RIETI Discussion Paper Series 26-E-062

Same Structure, Different Norms: Identifying regime shifts

INOSE, Junya
Toyo University

The Research Institute of Economy, Trade and Industry
https://www.rieti.go.jp/en/

RIETI Discussion Paper Series 26-E-062
August 2026

Same Structure, Different Norms: Identifying regime shifts*

Junya INOSE
Toyo University

Abstract

An inflation regime can shift through Phillips-curve slopes, shock variances, or the nominal reference levels that guide price and wage setting. This paper makes the three margins compete within a nested Bayesian New Keynesian DSGE model with separate price and wage anchors. For Japan in 2013 and 2022 and the Volcker disinflation in the U.S. in 1979, allowing both anchors and variances to switch delivers the best fit, whereas allowing either slope to switch does not improve model performance. The anchors move differently across episodes. Only the price anchor shifts in 2013; both move roughly one for one in 2022; and the wage anchor falls about twice as far as the price anchor in 1979. The estimated wage anchors also align with wage-setting data excluded from estimation.

Keywords: expectational anchors; inflation regime shift; NK-DSGE; Bayesian estimation; price and wage anchors; Japan
JEL classification: E31, E52, E24, C11, C52

*This study is conducted as a part of the Project "Heterogeneity of Economic Agents and Challenges for the Japanese Economy" undertaken at the Research Institute of Economy, Trade and Industry (RIETI).
The draft of this paper was presented at the RIETI DP seminar for the paper. I would like to thank participants of the RIETI DP Seminar for their helpful comments.

1. Introduction

A change in the inflation regime can arise on three margins.[1] It can occur in the propagation mechanism, that is, in the slope of the Phillips curve. It can occur in the stochastic environment, that is, in the variance of shocks. Or it can occur in the nominal reference level to which prices and wages are anchored. The three look much alike in the data. But their implications differ fundamentally. The first two represent changes in propagation or in the stochastic environment; the third represents a shift in nominal reference levels. Distinguishing among them matters for assessing the persistence and reversibility of a regime change, and hence the normalization of monetary policy. The framework offered here draws that distinction inside a structural model.

The difficulty is that the distinction cannot be drawn from observation alone. One can fit UCSV models, regime-switching models, or dynamic factor models to price and wage series, either separately or jointly. For Japan, Ueno (2024) estimates the common trend component of price and wage inflation with a dynamic factor model, and shows that the link between the two recovered after COVID. Such exercises show that both trends moved. They do not identify what moved them. Suppose nominal prices and nominal wages shift upward at the same time. Three quite different situations are consistent with this. Prices and wages may have turned independently. One may have driven the other, as in a wage–price spiral. Or a common anchor may have moved both at once. All three appear in the data as "both trends rose." The same ambiguity arises for the axis that carries the shift. Did the slope rise, did the variance widen, or did the level of the anchor move? Again the three look alike. Separating them requires one structure that ties prices and wages together. That structure must make two things explicit: the real-wage relation that links them, and the transmission mechanism.

The distinction is drawn by Bayesian estimation of a small New Keynesian (NK) DSGE model (Galí and Monacelli, 2005; Ichiue et al., 2013). The core is standard, with Calvo (1983) price and wage rigidities. To it are added regime-specific steady-state anchors: the price anchor $\bar{\pi}$, the wage anchor $\bar{\pi}^w$, and the steady-state real rate $\bar{r}$.[2] Energy enters as an intermediate input (Bruno and Sachs, 1985). At the core of the identification, the level shift in anchors competes against two rivals within a single nested structure: changing slopes and changing variances. The slopes are selected by comparing marginal likelihoods over a grid that contains the values of Hazell et al. (2022) and Druant et al. (2009). Their invariance across regimes is then tested. The problem of disentangling slopes from expectations is shared with Hazell et al. (2022). They separate the two using cross-state variation and time fixed effects; this paper separates them through the orthogonality of first and second moments (Proposition 1). The separating identification does not depend on the sign of the anchor shift. The same framework therefore applies to upward and to downward re-anchoring alike. Japan's upward shift in 2022 and the downward shift of the Volcker disinflation are two such cases.

Identifying shifts in trends or in the central bank's target within a structural model is not new. Most of the earlier literature fixes one axis in advance and estimates its magnitude. Cogley

[1]Replication code and files are openly available at `https://doi.org/10.5281/zenodo.21861626` (version v1.0.0).

[2]Here $\bar{r}$ denotes the steady-state real (policy) rate implied by the steady state of the Taylor rule. It is distinct from the natural rate of interest r^* defined by the flexible-price equilibrium.

and Sbordone (2008) log-linearize around a shifting steady state and absorb trend inflation into the intercept, for prices only. Schorfheide (2005) builds in regime switching in the target inflation rate π^* together with learning. Closest to this paper in placing time-varying targets on both prices and wages is Scott (2024). There, however, the two targets are policy objectives in the central bank's loss function. They follow independent AR processes without a common break, and are identified from the interest rate. Structural work that makes the axes compete is scarce. One instance is Liu et al. (2011), henceforth LWZ, who let the two axes of shock variances and the inflation target compete by model comparison, and conclude that variances play the leading role in the postwar United States. In the alternative specification they report as favored by the data, the target enters the indexation rules for both prices and wages. That target is nonetheless a single common one. The possibility that the price anchor and the wage anchor move by different amounts is not representable within that structure.

This paper makes three contributions, on three different levels.

The first contribution is methodological. Rather than assigning a regime shift to one structural margin in advance, the framework makes competing margins compete within a single nested structural model. Means, slopes, and variances are evaluated on the same likelihood, so the data determine which axis carries the shift. Most of the earlier literature fixes one axis and estimates its magnitude; here the choice among axes is itself an object of inference. Proposition 1 shows that regime means and centered propagation dynamics do not confound each other under the maintained assumptions, and Monte Carlo experiments with known true values evaluate the finite-sample performance of that separation. Unlike LWZ, who impose slope invariance, the comparison here treats a change in the slopes as a competing margin. The design is also portable. Because the anchor estimates do not depend on how the slopes are calibrated, the calibration used here can be carried to another country as it stands; only two exogenous values must be built from local data, the energy share from an input-output table and the energy persistence from a regression (Appendix C).

The second contribution is substantive. The framework separately recovers a price-setting anchor $\bar{\pi}$ and a wage-setting anchor $\bar{\pi}^w$, allowing the two to move by different amounts. Both are estimated within a single structural system, jointly with the steady-state real policy rate $\bar{r}$. These are not the central bank's policy targets. They are private steady-state components entering firms' price setting and households' wage setting; $\bar{r}$ is an auxiliary steady-state component that closes the nominal interest rate block. Separating the price and wage anchors reveals an asymmetry that a single common target cannot represent. In the United States in 1979 the downward shift in the wage anchor is about twice that in the price anchor, and the 90% interval for their difference excludes zero. In Japan in 2013 the price-anchor shift is detected, whereas the credible interval for the wage-anchor shift includes zero. The one-for-one shift in Japan in 2022 is likewise an estimation result, not an imposed restriction.

The third contribution is empirical. The same design is applied to two upward episodes in Japan and to the downward Volcker disinflation. The U.S. application therefore tests whether the inferred separation between nominal reference levels and propagation is specific to upward shifts, or also appears in a canonical disinflation.[3]

[3]On the question of which axis a regime change should be attributed to, there is an accumulated literature on both the structural and the reduced-form side. On the variance axis, the structural DSGE with stochastic

The framework is applied to three episodes: Japan in 2013 and 2022, and the U.S. disinflation of 1979. Three main results follow. First, regime switching is not supported for either the price or the wage Phillips-curve slope in any of the three episodes. Changing the treatment of the slopes also leaves the estimated anchor shifts unchanged in sign, and they move little over the range of calibrations the data support (Appendix C). A regime change is adequately explained neither by anchors alone nor by variances alone. In all three episodes the preferred specification allows both to switch. Which of the two dominates on its own varies by episode: anchors alone dominate in the two Japanese episodes, variances alone in the United States.

Second, the wage anchor exhibits a different estimated pattern in each episode. In 2013 only the rise in the price anchor is detected. In 2022 the price and wage anchors move together, close to one for one. In 1979 the wage anchor moves about twice as far as the price anchor. Each of these is a within-episode posterior; the differences across episodes are not themselves significant, given the short samples. The contrast nonetheless shows that a movement in the nominal price anchor need not be accompanied by a simultaneous movement of equal size in the wage anchor.

Third, the two-way wage–price feedback is not operating in a self-reinforcing manner. The estimated loop gain cannot be distinguished from zero in the four Japanese regimes, and is significantly negative in the two U.S. regimes. A positive self-amplifying loop appears in none of these episodes. The region that would permit strong positive feedback fails the determinacy condition under the common calibration used here. In addition, the estimated wage anchor matches, in level, wage-setting data that are not used in the structural estimation: the *shuntō* base pay increase in Japan and major collective bargaining settlements in the United States. This is evidence from outside the model, and it supports reading the estimated anchor as an economically meaningful benchmark for wage setting rather than as a mere intercept. The idea that the benchmark for wage setting shifts discontinuously between periods goes back to Perry (1980) and Mitchell (1985). Mitchell measured it as an intercept shift in a single-sector wage regression. The present framework extends that idea by identifying separate price and wage anchors jointly within a structural model, and by applying the same design to Japan and the United States.

The question posed here also connects to recent international debate. Whether a change in the slope reflects a permanent change in structural parameters, or a temporary amplification activated when labor markets are tight, has been raised for the post-pandemic U.S. Phillips curve as well (Jiang, 2025). In reduced form, Ball and Mazumder (2019) show that the anchoring of expectations to a target shifted the Phillips curve from an accelerationist form to a level form, and represent this as a shift in a single parameter. The present paper is the structural counterpart of that insight. For Japan, level shifts in trend inflation and flattening during the low-inflation period have been discussed (Kaihatsu and Nakajima, 2018; Kaihatsu et al., 2023; Kishaba and Okuda, 2025), as have changes in the anchoring of inflation expectations (Hogen and Okuma, 2025; Fukuda and Soma, 2025). Policymakers, too, have described the developments since 2022 as a historic shift away from the wage and price "norm" (Watanabe, 2025; Takata, 2026). This

volatility of Justiniano and Primiceri (2008) sits alongside the reduced-form contributions of Primiceri (2005); Cogley and Sargent (2005); Sims and Zha (2006). Attribution to regimes in policy coefficients is pursued by Bianchi (2013); Davig and Leeper (2007), and drift in the target by Cogley et al. (2010). Rather than treating these separately, this paper makes the three margins compete within a common regime-switching Kalman filter (Section 3.2).

paper approximates that colloquial "norm" operationally, by the level shift in the steady-state anchors of prices and wages (Section 4.2).

The rest of the paper is organized as follows. Section 2 sets out the model and the identification strategy. Section 3 presents the main results question by question. It asks three questions. How did the anchors move? Which axis moved: anchors, slopes, or variances? And is the wage–price feedback self-reinforcing? Each is answered across the three episodes. It then provides the link and decomposition connecting Japan 2013 to 2022, together with the comparison of the wage anchor against data outside the model. Section 4 discusses robustness, interpretation, and policy implications. Section 5 concludes.

2. Model and Identification

This section presents the New Keynesian (NK) DSGE model that underlies the empirical analysis. The exchange rate and import price block is omitted, and Section 4.1 checks robustness to its absence. The model has five structural equations, five observables, and five structural shocks. What distinguishes it is that a standard NK core carries regime-specific steady-state anchors, $\bar{\pi}$, $\bar{\pi}^w$, and $\bar{r}$. In the baseline specification a regime shift loads on these anchors alone, not on Phillips-curve slopes and not on shock variances. Alternative specifications that let slopes and variances switch are tested explicitly in the nested comparison of Section 3.2. The design separates a shift in the level of nominal variables from a change in the transmission mechanism. Section 2.9 sets out the estimation procedure.

2.1 Overview of the model

There are five endogenous variables, $\{x,\ \pi^p,\ \pi^w,\ \pi^{en},\ i\}$, and five structural shocks, $\{\varepsilon^x,\ \varepsilon^p,\ \varepsilon^w,\ \varepsilon^{\text{en}},\ \varepsilon^i\}$. The complete system is

$$x_t = \mathbb{E}_t x_{t+1} - \tfrac{1}{\sigma}\big(i_t - \mathbb{E}_t\pi^p_{t+1} - \bar{r}(s_t)\big) + \varepsilon^x_t \tag{1}$$

$$\pi^p_t = (1-\beta)\,\bar{\pi}(s_t) + \beta\,\mathbb{E}_t\pi^p_{t+1} + \kappa_p\,x_t + \zeta\,\pi^{en}_t + \varepsilon^p_t \tag{2}$$

$$\pi^w_t = (1-\beta)\,\bar{\pi}^w(s_t) + \beta\,\mathbb{E}_t\pi^w_{t+1} + \kappa_w\,x_t + \varepsilon^w_t \tag{3}$$

$$i_t = \rho_i\,i_{t-1} + (1-\rho_i)[\,\bar{r}(s_t) + \bar{\pi}(s_t) + \phi_\pi(\pi^p_t - \bar{\pi}(s_t)) + \phi_x\,x_t\,] + \varepsilon^i_t \tag{4}$$

$$\pi^{en}_t = \rho_{en}\,\pi^{en}_{t-1} + \varepsilon^{\text{en}}_t \tag{5}$$

Here $\bar{\pi}$, $\bar{\pi}^w$, and $\bar{r}$ are regime-specific anchors. At t_0 (January 2022) each switches once, as a step in level, from its PRE value to its POST value. Written with a step function,

$$\mu(s_t) \equiv \begin{bmatrix} \bar{\pi}(s_t) \\ \bar{\pi}^w(s_t) \\ \bar{r}(s_t) \end{bmatrix} = \mu_{\text{PRE}} + \Delta\mu\,H(t-t_0), \qquad \Delta\mu = \mu_{\text{POST}} - \mu_{\text{PRE}} \tag{6}$$

where $H(\cdot)$ is the unit step function ($H(u) = 0$ for $u < 0$ and $H(u) = 1$ for $u \geq 0$). The anchors switch once as a step in level, not as an impulse, and within each regime they are constant.

The transmission block, comprising the slopes κ, the persistence parameters ρ, and the shock variances, is held fixed across regimes. The intercept $(1-\beta)\bar{\pi}(s_t)$ is a quantity relative to a common reference point rather than an absolute level, and it does not accumulate along the horizon. This keeps the specification consistent with the standard NKPC. The discrete step approximates a continuous transition in beliefs, in which the share of agents who believe the new regime moves from zero to one (Section 2.3). The step uses only the flat ends of that transition. In estimation each equation is written in deviations from the anchors, so that $\mu(s_t)$ enters only through the regime means of the observables, $d(s_t) = [\bar{\pi}; \bar{\pi}^w; \bar{r} + \bar{\pi}; 0]$ (Definition 1 and Proposition 1). The argument (s_t) is suppressed below.

The nominal anchor is the central concept of the paper, and is defined before the blocks are described.

Definition 1 (Nominal anchor)**.** A nominal anchor is the deterministic component that fixes the steady-state level of a nominal variable, whether a price, a wage, or the nominal interest rate. It governs only the level to which that variable converges, and is separate from the transmission mechanism, that is, from Phillips-curve slopes, shock variances, and impulse responses. In this model the anchors are $(\bar{\pi}, \bar{\pi}^w, \bar{r})$.

Remark 1. The three anchors are not a single common target. Each is a separate mean parameter, and the price anchor and the wage anchor can move separately. But they are not unrelated. The steady-state real wage is given by $\bar{\pi}^w - \bar{\pi}$ and the steady-state policy rate by $\bar{r} + \bar{\pi}$. The three are therefore jointly restricted, through the real-wage relation and through the Fisher and policy-rule steady states. A regime shift is represented as a shift in the levels of these three.

The definition has two implications. First, an anchor resides in the nominal level to which resetting converges, not in the frequency of resetting, which is the slope κ. An anchor moves when the level that resetting aims at moves, not when resetting becomes easier. This separates a change in structure from a change in nominal levels. Second, an anchor is a steady-state component shared across several nominal variables. The real wage $\bar{\pi}^w - \bar{\pi}$, and the Fisher and Taylor steady state $i_{ss} = \bar{r} + \bar{\pi}$, tie the three to one another. They are therefore not free parameters that fit the mean of each series independently. Section 3.4 tests this restriction.

Which nominal levels can carry an anchor at all? Within the closed-economy New Keynesian framework adopted here, three nominal levels emerge naturally as candidates for regime-specific anchors. Firms set goods prices, households or unions set nominal wages, and the central bank sets the nominal interest rate. Each of these decisions determines the long-run nominal reference level of a different variable, and for that reason each can, in principle, carry its own anchor. This is the standard nominal block of the New Keynesian literature, in which price and wage rigidity are the two frictions and the policy rule closes the system (Erceg et al., 2000; Christiano et al., 2005; Smets and Wouters, 2007). Opening the economy adds a fourth, the exchange rate, together with import prices. Beyond that, asset prices, imported input prices such as energy, and price indices specific to particular agents, such as producer prices, carry nominal levels of their own.

This paper focuses on the three closed-economy anchors, and the remaining candidates are handled in three different ways. The exchange rate and import prices are not set aside silently.

Section 4.1 adds import price inflation as a sixth observable and estimates the direct pass-through to core CPI freely; the pass-through is economically negligible and the anchor estimates are unchanged. What is tested there is the pass-through channel; import price inflation is not given a regime-specific mean of its own. That channel is a rejected alternative rather than a maintained assumption. Energy enters as an exogenous cost input rather than as an anchor, and Appendix C.5 shows why. Estimating the level of the relative price of energy as a free term does absorb part of the shift in the price anchor. But the implied elasticity does not behave like a structural parameter. The estimate for Japan 2022 is nearly twice that for Japan 2013, 0.0117 against 0.0062, even though the two share a country and an input-output table. Neither lies inside the value that table implies. Demeaning within regimes then restores the baseline estimates. The level term acts as a regime dummy rather than as an anchor. Producer prices enter the model only through the energy cost term, and no separate producer-price anchor is carried. An anchor at the producer stage could differ from the one at the consumer stage, and the two indices did diverge in Japan in 2022; this is a limitation, recorded in Section 5. Asset prices are left outside the model. The object here is the nominal level to which goods-price and wage inflation settle, whereas asset valuations are determined largely by forward-looking arbitrage rather than by staggered nominal contracts. Whether an asset-price anchor is a useful concept, and how it would relate to the three studied here, is a separate question.

The colloquial “norm” corresponds to a level shift in this nominal anchor, and Section 4.2 discusses the interpretation.

2.2 Households and the IS curve

The representative household maximizes intertemporal utility over consumption C_t and labor supply N_t:

$$\max\ \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t [\, U(C_t) - V(N_t)\,], \qquad P_t C_t + B_t = W_t N_t + (1 + i_{t-1}) B_{t-1}, \tag{7}$$

where $\beta = 0.99$ is the monthly discount factor, P_t the consumer price level, W_t the nominal wage, i_t the nominal interest rate, and B_t nominal bonds. Log-linearizing the Euler equation around the steady state gives the IS curve (1). There, x_t is the output gap, that is, the log deviation of consumption from steady state; $\sigma = 1.0$ is the inverse elasticity of intertemporal substitution; $\bar{r}$ is the steady-state real (policy) rate; and ε_t^x is a demand shock. Placing $\bar{r}$ explicitly in both the IS curve and the Taylor rule separates the change in the level of the real rate across regimes from the structure.[4]

[4] The symbol $\bar{r}$ is conventional, but here it denotes the steady-state real policy rate $i_{ss} - \bar{\pi}$ implied by the steady state of the Taylor rule. It is not the natural rate of interest as a welfare concept, that is, the real rate r^* of the flexible-price equilibrium. In the POST period $\bar{r}$ turns negative (Section 3.1). That is an identity: the nominal policy rate was held roughly unchanged while the price anchor $\bar{\pi}$ shifted up, so the real policy rate fell. It is not a claim that r^* itself declined. How far the effective lower bound weakens the identification of $\bar{r}$ can be quantified by comparison with the United States (Appendix C.4).

2.3 Expectational and institutional foundations of the nominal anchor

The nominal anchor of Definition 1 appears in the price and wage Phillips curves, equations (2) and (3), as the intercept $(1-\beta)\bar{\pi}$. That intercept is not a mere constant. It is the long-run component of expected future inflation, which enters firms' price setting multiplicatively. A Calvo reset price discounts future marginal cost by the cumulated product of future inflation. When firms' inflation expectations are anchored to the rational-expectations anchor $\pi^*(s_t)$ of the equilibrium selected in regime s_t, log-linearization separates the log of that multiplicative expectational kernel into the additive intercept $(1-\beta)\bar{\pi}$. The intercept is therefore the log of the expected-inflation anchor in the selected equilibrium, and a regime shift is a discrete switch of the anchor. Appendix A gives the full derivation. Whether a level shift in the anchors also moved the slopes is not ruled out by assumption. It is tested as candidate A, a rise in the slope, in Section 3.2, and is rejected.[5]

Re-anchoring is formulated as the coexistence of agents who believe the new regime, a share $f_t \in [0,1]$, and agents who do not. After an exogenous trigger, f_t moves monotonically from zero to one. The effective anchor in each setting problem is then a mixture of the two ends:

$$\bar{\pi}_t^j = f_t\,\pi_{\mathrm{POST}}^{j*} + (1-f_t)\,\pi_{\mathrm{PRE}}^{j*}, \qquad j \in \{p, w\} \tag{8}$$

As the simplest benchmark, a common f_t is imposed on firms and households, which amounts to coordinated nominal adjustment. The terminal levels π_{POST}^{p*} and π_{POST}^{w*} may nonetheless differ. What matters is that estimation uses only the flat ends of f_t, with $f \approx 0$ in PRE and $f \approx 1$ in POST. The transition window over which f_t moves quickly is excluded, and each regime is estimated separately. The discrete shift is therefore an approximation that uses the flat ends of a continuous transition in beliefs. The anchor identification of Proposition 1 does not depend on the specific form of f_t (Appendix A). Endogenizing the economic mechanism that switches the anchor is not the aim here, and Section 4.2 discusses the interpretation.

Why can the price anchor and the wage anchor move separately? This is not endogenized here, but an institutional contrast and a corresponding modeling distinction motivate it. The institutional contrast is that wage setting is often more synchronized and institutionally coordinated than price setting. Japan's *shuntō* concentrates in almost the same weeks each year, and shares a public focal number: the demand put forward by Rengo, or the previous year's settlement. Major U.S. contracts also run for two to three years, so the dates of revision are discrete. Price setting, by contrast, is staggered and decentralized, and no focal number is shared by all firms. The focal number shared in *shuntō* and in major contracts serves as the nominal benchmark that individual agents consult when they revise wages. As long as that benchmark persists across bargaining rounds, it can serve as the observable counterpart of the steady-state level to which wage setting converges, that is, of the nominal anchor of Definition 1. This provides an institutional interpretation of the estimated wage anchor rather than a theoretical equivalence. The modeling distinction is that synchronization is a different dimension from the frequency

[5] This formulation differs from trend-inflation DSGE models (Cogley and Sbordone, 2008; Ascari and Sbordone, 2014) in where the linearization point is taken. They linearize around a time-varying trend $\bar{\pi}_t$, so the slope becomes a function of $\bar{\pi}_t$. Here the reference point is held common and the anchor is placed in the intercept, so the slope is invariant across regimes by construction.

of revision. Frequency determines how many agents revise per period, and corresponds to the slope κ. Synchronization and the focal number determine what revising agents consult, and correspond to the anchor. The first shows up in propagation, the second in the anchor. In an economy where wage setting depends on a synchronized focal number, a shift in the price anchor may therefore leave the wage anchor unchanged until the next bargaining round updates the focal reference. If the round agrees on a new level, the two move together. Where long contracts carry cost-of-living adjustment clauses, as in the United States in 1979, disinflation may require unwinding the built-in wage escalation, allowing the wage anchor to adjust by more than the price anchor. This view is consistent with Mitchell (1985), who argues that the benchmark for wage setting shifts discontinuously between periods.

But the two anchors are not independent in the long run. In steady state the growth of the real wage, $\bar{\pi}^w - \bar{\pi}$, equals the rate of productivity growth, so the two restrict each other (Remark 1). The question addressed here concerns the medium run, before that long-run link is realized: can the two anchors move separately? The three episodes of Section 3 can be read as three ways in which they do.

2.4 Price and wage Phillips curves and the energy block

Domestic goods firms face Calvo (1983) price rigidity and can reset their price with probability $(1-\theta_p)$ in each period. Log-linearization gives equation (2) for core CPI inflation π_t^p, defined as all items less food and energy. Here $\kappa_p = (1-\theta_p)(1-\beta\theta_p)/\theta_p$ is the price Phillips-curve slope and ε_t^p is a domestic price (markup) shock.

The model takes core CPI directly as the observable π^p. Energy prices enter price setting as an intermediate input. With the production function $Y_t = A_t N_t^{1-\gamma} E_t^{\gamma}$, where E_t is energy input, the contribution of energy to marginal cost adds the energy inflation term $\zeta\, \pi_t^{en}$ to the price Phillips curve. The coefficient is

$$\zeta = \gamma\, \kappa_p, \tag{9}$$

where γ is calibrated from the energy input share in each country's input-output table, 0.0263 for Japan and 0.0339 for the United States. Strictly, this cost term is proportional, through marginal cost, to the level of the relative price of energy (Appendix C.5). Observed energy inflation is used instead because the relative price series is close to a unit root. In a specification that uses the level, the difference in levels across regimes becomes observationally indistinguishable from the anchor. Appendix C.5 examines this point and shows that the estimated anchors are unchanged once identification comes from within-regime variation alone. At the selected κ_p, the coefficient $\zeta = \gamma\kappa_p$ is very small, so the pass-through from energy to underlying inflation is limited. The impulse responses in Section 3 confirm this.

On the wage side, households, or unions, face Calvo rigidity and reset the nominal wage with probability $(1-\theta_w)$, which gives the wage Phillips curve (3). Here κ_w is the wage Phillips-curve slope and ε_t^w is a nominal wage shock.

In this model the nominal wage π_t^w does not feed directly into core CPI π_t^p. The baseline shuts down the direct wage-to-price channel so that shifts in the two reference levels can be recovered separately. Section 3.3 then asks whether the data support relaxing this restriction by adding two-way cross terms; opening them leaves the loop gain economically negligible.

Energy inflation π_t^{en} is constructed from the energy-related components of the Corporate Goods Price Index (CGPI): petroleum and coal products, and electric power, city gas, and water. The U.S. counterpart is built in the same way from the corresponding energy price index (Appendix G). It is modeled as an exogenous AR(1) process, equation (5), in which $\varepsilon_t^{\text{en}}$ is an energy price shock. The persistence ρ_{en} is not estimated. It is calibrated to the OLS AR(1) coefficient of seasonally adjusted π^{en}, computed on each episode's own estimation window: 0.5361 for Japan 2022, 0.5689 for Japan 2013, and 0.5797 for the United States.[6]

2.5 Monetary policy rule and the output gap

The central bank follows a Taylor rule with interest rate smoothing, equation (4). The smoothing parameter is $\rho_i = 0.9244$, the monthly conversion $0.79^{1/3}$ of the quarterly estimate 0.79 in Clarida et al. (2000). The rule coefficients are $\phi_\pi = 1.50$ and $\phi_x = 0.50$ (Taylor, 1993); these are long-run coefficients and do not depend on the observation frequency. The term ε_t^i is a monetary policy shock. What matters is that policy responds to the deviation of inflation from its anchor, $(\pi_t^p - \bar{\pi})$, and that the intercept of the rule is the steady-state nominal rate $\bar{r} + \bar{\pi}$. When the inflation anchor $\bar{\pi}$ of a regime moves, the steady-state level of the policy rate moves with it. The output gap x_t is an observable, and Section 2.9 describes how the proxy series is built.

2.6 Steady state, anchors, and de-anchoring

The three anchors give the steady state, that is, the mean level, of each regime:

$$\pi^p = \bar{\pi}, \quad \pi^w = \bar{\pi}^w, \quad i = \bar{r} + \bar{\pi}, \quad x = \pi^{en} = 0. \tag{10}$$

Here $\bar{\pi}$ is the anchor of underlying inflation, $\bar{\pi}^w$ the anchor of nominal wage growth, and $\bar{r}$ the steady-state real rate. Each is estimated separately for the PRE and the POST regime. In the baseline anchor specification, labeled MC below, the Phillips-curve slopes (κ_p, κ_w) and the shock variances $\sigma_{(\cdot)}$ are held common across regimes, and only the anchors move. Whether the variances also change is examined separately in the nested comparison of Section 3.2. The preferred specification there is MBC, in which anchors and variances both move, while the slopes stay invariant in every specification.

This structure provides the basis for the de-anchoring procedure used in estimation.[7] Kalman smoothing and historical decomposition are applied to the deviations that remain after subtracting from each observed series the regime-specific anchor, that is, the steady-state level corresponding to $\bar{\pi}$, $\bar{\pi}^w$, or $\bar{r}$. This separates the deterministic anchor component from the structural shock component. In a specification that leaves the anchors implicit, the same shift instead appears as over-attribution to shocks. Section 2.8 discusses this as an alternative hypothesis, and Sections 2.9 and 3 give the details.

[6] Varying ρ_{en} around this value changes the anchor shifts by less than 0.01%.

[7] De-anchoring here refers to the estimation step that subtracts the regime-specific anchor level from each observed series. It does not refer to expectations becoming unanchored from a target, which is the usual sense of the term in the literature.

2.7 Observation equations and calibration

The model uses five observables for its five endogenous variables and five structural shocks (Table B.1). The number of observables matches the number of structural shocks, and the shock transmission matrix of the observation equations is constructed to have full rank, so stochastic singularity is avoided.

The structural parameters are fixed by calibration. Eight parameters are estimated by Bayesian methods: five shock variances and three anchors (Table B.2). The slopes (κ_p, κ_w) are not estimated within a given run. They are selected by comparing the log marginal data density over a grid that contains the values used in earlier work, $\kappa_p = 0.03$ from Hazell et al. (2022) and $\kappa_w = 0.02$ from Druant et al. (2009) (Section 2.9). The selected values are common to both regimes, and for Japan 2022 the grid selects $(\kappa_p, \kappa_w) = (0.010, 0.005)$. Invariance of the slopes across regimes is not tested by this selection. It is tested by the three-axis nested comparison, in which switching the slopes does not improve fit. Appendix C confirms that the anchor estimates do not depend on which cell of the grid is selected.

2.8 Three margins on which a regime shift can occur

This section sets out three margins of change as possible axes for the nominal shift since 2022. They are not mutually exclusive, and more than one can move at the same time. Which margin moved is determined by the nested model comparison of Section 3.2.

The first, candidate A, is a change in the slope of the Phillips curve. Under this margin κ_p, and possibly κ_w, rose from the PRE period to the POST period. In the nested comparison it is implemented as a switch in the price slope κ_p. It is equivalent to a fall in the Calvo parameters θ_p and θ_w, that is, to price and wage resetting becoming easier. It is tested by the three-axis nested comparison of Section 3.2 and by calibration invariance over the (κ_p, κ_w) grid (Appendix C).

The second, candidate B, is a change in shock variances. Under this margin σ_{ε^p}, σ_{ε^w}, and the rest widened in the POST period, which means greater dispersion in price and wage setting. It is tested by comparing the posteriors of the shock variances across PRE and POST.

The third, candidate C, is a level shift in the steady-state anchors. Under this margin $\bar{\pi}$, $\bar{\pi}^w$, and $\bar{r}$ shift in the POST period, upward for the first two and downward for $\bar{r}$, so that the mean nominal level moves. It is evaluated separately from any change in slopes or variances, through regime-by-regime estimation of the anchors and through the historical decomposition obtained after de-anchoring.

Candidate C attributes a systematic shift in regime levels to a deterministic anchor rather than to the mean of the shocks. If the anchors are left implicit, a systematic shift in regime levels instead appears as systematic drift in shocks that ought to be i.i.d., and produces over-attribution to particular shocks. In the nested comparison of Section 3, which margin moved differs by episode. In the two Japanese episodes anchors alone dominate variances alone in model fit, while in the United States variances alone dominate anchors alone. In all three episodes, however, the specification that switches both anchors and variances fits best. The evaluation is therefore not a choice among mutually exclusive hypotheses. It is a nested comparison among margins that can coexist.

2.9 Estimation strategy

The same structural core and the same estimation procedure are applied to three episodes. Table 1 reports, for each episode, the PRE period, the POST period, the transition window excluded from estimation, and the episode-specific settings. Common to all three is the framework itself. Three anchors are estimated, for prices, for wages, and for the steady-state real rate. The Phillips-curve slopes (κ_p, κ_w) are selected on a grid. The two regimes are then estimated separately, after excluding the transition window. The episodes need not be identical in every detail, and they are close to identical in practice. The observation specification is common to all three, with five observables including energy, $(\pi^p, \pi^w, \pi^{en}, i, x)$, and the output gap is real consumption in both countries, differing only in source. Two things change across episodes: the sample window, and two country-specific exogenous values. Those two are the energy share γ, from the input-output table, and the energy persistence ρ_{en}, from OLS. Both are constructed directly from each country's own data. Every other calibration, including the reference values for κ, the policy rule, and the discount factor, is common to all three episodes. Appendix C discusses what this invariance implies.

Table 1: Definition and settings of the three episodes

	Japan 2013	Japan 2022	U.S. 1979
PRE period	2000:1–2012:12	2014:1–2019:12	1970:1–1979:9
Transition (excl.)	2013:1–2013:12	2020:1–2021:12	1979:10–1982:12
POST period	2014:1–2019:12	2022:1–2026:4	1983:1–1990:12
Months, PRE / POST	156 / 72	72 / 52	117 / 96
Observables	5 (incl. energy)	5 (incl. energy)	5 (incl. energy)
Output gap	Real cons. (FIES)	Real cons. (FIES)	Real personal cons.
γ (energy share)	0.0263	0.0263	0.0339
ρ_{en} (OLS AR(1))	0.5689	0.5361	0.5797
Break trigger	2% target, QQE	Global supply shock	Volcker tightening

Notes: All three episodes use the common structural core of Section 2 and the common estimation procedure of this section. The set of observables is the same in each. The proxy for the output gap is real consumption in both countries, and only the source (the Family Income and Expenditure Survey, and the national accounts) and the country-specific exogenous values differ. The transition window is excluded on the same ground in each case, namely that the anchors are in transition and $f_t(1 - f_t)$ is non-zero (Section 2.9), and each regime is estimated separately. Only two exogenous values are country-specific, γ and ρ_{en}. The United States uses the same five observables as Japan, including energy and excluding the real wage. The POST period for Japan 2022 runs 52 months, from January 2022 to April 2026, but the nominal wage series ends in March 2026, so that series has 51 observations; the missing value is handled by the Kalman filter.

Estimation uses monthly data from January 2000 onward. The estimation itself is monthly, but the text reports anchor levels and shifts at annual rates, that is, monthly values multiplied by twelve and expressed in percent. The parameter tables in the appendix report both the monthly estimation units and the annual rates. The estimation window for the main episode, Japan 2022, runs from January 2014 to April 2026; the years 2000 to 2012 are used for the Japan 2013 episode.

The sample is split into a PRE period, January 2014 to December 2019, of 72 months, and a POST period, January 2022 to April 2026, of 52 months. Two transition windows are excluded: the year 2013, of 12 months, and the COVID period, January 2020 to December 2021, of 24 months.[8] Section 2.9 discusses the grounds for excluding the two windows. The 72 months of the PRE period exceed the 52 months of the POST period, and monthly data sharpen the identification of the break date (Section 2.9).

The estimation proceeds in two stages. A grid search over the slopes (κ_p, κ_w) selects the best cell by log marginal data density, and full MCMC is then run at that cell. Appendix B gives the construction of the series, the seasonal adjustment, the grid, the priors, and the MCMC settings. Japan 2022 serves as the example throughout; the other episodes follow the same procedure and differ only in the estimation window.

All results are reported at the 90% credible level. This level is standard in the Bayesian DSGE literature, and is used throughout given the trade-off with power in a short monthly POST sample. The Monte Carlo experiments evaluate the finite-sample properties of the same decision rule (Appendix E).

Because the anchors are made explicit as deterministic components (Section 2.6), the historical decomposition is carried out by de-anchoring. The procedure is as follows. From the anchors estimated separately for PRE and POST, the steady-state level of each observed series is constructed, namely $\pi^p = \bar{\pi}$, $\pi^w = \bar{\pi}^w$, and $i = \bar{r} + \bar{\pi}$. That level is subtracted from each series, and Kalman smoothing and shock decomposition are applied to the resulting deviations over the full sample, using the PRE parameters. The difference in contributions between POST and PRE is then attributed to a deterministic ANCHOR component and to the individual structural shock components. In a naive decomposition that does not subtract the anchors, a systematic shift in regime levels appears instead as systematic drift in smoothed shocks that ought to be i.i.d. The result is over-attribution to particular shocks. De-anchoring removes that contamination. It is a special case of the general separation result of Proposition 1, formalized as Corollary 1.

Proposition 1 (Non-confounding of anchors and centered dynamics under maintained assumptions)**.** Write the observed series as $y_t = \mu(s_t) + \tilde{y}_t$, where $\mu(s_t) = (\bar{\pi}, \bar{\pi}^w, \bar{r})$ is the deterministic nominal anchor that depends on the regime s_t (Definition 1), and $\tilde{y}_t$ is a mean-zero stationary component governed by the propagation parameters $\theta = (\kappa_p, \kappa_w, \Sigma, \dots)$, that is, the Phillips-curve slopes and the covariance of the structural shocks. The maintained assumptions are: (a) within each regime phase the state space is linear and time invariant, the transition matrix is stable, and $\tilde{y}_t$ is a Gaussian stationary process; (b) the anchor $\mu(s_t)$ enters only the deterministic level; (c) the propagation parameters θ enter only the centered dynamics $\tilde{y}_t$; (d) the two are additively separable; and (e) the usual regularity conditions hold, with the true values interior and the moment map differentiable. Anchors, slopes, and shock variances may all change across regimes, so structural invariance is not assumed. Then, in each regime s,

(i) the anchors $(\bar{\pi}_s, \bar{\pi}^w_s, \bar{r}_s)$ are identified from the first moments of the observed series: $\bar{\pi}_s$ from $\mathbb{E}\pi^p$, $\bar{\pi}^w_s$ from $\mathbb{E}\pi^w$, and $\bar{r}_s$ from $\mathbb{E}i - \mathbb{E}\pi^p$;

(ii) the propagation parameters θ, that is, the slopes and the variances, enter only the second moments of $\tilde{y}_t$, the autocovariances $\Gamma_s(h)$ for $h \geq 0$, and do not enter the first moments.

[8] The final month of the nominal wage series has not yet been published, so that series has 51 observations in the POST period. The missing value is handled by the Kalman filter.

The asymptotic information matrix is therefore block diagonal between the anchor block and the propagation block (Lemma 3). Anchors and propagation parameters appear in different moment blocks, so in the population a change in one need not be represented as a change in the other. Neither are $\Delta\bar{\pi}, \Delta\bar{\pi}^w, \Delta\bar{r}$ confounded with $\Delta\kappa_p, \Delta\Sigma$, nor the converse. This is a statement about non-confounding between the two blocks. Whether (κ_p, Σ) is uniquely pinned down within the propagation block is a separate question, taken up below and in Appendix A.2.

The proposition establishes separation of regime means from centered dynamics. It is not a claim of strong identification of every structural parameter; what is weakly identified within the propagation block is examined in Appendix C. Appendix A.2 gives the proof. The anchors enter only the first moments of the observables, that is, the regime means, and the propagation parameters only the second moments, the autocovariances. The asymptotic information matrix of the two blocks is block diagonal as a result.

Reading the anchor–variance block of Proposition 1 from the side of misspecification, in the context of the historical decomposition (Section 2.9), gives the following corollary.

Corollary 1 (Observational confounding of anchors and shocks: the need for de-anchoring)**.** Let the observed series follow $y_t = \mu(s_t) + \tilde{y}_t$, where $\mu(s_t)$ is a deterministic level that depends on the regime s_t, the nominal anchor of Definition 1, and $\tilde{y}_t$ is a mean-zero stationary component generated by i.i.d. structural shocks ε_t with $\mathbb{E}[\varepsilon_t] = 0$. Suppose smoothing and decomposition are carried out under a single steady state $\mu_0 \neq \mu(s_t)$ common to both regimes, so that regime-specific deterministic steady states are not allowed. Then the deterministic anchor component $b \equiv \mu(s_t) - \mu_0$ is observationally confounded with cumulated structural shocks. Specifically, the smoothed shocks $\hat{\varepsilon}_t$ acquire a non-zero, serially correlated component proportional to b, with $\hat{\varepsilon}_t \to G^{-1}b$, where G is the long-run response of the observables to innovations, and a deterministic level shift is misattributed to the accumulation of stochastic innovations. Conversely, subtracting the anchor $\mu(s_t)$ before smoothing, that is, de-anchoring, removes the confounding: $\mathbb{E}[\hat{\varepsilon}_t] = 0$ is restored and the anchor is identified as a deterministic component. The conclusion does not depend on the sign of the anchor difference b.

Appendix A.2 gives the proof.

3. Results: Identifying the Anchors Across Three Episodes

This section applies the common structure to three episodes: the upward shifts in Japan in 2013 and 2022, and the downward Volcker disinflation. Proposition 1 is invariant to the sign of the regime-mean shift, so the U.S. episode provides a natural counterpoint to the Japanese cases.[9] Sections 3.1 to 3.3 compare, across all three, how the anchors moved, which margin carried the shift, and whether the wage–price feedback is self-reinforcing. Section 3.4 then links the two Japanese episodes, and Section 3.5 assesses the wage-anchor interpretation using data from outside the model. Episode definitions and calibrations are collected in Table 1 and

[9] In models where trend inflation enters the price-setting slope materially (Cogley and Sbordone, 2008; Ascari and Sbordone, 2014), the separation is approximate, holding in a first-order approximation around each regime's own steady state. Each regime is approximated to first order independently here, so the separation holds within that approximation.

Appendix G. The scope of the framework is the class of linear state-space models in which a regime-dependent deterministic mean $\mu(s_t)$ and centered propagation dynamics θ are additively separable; where both move, de-anchoring still isolates the $\mu(s_t)$ component and any change in θ is tested separately.

The United States is treated as a closed economy and the exchange rate block is dropped, but energy is included among the observables in exactly the same way as for Japan.[10] There are five observables, $(\pi^p, \pi^w, \pi^{en}, i, x)$, and Appendix G (Table G.2) lists the FRED source and the transformation of each series. The real wage is dropped from the observables. Once nominal wages and prices are observed separately, the real wage is a linear combination of the two and carries no independent information; it is dropped in the Japanese estimation for the same reason. The PRE window, January 1970 to September 1979, and the POST window, January 1983 to December 1990, are taken independently, and the transition period that contains the change in operating procedure, October 1979 to December 1982, is excluded. The calibration follows Appendix G (Table G.1).[11]

All observed series are monthly series from FRED (Appendix Table G.2). Price and wage inflation are constructed as 100× the monthly log difference of the level series, and the policy rate is the effective federal funds rate, an annual rate, converted to monthly. Under these transformations the inflation anchors are expressed in monthly percent. The PRE-period mean of π^p, for example, is about 0.53% per month, or about 6.4% at an annual rate.

If the downward disinflation includes a shift in the anchors, the price and wage regime means should fall even once changes in slopes and variances are considered separately. Section 3.1 shows the downward shift in the U.S. anchors, and Section 3.2 evaluates the other two margins.

3.1 Anchor shifts in the three episodes

Table 2 places side by side, for the three episodes, the shift Δ in the price and wage anchors and the posterior of the difference between them. The price side moves significantly in all three episodes, in the sense that the 90% interval for Δ excludes zero. What differs is its relation to the wage side. In Japan in 2013 the wage shift cannot be distinguished from zero. In Japan in 2022 prices and wages move by almost the same amount, and their difference cannot be distinguished from zero. In the United States in 1979 wages move about twice as far as prices, and the difference excludes zero. Even when a change in the policy regime is associated with a shift in the price anchor, the wage anchor need not adjust at the same time or by the same amount. The three episodes are three ways in which this plays out. Each row, however, is a posterior that is complete within its own episode; tests of the differences between rows are reported in Section 3.4.

[10] Energy inflation π^{en} is added as an exogenous AR(1) observable, and passes into the price Phillips curve through the intermediate input cost term $\zeta\pi^{en}$, as in Japan. The Volcker period was a monetary disinflation, however, and the contribution of energy in the historical decomposition is close to zero (Appendix G). The oil shocks of the 1970s are absorbed into the PRE-period price shock ε^p. Since ε^p is a quantity on the propagation and shock side, and is identified separately from the anchor $\mu(s_t)$, it does not contaminate the anchor estimates directly.

[11] The value $\phi_\pi = 1.50$ is fixed to secure determinacy, that is, the existence of a unique rational expectations equilibrium. The policy response coefficient in the pre-Volcker period may have satisfied $\phi_\pi < 1$ (Clarida et al., 2000), in which case the model enters the indeterminacy region. The anchor identification here is identification under a fixed policy rule. It does not compete with the reading that the policy rule itself switched, through a restoration of the Taylor principle; the two are complementary. Estimating ϕ_π endogenously to discriminate between determinacy and indeterminacy is left for future work.

Table 2: Anchor shifts in the three episodes (annual %, 90% equal-tailed intervals)

Episode	$\Delta\bar{\pi}$ (price)	$\Delta\bar{\pi}^w$ (wage)	Difference	Pattern
Japan 2013	$+1.02$ $[+0.78, +1.27]$	$+0.74$ $[-0.01, +1.51]$	-0.29 $[-1.07, +0.51]$	Price only
Japan 2022	$+2.20$ $[+1.80, +2.60]$	$+1.88$ $[+0.83, +2.93]$	-0.32 $[-1.43, +0.79]$	One for one
U.S. 1979	-1.84 $[-2.37, -1.32]$	-3.55 $[-4.24, -2.85]$	-1.71 $[-2.58, -0.83]$	Wage $\approx 2\times$

Notes: The column headed Difference reports $\Delta\bar{\pi}^w - \Delta\bar{\pi}$. Δ is the posterior of POST minus PRE, at an annual rate, that is, the monthly value multiplied by twelve. Brackets give 90% equal-tailed intervals, $[5, 95]$. The price anchor excludes zero in all three episodes. In the final column the interval for the difference includes zero for Japan 2013 and Japan 2022, and excludes zero for the United States in 1979. Monthly levels and $P(\Delta > 0)$ for Japan 2022 appear in Appendix H.3; monthly levels for the United States appear in the text of this section.

The central result of the model is that the steady-state anchors shifted clearly from the PRE period to the POST period. Table 2 in Section 3 collects the anchor estimates for all three episodes and reports the values for Japan 2022. Appendix H.3 gives the monthly levels and intervals in detail.

The price anchor $\bar{\pi}$ rises from $+0.53\%$ to $+2.73\%$ at an annual rate, so $\Delta\bar{\pi} = +2.20\%$, or $+0.183$ per month.[12] The wage anchor $\bar{\pi}^w$ rises from $+0.47\%$ to $+2.34\%$, so $\Delta\bar{\pi}^w = +1.88\%$. The steady-state real rate $\bar{r}$ falls by $\Delta\bar{r} = -1.89\%$. For the price and wage anchors the 90% interval for Δ excludes zero, and $\bar{r}$ falls with $P(\Delta < 0) = 0.984$.

What matters is the relation between the two anchors. The posterior mean of the difference in shifts, $\Delta\bar{\pi}^w - \Delta\bar{\pi}$, is -0.027, with a 90% interval of $[-0.119, +0.066]$ that includes zero. The two anchors therefore shifted upward at the same time, by amounts that cannot be told apart statistically. It is worth stressing that this one-for-one movement is not imposed. The likelihood used here does not include the real wage among the observables (Section 2.9). The parameters $\bar{\pi}$ and $\bar{\pi}^w$ are identified independently, from the means of core CPI and of nominal wages, and no restriction anywhere sets them equal. The simultaneous one-for-one shift is what the data deliver.

Why this matters becomes clear when the same framework is applied to 2013 (Section 3.4). Under the announced price target and quantitative and qualitative easing, the price anchor rose significantly, but no rise in the wage anchor is detected. A rise in the price anchor does not automatically bring a rise in the wage anchor with it. The simultaneous one-for-one shift of 2022 is therefore a fact that calls for explanation. The United States in 1979 (Section 3) shows yet another pattern. The anchors move downward, but the fall in the wage anchor is about twice that in the price anchor, and the 90% interval for the difference excludes zero.

What is claimed here is not a comparison across episodes. Each of the three statements above is complete within its own episode. What can be said about 2013 is that the shift in the wage anchor cannot be distinguished from zero. What can be said about 2022 is that the difference between the two anchor shifts cannot be distinguished from zero. Tests of differences that cross the two episodes are reported, candidly, in Section 3.4.

Table B.3 reports posterior means and 90% intervals for all estimated parameters at the best cell of the POST period.

The anchor estimates for the United States in 1979 (Table 2) are as follows. All three anchors

[12] The values reported in the text are those of the baseline specification MC. The anchor shifts differ across specifications in the third decimal place, but apart from rounding this does not affect the conclusions.

shift significantly. The price anchor $\bar{\pi}$ turns down from +6.43% to +4.59% at an annual rate, so $\Delta = -1.84\%$ with a 90% interval of $[-2.37, -1.32]$. The wage anchor $\bar{\pi}^w$ turns down from +6.75% to +3.20%, so $\Delta = -3.55\%$ with an interval of $[-4.24, -2.85]$. The steady-state real rate $\bar{r}$ rises by $\Delta\bar{r} = +2.93\%$. The 90% interval for Δ excludes zero for all three anchors.

The U.S. episode reveals a different pattern. The fall in the wage anchor is about twice that in the price anchor. The posterior mean of the difference $\Delta\bar{\pi}^w - \Delta\bar{\pi}$ is -0.142, with a 90% interval of $[-0.215, -0.069]$ that excludes zero, and $P(\Delta\bar{\pi}^w < \Delta\bar{\pi}) = 0.999$. In the U.S. disinflation, then, the anchor for nominal wages fell further than the anchor for nominal prices. This estimated pattern differs both from Japan in 2022, where the 90% interval for the difference includes zero, and from Japan in 2013, where no significant shift in the wage anchor is detected. Across the three episodes the price anchor always moves significantly. It is the estimated pattern of the wage anchor that differs from episode to episode, although the differences across episodes are not themselves significant (Section 3.4).

The structure of the signs mirrors Japan in 2022 (Section 3.1). In Japan the price and wage anchors turned upward, by +2.20% and +1.88% at annual rates, and the steady-state real rate turned downward, by -1.89%. In the United States in 1979 all three components move with the opposite sign. Japan moved upward, re-anchoring from a low nominal equilibrium to a high one. The United States moved downward, from a high nominal equilibrium to a low one. The same identification framework captures both, with nothing but a change of sign. The mirror is not perfect, however. In Japan the two anchors moved one for one, whereas in the United States the fall in the wage anchor exceeded that in the price anchor. Episodes differ not only in the level of the norm but in how it moves.

The Volcker episode shows that the same separation between regime means and propagation applies to a downward disinflation. The 90% interval for Δ excludes zero for all three anchors, and the downward shift in the wage anchor is estimated sharply. The anchor shifts reverse sign relative to Japan in 2022, while the nested comparison of Section 3.2 evaluates separately whether the slopes or the variances also changed.

3.2 Which axis moved: a three-axis nested comparison

The previous section showed the shift in the anchors. This section determines formally, by nested model comparison, how that shift stands up against switches in slopes and variances. The PRE and POST periods are pooled into a single sample, January 2014 to April 2026, with the COVID period treated as missing, and a single break is placed at January 2022. Ten nested specifications are then evaluated with a common regime-switching Kalman filter. They differ only in which block of parameters switches at the break. M0 switches nothing; MA the price slope κ_p; MW the wage slope κ_w; MB the variances; and MC the anchors, which is the specification of this paper. The remaining five combine these: MAB, price slope and variances; MAC, anchors and price slope; MBC, anchors and variances; MABC, MBC plus the price slope; and MBCW, MBC plus the wage slope. A regime change is tested explicitly for the price slope and for the wage slope alike. All specifications are evaluated on the same data, so their log marginal data densities can be compared directly. The same comparison is applied unchanged to Japan 2013 (Section 3.4) and to the United States in 1979 (Section 3). Table 4 places the results for the three episodes side by side.

For Japan 2022, specifications that do not switch the anchors are rejected by the data. M0, MA, and MW all fall 17 to 42 below MC. A persistent shift in the mean of nominal variables cannot be explained by switching either the price slope or the wage slope.

Removing or adding one axis at a time from the best specification gives a clear structure (Table 5). Removing the anchor axis lowers the marginal likelihood by 20.8 to 47.0, and removing the variance axis lowers it by 12.1 to 74.8. Both axes are indispensable. Adding the price slope, by contrast, moves it by only -0.54 to -0.07, and adding the wage slope by only -0.11 to -0.03. Allowing a regime change in the slopes of the price or wage Phillips curve is not needed, in either Japanese episode or in the U.S. disinflation. This is not an assumption imposed here; it is the conclusion of three independent comparisons. What is tested, however, is κ_p, κ_w, and the shock variances. The policy rule coefficients, the elasticity of intertemporal substitution, shock persistence, and expectation formation are held fixed across regimes. The finding is therefore not that the entire propagation structure is invariant. It is that, within the model class considered, a regime change in the price and wage slopes is not required.

The treatment of the variance axis differs sharply by episode. In the United States in 1979, MB exceeds MC by +27.81, so the variance axis dominates. That attribution is correct. The PRE period of Japan 2013, from 2000 to 2012, contains the global financial crisis, and the POST period of the United States, from 1983 to 1990, coincides with the Great Moderation. The narrowing of variances in the latter is a fact established independently of this paper (Justiniano and Primiceri, 2008; Sims and Zha, 2006; Cogley and Sargent, 2005). The framework detected it.

The two Japanese episodes differ. MB falls 18.95 below MC in 2022 and 8.72 below MC in 2013. Switching variances alone does not reach switching anchors. The POST period lies in the middle of a global supply shock, so there is no reason for variances to shrink. Even so, the variance axis cannot beat the anchor axis.

In all three episodes the best specification is MBC, anchors plus variances. But MBC, MBCW, and MABC differ by only 0.03 to 0.54, so they cannot be told apart statistically. What the data identify is the anchor axis and the variance axis; the slope axis is identified for neither prices nor wages. The accurate reading is that MBC is selected as the most parsimonious member of the family of models that contain both axes. Table 3 compares the posterior estimates of the shock variances in the PRE and POST periods, both at the best cell $\kappa_p = 0.010$, $\kappa_w = 0.005$.

Table 3: Shock variances, PRE against POST (three episodes, posterior mode)

Episode	Shock	σ_{PRE}	σ_{POST}	POST/PRE
Japan 2022	ε^x	2.019	2.019	1.00
Japan 2022	ε^p	0.071	0.124	1.75
Japan 2022	ε^w	0.265	0.326	1.23
Japan 2022	ε^i	0.083	0.069	0.83
Japan 2022	ε^{en}	1.510	2.028	1.34
Japan 2013	ε^x	1.668	1.668	1.00
Japan 2013	ε^p	0.091	0.067	0.73
Japan 2013	ε^w	0.266	0.262	0.98
Japan 2013	ε^i	0.053	0.083	1.57
Japan 2013	ε^{en}	1.462	1.472	1.01
U.S. 1979	ε^x	1.775	1.775	1.00
U.S. 1979	ε^p	0.257	0.110	0.43
U.S. 1979	ε^w	0.284	0.214	0.76
U.S. 1979	ε^i	0.062	0.043	0.69
U.S. 1979	ε^{en}	1.293	2.844	2.20

Notes: Posterior modes under the best cell of each episode. Ratios above one indicate that variances widened in the POST period.

This is consistent, independently, with the Monte Carlo evidence. The size on the slope axis is 0.030 to 0.055 against a nominal 10%. In episodes where κ_p is small the coverage falls near a zero shift, to 0.715 for Japan 2013 and 0.635 for the United States (Appendix E). The main ranking is reproduced under marginal likelihoods computed by the modified harmonic mean (Appendix D).

The pattern in Table 5 is a sharp asymmetry across margins. Removing either the anchor axis or the variance axis costs between 12 and 75 in log marginal data density, while adding either slope axis changes it by less than 0.55. Two of the three margins are indispensable and the third adds nothing, and the same asymmetry appears in all three episodes.

Three points follow. First, the data do not support a regime change in the slopes, for prices or for wages. This is a finding that has been tested, not the result of assuming invariance. More important still, the treatment of the slopes does not propagate into the anchor estimates. Whether estimation uses MC, MBC, MAC, MABC, or MBCW, the anchor shift for Japan 2022 agrees to the third decimal place, at $(\Delta\bar{\pi}, \Delta\bar{\pi}^w)$ near $(+0.183, +0.156)$. That holds even though the estimated level of the slopes varies across specifications by a factor of 1.0 to 5.3 for κ_p and 0.90 to 0.93 for κ_w. The suspicion that the anchors are absorbing a change in the slopes is thus refuted directly.

Second, anchors and variances both switch. The best specification, MBC, exceeds both anchors alone (MC) and variances alone (MB).

Third, which of the two dominates on its own, in terms of model fit, differs by episode. In the two Japanese episodes anchors alone (MC) exceed variances alone, with MB at -18.95 and -8.72, while in the United States variances alone (MB $= +27.81$) exceed anchors alone. This is a ranking in marginal likelihood, however, and does not mean that variances explain the change in mean levels. In the historical decomposition, 94% to 97% of the mean shift in prices and wages is attributed to the anchor component in the United States as well (Appendix G). In every case MBC, which includes both axes, is best. The framework detects anchors when anchors move and variances when variances move. It is not a device that can see only anchors.

For Japan 2022 the advantage of MC over M0 is +36.46. By the usual Bayes factor standards this is strong evidence, but it is small next to the United States, where the figure is +42.29. Taking the PRE period from 2000 would lengthen the sample and widen the gap further. That longer PRE period, however, straddles the regime change of 2013 and is not a single regime (Section 3.4). The short but correct sample was chosen over the long but wrong one. The price of that choice is lower power, and it should be reported candidly.

Table 4: Three-axis nested comparison, with the price and wage slopes allowed to switch separately: log MDD relative to MC (three episodes)

Specification	Switching block	Japan 2013	Japan 2022	United States 1979
M0	None (single regime)	−17.26	−36.46	−42.29
MA	Price slope κ_p	−17.34	−34.84	−41.78
MW	Wage slope κ_w	−17.39	−36.50	−42.33
MB	Variances	−8.72	−18.95	+27.81
MC	Anchors	0.00	0.00	0.00
MAB	Price slope and variances	−8.69	−16.85	+27.58
MAC	Anchors and price slope	−0.45	−0.53	+0.07
MBC	Anchors and variances	**+12.12**	**+12.80**	**+74.77**
MABC	MBC plus price slope	+11.58	+12.36	+74.70
MBCW	MBC plus wage slope	+12.01	+12.74	+74.74

Notes: Each cell is the log MDD difference relative to MC. Bold marks the best specification in each episode. A regime change in a slope is written as the log ratio $\kappa^{\text{POST}} = \kappa \exp(d)$ (Section 2.9).

Table 5: Change in log MDD when one axis is removed from, or added to, the best specification (MBC)

Operation	Japan 2013	Japan 2022	United States 1979
Remove anchors (→MB)	−20.84	−31.75	−46.96
Remove variances (→MC)	−12.12	−12.80	−74.77
Add price slope (→MABC)	−0.54	−0.44	−0.07
Add wage slope (→MBCW)	−0.11	−0.06	−0.03

Notes: Computed from Table 4. Removing the anchor axis or the variance axis worsens fit substantially, so both are indispensable. Adding either the price or the wage slope axis changes fit only marginally, so neither is supported.

3.3 Wage–price feedback: a test for a spiral

The impulse responses of the previous section showed that, by design, the model does not pass wage shocks through to underlying inflation. The question is whether the data support that independence, or in other words whether the simultaneous shift in prices and wages was a wage–price spiral.

Much of the earlier literature adds a wage term to the price Phillips curve and asks whether its coefficient λ_w is zero. That is a test of pass-through, not a test of a spiral. One-way feedback does not generate self-amplification. If wages push prices up but prices do not push wages back up, the system neither diverges nor amplifies. A spiral is a loop, and the quantity that measures its strength is the product of the two coefficients, the loop gain $\lambda_w \lambda_p$. Contemporaneous cross terms are therefore placed in both the price and the wage equation, both coefficients are estimated jointly, and the posterior of their product is evaluated. The details of the specification, including why lagged cross terms are not used and why determinacy is not settled by the loop gain alone, are described in Appendix F.

The same design is applied to the PRE and POST periods of the three episodes, six regimes in all. Using the same specification, the same priors, and the same test statistic in Japan and the United States guarantees that the finding of no spiral in Japan does not come from a specification chosen for Japan. Table F.1 reports the results.

The estimated λ_w and λ_p include zero in all twelve intervals, six regimes by two coefficients. The loop gain cannot be distinguished from zero in the four Japanese regimes and is significantly negative in the two U.S. regimes. No positive self-amplifying loop exists in any of these episodes. In magnitude the largest absolute value is -0.017, in the U.S. PRE period, and every regime lies below 0.05, two orders of magnitude below the level at which self-amplification would matter. A confusion of terms must be avoided here. A loop gain that excludes zero is not a detected spiral. A spiral means a positive loop gain, that is, self-reinforcement. A negative loop gain is self-damping, the exact opposite. Prices and wages do not amplify one another; they offset one another.

Opening the two-way cross terms improves the marginal likelihood slightly, by -0.00 to $+0.96$ in Japan and by $+2.12$ to $+2.47$ in the United States. That is small as a return on two extra parameters, and not enough to decide model selection. What matters is that the sign of the loop gain does not change even where the improvement occurs. The United States has the two regimes with the largest improvement, and there the loop gain is significantly negative. Letting the data speak through open cross terms produces self-damping, not self-amplification.

Opening the cross terms also leaves the anchor estimates almost unchanged. For the POST period of Japan 2022, blocking the spiral gives $(\bar{\pi}, \bar{\pi}^w, \bar{r}) = (+0.228,\ +0.193,\ -0.195)$, and allowing it gives $(+0.228,\ +0.194,\ -0.195)$. Here $\bar{\pi}$ and $\bar{r}$ are unchanged to the third decimal place, and the difference in $\bar{\pi}^w$ is one unit in the third decimal. The conclusion that attributes the shift in nominal levels to the anchors holds whether the spiral channel is closed or open.

Together with the determinacy result, namely that the region of strong positive feedback is indeterminate under the calibration used here, these three points direct attention to the same conclusion from four directions. The simultaneous shift in prices and wages after 2022 is not the result of one driving the other. The finding is also consistent with earlier work reporting weak pass-through from wages to prices in Japan (Kishaba and Okuda, 2025).

3.4 The Japanese narrative: linking and decomposing 2013 and 2022

This section links the two upward re-anchorings in Japan, in 2013 and in 2022, as a sequence, and provides the propagation and the decomposition.

Three stylized facts about Japanese inflation dynamics after 2022 are presented first. The series used below are core CPI inflation π^p, nominal wage growth π^w, real wage growth $\pi^{w,\text{real}}$, energy price inflation π^{en}, and the output gap x. The sample is split into a PRE period, January 2014 to December 2019, of 72 months, and a POST period, January 2022 to April 2026, of 52 months, with the COVID period of 2020 and 2021 excluded. The PRE period starts in January

2014 because the nominal anchor itself moved around the point at which the Bank of Japan announced a 2% price target and began quantitative and qualitative easing in 2013. The years 2000 to 2012 and 2014 to 2019 cannot be treated as one regime (Section 3.4). The year 2013 is excluded as a transition window.

The first fact is that the mean level of nominal variables shifted up in the POST period while real variables were unchanged. Table 6 reports PRE and POST means of month-on-month rates, in percent, together with tests of the difference.

Table 6: PRE and POST means of nominal and real variables (month-on-month, %)

Variable	PRE mean	POST mean	Difference	Annual (%)	t	Verdict
π^p (core CPI)	+0.05	+0.23	+0.19	+2.2	9.3	***
π^w (nominal wage)	+0.04	+0.20	+0.16	+1.9	2.9	**
$\pi^{w,\text{real}}$ (real wage)	−0.01	−0.07	−0.06	−0.8	−0.9	n.s.
x (output gap)	+0.18	+0.20	+0.02	—	+0.1	n.s.

Notes: Means of month-on-month rates over the estimation windows, PRE 2014:1–2019:12 and POST 2022:1–2026:4. Because the months of consumption tax rate changes are treated as missing in the price series, the effective PRE sample for those series is $n = 70$. For nominal wages $n_{\text{POST}} = 51$, and for the remaining series $n_{\text{PRE}} = 72$ and $n_{\text{POST}} = 52$. Here *** denotes $p < 0.001$ and ** denotes $p < 0.01$; n.s. denotes not significant ($p = 0.37$ for $\pi^{w,\text{real}}$ and $p = 0.95$ for x). The difference for π^p is +0.19 per month, or +2.2% at an annual rate, which matches the total of the historical decomposition in Table 7, +2.23% at an annual rate.

Core CPI shows a significant upward shift of +2.2% at an annual rate and nominal wages one of +1.9%, with $p < 0.001$ and $p < 0.01$ respectively. Real wages and the output gap, by contrast, show no significant change. The shift is confined to nominal variables. These descriptive shifts also match closely the change in anchors delivered by the structural estimation below, where the price anchor $\bar{\pi}$ moves from +0.04 to +0.23 (Section 3). Energy prices π^{en} show a high monthly mean in the POST period, about +0.43, but the change is volatile and not significant ($p \approx 0.64$). Energy is a temporary cost factor, not the source of the shift in nominal levels.

The second fact is that the acceleration in nominal wages was not accompanied by an improvement in real wages. As the first fact shows, nominal wages π^w rose by about +1.9% at an annual rate, while real wages $\pi^{w,\text{real}}$ fell slightly, by −0.8% at an annual rate, and not significantly. Fukuda and Soma (2025) notes that real wages on a total cash earnings basis were negative from April 2022 through May 2024, and the IMF reports a cumulative fall in real wages of about 6% since 2022 (International Monetary Fund, 2026). Most of the nominal wage increase was compensation for inflation, and did not raise real purchasing power.

Third, the timing of the upward shift is examined endogenously. A sup-Wald (Quandt–Andrews) test is applied to the X-13 seasonally adjusted series. For prices the estimated break date is February 2022 whichever treatment of the standard errors is used ($\sup F = 109.8$ under homoskedasticity and 34.6 under HAC, Newey–West with twelve lags). For wages it is January 2022 under homoskedasticity (8.9) but August 2019 under HAC standard errors, so the wage break date is not sharply identified from a univariate test. The claims of this paper about wage setting rest not on that test but on the structural anchor and its comparison with *shuntō* data (Section 3.5).

It is also worth asking what happens when the regime is estimated as a latent state rather than imposed. Selecting the number of states of a Markov-switching model by BIC gives three states for core CPI. The third is not a regime, however. Its mean is -1.20% per month, an extreme negative value, and the mass of its posterior probability concentrates inside the excluded transition window, amounting to roughly one month. It is an outlier state that absorbs the temporary collapse in prices during COVID. The remaining two states reproduce the split imposed here, with posterior probability 1.00 in the PRE period and 0.89 in the POST period. The same procedure applied to the other two episodes leads to the same conclusion. For Japan 2013 two states are selected (PRE 0.88, POST 0.99), and the transition window splits 0.55 to 0.45, which is what a window in transition should look like. For the United States three states are selected, but the mass of the third (mean $+1.03$ per month) concentrates in the transition window (0.51), so the data themselves select the three-phase design of PRE, transition, and POST used here. Making the break date latent therefore delivers the same two regimes and the same boundary. Where an additional state appears, it corresponds not to a third regime but to the transition window or to the COVID outlier that this paper excludes.

The wage side has a further, independent corroboration. How sharply the univariate sup-Wald test detects a wage break follows the same ordering across the three episodes as the movement of the structural wage anchor. In Japan in 2013, where the shift in the wage anchor cannot be distinguished from zero, no wage break is detected and no break date is identified ($\sup F = 3.7$). In Japan in 2022, where the wage anchor moves significantly, a break is detected in January 2022 (8.9, though unstable under HAC). In the United States in 1979, where the wage anchor moves about twice as far as the price anchor, the break is detected with the largest statistic in any of these tests ($\sup F = 76.8$ under homoskedasticity and 110.5 under HAC, both in February 1983). The endogenous intercept of the structural model and a univariate reduced-form test that uses no structure at all deliver the same ordering, independently. This is a third line of support, after the comparison with *shuntō* and collective bargaining data in Section 3.5, and it speaks directly to the concern that the wage anchor might be an artifact of fitting.

These three facts lead to a working hypothesis. Japanese inflation after 2022 is an upward shift in the level, that is, the regime, of nominal variables, unaccompanied by structural change in the real economy. The sections that follow test this hypothesis quantitatively within a structural DSGE. Specifically, it is formulated as invariance of the slopes together with a level shift in the anchors, and whether variances change at the same time is also examined.

This section applies the same structural core, the same calibration, and the same three-axis nested comparison to another date in Japan; the two Japanese episodes also share the same observation specification. In January 2013 the Bank of Japan introduced a 2% "price stability target" in a joint statement with the government, and in April it began quantitative and qualitative easing on an unprecedented scale. If the nominal anchor ever moved, it should have moved here.

The price anchor rises from -0.48% to $+0.54\%$ at an annual rate, so $\Delta\bar{\pi} = +1.02\%$ with a 90% interval of $[+0.78, +1.27]$. The interval excludes zero. Under an announced price target and quantitative and qualitative easing, the anchor for prices did move upward.

The wage anchor is different. It moves from -0.31% to $+0.42\%$ at an annual rate, so $\Delta\bar{\pi}^w = +0.74\%$, but the 90% interval, $[-0.01, +1.51]$, includes zero. The data do not permit the statement that the anchor for wages moved.

Even in an episode in which the central bank announced a price target, purchased assets on an unprecedented scale, and did move the price anchor upward, no rise in the wage anchor was detected. In the central-bank-led episode of 2013, the price anchor was the only anchor that moved clearly. That prices and wages moved together in 2022, by almost equal amounts (Section 3.1), is therefore not an automatic consequence but a fact that calls for explanation. Section 4.2 reads it as coordinated nominal adjustment.

The estimation window, the nested comparison, and the cross-episode tests for this episode are reported in Appendix H.1. Two points matter for the argument here. The claims of this paper are complete within each episode: the difference between 2013 and 2022 cannot itself be tested, because both POST windows are short and the wage side lacks power. And the framework contains no mechanism that generates the one-for-one shift of 2022; coordination through *shuntō*, a common signal, and mechanical indexation are observationally equivalent in aggregate data.

The transmission mechanism is examined through impulse response functions (IRFs) to unit structural shocks. The IRFs are in deviation form and do not depend on the anchors, so they are common to the PRE and POST periods. Figure 1 shows the IRFs for every shock and every variable.

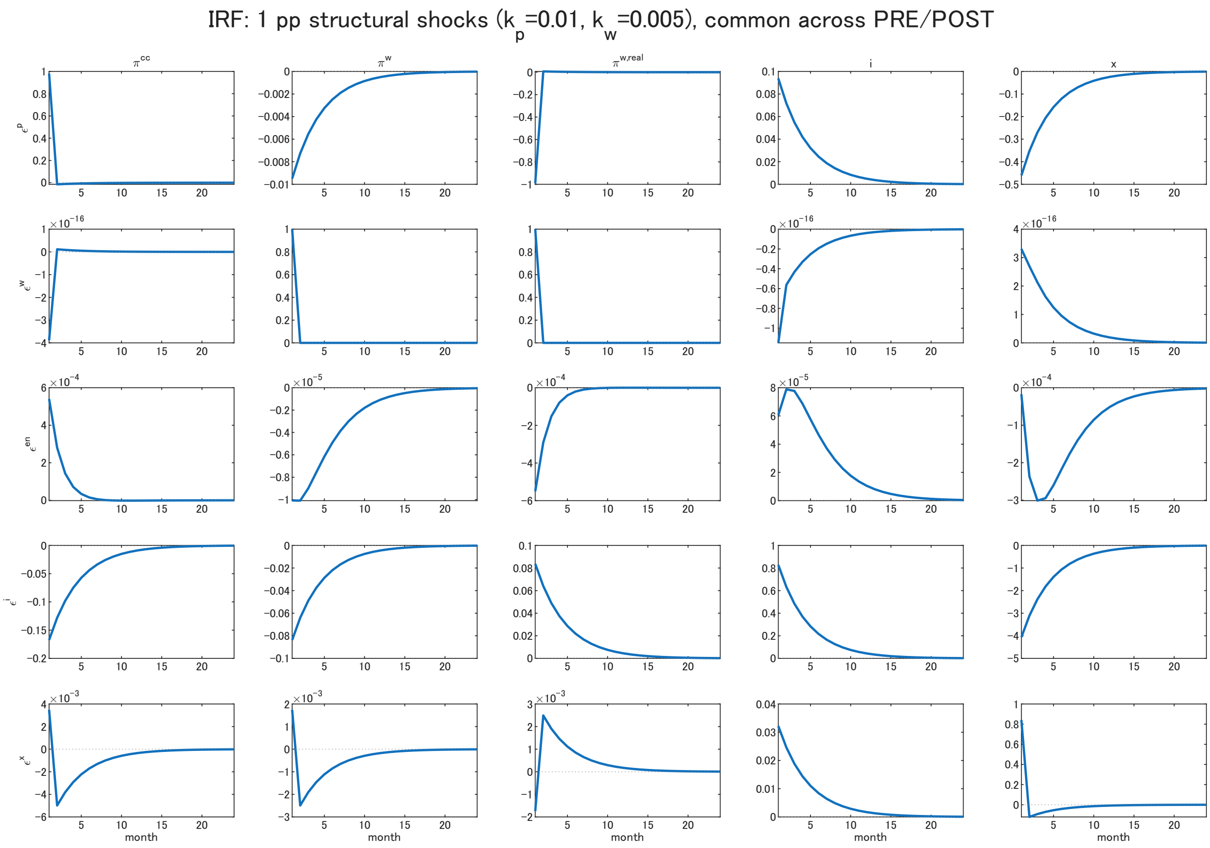

Figure 1: Impulse responses to 1% structural shocks (best cell $\kappa_p = 0.010, \kappa_w = 0.005$; common to PRE and POST)

Three things can be read from the IRFs, with the first-month response given in parentheses.

First, the wage shock ε^w does not pass through to core CPI at all; the response of π^p is exactly zero, and only nominal and real wages respond. This holds by construction: the baseline model shuts down that channel (Figure 1). Section 3.3 asks whether the data support relaxing the restriction.

Second, the domestic price shock ε^p pushes real wages down sharply. It raises π^p by $+0.98$ while nominal wages barely move (-0.01), so real wages fall by -0.99 (Figure 1). This asymmetry, in which a rise in prices lowers real purchasing power without an accompanying wage adjustment, matches the observation that the POST period brought nominal wage increases but little real improvement.

Third, the pass-through from energy to underlying inflation is negligible. The first-month response of π^p to the energy shock ε^{en} is only $+0.0005$, an order of magnitude of 10^{-4} (Figure 1). This reflects how small the intermediate input coefficient $\zeta = \gamma\kappa_p$ is.

The price block and the wage block therefore move almost independently, and energy is not a main driver of underlying inflation.

The de-anchoring procedure of Section 2.9 decomposes the POST minus PRE difference in contributions into a deterministic ANCHOR component and the individual structural shock components. Table 7 reports the monthly difference in contributions for the main variables.

Table 7: Historical decomposition: POST minus PRE difference in contributions (annualized %, Japan 2022)

Component	π^p	π^w
ANCHOR	$+2.20$	$+1.88$
ε^x	0.00	0.00
ε^p	0.00	-0.01
ε^w	0.00	$+0.04$
ε^i	$+0.04$	$+0.02$
ε^{en}	0.00	0.00
init	-0.01	0.00
Residual	0.00	0.00
Total	$+2.23$	$+1.93$
ANCHOR share	98.5%	97.2%

Notes: POST minus PRE difference in the contribution of each component, annualized as the monthly value times twelve. ANCHOR is the contribution of the estimated regime-conditional mean $\mu(s_t)$; for Japan 2022 the steady-state level of prices moves from +0.53% to +2.73% at an annual rate. The remaining rows are the cumulated contribution of each structural shock and the initial condition (init). Because the model is stationary, init has decayed to effectively zero by the POST period, and the residual is effectively zero as well, so the decomposition closes cleanly. This decomposition attributes the mean shift under de-anchoring accounting. The choice among axes, anchors against cost push against variances, rests on model comparison (Table 4) rather than on decomposition shares.

In the steady-state level of the anchor, the underlying trend of core CPI rises from +0.53% to +2.73% at an annual rate, a shift of +2.20%. Under de-anchoring accounting, 98.5% of that shift is attributed to the ANCHOR component. The difference in the contribution of the domestic price shock ε^p is close to zero.

The rise in the underlying trend of nominal wages, a shift of +1.88% at an annual rate, is likewise attributed to the ANCHOR component in 97.2% of its magnitude. The contribution of the wage shock ε^w is only 2.3%, and the residual is close to zero. Under de-anchored accounting, then, almost all of the rise in both prices and wages is described as a level shift in the anchors.

The size of these shares should not itself be read as evidence against cost push. Attributing the mean shift to ANCHOR is a definitional consequence of de-anchoring (Corollary 1). The rejection of the competing hypothesis, cost push on the variance axis, rests on three other things. It rests on the three-axis nested comparison (Section 3.2 and Table 4), where for Japan 2022 the variances-only specification MB falls 18.95 below the anchors-only specification MC. It rests on the propagation IRFs, where ε^p pushes real wages down but does not carry the shift in the underlying level. And it rests on the fact that the price and wage anchors moved one for one, with $\Delta(\bar{\pi}^w - \bar{\pi}) = -0.32\%$ per year and a 90% interval that includes zero.

Because the price and wage anchors moved one for one, the shift in the underlying real wage implied by the anchors is small. It is -0.32% at an annual rate, with a 90% interval of $[-1.43, +0.79]$ that includes zero. This is consistent with the realized change in real wages being statistically indistinguishable from zero (Table 6). Figure 2 shows the historical decomposition of π^p and π^w, in which the ANCHOR component can be seen rising in the POST period from January 2022.

It would be wrong to read the de-anchoring result as merely subtracting means and then recovering those same means. The three anchors $(\bar{\pi}, \bar{\pi}^w, \bar{r})$ are not free parameters that fit the mean of each series separately. They are structural quantities that must reconcile four regime means at once: (i) the mean of core CPI, (ii) the mean of nominal wages, (iii) the level of the policy rate through the Fisher and Taylor steady state $i = \bar{r} + \bar{\pi}$, and (iv) the mean of the real wage. The steady-state real wage in particular is not free. It is a predicted quantity, $\bar{\pi}^w - \bar{\pi}$, and one additional restriction is imposed there. The POST minus PRE shift in the real wage implied by the anchors is $\Delta(\bar{\pi}^w - \bar{\pi}) = -0.32\%$ per year, obtained from the ANCHOR row of Table 7 as $+1.88 - 2.20$. Like the realized value in the descriptive statistics (Table 6), it is statistically indistinguishable from zero. Three parameters fit four means, and the mean of the real wage, which was not fitted separately, is reproduced approximately as a result. The fit of the deterministic component is therefore not an identity. Real wages are constructed mechanically from nominal wages and prices, however, so this is a consistency check on an implication left out of the likelihood, not external validation against independent data. External validation is the comparison with wage-setting data in Section 3.5. In addition, the shift in the structural anchor $\bar{\pi}$, $+2.20\%$ at an annual rate, matches the mean shift in the descriptive statistics, $+2.2\%$ at an annual rate (Table 6). This too is not because the two are equal by definition. It verifies that they remain compatible within a system that includes the Fisher restriction on the policy rate.

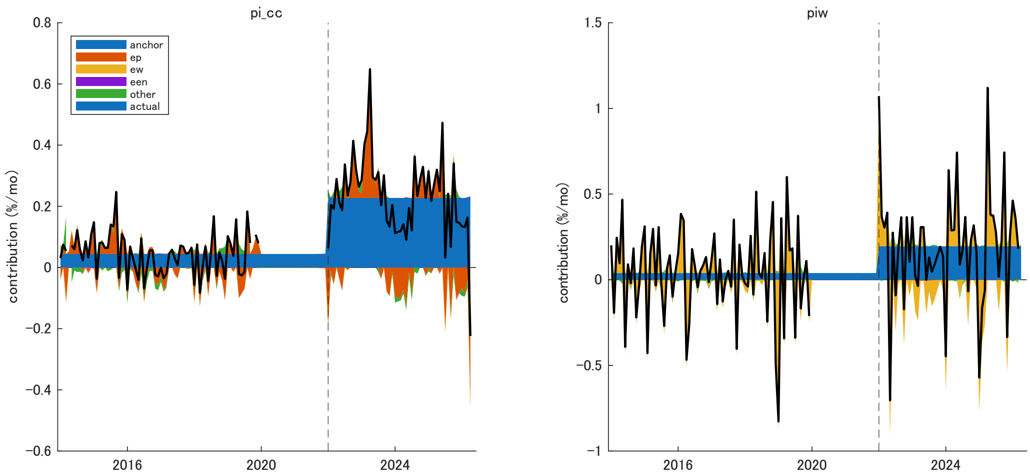

Figure 2: Japan 2022: historical decomposition of π^p (left) and π^w (right)

3.5 External validation: the downward mirror

Up to this point the wage anchor $\bar{\pi}^w$ has been estimated as an endogenous intercept in a state-space model. Whether that value corresponds to a benchmark that actually governs wage setting, or is merely an artifact of fitting the model, cannot be judged from inside the model. This section compares $\bar{\pi}^w$ with wage-setting data that come from outside the model. For Japan these are the base pay increases negotiated in the Rengo *shuntō* rounds, and for the United States the average wage increase over the contract term in the BLS Major Collective Bargaining Settlements. Neither is used anywhere in the structural estimation.

The wage series observed in this paper is the realized nominal wage. For nominal wages per worker in Japan, scheduled pay increases are largely offset at the macro level by composition effects, so the counterpart of $\bar{\pi}^w$ is the base pay increase, not the total wage increase including scheduled increments.[13] In the United States, the closest counterpart to realized nominal wages is the over-life average for newly negotiated contracts, which is conceptually parallel to the Japanese base pay increase in being a rate that is about to take effect.

Figure 3 places the wage-setting data of the two countries side by side with the posterior of $\bar{\pi}^w$. The correspondence holds at three levels. First, the levels match in the PRE period. The average *shuntō* base pay increase in Japan's PRE period, 2015 to 2019, is 0.54%, close to the posterior mean of $\bar{\pi}^w$ at +0.47%, with a 90% interval of $[-0.22, +1.15]$. For the U.S. PRE period, 1970 to 1979, the over-life average is 6.82% against a posterior mean of +6.75%, with an interval of $[+6.28, +7.23]$. Agreement in direction alone would amount to matching a sign; agreement in level is harder to obtain by chance. This supports reading the intercept of the state space as a benchmark for nominal wage setting. Second, the shifts run in opposite directions. Japan moves upward, from 0.54 to 2.70%, and the United States downward, from 6.82 to 2.60%. Under one and the same framework, this is a downward mirror in which only the sign is reversed. Third, the sizes of the shifts are consistent as well.

[13] Government energy subsidies (on fuel from January 2022, on electricity and gas from January 2023) held down measured energy and headline inflation, and through the CPI they raised measured real wages. Whether bargaining internalized this relief in setting base pay, so that the subsidies bore on $\bar{\pi}^w$ itself, is not separately identified: $\bar{\pi}^w$ is estimated as the persistent component of base-pay setting.

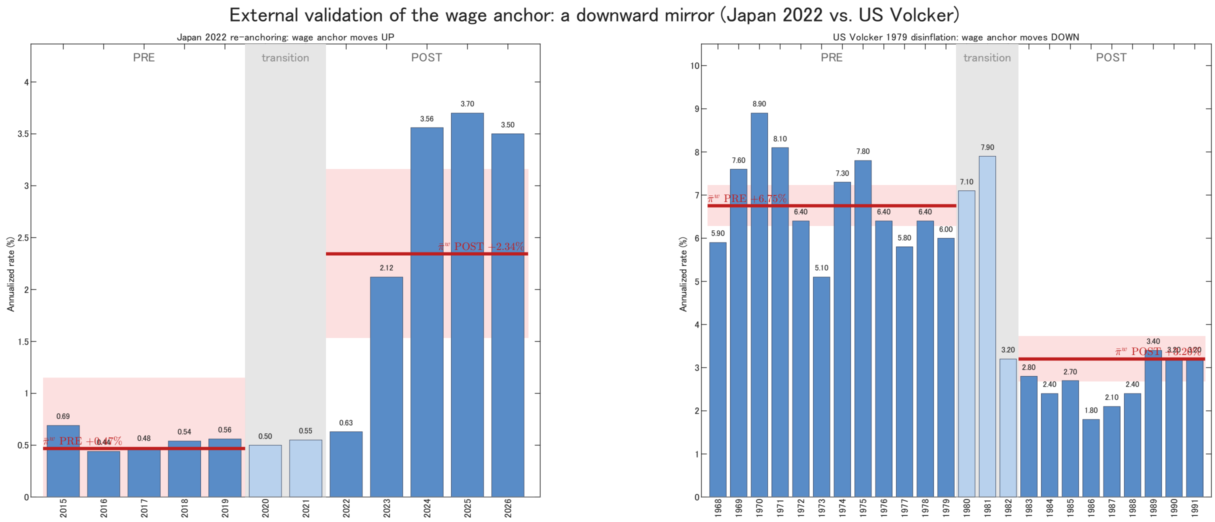

Figure 3: External validation of the wage anchor $\bar{\pi}^w$: the downward mirror (Japan 2022 against the United States 1979). Bars show wage-setting data: for Japan, the Rengo *shuntō* base pay increase (wage increase component, all firm sizes, weighted average); for the United States, the over-life average wage increase in BLS Major Collective Bargaining Settlements (private sector, 1,000 workers or more). The solid red line is the posterior mean of the DSGE wage anchor $\bar{\pi}^w$, and the shaded band is the 90% interval. Grey marks the transition window, excluded from estimation. The wage anchor shifts upward in Japan and downward in the United States in 1979, with only the sign reversed under one identification framework.

Read closely, the figure also shows two gaps. Both are consistent with the interpretation offered here. First, the *shuntō* base pay increase in Japan's POST period, around 3.5%, runs somewhat above the posterior band of $\bar{\pi}^w$, whose mean is +2.34% with an upper bound of +3.16%. This is a coverage premium: *shuntō* covers unionized workers at large firms, while the Monthly Labour Survey covers the whole economy. The gap is observed directly, with scheduled wages in the Monthly Labour Survey rising 1.7% in 2023 against a realized *shuntō* figure of 2.1%. Second, the U.S. over-life average in the POST period, 2.60%, runs somewhat below the posterior mean of +3.20%. The reason is that newly negotiated settlements are a rate that is about to take effect. Effective wage adjustments, which include deferred increases and cost-of-living escalation, averaged 3.17% over 1983 to 1988 (Davis, 1989), and that figure is close to $\bar{\pi}^w$. The decline in COLA coverage, from 61% in 1977 to 40% in 1988, supports the mechanism by which realized wages retained inertia for some time after settlement rates fell sharply.

Mitchell (1985) argued that in the wage equation $w = a + bU^{-1} + cp$ the intercept a shifted down in the early 1980s while the slopes (b, c) remained stable. That is the single-sector, regression-based precursor of the claim made here, that anchors move while slopes stay invariant. The external validation of this paper shows that the intercept shift corresponds, in estimation as well, to actual wage-setting data in both countries.

4. Robustness, Interpretation, and Policy Implications

4.1 Robustness

The central conclusions of this paper, that the shift in nominal levels is attributable to the anchors (Section 3.2) and that no self-reinforcing loop is operating (Section 3.3), were obtained for Japan in 2022. This section confirms that they survive once the exchange rate and import price channel, omitted from the baseline, is added to the observables (Section 4.1).

The baseline model treats energy as an exogenous cost but does not include exchange rates or import prices in the structure. That is a design choice, made to give priority to estimation stability and to the identification of the anchors. Yen depreciation and rising import costs have nonetheless been discussed as one of the main drivers of Japanese inflation after 2022. It is therefore necessary to check whether the result that almost all of the rise in core CPI is attributable to the anchors is an artifact of excluding the exchange rate and import price channel from the structure. This section shows that the dominance of the anchors survives when that channel is opened explicitly. The design of the check mirrors the test for a wage–price spiral in Section 3.3.

Import price inflation π_t^{im} is added as an observable, giving six observables, the five-observable core plus π^{im}, and a direct pass-through term $\chi\,\pi_t^{im}$ is added to the price Phillips curve:

$$\pi_t^p = (1-\beta)\bar{\pi} + \beta\, E_t\pi_{t+1}^p + \kappa_p x_t + \zeta\,\pi_t^{en} + \chi\,\pi_t^{im} + \varepsilon_t^p, \tag{11}$$

$$\pi_t^{im} = \rho_{im}\,\pi_{t-1}^{im} + \varepsilon_t^{im}. \tag{12}$$

Import prices are exogenous AR(1) and no separate exchange rate observable is used; Appendix C.5 gives the calibration of ρ_{im} and $\sigma_{\varepsilon^{im}}$ and the prior on χ.

First, the direct pass-through from import prices to core CPI is economically negligible. Estimating χ freely gives a posterior mean of -0.007 and a 90% interval of $[-0.012, -0.002]$. The interval just excludes zero, but the sign is negative and the magnitude minute: a one-point rise in import prices contributes less than 0.01 points directly to core CPI. Freeing χ does not improve fit either, with a log MDD loss of -1.0 and $\chi = 0$ best, and raising χ exogenously worsens fit monotonically (Table 8). The data actively reject a large pass-through.

Second, the anchor shift barely moves whatever value χ is fixed at. As Table 8 shows, moving χ from 0 to 0.40 leaves the annual shift in the price anchor $\Delta\bar{\pi}$ within [+1.91, +2.33]%/yr, flat around the five-observable canonical baseline of +2.20%/yr. The POST levels of the wage anchor and of the steady-state real rate also move only slightly. In addition, the POST anchor level from the $\chi = 0$ estimation with the import observable added, $\bar{\pi} = +0.233$, is close to the five-observable canonical value of +0.227, which confirms that adding the import block does not distort the estimates.

Table 8: Robustness of the anchors to import price pass-through χ (Japan 2022; six observables, the five-observable core plus π^{im})

χ	log MDD	$\Delta\bar{\pi}$ (%/yr)	$\bar{\pi}^{w}_{\mathrm{POST}}$ monthly (annual)	$\bar{r}_{\mathrm{POST}}$ monthly (annual)
0.00	−248.6	+2.33	+0.193 (+2.32)	−0.203 (−2.44)
0.05	−298.6	+1.99	+0.194 (+2.33)	−0.183 (−2.20)
0.10	−329.7	+1.91	+0.195 (+2.34)	−0.179 (−2.15)
0.20	−366.0	+2.03	+0.195 (+2.34)	−0.189 (−2.27)
0.40	−408.0	+2.17	+0.197 (+2.37)	−0.210 (−2.52)
Free $\hat{\chi}$	−249.6	—	—	—

Notes: Estimates with the import price observable added, six observables in total. Fixing χ over 0 to 0.40 leaves $\Delta\bar{\pi}$ within $[+1.91, +2.33]$%/yr, flat around the five-observable canonical baseline of +2.20%/yr. Freeing χ does not improve fit.

Adding the exchange rate and import price channel to the observables and allowing direct pass-through from import prices to core CPI therefore leaves the conclusion intact. The pass-through coefficient is economically negligible, and freeing it does not improve fit, so the level shift in core CPI is still attributed to the anchors. But $\chi \approx 0$ shows that import prices do not pass through directly to underlying inflation. It does not rule out the possibility that rising import prices were the trigger for the switch in anchors, the signal around which coordination formed. Nor, conversely, does $\chi \approx 0$ actively establish that reading. The accurate statement is that the two are consistent with each other. The first is a question about propagation, treated in this section; the second is a question about what triggers equilibrium selection, treated in Section 4.2 as candidate 2. This is also consistent with core CPI being, by definition, the index least exposed to direct import costs, since it excludes fresh food and energy, and with the understanding that import cost push operates on headline, energy, and food. Like the wage–price spiral check of Section 3.3, however, this test is confined to linear, time-invariant, contemporaneous pass-through. Pass-through with lags, or state-dependent pass-through that operates only during large depreciations, lies outside the framework used here. That is a limitation of the design, and Section 5 collects it with the others.

4.2 Interpretation

The central finding of this paper is that the inflation regime shift in Japan after 2022 is characterized as a level shift in steady-state anchors rather than as a change in structure. The view that the structure changed, through a rise in the Phillips-curve slope κ or a fall in the Calvo parameter θ, is not consistent with these results. The slopes are invariant, and the switch in regime levels is best explained by a specification that allows both anchors and variances to move (Section 3). Compared one against the other, anchors alone beat variances alone in the two Japanese episodes, and the reverse holds in the United States. What changed is the level of the steady-state anchors for prices, wages, and the real rate. Prices and wages, moreover, shifted upward at the same time and by amounts that cannot be distinguished statistically (Section 3.1).

This paper does not define the colloquial concept of a "norm." It approximates the concept operationally, by the nominal anchor of Definition 1, that is, by the steady-state level to which price setting and wage setting converge. The "different norms" of the title mean that the level of this anchor moved between regimes. They do not mean that the transmission mechanism changed. This usage differs from the classical wage norm. The norms of Hicks (1955) and Tobin (1972) are social conventions that discipline the structure of relative wages, whereas what is measured here is the anchor for the nominal level of prices and of wages separately. Under this operational definition, the level shift documented here connects naturally to the "norm" long discussed in Japan. One view holds that not raising prices was the default posture of firms (Watanabe and Watanabe, 2018; Watanabe, 2025). More recently it has been argued that this social convention is dissolving against a background of labor shortage (Fukunaga et al., 2024). Policymakers, too, have described developments since 2022 as a historic shift away from the wage and price norm, and as a move from one equilibrium to another (Takata, 2026).

What is measured here is a level shift in regime means that requires no switch in the price or wage slopes. It is the same kind of object as the intercept shift in a Mitchell-style wage equation (Section 3.5).

Organizing the results along two axes, nominal against real and the price market against the wage market, gives the 2×2 matrix of Table 9. Every entry in the nominal column shifts significantly, and every entry in the real column does not.

Table 9: Asymmetry between nominal and real, and between the price and wage markets

	Nominal	Real
Price market	π^p : shift***	x : unchanged (n.s.)
Wage market	π^w : shift**	$\pi^{w,\text{real}}$: unchanged (n.s.)

The internal structure of this asymmetry matches the anchor estimates exactly. The price anchor and the wage anchor shifted upward by almost equal amounts, +2.20% and +1.88% at annual rates, with a 90% interval for the difference that includes zero. If the two move equally, the real wage anchor $\bar{\pi}^w - \bar{\pi}$ barely moves by definition. That real wages stayed flat, the second fact of Section 3.4, is precisely this implication. What matters here is that real wages do not enter the likelihood (Section 2.9). The two anchors are identified independently, from the means of core CPI and of nominal wages. Even so, the steady-state real wage implied by their difference reproduces a realized value that was never fitted. This is not an identity. It is a consistency check against an implication left out of the likelihood (Section 3.4). The regime shift is thus a purely nominal switch, in which nominal variables adjusted one for one and real variables did not move. Nominal wages accelerated, but real purchasing power did not improve.

What the model identifies directly is a level shift in the expectational anchor under an invariant structure, which is candidate 1. The remaining question is why that anchor switched discretely. The first interpretation offered here is a change in equilibrium selection, candidate 2. It is consistent with the structural formulation of Section 2.3 and with the exchange rate and import price results of Section 4.1, and it stays within rational expectations.

Price and wage setting are strategic complements: each firm, and each worker, moves only when it expects others to revise upward at the same time. This generates multiple equilibria, a low-nominal one in which nobody moves and a high-nominal one in which everyone moves together. Japan was anchored on the former for a quarter of a century. The rise in import prices after COVID acted as a cost shock common to all firms, and may have supplied the common knowledge that moving now means others will move too. On this view, the common shock made it observable that prices and wages could be revised upward simultaneously, and the selected equilibrium moved from the low-nominal one to the high-nominal one. The estimated signature is consistent with that reading: the anchors for prices and wages move at the same time and by almost the same amount, with no switch in the transmission mechanism required.

This remains one interpretation. The model does not solve multiple equilibria explicitly, the timing of the anchor shift is exogenous, and neither the trigger nor the mechanism of equilibrium selection is identified. What can be said is only that the interpretation is consistent with the results of Section 4.1, where the direct pass-through χ from import prices to underlying inflation was economically negligible. Under the equilibrium selection view, import prices worked not as a transmission channel that pushed underlying inflation up, but as a signal that made it possible for firms and workers to move at the same time. The two are consistent, but $\chi \approx 0$ does not actively establish the interpretation.

The same aggregate shift in anchor levels could also arise from micro mechanisms other than equilibrium selection. Reference-dependent price setting, a reversal of information cascades, and selection effects under state-dependent pricing (Kaihatsu et al., 2023) are all candidates. The model identifies the reduced-form signature common to all of them, a level shift in the nominal anchor under an invariant structure. Determining which micro mechanism operated requires micro data, such as the distribution of price changes or firm surveys, and is left for future work. The ANCHOR component is attributed candidly, as the deterministic component that absorbs a systematic nominal shift in regime levels.

What a change in variances might mean is a separate question, on which the estimates do not point to a single reading. The nested comparison shows that allowing the shock variances to switch improves fit in every episode, but it does not by itself say what such a change represents. The direction of the switch differs by shock and by episode: the domestic price shock widens in Japan 2022 but narrows in both Japan 2013 and the United States, while the energy shock widens in Japan 2022 and, most sharply, in the United States. This pattern is difficult to reconcile with a single omitted variable that a mildly misspecified model would absorb into the variances, since a common omission would tend to move the same shock in the same direction across episodes. It is more naturally read as a change in the external environment specific to each episode, in the spirit of the "good luck" account of the Great Moderation (Stock and Watson, 2002; Justiniano and Primiceri, 2008): the wider energy variance in the United States coincides with the second oil shock of 1979, and the wider price and energy variances in Japan 2022 with the cost pressures of a weak yen. Because the anchors are made explicit, a systematic shift in regime levels is not misattributed to drifting shocks (Corollary 1), so the variance movement that remains is better read as a change in the stochastic environment than as an artifact of the level shift. Why anchors and variances switch together in every episode is left open.

Historically, the view that expectations re-anchor discontinuously at a regime change goes back to Sargent (1982). He showed, historically and qualitatively, that the end of several hyperinflations can be understood as a re-anchoring of expectations through a change in the fiscal and monetary regime, without a large contraction in the real economy. The present paper does not replace that insight. Its contribution is to make the insight structurally measurable, by identifying the price anchor and the wage anchor separately and simultaneously, by testing the invariance of the transmission mechanism, and by decomposing contributions quantitatively through de-anchoring. It supplies a quantitative measuring device, the simultaneous identification of price and wage anchors, for a regime change that has been described qualitatively.

The framework is not confined to cases in which only the anchors move. Where a level shift coexists with a change in propagation, whether in slopes or in variances, de-anchoring still separates the two components. Japan 2022 is an anchors-only episode, in which structural steepening (candidate A) is rejected and the level shift explains almost everything (Section 3.2). The Volcker disinflation of 1979 to 1983, by contrast, has usually been read as a demand contraction working through the slope. In the results here no switch in the price slope is required, and the downward shift in regime means is attributed to the anchor component. In the United States, however, switching variances also contributes substantially to model fit, with the variances-only specification exceeding the anchors-only specification by 27.81 in marginal likelihood, and the specification containing both is best (Section 3.1 and Appendix G). That the downward shift is significant on its own, without requiring a switch in slopes, is consistent with a Sargent-style reading of expectational re-anchoring.

Similar analysis could be applied to the adoption of inflation targeting in New Zealand, Canada, and the United Kingdom, or to post-COVID inflation in the euro area and the United States. Empirical application across many countries and periods lies beyond the scope of this paper.

4.3 Policy implications

The results indicate that the demand channel is weak, within the range this paper captures. The proxy for the output gap in the POST period, an HP-filtered gap in real consumption from the Family Income and Expenditure Survey for Japan, is close to zero and shows no significant shift, and the slope of the price Phillips curve is small at $\kappa_p = 0.010$. This is a statement about the demand channel as captured by a consumption-based proxy. It does not deny effects operating through an aggregate output gap or through credit. Subject to that qualification, what can be said is that the path from a rate rise to a lower output gap to lower inflation appears small within the range estimated here. Since this regime shift is a level shift in the expectational anchor rather than a movement in demand, the role of monetary policy lies instead in re-fixing the new nominal anchor at a level consistent with the 2% target. The same reading applies to the Volcker period, which has been treated as a textbook disinflation through demand contraction. As Section 3.2 shows, the Volcker disinflation too is a downward shift in anchors under invariant slopes, and decomposes structurally as re-anchoring rather than as a demand channel operating through a steeper slope.

The results place two reservations on the idea of a virtuous cycle of wages and prices. First, nominal wages rose but real wages did not improve significantly, at -0.8% per year with $p = 0.37$. The output gap shows no significant shift either ($p = 0.95$). Given that the price anchor and the wage anchor shifted upward by almost equal amounts, this is not a coincidence but an implication (Section 3.1). A one-for-one nominal switch leaves the real wage anchor unmoved by definition.

For a virtuous cycle to deliver a durable improvement in welfare, real wages must rise alongside productivity. An upward shift in the nominal anchor is not enough.

Second, the cycle does not close. Estimating two-way contemporaneous feedback leaves the coefficient from prices to wages, λ_p, and the coefficient from wages to prices, λ_w, with 90% intervals that include zero in all four Japanese regimes (Section 3.3 and Appendix F). The loop gain likewise cannot be distinguished from zero. Even though the levels of prices and wages both moved upward, the contemporaneous interaction linking them is not detected at all. The mutually reinforcing mechanism that policymakers have described as a virtuous cycle cannot be confirmed in the data, at least over the sample and under the specification used here. That the two moved in the same direction is better read, on these estimates, as their common anchor moving at the same time, rather than as each pushing the other.

Reading the U.S. disinflation as a downward shift in the wage norm faces a long-standing counter-hypothesis. Against Mitchell (1985), Fischer argued that what occurred was not a fall in norms but the unwinding of an overshoot in union wage levels during the 1970s, and Flanagan (1984) explained the same wage concessions as long-run flexibility in response to market conditions rather than as a shift in norms. The evidence here is hard to reconcile with these level-unwinding and market-response accounts. First, the invariance of the slopes (κ_p, κ_w) is shown competitively in estimation (Section 3.2), so increased flexibility as a stronger market response is not supported. Second, that the price anchor and the wage anchor move at the same time and by the same amount, fact B for Japan 2022, points to a move in the nominal anchor itself rather than to a wage-specific overshoot in levels. Those counter-hypotheses are not rejected decisively here, however. What it can distinguish is one dimension, the invariance of the propagation structure.

The finding that no switch in the price or wage slopes is required suggests that the current inflation environment is not necessarily irreversible. What moved is the level of the nominal anchor, and there is no evidence of change on the side of the transmission mechanism that supports it. Since the level of the anchor is a quantity that can move between regimes, a return to the earlier low level is possible in principle, through normalization of external factors or through a re-fixing of expectations. The formation of anchors is not endogenized here, however, and the causal effect of policy on them is not identified. What follows is therefore limited: policy diagnosis needs to observe and evaluate the anchors for prices and for wages separately. That an anchor can switch in the opposite direction is demonstrated by the downward shift in the United States in 1979 (Section 3). Re-anchoring is a phenomenon that can run in both directions, and the persistence of an upward shift is not automatically guaranteed.

5. Conclusion

The starting point of this paper was an observation about Japan after 2022: prices and wages shifted upward at the same time while the real side, real wages and the output gap, stayed still. Whether that simultaneous shift was an independent switch, a wage–price spiral, or a shift in a common anchor cannot be identified by extracting trends series by series. This paper answered the question by Bayesian estimation of an NK-DSGE model containing steady-state anchors for prices, wages, and the real rate, by determining which axis a regime shift loads on, against slopes and variances, and by applying the same framework to three episodes.

Four main findings emerge. First, allowing a regime switch in either the price or the wage slope does not improve fit. Adding the slope axis to the best specification moves the marginal likelihood by no more than 0.54. Moreover, even when the estimated slopes vary by a factor of several, the estimated anchor shifts do not change sign and move little over the range of calibrations the data support (Appendix C). Within the model class considered, a regime change in the price and wage slopes is not required, and the anchors are not absorbing a change in the slopes.

Second, the axis on which a regime change loads differs by episode. Only in the United States in 1979 do variances alone beat anchors alone in model fit; in the two Japanese episodes the reverse holds. The framework detects variances when variances move.

Third, the wage anchor shows a different estimated pattern in each of the three episodes. In 2013 the only significant detection is a rise in the price anchor. In 2022 prices and wages move at the same time by almost equal amounts, and their difference cannot be distinguished from zero. In the United States in 1979 wages move down about twice as far as prices. These are facts complete within each episode; the differences across episodes are not themselves significant, given how short the samples are (Section 3.4).

Fourth, in estimates of two-way feedback the self-reinforcing loop gain is never positive. It cannot be distinguished from zero in the four Japanese regimes and is significantly negative in the two U.S. regimes. The individual coefficients λ_w and λ_p also include zero in all twelve intervals.

Together these results characterize Japanese inflation after 2022 not as a change in structure but as a collective, synchronized level shift in nominal anchors. With the transmission mechanism held invariant, the nominal anchors for prices and wages switched upward one for one, and the real side did not move. The contribution of this paper is to separate and quantify that phenomenon for prices and wages simultaneously, within a structural DSGE, and then to apply the same measuring device to three episodes and show that the results differ. What the central bank could move on its own was the price anchor alone, in 2013. What happened in 2022 was something else: the anchors for prices and for wages switched at the same time and by the same amount.

The design has the following limitations. The test for two-way price and wage feedback (Section 3.3) and the test for import price pass-through (Section 4.1) are both confined to a linear, time-invariant, contemporaneous specification. The loop gain is never positive, but that is a statement within the estimated linear approximation. It does not rule out different dynamics at far higher rates of inflation. Likewise, $\chi \approx 0$ shows that contemporaneous linear pass-through is negligible; it does not deny state-dependent pass-through during a large depreciation. In addition, producer prices enter the model only through the energy cost term, and no separate producer-price anchor is carried. Producer and consumer prices diverged substantially in Japan in 2022, and whether the nominal anchor at the producer stage differs from the one at the consumer stage cannot be examined within the present framework. In addition, with the Japanese policy rate pinned at the effective lower bound, the identification of $\bar{r}$ is relatively weak, and that weakness comes from the nature of the data rather than from the model (Appendix C.4).

The questions that follow arise from the facts themselves. The first is why a level shift in anchors and a change in shock variances occur together. In the nested comparison, the best specification in all three episodes was MBC, which switches anchors and variances at once. The co-occurrence is systematic rather than accidental. The belief mixture of Section 2.3 offers a natural candidate explanation. During a transition in which the share f_t of agents who believe the new regime lies between zero and one, heterogeneity of beliefs generates a term proportional to $f_t(1-f_t)$ in second moments, raising variances just after a regime change. Once credibility settles and f_t approaches one, the term vanishes. This paper chose to exclude that phase from estimation (Section 2.9), but the same window could instead be used to estimate the path of f_t and to test whether the fall in variances tracks the convergence of beliefs. The view carries a testable implication. The U.S. POST period, 1983 to 1990, sits at the entrance to the Great Moderation, and the POST period for Japan 2022 is still only 52 months. If convergence of beliefs drives the fall in variances, Japanese variances should decline going forward.

The second is to endogenize the anchors themselves. The anchors here are exogenous parameters that move at a date specified by the researcher. It would be desirable to incorporate a learning approach (Hogen and Okuma, 2025) and to endogenize how a level shift in anchors forms and settles as an equilibrium object of belief updating. A full regime-switching DSGE that treats the regime as a latent state would accommodate recurrent regime transitions. Whether the break date itself is the binding constraint is less clear: making the regime latent in a univariate Markov-switching model already reproduces the split imposed here in all three episodes (Section 3.4). The gain from a latent-state DSGE would lie in the joint treatment of transitions rather than in the location of the boundary.

The third is to identify the mechanism that produces a one-for-one simultaneous shift. That the price and wage anchors moved by statistically indistinguishable amounts in Japan in 2022 is observationally equivalent, in aggregate data, to at least three readings: genuine coordination through the *shuntō* mechanism; both anchors responding independently to a common signal, namely the establishment of credibility for the 2% target; and mechanical indexation. Distinguishing among them requires micro data, such as the distribution of *shuntō* settlements across firms, price expectations from firm surveys, and the timing of wage negotiations relative to price revisions.

The fourth is to turn the episodes into a panel. The estimation here applies to three episodes by substituting episode-specific settings into common code, and the calibration invariance of Appendix C is what makes that substitution legitimate rather than merely convenient. Exploiting that portability to cover the euro area in 2021 and 2022, the United Kingdom, or the countries that adopted inflation targeting in the early 1990s would raise the contrast between Japan 2013, where only the price anchor moved, and Japan 2022, where both anchors moved one for one, from a pair of anecdotes to a systematic comparison. The conditions that determine whether the wage anchor follows the price anchor, such as the centralization of wage bargaining, the credibility of the central bank, and the tightness of the labor market, become testable only once the number of episodes grows. Finally, this paper is purely positive and makes no normative assessment. In Japan in 2022 the nominal anchors moved one for one while real wages did not improve. For whom coordinated nominal adjustment is desirable requires evaluation in a framework that makes distribution explicit.

More fundamentally, there is the generalization of the identification framework itself. The identification here closes cleanly, as a first formulation, under strong assumptions (Appendix A.2). Those assumptions are that the belief share f_t is deterministic and monotone, that the wage slope κ_w is calibrated, and that the structural core carries no exchange rate or import price block. Relaxing them is the natural direction of development: stochastic and non-monotone paths for f_t; a generalization that allows heterogeneous speeds of belief propagation on the price and wage sides, with $f_t^p \neq f_t^w$; joint identification of κ_w; extensions to an open economy, to financial frictions, and to multiple sectors; and application to re-anchoring episodes beyond Japan and Volcker, such as the adoption of inflation targeting or the introduction of a common currency. What matters is that Proposition 1 delivers separating identification without assuming structural invariance. That makes the framework a direction-neutral and general way of measuring regime shifts, decomposing them into changes in the transmission mechanism and level shifts in anchors, and it is on this basis that the extensions can build.

A. Theoretical Appendix: Derivations and Proofs

A.1 Deriving the nominal anchor under Calvo pricing

This appendix derives the result stated in Section 2.3, that the intercept $(1-\beta)\bar{\pi}$ of the price Phillips curve (2) appears as the log of a multiplicative expected-inflation anchor. The wage side, $(1-\beta)\bar{\pi}^w$ in (3), follows in the same way.

Consider a firm that can reset its price with probability $1-\theta_p$ in each period. A firm with a resetting opportunity maximizes expected discounted profit under a probability θ_p of keeping the price fixed, and the reset price X_t satisfies

$$X_t = \frac{\varepsilon}{\varepsilon-1}\,\frac{\mathbb{E}_t \sum_{j=0}^{\infty}(\beta\theta_p)^j \Lambda_{t,t+j} Y_{t+j} P_{t+j}^{\varepsilon} MC_{t+j}}{\mathbb{E}_t \sum_{j=0}^{\infty}(\beta\theta_p)^j \Lambda_{t,t+j} Y_{t+j} P_{t+j}^{\varepsilon-1}} \tag{13}$$

where $\Lambda_{t,t+j}$ is the stochastic discount factor, MC is nominal marginal cost, and ε is the elasticity of demand. Normalizing by the current price P_t and using $\Pi_{t,t+j} = P_{t+j}/P_t = \prod_{k=1}^{j}\Pi_{t+k}$, future inflation enters multiplicatively, as $\Pi_{t,t+j}^{\varepsilon}$ and $\Pi_{t,t+j}^{\varepsilon-1}$.

Suppose that, under the equilibrium selected in regime s_t, firms' expectations of future inflation are anchored to $\Pi^*(s_t) = e^{\pi^*(s_t)}$:

$$\mathbb{E}_t[\Pi_{t,t+j} \mid s_t] = \Pi^*(s_t)^{\,j}. \tag{14}$$

This is not a subjective distortion. It is rational expectations, taking the equilibrium s_t as given. Decompose expected cumulative inflation as

$$\Pi_{t,t+j} = \Pi^*(s_t)^{\,j}\,\exp\Big(\textstyle\sum_{k=1}^{j}\hat{\pi}_{t+k}\Big), \qquad \hat{\pi}_{t+k} \equiv \pi_{t+k} - \pi^*(s_t), \tag{15}$$

where $\hat{\pi}$ is the stationary deviation around the equilibrium, with $\mathbb{E}_t\hat{\pi}_{t+k} = 0$. The anchor component $\Pi^*(s_t)^{\,j}$ factors out as a multiplicative constant kernel.

Substituting this decomposition into the reset price gives the multiplicative factorization

$X_t = A^{\pi}(s_t)\,\tilde{X}_t$, where $A^{\pi}(s_t)$ is the anchor kernel, a constant determined by the power sums of $\Pi^*(s_t)$, and $\tilde{X}_t$ is the stationary deviation component. Log-linearizing this together with the price index $P_t^{1-\varepsilon} = (1-\theta_p)X_t^{1-\varepsilon} + \theta_p P_{t-1}^{1-\varepsilon}$, the log of the anchor kernel becomes

$$\log A^{\pi}(s_t) \implies (1-\beta)\,\pi^*(s_t) \tag{16}$$

and appears as the additive intercept of the price Phillips curve. The deviation component $\tilde{X}_t$ delivers the standard slope term $\kappa_p = \dfrac{(1-\theta_p)(1-\beta\theta_p)}{\theta_p}$ and the forward-looking term $\beta\,\mathbb{E}_t\pi^p_{t+1}$. Equation (2) is recovered, and $\pi^*(s_t) = \bar{\pi}$ is identified.

The slope κ_p depends only on the fixed-price probability θ_p, the discount factor β, and the elasticity ε. To first order it does not depend on the anchor $\pi^*(s_t)$. In the generalized NKPC of Ascari and Sbordone (2014) the linearization is taken around a time-varying trend, so the slope becomes a function of $\bar{\pi}$; under the reference point used here, that dependence is pushed into higher order, of about $O(\pi^{*2})$. This paper does not simply assume that higher-order slope dependence away. It tests explicitly, as candidate A (Section 3.2), whether the slopes moved across regimes, and rejects it. The first-order separation into the intercept, in which the anchor is the log of the expectational kernel, is therefore exact, and slope invariance is an empirically supported restriction rather than an approximation imposed by hand.

The mixed anchor of the text, equation (8), is obtained by Calvo aggregation over price setters with heterogeneous beliefs. Let the Calvo fixed-price probability be θ and the elasticity of demand $\epsilon > 1$, with no indexation. Among resetting firms, a share f_t believes the new regime, with reference inflation $\bar{\pi}_{\text{POST}}$, and a share $1-f_t$ does not, with $\bar{\pi}_{\text{PRE}}$. The log reset price of type j is $x_{j,t} = p_{t-1} + \bar{\pi}_j + \mu_t$, so the difference in beliefs enters only through the reference inflation $\bar{\pi}_j$, while the common real term μ_t, the discounted present value of marginal cost, is independent of type. Aggregation of the CES price index gives

$$P_t^{1-\epsilon} = \theta\,P_{t-1}^{1-\epsilon} + (1-\theta)\big[f_t\,X_{\text{POST},t}^{1-\epsilon} + (1-f_t)\,X_{\text{PRE},t}^{1-\epsilon}\big], \qquad X_{j,t} = e^{x_{j,t}}, \tag{17}$$

and expanding to second order around the steady state gives the following structure. At first order the intercept of aggregate inflation is exactly the anchor mixed by the share of believers,

$$\pi_t^{p,(1)} = (1-\theta)\big[\,f_t\bar{\pi}_{\text{POST}} + (1-f_t)\bar{\pi}_{\text{PRE}} + \mu_t\,\big]$$

which reproduces the linear mixed anchor of equation (8). Wage setting under the same f_t gives $\bar{\pi}_t^w = f_t\bar{\pi}^w_{\text{POST}} + (1-f_t)\bar{\pi}^w_{\text{PRE}}$. At second order, exact aggregation carries a correction relative to homogeneous aggregation, in which everyone uses the same effective anchor. The correction is proportional to the cross-type variance of reset prices, $\overline{x^2} - \bar{x}^2 = f_t(1-f_t)\,s^2$ with $s \equiv \bar{\pi}_{\text{POST}} - \bar{\pi}_{\text{PRE}}$:

$$\Delta\pi_t^{p,(2)} = \frac{(\epsilon-1)(\theta-1)}{2}\,f_t(1-f_t)\,s^2 \tag{18}$$

The correction term (18) has two properties. First, it is deterministic, a function of f_t and the constant s with no stochastic shock, so it does not enter the autocovariances $\Gamma(h)$ and does not contaminate the second-moment block of slopes and variances. Second, being proportional to $f_t(1-f_t)$, it vanishes at both plateaus, $f_t \in \{0,1\}$. The term is therefore non-zero only during a transition in which beliefs are mixed, and it does not contribute to estimation under the design used here, which estimates on the PRE and POST plateaus and excludes the transition window (Sections 2.9 and 3). On the plateaus the linear mixed anchor holds exactly rather than approximately, and the second-order heterogeneity term does not damage the separating identification of Proposition 1. Numerically, at the monthly anchor difference s, the magnitude of (18) is negligible. That heterogeneous beliefs introduce a dispersion-dependent additional term into the NKPC is shown for general belief distributions by Meeks and Monti (2023), in their Δ_t term. The correction (18) specializes that result to a two-point distribution of believers and non-believers, and gives it the closed form $\frac{(\epsilon-1)(\theta-1)}{2} f_t(1-f_t)s^2$. Where Meeks and Monti (2023) estimate the additional term from distributional data, however, this paper uses the fact that it vanishes at both plateaus, through $f_t(1-f_t)$, to justify excluding the transition window. The term is used to preserve identification, not as an object of estimation.

Using the same share f_t for firms and households is a strong assumption. Firms that set prices and households that bargain over wages are not necessarily the same decision makers. It is adopted here as the simplest benchmark. If beliefs about the new regime form in an information environment shared across society, a single f_t is a natural starting point. What differs between the two sides is the terminal level, with $\pi^{p*}_{\mathrm{POST}} \neq \pi^{w*}_{\mathrm{POST}}$, not the path of beliefs f_t itself. This is where the phrase coordinated, through a common f_t, with different norms, through different levels, comes from. Generalizing to different speeds of belief propagation on the price and wage sides, with $f^p_t \neq f^w_t$, as when firms come to believe before households, is left for future work.

One concern deserves an answer: whether the assumption of a common f_t is driving the results. Estimation here uses only the two ends at which f_t is flat, $f \approx 0$ in PRE and $f \approx 1$ in POST, and excludes from estimation the transition window over which f_t moves quickly (Section 2.9). At both plateaus $f_t \in \{0, 1\}$, so the effective anchor (8) reduces exactly to π^{j*}_{PRE} at $f = 0$ and to π^{j*}_{POST} at $f = 1$, and the common f_t term drops out at both ends. The terminal levels on the price and wage sides, π^{p*}_{POST} and π^{w*}_{POST}, together with their PRE counterparts, are therefore free parameters estimated independently in each regime, and the common f_t plays no part in identifying them. The common f_t is an interpretation of the transition process, and does not restrict identification of the anchor levels at the two ends. As a result, the finding that the price and wage anchors shift by different amounts (Section 3.1) is not a product of the common f_t. It is simply the difference between terminal levels estimated independently at the two ends. Generalizing to $f^p_t \neq f^w_t$ affects only the path of the transition, and does not change the identification results at the plateaus, and hence does not change the estimated anchors or their differences.

Writing the law of motion for f_t as social learning, in which believers increase in proportion to their contact with existing believers, gives the logistic equation

$$\dot{f}_t = \gamma\, f_t(1 - f_t) \quad \Longrightarrow \quad f_t = \frac{1}{1 + e^{-\gamma(t-t_0)}} \tag{19}$$

where γ is the speed of propagation and t_0 the inflection point. Since $f = 0$ cannot take off on its own, the phase transition begins only once an exogenous trigger injects a small initial spark $f_0 > 0$. This is consistent with treating the trigger for re-anchoring as exogenous. The path f_t is a sigmoid: monotone, deterministic, and one-directional, and flat at both ends, $f \approx 0$ and $f \approx 1$, outside a short window around t_0.

A deterministic, one-directional f_t is orthogonal in the spectrum to stationary stochastic shocks, as a deterministic component against a stochastic one. This is the basis for the claim in Proposition 1 that the anchor transition is not absorbed into the variances, and the Monte Carlo experiments of Appendix E confirm the separation in finite samples. Because estimation uses only the two flat phases, neither γ nor t_0 needs to be estimated. The anchor identification of this paper depends only on f_t being monotone and deterministic, not on its specific functional form. Identification closes cleanly under the strong assumption that f_t is logistic and deterministic; generalizing to stochastic variation in f_t, to non-monotone transitions, and to direct estimation of γ from belief panels is left for future work.

A.2 Proof of the separating identification result

This appendix gives the complete proof of Proposition 1, the non-confounding of anchors and centered dynamics.

In each regime phase, that is, the PRE phase with $f \approx 0$ and the POST phase with $f \approx 1$, where f_t is flat (Appendix A), the model is a time-invariant state space,

$$s_t = G(\theta)\, s_{t-1} + B\, \varepsilon_t, \qquad \varepsilon_t \sim (0, \Sigma), \tag{20}$$

$$y_t = d(\mu) + Z\, s_t, \tag{21}$$

where s_t is the state vector, $y_t = (\pi_t^p, \pi_t^w, \pi_t^{en}, i_t, x_t)'$ the observables, Z a selection matrix, $d(\mu) = [\bar{\pi}; \bar{\pi}^w; 0; \bar{r} + \bar{\pi}; 0]$ the deterministic observation level implied by the anchors (Section 2.9), and $\theta = (\kappa_p, \kappa_w, \Sigma)$. Energy π^{en} is a mean-zero exogenous AR(1) block with no deterministic level, so the corresponding element of $d(\mu)$ is zero. The matrix $G(\theta)$ is assumed stable. The anchors μ enter only $d(\mu)$, and the slopes and variances θ enter only $(G, B\Sigma B')$.

Lemma 1 (First-moment block)**.** In each phase, $\mathbb{E}[y_t] = d(\mu)$, and the map $\mu \mapsto d(\mu)$ is bijective. The anchors $\mu = (\bar{\pi}, \bar{\pi}^w, \bar{r})$ are therefore identified uniquely from the stationary means: $\bar{\pi} = \mathbb{E}\pi^p$, $\bar{\pi}^w = \mathbb{E}\pi^w$, and $\bar{r} = \mathbb{E}i - \mathbb{E}\pi^p$.

Proof. Since G is stable, s_t is mean-zero stationary with $\mathbb{E}s_t = 0$, so $\mathbb{E}[y_t] = d(\mu) + Z\,\mathbb{E}s_t = d(\mu)$. The first three elements of $d(\mu)$ are $(\bar{\pi}, \bar{\pi}^w, \bar{r} + \bar{\pi})$, which is triangular and nonsingular in μ, so the inverse map is as stated. □

Lemma 2 (Second-moment block and invariance of the mean)**.** The autocovariances $\Gamma(h) \equiv \mathrm{Cov}(y_t, y_{t-h})$ for $h \geq 0$ depend only on $(G(\theta), B\Sigma B')$ and are invariant to μ. With P solving the Lyapunov equation $P = GPG' + B\Sigma B'$, one has $\Gamma(0) = ZPZ'$ and $\Gamma(h) = ZG^hPZ'$ for $h \geq 1$.

Proof. Since $y_t - \mathbb{E}y_t = Zs_t$, the deterministic level $d(\mu)$ is subtracted out. The stationary covariance P of the state solves the Lyapunov equation and depends only on $(G, B\Sigma B')$, and $\mathrm{Cov}(s_t, s_{t-h}) = G^hP$ gives the result. The anchor μ appears nowhere. □

Lemma 3 (Block diagonal information: orthogonality of the anchor and propagation blocks)**.** The log likelihood of a Gaussian stationary process decomposes asymptotically into a mean part and an autocovariance part, and the asymptotic expected information matrix in (μ, θ) is block diagonal, $\mathcal{I} = \operatorname{diag}(\mathcal{I}_\mu, \mathcal{I}_\theta)$. This does not claim that the cross second derivative of the realized finite-sample likelihood is exactly zero.

Proof. Under the Whittle approximation, the log likelihood of a stationary Gaussian process separates into a quadratic form in the mean and a term in the autocovariances, or equivalently the spectrum. By Lemma 1 the former is a function of μ alone, and by Lemma 2 the latter of θ alone. The expected cross second derivative $\mathbb{E}[\partial^2/\partial\mu\,\partial\theta]$ therefore vanishes, and the asymptotic information matrix is block diagonal. □

From Lemmas 1 to 3, μ is identified in each phase from first moments while $\theta = (\kappa_p, \kappa_w, \Sigma)$ enters only second moments, so the information about the two is block diagonal, that is, orthogonal. A perturbation of θ, the slopes and variances, therefore does not move the identified μ, and the converse holds as well. Taking differences between the two phases, $\Delta\bar{\pi}$, $\Delta\bar{\pi}^w$, and $\Delta\bar{r}$ are identified separately from $\Delta\kappa_p$ and $\Delta\Sigma$, without confounding.

Over the full sample, from PRE through the transition to POST, the anchor path is deterministic, $\mu_t = f_t\mu_{\text{POST}} + (1 - f_t)\mu_{\text{PRE}}$ as in (8), with f_t monotone and one-directional (Appendix A). A deterministic, one-directional path in levels has a discrete spectrum and is orthogonal to the absolutely continuous spectrum of the stationary stochastic component $\tilde{y}_t$. Equivalently, no realization of mean-zero stationary shocks ε generates a non-zero deterministic mean path. The anchor path therefore cannot be absorbed into Σ. Decomposing under a common level μ_0 forces that deterministic level onto the smoothed innovations as $G^{-1}b$ (Corollary 1).

Proof of Corollary 1. Kalman smoothing gives the minimum mean squared error estimate of the innovations, taking the model as given. Under a misspecified constant mean μ_0, the observation residuals contain a constant bias $b = \mu(s_t) - \mu_0$ over the regime interval. The only stochastic input is the innovations, and the model maps them into the observables through a stable, that is, stationary, filter. Reproducing a constant bias b over an interval therefore requires a sequence of innovations with non-zero mean, $\hat{\varepsilon}_t \to G^{-1}b$. That sequence is serially correlated within the regime window and satisfies $\mathbb{E}[\hat{\varepsilon}_t \mid s_t] \neq 0$, contradicting the assumption of i.i.d. mean-zero shocks. The level shift is thus loaded onto the innovations. Subtracting $\mu(s_t)$ sets $b = 0$ and restores $\mathbb{E}[\hat{\varepsilon}_t] = 0$. The argument uses only $|b| > 0$ and does not depend on its sign. □

By Lemma 2, identification of θ reduces to the map into the autocovariances $\{\Gamma(h)\}_{h=0}^{H}$. In the five-observable core used here, calibrating κ_w makes the Jacobian $\partial\operatorname{vec}\{\Gamma(h)\}/\partial(\kappa_p, \operatorname{vech}\Sigma)$ of full column rank at the true values, so local identification holds (Iskrev, 2010; Komunjer and Ng, 2011). Intuitively, κ_p moves the shape of the autocovariances, the lag structure between inflation and the gap, through the price Phillips curve, while Σ moves their scale, so the two are recovered separately. This local non-degeneracy is a population property. The finite-sample power for κ_p is nonetheless modest, because $x_t \approx 0$ in the POST period (Appendix E), which is why κ_p is handled on a grid rather than estimated. Non-confounding itself, the block diagonality, is confirmed in finite samples as well: cross-axis leakage between anchors and propagation is small and does not grow with the size of the true change (Appendix E). Proposition 1 follows.

This proof closes non-confounding cleanly under strong assumptions: f_t deterministic and monotone, κ_w calibrated, and a structural core with no exchange rate or import price block.

These were placed deliberately, as a first formulation. Stochastic variation and non-monotone transitions in f_t, joint identification of κ_w, and extension to an open economy or to multiple sectors are left to later generalization. What deserves emphasis is that structural invariance is not assumed. The proof delivers separating identification in a framework where all three components may change, so the empirical finding $\Delta\kappa_p \approx 0$ is an identified result rather than an imposed assumption.

B. Data and Estimation Details

B.1 Observation equations and calibration

This appendix collects the observation equations (Table B.1) and the calibration and estimation classification of parameters (Table B.2) referred to in Sections 2 and 2.9.

Table B.1: Observables and their model counterparts

Observable	Data source (index level input; SA, month-on-month %)	Model counterpart
`pi_cc`	Core CPI, all items less fresh food and energy	π_t^p
`piw`	Nominal wages, total cash earnings (Monthly Labour Survey)	π_t^w
`pi_en`	Energy prices (energy components of the CGPI)	π_t^{en}
`ir`	Policy rate (level)	i_t
`ipigap`	HP-filtered gap in real consumption	x_t

Table B.2: Classification of parameters: calibrated and estimated

Class	Parameter	Value / prior	Source and notes
Calibrated (fixed)	β	0.99	Monthly discount factor
	σ	1.0	Inverse elasticity of intertemporal substitution
	ρ_i	0.9244	Clarida et al. (2000): quarterly 0.79 converted to monthly as $0.79^{1/3}$
	$\phi_\pi,\ \phi_x$	1.50, 0.50	Taylor (1993): long-run coefficients, hence frequency independent
	γ (energy share)	0.0263	2020 input-output table, 108 sectors (petroleum and coal products, electric power, gas); $\zeta = \gamma\kappa_p$
	ρ_{en}	0.5361	OLS AR(1) on seasonally adjusted π^{en} (Japan 2022 estimation window)
	$\kappa_p,\ \kappa_w$	0.010, 0.005	Best cell of the 35-cell grid (Japan 2022)
Estimated (variances)	σ_{ε^x}	Inv-Γ(1.00, 0.70)	Demand shock
	σ_{ε^p}	Inv-Γ(0.15, 0.10)	Domestic price shock
	σ_{ε^w}	Inv-Γ(0.30, 0.20)	Nominal wage shock
	σ_{ε^i}	Inv-Γ(0.02, 0.015)	Monetary policy shock
	$\sigma_{\varepsilon^{\text{en}}}$	Inv-Γ(1.00, 0.70)	Energy shock
Estimated (anchors)	$\bar\pi$	$\mathcal{N}(0,\ 0.30)$	Price anchor
	$\bar\pi^w$	$\mathcal{N}(0,\ 0.30)$	Wage anchor
	$\bar r$	$\mathcal{N}(0,\ 0.30)$	Steady-state real rate

B.2 Data and estimation methods

This appendix supplements the estimation methods used in Sections 3.4 and 2.9.

The four price and wage series are held as index levels, seasonally adjusted with X-13ARIMA-SEATS (X-11 method, log transformation, automatic ARIMA identification), and then converted to month-on-month rates. The implementation calls the Census Bureau `x13as` binary, the ASCII version, through Lengwiler's MATLAB Toolbox. The months of consumption tax rate changes, April 2014 and October 2019, are excluded from the window used to estimate seasonal factors, as irregular variation, and the COVID period of 2020 and 2021 is excluded as an outlier. In the structural estimation those two months are also treated as missing for `pi_cc` and `pi_en` (Section 2.9). The effective sample for the Japanese price series is therefore $n = 70$ in the PRE period of Japan 2022, excluding April 2014 and October 2019. Where X-13 is unavailable, the procedure falls back on stable seasonal adjustment of the log level. Seasonal adjustment reduces the standard deviation of each series, by -63% for `piw`, -39% for `pi_cc`, and -17% for `pi_en`, and the R^2 of the seasonal component after adjustment falls to about 0.01 to 0.02.

The mean shifts and their significance in the stylized facts (Section 3.4) are estimated as follows. For each variable, the X-13ARIMA-SEATS seasonally adjusted month-on-month rate, in percent, is regressed on a POST dummy, and the coefficient on that dummy is reported as the mean shift. Standard errors are Newey–West, with a HAC lag of twelve. The price series are controlled for the consumption tax rate changes of April 2014 and October 2019 with dummies. The output gap x is obtained by applying an HP filter ($\lambda = 129{,}600$, the standard monthly value) to the log level formed by cumulating seasonally adjusted month-on-month real consumption. The estimation sample runs from January 2014 to April 2026, excluding 2013, the transition window, and the COVID period.

A note on identification of the break date (Fact 3). The price series are controlled for the months of tax rate changes with dummies. Because the excluded transition windows, 2013 and the COVID period of 2020 and 2021, contain no observations, the break date could formally be identified anywhere inside those gaps. The estimates given by the sup-Wald (Quandt–Andrews) statistic nonetheless converge cleanly on the start of 2022 for both series, and do not settle in the gaps. That the breaks in prices and in wages are synchronized at the start of 2022 supports treating the two under a common regime boundary.

B.3 Construction of the series and the estimation procedure

The price and wage series (`pi_cc`, `piw`, `pi_en`) are entered as index levels, seasonally adjusted with X-13ARIMA-SEATS, and then converted to month-on-month rates in percent (Appendix B.2 gives the implementation). For the months in which the consumption tax rate changed, April 2014 and October 2019, the corresponding observations of `pi_cc` and `pi_en` are treated as missing and left to the missing-data handling of the Kalman filter. A change in the tax rate is a statutory across-the-board change in prices, not a change in how firms set prices. What this paper identifies is price-setting behavior, and a one-off price change caused by an institutional reform is neither a structural shock nor a change in an anchor. It is therefore excluded from estimation. Leaving it in the observations biases both the regime-specific price shock variance and the price anchor upward. Controlling for it with a dummy is an alternative, but a dummy with a free coefficient at a single date is, in a linear Gaussian state space, effectively equivalent to treating that observation as missing. Missing-data handling is the more transparent of the two, because it requires no prior assumption about the rate of pass-through.[14] Wages, the policy rate, and the output gap have no direct counterpart in the tax change, and none of them is adjusted. The U.S. episode contains no change in tax rates, so the treatment does not apply there.

The policy rate (`ir`) is used in levels. The output gap (`ipigap`) is built from the real index of household consumption expenditure excluding housing and related items, taken from the Family Income and Expenditure Survey. A Hodrick–Prescott filter ($\lambda = 129{,}600$, the standard monthly value) is applied to its log level. The data sources are the Statistics Bureau of Japan (CPI and the Family Income and Expenditure Survey), the Ministry of Health, Labour and Welfare (Monthly Labour Survey), and the Bank of Japan (Corporate Goods Price Index and the policy rate). All U.S. series are taken from FRED, with the same five observables as for Japan; the series and their FRED identifiers appear in Appendix G.

Seasonal adjustment of the wage series matters in particular for the identification of the wage anchor $\bar{\pi}^w$. Before adjustment the 90% interval for $\bar{\pi}^w$ straddled zero; after adjustment it lies

[14] In the descriptive statistics of Section 3.4 the object of interest is a mean shift as a regression coefficient, so a dummy is used there in the usual way (Appendix B.2).

clearly above zero (Section 3).

No HP filter is applied to the price and wage series. Month-on-month rates are already stationary, and applying an HP filter on top would amount to double filtering. It would risk removing artificially the very kind of persistent level change that the rise in inflation after 2022 represents.

The first stage is a grid search over the slopes (κ_p, κ_w). The grid has $7 \times 5 = 35$ cells, with $\kappa_p \in \{0.005, 0.01, 0.02, 0.03, 0.05, 0.10, 0.20\}$ and $\kappa_w \in \{0.005, 0.01, 0.02, 0.05, 0.10\}$. For each cell the pooled sample is estimated under the specification in which only the anchors switch across regimes, labeled MC below. The best cell is then chosen by comparing the log marginal data density (log MDD), obtained by Laplace approximation. The selected (κ_p, κ_w) is common to both regimes. Whether the slopes change across regimes is tested not on this grid but on the slope axis of the nested comparison (Section 3.2). The grid search uses the Laplace approximation for computational reasons, but the ranking of the main nested specifications is the same under the modified harmonic mean computed after full MCMC (Appendix D). For Japan 2022 the grid selects $(\kappa_p, \kappa_w) = (0.010,\ 0.005)$. Whichever cell of the grid is selected, the anchor shifts keep the same sign and the same order of magnitude (Appendix C).

A regime change in a slope is written as a log ratio, $\kappa^{\mathrm{POST}} = \kappa \exp(d)$. The additive form $\kappa + d$ is avoided because, when the selected κ lies at the lower end of the calibration grid, the positivity constraint falls near the mode and corrupts the Hessian of the Laplace approximation. The variance axis uses a log ratio for the same reason. The prior on both slope axes is $d \sim N(0, 1)$, chosen so that under the reference calibration it is about as wide as the earlier additive prior. The purpose is to remove the boundary, not to make the test looser or stricter.

The second stage is full MCMC at the selected cell. Five shock variances (σ_{ε^x}, σ_{ε^p}, σ_{ε^w}, σ_{ε^i}, $\sigma_{\varepsilon^{\mathrm{en}}}$) and three regime-specific anchors $(\bar{\pi}, \bar{\pi}^w, \bar{r})$ are re-estimated by Metropolis–Hastings, with 200,000 draws on each of two chains. Posterior means, 90% intervals, and the log marginal likelihood from the modified harmonic mean are obtained. All chains converge well; the diagnostics appear in Appendix B.4, and the fit of the model in Appendix B.5.

B.4 Full posterior for the POST period

Table H.2 in Section 3.1 focused on the regime shift in the three anchors, $\bar{\pi}$, $\bar{\pi}^w$, and $\bar{r}$. This appendix reports posterior means and 90% intervals for all eight parameters, the three anchors and the five shock variances, at the best cell of the POST period, $\kappa_p = 0.010$ and $\kappa_w = 0.005$ (Table B.3).

The MCMC diagnostics are as follows. Metropolis–Hastings is run on two chains with 200,000 draws each and a burn-in rate of 50%. The acceptance rate is about 0.42, the effective sample size is roughly 4,900, and the Gelman–Rubin $\hat{R}$ is at most 1.0004. Table D.3 lists the diagnostics for all three episodes.

The posterior of each shock variance is tighter than its prior, the inverse gamma of Table B.2, which confirms that the data update them informatively. The variances of the domestic price and monetary policy shocks in particular are small and precisely estimated, while the variance of the energy shock is large. These variances play only a secondary role in the central conclusion of the paper, the level shift in the anchors (Section 3.2).

Table B.3: Posterior estimates for the POST period (best cell $\kappa_p = 0.010$, $\kappa_w = 0.005$; full MCMC)

Parameter	Posterior mean (monthly)	90% interval [5,95] (monthly)	Annual (%)
$\bar{\pi}$ (POST)	+0.2272	[+0.2023, +0.2524]	+2.73
$\bar{\pi}^w$ (POST)	+0.1953	[+0.1278, +0.2633]	+2.34
$\bar{r}$ (POST)	−0.1957	[−0.2892, −0.1045]	−2.35
σ_{ε^x}	2.0431	[1.8417, 2.2714]	—
σ_{ε^p}	0.1015	[0.0913, 0.1128]	—
σ_{ε^w}	0.2973	[0.2677, 0.3303]	—
σ_{ε^i}	0.0789	[0.0712, 0.0875]	—
$\sigma_{\varepsilon^{\mathrm{en}}}$	1.7700	[1.5958, 1.9638]	—

Notes: Eight parameters in total, three POST-level anchors and five shock standard deviations. The estimation unit is monthly, and annual rates ($\times 12$) are shown alongside for the anchors. Standard deviations measure the dispersion of monthly shocks and are not annualized. The POST level of an anchor is PRE $+\Delta$ within the same chain. Under the MC specification the variances are regime invariant, so each standard deviation takes a single value with POST = PRE. Intervals are equal-tailed, $[5, 95]$.

B.5 Fit of the model

This appendix reports how well the estimated model fits the observed series, using the root mean square (RMS) of standardized one-step-ahead forecast errors, series by series. A value of 1.0 is ideal. Values well above 1 indicate undercoverage, meaning the model understates its own forecast variance, and values well below 1 indicate overcoverage.

An earlier specification of this paper fixed the standard deviation of the demand shock, σ_{ε^x}, at 0.05. That fixed value was far below the true one, and the model understated the one-step-ahead forecast variance of the output gap. As a result most of the likelihood came from forecast errors in the gap, and standardized forecast errors reached several σ. Moving σ_{ε^x} into the set of estimated parameters (Table B.2) brings the RMS for the gap close to 1 and improves log MDD per observation substantially. What matters is that the correction barely moves the anchor estimates. The anchors are identified from first moments, the mean level of each series, and are separated from misspecification of second moments such as shock variances (Proposition 1). That the fit improves while the anchors stay put is itself evidence that the identification strategy works as the theory says it should.

Series by series (Table B.4), the RMS for prices, wages, and energy is close to 1 in both countries. The gap also settles near 1. The exception is the Japanese policy rate, where the RMS falls well below 1. The reason is clear. Japan's PRE period sits at the effective lower bound, and the policy rate barely moves, with a monthly standard deviation of 0.012. The Taylor rule used here has no lower bound, so the model forecasts interest rate variation that never materialized. This is overcoverage, and it is not an error that distorts the likelihood. The serious form of misspecification is undercoverage, as in the earlier example of the gap, where standardized errors reached several σ. The consequence of overcoverage is no effect on $\bar{\pi}$ and $\bar{\pi}^w$, since both are identified from the means of prices and wages, and only a loss of precision for $\bar{r}$ (Appendix C.4). Modeling the effective lower bound explicitly would require a nonlinear solution method and lies beyond the scope of this paper. In the United States the interest rate was not constrained by a lower bound and did vary substantially, so the RMS is 0.65, only slightly below 1.

Table B.4: RMS of standardized one-step-ahead forecast errors by series (1.0 is ideal; < 1 is overcoverage, > 1 undercoverage)

Observable	Japan 2022	Japan 2013	U.S. 1979
π^p (core CPI)	0.985	1.001	0.999
π^w (wages)	0.985	0.992	0.993
π^{en} (energy)	0.999	0.998	0.998
i (policy rate)	0.058	0.053	0.651
x (gap)	1.172	1.163	1.070

C. Robustness of Identification: What Is Identified and What Is Not

This appendix and those through Appendix E all support the same question, whether identification really works, with different kinds of evidence. The division of labor is as follows. This appendix shows robustness on the actual data, that is, that the anchor estimates do not move under changes in calibration or across sweeps. Appendix D shows robustness to the choice of marginal likelihood estimator. Appendix E provides finite-sample verification by simulation with known true values, confirming that the theoretical non-confounding of Proposition 1 works at the sample lengths used here. Four layers support the identification of the anchors independently: theory (Proposition 1), robustness on actual data (this appendix), finite samples (Appendix E), and validation from outside the model (Section 3.5).

C.1 Calibration invariance of the anchor estimates

Among the calibrated values of this paper (Table B.2), those that could bear on the identification of the anchors are the Phillips-curve slopes (κ_p, κ_w) and the policy rule parameters (ϕ_x, ρ_i). This appendix shows, for both the Japanese and the U.S. cases, that moving these over a wide range leaves the anchor shifts Δ unchanged in sign and in order of magnitude.

The sweep moves (κ_p, κ_w) over ranges of 40 and 20 times their reference values, ϕ_x over a range of 20 times, and ρ_i from 0.70 to 0.95, re-estimating the three anchors at each point. What is reported is the range of variation in $\Delta\bar{\pi}$, $\Delta\bar{\pi}^w$, and $\Delta\bar{r}$ over the whole sweep, as a percentage of the reference value.

In every sweep the shifts in all three anchors keep their sign and their order of magnitude (Table C.1).

Table C.1: Calibration invariance of the anchor estimates: shifts Δ over the (κ_p, κ_w) grid (monthly and annual; value at the selected cell, range of variation, and sign stability)

Episode	Anchor	Δ selected	Annual (%)	Min	Max	% var.	Sign
Japan 2022	$\bar{\pi}$	$+0.1827$	$+2.19$	$+0.1598$	$+0.1835$	13.5	Y
Japan 2022	$\bar{\pi}^w$	$+0.1555$	$+1.87$	$+0.1217$	$+0.1561$	23.2	Y
Japan 2022	$\bar{r}$	-0.1559	-1.87	-0.1570	-0.1356	14.5	Y
Japan 2013	$\bar{\pi}$	$+0.0852$	$+1.02$	$+0.0673$	$+0.0852$	23.5	Y
Japan 2013	$\bar{\pi}^w$	$+0.0610$	$+0.73$	$+0.0094$	$+0.0655$	108.5	Y
Japan 2013	$\bar{r}$	-0.0757	-0.91	-0.0852	-0.0746	13.8	Y
U.S. 1979	$\bar{\pi}$	-0.1539	-1.85	-0.1547	-0.1204	24.2	Y
U.S. 1979	$\bar{\pi}^w$	-0.2973	-3.57	-0.2983	-0.2448	18.6	Y
U.S. 1979	$\bar{r}$	$+0.2452$	$+2.94$	$+0.2160$	$+0.2452$	12.6	Y

One point about measurement deserves care here. Measured over the whole grid of the sweep, $\Delta\bar{\pi}^w$ for Japan 2013 moves by 108.5%. That figure, however, includes a calibration the data reject by 32 in log MDD, namely $\kappa_w = 0.10$. Restricted to the grid the data support, defined as cells within 5 in log MDD of the best cell, the variation is only 14.6%, over 6 cells. Measuring all nine combinations of three episodes and three anchors on the same criterion, this is the only one that exceeds 10%; the remaining eight are all at or below 3%. The number of cells inside the identified region differs by episode, however, with 9 for Japan 2022, 6 for Japan 2013, and 2 for the United States, and fewer cells mechanically narrow the range. The very small U.S. figures, at or below 0.5%, include that effect. Calibration invariance should be assessed on the cells the data support. The spread of estimates at rejected calibrations reflects the nature of a rejected specification, not weak identification.

More important is that the quantity on which this paper rests its claims does not depend on the calibration at all. For Japan 2013, $\Delta\bar{\pi}^w - \Delta\bar{\pi}$ is negative in all 35 cells. The sign of the statement that the shift in the wage anchor falls short of the shift in the price anchor therefore holds for every combination of (κ_p, κ_w).

This invariance has an implication for using the paper as a measurement tool in other countries and periods. Even where literature values for $(\kappa_p, \kappa_w, \phi_\pi, \phi_x, \rho_i)$ are unavailable for the country of application, the values used here can be adopted as they stand. The only exogenous values a reader must construct from their own data are two: the energy share γ, from an input-output table, and the energy persistence ρ_{en}, from OLS. Conversely, the estimate of κ_p cannot be read as the slope of that country's Phillips curve (Appendix G). The slopes are weakly identified, and that weak identification is exactly why the identification of the anchors does not depend on how the slopes are set.

Table C.2 shows the range of variation, in percent, of the anchor shifts when (ϕ_x, ρ_i) and the scale of the variances are moved. In every sweep the sign and order of magnitude of the anchors are unchanged, and the only relatively large variation is in $\bar{r}$, which is sensitive to ϕ_x and ρ_i.

Table C.2: Range of variation in the anchor shifts under policy rule and variance sweeps (%)

Episode	Sweep	% var. $\bar{\pi}$	% var. $\bar{\pi}^w$	% var. $\bar{r}$
Japan 2022	ϕ_x	0.8	0.2	8.5
Japan 2022	ρ_i	1.1	0.6	8.1
Japan 2022	σ	0.6	0.4	12.9
Japan 2013	ϕ_x	7.1	3.6	**51.9**
Japan 2013	ρ_i	2.6	7.6	21.2
Japan 2013	σ	1.8	6.4	**32.1**
U.S. 1979	ϕ_x	3.0	1.2	3.2
U.S. 1979	ρ_i	4.2	0.7	4.7
U.S. 1979	σ	1.5	0.3	9.7

C.2 The wage slope κ_w: testing for a regime change and non-identification of its level

Since the wage anchor is the centerpiece of this paper, the treatment of the wage slope κ_w invites more scrutiny than the other propagation parameters. This appendix shows two things about κ_w. First, a regime change in κ_w is tested in the same nested comparison used in the text, and the data do not support it. Second, the level of κ_w is not identified from the data at all, and that does not affect the conclusions.

Take the first point. The nested comparison of the text (Section 3.2 and Table 4) includes a switch in the wage slope κ_w, symmetrically with the switch in the price slope κ_p across MA, MAB, MAC, and MABC. The specification MW, which switches only κ_w, is almost indistinguishable from M0, which switches nothing, in all three episodes, and falls 17 to 42 below MC. Adding a κ_w switch to the best specification MBC, giving MBCW, falls only 0.03 to 0.11 below MBC. A regime change in the wage slope is therefore not supported by the data. This is a finding that has been tested, not the result of assuming κ_w invariant. The asymmetry, in which the anchor and variance axes matter greatly while neither slope axis adds anything, does not depend on the choice of marginal likelihood estimator, as verification by the modified harmonic mean confirms (Appendix D).

Now the second point. Not only is a switch unsupported; the level of κ_w itself is not determined by the data. The reason is that the slope term $\kappa_w x_t$ in the wage Phillips curve carries information only through the output gap x_t, and x_t stays close to zero throughout the POST period, so almost no information comes from that term. Figure C.1 shows the profile likelihood in log MDD as κ_w varies under the MC specification. The profile is highest at the lower end of the grid, $\kappa_w = 0.0005$, and declines monotonically from there. Over the range $\kappa_w = 0.0005$ to 0.0075 the log MDD difference is less than 0.5 and the profile is flat, so κ_w cannot be pinned down informatively by the data.

This non-identification does not affect the conclusions. Varying the fixed value over $\{0.005, 0.010, 0.020, 0.030, 0.050\}$ leaves unchanged the ordering in which the anchor and variance axes contribute substantially and the price slope axis does not (Table C.3). At the selected cell the best specification is MBC. Moving κ_w away from the selected value lets MBCW overtake it, but the ratio for κ_w estimated there is 0.17 to 0.53, well below one. That shows the W axis is not capturing a regime change; it is correcting a misspecified κ_w by pulling it down.

The choice among axes in this paper does not depend on the fixed value of κ_w. Similar flat regions appear in the other episodes: cells within 0.5 of the peak cover $\kappa_w \in [0.0005, 0.0200]$ for Japan 2013, 9 cells, and [0.0005, 0.0030] for the United States, 4 cells. The selected U.S. value of 0.005 sits just outside that flat region, 0.72 from the peak, on the edge that a threshold of 1.0 would include. In other words κ_w is not identified from the data over this range. What matters is that the anchor shifts barely move within the flat region. For Japan 2022 the variation is 0.0001 in $\Delta\bar{\pi}$, 0.0012 in $\Delta\bar{\pi}^w$, and 0.0007 in $\Delta\bar{r}$, all monthly. Estimation that allows κ_w to switch (MBCW) also leaves the anchor shift estimates unchanged: for Japan 2022, $(\Delta\bar{\pi}, \Delta\bar{\pi}^w)$ agrees near $(+0.183, +0.156)$ under MC, MBC, and MBCW alike. Neither the level of κ_w nor a change in it propagates into the anchor estimates. That the flat region lies below the literature value of 0.02 accords with earlier work on the flattening of Japanese wage and service price Phillips curves (Kishaba and Okuda, 2025).

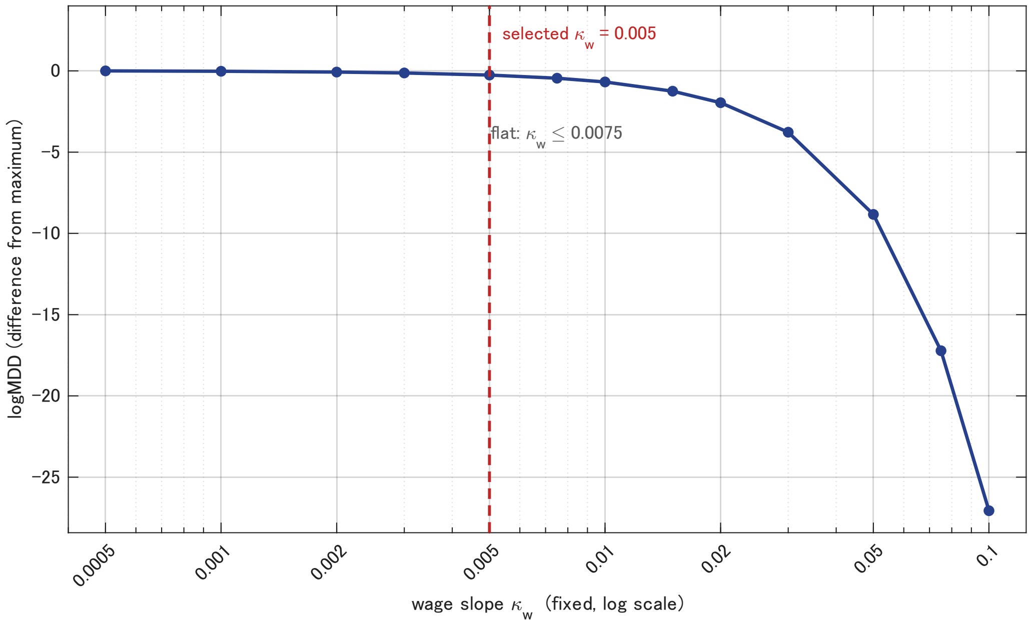

Figure C.1: Profile likelihood for κ_w (Japan 2022, MC specification, $\kappa_p = 0.010$ fixed, log scale). Over the range $\kappa_w \leq 0.0075$ the log MDD difference is less than 0.5 and the profile is flat, so κ_w is not identified from the data there. This is the basis for fixing κ_w. The value selected by the grid, $\kappa_w = 0.005$, lies inside the flat region.

Table C.3: Robustness to the fixed value of κ_w (log MDD relative to MC)

κ_w	M0	MA	MB	MAB	MAC	MBC	MABC	MBCW
0.005 (selected)	−36.46	−34.84	−18.95	−16.85	−0.53	**+12.80**	+12.36	+12.74
0.010	−36.49	−34.90	−19.00	−16.93	−0.54	**+12.78**	+12.34	+12.67
0.020	−36.51	−34.98	−19.11	−17.11	−0.54	**+12.70**	+12.26	+12.64
0.030	−36.50	−35.03	−19.23	−17.29	−0.54	+12.58	+12.14	**+12.74**
0.050	−36.39	−35.02	−19.47	−17.65	−0.54	+12.27	+11.83	**+13.33**

Notes: Each row gives the log MDD difference relative to MC at the stated κ_w, with $\kappa_p = 0.010$ fixed. Bold marks the best specification in each row. At the selected cell, $\kappa_w = 0.005$, the best specification is MBC. Moving κ_w away from the selected value lets MBCW overtake it, but this is not the W axis detecting a regime change. It is a correction that pulls a misspecified κ_w down in the POST period, and the estimated ratio, 0.17 to 0.53, is well below one. At every value of κ_w the ordering is unchanged: the anchor and variance axes contribute substantially and the price slope axis does not.

C.3 Calibration of the import price block

Import prices are exogenous AR(1), and the pass-through of exchange rate depreciation is folded into the shock ε^{im}; no separate exchange rate observable is used. As with the energy persistence ρ_{en}, the parameters ρ_{im} and $\sigma_{\varepsilon^{im}}$ are calibrated by OLS, on the same estimation window as in the text and excluding the months of tax rate changes and the COVID period. The calibrated values are $\rho_{im} = 0.498$ and $\sigma_{\varepsilon^{im}} = 2.217$. The size of the second reflects how large movements in exchange rates and import prices are, and shows that this check does not shut the import channel arbitrarily. What is at issue is the next step, the pass-through χ from import prices to underlying inflation. The prior on χ is the same weakly informative $N(0, 0.10)$ used for λ_w in Section 3.3.

C.4 Identification of the steady-state real rate $\bar{r}$ and the effective lower bound

Of the three anchors, $\bar{r}$ depends most heavily on the interest rate series. Japan's policy rate barely moves in either the PRE or the POST period, staying near the effective lower bound (ELB), which raises the concern that $\bar{r}$ may be weakly identified. This appendix addresses that concern. The point is that the weak identification of $\bar{r}$ is not a defect of the model but a property of the Japanese data.

The exercise re-estimates the model with the standard deviation of the monetary policy shock, σ_{ε^i}, set to 33 times its estimated value. Making σ_{ε^i} extremely large lets the policy shock absorb almost all observed variation in the interest rate, so the rate contributes almost no information about the dynamics of the model. Only its level remains. How far the estimate of each anchor moves under this operation is then compared between Japan and the United States. Both countries use a common structural core, a common calibration, and a common estimation protocol (Appendix C). Only the data differ.

In the United States, $\Delta\bar{\pi}$, $\Delta\bar{\pi}^w$, and $\Delta\bar{r}$ all barely move. The estimates are robust, $\bar{r}$ included, even after the dynamic information in the interest rate is destroyed. In Japan, $\Delta\bar{\pi}$ and $\Delta\bar{\pi}^w$ are robust while $\Delta\bar{r}$ moves relatively more. Under the same model and the same method, the only thing that produces this difference is the presence or absence of the ELB.

There are two implications. First, identification of the price and wage anchors does not depend on the interest rate, because both are identified from the means of prices and wages (Proposition 1). The central claims of this paper concern those anchors and are unaffected by the weak identification of $\bar{r}$. Second, the identification of $\bar{r}$ becomes weak in Japan not because the model is wrong but because the Japanese policy rate is pinned at the lower bound and carries no information. The same model identifies $\bar{r}$ robustly in the United States, where the rate is not constrained by the bound. This is a limitation of the data, not of the model.

C.5 The energy cost term: the level of the relative price and identification of the anchors

Cost minimization under the production function $Y_t = A_t N_t^{1-\gamma} E_t^{\gamma}$ gives nominal marginal cost $MC_t = A_t^{-1}\{W_t/(1-\gamma)\}^{1-\gamma}(P_t^E/\gamma)^{\gamma}$. Expressing this in real terms and log-linearizing,

$$\widehat{mc}_t = (1-\gamma)(\widehat{w_t - p_t}) + \gamma\,\hat{s}_t - \hat{a}_t, \qquad s_t \equiv p_t^E - p_t \tag{22}$$

so what enters the price Phillips curve is the level of the relative price, $\hat{s}_t$, with coefficient $\zeta = \gamma\kappa_p$. The form of the coefficient matches the one in the text, but the argument is a flow there. The two are linked by $s_t = s_{t-1} + \pi_t^{en} - \pi_t^p$, and in general they are not proportional.

The difference cannot be ignored. The relative price of energy in Japan rose after 2022 and stayed high, so using the level leaves a large cost wedge in place throughout the POST period. That is exactly where the price anchor sits. And γ is only a calibrated value taken from an input-output table. This raises the concern that with a smaller γ, energy alone might account for the shift in the anchors.

The level term was therefore estimated explicitly. The series s_t enters the forward solution as an AR(1) with ρ_s calibrated by OLS, and the coefficient ζ_s is estimated rather than calibrated, under a log-normal prior centered on the value implied by the input-output table with a standard deviation of 1.5. Since the concern is precisely about the calibration of γ, this is where the data must speak. The series s_t is not added to the observables. Because $s_t = \sum(\pi^{en} - \pi^p)$ is an exact function of existing observables, adding it would make the likelihood singular. It is treated as a known exogenous path, and the number of observables stays at five.

The results appear in Table C.4. The coefficient ζ_s excludes zero in all three episodes and improves the marginal likelihood, so it cannot be dismissed as ineffective. And for Japan 2022 the shift in the price anchor shrinks by about 37%, from +0.183 to +0.115; for Japan 2013 it shrinks by about 26%, from +0.085 to +0.063. Even after shrinking, however, the 90% intervals exclude zero, at $[+0.071, +0.158]$ and $[+0.040, +0.086]$ respectively (Table C.5). The level term does not destroy identification of the anchors. It absorbs part of the shift.

But this is not structural cost pass-through. The series s_t is close to a unit root, with $\rho_s \approx 0.99$, and is low in the PRE period and high in the POST period. Within the estimation window it behaves in effect as a regime dummy, and takes over part of the anchor's job mechanically. Four pieces of evidence show this.

First, the implied elasticity does not behave like a structural parameter. It ranges from 0.0062 to 0.0140 across the three episodes. Even the two Japanese episodes, which share a country and an input-output table, differ by nearly a factor of two in point estimate, 0.0117 against 0.0062, and their intervals barely overlap, at $[0.0066, 0.0169]$ and $[0.0031, 0.0094]$. Neither Japanese interval contains the value implied by the input-output table, 0.0263. An elasticity derived from a production function does not behave this way. The same conclusion follows from fixing ζ_s at a constant multiple of the value implied by the input-output table, rather than estimating it, and comparing marginal likelihoods. Imposing the implied value itself, a multiple of one, makes the marginal likelihood worse than the baseline specification that drops the level term altogether: -289.89 against -288.06 for Japan 2022, -358.86 against -320.10 for Japan 2013, and -576.26

against -556.02 for the United States. Fit is best at one quarter to one half of the implied value, at -280.92, -315.07, and -535.42 respectively, and the optimum appears at the same place in all three episodes. If the level term captured structural cost pass-through, imposing the elasticity from the input-output table should improve fit. In fact imposing it does worse than dropping the term.

Second, demeaning s_t within each regime restores both the anchors and the coordination to their baseline estimates. That operation removes the difference in levels across regimes entirely, and identifies ζ_s from within-regime covariation alone. For Japan 2022, $\Delta\bar{\pi}$ returns to $+0.182$, or 99.6% of the baseline, and the coordination measure $\Delta\bar{\pi}^w - \Delta\bar{\pi}$ moves $-0.027 \to +0.059 \to -0.025$, recovering the baseline (Table C.5). Japan 2013 recovers in the same way, $+0.085 \to +0.063 \to +0.086$. The U.S. coordination measure loses significance under full-sample demeaning but recovers under within-regime demeaning.

Third, and most important, information across regimes contributes almost nothing to fit. Comparing marginal likelihoods in Table C.4, removing the level difference across regimes entirely leaves the improvement almost unchanged, with a gap of 0.2 to 2.0. For Japan 2013 within-regime demeaning fits better, at -316.68 against -317.07. For Japan 2022, in exchange for an improvement in fit of at most 1.99, 37% of the shift in the price anchor was being removed. What the level term buys is not explanatory power but a reallocation of the anchor.

Fourth, it has no explanatory power across episodes. The PRE-to-POST gap in the relative price, $\Delta\bar{s}$, is almost identical across the three episodes, at $+20.8$, $+21.9$, and $+23.0$, yet $\Delta\bar{\pi}$ is $+0.085$, $+0.183$, and -0.154, reversing in sign. In the two Japanese episodes, where $\Delta\bar{s}$ is nearly equal, the behavior of the wage anchor is the opposite: it does not follow in 2013 and moves one for one in 2022. Energy enters only the price Phillips curve, so it cannot produce that difference.

At the same time, short-run pass-through is real once it is identified correctly. Under the within-regime demeaned specification ζ_s remains significant, and the elasticity implied for Japan 2022, 0.0212 with a 90% interval of $[0.0087, 0.0340]$, contains the input-output value of 0.0263. Cost pass-through and a permanent shift in regime means are identified as separate phenomena.

One further point. Within the price equation, the regime-mean effect of the level term, $\zeta_s \bar{s}_r/(1-\beta\rho_s)$, and the price anchor $\bar{\pi}_r$ are observationally equivalent, and only their sum is identified from across-regime information. What the test in this appendix answers is a different question: when the relative price drifted within a regime, did inflation follow it? The four points above show that the conclusion is reached without needing a decomposition across regimes.

Table C.4: Estimating the level term in the relative price of energy

Episode	Demeaning	ζ_s ($\times 10^{-3}$, 90%)	Implied γ (90%)	I-O γ	MHM
Japan 2022	Full sample	0.117 $[0.066, 0.169]$	0.0117 $[0.0066, 0.0169]$	0.0263	-282.76
	Within regime	0.214 $[0.083, 0.341]$	0.0214 $[0.0083, 0.0341]$	0.0263	-284.75
Japan 2013	Full sample	0.031 $[0.015, 0.047]$	0.0062 $[0.0031, 0.0094]$	0.0263	-317.07
	Within regime	0.042 $[0.021, 0.063]$	0.0084 $[0.0042, 0.0127]$	0.0263	-316.68
U.S. 1979	Full sample	0.070 $[0.053, 0.087]$	0.0140 $[0.0106, 0.0175]$	0.0339	-537.08
	Within regime	0.069 $[0.052, 0.086]$	0.0138 $[0.0104, 0.0171]$	0.0339	-537.27

Notes: The MHM of the baseline specification is -288.04 for Japan 2022, -320.09 for Japan 2013, and -556.00 for the United States. The implied $\gamma = \zeta_s/\kappa_p$. Full sample demeans s_t over the whole estimation window; within regime demeans it inside each regime, which removes the difference in levels across regimes entirely. The improvement in fit is almost the same either way: $+5.28$ and $+3.29$ for Japan 2022, $+3.02$ and $+3.41$ for Japan 2013, and $+18.92$ and $+18.73$ for the United States. Differences in level across regimes therefore contribute almost nothing to fit.

Table C.5: Effect of the level term on the anchors and on coordination

Episode		Baseline	Level term (full sample)	Level term (within regime)
Japan 2022	$\Delta\bar{\pi}$	$+0.183^*$	$+0.115^*$	$+0.182^*$
	$\Delta\bar{\pi}^w - \Delta\bar{\pi}$	-0.027	$+0.059$	-0.025
Japan 2013	$\Delta\bar{\pi}$	$+0.085^*$	$+0.063^*$	$+0.086^*$
	$\Delta\bar{\pi}^w - \Delta\bar{\pi}$	-0.024	$+0.010$	-0.025
U.S. 1979	$\Delta\bar{\pi}$	-0.154^*	-0.219^*	-0.201^*
	$\Delta\bar{\pi}^w - \Delta\bar{\pi}$	-0.142^*	-0.058	-0.082^*

Notes: Monthly. An asterisk marks a 90% interval that excludes zero. Under full-sample demeaning the price anchor for Japan 2022 shrinks by 37% but keeps its significance, and the coordination measure reverses sign and moves outside the baseline 90% interval of $[-0.119, +0.066]$. The U.S. coordination measure loses significance. Under within-regime demeaning all of them return to the baseline estimates.

D. Details of the Model Comparison

D.1 Composition of the variance switch (MBC) and stability of the MCMC

In the nested model comparison of Section 3.2 (Table 4), the best specification was MBC, which switches anchors and variances together. This appendix reports the composition of that variance switch and its robustness to the choice of marginal likelihood estimator.

Table D.1 gives the composition of the variance switch under MBC, as POST over PRE ratios of shock standard deviations at the posterior mode. The widening of variances extends to domestic prices (+75%), energy (+34%), and wages (+23%). Only the monetary policy shock narrows, with a ratio of 0.83. The significance of a change in variances is assessed by model comparison (Table 4) rather than by intervals. That the widening spans domestic prices, energy, and wages while only the monetary policy shock narrows indicates that disturbances in the POST period intensified broadly, on the cost side and in price and wage setting alike, while disturbances originating in monetary policy subsided. This is a recomposition of disturbances, independent of the level shift in the mean, that is, in the anchors.

The nested comparison in the text rests on marginal likelihoods from the Laplace approximation. To check the accuracy of that approximation, the top three specifications, MBC, MABC, and MBCW, were compared against the modified harmonic mean (MHM) computed with the same sampler as the main estimation, RWMH with 100,000 draws on each of two chains (Table D.2). In all nine comparisons the two estimators differ by at most 0.09, and the best specification agrees in all three episodes.

What this confirms is not that MBC wins, but that the Laplace approximation is accurate in this model. The fact supports the reading of the text in two ways. The loss from removing the anchor axis (20.8 to 47.0) and from removing the variance axis (12.1 to 74.8) is 140 to 860 times the approximation error, whose maximum is 0.087, so the main conclusions do not depend on the choice of estimator. On the other side, the change from adding the wage slope (0.03 to 0.11) is of the same order as the approximation error, and the change from adding the price slope (0.07 to 0.54) is at most a few times that error. Both are orders of magnitude away from the loss of removing an axis, which is consistent with the reading that the data do not identify a regime change in the slopes. What verification by MHM strengthens, then, is not the ranking but the asymmetry: the slope axes add nothing while the anchor and variance axes matter greatly, and that asymmetry does not depend on the estimator. Table D.3 lists the diagnostics, acceptance rates, ESS, and $\hat{R}$, for all three episodes.

Table D.1: Composition of the variance switch under MBC (POST/PRE ratios of shock standard deviations, posterior mode)

Shock	POST/PRE ratio
σ_{ε^p} (domestic price)	1.75
σ_{ε^w} (wage)	1.23
σ_{ε^i} (monetary policy)	0.83
$\sigma_{\varepsilon^{en}}$ (energy)	1.34

Notes: Posterior mode. The widening of variances covers domestic prices, energy, and wages, and only monetary policy narrows. Significance is assessed by model comparison (Table 4).

Table D.2: Comparison of marginal likelihood estimators: Laplace approximation and modified harmonic mean (RWMH, 100,000 draws on each of two chains)

Episode	Specification	Laplace	MHM	Difference
Japan 2022	**MBC**	−275.26	−275.22	+0.04
	MABC	−275.70	−275.61	+0.09
	MBCW	−275.32	−275.37	−0.05
Japan 2013	**MBC**	−307.98	−307.99	−0.01
	MABC	−308.53	−308.50	+0.03
	MBCW	−308.09	−308.15	−0.06
U.S. 1979	**MBC**	−481.25	−481.29	−0.04
	MABC	−481.32	−481.38	−0.06
	MBCW	−481.28	−481.33	−0.05

Notes: Acceptance rates are 0.41 to 0.43. In all nine comparisons the absolute difference is at most 0.09, so the Laplace approximation is accurate in this model. The best specification is the same under both estimators in all three episodes (MBC). The mutual differences among MBC, MABC, and MBCW, at 0.03 to 0.54, are of the same order as the error of the estimators, so no ranking among them can be argued at this resolution.

Table D.3: MCMC diagnostics (three episodes, best cell of the POST period)

Episode	MHM	Acceptance	ESS	$\hat{R}_{\max}$	Draws
Japan 2022	−288.04	0.421	5115	1.0002	200,000
Japan 2013	−320.09	0.420	4865	1.0006	200,000
U.S. 1979	−556.00	0.420	4992	1.0001	200,000

E. Verifying Identification: Monte Carlo Experiments

Each episode is estimated on five observables including energy ($\pi^p, \pi^w, \pi^{\text{en}}, i$, and the output gap), with a short POST window: 52 months for 2022, 72 for 2013, and 96 for the United States. Whether the non-confounding of Proposition 1 works in finite samples for this very specification is verified by repeated experiments on the slope and anchor axes, under the calibration of both Japanese episodes and of the U.S. Volcker period. The design was fixed in advance. The data-generating core includes energy as an exogenous AR(1), with $\rho_{\text{en}} = 0.5361$ for Japan 2022, and passes the intermediate input cost term $\zeta\pi^{\text{en}}$ into the price Phillips curve, with $\zeta = \gamma\kappa_p$ and $\gamma = 0.0263$ for Japan. True values of the parameters that do not change are calibrated to the PRE fit of each episode, with $(\sigma_{\varepsilon^x}, \sigma_{\varepsilon^p}, \sigma_{\varepsilon^w}, \sigma_{\varepsilon^i}, \sigma_{\varepsilon^{\text{en}}}) = (2.02, 0.07, 0.27, 0.08, 1.51)$ and anchors ≈ 0 for the 2022 PRE period, and the slopes are fixed at the selected cell of each episode, $(\kappa_p, \kappa_w) = (0.010,\ 0.005)$ for Japan 2022. As in the main estimation, the standard deviation of the demand shock is estimated, with a true value of about 2.0, and the set of estimated parameters is the same as in the main estimation, five standard deviations plus three anchors, with κ_p added on the slope axis. The true change in each component is moved over a grid, and the firing rate is measured as the share of replications in which the 90% highest posterior density interval for the difference Δ in that parameter excludes zero, with 200 replications per axis and the same decision rule as in the text. At $\Delta = 0$ this is size; elsewhere it is power; and the firing rate of the test on the other component is cross-axis leakage.

Power reaches one over the range of empirical shifts. The empirical shift for 2022, $\Delta\bar{\pi} = 0.183$, is bracketed by the grid points 0.155 and 0.241, and power is 1.000 at both. The empirical shift for 2013, $\Delta\bar{\pi} = 0.085$, sits near the grid point 0.080, where power is 1.000, and at the U.S. empirical shift of $|\Delta\bar{\pi}| = 0.154$ power is 1.000 as well (Table E.1). The shift at which 80% power is reached is $\Delta\bar{\pi} \approx 0.054$ for 2022, 0.040 for 2013, and 0.072 for the United States, all far below the empirical shifts. The anchor shifts detected in this paper lie in the region where power is saturated.

When a true change is imposed on one axis, the false firing of the other axis, the cross-axis leakage, stays at or below 0.125 in both directions and in all three episodes. Leakage onto the slope side when a true change is imposed on the anchor axis is especially small, at 0.01 to 0.08. Leakage does not increase systematically as the true Δ grows. A true change on one axis, in other words, does not turn into a false change on the other. This is direct evidence that the attribution of a regime change to an axis (Section 3.2) works cleanly in finite samples, and it shows that the two main competing hypotheses a referee would raise, a change in the propagation slope and a change in stochastic variances, can be separated and evaluated on the specification actually estimated.

Because decisions use 90% intervals, the nominal size is 0.10. The size on the anchor axis, the firing rate at $\Delta = 0$, is 0.070 for 2022, 0.085 for 2013, and 0.085 for the United States. With a binomial standard error of 0.021 at 200 replications, these sit slightly below the nominal 0.10, at a conservative level. Coverage is 0.88 to 0.96, close to the nominal 0.90. The size on the slope axis is 0.030 to 0.055, below nominal, so changes in the slope are not over-detected. In episodes where κ_p is small, however, the slope estimate is biased upward, by +0.004 in 2022, +0.005 in 2013, and +0.011 in the United States, and coverage near a zero shift falls to 0.82, 0.72, and 0.64. The claim that the slopes are invariant therefore rests not on Monte Carlo power but on model comparison, in which adding the slope axis worsens log MDD (Section 3.2).

The results on the slope axis are easier to interpret as multiples rather than in absolute terms. The reference κ_p is 0.010 for Japan 2022 and 0.005 for 2013, so the slope rows of Table E.1 show what multiple of the reference each $\Delta\kappa_p$ represents. Power against a change of twice the reference is low, at 0.23 to 0.33, but it reaches 0.48 to 0.72 at three times and 0.79 to 0.94 at four times. Had the slopes switched by a factor of four, the sample lengths used here would have detected it with high probability. Against that, the POST over PRE slope ratio estimated when a regime change in the slopes is allowed is only 0.985 for Japan 2013 and 1.28 for 2022. The estimates lie far below the detectable region. The finding of Section 3.2 is therefore read accurately not as too small to detect but as estimated not to have changed. Power against small changes of about twice the reference is nonetheless insufficient, and no claim is made that changes of that size can be ruled out.[15]

All three episodes are verified on the same five-observable core including energy and with the same estimator (Table E.1). Variation in the energy series itself differs greatly across episodes, but the contribution of the energy shock in the historical decomposition is small in every episode, about 0.01% at an annual rate (Table 7 and Appendix G). Identification of the secondary movement in variances in the United States rests not on the variance axis of these Monte Carlo experiments but on the nested comparison in the text, in which variances alone (MB) exceed anchors alone (MC) and anchors plus variances (MBC) is best (Section 3.2).

F. Details of the Wage–Price Feedback

This appendix gives the details of the specification behind the two-way test of Section 3.3, together with the posteriors of the individual coefficients. What had to be tested was neither λ_w nor λ_p on its own, but their product once both are opened at the same time. The paper estimates a two-way specification with cross terms in both the price and the wage Phillips curve,

$$\pi_t^p = (1-\beta)\bar{\pi} + \beta\, E_t \pi_{t+1}^p + \kappa_p x_t + \zeta\, \pi_t^{en} + \lambda_w\, \pi_t^w + \varepsilon_t^p, \tag{23}$$

$$\pi_t^w = (1-\beta)\bar{\pi}^w + \beta\, E_t \pi_{t+1}^w + \kappa_w x_t + \lambda_p\, \pi_t^p + \varepsilon_t^w \tag{24}$$

and reports the posterior of $\lambda_w \lambda_p$ together with the largest eigenvalue of the wage–price system. Both coefficients carry the same weakly informative prior, $N(0, 0.10)$.

[15] The decision rule is kept at the 90% interval used in the text. Tightening only the Monte Carlo to a 95% interval would lower size, but it would change the object being verified, namely the validity of the procedure actually used in the text, and is therefore not adopted.

Table E.1: Verifying identification: power, size, and cross-axis leakage (three episodes; simulation with known true values)

Δ	P_{detect}	P_{leak}	cover
(A) Japan 2022 ($n_{\text{PRE}} = 72,\ n_{\text{POST}} = 52$)			
Anchor $\Delta\bar{\pi}$ (leakage onto slope)			
0 (size)	0.070	0.055	0.930
0.020	0.245	0.060	0.930
0.050	0.770	0.040	0.910
0.080	0.980	0.045	0.905
0.120	1.000	0.070	0.895
0.155	1.000	0.080	0.915
0.241	1.000	0.070	0.915
Slope $\Delta\kappa_p$ (leakage onto anchor)			
0 (size)	0.045	0.090	0.820
0.005	0.125	0.080	0.905
0.010	0.330	0.080	0.880
0.015	0.455	0.115	0.905
0.020	0.715	0.085	0.895
0.030	0.935	0.125	0.870
0.050	1.000	0.090	0.880
(B) Japan 2013 ($n_{\text{PRE}} = 156,\ n_{\text{POST}} = 72$)			
Anchor $\Delta\bar{\pi}$ (leakage onto slope)			
0 (size)	0.085	0.070	0.920
0.020	0.430	0.060	0.955
0.050	0.985	0.030	0.925
0.080 (near empirical)	1.000	0.025	0.905
0.120	1.000	0.070	0.950
0.155	1.000	0.065	0.945
0.241	1.000	0.040	0.940
Slope $\Delta\kappa_p$ (leakage onto anchor)			
0 (size)	0.055	0.065	0.715
0.005	0.230	0.060	0.865
0.010	0.475	0.085	0.900
0.015	0.785	0.075	0.920
0.020	0.950	0.095	0.945
0.030	0.995	0.070	0.880
0.050	1.000	0.125	0.870
(C) United States 1979 ($n_{\text{PRE}} = 117,\ n_{\text{POST}} = 96$)			
Anchor $\Delta\bar{\pi}$ (leakage onto slope)			
0 (size)	0.085	0.020	0.925
0.050	0.540	0.010	0.910
0.080	0.900	0.025	0.875
0.120	1.000	0.030	0.910
0.155 (empirical)	1.000	0.050	0.880
Slope $\Delta\kappa_p$ (leakage onto anchor)			
0 (size)	0.030	0.100	0.635
0.010	0.070	0.100	0.855
0.020	0.220	0.095	0.910
0.030	0.420	0.080	0.905
0.050	0.810	0.110	0.885

Notes: The five-observable core including energy is used. A decision fires when the 90% HPD interval for the difference Δ on that axis excludes zero. $\Delta = 0$ gives size; 200 replications per axis. P_{leak} is the false firing rate of the test on the other axis, and cover is coverage on the axis in question.

Placing a lag on the cross term, as in $\lambda_w \pi^w_{t-L}$, looks natural but measures something other than contemporaneous interaction in this model. A purely forward-looking Phillips curve implies a series whose inflation is almost serially uncorrelated. Inflation in actual data, however, is persistent. Opening a lagged cross term against that gap makes the term serve as a proxy for the omitted lagged own term, absorbing the persistence of the other variable. The coefficient λ then becomes large and significant and the likelihood improves substantially, but that is evidence of persistence rather than of a spiral. The Japanese data show this directly. With a lag, the cross coefficient becomes large and significant and the likelihood improves greatly, but the anchor estimates are contaminated as well. Estimating the same data under the contemporaneous specification gives a small cross coefficient, a slight improvement in likelihood, and unchanged anchors. That the anchors move at all is itself a sign that the lagged estimator is measuring something else, namely persistence. The contemporaneous two-way specification is therefore the standard design here.

Going two-way raises the question of determinacy, the uniqueness of the rational expectations equilibrium. Determinacy is not settled by the loop gain $\lambda_w \lambda_p$ alone. It depends on (λ_w, λ_p) and on the eigenvalue structure of the whole system, through the Blanchard–Kahn conditions. The simple proposition that a positive loop means indeterminacy therefore does not hold. Determinate cells in fact spread over the region near the origin and over the region where λ_w and λ_p have opposite signs. The result obtained is narrower but accurate. Scanning the (λ_w, λ_p) plane on a grid of $13 \times 13 = 169$ points, the number of determinate cells is 92/169 for Japan 2022, 88/169 for Japan 2013, and 92/169 for the United States in 1979. Within those, the region in which λ_w and λ_p are both sufficiently positive, that is, strong self-reinforcing feedback, lies outside the determinate region in every episode. The parameter region that would permit strong positive feedback is thus incompatible with the determinacy condition that the observed stable nominal dynamics require of the model. This is a property under the common calibration used here and over the range of slopes considered.

There is no single threshold, however. Determinacy is not a function of the loop gain; it depends separately on λ_w and on λ_p. For instance $(\lambda_w, \lambda_p) = (+0.05, +0.05)$ and $(-0.05, -0.05)$ give the same loop gain of $+0.0025$, yet the first is determinate and the second is not. No threshold is therefore reported. Instead the determinate region is shown as a map on the (λ_w, λ_p) plane (Appendix F.2). What this paper tests is whether a spiral is operating in the steady-state equilibrium of each regime. The transition between regimes is excluded from estimation, and what happened there is out of scope. Estimation is carried out inside the determinate region, and the posterior of the loop gain is evaluated there.

Table F.1: Two-way contemporaneous test for a wage–price spiral (six regimes)

Episode	Regime	Loop gain ($\times 10^{-3}$, 90%)	Mean	Δlog MDD	Verdict
Japan 2013	PRE	$[-8.59,\ +0.03]$	-2.39	-0.00	Includes zero
	POST	$[-11.79,\ +0.20]$	-3.43	$+0.96$	Includes zero
Japan 2022	PRE	$[-10.91,\ +0.11]$	-3.17	$+0.81$	Includes zero
	POST	$[-12.99,\ +0.04]$	-3.58	$+0.45$	Includes zero
U.S. 1979	PRE	$[-45.01,\ -0.66]$	-16.81	$+2.47$	Significantly negative
	POST	$[-30.52,\ -0.03]$	-9.93	$+2.12$	Significantly negative

Notes: A spiral requires a positive loop gain $\lambda_w\lambda_p$. Intervals are 90%, in units of $\times 10^{-3}$ for the loop gain. Includes zero means the interval contains zero; Negative means it lies entirely below zero.

F.1 Details of the wage–price spiral test

Finally, the posteriors of the individual coefficients λ_w (wage to price) and λ_p (price to wage) are shown (Table F.2). There was a reason for reporting only the loop gain in the text. An individual coefficient does not imply a spiral unless the loop is closed, even if one of the two is significant. Looking at individual values alone invites the misreading that significant pass-through exists.

Table F.2: Wage–price spiral: 90% intervals for the individual coefficients λ_w (wage $\to$ price) and λ_p (price $\to$ wage) and for the loop gain (six regimes)

Episode	Regime	λ_w (90%)	λ_p (90%)	Loop ($\times 10^{-3}$)	Positive?
Japan 2022	PRE	$[-0.012, +0.093]$	$[-0.180, +0.098]$	$[-10.91, +0.11]$	Includes zero
Japan 2022	POST	$[-0.105, +0.076]$	$[-0.167, +0.166]$	$[-12.99, +0.04]$	Includes zero
Japan 2013	PRE	$[-0.067, +0.033]$	$[-0.165, +0.169]$	$[-8.59, +0.03]$	Includes zero
Japan 2013	POST	$[-0.007, +0.095]$	$[-0.183, +0.075]$	$[-11.79, +0.20]$	Includes zero
U.S. 1979	PRE	$[-0.179, +0.038]$	$[-0.016, +0.304]$	$[-45.01, -0.66]$	Negative
U.S. 1979	POST	$[-0.004, +0.189]$	$[-0.203, +0.045]$	$[-30.52, -0.03]$	Negative

Notes: A positive loop gain $\lambda_w\lambda_p$ is the condition for a spiral. All twelve intervals for λ_w and λ_p include zero, so one-way pass-through is not detected even at the level of individual coefficients. Convergence diagnostics are sound, with effective sample sizes for λ of 1,254 to 5,679, $\hat{R} \leq 1.0035$, and a between-chain standard deviation of the modified harmonic mean of at most 0.10. These are stated explicitly because the parameters are suspected of weak identification.

For the POST period of Japan 2022, the coefficient from prices to wages, λ_p, has a 90% interval of $[-0.167, +0.166]$, and the coefficient from wages to prices, λ_w, has one of $[-0.105, +0.076]$. Both include zero. Since not even one-way pass-through is detected, their product, the loop gain, cannot be distinguished from zero either (Table F.1).

This asymmetry matters for policy (Section 4.3). The mutual reinforcement presumed by a virtuous cycle of wages and prices is not detected in either direction. Since not even one-way pass-through is identified, the individual edges that would make up the cycle cannot be confirmed, let alone whether the cycle closes. A rise in wages is not producing further price increases. The POST period is only 52 months, however, so caution is needed before reading non-detection as evidence of absence.

In the United States, both λ_w and λ_p have 90% intervals that include zero in both regimes. An earlier version of this paper used lagged cross terms and reported a significant λ_w in the POST period, at +0.12 to 0.14. That result is withdrawn. Significance disappears under the contemporaneous specification. The earlier significance was a product of the lagged cross term absorbing the persistence of inflation (Section 3.3). U.S. core CPI is more persistent than its Japanese counterpart, with autocorrelation of +0.16 to 0.32 against a model-implied value near zero, so the proxy effect appears more strongly there.

F.2 Map of the determinate region

As stated in Section 3.3, there is no single threshold for determinacy under two-way feedback. Determinacy is not a function of the loop gain $\lambda_w \lambda_p$; it depends separately on λ_w and on λ_p. The pairs $(+0.05, +0.05)$ and $(-0.05, -0.05)$ give the same loop gain, yet the first is determinate and the second is not. A statement of the form "indeterminacy sets in once the loop gain exceeds c" is therefore wrong, however convenient.

This appendix scans the (λ_w, λ_p) plane on a grid of $13 \times 13 = 169$ points and shows, as a map, the uniqueness of equilibrium at each point. The assessment is made under the estimated parameters of each episode.

The number of determinate cells is 92/169 for Japan 2022, 88/169 for Japan 2013, and 92/169 for the United States in 1979 (Figure F.1). The counts are almost the same across the three episodes, and the shapes of the determinate regions are similar. The determinate region spreads around the origin and into the quadrants where λ_w and λ_p have opposite signs.

The decisive fact is this. The region in which λ_w and λ_p are both sufficiently positive, that is, strong self-reinforcing feedback, lies outside the determinate region in all three episodes. Note that it is not positive loop gain in general that is indeterminate. Determinacy is settled by (λ_w, λ_p) and the eigenvalue structure of the whole system, and determinate cells spread near the origin and in the opposite-sign regions. What is indeterminate is the region of strong positive feedback. This is a structural property of the model rather than a result of estimation, and in that sense it does not depend on the calibration or on the data.

The implication is clear. As long as the observed nominal dynamics are generated by a stable, that is, determinate, equilibrium, strong positive self-reinforcing feedback is not admissible in the first place. On top of that, the estimated loop gain cannot be distinguished from zero. A spiral is neither detected in estimation nor, in its strong form, compatible with the observed stability.

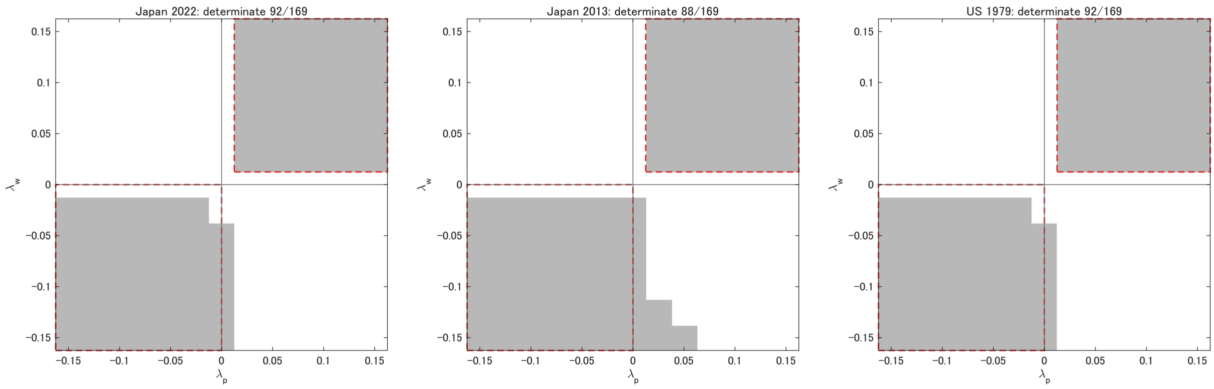

Figure F.1: Determinate region on the (λ_w, λ_p) plane (left: Japan 2022; center: Japan 2013; right: United States 1979). The region of strong positive self-reinforcing feedback lies outside the determinate region in all three episodes.

G. The Volcker Disinflation: Structural Invariance and Robustness

This appendix gives the implementation details of the Volcker comparison of Section 3. The calibration appears in Table G.1 and the observed series and their transformations in Table G.2. It then covers the two-way test for a wage–price spiral, the verification of identification, and the historical decomposition.

Table G.1: Calibration of the closed-economy NK model for the United States (Volcker period)

Parameter	Value	Notes and source
Discount factor β	0.99	Monthly
Elasticity of intertemporal substitution σ	1.00	Log utility
Interest rate smoothing ρ_i	0.9244	$0.79^{1/3}$ (monthly conversion of the quarterly 0.79 in Clarida et al., 2000)
Inflation response ϕ_π	1.50	Taylor (1993), fixed to secure determinacy (see footnote in the text)
Gap response ϕ_x	0.50	Taylor (1993)
Energy input share γ	0.0339	BEA 1977 benchmark input-output table, 366 sectors: petroleum refining, electric power, and gas
Energy price persistence ρ_{en}	0.5797	OLS AR(1) on π^{en} over the estimation window
Price slope κ_p	0.005	Selected by grid search (reference value from Hazell et al., 2022; grid 0.005–0.20)
Wage slope κ_w	0.005	Selected by grid search (reference value from Druant et al., 2009; grid 0.005–0.10)

Table G.2: Observed series and transformations (Volcker period, FRED data)

Model variable	FRED series	Description	Transformation
π^p	CPILFENS	Price inflation: core CPI, all items less food and energy, NSA	X-13 seasonally adjusted, 100× log difference
π^w	AHETPI	Wage inflation: average hourly earnings, production and nonsupervisory workers, SA	100× log difference
π^{en}	PPIENG	Energy inflation: PPI fuels and related products and power, 1982=100, NSA	X-13 seasonally adjusted, 100× log difference
i	FEDFUNDS	Policy rate: effective federal funds rate	Annual rate/12
x	DPCERA3M086SBEA	Output gap: real personal consumption expenditures, chained quantity index, SA	Log detrended with an HP filter ($\lambda = 129600$)

Notes: see the corresponding discussion in the text.

The same three-axis nested comparison used for Japan 2022 (Section 3.2) is applied unchanged to the United States in 1979 (Table 4). The result differs from Japan 2022. Switching the anchors (MC) exceeds M0 and MA by +42 and +41 and is necessary, but switching variances alone (MB) exceeds MC by +27.81. In the United States, then, the explanatory power of the variance axis exceeds that of the anchor axis.

This is not a weakness of the framework. It is a check on the external validity of the tool. The POST window of the Volcker period, 1983 to 1990, coincides with the Great Moderation. That shock variances of macro variables narrowed over that period is a fact established independently of this paper (Justiniano and Primiceri, 2008; Sims and Zha, 2006; Cogley and Sargent, 2005). The three-axis decomposition detected that narrowing. It detects variances in an episode where variances actually moved, and does not detect them in an episode where they did not, namely Japan 2022. The framework is not a device that looks only at anchors.

The same logic applies to Japan 2013 (Section 3.4). There too both the anchor axis and the variance axis are needed, and removing either from MBC lowers the marginal likelihood substantially, by 20.8 and 12.1 respectively. The variance axis exceeds the anchor axis on its own in only one of the three episodes, the United States in 1979, by +27.81. The POST period of Japan 2022 lies in the middle of a global supply shock, so there is no reason for variances to shrink. Even so, the specification that switches variances alone falls 18.95 short of the one that switches anchors alone. The central claim of this paper, that shocks of the same size came to occur around a different level, takes its purest form in this contrast.

The anchors did move, however. All three U.S. anchors shift downward significantly (Table 2). Variances moved as well, and simply have greater explanatory power. Identification of the anchors is not contaminated by the presence of the variance axis. Proposition 1 guarantees non-confounding between the anchors and the propagation block, and the Monte Carlo experiments confirm it in finite samples (Appendix E). Indeed, adding the price slope to the best specification moves fit by only −0.07 (Table 5). The axis that carries the change differs, but a switch in the price or wage slopes is not required.

Section 3.2 showed that the U.S. regime change loads on the anchors and the variances rather than on the slopes. What remains is the wage–price spiral, which is often said to characterize the persistence of inflation in the 1970s.

This section runs the two-way spiral test for the United States. The specification, the priors, and the test statistic are exactly those of Section 3.3 for Japan: contemporaneous cross terms λ_w and λ_p, with weakly informative $N(0, 0.10)$ priors, are placed in both the price and the wage Phillips curve, and the posterior of the loop gain $\lambda_w \lambda_p$ is reported. Matching the specification across the two countries shows that the conclusion about whether a spiral exists is not a product of the specification chosen. The results appear in Table F.1 in the text. The U.S. loop gain is significantly negative in both the PRE and the POST period, which is self-damping rather than self-reinforcing, and the individual coefficients λ_w and λ_p include zero in both regimes (Appendix F). An earlier version of this paper placed a lag on the cross terms, but under a purely forward-looking model a lagged cross term absorbs the persistence of the other variable and measures something other than contemporaneous interaction (Section 3.3). Since U.S. core CPI is highly persistent, that effect is strong there, and the earlier results cannot be interpreted as evidence of a spiral. This section uses the contemporaneous specification throughout.

Verification of identification for the U.S. core, by simulation with known true values, is reported in Panel C of Table E.1, unified with the two Japanese episodes. On both the slope and the anchor axis, the power at the empirical shift (for the anchor, 1.000 at $|\Delta\bar{\pi}| \approx 0.155$), the size, and the cross-axis leakage take the same form as in Japan, and a true change on one axis does not turn into a false change on the other. Identification of the secondary movement in variances in the United States is carried by the nested comparison in the text (Section 3.2).

The U.S. anchor shift is downward, and the downward shift in the wage anchor, −3.55% at an annual rate, is about twice that in the price anchor, −1.84% (Table G.3 and Figure G.1). Almost all of the underlying disinflation, 94.1% for prices and 96.7% for wages, is attributed to the ANCHOR component. In contrast with Japan 2022, where prices and wages moved one for one, wages adjusted downward by more than prices in the United States. This is the counterpart, in the decomposition, of the fact that the coordinated nominal adjustment of the paper's title does not hold for the United States.

These results complement, structurally, the view that discusses inflation in 2021 to 2023 in terms of wage–price dynamics and attributes the persistence of the 1970s to wage indexation and to expectations that were not firmly anchored (Bernanke and Blanchard, 2023). What is detected here is not a self-reinforcing spiral but movement in the steady-state level of inflation expectations, that is, in the anchors.

Table G.3: Historical decomposition: POST minus PRE monthly difference in contributions (%/month, United States 1979; annualize by ×12, so ANCHOR is −1.84% for prices and −3.55% for wages)

Component	π^p	π^w
ANCHOR	−0.154	−0.296
ε^x	−0.003	−0.003
ε^p	−0.008	0.000
ε^w	0.000	−0.009
ε^i	+0.002	+0.002
ε^{en}	0.000	0.000
init	0.000	0.000
Residual	0.000	0.000
Total	−0.163	−0.306
ANCHOR share	94.1%	96.7%

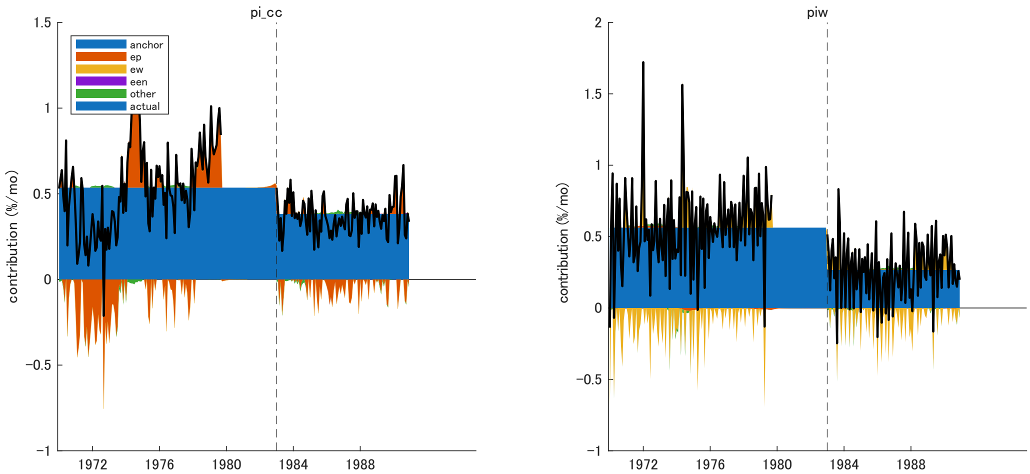

Figure G.1: United States 1979: historical decomposition of π^p (left) and π^w (right)

H. External Validation and Monthly Detail

H.1 The 2013 episode in detail

The PRE period runs from January 2000 to December 2012, 156 months, and the POST period from January 2014 to December 2019, 72 months, with January to December 2013 excluded as a transition window. The ground for exclusion is the same as for the COVID period in Section 3 (Section 2.9). The model, the observation equations, the calibration, the priors, and the estimation procedure are all unchanged from Sections 2 and 2.9. Only the sample window and the break date differ.

In the three-axis nested comparison (Table 4), the anchors-only specification (MC) exceeds the variances-only specification (MB) by +8.72, and anchors plus variances (MBC) is best. Here too both axes are needed: removing the anchor axis from MBC lowers the marginal likelihood by 20.8, and removing the variance axis lowers it by 12.1. That the framework detects variances as dominant when variances move is shown by the United States in 1979, where variances alone

(MB) exceed anchors alone (MC) by +27.81. The criticism that anchors were selected for Japan 2022 because the framework can see nothing but anchors therefore does not hold. As in the other two episodes, adding either the price or the wage slope to the best specification fails to improve the marginal likelihood, by −0.54 and −0.11 respectively.

The POST period of this episode, 2014 to 2019, is the same window as the PRE period of Japan 2022. The two are estimated independently. The resulting anchors are +0.045 against +0.044 for prices, and +0.035 against +0.039 for wages, agreeing to roughly the second decimal place. Recovering the same level from a different sample, a different break date, and a separate estimation shows that the regime split used here is not an artifact of estimation.

That the claims of this paper are complete within each episode is a matter of methodological caution, and also a choice forced by circumstance. Three points are worth stating explicitly.

First, evaluating the posterior of $\Delta\bar{\pi}^w - \Delta\bar{\pi}$ for 2013 gives a posterior mean of −0.024 and a 90% interval of $[-0.089, +0.043]$ that includes zero, with $P(\Delta\bar{\pi}^w < \Delta\bar{\pi}) = 0.727$. The difference itself, the claim that wages did not follow prices, is therefore not significant. What can be claimed is that the shift in the wage anchor cannot be distinguished from zero.

Second, the difference between 2013 and 2022 cannot be tested directly across episodes. The two episodes are independent samples, and a posterior for the difference can be computed by pairing draws. The difference in the shift of the price anchor, 2022 minus 2013, has a posterior mean of +0.098 and a 90% interval of $[+0.059, +0.137]$ that excludes zero. For the price anchor, then, the movement in 2022 can be said to exceed that in 2013. But the difference in the shift of the wage anchor has a posterior mean of +0.095 and a 90% interval of $[-0.013, +0.202]$ that includes zero, with $P(> 0) = 0.926$. The POST periods of 2013, 72 months, and of 2022, 52 months, are both short, and on the wage side there is not enough power. The sentence "wages did not move in 2013 and did move in 2022" therefore cannot be read as a test of the difference across episodes. What can be placed side by side are facts that are complete within each episode. The distinction is essential if a difference in significance is not to be confused with a significant difference.

Third, the framework of this paper contains no mechanism that generates the one-for-one simultaneous shift of 2022. Genuine coordination through the *shuntō* mechanism, both anchors responding independently to a common signal, mechanical indexation, threshold effects, and a tight labor market are all observationally equivalent in aggregate data. Distinguishing among them requires micro data (Section 5).

The same historical decomposition is provided for 2013 (Table H.3 in the appendix). The shift in the price anchor, +1.02% at an annual rate, accounts for almost all of the rise in the underlying trend, 99.0%. The shift in the wage anchor, +0.74% at an annual rate, is relatively small, with a share of 90.5% and a wage total of +0.82% at an annual rate. This contrasts with 2022, when prices and wages moved one for one.

H.2 Supplement on the validity of the trend

This appendix confirms that the deterministic anchor component attributed by de-anchoring is a level shift in trend inflation observed independently, rather than an artifact of parameterizing the model in terms of anchors. The anchors of this paper are compared against reduced-form trend extraction that assumes no structure. This is a check on external validity rather than a competition on goodness of fit, and because the model classes differ, log MDD is not compared directly.

The main benchmark is Stock and Watson (2007), hereafter UCSV, the standard in the trend inflation literature. UCSV is an unobserved components model, $\pi_t = \tau_t + \sigma_{\varepsilon,t}\,\varepsilon_t$ and $\tau_t = \tau_{t-1} + \sigma_{\eta,t}\,\eta_t$, with stochastic volatility in both the trend and the transitory component. It assumes no discrete regimes and extracts a continuous stochastic trend τ_t. UCSV is fitted to the same seasonally adjusted series used here, π^p and π^w, and the trend means over the PRE and POST windows are compared with the structural anchors (Figure H.1).

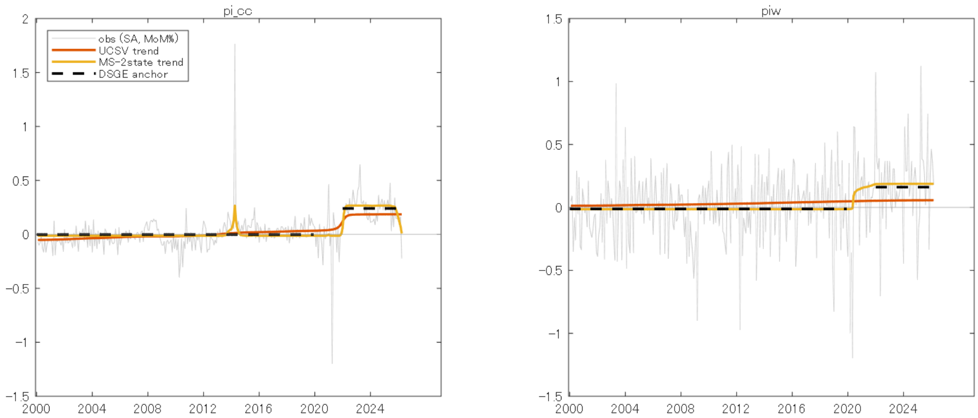

Figure H.1: Reduced-form trends (UCSV and a two-state Markov-switching trend) against the DSGE anchors (month-on-month, %). Grey: seasonally adjusted observations; blue: UCSV trend; orange: two-state Markov-switching trend; black dashed: DSGE anchors for PRE and POST.

Agreement on the price side is clear. The UCSV trend for core CPI shifts up from a monthly +0.046 in the PRE period to +0.230 in the POST period, a shift of +0.184, or +2.21% at an annual rate. It agrees with the shift in the DSGE price anchor, +0.044 to +0.227, a shift of +0.183 or +2.20% annually, in sign, in timing, and in magnitude, with UCSV at about 100% of the DSGE value and POST levels of +0.230 against +0.227. Even a continuous trend extraction that presumes no discrete break shows a level shift of the same direction and size at the start of 2022. This externally supports the discrete two-regime anchor setting imposed here as consistent with the data, and confirms that the anchor component is a genuine shift in trend inflation.

The wage side is similar. The UCSV trend for nominal wages rises from +0.040 in the PRE period to +0.216 in the POST period, a shift of +0.176 or +2.11% annually, which is about 112% of the shift in the DSGE wage anchor, +0.039 to +0.195, a shift of +0.156 or +1.88% annually. What matters is that UCSV extracts the trend from the single nominal wage series alone, using no structure and no cross-equation restriction. Even so it delivers almost the same level shift as the structural anchor. The shift in the wage anchor is therefore information contained in the wage series itself, not a quantity that appears only once the price anchor and the Fisher restriction are solved jointly.

Continuous trend extraction that uses no structure (UCSV) thus agrees with the DSGE anchor shifts in all six series, three episodes by prices and wages, with residuals of ten to twenty percent of the shift. A discrete two-state Markov-switching trend agrees with the DSGE for prices in every episode, and for wages in the United States (-0.322 against a DSGE value of -0.296), but understates the wage shift in Japan alone ($+0.040$ against $+0.156$ in 2022). The reason is that Japanese wages rise discontinuously, concentrated in the *shuntō*, so a two-state discrete approximation absorbs the transitory spike into the high state and fails to capture a sustained rise at an intermediate level. That limitation reinforces the case for using UCSV, which assumes no regimes, as the main benchmark. Either way the anchor component is not an artifact of parameterization. The approximate reproduction of the mean real wage, which the three anchors do not fit independently, is separate evidence of a different kind (the consistency check in Section 3.4).

Table H.1 summarizes these results as verdicts on the three hypotheses of Section 2.8.

Table H.1: Summary of the hypothesis tests

Hypothesis	Test	Result	Verdict
Candidate A: steeper PC slope	Three-axis nested comparison; calibration invariance	Adding the price slope to the best specification moves fit by only -0.54 to -0.07, and the wage slope by only -0.11 to -0.03	Rejected
Candidate B: wider shock variances	Nested model comparison	Variances alone (MB) fall short of anchors alone (MC), but anchors plus variances (MBC, $+12.80$) is best	Rejected as an account of the level shift; supported as a description of volatility
Candidate C: level shift in anchors	Regime-by-regime estimation; de-anchoring decomposition	$\bar{\pi}$ and $\bar{\pi}^w$ exceed zero; $\sim$100% of the rise in π^p is ANCHOR	Supported (dominant)

Table H.2: Anchor levels for Japan 2022 (PRE to POST, monthly, %)

	PRE	POST	Δ (POST−PRE) [90% interval]	$P(\Delta > 0)$	Annual
Price anchor $\bar{\pi}$	$+0.044$	$+0.227$	$+0.183$ $[+0.150, +0.216]$	1.000	+2.20%
Wage anchor $\bar{\pi}^w$	$+0.039$	$+0.195$	$+0.156$ $[+0.069, +0.244]$	0.998	+1.88%
Steady-state real rate $\bar{r}$	-0.038	-0.196	-0.158 $[-0.278, -0.037]$	0.016	−1.89%
Difference $\Delta\bar{\pi}^w - \Delta\bar{\pi}$			-0.027 $[-0.119, +0.066]$	0.314	—

Notes: see the corresponding discussion in the text.

Table H.3: Historical decomposition: POST minus PRE difference in contributions (annualized %, Japan 2013)

Component	π^p	π^w
ANCHOR	+1.02	+0.74
ε^x	−0.01	−0.02
ε^p	−0.02	0.00
ε^w	0.00	+0.01
ε^i	+0.04	+0.08
$\varepsilon^{\mathrm{en}}$	0.00	0.00
init	0.00	0.00
Residual	0.00	0.00
Total	+1.03	+0.82
ANCHOR share	99.0%	90.5%

The PC slopes are invariant, and the rise in the level of nominal variables after 2022 is attributed to a level shift in the steady-state anchors. Shock variances change at the same time, but variances alone do not reach the anchors (Section 3.2). This is not a change in the propagation structure. It is a switch in the mean nominal level.

The main benchmark in this section is UCSV, a continuous trend that assumes no discrete regimes. A second comparison, a two-state mean-switching trend of the Kaihatsu and Nakajima (2018) type estimated by EM under Markov switching, is also shown in Figure H.1. For prices, π^p moves from a monthly +0.064 to +0.211, a shift of +0.147, close to the DSGE anchor of +0.183. For Japanese wages in 2022, however, it moves only from +0.092 to +0.132, a shift of +0.040, far below the DSGE value of +0.156 and the UCSV value of +0.176. Following the Markov trend for π^w, it jumps temporarily to the high state ($\mu_{\mathrm{hi}} \approx 0.73$) right after the 2022 *shuntō* and then returns to the low state (≈ 0.09), so the sustained POST level (≈ 0.16) falls between the two states and is classified into the low one. This follows from Japanese wage revisions being lumpy and concentrated in the *shuntō*. In the United States, where wages adjust smoothly, the two-state trend agrees with the DSGE (−0.322 against −0.296). Beyond the general problem that a reduced form does not supply the economic content of the two states, a discrete approximation has difficulty capturing the Japanese wage trend. This is a further reason for taking UCSV, which assumes no regimes, as the main benchmark.

H.3 Monthly detail of the anchors for Japan 2022

Section 3 presented the anchor estimates for Japan 2022 within the cross-episode table (Table 2). This appendix gives the monthly levels, shifts, and 90% intervals in detail (Table H.2).

Declaration of generative AI and AI-assisted technologies

The research design, model specification, identification strategy, estimation strategy, and the interpretation of all results are entirely the author's own. The author used Anthropic's Claude to translate the manuscript from the author's Japanese draft into English and to assist with editing for language and readability; and, separately, to assist in writing a substantial portion of the MATLAB/Dynare code implementing the estimation and the auxiliary code for LaTeX preparation, output formatting, and validation. The author directed, reviewed, and validated all AI-assisted code and output, and takes full responsibility for the content of this paper. Replication code is available on request.

References

Guido Ascari and Argia M. Sbordone. The macroeconomics of trend inflation. *Journal of Economic Literature*, 52(3):679–739, 2014.

Laurence Ball and Sandeep Mazumder. A Phillips curve with anchored expectations and short-term unemployment. *Journal of Money, Credit and Banking*, 51(1):111–137, 2019.

Ben S. Bernanke and Olivier J. Blanchard. What caused the U.S. pandemic-era inflation? NBER Working Paper 31417, National Bureau of Economic Research, 2023.

Francesco Bianchi. Regime switches, agents' beliefs, and post-world war II U.S. macroeconomic dynamics. *Review of Economic Studies*, 80(2):463–490, 2013.

Michael Bruno and Jeffrey D. Sachs. *Economics of Worldwide Stagflation*. Harvard University Press, Cambridge, MA, 1985.

Guillermo A. Calvo. Staggered prices in a utility-maximizing framework. *Journal of Monetary Economics*, 12(3):383–398, 1983.

Lawrence J. Christiano, Martin Eichenbaum, and Charles L. Evans. Nominal rigidities and the dynamic effects of a shock to monetary policy. *Journal of Political Economy*, 113(1):1–45, 2005.

Richard Clarida, Jordi Galí, and Mark Gertler. Monetary policy rules and macroeconomic stability: Evidence and some theory. *Quarterly Journal of Economics*, 115(1):147–180, 2000. doi: 10.1162/003355300554692.

Timothy Cogley and Thomas J. Sargent. Drifts and volatilities: Monetary policies and outcomes in the post WWII U.S. *Review of Economic Dynamics*, 8(2):262–302, 2005.

Timothy Cogley and Argia M. Sbordone. Trend inflation, indexation, and inflation persistence in the new keynesian phillips curve. *American Economic Review*, 98(5):2101–2126, 2008.

Timothy Cogley, Giorgio E. Primiceri, and Thomas J. Sargent. Inflation-gap persistence in the US. *American Economic Journal: Macroeconomics*, 2(1):43–69, 2010.

Troy Davig and Eric M. Leeper. Generalizing the taylor principle. *American Economic Review*, 97(3):607–635, 2007.

William M. Davis. Major collective bargaining settlements in private industry in 1988. *Monthly Labor Review*, 112(5):34–43, 1989.

Martine Druant, Silvia Fabiani, Gabor Kézdi, Ana Lamo, Fernando Martins, and Roberto Sabbatini. How are firms' wages and prices linked: Survey evidence in Europe. Working Paper Series 1084, European Central Bank, 2009.

Christopher J. Erceg, Dale W. Henderson, and Andrew T. Levin. Optimal monetary policy with staggered wage and price contracts. *Journal of Monetary Economics*, 46(2):281–313, 2000.

Robert J. Flanagan. Wage concessions and long-term union wage flexibility. *Brookings Papers on Economic Activity*, 1984(1):183–221, 1984.

Shin-ichi Fukuda and Naoto Soma. The evolution of inflation expectations in Japan. *Asian Economic Policy Review*, 20(2):207–217, 2025. doi: 10.1111/aepr.12504. Earlier version: CIRJE Discussion Paper F-1238.

Ichiro Fukunaga, Yoshihiko Hogen, and Yoichi Ueno. Japan's economy and prices over the past 25 years: Past discussions and recent issues. Working Paper Series 24-E-14, Bank of Japan, 2024.

Jordi Galí and Tommaso Monacelli. Monetary policy and exchange rate volatility in a small open economy. *Review of Economic Studies*, 72(3):707–734, 2005.

Jonathon Hazell, Juan Herreño, Emi Nakamura, and Jón Steinsson. The slope of the Phillips curve: Evidence from U.S. states. *Quarterly Journal of Economics*, 137(3):1299–1344, 2022.

John R. Hicks. Economic foundations of wage policy. *The Economic Journal*, 65(259):389–404, 1955.

Yoshihiko Hogen and Ryoichi Okuma. The anchoring of inflation expectations in Japan: A learning-approach perspective. *Japan and the World Economy*, 73:101293, 2025.

Hibiki Ichiue, Takushi Kurozumi, and Takeki Sunakawa. Inflation dynamics and labor market specifications: A Bayesian dynamic stochastic general equilibrium approach for Japan's economy. *Economic Inquiry*, 51(1):273–287, 2013.

International Monetary Fund. Japan: 2026 article IV consultation—staff report. IMF Country Report 26/75, International Monetary Fund, 2026.

Nikolay Iskrev. Local identification in DSGE models. *Journal of Monetary Economics*, 57(2): 189–202, 2010.

Hanyuan Jiang. The bones and shapes of the Phillips curve. arXiv preprint arXiv:2506.14030, 2025. London School of Economics and Political Science.

Alejandro Justiniano and Giorgio E. Primiceri. The time-varying volatility of macroeconomic fluctuations. *American Economic Review*, 98(3):604–641, 2008.

Sohei Kaihatsu and Jouchi Nakajima. Has trend inflation shifted?: An empirical analysis with an equally-spaced regime-switching model. *Economic Analysis and Policy*, 59:69–83, 2018. doi: 10.1016/j.eap.2018.04.003.

Sohei Kaihatsu, Mitsuru Katagiri, and Noriyuki Shiraki. Phillips correlation and price-change distributions under declining trend inflation. *Journal of Money, Credit and Banking*, 55(5): 1271–1305, 2023.

Yui Kishaba and Tatsushi Okuda. The slope of the Phillips curve for service prices in Japan: Regional panel data approach. *Journal of the Japanese and International Economies*, 78:

101388, 2025. doi: 10.1016/j.jjie.2025.101388. Earlier version: Bank of Japan Working Paper Series 23-E-8 (2023).

Ivana Komunjer and Serena Ng. Dynamic identification of dynamic stochastic general equilibrium models. *Econometrica*, 79(6):1995–2032, 2011.

Zheng Liu, Daniel F. Waggoner, and Tao Zha. Sources of macroeconomic fluctuations: A regime-switching DSGE approach. *Quantitative Economics*, 2(2):251–301, 2011.

Roland Meeks and Francesca Monti. Heterogeneous beliefs and the Phillips curve. *Journal of Monetary Economics*, 139:41–54, 2023.

Daniel J. B. Mitchell. Shifting norms in wage determination. *Brookings Papers on Economic Activity*, 1985(2):575–608, 1985.

George L. Perry. Inflation in theory and practice. *Brookings Papers on Economic Activity*, 1980 (1):207–260, 1980.

Giorgio E. Primiceri. Time varying structural vector autoregressions and monetary policy. *Review of Economic Studies*, 72(3):821–852, 2005.

Thomas J. Sargent. The ends of four big inflations. In Robert E. Hall, editor, *Inflation: Causes and Effects*, pages 41–98. University of Chicago Press, Chicago, 1982.

Frank Schorfheide. Learning and monetary policy shifts. *Review of Economic Dynamics*, 8(2): 392–419, 2005.

C. Patrick Scott. Commitment issues: Does the Fed have an inflation incentive to commit? *Studies in Nonlinear Dynamics & Econometrics*, 28(5):767–784, 2024. doi: 10.1515/snde-2022-0034.

Christopher A. Sims and Tao Zha. Were there regime switches in U.S. monetary policy? *American Economic Review*, 96(1):54–81, 2006.

Frank Smets and Rafael Wouters. Shocks and frictions in US business cycles: A Bayesian DSGE approach. *American Economic Review*, 97(3):586–606, 2007.

James H. Stock and Mark W. Watson. Has the business cycle changed and why? *NBER Macroeconomics Annual*, 17:159–218, 2002.

James H. Stock and Mark W. Watson. Why has U.S. inflation become harder to forecast? *Journal of Money, Credit and Banking*, 39(s1):3–33, 2007.

Hajime Takata. Economic activity, prices, and monetary policy in Japan. Speech at a Meeting with Local Leaders in Kyoto, Bank of Japan, February 2026. February 26, 2026.

John B. Taylor. Discretion versus policy rules in practice. *Carnegie-Rochester Conference Series on Public Policy*, 39:195–214, 1993.

James Tobin. Inflation and unemployment. *The American Economic Review*, 62(1):1–18, 1972.

Yoichi Ueno. Linkage between wage and price inflation in Japan. Bank of Japan Working Paper Series 24-E-7, Bank of Japan, 2024.

Kota Watanabe and Tsutomu Watanabe. Why has Japan failed to escape from deflation? *Asian Economic Policy Review*, 13(1):23–41, 2018.

Tsutomu Watanabe. Japan's chronic deflation: Causes and consequences. Working Paper 393, Center on Japanese Economy and Business, Columbia University, 2025.