for the Promotion  
of Economic Research - CESifo GmbH

Poschingerstr. 5, 81679 Munich, Germany  
Telephone +49 (0)89 2180-2740

Email [office@cesifo.de](mailto:office@cesifo.de)  
<https://www.cesifo.org>

Editor: Clemens Fuest

An electronic version of the paper may be downloaded free of charge

- from the CESifo website: [www.ifo.de/en/cesifo/publications/cesifo-working-papers](http://www.ifo.de/en/cesifo/publications/cesifo-working-papers)
- from the SSRN website: [www.ssrn.com/index.cfm/en/cesifo/](http://www.ssrn.com/index.cfm/en/cesifo/)
- from the RePEc website: <https://ideas.repec.org/s/ces/ceswps.html>

# The Natural Rate of Inflation

Engin Kara\*

---

## Abstract

I identify the *natural rate of inflation*: the threshold where signal-to-noise crosses unity. Extending Lucas (1972), I show that when inflation falls below sectoral volatility, firms cannot distinguish aggregate from idiosyncratic shocks. Sticky-price firms rationally filter out competitor price movements, creating endogenous stability. When inflation exceeds the natural rate, filtering ceases, coordination increases, and inflation becomes self-reinforcing. Using six million UK micro-price observations, I estimate this threshold at 1.9%. Crossing it triggers a regime shift: price dispersion falls 10%, inflation persistence jumps from near-zero to 0.44, and monetary policy loses traction. These findings validate 2% targets as the limit of self-correction.

*Keywords:* Inflation targeting, strategic pricing, behavioural thresholds, sectoral volatility, monetary policy, price complementarity.

*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

The inflation process is not uniform. Its persistence shifts dramatically across time and countries—a fundamental nonlinearity that standard Phillips curves cannot explain. Periods of stable, anchored inflation give way to episodes where price pressures become self-reinforcing. Yet we lack a microfoundation for this nonlinearity—a theory grounded in firm behaviour that explains why inflation dynamics differ across regimes.

This paper provides one. I identify the *natural rate of inflation*: the threshold where signal equals noise. Below it, firms filter out transitory price movements and disturbances dissipate. Above it, firms pass through common shocks and inflation persists.

The mechanism is signal extraction combined with noise filtering. Extending Lucas (1972), I show that firms' response to competitor prices depends on their inference about shock origins. When a competitor raises prices, a firm asks: is this common cost pressure or firm-specific? If likely idiosyncratic, it will revert—and a sticky-price firm that incorporates this noise into its own price will be mispriced when reversion occurs. Profit maximisation requires filtering out transitory movements.

The critical insight is that inflation magnitude itself reveals information. When inflation is low, it falls within typical sectoral noise. Firms cannot distinguish aggregate from idiosyncratic shocks, so they filter—and price disturbances dissipate. When inflation exceeds sectoral noise, its magnitude reveals aggregate conditions directly. There is no noise to filter. Firms pass through common shocks, and disturbances persist.

The inference weight  $\theta$ —the probability a shock is idiosyncratic—varies continuously with inflation. Yet this continuous variation generates what appears to be discrete regime change. The reason is that inference is most sensitive near the threshold. When sectoral inflation is well below  $\sigma_c$ , firms are already confident shocks are noise—a small increase barely changes their view. When sectoral inflation is well

above  $\sigma_c$ , firms are already confident shocks are aggregate—again, small changes matter little. But near the threshold, firms are genuinely uncertain. A small increase in inflation can shift their belief from “probably noise” to “probably aggregate,” and with it, their pricing strategy. Most of the behavioural change therefore occurs in a narrow band around  $|\pi_c| = \sigma_c$ , creating patterns empirically indistinguishable from a discrete regime shift.

I embed this signal extraction mechanism in a heterogeneous menu-cost model. Crucially, the model shows that price stickiness is not random but endogenous to the volatility structure. High-volatility sectors experience frequent shocks that push prices outside the inaction band, forcing firms to adjust frequently. Low-volatility sectors experience infrequent shocks, allowing firms to remain sticky. The data confirm this prediction: the expenditure-weighted correlation between sectoral volatility and adjustment frequency is 0.73. This explains why a common threshold works: sticky-price sectors—which account for 90% of expenditure—cluster tightly around a mean volatility of 1.85%, essentially the 1.9% threshold. Flexible-price sectors have much higher thresholds, but since they do not filter regardless, their threshold location is irrelevant for aggregate dynamics.

I test these predictions using six million UK price records from 2003–2021, covering 67 sectors and 350,000 products. The results confirm the theory through two distinct channels: micro-level strategic interaction and macro-level dynamic persistence.

**Micro-evidence: strategic filtering.** Below the 1.9% threshold, sticky-price firms rationally filter competitor price movements. The baseline estimate shows a significant negative response to competitor price increases ( $\beta = -0.89$ ): sticky firms treat the signal as noise and lean against it. Flexible firms, by contrast, show no such filtering ( $\beta = +0.14$ )—they can re-adjust next period, so filtering provides no benefit. Above threshold, this heterogeneity disappears. Both firm types converge to near-zero responses as filtering shuts off.

To establish causality, I address the reflection problem (Manski, 1993) by ex-

exploiting the market structure of the UK retail sector, which is dominated by national chains with centralised pricing. I instrument local competitor price changes with price changes in the same category from *other regions*. Because national chains respond to aggregate cost shocks uniformly across regions but do not adjust prices in response to local demand shocks, this instrument isolates exogenous aggregate variation. The IV estimates confirm the fixed effects results: sticky-price firms below threshold show a significant negative coefficient ( $\beta_{IV} = -0.51$ ), while flexible firms show positive pass-through (+0.08). A pure common-cost shock would produce positive comovement for both; the negative sign for sticky firms confirms rational filtering. Above threshold, both firm types converge toward zero. This heterogeneity, and its disappearance at 1.9%, is inconsistent with common shocks but predicted by noise filtering. Model comparison confirms the threshold is discrete: smooth nonlinear specifications capture less than 1% of the behavioural change that sticky firms exhibit when crossing 1.9%.

**Macro-evidence: aggregate persistence.** The micro-behaviour drives aggregate dynamics through two mechanisms. First, sticky firms’ filtering dampens shock propagation. Second, flexible firms exhibit strong mean reversion below the threshold ( $\rho = -0.33$ ): they overshoot but quickly reverse, providing additional stabilisation. Using high-frequency identified monetary shocks, I find that below the threshold, sticky firms exhibit negative price responses to expansionary shocks—rationalising the classic “Price Puzzle” as a filtering phenomenon. Above the threshold, this anomaly vanishes, prices respond conventionally, and monetary policy loses its real traction. When the economy crosses 1.9%, these stabilising frictions break simultaneously. Consequently, aggregate inflation persistence jumps from near-zero to 0.44, and price dispersion falls by 10%.

The natural rate of inflation completes a trilogy of structural boundaries in monetary economics. Wicksell (1898) identified the natural rate of interest: the threshold where saving equals investment and prices are stable. Friedman (1968) identified the

natural rate of unemployment: the threshold below which inflation accelerates. Both represent structural boundaries determined by real frictions. I identify a third: the natural rate of inflation, where signal equals noise. Below it, markets self-correct. Above it, self-correction ceases. Like Wicksell’s and Friedman’s thresholds, mine is not a policy choice but a structural feature of market economies.

This speaks to a longstanding puzzle: why have central banks worldwide converged on 2% targets? Table 1 documents this convergence across diverse economies and initial conditions—from New Zealand’s pioneering 0–2% range in 1988 to the Federal Reserve’s formal adoption of 2% in 2012. The UK natural rate of 1.9% suggests a structural interpretation. If similar market structures generate similar idiosyncratic volatility across developed economies, central banks may have independently discovered the boundary where self-correction breaks down. The threshold is not arbitrary: it equals the expenditure-weighted mean of sectoral volatilities. Low-volatility sectors—which account for 90% of expenditure—cluster tightly around this value. High-volatility sectors (energy, fresh food) represent only 10% of expenditure and do not filter regardless.

Table 1: Evolution of Inflation Targets in Major Economies

Country	Year Adopted	Initial Target	Current Target
New Zealand	1988	0–2%	1–3%
Canada	1991	3–5%	2% ± 1%
UK	1992	1–4%	2%
Australia	1993	2–3%	2–3%
Sweden	1993	2% ± 1%	2% ± 1%
ECB	1999	Below 2%	2%
USA	2012	Implicit	2%

*Notes:* Initial targets varied but converged toward 2%. Sources: Central bank websites.

This framework suggests an asymmetry in the costs of missing the target. An overshoot risks crossing into a high-persistence regime where self-correction shuts off. An undershoot remains in the self-correcting regime. Policies that tolerate overshoots—such as average inflation targeting—may carry greater structural risk

than those tolerating undershoots. The costs of overshooting are potentially self-reinforcing; the costs of undershooting are self-correcting.

This paper bridges three literatures. First, it contributes to research on optimal inflation targets with price rigidities. While Coibion et al. (2012) find welfare-optimal inflation around 1.5% using Calvo pricing, and Blanco (2021) show that accounting for heterogeneous price adjustments doubles this optimal rate, I identify a different object: the natural rate at 1.9% where pricing behaviour fundamentally changes. This structural threshold—not necessarily optimal—represents a constraint that any welfare analysis must respect.

Second, I contribute to research on strategic complementarities in price setting, pioneered by Amiti et al. (2019) who developed the competitor price index approach. Existing work tends to treat strategic complementarities as constant structural parameters (see, for example, Gagliardone et al. (2025)). My findings offer a reinterpretation. The negative coefficient I estimate below threshold ( $\beta_{IV} = -0.51$ ) is not strategic substitution in the traditional sense—it reflects noise filtering: sticky-price firms do not incorporate competitor movements they expect to revert. The near-zero coefficient above threshold is not strategic complementarity—it reflects the absence of noise to filter.

Third, I contribute to research on information frictions in inflation dynamics. My mechanism builds on the rational inattention literature pioneered by Reis (2006) and Sims (2003). Afrouzi et al. (2023) show that the information structure facing firms is central for monetary policy transmission.<sup>1</sup> While Reis focuses on the timing of discrete information updates, I emphasise how the quality of continuously available information varies with the aggregate state. Firms always observe signals, but their informativeness depends on inflation magnitude. When inflation is low, signals are noisy—firms cannot distinguish aggregate from idiosyncratic shocks, so sticky-price

---

<sup>1</sup>Pasten and Schoenle (2016) provide related evidence on information frictions in a multi-product setting, showing how economies of scope in information processing affect monetary neutrality.

firms filter out transitory movements. When inflation is high, its magnitude reveals aggregate conditions directly—there is nothing to filter. This creates endogenous state-dependence: not because firms choose to ignore signals, but because the signal-to-noise ratio varies with inflation. I therefore provide a microfoundation for why pricing behaviour varies with macroeconomic conditions, complementing work on state-dependent pricing (Golosov and Lucas (2007), Nakamura and Steinsson (2010)) by showing that information quality itself is state-dependent.

The remainder of the paper proceeds as follows. Section 2 describes the data and measurement. Section 3 presents motivating evidence on the signal-to-noise framework and establishes why the threshold emerges at 1.9%. Section 4 develops the theoretical mechanism, extending Lucas (1972) to derive state-dependent signal extraction. Section 5 embeds this mechanism in a heterogeneous menu-cost model. Section 6 describes the measurement of competitor price movements and identification strategy. Section 7 tests the threshold prediction and presents the main empirical results. Section 8 addresses the reflection problem using regional instruments. Section 9 compares threshold versus smooth nonlinearity specifications. Section 10 presents robustness checks. Section 11 derives aggregation implications and examines implications for inflation dynamics as well as monetary policy transmission. Section 12 concludes.

## 2. Data and Measurement

I employ the UK CPI micro-price dataset compiled by Davies (2021). The dataset covers 2003-2021 and encompasses approximately 60% of the UK CPI basket. It provides quarterly price changes for over 6 million observations across 340,000 unique products.<sup>2</sup> I exclude temporary sales to focus on fundamental pricing decisions and

---

<sup>2</sup>I aggregate to quarterly frequency to match the monetary policy shocks from the UK Monetary Policy Event-Study Database (UKMPD), which identifies shocks from quarterly Bank of England MPC announcements. Additionally, quarterly aggregation captures pricing dynamics that unfold over multiple periods while ensuring sufficient price variation given typical adjustment frequencies.

winsorise extreme price changes at the 1<sup>st</sup> and 99<sup>th</sup> percentiles to mitigate the influence of outliers while preserving distributional asymmetry.<sup>3</sup>

The dependent variable is the non-sale price change ( $\Delta p_{i,t}$ ), which excludes temporary price reductions. The key independent variable—competitors’ average price change within each product category—is constructed as described in Section 6. I construct two key inflation measures:

$$\pi_t = \sum_i \left( \pi_{i,c,t} \times \frac{w_{i,t}}{\sum_j w_{j,t}} \right) \quad (1)$$

$$\pi_{c,t} = \sum_{i \in c} \left( \pi_{i,c,t} \times \frac{w_{i,t}}{\sum_{j \in c} w_{j,t}} \right) \quad (2)$$

where  $w_{i,t}$  represents the expenditure weight of product  $i$  at time  $t$ , and  $c$  denotes the COICOP category. COICOP categories include Food, Alcohol, Clothing, Housing, Furnishing, Health, Transport, Communication, Recreation, Restaurants, and Miscellaneous goods and services. While  $\pi_t$  captures economy-wide inflation,  $\pi_{c,t}$  isolates sector-specific dynamics, allowing to study how firms respond to local versus global price pressures.

I control for dispersion to ensure results reflect noise filtering rather than mechanical effects of price dispersion, which is measured as the weighted variance of price changes:

$$\sigma_t^2 = \sum_i w_{i,t} (\Delta p_{i,t} - \pi_t)^2 \quad (3)$$

This captures the heterogeneity in pricing decisions across firms. I also control for skewness to ensure that responses are not confounded by asymmetric distributional shifts:

$$\text{Skew}_t = \frac{\sum_i w_{i,t} (\Delta p_{i,t} - \pi_t)^3}{(\sum_i w_{i,t} (\Delta p_{i,t} - \pi_t)^2)^{3/2}} \quad (4)$$

---

<sup>3</sup>I compare aggregate inflation implied by micro data with the official inflation data. This comparison shows that the inflation measure implied by the micro-data comoves strongly with the official inflation data. On the other hand, the series with sales introduces at times extreme volatility. The correlation with official data is 0.60 excluding sales and 0.39 including sales.

This variable is lagged one period and standardised to zero mean and unit variance. Results are robust to excluding these controls.

These measures capture the pricing environment firms navigate. The key question is whether this environment fundamentally changes at an identifiable threshold.

### 3. Motivating Evidence: The Information Threshold

A sticky-price firm faces a decision: should it incorporate competitor price movements into its own price? The answer depends on persistence. If competitor movements will persist, the firm should follow—otherwise it becomes mispriced relative to the market. If competitor movements will revert, the firm should not follow—otherwise it is left overpriced when reversion occurs, losing customers. Filtering—not incorporating movements expected to revert—helps protect the firm’s competitive position.

The challenge for the firm is that it cannot observe persistence directly and must infer it. The key insight is that persistence depends on whether price movements are aggregate or idiosyncratic. Aggregate shocks affect all firms and tend to persist; idiosyncratic shocks affect individual firms and tend to revert. A firm observing competitors raise prices asks: is this aggregate cost pressure that will persist, or idiosyncratic circumstances that will revert?

#### 3.1. The Signal-to-Noise Ratio

The magnitude of sectoral price movements reveals their likely source. Large movements are unlikely to arise from idiosyncratic factors alone—aggregate forces must be present.

Sectoral inflation can be decomposed as:

$$\pi_c = \pi^{agg} + \varepsilon_c \tag{5}$$

where  $\pi^{agg}$  is aggregate inflation and  $\varepsilon_c$  is the sector-specific component. The firm observes both  $\pi_c$  and  $\pi^{agg}$ , but this does not resolve the inference problem. The ques-

tion is not “what is aggregate inflation?” but “will competitor price movements in my sector persist?” Even knowing  $\pi^{agg}$ , the firm cannot tell whether a competitor raising prices is responding to aggregate pressure (and will stay high) or to idiosyncratic circumstances (and will revert).

The magnitude provides a signal. When sectoral inflation is large relative to typical idiosyncratic variation, aggregate forces are likely driving competitor behaviour—and aggregate-driven movements persist.

I define the signal-to-noise ratio as:

$$SNR_{c,t} = \frac{|\pi_{c,t}^{ann}|}{\sigma_c} \quad (6)$$

where the numerator is absolute annualised sectoral inflation and the denominator is the sector’s typical idiosyncratic volatility:

$$\sigma_c = \text{SD}(\pi_c - \pi^{agg}) \quad (7)$$

This measures how much sector  $c$  normally deviates from aggregate conditions—a structural parameter firms learn from experience.

When  $SNR < 1$ , sectoral inflation falls within typical variation—competitor movements could easily reflect idiosyncratic factors that will revert. When  $SNR \geq 1$ , sectoral inflation exceeds typical noise—idiosyncratic factors alone are unlikely to generate movements this large, so aggregate forces must be present, and aggregate-driven movements persist.

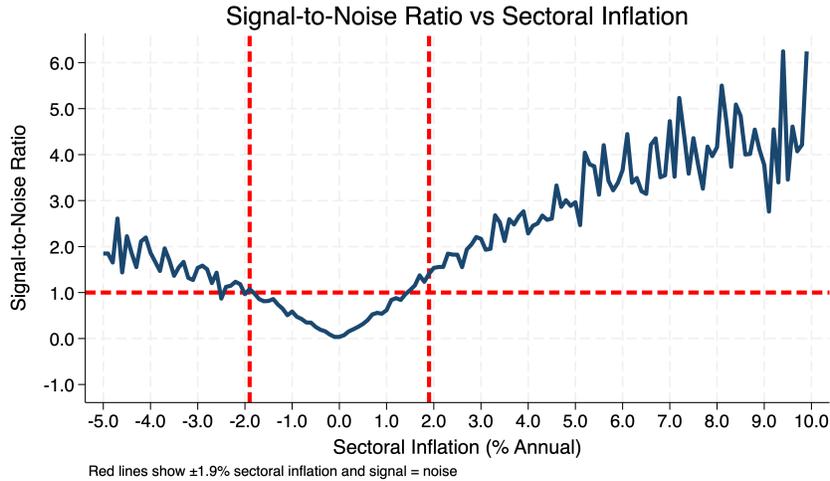


Figure 1: Signal-to-Noise Ratio and the Information Threshold

*Notes:* Signal-to-noise ratio (absolute sectoral inflation divided by sector-specific idiosyncratic volatility  $\sigma_c$ ) averaged across sectors. SNR crosses unity at approximately  $\pm 1.9\%$  (vertical lines). Below this threshold, competitor price movements are indistinguishable from idiosyncratic noise; above it, their magnitude reveals aggregate-driven behaviour.

Figure 1 plots SNR against sectoral inflation. The ratio crosses unity at approximately  $\pm 1.9\%$ —reflecting average sectoral volatility ( $\bar{\sigma}_c \approx 1.9\%$ ). The symmetry at  $\pm 1.9\%$  reinforces the insight that this is an information boundary, not an artefact of inflation-specific phenomena.

### 3.2. Why a Common Threshold Works

In principle, each sector has its own threshold at  $|\pi_c| = \sigma_c$ . Energy, with  $\sigma_c = 22\%$ , crosses at 22%. Hairdressing, with  $\sigma_c = 1\%$ , crosses at 1%. Why does a common 1.9% threshold work?

Two facts provide the answer.

**Fact 1: Sectoral volatility and price flexibility are highly correlated.**

Figure 2 plots  $\sigma_c$  against average adjustment frequency by sector. The expenditure-weighted correlation is 0.73. High-volatility sectors adjust prices frequently; low-volatility sectors adjust infrequently.

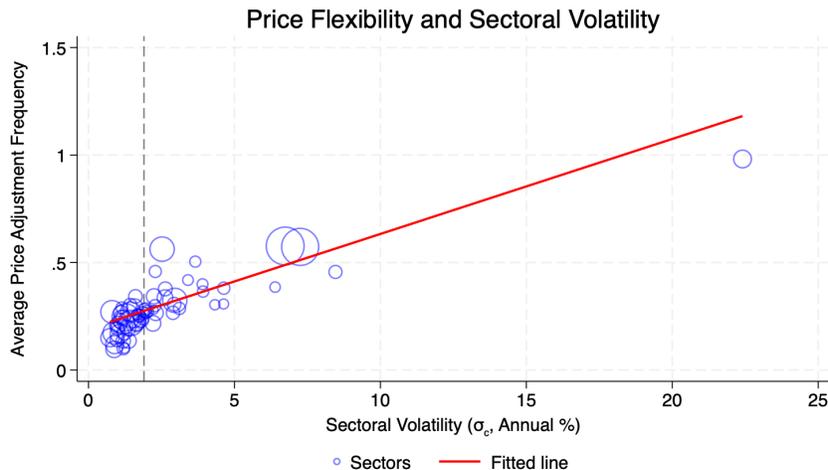


Figure 2: Price Flexibility and Sectoral Volatility

*Notes:* Each circle represents a COICOP sector. Circle size proportional to expenditure weight. Horizontal axis: sectoral inflation volatility (annualised). Vertical axis: average price adjustment frequency. Dashed line indicates the natural rate (1.9%).

Table 2: Sectoral Volatility and Price Flexibility

	Low Volatility ( $\sigma_c < 3\%$ )	High Volatility ( $\sigma_c \geq 3\%$ )	Difference
Mean Volatility ( $\sigma_c$ )	1.6%	6.4%	+4.8 pp
Mean Adjustment Frequency	0.28	0.46	+0.18
Expenditure Share	90%	10%	—
Number of Sectors	54	13	—
Correlation (freq, $\sigma_c$ ): 0.73 (expenditure-weighted)			
Expenditure-weighted mean $\sigma_c$ : 1.85%			

*Notes:* Low-volatility sectors defined as  $\sigma_c < 3\%$  (annualised). Adjustment frequency is the proportion of quarters with a non-zero price change. All statistics are expenditure-weighted.

Table 2 quantifies this relationship. Low-volatility sectors ( $\sigma_c < 3\%$ ) account for 90% of expenditure and adjust prices in 28% of quarters. High-volatility sectors adjust nearly twice as often (46%) but represent only 10% of the consumption basket. The expenditure-weighted mean of sectoral volatility is 1.85%—essentially the 1.9% threshold identified in the empirical analysis.

**Fact 2: Only sticky-price firms need to filter.** Flexible-price firms can re-

adjust next period. If they respond to noise today, they correct tomorrow. The cost of error is low, so filtering provides little benefit.

The common 1.9% threshold is therefore not arbitrary but an emergent property of the data. For sticky-price sectors ( $\sigma_c \approx 1\text{--}2\%$ ), the common threshold roughly equals their sector-specific thresholds. For flexible-price sectors ( $\sigma_c > 3\%$ ), the threshold is irrelevant—they do not filter regardless. Since sticky-price sectors account for 90% of expenditure, the common threshold  $\bar{\sigma}_c \approx 1.9\%$  captures when the bulk of the economy shifts regime. Section 4 formalises this as the natural rate of inflation and derives testable Predictions.

#### 4. Theoretical Framework: Extending Lucas

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  $\pi$  but cannot directly distinguish its source. The signal is either aggregate (common to all agents) or idiosyncratic (specific to this market). Lucas showed the optimal inference weight is:

$$\theta = \frac{\sigma_c^2}{\sigma_c^2 + \sigma_{agg}^2} \tag{8}$$

where  $\theta$  is the probability the shock is idiosyncratic,  $\sigma_c^2$  is the variance of idiosyncratic shocks, and  $\sigma_{agg}^2$  is the variance of aggregate shocks.

A key limitation of Lucas’s framework is that  $\theta$  depends only on structural variances—it is constant regardless of the realised signal. A sector experiencing 5% inflation faces the same inference problem as one experiencing 0.5% inflation. Agents always face the same degree of confusion.

##### 4.1. State-Dependent Signal Extraction

I extend Lucas’s model by allowing the inference weight to depend on the *realised* signal. The intuition is straightforward: large realisations are informative precisely

because they are unlikely to arise from idiosyncratic factors alone.

A firm in sector  $c$  observes sectoral inflation  $\pi_c$  but faces uncertainty about whether competitor price movements reflect aggregate or idiosyncratic forces:

$$\pi_c = \pi^{agg} + \varepsilon_c \tag{9}$$

where  $\pi^{agg}$  is aggregate inflation and  $\varepsilon_c$  is the sector-specific component. The firm observes both  $\pi_c$  and  $\pi^{agg}$ , but this does not resolve the inference problem. The question is not “what is aggregate inflation?” but “will competitor price movements persist?” Even knowing  $\pi^{agg}$ , the firm cannot tell whether a competitor raising prices is responding to aggregate pressure (and will stay high) or to idiosyncratic circumstances (and will revert).

The firm knows its sector’s typical idiosyncratic volatility  $\sigma_c \equiv \text{SD}(\pi_c - \pi^{agg})$ —how much its sector normally deviates from aggregate conditions.

The firm’s inference problem: when  $|\pi_c|$  is small relative to  $\sigma_c$ , competitor movements could easily reflect idiosyncratic factors that will revert. When  $|\pi_c|$  is large relative to  $\sigma_c$ , idiosyncratic factors alone are unlikely to generate such movements—aggregate forces must be present, and aggregate-driven movements persist.

**Assumption 1** (Distributional Assumptions). *Shocks are idiosyncratic or aggregate with equal prior probability. Idiosyncratic shocks are normally distributed with mean zero and variance  $\sigma_c^2$ . Aggregate shocks have heavier tails, with likelihood ratio:*

$$\frac{f_A(\pi_c)}{f_I(\pi_c)} = \frac{\pi_c^2}{\sigma_c^2} \tag{10}$$

This functional form is chosen for tractability. The key qualitative feature—that large realisations shift inference toward aggregate—holds for any specification where aggregate shocks have heavier tails. The specific  $\theta(\pi_c)$  formula delivers a clean closed-form solution, but the threshold location ( $|\pi_c| = \sigma_c$ ) is identified from data, not derived from this assumption. In the UK, average sectoral volatility is 1.9%, so the threshold is 1.9%—this would hold under any reasonable distributional specification that preserves the monotonicity of  $\theta(\pi_c)$ .

**Proposition 1** (State-Dependent Inference). *Under Assumption 1, the posterior probability that a shock is idiosyncratic, given realisation  $\pi_c$ , is:*

$$\theta(\pi_c) = \frac{\sigma_c^2}{\sigma_c^2 + \pi_c^2} \quad (11)$$

*Proof.* By Bayes' rule:

$$\theta(\pi_c) = \frac{f_I(\pi_c)}{f_I(\pi_c) + f_A(\pi_c)} = \frac{f_I(\pi_c)}{f_I(\pi_c) + \frac{\pi_c^2}{\sigma_c^2} f_I(\pi_c)} = \frac{1}{1 + \frac{\pi_c^2}{\sigma_c^2}} = \frac{\sigma_c^2}{\sigma_c^2 + \pi_c^2} \quad (12)$$

□

The equal prior assumption ( $P(I) = P(A) = \frac{1}{2}$ ) reflects the uninformative case: absent any signal, firms have no basis to distinguish shock types. The signal magnitude then determines inference entirely.

#### 4.2. Connection to Signal-to-Noise Ratio

The inference weight maps directly to the signal-to-noise ratio. Define:

$$SNR = \frac{|\pi|}{\sigma_c} \quad (13)$$

Then Proposition 1 can be rewritten as:

$$\theta = \frac{1}{1 + SNR^2} \quad (14)$$

This yields a clear mapping between signal strength and inference:

Table 3: Signal-to-Noise Ratio and Inference

$SNR$	$\theta$	Interpretation
0	1.0	Certainly idiosyncratic
0.5	0.8	Probably idiosyncratic
<b>1</b>	<b>0.5</b>	<b>Maximum uncertainty</b>
2	0.2	Probably aggregate
$\rightarrow \infty$	0	Certainly aggregate

### 4.3. The Natural Rate of Inflation

The threshold occurs where  $\theta = 0.5$ —the point of maximum uncertainty:

$$\theta(\pi^N) = 0.5 \implies |\pi^N| = \sigma_c \quad (15)$$

I define this as the *natural rate of inflation*: the boundary where inference shifts from “probably idiosyncratic” to “probably aggregate.”

As shown in Figure 1, the signal-to-noise ratio crosses unity at approximately  $\pm 1.9\%$  in UK data—reflecting the average sectoral idiosyncratic volatility ( $\bar{\sigma}_c = 1.9\%$ ). This is not a calibration choice; it emerges from the data.

### 4.4. From Inference to Behaviour: Filtering

How does inference translate into behaviour? A sticky-price firm must set a price that remains fixed for horizon  $\Phi$ , requiring a price optimal *on average* over this horizon. Whether competitor movements persist therefore matters.

Consider a sticky firm observing competitors raise prices. If the movement is idiosyncratic (probability  $\theta$ ), competitors will revert—a firm that incorporates this movement becomes mispriced. If the movement is aggregate (probability  $1 - \theta$ ), it reflects economy-wide conditions already captured by aggregate inflation. Competitor prices add no information.

This generates two regimes:

**Below threshold** ( $|\pi_c| < \sigma_c, \theta > 0.5$ ): Competitor movements are likely idiosyncratic and will revert. Incorporating them leads to mispricing. The optimal response is to filter:  $\beta < 0$ .

**Above threshold** ( $|\pi_c| > \sigma_c, \theta < 0.5$ ): Competitor movements are likely aggregate-driven, but aggregate conditions are already observed directly. Competitor prices convey no additional information:  $\beta \approx 0$ .

Why do sticky and flexible firms differ? Flexible-price firms need not filter—if competitor movements prove transitory, they re-adjust next period at low cost.

Sticky-price firms lack this option. Committed to their price for horizon  $\Phi$ , they must filter today or remain mispriced until their next adjustment.

#### 4.5. *Equilibrium Coordination*

Sectoral inflation is publicly observed, creating common knowledge about which side of the threshold the sector occupies. This coordinates expectations. Below threshold, even a firm suspecting an aggregate shock knows other sticky firms in its sector face the same weak signal. They will filter. If they filter, their movements revert. So this firm filters too. Above threshold, the signal is strong. Other firms will pass through. If they pass through, movements persist. So this firm passes through too.

The result is self-fulfilling dynamics: filtering below threshold creates the reversion that justifies filtering; pass-through above threshold creates the persistence that justifies pass-through. The natural rate is the boundary between these equilibria.

This explains why even identified aggregate shocks generate filtering below threshold. The firm’s decision depends not on whether *this* shock is aggregate, but on whether the market coordinates on filtering or pass-through. Below threshold, the equilibrium is filtering—regardless of any individual shock’s origin.

#### 4.6. *Testable Predictions*

The framework generates five predictions: (1) sticky-price firms filter competitor movements below threshold ( $\beta < 0$ ); (2) flexible-price firms do not filter ( $\beta \approx 0$ ); (3) the sticky-flexible gap disappears above threshold; (4) aggregate inflation persistence changes at threshold; (5) even identified aggregate shocks generate filtering below threshold.

Prediction (5) is particularly sharp. If filtering reflected pure inference—“I don’t know if this is aggregate”—then identifying an aggregate shock should eliminate filtering. But if filtering reflects equilibrium coordination—“others will filter regardless”—then identified shocks should still generate filtering below threshold. The negative coefficient below threshold—what the literature labels “strategic substitution”—has

a microfoundation: sticky firms filtering transitory movements to protect competitiveness.

## 5. From Signal Extraction to Pricing Decisions

This section embeds the signal extraction mechanism in a heterogeneous menu-cost model to derive its implications for pricing formally.

### 5.1. Model Environment

The model departs from standard menu cost frameworks in two ways. First, firms differ in adjustment costs: each firm  $i$  draws a permanent menu cost parameter  $\psi_i$  from a distribution  $F(\psi)$  with support  $[\underline{\psi}, \bar{\psi}]$ , implying a real adjustment cost of  $\psi_i \cdot w_t$ . Second, firms face the signal extraction problem developed above, making their inference weight  $\theta$  depend on inflation magnitude.

A continuum of monopolistically competitive firms operate linear technologies  $y_{it} = A_t n_{it}$ , facing demand (16) and common nominal marginal cost  $MC_t = W_t/A_t$ . Here  $y_{it}$  denotes output,  $A_t$  productivity,  $W_t$  wages, and  $n_{it}$  labour. Firms differ only in  $\psi_i$ : high- $\psi$  firms adjust infrequently (sticky prices); low- $\psi$  firms adjust often (flexible prices).

A representative household has CES preferences over differentiated goods with elasticity  $\epsilon > 1$ . Standard optimisation yields the demand system:

$$y_{it} = \left( \frac{P_{it}}{P_t} \right)^{-\epsilon} Y_t, \quad P_t = \left[ \int_0^1 P_{it}^{1-\epsilon} di \right]^{\frac{1}{1-\epsilon}} \quad (16)$$

where  $P_{it}$  is firm  $i$ 's price and  $P_t$  is the aggregate price index. The timing within each period unfolds as follows: (i) firms observe the aggregate state  $\Omega_t = \{A_t, P_t, Y_t, \pi_{c,t}, \pi_t\}$  and form beliefs via signal extraction; (ii) each firm decides whether to adjust based on its menu cost and distance from optimal price; (iii) adjusting firms pay  $\psi_i \cdot w_t$  and set new prices; (iv) production and consumption occur.

The key innovation is that all firms face the same signal extraction problem, generating a common  $\theta(\pi_c)$ , but heterogeneous menu costs imply different pricing

horizons  $\Phi(\psi_i)$ . This interaction determines noise filtering: sticky firms (high  $\Phi$ ) filter aggressively below threshold; flexible firms (low  $\Phi$ ) need not filter.

### 5.2. The Firm's Problem

Each firm  $i$  chooses its price to maximise the expected discounted value of profits. The firm's state is characterised by its current price  $P_{it}$  (or equivalently, the log relative price  $\hat{p}_{it} \equiv \log(P_{it}/P_{c,t})$ , defined relative to the sectoral price index) and the aggregate state  $\Omega_t$ .<sup>4</sup> The Bellman equation is:

$$V(\hat{p}_{it}, \Omega_t; \psi_i) = \max \{V^n(\hat{p}_{it}, \Omega_t; \psi_i), \quad V^a(\Omega_t; \psi_i) - \psi_i w_t\} \quad (17)$$

where  $V^n$  denotes the value of not adjusting and  $V^a$  denotes the value after adjustment.

If the firm does not adjust, it earns current profits and its relative price erodes with sectoral inflation:

$$V^n(\hat{p}_{it}, \Omega_t; \psi_i) = \Pi(\hat{p}_{it}, \Omega_t) + \beta \mathbb{E}_t[V(\hat{p}_{it} - \pi_{c,t+1}, \Omega_{t+1}; \psi_i)] \quad (18)$$

If the firm adjusts, it chooses the optimal reset price  $\hat{p}_t^R(\psi_i)$ :

$$V^a(\Omega_t; \psi_i) = \max_{\hat{p}} \{\Pi(\hat{p}, \Omega_t) + \beta \mathbb{E}_t[V(\hat{p} - \pi_{c,t+1}, \Omega_{t+1}; \psi_i)]\} \quad (19)$$

The profit function in terms of log relative prices is:

$$\Pi(\hat{p}_{it}, \Omega_t) = Y_t [\exp((1 - \epsilon)\hat{p}_{it}) - mc_t \exp(-\epsilon\hat{p}_{it})] \quad (20)$$

where  $mc_t = W_t/(P_t A_t)$  is the real marginal cost.

---

<sup>4</sup>For tractability and consistency with the identification strategy, I define the relative price vis-à-vis the sectoral price index rather than the aggregate price index. Since the filtering mechanism is driven by the signal-to-noise ratio of sectoral prices ( $|\pi_c| \leq \sigma_c$ ), this provides a coherent mapping to the empirical specification. See Appendix A for details.

### 5.3. Endogenous Price Flexibility

The effective pricing horizon  $\Phi(\psi)$  is determined endogenously by the interaction of menu costs and sectoral volatility. As shown in Appendix A, the adjustment hazard is:

$$\lambda(\psi) = \frac{\kappa_0 \sigma_c^2}{2\psi} \quad (21)$$

where  $\kappa_0$  is the profit function curvature and  $\sigma_c^2$  is sectoral volatility.

This expression reveals a key connection: the adjustment hazard depends on sectoral volatility. High-volatility sectors experience frequent shocks that push prices outside the inaction band, generating high  $\lambda$  and short pricing horizons  $\Phi$ . Low-volatility sectors experience infrequent shocks, generating low  $\lambda$  and long horizons.

The model thus endogenously generates the strong correlation between sectoral volatility and price flexibility documented in Section 2 (Figure 2). This is not assumed—it emerges from the optimising behaviour of firms facing heterogeneous sectoral environments.

### 5.4. The Filtering Problem

A sticky-price firm must set a price that will remain fixed for an extended horizon  $\Phi$ . Profit maximisation requires setting the price that is optimal *on average* over this horizon—not the price that is optimal today. This creates a filtering problem: the firm observes current sectoral prices, but these contain noise—idiosyncratic movements that will revert before the pricing spell ends. Incorporating this noise would leave the firm mispriced when reversion occurs.

Sectoral price movements contain two components. The **aggregate component** ( $\pi^{agg}$ ) reflects common cost pressures affecting all firms and will persist. The **idiosyncratic component** ( $\varepsilon_c$ ) reflects sector-specific factors—temporary demand shifts, inventory adjustments, or pricing errors—and will revert. The firm observes sectoral inflation:

$$\pi_c = \pi^{agg} + \varepsilon_c \quad (22)$$

If idiosyncratic shocks revert with persistence  $\rho_\varepsilon$ , the expected level at horizon  $k$  is  $\mathbb{E}_t[\varepsilon_{c,t+k}] = \rho_\varepsilon^k \cdot \varepsilon_{c,t}$ . A firm with pricing horizon  $\Phi$  cares about the *average* level over its entire spell:

$$\mathbb{E}_t \left[ \frac{1}{\Phi} \sum_{k=0}^{\Phi-1} \varepsilon_{c,t+k} \right] = \Lambda \cdot \varepsilon_{c,t} \quad (23)$$

where the **horizon filter** is:

$$\Lambda(\Phi, \rho_\varepsilon) = \frac{1 - \rho_\varepsilon^\Phi}{\Phi(1 - \rho_\varepsilon)} \quad (24)$$

The horizon filter  $\Lambda \in (0, 1]$  determines how much of the current idiosyncratic movement persists over the pricing horizon. Flexible firms ( $\Phi = 1$ ) have  $\Lambda \approx 1$ —if they respond to noise today, they correct tomorrow. Sticky firms (high  $\Phi$ ) have  $\Lambda$  small—noise washes out over their pricing horizon, so incorporating it today leaves them mispriced.

From Section 4, a firm assigns probability  $\theta$  that observed sectoral movements are idiosyncratic:

$$\theta(\pi_c) = \frac{\sigma_c^2}{\sigma_c^2 + \pi_c^2} \quad (25)$$

The firm’s optimal response combines signal extraction and filtering:

- With probability  $\theta$  (idiosyncratic), the movement will partially revert—respond with weight  $\Lambda$
- With probability  $1 - \theta$  (aggregate), the movement will persist—respond fully

This defines the **effective pass-through**:

$$\rho(\theta) = \theta \cdot \Lambda + (1 - \theta) \cdot 1 = 1 + \theta(\Lambda - 1) \quad (26)$$

Below threshold ( $\theta \rightarrow 1$ ), effective pass-through approaches  $\Lambda$ . For sticky firms with long pricing horizons,  $\Lambda$  is small, so they filter aggressively. Above threshold ( $\theta \rightarrow 0$ ), effective pass-through approaches unity and filtering shuts off.

### 5.5. Optimal Reset Price

**Proposition 2** (Optimal Reset Price with Filtering). *The optimal reset price is:*

$$\hat{p}_t^R = \hat{p}_t^* + \alpha_1(\Phi, \theta) \mathbb{E}_t[\pi_{c,t+1}] \quad (27)$$

where:

$$\alpha_1(\Phi, \theta) = \frac{\beta\Phi(1 + \rho(\theta))}{1 + \beta\Phi} \quad (28)$$

*Proof:* See Appendix B.

Filtering reduces the coefficient on expected sectoral inflation. Below threshold ( $\theta \rightarrow 1$ ),  $\rho(\theta) \rightarrow \Lambda$  and  $1 + \rho(\theta) \rightarrow 1 + \Lambda < 2$ . For sticky firms with  $\Lambda$  small,  $\alpha_1$  is substantially reduced—firms do not raise prices much in response to expected inflation because they expect competitor movements to revert. Above threshold ( $\theta \rightarrow 0$ ),  $\rho(\theta) \rightarrow 1$  and  $\alpha_1 \rightarrow \frac{2\beta\Phi}{1+\beta\Phi}$ —full pass-through is restored.

### 5.6. From Micro Filtering to Macro Persistence

The filtering mechanism has direct implications for aggregate inflation persistence. Consider a shock that raises one firm’s price. How this shock propagates depends on the inflation regime.

**Below threshold** ( $|\pi_c| < \sigma_c$ ): Competitors observe the price increase but assign high probability that it is idiosyncratic ( $\theta \approx 1$ ). They filter it out, expecting reversion. The shock remains localised and dissipates.

**Above threshold** ( $|\pi_c| > \sigma_c$ ): Competitors observe the price increase and recognise it as reflecting aggregate conditions ( $\theta \approx 0$ ). They pass it through to their own prices. The shock spreads across firms and becomes embedded in the aggregate price level.

The micro mechanism suggests aggregate inflation inherits persistence from pass-through behaviour. A reduced-form representation is:

$$\pi_{c,t} = \rho(\theta) \cdot \pi_{c,t-1} + \varepsilon_t \quad (29)$$

where  $\rho(\theta) = 1 + \theta(\Lambda - 1)$ . Below threshold,  $\rho(\theta) \approx \Lambda$ , which is small for sticky firms—shocks die out. Above threshold,  $\rho(\theta) \approx 1$ —shocks persist.

This connects micro behaviour to macro dynamics. The same mechanism that generates filtering below threshold also generates low aggregate persistence. The same mechanism that generates pass-through above threshold also generates high aggregate persistence.

The empirical analysis tests the model’s predictions directly using micro-price data.

### 5.7. Testable Predictions

The model generates predictions about the coefficient on competitor price movements. Define  $\beta_1$  as the firm’s response to competitor prices relative to full pass-through. Full pass-through ( $\beta_1 = 0$ ) means incorporating competitor movements one-for-one. Filtering ( $\beta_1 < 0$ ) means discounting competitor movements—responding less than one-for-one.

When  $\theta \approx 1$  (below threshold) and  $\Lambda \approx 0$  (sticky firm), the firm heavily discounts competitor movements, generating  $\beta_1 < 0$ . The theoretical coefficient is:

$$\beta_1 = \theta(\Lambda - 1) \tag{30}$$

This yields four predictions:

1. **Sticky firms, below threshold:**  $\theta \approx 1$ ,  $\Lambda$  small  $\Rightarrow \beta_1$  strongly negative
2. **Flexible firms, below threshold:**  $\theta \approx 1$ ,  $\Lambda \approx 1 \Rightarrow \beta_1 \approx 0$
3. **All firms, above threshold:** Filtering shuts off ( $\theta \rightarrow 0$ )  $\Rightarrow \beta_1 \approx 0$
4. **Aggregate persistence:** Low below threshold (shocks filtered), high above threshold (shocks passed through)

Only sticky-price firms, and only below threshold, should exhibit filtering. The negative coefficient below threshold—what the literature labels “strategic substitution”—reflects sticky firms discounting competitor movements they expect to revert. I now take these predictions to the data.

## 6. Measuring Competitor Price Movements and Identification

I construct each firm’s competitors’ average price change as the expenditure-weighted mean of all other products in the same COICOP category:

$$\Delta \bar{p}_{-i,c,t} = \frac{\sum_{j \in c, j \neq i} w_{j,t} \cdot \Delta p_{j,c,t}}{\sum_{j \in c, j \neq i} w_{j,t}} \quad (31)$$

where weights  $w_{j,t}$  reflect expenditure shares and the subscript  $-i$  indicates exclusion of product  $i$  from both numerator and denominator. This captures the competitive environment each firm faces without mechanical correlation with its own pricing decision. I use lagged values ( $\Delta \bar{p}_{-i,c,t-1}$ ) to mitigate simultaneity.

### 6.1. Identification Strategy

Identifying causal responses to competitor prices faces two fundamental challenges: common shocks and the reflection problem (Manski, 1993). Common shocks could generate spurious correlation if firms in a category face correlated cost pressures. The reflection problem arises because firm  $i$  and its competitors respond to each other simultaneously, making causal inference difficult.

I address these challenges through two complementary strategies.

**Pattern-based identification.** Standard common shocks (e.g., input costs) predict positive comovement. Yet the estimated coefficient for sticky firms is *negative* ( $\beta_{FE} = -0.89$ )—they move opposite to competitors below threshold. The sign itself rules out standard cost shocks. Furthermore, sticky and flexible firms behave fundamentally differently within the *same sectors and time periods*: flexible firms show near-zero coefficients ( $\beta = +0.14$ ). If omitted variables drove the results, they would affect both firm types similarly. This heterogeneity is uniquely predicted by filtering: flexible firms need not filter because they can re-adjust next period. Finally, the sticky-flexible difference disappears above 1.9%, ruling out time-invariant confounders.

**Instrumental variables.** To address the reflection problem directly, Section 8 instruments own-region competitor prices with prices in the same category from *other*

regions. This isolates variation driven by national shocks that are exogenous to the local firm. The IV results confirm the filtering pattern: sticky firms show significant negative coefficients below threshold ( $\beta_{IV} = -0.51$ ), flexible firms show near-zero ( $\beta_{IV} = +0.08$ ), and the difference attenuates above 1.9%.

The negative coefficient below threshold—what the literature labels “strategic substitution”—is therefore robust to simultaneity. It reflects sticky firms filtering competitor movements that appear indistinguishable from noise in low-inflation environments.

## 6.2. Summary Statistics

Table 4 describes the data.

Table 4: Summary Statistics

Variable	Mean	SD	Min	p25	p50	p75	Max
Price Change (%)	0.45	10.81	-59.78	0.00	0.00	0.00	59.42
Competitors’ Price Change (%)	0.03	0.40	-16.19	0.00	0.01	0.04	21.09
Adjustment Frequency	0.29	0.19	0.00	0.16	0.27	0.40	1.00
Sectoral Inflation (%)	0.39	1.23	-31.44	-0.16	0.38	0.85	27.73
Aggregate Inflation (%)	0.45	0.35	-0.37	0.20	0.43	0.67	1.58

*Notes:* Six million quarterly observations, UK CPI 2003–2021. Price changes exclude temporary sales. Competitors’ price change is expenditure-weighted within COICOP categories, excluding own price. Adjustment frequency is the proportion of quarters with non-zero price changes—the empirical proxy for pricing horizon  $\Phi$  (see Section 4). Low frequency firms are “sticky” (long  $\Phi$ ); high frequency firms are “flexible” (short  $\Phi$ ). All inflation rates are quarterly.

The data show substantial price rigidity—over 75% of observations show no price change—combined with heterogeneity in adjustment frequency (mean 29%, SD 19%). This heterogeneity allows testing whether filtering differs systematically between sticky and flexible firms. The sample spans diverse inflation environments (quarterly aggregate rates from  $-0.37\%$  to  $1.58\%$ ), providing variation to identify behavioural changes at the threshold.

## 7. Testing the Threshold Prediction

The model predicts that filtering shuts off at  $|\pi_c| = \sigma_c \approx 1.9\%$ . I now test this prediction. To identify whether a threshold exists and where it occurs, I estimate equation (32) for candidate thresholds  $\tau \in \{-2\%, -1.75\%, \dots, 6\%\}$  and compute F-statistics for  $H_0 : \beta_4 = 0$ . If the model is correct, the F-statistic should peak at 1.9%.

$$\begin{aligned}
\Delta p_{i,c,t} = & \alpha_i + \beta_1 \Delta \bar{p}_{-i,c,t-1} + \beta_2 \Delta \bar{p}_{-i,c,t-1} \times \mathbb{I}(\pi_{c,t-1}^{ann} \geq \tau) \\
& + \beta_3 \Delta \bar{p}_{-i,c,t-1} \times \text{freq}_i \\
& + \beta_4 \Delta \bar{p}_{-i,c,t-1} \times \mathbb{I}(\pi_{c,t-1}^{ann} \geq \tau) \times \text{freq}_i \\
& + \gamma_1 \text{Skew}_{t-1} + \gamma_2 \text{Disp}_{t-1} + \gamma_3 \text{Skew}_{t-1} \times \text{freq}_i \\
& + \gamma_4 \text{Disp}_{t-1} \times \text{freq}_i + \delta_1 \Delta p_{i,c,t-1} \\
& + \delta_2 \pi_t + \delta_3 \pi_{c,t-1} + \varepsilon_{i,c,t}
\end{aligned} \tag{32}$$

where  $\Delta p_{i,c,t}$  is firm  $i$ 's price change,  $\Delta \bar{p}_{-i,c,t-1}$  is competitors' average price change, and  $\mathbb{I}(\pi_{c,t-1}^{ann} \geq \tau)$  indicates when annualised sectoral inflation exceeds threshold  $\tau$ . Controls include lagged skewness and dispersion (interacted with frequency), lagged own price change, and aggregate and sectoral inflation.

I measure firm type using price adjustment frequency ( $\text{freq}_i$ ), calculated as the proportion of quarters with non-zero price changes over the sample period. This serves as the empirical proxy for pricing horizon  $\Phi$ —low-frequency firms are “sticky” (long horizon), high-frequency firms are “flexible” (short horizon).

This maps directly to the model. The baseline coefficient  $\beta_1$  captures the response to competitor price movements for sticky firms ( $\text{freq} = 0$ ) below threshold. The threshold interaction  $\beta_2$  tests whether this response changes when inflation crosses  $\tau$ . The frequency interaction  $\beta_3$  captures how pricing horizon affects filtering—flexible firms (high freq) filter less.

The triple interaction coefficient  $\beta_4$  tests the core prediction: whether filtering by

sticky firms shuts off at the threshold. Below threshold, sticky firms filter ( $\beta_1 < 0$ ) while flexible firms do not ( $\beta_1 + \beta_3 \approx 0$ ). Above threshold, filtering is unnecessary—the gap between sticky and flexible firms should disappear. If the model is correct, I should observe: (1) F-statistics peaking at  $\tau = 1.9\%$ , and (2) significantly negative  $\beta_4$ , indicating that the sticky-flexible gap narrows above threshold.

### 7.1. Threshold Identification

Figure 3 presents the threshold search results. Recall that theory makes a precise prediction: behavioural change should occur at  $|\pi_c| = \sigma_c \approx 1.9\%$ —where signal-to-noise crosses unity. To identify the threshold, I employ an F-test for the significance of  $\beta_4$ . Since I test 33 thresholds, I apply Bonferroni correction ( $\alpha = 0.0015$ ) to guard against spurious significance.

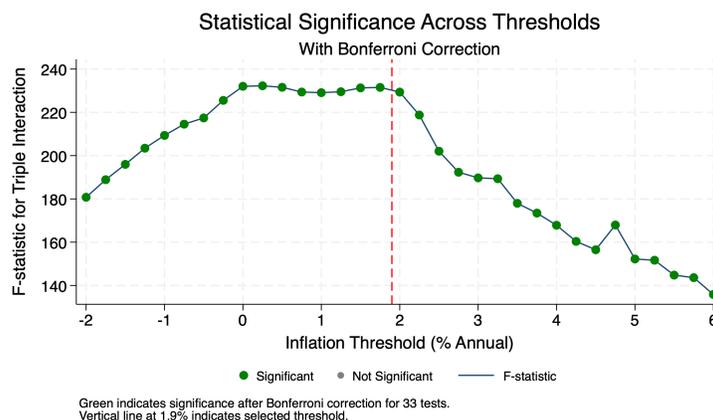


Figure 3: Statistical Significance of Threshold Effects Across Alternative Inflation Rates  
*Notes:* F-statistics testing  $\beta_4$  at thresholds from  $-2\%$  to  $6\%$  ( $0.25\text{pp}$  increments). Green dots survive Bonferroni correction for 33 tests. F-statistics plateau at 230–235 from  $0\%$  to  $2\%$ , then decline sharply. Vertical line marks selected threshold at  $1.9\%$ .

The results confirm the theory. F-statistics plateau at 230–235 from  $0\%$  to  $1.9\%$ , then decline sharply—falling below 180 by  $4\%$ . Even under Bonferroni correction, the F-statistics remain above the corrected cutoff across a wide range.

Why a plateau rather than a sharp peak? Throughout the  $0\text{--}2\%$  range,  $|\pi_c| < \sigma_c$ , so  $SNR < 1$ . Sticky firms filter throughout this range—the behavioural pattern is

similar whether inflation is 0.5% or 1.5%. The stable, high F-statistics reflect the same filtering regime across all low-inflation values.

While all thresholds from  $-2\%$  to  $4.5\%$  yield statistically significant results after Bonferroni correction, the variation in F-statistics reveals that not all thresholds are economically equivalent. The contrast between stable high values below  $2\%$  and declining values above suggests that the  $0-1.9\%$  range represents a distinct behavioural regime. I select  $1.9\%$  as the threshold, matching the theoretical prediction.<sup>5</sup>

## 7.2. Validation

A natural concern is whether the  $1.9\%$  threshold is genuinely special or an artifact of the analysis. I address this with two tests. First, I test 100 alternative thresholds randomly selected between  $1.5\%$  and  $2.5\%$ , re-estimating my main specification for each one. Figure 4 presents the results. F-statistics remain stable around 230 between  $1.5\%$  and  $1.9\%$ , then decline steadily above  $2\%$ , falling below 200 by  $2.5\%$ . The  $1.9\%$  threshold is not arbitrary—it marks the boundary where filtering shuts off.

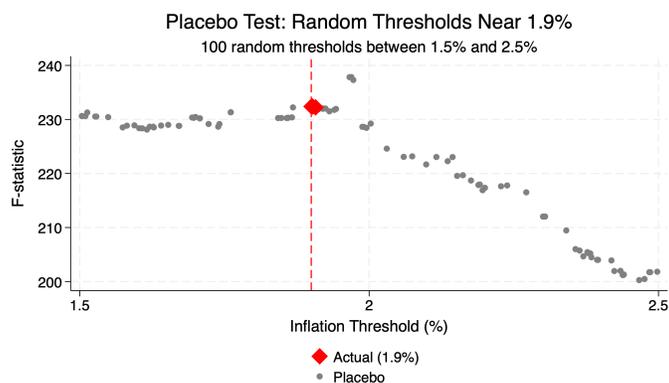


Figure 4: Placebo Test: Random Thresholds Near 1.9%

*Notes:* F-statistics for 100 random thresholds between  $1.5\%$  and  $2.5\%$ . Gray circles are placebo thresholds; red diamond is the actual  $1.9\%$  threshold. F-statistics remain stable through  $1.9\%$ , then decline—confirming  $1.9\%$  marks where filtering begins to shut off.

<sup>5</sup>The Akaike Information Criterion is minimised at  $2.75\%$  rather than  $1.9\%$ . The F-statistic and AIC optimise different objectives: the F-statistic maximises the statistical significance of behavioural differences between firm types, while the AIC optimises overall model fit. Since my theory concerns how firm types respond differently across regimes, the F-statistic provides a more targeted test. The AIC result suggests the  $1.9\%-2.75\%$  range represents a transition zone.

Second, I examine whether the pattern is visible in raw data without regression. Figure 5 shows average within-category price dispersion below versus above 1.9%. Dispersion falls from 10.83 to 9.74 percentage points—a 10% reduction. Below threshold, sticky firms filter, moving opposite to competitors and increasing dispersion. Above threshold, filtering shuts off, firms move together, and dispersion falls. This pattern emerges from descriptive statistics directly, without controls or functional form assumptions.

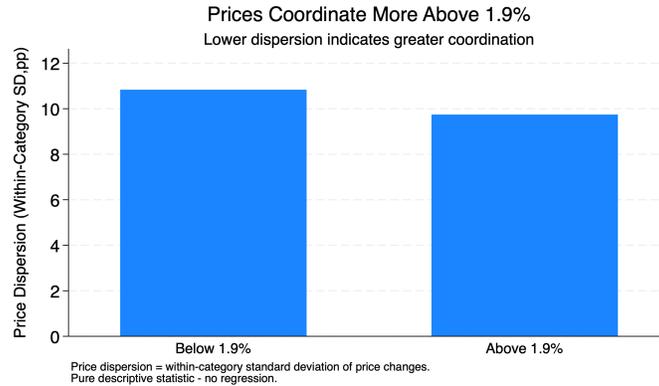


Figure 5: Descriptive Evidence: Price Dispersion Falls at 1.9%

*Notes:* Mean within-category price dispersion below versus above 1.9% sectoral inflation. Dispersion is the standard deviation of price changes within each COICOP category-quarter cell.

The empirical break occurs where theory predicts:  $|\pi_c| = \sigma_c$ .

### 7.3. Testing Threshold Effects

Having identified 1.9% as the critical threshold, I now test whether firms' responses to competitors differ across inflation regimes. Specifically, I ask: once I fix  $\tau = 1.9\%$ , do sticky and flexible firms behave differently?

Table 5 reports four nested specifications of equation (32), each adding one set of interactions. The baseline coefficient  $\beta_1$  measures the response to competitors' price changes when inflation is below 1.9% for firms at the frequency baseline (frequency = 0). A negative value indicates filtering—firms do not follow competitors. The threshold interaction  $\beta_2$  captures how this behaviour changes when inflation exceeds

1.9%. The frequency interaction  $\beta_3$  reveals the difference between flexible and sticky-price firms in low-inflation environments. The triple interaction  $\beta_4$  is the key test: whether crossing the inflation threshold affects sticky and flexible firms differently.

Table 5: Pricing Responses and Inflation Thresholds

	(1)	(2)	(3)	(4)
	Baseline	Threshold	Frequency	Full Model
Competitors' Price Change ( $\beta_1$ )	0.006 (0.017)	-0.028 (0.032)	-0.350*** (0.024)	-1.034*** (0.049)
Threshold $\times$ Competitors' Price ( $\beta_2$ )		0.052 (0.038)		1.004*** (0.060)
Competitors' Price $\times$ Frequency ( $\beta_3$ )			0.542*** (0.029)	1.472*** (0.066)
Threshold $\times$ Competitors' Price $\times$ Frequency ( $\beta_4$ )				-1.431*** (0.094)
Skewness (Std.)	0.048*** (0.004)	0.048*** (0.004)	0.050*** (0.004)	-0.014* (0.007)
Dispersion (Std.)	0.002 (0.006)	0.001 (0.006)	0.000 (0.006)	0.129*** (0.010)
Skewness $\times$ Frequency				0.229*** (0.032)
Dispersion $\times$ Frequency				-0.444*** (0.038)
Lagged Price Change	-0.271*** (0.001)	-0.271*** (0.001)	-0.271*** (0.001)	-0.271*** (0.001)
Aggregate Inflation	1.012*** (0.013)	1.012*** (0.013)	1.015*** (0.013)	1.010*** (0.013)
Sectoral Inflation (t-1)	0.313*** (0.006)	0.314*** (0.006)	0.317*** (0.006)	0.316*** (0.006)
Observations	5,802,099	5,802,099	5,802,099	5,802,099
Number of firms	350,304	350,304	350,304	350,304
R-squared (within)	0.0736	0.0736	0.0736	0.0737

*Notes:* The dependent variable is the firm's quarterly price change (in percent). All specifications include firm fixed effects. Standard errors clustered at the firm level in parentheses. The threshold indicator equals one when lagged annual sectoral inflation exceeds 1.9%. Frequency measures the firm's historical price adjustment frequency (proportion of periods with price changes). Competitors' Price Change is the weighted average price change of other firms in the same COICOP category. \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.1$ .

Column (1) establishes a crucial baseline: without accounting for heterogeneity, firms show no response to competitors ( $\beta_1 = 0.006$ , not significant). This null result explains why previous studies may have missed filtering behaviour—the average masks opposing responses. Column (2) adds the inflation threshold but maintains homogeneous firms. The threshold effect remains insignificant, confirming that inflation alone does not change average behaviour. Column (3) introduces firm heterogeneity through the frequency interaction. Now patterns emerge: flexible-price firms respond positively to competitors ( $\beta_3 = 0.542$ ), while the baseline remains near zero. Aggregation was hiding heterogeneous responses.

Column (4) presents the complete model. Four key findings emerge: First, in low-inflation regimes, sticky-price firms filter aggressively ( $\beta_1 = -1.034$ ): when competitors raise prices by 1 percentage point, sticky firms raise their own prices by 1 percentage point *less*. Second, flexible-price firms show weak positive response ( $\beta_1 + \beta_3 = 0.438$ )—they need not filter because they can re-adjust next period. Third, the large negative coefficient on the triple interaction ( $\beta_4 = -1.431$ ) means sticky firms’ filtering shuts off above the threshold, nearly eliminating the gap with flexible firms. Finally, above threshold, both types converge to near-zero responses ( $-0.03$  for sticky,  $0.01$  for flexible). The gap collapses from 1.47pp to 0.04pp. These fixed effects estimates are confirmed by instrumental variables analysis in Section 8.

These results confirm the model’s predictions. Below threshold, sticky firms filter noise ( $\beta_1 = -1.034$ ) while flexible firms do not need to ( $\beta_1 + \beta_3 = 0.438$ ). Above threshold, filtering shuts off and both types converge.

This has two implications for monetary policy. First, below the natural rate, sticky firms’ filtering provides automatic stabilisation—when some firms raise prices, sticky firms do not follow, dampening the shock. Above the natural rate, this self-correction mechanism shuts off. Second, the differential behaviour across firm types creates richer dynamics below threshold, potentially amplifying policy transmission.

#### 7.4. Economic Interpretation

To interpret these results, I calculate the response for representative firm types:

Table 6: Pricing Responses by Firm Type and Inflation Regime

Firm Type	Response to Competitors		
	Below 1.9%	Above 1.9%	Threshold Effect
<b>Panel A: Representative Firms</b>			
Sticky (freq = 0.1)	-0.887*** (0.054)	-0.026 (0.048)	+0.861*** (0.072)
Flexible (freq = 0.8)	0.144*** (0.041)	-0.027 (0.039)	-0.171** (0.057)
Difference	-1.031*** (0.068)	0.001 (0.062)	+1.032*** (0.092)
<b>Panel B: Alternative Cutoffs</b>			
Very Sticky (freq = 0.05)	-0.961***	-0.028	+0.933***
Median (freq = 0.23)	-0.695***	-0.022	+0.673***
Very Flexible (freq = 0.95)	0.365***	-0.031	-0.396***

*Notes:* This table reports pricing responses calculated from the regression coefficients in Table 5, Column (4). Response measures how a firm's price changes in response to a 1 percentage point increase in competitors' prices. Below 1.9% uses:  $\beta_1 + \beta_3 \times \text{freq}$ . Above 1.9% uses:  $(\beta_1 + \beta_2) + (\beta_3 + \beta_4) \times \text{freq}$ . Threshold effect is the difference between above and below responses. Standard errors calculated using the delta method. Panel A shows effects for sticky (10th percentile) and flexible (80th percentile) firms. Panel B shows robustness across the frequency distribution. \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.1$ .

Table 6 translates the statistical estimates into percentage-point responses when competitors raise prices by 1 percentage point. Panel A focuses on representative sticky-price firms (freq = 0.1) and flexible-price firms (freq = 0.8).

The results reveal striking heterogeneity. Below the 1.9% threshold, sticky-price firms filter aggressively (-0.887): when competitors raise prices by 1 percentage point, sticky firms raise by 0.89 percentage points *less*. This filtering—not incorporating competitor movements expected to revert—provides a natural stabilising force in low-inflation environments. Flexible-price firms show a weak positive response (0.144); they need not filter because they can re-adjust next period.

Above the threshold, both firm types converge to near-zero responses (-0.026 and -0.027), confirming that filtering shuts off. The near-identical responses are not statistically different from each other or from zero. The threshold effects are large

and opposite in sign: sticky-price firms experience a +0.861 percentage point increase (filtering shuts off), while flexible-price firms experience a  $-0.171$  percentage point decrease.

The bottom row of Panel A shows that the difference between sticky and flexible firms essentially disappears above the threshold (0.001 versus  $-1.031$  below). Panel B demonstrates robustness across the frequency distribution, with threshold effects monotonically increasing in stickiness—the stickier the firm, the more it was filtering, and the larger the change when filtering shuts off.

These calculations reveal the core mechanism: below threshold, sticky firms filter noise while flexible firms do not need to. Above threshold, filtering shuts off and both types converge. The dramatic difference in threshold effects (0.861 for sticky versus  $-0.171$  for flexible) explains why inflation dynamics fundamentally change at 1.9%.

## 8. Addressing the Reflection Problem

The reflection problem (Manski, 1993) poses a fundamental challenge for identifying responses to competitor prices. When firm  $i$ 's price and its competitors' prices are jointly determined, we cannot distinguish whether firm  $i$  responds to competitors, whether competitors respond to firm  $i$ , or whether both respond to common factors.

The patterns documented in Section 7—negative coefficients for sticky firms, near-zero for flexible firms, convergence above threshold—are consistent with filtering. But a sceptic might argue these patterns arise from simultaneity bias that differentially affects firm types.

### 8.1. *The Instrument*

I address this using variation in competitor prices that firm  $i$  cannot cause: prices in the same product category but in other regions.

Consider a clothing firm in London. Its competitors are other London clothing firms. When a national shock hits the clothing sector—rising cotton prices, supply chain disruptions, changing fashion—it affects clothing prices in Manchester and

Birmingham as well as London. But the London firm does not cause Manchester prices to move.

The instrument is the average price change in the same category but in other regions:

$$Z_{i,c,r,t} = \Delta \bar{p}_{c,-r,t} \quad (33)$$

where  $-r$  denotes all regions except firm  $i$ 's region. This isolates variation in competitor prices that originates outside the firm's market.

The identifying assumption is that regional demand shocks do not spill over within category, while national category shocks do. A shock to London clothing demand affects London clothing prices but not Manchester clothing prices. A shock to cotton prices affects clothing prices in all regions. The instrument captures the latter while excluding the former.

## 8.2. First Stage: Relevance

For the instrument to work, other-region prices must predict own-region competitor prices. Table 7 reports the first stage.

Table 7: First Stage: Other-Region Prices Predict Own-Region Competitor Prices

	Own-Region Competitor Prices
Other-Region Prices	0.922*** (0.006)
F-statistic	22,413
Observations	5,801,956

*Notes:* Robust standard errors clustered by product in parentheses. \*\*\*  $p < 0.01$ .

The first stage is strong. When other-region prices rise by 1 percentage point, own-region competitor prices rise by 0.92 percentage points. The F-statistic of 22,413 far exceeds conventional thresholds for instrument strength (e.g.,  $F > 10$ ). National category-level shocks transmit strongly across regions, as expected.

### 8.3. Institutional Validity

The validity of the instrument rests on the assumption that local demand shocks do not contaminate other-region prices. The structure of UK retail strongly supports this assumption.

I examine the geographic presence of firms using unique shop identifiers. Table 8 reveals that the sample is dominated by national chains operating across multiple regions.

Table 8: Geographic Presence of Retail Firms

Number of Regions	Share of Shops	Cumulative
2–10	26.5%	26.5%
11–50	41.2%	67.7%
51–100	10.2%	77.9%
>100	22.1%	100.0%
Median regions per shop	21	

*Notes:* Distribution of shop codes by number of UK regions in which they appear. No shop appears in fewer than 2 regions. The sample is dominated by national retail chains.

No shop in the sample operates in a single region. The median shop operates in 21 regions; over 20% operate in more than 100 regions. This confirms that pricing decisions are made by national headquarters responding to aggregate conditions—not by independent local agents responding to local demand.

This market structure strengthens the instrument. A local demand shock in London—a heatwave, a festival, a new competitor—does not cause Tesco headquarters to change prices nationally. National chains set prices based on aggregate factors: supply chains, exchange rates, wholesale costs, national wage floors. When prices move in Manchester, it reflects these aggregate conditions, not London-specific demand.

Consequently, other-region prices provide a clean proxy for national shocks. The instrument isolates variation in competitor prices driven by aggregate factors while excluding local demand contamination. The reflection problem is broken: firm  $i$  in London cannot cause the national pricing decisions that move Manchester prices.

#### 8.4. Second Stage: The Filtering Pattern Under Exogenous Variation

I instrument own-region competitor prices and all interactions (with threshold and frequency) using other-region prices and corresponding interactions. Table 9 reports IV estimates alongside the fixed effects estimates from Section 7.

Table 9: Fixed Effects and IV Estimates of Response to Competitor Prices

	Fixed Effects		IV	
	Below 1.9%	Above 1.9%	Below 1.9%	Above 1.9%
Sticky firms (freq = 0.1)	-0.89***	-0.03	-0.51***	-0.12**
Flexible firms (freq = 0.8)	+0.14**	+0.02	+0.08	-0.02
Difference (Sticky - Flexible)	-1.03***	-0.05	-0.59***	-0.10

*Notes:* Coefficients represent the response to a 1 percentage point increase in competitor prices. Fixed effects estimates from Table 5, Column (4). IV instruments own-region competitor prices with other-region prices in same category. Both specifications include product fixed effects and controls for lagged own price, aggregate inflation, sectoral inflation, skewness, and dispersion. Robust standard errors clustered by product. \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ .

The filtering pattern survives instrumentation:

1. **Sticky firms filter below threshold.** The IV coefficient is  $-0.51$ —smaller than fixed effects but strongly negative and highly significant. When competitor prices rise for reasons originating outside the firm’s region, sticky firms still lean against the movement.
2. **Flexible firms do not filter.** The IV coefficient is  $+0.08$ , small and statistically indistinguishable from zero. Flexible firms show no systematic response to competitor movements.
3. **The difference attenuates above threshold.** The gap between sticky and flexible firms narrows from  $-0.59$  below threshold to  $-0.10$  above—an 83% reduction. Both firm types converge toward zero response above 1.9%.

#### 8.5. Why IV Coefficients Are Smaller—and Still Negative

The IV estimate ( $-0.51$ ) is smaller in magnitude than the fixed effects estimate ( $-0.89$ ), but remains strongly negative. Both facts are informative.

The attenuation reflects that IV isolates aggregate variation—national shocks that spill across regions. Fixed effects estimates include all variation, both aggregate and idiosyncratic. Since firms filter idiosyncratic movements more aggressively than aggregate ones, the fixed effects coefficient is more negative.

But a question remains: if IV isolates aggregate shocks, and aggregate shocks persist, why do sticky firms still filter? The theory predicts pass-through of aggregate shocks.

The answer reveals the core mechanism. Below threshold, the signal-to-noise ratio is low. Firms cannot reliably distinguish aggregate from idiosyncratic movements—regardless of the shock’s true origin. A national cotton price shock that raises competitor prices by 0.5% looks identical to idiosyncratic noise when sectoral volatility is 1.9%. The firm does not observe that the shock is aggregate. It filters.

This is the key insight: *the regime determines behaviour, not the shock*. Below threshold, firms filter because the information environment is noisy—even when the underlying shock happens to be aggregate. Above threshold, the signal is strong enough to reveal aggregate conditions, and filtering stops.

The IV results thus confirm that information frictions dominate. Sticky firms filter below threshold not because shocks are actually idiosyncratic, but because they *appear* idiosyncratic in a low-inflation environment. The natural rate of inflation marks where signal overcomes noise—where firms can distinguish aggregate from transitory movements.

Taken together, the reflection problem cannot explain these results. Reflection would bias sticky and flexible firms equally—both face identical simultaneity within their markets. Yet sticky firms show strongly negative coefficients while flexible firms show near-zero, and this heterogeneity emerges only below threshold.

The IV confirms that the filtering pattern reflects firm behaviour, not statistical artifact. When competitor prices rise due to shocks originating in other regions—variation that firm  $i$  could not cause—sticky firms still filter below threshold. The

causal interpretation holds: sticky-price firms discount competitor movements they expect to revert, and this filtering shuts off when inflation crosses the natural rate.

## 9. Threshold or Smooth Nonlinearity?

The behavioural change documented at 1.9% sectoral inflation could represent either a true discontinuous threshold or a region of steep but smooth nonlinearity. This distinction matters: a threshold implies regime change, while smooth nonlinearity suggests gradual intensification.

To compare, I estimate three models with identical control structures: (1) **Threshold Model**: Discrete jump when sectoral inflation exceeds 1.9%, (2) **Linear Interaction**: Responses vary linearly with sectoral inflation and (3) **Quadratic**: Smooth nonlinear variation with sectoral inflation.

All three achieve virtually identical aggregate fit ( $AIC \approx 4.32 \times 10^7$ ), yet threshold effects remain jointly significant ( $F = 144.11$ ,  $p < 0.001$ ). Smooth models approximate average patterns across sectors and firms, but this aggregate equivalence masks differences at the firm level.

### 9.1. The Micro-Macro Disconnect

The critical test: can smooth models capture what happens to sticky firms when their sector crosses 1.9%? I calculate the implied behavioural change for sticky firms (frequency = 0.2) when sectoral inflation moves from 1.8% to 2.0%.

Table 10: Sticky Firms' Response When Sectoral Inflation Crosses 1.9%

Model	Behavioural Change (pp)	% of Threshold Effect Captured	Interpretation
Threshold Model	0.718	100.0%	Discrete jump
Linear Interaction	0.001	0.2%	Misses 99.8%
Quadratic	0.003	0.4%	Misses 99.6%

*Notes:* Predicted change in response to competitors' prices for sticky-price firms (frequency = 0.2) when sectoral inflation crosses 1.9%. Behavioural change measured as the difference in response when sectoral inflation moves from 1.8% to 2.0%.

Smooth models fail because they spread behavioural change across the entire inflation range. Moving from 1.8% to 2.0% is just 0.2 percentage points—a tiny step

on a smooth gradient. The threshold model concentrates all change at 1.9%, capturing the discrete shift from filtering to pass-through that sticky firms experience.

This is precisely what theory predicts. At  $|\pi_c| = \sigma_c$ , inference flips from "probably noise" to "probably aggregate." Smooth models, by construction, cannot capture this.

## 10. Robustness

The main results establish that sticky-price firms filter competitor signals below the 1.9% threshold, and that this filtering shuts off above threshold. Two potential concerns arise: (i) price adjustment frequency may be endogenous to expected inflation, and (ii) the statistical inference may be sensitive to clustering assumptions. I address each in turn.

### 10.1. Pre-Determined Frequency

A potential concern is that price adjustment frequency may be endogenous—firms expecting high inflation might choose greater flexibility. To address this, I re-estimate the main specification using frequency measured only from a pre-sample period (2003–2006) to predict behaviour in the post-sample (2007–2021). The correlation between pre-sample and full-sample frequency is 0.68, indicating substantial persistence in firm-level pricing behaviour.

The triple interaction coefficient remains strongly significant ( $\beta_4 = -1.16$ ,  $t = -8.29$ ), with magnitude comparable to the baseline ( $-1.43$ ). Firms classified as "sticky" based on their 2003–2006 behaviour continue to exhibit filtering in subsequent years when sectoral inflation falls below 1.9%. This confirms that the regime-dependent filtering pattern is not driven by endogenous frequency choice.

### 10.2. Clustering and Specification

Table 11 examines the sensitivity of the results to alternative econometric specifications. Column (1) reproduces the baseline with firm-level clustering. Column (2) implements two-way clustering at the firm and category-time levels (Cameron et al., 2011). This is a highly conservative specification that accounts for the fact that

the key regressor (competitor price changes) and the threshold variable vary at the category-time level, allowing for arbitrary correlation of errors within markets. Column (3) further saturates the model with time fixed effects, absorbing all aggregate variation.

Table 11: Robustness to Alternative Specifications

	(1)	(2)	(3)
	Baseline	Two-way Clustering	Two-way + Time FE
Competitors' Price Change ( $\beta_1$ )	-1.034*** (0.049)	-1.034*** (0.231)	-0.991*** (0.239)
Threshold $\times$ Competitors' Price ( $\beta_2$ )	1.004*** (0.060)	1.004*** (0.287)	0.963*** (0.289)
Competitors' Price $\times$ Frequency ( $\beta_3$ )	1.472*** (0.066)	1.472*** (0.528)	1.481*** (0.535)
Threshold $\times$ Competitors' Price $\times$ Frequency ( $\beta_4$ )	-1.431*** (0.094)	-1.431* (0.776)	-1.404* (0.785)
Controls	Yes	Yes	Yes
Aggregate Inflation	Yes	Yes	Absorbed
Firm FE	Yes	Yes	Yes
Time FE	No	No	Yes
Observations	5,802,099	5,802,099	5,676,549
R-squared	0.074	0.074	0.104
Clusters (firm)	350,304	350,304	330,541
Clusters (category-time)	–	4,792	4,753

*Notes:* Column (1) reproduces baseline with firm-level clustering. Column (2) implements two-way clustering at firm and category-time levels following Cameron et al. (2011). Column (3) adds time fixed effects, absorbing aggregate inflation. Threshold is an indicator equal to 1 if sectoral inflation exceeds 1.9%. Controls include lagged own price change, sectoral inflation, skewness and dispersion interactions. Standard errors in parentheses. \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.1$ .

The coefficient estimates are remarkably stable across all specifications. The key triple interaction term ( $\beta_4$ ) remains effectively unchanged (moving from  $-1.431$  to  $-1.404$ ), indicating that the estimated effect is not driven by time-varying aggregate

confounders or specific error structures.

As expected, accounting for intra-market correlation in Column (2) increases standard errors, as the effective degrees of freedom are determined by the number of category-time clusters (4,792) rather than observations. Despite this rigorous penalty, the triple interaction remains statistically significant ( $p < 0.10$ ) and economically large.

The magnitude of the effect is robust. The coefficient of  $-1.4$  implies that crossing the threshold induces a substantial rotation in pricing behaviour. Specifically, sticky firms' filtering coefficient shifts by  $+0.86$  percentage points (from  $-1.03$  to  $-0.17$ ), effectively shutting down the strategic substitution observed in the low-inflation regime. Flexible firms, in contrast, show minimal change. This robustness to two-way clustering confirms that the regime-dependent heterogeneity is a systematic feature of the data, not an artefact of precision.

## 11. Micro Filters, Macro Regimes

The micro-evidence establishes that sticky-price firms filter competitor signals when sectoral inflation falls below sectoral volatility. I now derive the implications for aggregate dynamics.

**Aggregation.** Let  $\omega_c$  denote sector  $c$ 's expenditure weight. Aggregate inflation persistence can be approximated as the weighted average of sectoral pass-through rates:

$$\rho^{agg} \approx \sum_c \omega_c \cdot \rho_c(\pi_c, \sigma_c) \quad (34)$$

where the effective pass-through  $\rho_c$  is state-dependent. For sticky sectors,  $\rho_c \approx 0$  (filtering) when  $|\pi_c| < \sigma_c$ , and  $\rho_c \rightarrow 1$  (pass-through) when  $|\pi_c| > \sigma_c$ . For flexible sectors,  $\rho_c \approx 1$  regardless of the regime, as the option to re-adjust renders filtering unnecessary.

**Why a common threshold works.** If every sector had a different  $\sigma_c$ , aggregate dynamics would be smooth, with no distinct threshold. However, the data reveals a

structural clustering.

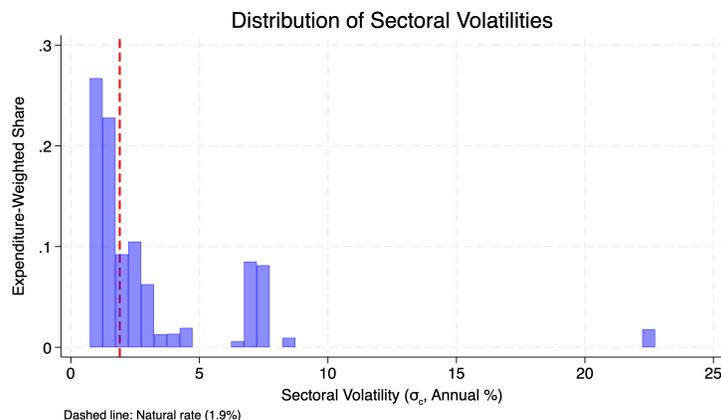


Figure 6: Distribution of Sectoral Volatilities

*Notes:* Expenditure-weighted distribution of sectoral inflation volatilities ( $\sigma_c$ ). Dashed line indicates the natural rate (1.9%).

Figure 6 plots the expenditure-weighted distribution of sectoral volatilities ( $\sigma_c$ ). Sectors with  $\sigma_c < 3\%$  account for approximately 90% of expenditure; their weighted-mean volatility is 1.9%. The high-volatility sectors in the right tail ( $\sigma_c > 5\%$ )—energy, fresh food, and other commodities—are precisely those with flexible prices, consistent with Figure 2. Because flexible-price sectors do not filter, their threshold location is irrelevant for persistence dynamics.

**The Aggregate Regime Shift.** This structural clustering implies that when aggregate inflation crosses 1.9%, the bulk of the sticky-price economy simultaneously shifts from a “filtering” regime to a “pass-through” regime. The aggregate threshold is not an additional free parameter; it is an emergent property derived from the distribution of sectoral volatilities weighted by their macroeconomic importance. Under the current sectoral structure, aggregate inflation serves as an effective summary statistic for the monetary transmission mechanism analysed in Section 11.2.

**Implications for Core Inflation.** This framework provides a structural rationale for monitoring core inflation. Central banks typically strip out volatile components (energy, fresh food) on the grounds that they are “noisy.” The analysis reveals

a deeper justification: high-volatility sectors are precisely those with flexible prices that do not filter. Their inflation dynamics are irrelevant for aggregate persistence. Core inflation—which tracks the low-volatility, sticky-price sectors clustered around  $\sigma_c \approx 1.9\%$ —is the appropriate measure for assessing whether the economy has crossed the natural rate threshold. Headline inflation exceeding the natural rate due solely to an energy spike does not trigger the regime shift; core inflation crossing 1.9% does.

### 11.1. Implications for Inflation Dynamics and Monetary Transmission

If the 1.9% threshold marks a shift in firm behaviour, we should observe three consequences: (i) lower price dispersion as filtering shuts off, (ii) higher inflation persistence as shocks propagate, and (iii) weaker monetary policy transmission as prices coordinate. I test each prediction.

#### 11.1.1. Price Dispersion

When sticky firms stop filtering above 1.9%, pricing patterns should become more uniform. For each COICOP category  $c$  at time  $t$ , I calculate weighted price dispersion:

$$\sigma_{c,t} = \sqrt{\sum_{i \in c} \omega_{i,t} (\Delta p_{i,t} - \Delta \bar{p}_{c,t})^2} \quad (35)$$

Table 12: Price Dispersion at the 1.9% Threshold

	(1) Price Dispersion (Category-level)	(2) Price Synchronisation (Firm-level deviation)
Above 1.9% Threshold	-1.010*** (0.005)	-0.416*** (0.009)
Constant	10.629*** (0.007)	5.077*** (0.009)
Percentage Change	-9.5%	-8.2%
Observations	6,403,124	6,403,124

*Notes:* Column (1): expenditure-weighted standard deviation of price changes within COICOP category-date cells. Column (2): absolute deviation of individual price changes from the weighted category-date mean. Both specifications include category and time fixed effects. Standard errors clustered at firm level. \*\*\*  $p < 0.01$ .

Table 12 confirms the prediction. Above threshold, price dispersion falls by 9.5%. Below threshold, sticky firms filter noise, moving opposite to competitors and creating

dispersion. Above threshold, filtering shuts off and firms move together.

### 11.1.2. Inflation Persistence

Lower dispersion should translate into higher inflation persistence. Table 13 tests this by splitting the sample based on dispersion persistence.

Table 13: Dispersion Persistence and Inflation Persistence

	Low Dispersion Persistence	High Dispersion Persistence	Difference
<b>Dispersion Dynamics</b>			
AR(1) coefficient ( $\rho_{disp}$ )	0.407** (0.165)	0.787*** (0.080)	+0.380 ( $p < 0.01$ )
Half-life (quarters)	0.8	2.9	3.6×
<b>Inflation Dynamics</b>			
AR(1) coefficient ( $\rho_{\pi}$ )	-0.080 (0.185)	0.444*** (0.096)	+0.524 ( $p < 0.001$ )
Half-life (quarters)	0 (mean-rev)	0.9	Persistent

*Notes:* Sample split by median of dispersion persistence. Dispersion measured as weighted cross-sectional standard deviation. Robust standard errors in parentheses. \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.10$ .

When dispersion is transient ( $\rho_{disp} = 0.41$ ), inflation exhibits no persistence ( $\rho_{\pi} = -0.08$ )—the below-threshold regime. When dispersion is persistent ( $\rho_{disp} = 0.79$ ), inflation inherits this persistence ( $\rho_{\pi} = 0.44$ ). The jump from mean-reversion to 0.44 matches the regime change at 1.9%.

### 11.2. Monetary Policy Transmission

The filtering mechanism implies that monetary shocks should transmit differently across regimes. Unlike the micro-level analysis, which conditions on sectoral inflation, monetary policy operates economy-wide. The structural clustering of sticky-price sectors around  $\sigma_c \approx 1.9\%$  means that when aggregate inflation crosses this threshold, the bulk of the economy shifts regime simultaneously. I therefore condition on aggregate inflation when analysing monetary transmission.

I test this using high-frequency identified monetary policy shocks from Braun et al.

(2025), estimating state-dependent local projections:

$$\begin{aligned}
 \Delta p_{i,c,t+h} = & \alpha_{i,h} + \beta_h^1 \text{MP}_t + \beta_h^2 \text{MP}_t \times \mathbb{I}(\pi_{t-1}^{agg} \geq 1.9\%) \\
 & + \beta_h^3 \text{MP}_t \times \text{freq}_i \\
 & + \beta_h^4 \text{MP}_t \times \mathbb{I}(\pi_{t-1}^{agg} \geq 1.9\%) \times \text{freq}_i \\
 & + \Gamma_h X_{i,c,t} + \varepsilon_{i,c,t+h}
 \end{aligned} \tag{36}$$

The specification interacts the monetary shock with both the inflation regime and firm-level price flexibility, allowing transmission to vary along both dimensions.

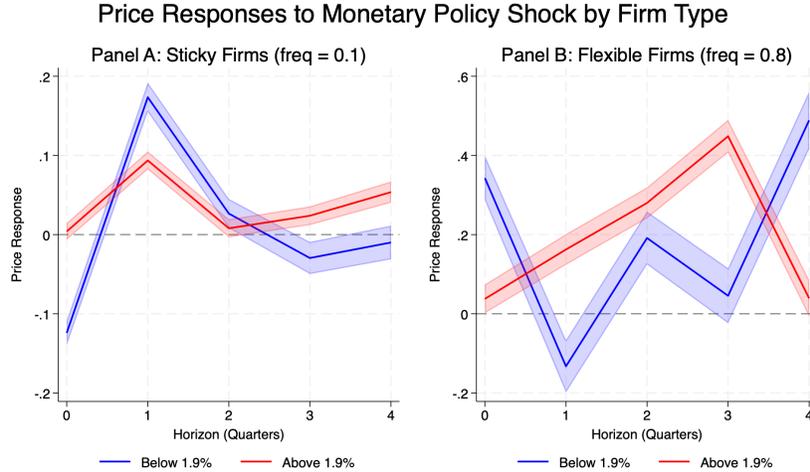


Figure 7: Price Responses to Monetary Policy Shocks by Firm Type

*Notes:* Local projection impulse responses to a one standard deviation expansionary monetary policy shock. Panel A: sticky firms (freq=0.1). Panel B: flexible firms (freq=0.8). Blue: aggregate inflation below 1.9%. Red: aggregate inflation above 1.9%. Shaded areas: 90% confidence intervals. Shocks identified using high-frequency surprises in short-sterling futures around Bank of England announcements. Sample: 2003–2021.

Figure 7 reveals the filtering mechanism in action. Panel A shows sticky firms. Below threshold (blue), these firms exhibit a negative response on impact—they move *opposite* to the expansionary shock. Unable to distinguish the monetary signal from noise, sticky firms treat it as idiosyncratic and lean against it. Above threshold (red), filtering disappears: sticky firms respond conventionally, with a small positive impact that persists.

Panel B shows flexible firms. Below threshold, responses are volatile—large positive movements followed by sharp reversals. This is the self-correction mechanism: flexible firms overshoot but quickly reverse. Above threshold, responses smooth out. Flexible firms no longer self-correct; they coordinate with the broader market.

Table 14: State-Dependent Inflation Persistence by Firm Type

	Q1 Sticky	Q2 Med-Low	Q3 Med-High	Q4 Flexible
<b>Panel A: Low Inflation Regime (<math>\pi_t &lt; 1.9\%</math> annual)</b>				
Persistence ( $\beta$ )	0.328** (0.132)	0.246** (0.119)	0.405*** (0.110)	-0.329*** (0.114)
<b>Panel B: High Inflation Regime (<math>\pi_t \geq 1.9\%</math> annual)</b>				
Persistence ( $\beta$ )	0.552*** (0.123)	0.324** (0.145)	0.318** (0.155)	-0.017 (0.166)
<b>Change (High – Low)</b>	+0.224	+0.078	-0.087	+0.312

*Notes:* AR(1) persistence estimates by frequency quartile and inflation regime. Low regime: annual inflation  $< 1.9\%$ . High regime:  $\geq 1.9\%$ . Standard errors from Newey-West with 4 lags. \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.10$ .

Table 14 quantifies this dual shift. Below threshold, flexible firms exhibit strong mean reversion ( $\beta = -0.329$ ): price shocks dissipate as firms correct overshooting. Above threshold, self-correction vanishes ( $\beta = -0.017$ ). Simultaneously, sticky firms become stickier: persistence rises from 0.328 to 0.552, a 68% increase.

The results reveal two distinct mechanisms operating below the natural rate. First, sticky firms filter shocks they cannot distinguish from noise, dampening price transmission. Second, flexible firms self-correct, reversing overshooting within 2–3 quarters. Together, these mechanisms create rich, heterogeneous price dynamics that allow monetary policy to gain traction on real activity.

Above the natural rate, both mechanisms shut off. Firms recognise aggregate shocks and coordinate their responses. Prices move together, dispersion falls, and the real effects of monetary policy diminish.

The policy implication is direct: monetary transmission operates primarily below

the natural rate of inflation. Targeting above 1.9% does not simply raise average inflation—it undermines the central bank’s ability to influence real activity.

## 12. Conclusion

This paper identifies a structural boundary in inflation dynamics: the natural rate of inflation. Using 6 million UK micro-price observations, I document that firm pricing behaviour fundamentally changes at 1.9%. Below this threshold, sticky-price firms filter competitor price movements, creating endogenous stability. Above it, filtering shuts off: price dispersion falls, inflation persistence rises, and monetary policy transmission weakens.

The threshold emerges from signal extraction. When sectoral inflation falls below typical idiosyncratic volatility ( $|\pi_c| < \sigma_c$ ), firms cannot distinguish aggregate from sector-specific shocks and optimally filter movements likely to revert. Once inflation exceeds sectoral noise, its magnitude reveals aggregate conditions. The natural rate is this boundary:  $|\pi^N| = \bar{\sigma}_c \approx 1.9\%$ .

This definition implies the natural rate is structurally determined by sectoral volatility—not an invariant optimum. The 2% target that central banks have converged upon likely reflects the typical volatility structure of advanced economies, providing a structural answer to why this seemingly arbitrary number has proved remarkably stable across countries and decades. Economies with greater sectoral heterogeneity should exhibit higher natural rates; those with more homogeneous structures, lower ones. Norway’s 2% target provides a natural experiment: if commodity exporters exhibit higher sectoral volatility, a target that sits comfortably above the UK’s natural rate might sit *below* Norway’s—implying Norwegian monetary policy fights against market forces that already stabilise prices.

For policymakers, the framework reveals a critical asymmetry. Operating below the natural rate wastes effort on stabilisation that markets already provide. Operating substantially above it activates persistence that is costly to reverse. Near the threshold, aggressive action is warranted to prevent crossing into the high-persistence

regime. Once inflation substantially exceeds the natural rate, policymakers face a different calculus: fighting persistence with tools whose effectiveness has diminished.

Formalising optimal policy in the presence of this structural threshold—including the welfare costs of crossing it—remains for future work.

## Appendix A. Appendix A: Value Function Approximation and Effective Horizon

I derive the quadratic approximation of the value function. Consistent with Appendix B and the identification strategy, I model the firm's problem in partial equilibrium relative to its sector. Time subscripts are suppressed where no ambiguity arises.

Define the firm's relative price as  $\hat{p} \equiv \log(P_i) - \log(P_c)$ , where  $P_c$  is the sectoral price index. The relevant inflation rate eroding this relative price is sectoral inflation  $\pi_c$ .

**Step 1: The Bellman Equation.** The value function for a firm with menu cost  $\psi$  and relative price  $\hat{p}$  satisfies:

$$V(\hat{p}, \Omega, \psi) = \max\{V^n(\hat{p}, \Omega, \psi), V^a(\Omega, \psi) - \psi w\} \quad (\text{A.1})$$

where  $V^n$  is the value of non-adjustment and  $V^a$  is the value of adjustment:

$$V^n(\hat{p}, \Omega, \psi) = \Pi(\hat{p}, \Omega) + \beta \mathbb{E}[V(\hat{p} - \pi'_c, \Omega', \psi)] \quad (\text{A.2})$$

where  $\pi'_c$  denotes next period's sectoral inflation.

**Step 2: The Non-Adjustment Region.** For prices sufficiently close to the optimal price  $\hat{p}^*$ , the firm optimally chooses not to adjust:

$$V(\hat{p}, \Omega, \psi) = \Pi(\hat{p}, \Omega) + \beta \mathbb{E}[V(\hat{p} - \pi'_c, \Omega', \psi)] \quad (\text{A.3})$$

**Step 3: Taylor Expansion.** Let  $V^*(\Omega, \psi)$  denote the frictionless value. For  $\hat{p}$  near  $\hat{p}^*$ :

$$V(\hat{p}, \Omega, \psi) \approx V(\hat{p}^*, \Omega, \psi) + V_{\hat{p}}(\hat{p}^*, \Omega, \psi)(\hat{p} - \hat{p}^*) + \frac{1}{2} V_{\hat{p}\hat{p}}(\hat{p}^*, \Omega, \psi)(\hat{p} - \hat{p}^*)^2 \quad (\text{A.4})$$

By the envelope theorem,  $V_{\hat{p}}(\hat{p}^*, \Omega, \psi) = 0$  at the optimum.

**Step 4: The Second Derivative.** Differentiating the Bellman equation twice

with respect to  $\hat{p}$ :

$$V_{\hat{p}\hat{p}}(\hat{p}, \Omega, \psi) = \Pi_{\hat{p}\hat{p}}(\hat{p}, \Omega) + \beta \mathbb{E}[V_{\hat{p}\hat{p}}(\hat{p} - \pi'_c, \Omega', \psi)] \quad (\text{A.5})$$

**Step 5: The Adjustment Probability.** With probability  $\lambda(\psi)$ , the firm hits the adjustment boundary next period and resets. At that point, the continuation value is independent of the current price deviation:

$$V_{\hat{p}\hat{p}}(\hat{p}^*, \Omega, \psi) = \Pi_{\hat{p}\hat{p}}(\hat{p}^*, \Omega) + \beta(1 - \lambda(\psi))V_{\hat{p}\hat{p}}(\hat{p}^*, \Omega, \psi) \quad (\text{A.6})$$

where the second equality assumes approximate stationarity. Solving:

$$V_{\hat{p}\hat{p}}(\hat{p}^*, \Omega, \psi) = \frac{\Pi_{\hat{p}\hat{p}}(\hat{p}^*, \Omega)}{1 - \beta(1 - \lambda(\psi))} \quad (\text{A.7})$$

**Step 6: The Effective Pricing Horizon.** Define the profit function curvature as  $\kappa_0 \equiv -\Pi_{\hat{p}\hat{p}}(\hat{p}^*, \Omega) > 0$ . The **effective pricing horizon** is:

$$\Phi(\psi) \equiv \frac{1}{1 - \beta(1 - \lambda(\psi))} \quad (\text{A.8})$$

The value function curvature is therefore:

$$V_{\hat{p}\hat{p}}(\hat{p}^*, \Omega, \psi) = -\kappa_0 \Phi(\psi) \quad (\text{A.9})$$

Firms with high menu costs adjust infrequently (low  $\lambda$ ), giving them longer effective horizons (high  $\Phi$ ). They weight price deviations more heavily because they are stuck with them longer.

**Step 7: The Adjustment Hazard.** The hazard  $\lambda(\psi)$  is determined by the  $(S, s)$  bands. At the adjustment boundary, indifference requires (Dixit, 1991, equation 1):

$$\psi = \frac{1}{2}\kappa_0 \bar{x}(\psi)^2 \implies \bar{x}(\psi) = \sqrt{\frac{2\psi}{\kappa_0}} \quad (\text{A.10})$$

where  $\bar{x}(\psi)$  is the width of the inaction region. The hazard rate is the inverse of the

expected hitting time for a Brownian motion to reach the barrier:

$$\lambda(\psi) = \frac{\sigma_c^2}{\bar{x}(\psi)^2} = \frac{\kappa_0 \sigma_c^2}{2\psi} \quad (\text{A.11})$$

This expression reveals a key connection: the adjustment hazard depends on sectoral volatility  $\sigma_c^2$ . High-volatility sectors experience frequent shocks that push prices outside the inaction band, generating high  $\lambda$  and short pricing horizons  $\Phi$ . Low-volatility sectors experience infrequent shocks, generating low  $\lambda$  and long horizons. The model thus endogenously generates the strong correlation between sectoral volatility and price flexibility documented in Section 2.

**Final Approximation.** Substituting into the Taylor expansion:

$$V(\hat{p}, \Omega, \psi) \approx V^*(\Omega, \psi) - \frac{1}{2} \kappa_0 \Phi(\psi) (\hat{p} - \hat{p}^*)^2 \quad (\text{A.12})$$

The value function derivative, used in Appendix B, is:

$$V_{\hat{p}}(\hat{p}, \Omega, \psi) \approx -\kappa_0 \Phi(\psi) (\hat{p} - \hat{p}^*) \quad (\text{A.13})$$

## Appendix B. Appendix B: Derivation of Optimal Reset Price

The firm with menu cost  $\psi$  chooses a reset price  $\hat{p}_t^R$  to maximise value. Using the definitions and approximations from Appendix A, the firm solves:

$$\max_{\hat{p}_t^R} \{ \Pi(\hat{p}_t^R, \Omega_t) + \beta \mathbb{E}_t [V(\hat{p}_t^R - \pi_{c,t+1}, \Omega_{t+1}, \psi)] \} \quad (\text{B.1})$$

The first-order condition is:

$$\Pi_{\hat{p}}(\hat{p}_t^R, \Omega_t) + \beta \mathbb{E}_t [V_{\hat{p}}(\hat{p}_t^R - \pi_{c,t+1}, \Omega_{t+1}, \psi)] = 0 \quad (\text{B.2})$$

**Value Function Approximation.** I approximate the profit function  $\Pi$  as quadratic around the frictionless optimal price  $\hat{p}_t^*$ . From Appendix A, the profit function curvature is  $\kappa_0 \equiv -\Pi_{\hat{p}\hat{p}}(\hat{p}_t^*) > 0$ . The marginal profit is approximately:

$$\Pi_{\hat{p}}(\hat{p}_t^R, \Omega_t) \approx -\kappa_0 (\hat{p}_t^R - \hat{p}_t^*) \quad (\text{B.3})$$

From equation (A.13), the derivative of the value function is:

$$V_{\hat{p}}(\hat{p}, \Omega, \psi) \approx -\kappa_0 \Phi(\psi)(\hat{p} - \hat{p}^*) \quad (\text{B.4})$$

Define the reset price gap as  $x_t \equiv \hat{p}_t^R - \hat{p}_t^*$ .

**Evolution of the Target.** The firm anticipates that its optimal target  $\hat{p}_t^*$  will evolve in response to sectoral inflation. This response is governed by the effective pass-through  $\rho(\theta)$  derived in Section 4:

$$\rho(\theta) = 1 + \theta(\Lambda - 1) \quad (\text{B.5})$$

where  $\theta(\pi_c) = \frac{\sigma_c^2}{\sigma_c^2 + \pi_c^2}$  is the probability the shock is idiosyncratic, and  $\Lambda \in (0, 1]$  is the horizon filter from Section 5 (equation 24). The target evolves as:

$$\hat{p}_{t+1}^* - \hat{p}_t^* \approx \rho \cdot \pi_{c,t+1} \quad (\text{B.6})$$

The intuition:

- If the firm perceives sectoral inflation as aggregate (high signal,  $\theta \rightarrow 0$ ), it adjusts its target one-for-one with competitors ( $\rho \rightarrow 1$ ).
- If the firm perceives sectoral inflation as idiosyncratic noise (low signal,  $\theta \rightarrow 1$ ), it filters the movement, adjusting its target by less ( $\rho \rightarrow \Lambda$ ).

**The Price Gap.** The expected price gap at  $t + 1$  is the reset price, eroded by sectoral inflation, minus the new target:

$$\begin{aligned} \text{Gap}_{t+1} &= (\hat{p}_t^R - \pi_{c,t+1}) - \hat{p}_{t+1}^* \\ &= (\hat{p}_t^R - \pi_{c,t+1}) - (\hat{p}_t^* + \rho\pi_{c,t+1}) \\ &= (\hat{p}_t^R - \hat{p}_t^*) - (1 + \rho)\pi_{c,t+1} \\ &= x_t - (1 + \rho)\pi_{c,t+1} \end{aligned} \quad (\text{B.7})$$

**Solution.** Substituting the marginal profit and value function approximations

into the first-order condition:

$$-\kappa_0 x_t - \beta \kappa_0 \Phi(\psi) \mathbb{E}_t[x_t - (1 + \rho)\pi_{c,t+1}] = 0 \quad (\text{B.8})$$

Since  $x_t$  is determined at time  $t$ , I can pull it out of the expectation. Dividing by  $-\kappa_0$ :

$$x_t + \beta \Phi(\psi) x_t - \beta \Phi(\psi) (1 + \rho) \mathbb{E}_t[\pi_{c,t+1}] = 0 \quad (\text{B.9})$$

Rearranging:

$$x_t (1 + \beta \Phi(\psi)) = \beta \Phi(\psi) (1 + \rho) \mathbb{E}_t[\pi_{c,t+1}] \quad (\text{B.10})$$

This yields the optimal reset price  $\hat{p}_t^R = \hat{p}_t^* + x_t$ :

$$\hat{p}_t^R = \hat{p}_t^* + \underbrace{\frac{\beta \Phi(\psi) (1 + \rho)}{1 + \beta \Phi(\psi)}}_{\alpha_1(\Phi, \rho)} \mathbb{E}_t[\pi_{c,t+1}] \quad (\text{B.11})$$

**Interpretation.** This result formalises the interaction between stickiness ( $\Phi$ ) and filtering ( $\rho$ ):

- **Filtering Regime** ( $\theta \rightarrow 1, \rho \rightarrow \Lambda$ ): Below threshold, firms filter signals. For sticky firms with long pricing horizons ( $\Phi$  large,  $\Lambda$  small),  $\rho \approx \Lambda$  is close to zero, so  $(1 + \rho) \approx 1$ . The firm accounts only for mechanical erosion of its relative price, ignoring expected competitor movements because they are likely to revert. For flexible firms ( $\Lambda \approx 1$ ), filtering provides little benefit—they can re-adjust next period.
- **Pass-Through Regime** ( $\theta \rightarrow 0, \rho \rightarrow 1$ ): Above threshold, signals are clear. The term  $(1 + \rho) \approx 2$ . All firms respond aggressively to expected sectoral inflation, accounting for both mechanical erosion and persistent competitor price increases.

## References

- Afrouzi, H., Halac, M., Rogoff, K.S., Yared, P., 2023. Monetary Policy without Commitment. NBER Working Paper 31207. National Bureau of Economic Research. URL: <https://doi.org/10.3386/w31207>, doi:10.3386/w31207.
- Amiti, M., Itskhoki, O., Konings, J., 2019. International shocks, variable markups, and domestic prices. *The Review of Economic Studies* 86, 2356–2402.
- Blanco, A., 2021. Optimal inflation target in an economy with menu costs and a zero lower bound. *American Economic Journal: Macroeconomics* 13, 108–41. URL: <https://www.aeaweb.org/articles?id=10.1257/mac.20180198>, doi:10.1257/mac.20180198.
- Braun, R., Miranda-Agrippino, S., Saha, T., 2025. Measuring monetary policy in the uk: The uk monetary policy event-study database. *Journal of Monetary Economics* 149, 103645.
- Cameron, A.C., Gelbach, J.B., Miller, D.L., 2011. Robust inference with multiway clustering. *Journal of Business & Economic Statistics* 29, 238–249.
- Coibion, O., Gorodnichenko, Y., Wieland, J., 2012. The optimal inflation rate in new keynesian models: Should central banks raise their inflation targets in light of the zero lower bound? *The Review of Economic Studies* 79, 1371–1406.
- Davies, R., 2021. Prices and inflation in a pandemic: A micro data approach. URL: <http://cep.lse.ac.uk/pubs/download/dp017.pdf>.
- Dixit, A., 1991. Analytical approximations in models of hysteresis. *Review of Economic Studies* 58, 141–151.
- Friedman, M., 1968. The role of monetary policy. *The American Economic Review* 58, 1–17.
- Gagliardone, L., Gertler, M., Lenzu, S., Tielens, J., 2025. Anatomy of the phillips curve: Micro evidence and macro implications. *American Economic Review* Forthcoming.

- Golosov, M., Lucas, R.E., 2007. Menu costs and phillips curves. *Journal of Political Economy* 115, 171–199.
- 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](https://doi.org/10.1016/0022-0531(72)90142-1).
- Manski, C.F., 1993. Identification of endogenous social effects: The reflection problem. *Review of Economic Studies* 60, 531–542.
- Nakamura, E., Steinsson, J., 2010. Monetary non-neutrality in a multi-sector menu cost model. *Quarterly Journal of Economics* 125, 961–1013.
- Pasten, E., Schoenle, R., 2016. Rational inattention, multi-product firms and the neutrality of money. *Journal of Monetary Economics* 80, 1–16. doi:<https://doi.org/10.1016/j.jmoneco.2016.04.004>.
- Reis, R., 2006. Inattentive producers. *Review of Economic Studies* 73, 793–821.
- Sims, C.A., 2003. Implications of rational inattention. *Journal of Monetary Economics* 50, 665–690.
- Wicksell, K., 1898. *Interest and Prices: A Study of the Causes Regulating the Value of Money*. Macmillan, London. Translated by R.F. Kahn, 1936.