Σrsa
Economic
Research
Southern
Africa

Trend Inflation and the Costs of Price Dispersion in a Fiscal DSGE Model

Hylton Hollander and Clinton Joel

ERSA working paper 923

December 2025

Trend Inflation and the Costs of Price Dispersion in a Fiscal DSGE Model

Clinton Joel[*] and Hylton Hollander[†]

[*]Email: clintonjoel91@gmail.com. School of Economics, University of Cape Town.
[†]Email: hylton.hollander@uct.ac.za. School of Economics, University of Cape Town.

November 25, 2025

Abstract

Most inflation-targeting frameworks allow for a positive trend inflation rate, yet its appropriate level remains uncertain. The extended deliberation in South Africa to move from a 3–6 percent target band to a 3 percent target (with a 1 percent tolerance band) illustrates this tension. Using a two-agent New Keynesian DSGE model with nominal price and wage rigidities and fiscal dynamics, this paper shows that even moderate trend inflation causes significant resource misallocation through price and wage dispersion, flattens the Phillips curve, raises welfare losses and sacrifice ratios, and delays fiscal stabilisation. Omitting trend inflation in a Taylor-rule framework overstates policy inertia and understates the responsiveness needed to stabilise prices. Anchoring expectations closer to lower bound of the target band improves stability: lowering trend inflation from 6 to 3 percent reduces the sacrifice ratio by 0.67 percentage points.

JEL classification: E17, E31, E52, E58, E62.

Keywords: Trend inflation; Price dispersion; DSGE model; Fiscal–monetary policy; South Africa.

Contents

1 Introduction

Most central banks target a small but positive rate of trend inflation, making it essential to understand the implications of steady-state inflation when nominal rigidities are present. Standard New Keynesian models, typically characterised by Calvo-style price rigidity, often assume a zero-inflation steady state—thereby overlooking an important channel of inefficiency: price dispersion. When firms adjust prices asynchronously, some update mark-ups optimally while others lag, creating relative-price misalignments that raise the aggregate price level and generate welfare losses through inefficient resource allocation.

Building on the New Keynesian framework of Garín et al. (2016) and incorporating the two-household structure of Bhatnagar (2023), this paper extends the model to include fiscal dynamics calibrated to South African data. Drawing on the insights of Ascari and Sbordone (2014), we then introduce positive—and potentially time-varying—trend inflation to examine how it amplifies price dispersion and affects macroeconomic stability, welfare, and monetary-policy trade-offs. Unlike standard analyses, the model integrates fiscal interactions and trend inflation within a unified medium-scale DSGE framework, offering a more realistic view of inflation dynamics under realistic policy constraints.

Our results show that even moderate trend inflation markedly increases resource misallocation costs. Higher trend inflation flattens the Phillips curve, amplifies mark-up variability, and reduces the economy's flexibility to absorb shocks—requiring larger and more prolonged output adjustments to achieve disinflation. Consequently, the sacrifice ratio—the output cost of lowering inflation—rises with trend inflation, underscoring the growing real cost of stabilisation in high-inflation regimes. We find that lowering trend inflation from 6 to 3 percent reduces the sacrifice ratio by 0.67 percentage points. Bayesian estimations on South African data underline the importance of explicitly modelling positive trend inflation, showing that ignoring it overstates interest rate smoothing and leads policymakers to underestimate the required policy tightening.

The remainder of the paper proceeds as follows. Section 2 reviews the literature, focusing on trend inflation in New Keynesian frameworks. Section 3 outlines the DSGE model structure and discusses the critical role of price dispersion. Section 3.3 presents the calibration strategy and Bayesian estimation approach. Section 4 discusses key results concerning the economic effects of price dispersion, long-run output, welfare, and dynamic adjustments to various shocks. In addition, this section studies how trend inflation shapes the short-run steepness of the Phillips curve and the implied sacrifice ratio. Section 5 provides sensitivity analyses and explores the policy implications derived from our estimation results. Section 6 concludes with key insights and avenues for future research.

2 Literature Review

In the standard New Keynesian framework featuring only price stickiness, monetary policy benefits from the so-called 'divine coincidence'[1], whereby stabilising inflation simultaneously stabilises the output gap (Blanchard and Galí, 2007). This property makes inflation targeting particularly appealing as a policy rule. However, in more empirically realistic settings this divine coincidence no longer holds (Blanchard and Galí, 2007; Erceg et al., 2000; Ascari and Sbordone, 2014).[2] In these richer frameworks, inflation targeting alone can lead to suboptimal welfare outcomes, making the task of stabilising the output gap more complex and crucial (Garín et al., 2016).

To address these richer dynamics, we build a model with three key features highlighted in the recent literature: (i) positive trend inflation, (ii) nominal rigidities, and (iii) heterogeneous households. We are particularly interested in how each of these dimensions influences both the cyclical dynamics of inflation and output and the long-run monetary policy trade-offs faced by central banks.

We distinguish Ricardian (optimising) from non-Ricardian (rule-of-thumb) households to capture both macroeconomic transmission and redistributive effects. Allowing for rule-of-thumb households alters the propagation of shocks and welfare evaluation (Mohimont, 2022). This distinction is especially pertinent for South Africa, where transfers cushion liquidity-constrained "poor" households while forward-looking "rich" households internalise future taxes; as Hollander et al. (2024) show, capturing that split is necessary to evaluate macro-stabilisation and redistributive policy together. From a policy angle, Bhatnagar (2023) demonstrate that non-Ricardian households amplify monetary-policy effects on output and inflation, especially through a transfer channel that intensifies when steady-state debt is positive; they further show that contractionary shocks raise consumption inequality. These findings underline the need to model household heterogeneity when evaluating welfare and policy trade-offs.

The interaction of household heterogeneity with nominal rigidities is central to fiscal transmission. In a sticky-price New-Keynesian model, Galí et al. (2007) show that debt-financed government spending can crowd in aggregate consumption when rule-of-thumb households are present—an effect that disappears with purely Ricardian agents. Liquidity-constrained households do not internalise future taxes; the fiscal expansion therefore raises their current disposable income, which they spend immediately. Because prices are sticky, real wages fall only slowly, labour demand (and thus labour income) increases, and the consumption response of non-Ricardians is amplified. Fiscal multipliers are thus sensitive to the share of non-Ricardians and to the design of transfer programmes. Galí et al. (2007) conclude that monetary- and fiscal-policy responses have impor-

[1] An optimal policy that closes the output gap also stabilises inflation; this was termed the 'divine coincidence' by Blanchard and Galí (2007).

[2] For example, wage setting rigidities, trend inflation, or heterogenous households.

tant distributional consequences that must be considered in policy design. Embedding a two-household structure in our model therefore enriches the fiscal block and allows us to assess both aggregate and distributional effects of alternative policy mixes.

A growing body of research highlights the importance of explicitly modelling trend inflation rather than implicitly assuming it away. Traditional monetary frameworks often set steady-state inflation to zero or mitigate the influence of positive trend inflation through indexation mechanisms. However, Ascari and Sbordone (2014) argue that accurately capturing trend inflation is both empirically and theoretically essential. In their framework, trend inflation captures the persistent stance of monetary policy, while short-run fluctuations arise from factors such as price-setting frictions, external shocks, and temporary monetary interventions.

Crucially, they highlight that trend inflation can evolve due to shifting inflation targets, gradual policy-driven adjustments, or learning by economic agents. Carefully distinguishing short-run fluctuations from long-run trends deepens our understanding of inflation persistence, a key element in effective monetary policy design. They also show that purely forward-looking policy rules perform poorly in models with intrinsic persistence. Thus, the assumed trend inflation rate significantly shapes the economy's cyclical dynamics and monetary policy trade-offs (Ascari and Sbordone, 2014).

Building on these insights, our findings emphasise the practical importance of explicitly incorporating positive steady-state inflation into New Keynesian models. While standard theoretical specifications frequently treat zero inflation as the optimal long-run benchmark, Ascari and Sbordone (2014) note that central banks typically pursue price stability defined as a small but positive inflation rate. Assuming zero steady-state inflation omits critical dynamics that arise from the interplay between trend inflation, relative price distortions, and monopolistic competition. Therefore, accounting explicitly for positive trend inflation is essential for accurately capturing these dynamics and for informing effective monetary policy design.

Overall, explicitly modelling positive trend inflation is crucial for capturing important macroeconomic dynamics overlooked in standard frameworks. Higher trend inflation not only increases persistence and volatility of macroeconomic variables but also exacerbates the costs of price dispersion, complicating effective monetary policy management. As the Phillips curve becomes flatter with rising trend inflation, monetary policy effectiveness diminishes, forcing policymakers to navigate more challenging trade-offs between stabilising inflation and output (Ascari and Sbordone, 2014). Understanding these interactions is therefore essential for formulating robust policy responses under realistic inflation conditions.

Recent empirical work strengthens the case for explicitly incorporating trend inflation in policy models, especially in economies with a legacy of high inflation. Jacome et al. (2025) show that monetary policy rules are path-dependent: central banks with histories of high inflation tend to maintain more hawkish stances even after stabilisation and credibility gains. This persistent conservatism reflects a form of "experienced learning," whereby expectations and policy preferences remain influenced by past inflation rather than current targets. Complementing this, Walsh (2009) finds that inflation targeting has markedly improved macroeconomic performance in emerging and developing economies—delivering both lower inflation and greater stability. Together, these findings highlight the importance of anchoring expectations and clarifying long-run targets, reinforcing the need to model trend inflation explicitly in monetary frameworks for countries such as South Africa.

Building on this international evidence, South African work on trend inflation is only beginning to emerge. Horn et al. (2025) follow the approach of Eo et al. (2023), where *trend inflation* is treated primarily as a low-frequency component of inflation that helps to strip out short-run noise. This is conceptually distinct from the strand of the literature we build on, following Ascari and Sbordone (2014), which interprets trend inflation as a time-varying policy target within a New Keynesian framework. Nonetheless, Horn et al. (2025) show that South Africa's trend inflation has remained stubbornly above the midpoint of the former target band and that inflation dispersion across price categories is high, which is consistent with the mechanism we highlight.

Adding to this, Foresto et al. (2025) find that the South-African Phillips curve is non-linear and state-dependent. During the pandemic, inflation fell by less than expected despite a productivity collapse, reflecting downward nominal rigidity. They estimate a regime-switch threshold between 4.3% and 9.3% inflation (midpoint 5.6%) and argue that anchoring trend inflation near 3.4% provides room for routine supply shocks without risking de-anchoring or steeper trade-offs.

Collectively, these studies motivate our model: a medium-scale New Keynesian framework with trend inflation, Calvo price and wage rigidity, and heterogeneous households, calibrated to South African data to reassess welfare costs and policy implications when the divine coincidence no longer holds.

3 The New Keynesian Model

We nest the generalised New-Keynesian trend-inflation framework of Ascari and Sbordone (2014) within the sticky-price structure of Garín et al. (2016) to examine how it alters price dispersion and welfare. Consistent with Galí et al. (2007) and Bhatnagar (2023), the model includes Ricardian and non-Ricardian households to capture redistribution and fiscal-policy channels relevant to South Africa (Hollander et al., 2024).

Throughout the core analysis we keep wages flexible by setting the Calvo parameters to zero ($\phi_w = \zeta_w = 0$), so every nominal friction—and therefore every steady-state misallocation—comes from the price-setting dynamics of firms. This restriction lets us isolate how positive trend inflation influences long-run output, consumption, welfare and the price-dispersion wedge. Section 5 relaxes the assumption by reinstating Calvo wage stickiness to gauge how strongly our results depend on wage-setting frictions. Table 1 summarises the model's building blocks. Full derivations and equilibrium conditions appear in the Technical Appendix.

Table 1: Core features and roles of the Two-Agent New Keynesian Model

Block	Key features	Purpose
Households	Two types: Ricardian (optimising) and non-Ricardian (rule-of-thumb). Ricardians exhibit habit formation and labour disutility.	Capture realistic consumption–saving behaviour and a channel for fiscal transfers affecting welfare.
Labour market	Calvo-style staggered wage setting with optional partial indexation ($\phi_w = \zeta_w = 0$ in the baseline). When active, it generates a wage-dispersion wedge υ_w.	Provide a second nominal rigidity that amplifies misallocation and alters monetary-policy transmission under trend inflation.
Firms	Monopolistically competitive producers set prices à la Calvo with partial indexation to past inflation. A competitive aggregator combines into final good.	Generate a price-dispersion wedge and a time-varying mark-up linked to trend inflation.
Trend inflation	Positive steady-state inflation; short-run inflation fluctuates around the target.	Allows price dispersion wedge to arise in the steady state and determine the Phillips curve slope.
Monetary policy	Taylor-type rule with interest-rate smoothing and responses to deviations of inflation from target and output growth.	Stabilise inflation and output in the short run; in the long run the policy rate converges to the steady state target consistent with $\bar{\pi}$.
Fiscal policy	Distortionary consumption, labour taxes, expenditure and transfers to non-Ricardians respond to deviations of debt and output. The government issues one-period bonds facing a risk premium.	Link debt dynamics to real activity, provide redistribution, and interact with monetary policy through the risk premium and tax–spending rules.
Exogenous shocks	Monetary policy, technology, price mark-up, risk premium and government spending shocks.	Trace how disturbances propagate through the sticky-price, trend-inflation economy, shaping fluctuations and policy trade-offs.

3.1 The role of price dispersion

A central by-product of Calvo price staggering is price dispersion—the fact that otherwise identical intermediate firms charge different prices because they reset at different times. We measure this misalignment by

$$v_t^p = \int_0^1 \left(\frac{P_t(j)}{P_t}\right)^{-\epsilon_p} dj \geq 1, \tag{1}$$

which equals one under perfectly flexible prices and rises with the degree of nominal rigidity or the level of trend inflation. Dispersion matters because it scales down effective productivity in the aggregate production function,

$$Y_t = \frac{A_t N_t^d}{v_t^p}, \tag{2}$$

and thus creates a first-order wedge between potential and actual output. It also enters household welfare through higher marginal costs and (in our model) feeds back into wage decisions, making the output–inflation trade-off steeper when trend inflation is high. Full algebraic expressions for v_t^p and its steady-state solution appear in the Technical Appendix.

In steady state, a positive trend inflation $\bar{\pi}$ prevents v^p from collapsing to unity and pins down both the reset price and dispersion in closed form:

$$p^\# = \left(\frac{1-\phi_p(1+\bar{\pi})^{(\epsilon_p-1)(1-\zeta_p)}}{1-\phi_p}\right)^{\frac{1}{1-\epsilon_p}}, \qquad v^p = \frac{1-\phi_p}{1-\phi_p(1+\bar{\pi})^{\epsilon_p(1-\zeta_p)}}\,(p^\#)^{-\epsilon_p}. \tag{3}$$

Following King and Wolman (1996) and Ascari and Sbordone (2014) we show the steady-state markup decomposes into two interpretable components. The closed-form expression decomposes into a *price-adjustment gap* and a *marginal markup*:

$$\mu = \frac{1}{mc} = \underbrace{\left[\frac{1-\phi_p(1+\bar{\pi})^{\epsilon_p-1}}{1-\phi_p}\right]^{-1/(1-\epsilon_p)}}_{\text{price-adjustment gap}} \underbrace{\left[\frac{\epsilon_p}{\epsilon_p-1}\frac{1-\phi_p\beta(1+\bar{\pi})^{(\epsilon_p-1)(1-\zeta_p)}}{1-\phi_p\beta(1+\bar{\pi})^{\epsilon_p(1-\zeta_p)}}\right]}_{\text{marginal markup}}. \tag{4}$$

The quantitative implications of these markup and dispersion effects for steady-state output and welfare are examined in Section 4.2. In addition, trend inflation affects short-run dynamics through the Generalised New Keynesian Phillips Curve slope $\kappa(\bar{\pi})$ (see Section 4.4) and long-run levels through the dispersion wedge $v^p(\bar{\pi})$ and the associated markup components above.

Although the present paper focuses on price dispersion, an analogous wedge emerges on the labour side once wages are sticky. With Calvo wage frictions otherwise identical households post different real wages, giving rise to

$$v_t^w = \int_0^1 \left(\frac{w_t(l)}{w_t}\right)^{-\epsilon_w(1+\eta)} dl \geq 1, \tag{5}$$

where $w_t(l)$ is household l's wage, w_t the average wage, ϵ_w the elasticity of substitution across labour types, and η the Frisch elasticity of labour supply. This wage-dispersion term creates a second resource wedge,

$$N_t^d = \frac{N_t}{v_t^w}, \tag{6}$$

so that higher v_t^w reduces effective labour input just as price dispersion lowers effective productivity. While we set $\phi_w = 0$ in the baseline to isolate price effects, Section 5 reinstates wage stickiness and shows how the two wedges interact under positive trend inflation. For completeness, we show the steady state optimal reset wage relative to the real wage and how incorporating positive trend inflation $\bar{\pi}$ prevents wage dispersion in the steady state from collapsing to unity:

$$w^\# = \left(\frac{1 - \phi_w (1+\bar{\pi})^{(\epsilon_w - 1)(1-\zeta_w)}}{1-\phi_w} \right)^{\frac{1}{1-\epsilon_w}} w, \qquad v^w = \frac{(1-\phi_w)\left(\frac{w^\#}{w}\right)^{-\epsilon_w(1+\eta)}}{1 - \phi_w (1+\bar{\pi})^{(1-\zeta_w)(\epsilon_w(1+\eta))}}. \tag{7}$$

3.2 Empirical motivation

Figure 1 plots South African CPI inflation against three measures of price dispersion: (i) a model-based nonlinear index derived from the exact law of motion, (ii) its log-linearised approximation, and (iii) a diffusion-index proxy constructed from micro-price data by Horn et al. (2025)[3]

To interpret the dynamics of model-implied price dispersion, we make use of a simplified log-linear approximation around the deterministic steady state:

$$\hat{v}_{p,t} = A\tilde{\pi}_t + B\hat{v}_{p,t-1}, \tag{8}$$

where $\hat{v}_{p,t}$ denotes the percentage deviation of price dispersion from steady state, and $\tilde{\pi}_t \equiv \ln(1+\pi_t) - \ln(1+\bar{\pi})$ is the log-deviation of gross inflation from its trend. The coefficients A and B capture the contemporaneous and lagged inflation effects on dispersion, and are both increasing in the trend inflation rate $\bar{\pi}$, the Calvo stickiness parameter ϕ_p, and the elasticity of substitution ϵ_p. (The full derivation is provided in the Technical Appendix.)

Three stylised facts emerge from the chart:

- Positive comovement: inflation and price dispersion move together, with correlations in the 0.54–0.68 range. This supports the model's prediction that higher inflation is associated with greater relative price misalignment (Ascari and Sbordone, 2014).
- Persistence amplification: the lagged term B rises with trend inflation, meaning that shocks

[3] The CPI diffusion index measures the breadth of price pressures in South Africa. For each month, it reports the share of 8-digit COICOP CPI components whose year-on-year inflation exceeds the mid-point of the South African Reserve Bank's (SARB) former target band (4.5%). Two versions are calculated: (i) an unweighted index that assigns equal weight to every component and (ii) a weighted index that uses official CPI expenditure weights. The series is constructed from Statistics SA micro-price data and published by Codera Analytics.

Figure 1: Inflation and model-based price-dispersion measures (2009Q1–2025Q1)

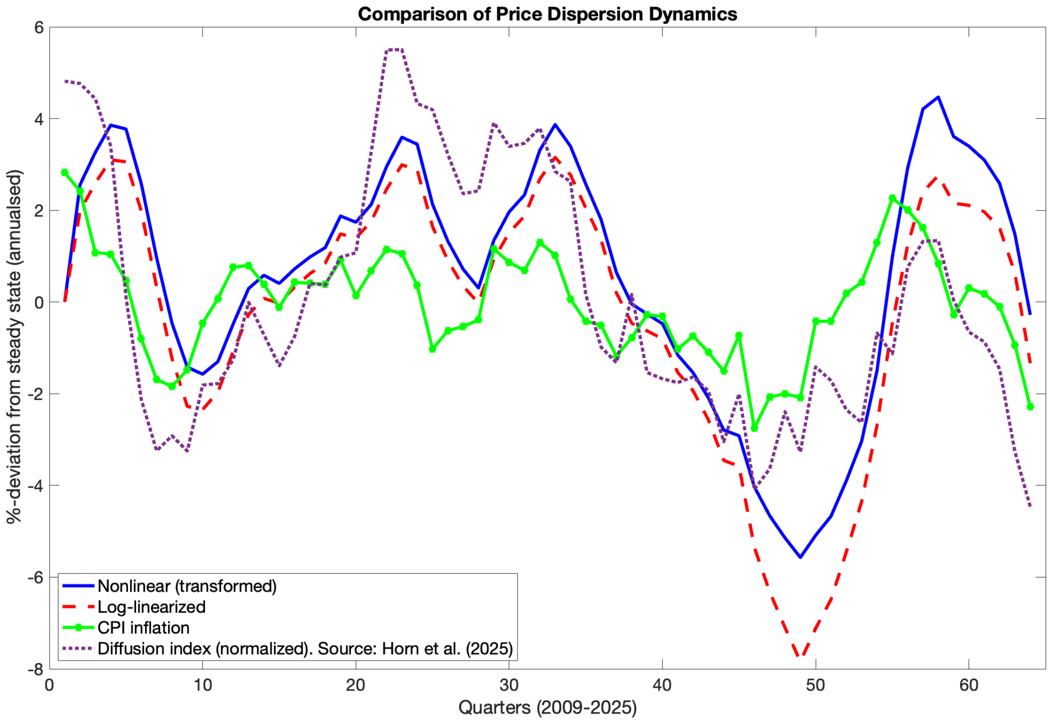

Sources: Diffusion index obtained from Horn et al. (2025). Author's calculations.

to dispersion decay more slowly in high-inflation environments. This persistence is endogenous to the model and aligns with observed inertia in the data.

- Non-linear costs: the nonlinear dispersion measure (solid blue) diverges more sharply from steady state than the linearised version (dashed red), especially during large inflation shocks. This illustrates that linear approximations understate the distortionary effects of inflation away from the steady state.

Taken together, these findings support the model's central policy implication: lowering trend inflation reduces price-dispersion distortions, steepens the Phillips curve, and reduces the output cost of disinflation. Price dispersion weakens allocative efficiency by distorting relative prices and reallocating resources away from their optimal use–resulting in output losses and higher welfare costs in high-inflation regimes.

3.3 Model calibration

We adopt a standard calibration approach commonly used in the New Keynesian literature, drawing primarily from Garín et al. (2016) and Ascari and Sbordone (2014), with targeted adjustments to reflect the structure and focus of our model. All parameter values are summarised in the Technical Appendix, where we group them by block: (i) household preferences, price setting, and monetary policy, (ii) fiscal policy and steady-state targets, and (iii) shock processes.

Key preference and technology parameters–such as the discount factor, the Frisch elasticity of labour supply, and the elasticity of substitution across goods and labour types–are set to conventional values to ensure comparability with prior studies. Nominal rigidities are calibrated based on microeconomic and macroeconomic evidence reported in Ascari and Sbordone (2014), including price and wage stickiness parameters and indexation coefficients.

The monetary policy rule is parameterised to ensure determinacy and capture realistic policy responses to inflation and output, following guidance from Ascari and Sbordone (2014). Fiscal feedback rules for taxes, transfers, and government spending are standard and guided by estimation results from Kemp and Hollander (2020); Havemann and Hollander (2024), with debt- and output-responsiveness coefficients chosen to stabilise long-run debt dynamics.

Steady-state calibration targets, such as the public debt-to-GDP ratio, government spending share, and tax rates, are aligned with empirical averages for emerging and advanced economies commonly used in the New Keynesian literature. Finally, the persistence parameters of exogenous shocks—including monetary policy, productivity, risk premia, mark-up and government spending shocks—are calibrated to match observed macroeconomic persistence, drawing from Garín et al. (2016), Ascari and Sbordone (2014), and Kemp and Hollander (2020).

This calibration strategy ensures that the model replicates realistic macroeconomic dynamics while maintaining internal consistency and alignment with the existing DSGE literature. In section 5, we estimate key structural parameters with Bayesian techniques, using South African quarterly data for 2009Q1–2019Q4. The observables–real output, CPI inflation, and the short-term policy rate–mirror the information set in the Taylor rule itself. Our analysis here focuses on the extent the monetary authority's policy reaction function is misidentified when a New-Keynesian model ignores positive trend inflation.

4 Main findings

This section separates our findings by time horizon. We first present the steady-state (long-run) effects: (i) the output cost of price dispersion and (ii) how positive trend inflation lowers long-run output and welfare. We then turn to the short-run: we explore the transition dynamics of monetary tightening and of four other canonical disturbances—technology, price mark-up, risk-premium, and government-spending shocks—to show how the economy's response changes as the trend-inflation target shifts.

4.1 The economics of price dispersion

As highlighted in the empirical motivation, price dispersion is increasing in the trend inflation rate $\bar{\pi}$, the Calvo stickiness parameter ϕ_p, and the elasticity of substitution ϵ_p. We quantify the macroeconomic cost of price dispersion by tracking the decline in effective aggregate productivity, defined as $\tilde{A}_t = \frac{A_t}{v_t^p}$, where v_t^p denotes price dispersion. This follows the approach of Damjanovic and Nolan (2010) and Ascari and Sbordone (2014).

In the intermediate-goods sector, a continuum of monopolistically competitive firms $j \in [0,1]$ set prices under Calvo stickiness. Each period, a fraction $1-\phi_p$ of firms can re-optimise their prices, while the remaining ϕ_p are stuck at past prices. As a result, even though firms are ex ante identical, they charge different prices in equilibrium due to staggered price-setting. This heterogeneity in pricing introduces a misalignment between actual and optimal relative prices, generating output losses.

Figure 2: The Cost of Price Dispersion under Alternative Calvo Parameters and Trend Inflation

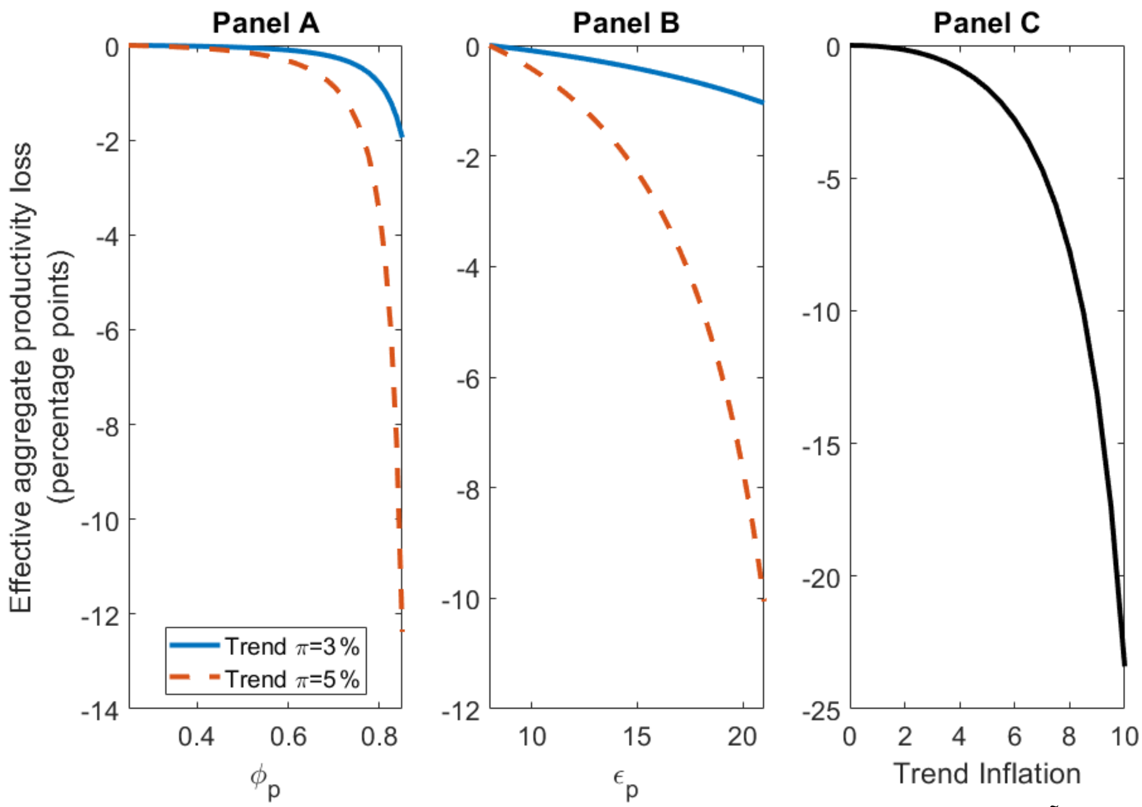

Notes: *Panel A* plots the percentage-point loss in effective aggregate productivity, defined as $100(\tilde{A}-1)$, where $\tilde{A}_t = A_t/v_t^p$, against the Calvo price-stickiness parameter ϕ_p. Each series is normalised to zero at $\phi_p = 0.5$, with ϕ_p ranging from 0.5 to 0.85 in evenly spaced steps. Solid lines are for trend inflation $\bar{\pi} = 3\%$ and dashed lines for $\bar{\pi} = 5\%$. *Panel B* shows the same productivity loss as a function of the markup-shock parameter ϵ_p, normalised to zero at $\epsilon_p = 8$ and varying ϵ_p from 8 to 21 in evenly spaced steps, again under $\bar{\pi} = 3\%$ (solid) and 5% (dashed). *Panel C* displays the un-normalised loss, also computed as $100(\tilde{A}-1)$, against trend inflation $\bar{\pi} \in \{0, 0.5, \ldots, 10\}\%$, holding $\phi_p = 0.75$ and $\epsilon_p = 11$. All results are from the author's Dynare simulations of a New Keynesian DSGE model with Calvo price setting, flexible wages, and positive trend inflation.

Source: Author's calculations.

Panel A of Figure 2 illustrates that as ϕ_p increases, output losses rise sharply—especially when trend inflation is high. With greater stickiness, fewer firms are able to reset prices, and those that are stuck incur a loss in profit relative to the optimal resetters. This inefficiency depresses aggregate output. The losses are amplified nonlinearly at higher $\bar{\pi}$, as more frequent deviations from optimal pricing compound the dispersion effect. These results underscore the costs of a high $\bar{\pi}$: a higher trend inflation rate widens average output losses and increases the dispersion of relative prices across firms.

Panel B considers the role of the elasticity of substitution across differentiated goods, ϵ_p. In a Dixit–Stiglitz setting, a higher ϵ_p implies that consumers are more responsive to relative price changes. When firms raise prices relative to others, demand drops off more sharply. As a result, price misalignments are more costly in terms of lost output. We observe that output losses increase with ϵ_p, and again the effect is more pronounced at higher trend inflation rates. This is because inflation-induced price gaps are more distortive in economies where substitution across varieties is easier.

This result is intuitive under monopolistic competition with constant elasticity of substitution. A firm's steady-state markup over marginal cost is given by $\mu = \frac{\epsilon_p}{\epsilon_p - 1}$. As ϵ_p increases, this markup declines, approaching $\mu \to 1$ in the limit. In this case, consumers view differentiated goods as nearly perfect substitutes, severely limiting each firm's ability to sustain prices above marginal cost. Consequently, even small price deviations from the optimal markup lead to substantial shifts in demand, making price misalignments more damaging for aggregate efficiency.

Panel C isolates the role of trend inflation itself. As $\bar{\pi}$ rises, price dispersion increases and reduces effective productivity even when ϕ_p and ϵ_p are held constant. Since firms not resetting their prices fall further out of line with optimal markups under high inflation, the distortion accumulates. The result is a reduction in the output produced per unit of labour, tightening the aggregate resource constraint. This confirms that price dispersion acts as a direct channel of allocative inefficiency, lowering both productivity and steady-state output in high-inflation regimes.

Taken together, these results underscore the macroeconomic cost of allowing trend inflation to drift higher. Price dispersion, as an endogenous by-product of nominal rigidities, lowers output even in the absence of shocks and undermines the allocative role of prices. Conversely, a higher $\bar{\pi}$ intensifies these distortions, eroding overall economic efficiency.

4.2 Trend inflation, long-run output and welfare

As already shown in Figure 2 Panel C, higher trend inflation ($\bar{\pi}$) raises price dispersion and depresses effective productivity. Figure 3 reports key steady-state variables as percentage deviations from the zero-inflation benchmark. Panel A confirms the corresponding transmission mechanism: lowering $\bar{\pi}$ reduces dispersion, alleviates the productivity wedge, and shrinks the steady-state output loss; the disinflationary gains are consistent with the transmission mechanisms documented by Ascari and Sbordone (2014).

Figure 3: Trend Inflation and Selected Steady-State Variables

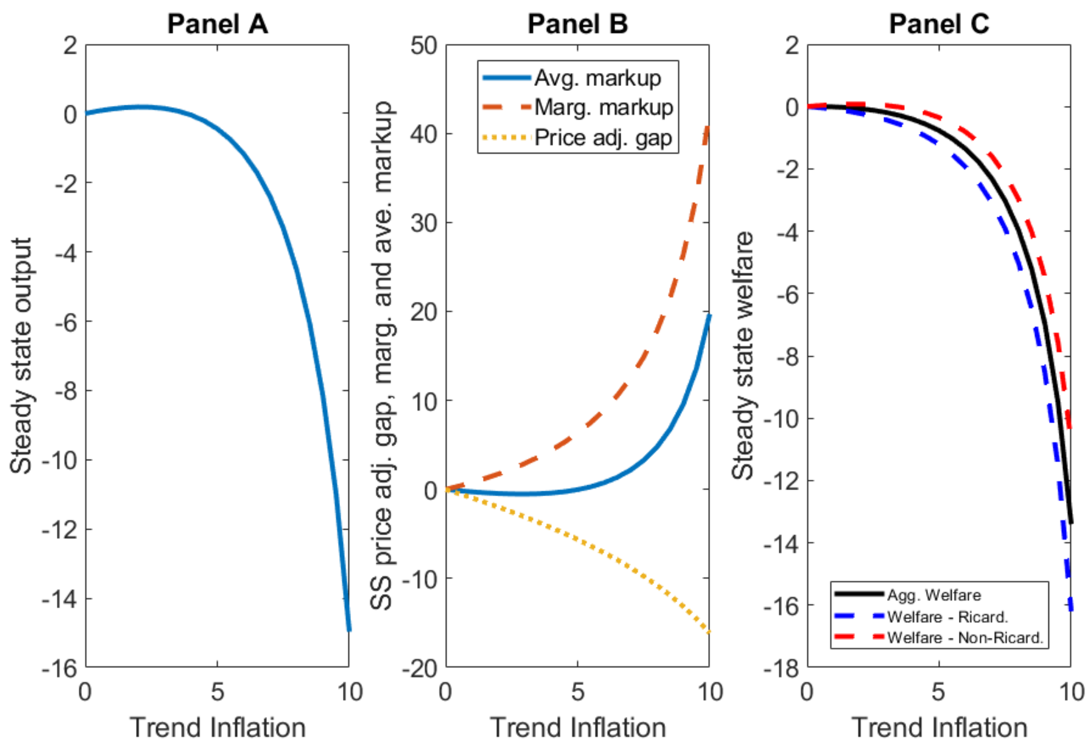

Notes: *Panel A* plots the percent deviation from the baseline steady state at $\bar{\pi} = 0\%$ of aggregate output Y and consumption by Ricardian (C^R) and non-Ricardian (C^{NR}) households. Specifically, each series is computed as $100\,\frac{X(\bar{\pi})-X(0)}{|X(0)|}$, where $X \in \{Y, C^R, C^{NR}\}$, and is shown against trend inflation $\bar{\pi} \in \{0, 0.5, \ldots, 10\}\%$. *Panel B* shows the analogous percent deviations of the average markup, the marginal markup, and the price-adjustment gap (see equation 4 for further details), each normalised to zero at $\bar{\pi} = 0\%$ and plotted over the same grid of $\bar{\pi}$. *Panel C* displays percent deviations of household welfare W, Ricardian welfare W^R, and non-Ricardian welfare W^{NR}, again computed as $100\,(W(\bar{\pi}) - W(0))/|W(0)|$ versus $\bar{\pi} \in \{0, 0.5, \ldots, 8\}\%$. All results derive from author's Dynare simulations of a New Keynesian DSGE model with Calvo price setting, flexible wages, and positive trend inflation.

Source: Author's calculations.

Average markup As shown in Equation 4 the closed-form expression decomposes into a *price-adjustment gap* and a *marginal markup.* The price-adjustment gap falls with $\bar{\pi}$: persistent inflation rapidly erodes previously-set prices, prompting frequent resets and amplifying dispersion (Figure 3 Panel B). The marginal markup, by contrast, rises with $\bar{\pi}$ because forward-looking firms anticipate higher future costs. The net effect on μ depends on plausible calibrations for the elasticity of substitution ϵ_p. A higher trend inflation raises the average markup and the associated output distortion.

Welfare Steady-state welfare for a representative household (Ricardian or non-Ricardian) is

$$W(\pi) = v(\pi)\,\ln\big[(1-b)C(\pi)\big] - \psi\,\frac{N(\pi)^{1+\eta}}{1+\eta}, \tag{9}$$

where $C(\pi)$, $v(\pi)$, and $N(\pi)$ denote consumption, the preference shock, and labour at trend inflation π. Defining $\pi^0 = 0$ as the baseline, the percentage welfare loss is

$$\mathscr{L}(\pi) = \frac{W(\pi) - W(\pi^0)}{|W(\pi^0)|} \times 100. \tag{10}$$

Panel C confirms that aggregate welfare falls steadily as $\bar{\pi}$ rises. Price dispersion drives this result: a larger v_p lowers effective productivity, so firms would need more efficient labour input N^d to maintain the same output. Yet with flexible wages ($v_w = 1$ in our baseline), the wage-dispersion wedge vanishes and total labour supplied simply tracks labour demand, $N = N^d v_w = N^d$. Once we relax the flexible-wage assumption ($v_w \neq 1$), dispersion drives a wedge between the labour households are willing to supply and firms' demand, decoupling N from N^d.

Ricardian households, who internalise lower real returns, cut consumption and labour, further depressing their welfare. Non-Ricardians do not react to real rates and are cushioned by rising transfers, so their welfare slips more slowly. Even so, the productivity loss from price dispersion dominates, and economy-wide welfare declines monotonically with trend inflation.

Overall, we find that allowing inflation to settle at a higher trend rate magnifies price-dispersion distortions, depresses long-run output, and erodes welfare.

4.3 Trend inflation and short-run dynamics

Trend inflation not only alters steady-state outcomes; it also reshapes how the economy absorbs short-run disturbances. The sections that follow trace impulse responses to five canonical shocks–monetary policy, technology, mark-up (cost-push), risk-premium, and government spending–under alternative steady-state inflation rates, $\bar{\pi} \in 0, 2, 4, 6\%$. Juxtaposing these paths reveals a common pattern: as $\bar{\pi}$ rises, price dispersion widens, the Phillips curve flattens, and automatic stabilisers weaken, so disinflation requires larger output sacrifices, positive demand shocks lose traction, and the debt-to-GDP ratio tends to rise more persistently.

Monetary policy shock Figure 4 plots the responses to a one-percentage-point (annualised) innovation to the policy rate, $\varepsilon_t^i \sim$ i.i.d.$(0, \sigma_i^2)$, under four steady-state inflation rates, $\bar{\pi} \in \{0, 2, 4, 6\}\%$. Each panel reports the percentage deviations of output, inflation, the nominal and real interest rate, the price-dispersion wedge, and the debt-to-GDP ratio from their pre-shock steady states.

We use the following Taylor rule specification (shown here in log-linearised form):

$$\hat{i}_t = \rho_i \, \hat{i}_{t-1} + \left(\phi_\pi \, \tilde{\pi}_t + \phi_y \, \Delta \hat{y}_t\right) + \varepsilon_t^i \,, \tag{11}$$

where $\hat{i}_t$ is the deviation of the nominal interest rate from its steady state, $\tilde{\pi}_t = \hat{\pi}_t - \hat{\bar{\pi}}$ is inflation relative to its target (trend inflation rate in steady-state), $\Delta\hat{y}_t$ is output growth, and ρ_i captures interest-rate smoothing (persistence in the policy rate).[4]

The monetary policy shock raises the nominal policy rate on impact, but because expected inflation adjusts only gradually the real interest rate rises by slightly more than the nominal rate, tightening intertemporal consumption conditions. Ricardian households therefore postpone expenditure, and output contracts persistently. This lower aggregate demand reduces marginal costs causing a fall in inflation. As the pace of price increases slow, firms that cannot re-optimise drift slightly towards their desired mark-ups and the price-dispersion wedge v_p therefore narrows. Less dispersion raises effective productivity, cushioning–but not offsetting–the demand-driven fall in output.

Figure 4: Impulse responses to a 1 pp monetary-policy shock under alternative trend-inflation targets

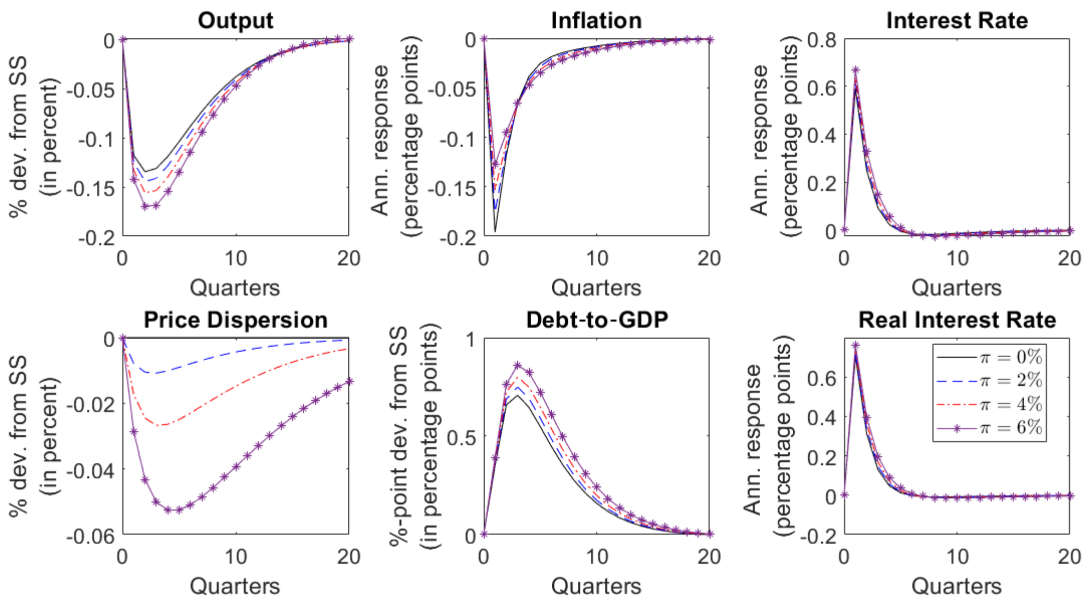

Notes: The figure shows first-order IRFs over 20 quarters following a one-percentage-point annualised shock to the Taylor rule (quarterly s.d. 0.0025, persistence $\rho_i = 0.75$). Output and price dispersion are reported as percentage deviations from steady state. Inflation, the nominal policy rate and the real rate are annualised percentage-point deviations. The debt-to-GDP panel shows percentage-point changes in the ratio. Lines correspond to trend-inflation targets $\bar{\pi} \in \{0, 2, 4, 6\}\%$ (black solid, blue dashed, red dash-dot, purple circles).

Source: Author's Dynare simulations.

On the fiscal side, lower output and higher interest payments push the debt-to-GDP ratio up by roughly one percentage point. Automatic stabilisers behave as expected: the output contraction temporarily pushes the government-spending-to-output ratio above its steady state, but as activity recovers the ratio falls back, reflecting the endogenous unwinding of the fiscal response.

[4]For calibration we adopt the Taylor rule without an extra $(1 - \rho_i)$ smoothing weight on the contemporaneous inflation and output terms, namely with $\rho_i = 0.75$. We set $\phi_\pi = 1.5$ in the simulations and deliver comparable monetary responses over a trend-inflation range of 0–10%. Ascari and Sbordone (2014) set $\phi_\pi = 2$ for their monetary policy shock to allow for determinacy at the 6% trend inflation case, which differs from their baseline calibration where they set $\phi_\pi = 1.5$.

Raising trend inflation from $\bar{\pi} = 0\%$ to 6% amplifies the initial output loss and slows the recovery. When $\bar{\pi}$ is higher, Calvo-pricing firms reset less aggressively to the fall in marginal costs, so the Phillips curve is flatter; a given contraction in output therefore produces a smaller disinflation. Evaluating the Taylor rule at a higher steady state inflation level produces a mechanically larger jump in the nominal policy rate. Yet the ex-ante real rate rises only a little more, because expected inflation falls by less when the Phillips curve is flatter. Consequently, the fall in demand dominates and the output trough deepens. At the same time, the narrowing of the price-dispersion wedge is larger and more persistent, because the shock moves prices closer to desired mark-ups in an environment already characterised by greater steady-state dispersion. These mechanisms are consistent with Ascari and Sbordone (2014), who show that higher trend inflation weakens the price-stability channel and raises the output-stability cost of monetary tightening.

We find, although an interest-rate hike succeeds in reducing inflation, its efficiency deteriorates as $\bar{\pi}$ rises: the disinflation gain per unit of output loss falls, and the debt burden worsens. Policymakers therefore face a starker output–inflation trade-off in high-inflation regimes, suggesting a more cautious tightening stance than would be optimal at low trend inflation.

Technology shock A one-percentage-point innovation to total factor productivity (ε_a) shifts the production function outward, enabling the economy to produce more output for a given amount of labour input. Figure 5 shows real output rising immediately—peaking almost 0.4% above steady state—before returning over the next twenty quarters. Lower marginal costs pass through to prices, generating a brief disinflation.

The monetary authority's Taylor rule prescribes an immediate cut in the policy rate. Given expected inflation drops, the real rate also falls. The resulting reduction in real borrowing costs stimulates aggregate demand, reinforcing the supply-side expansion triggered by the productivity shock.

The disinflation also narrows price dispersion, indicating that the misallocation created by staggered price-setting remains quantitatively negligible. On the fiscal side, the automatic stabiliser in government spending, together with the larger output denominator, compresses the debt-to-GDP ratio by nearly one percentage point after 8 quarters.

Increasing trend inflation from $\bar{\pi} = 0\%$ to 6% monotonically attenuates every response. Higher trend inflation weakens the pass-through from marginal cost to prices, so the fall in inflation is smaller, the policy-rate cut is shallower, and the associated gains in output and debt sustainability are correspondingly muted. Hence, elevated trend inflation dampens the real effects of technology shocks and reduces the corrective burden placed on monetary policy.

Figure 5: Impulse responses to a 1 pp technology shock under alternative trend-inflation targets

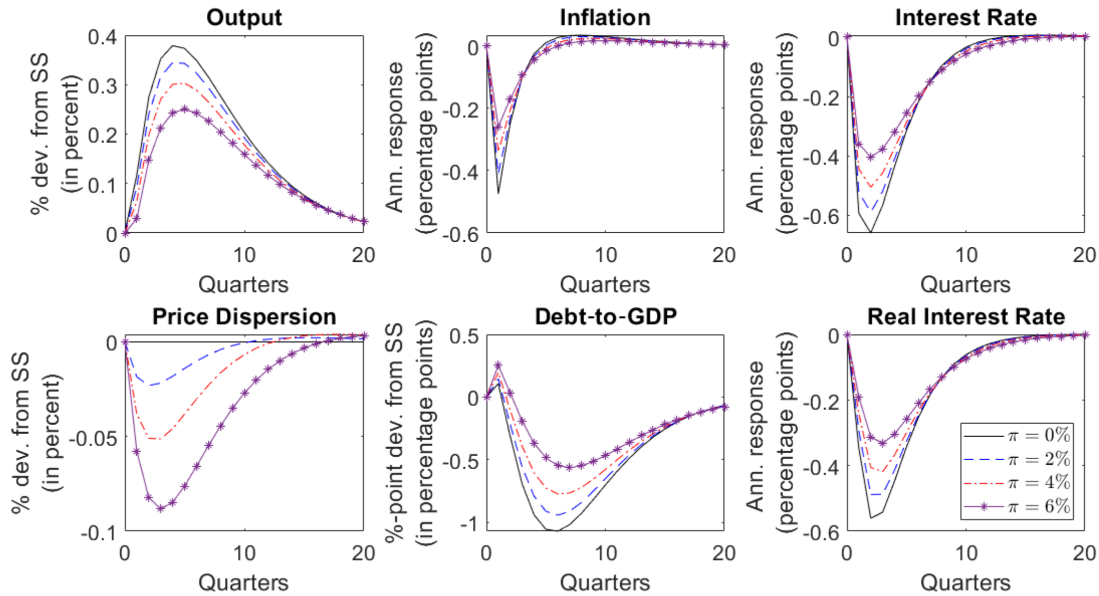

Notes: First-order IRFs over 20 quarters following a one-percentage-point annualised TFP shock (quarterly s.d. calibrated to match the 1 pp size, persistence ρ_a). Output and price dispersion are reported as percentage deviations from steady state; inflation, the nominal policy rate, and the real rate are annualised percentage-point deviations; the debt-to-GDP panel shows percentage-point changes in the ratio. Lines correspond to trend-inflation targets $\bar{\pi} \in \{0, 2, 4, 6\}\%$ (black solid, blue dashed, red dash-dot, purple circles).
Source: Author's Dynare simulations.

Price elasticity (mark-up) shock A one–percentage–point innovation to the price mark-up parameter, $\varepsilon_{p,t} \sim \text{i.i.d.}(0, \sigma_p^2)$, lowers firms' desired mark-ups in the reset-price condition. Firms that can re-optimise will cut prices, while those that cannot remain temporarily overpriced, so price dispersion v_p rises (Panel D in Figure 6). The lower average mark-up raises real marginal cost and stimulates demand, lifting output on impact. Inflation briefly turns negative on impact. The Taylor rule therefore cuts the policy rate immediately, the ex-ante real rate follows, and the demand boost is reinforced. Debt-to-GDP initially falls given higher output and lower interest payments and then drifts back as the stimulus fades.

Across trend–inflation targets $\bar{\pi} \in \{0, 2, 4, 6\}$ percent, the differences are muted. The price–dispersion response v_p rises on impact in all cases and is slightly larger and more persistent when $\bar{\pi}$ is higher: re-optimising firms cut prices while non-adjusting firms remain temporarily overpriced, and the inertia of v_p increases with trend inflation. Despite the short-run boost to output the underlying price-dispersion wedge v_p is larger at high trend inflation, and so the economy continues to bear a permanently higher misallocation despite the stronger short-run expansion.

Risk premium shock Panel F of Figure 7 shows a one–percentage-point innovation to the risk premium which raises the risk-adjusted real rate $r_t^{\text{adj}} = \frac{(1+i_t)(1+rp_t)}{1+\mathbb{E}_t \pi_{t+1}} - 1$. Higher real borrowing costs tighten households' intertemporal condition, so Ricardian consumers defer expenditure and output contracts persistently. Lower production reduces labour demand, pushing marginal costs

Figure 6: Impulse responses to a 1 pp mark-up shock under alternative trend-inflation targets

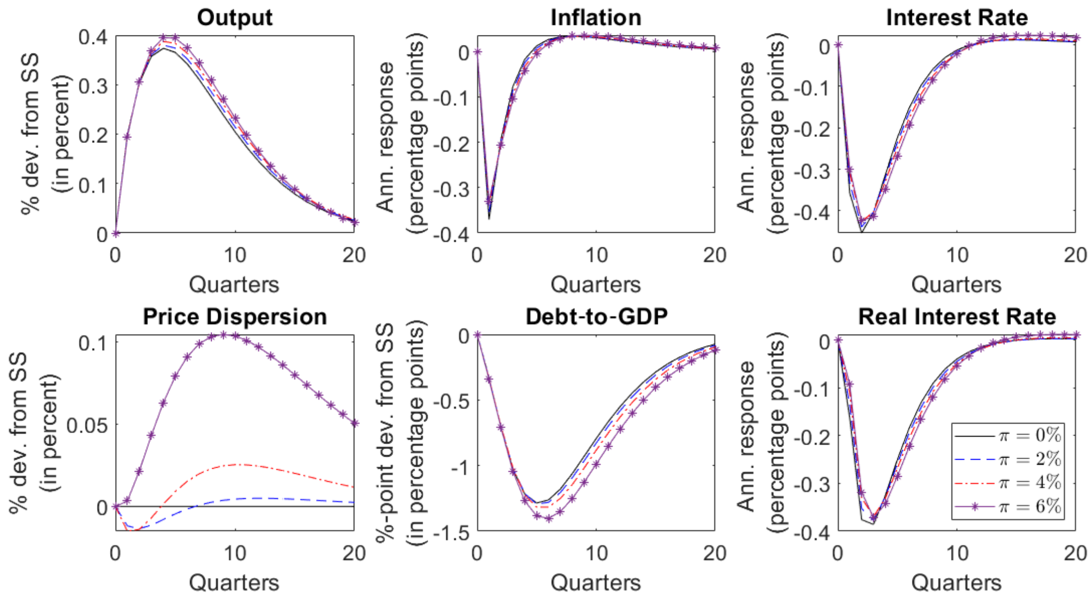

Notes: First-order IRFs over 20 quarters after a one-percentage-point annualised mark-up shock (quarterly s.d. 0.0025). Output and price dispersion are expressed as percentage deviations from steady state; inflation, the nominal rate and the real rate are annualised percentage-point deviations; the debt-to-GDP panel shows percentage-point changes in the ratio. Lines correspond to trend-inflation targets $\bar{\pi} \in \{0, 2, 4, 6\}\%$ (black solid, blue dashed, red dash-dot, purple circles).

Source: Author's Dynare simulations.

Figure 7: Impulse responses to a 1 pp risk-premium shock under alternative trend-inflation targets

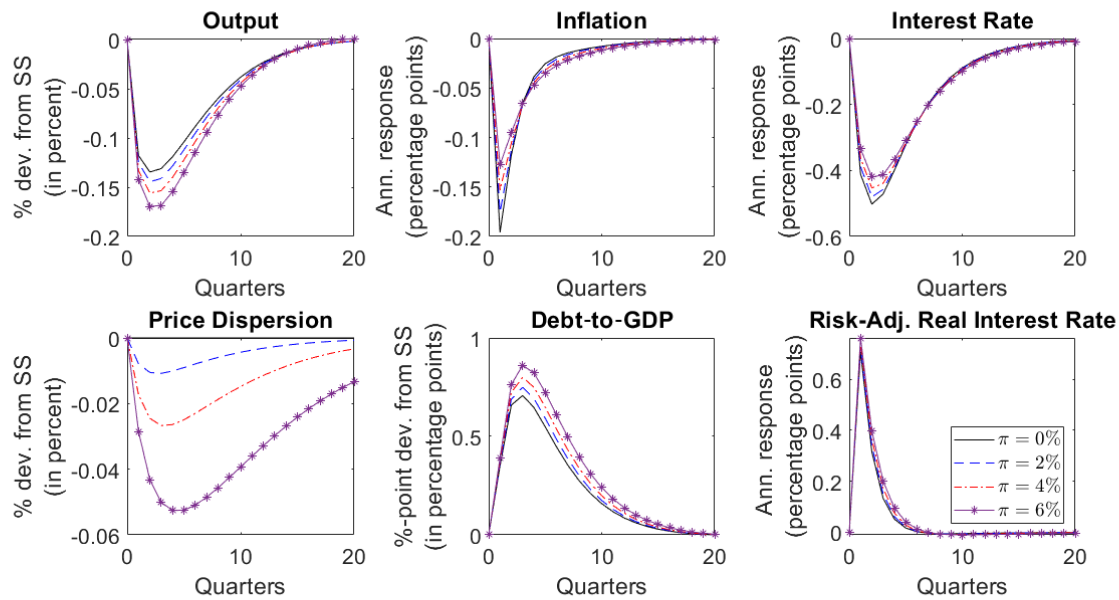

Notes: First-order IRFs over 20 quarters following a one-percentage-point *annualised* risk-premium shock (quarterly s.d. 0.0025, persistence ρ_{rp}). Output and price dispersion are reported as percentage deviations from steady state; inflation, the nominal rate, and the risk adjusted real rate are annualised percentage-point deviations; the debt-to-GDP panel shows percentage-point changes in the ratio. Lines correspond to trend-inflation targets $\bar{\pi} \in \{0, 2, 4, 6\}\%$ (black solid, blue dashed, red dash-dot, purple circles).

Source: Author's Dynare simulations.

and inflation down. The optimal reset price falls and previously fixed prices move closer to their desired mark-ups. Price dispersion therefore narrows, which raises effective productivity but not enough to offset the demand-driven output decline.

The joint fall in output and inflation feeds into the Taylor rule, inducing a cut in the nominal policy rate. Nevertheless, with the risk premium still elevated, the net real rate remains above its steady-state level, so the contraction continues. On the fiscal side, the automatic stabiliser in government spending, the smaller output denominator, and higher real interest payments jointly raise the debt-to-GDP ratio.

The increase in steady-state inflation from $\bar{\pi} = 0\%$ to 6% raises the initial output loss and slows the recovery. At higher $\bar{\pi}$, Calvo-pricing firms adjust prices less, so the disinflation is smaller. While the Taylor rule responds to the milder fall in inflation, the policy-rate cut is attenuated, leaving the risk-adjusted real rate slightly higher for longer. Consequently, demand is squeezed more persistently and the output trough deepens, mirroring the pattern observed under the monetary-policy shock.

Government spending shock Figure 8 reports the effects of an i.i.d. spending innovation, $\varepsilon_{G,t} \sim \text{i.i.d.}(0, \sigma_G^2)$. The exogenous rise in real public expenditure appears directly in the resource constraint and lifts aggregate demand. Output consequently jumps on impact and then decays toward steady state. Higher demand raises firms' marginal costs, so the optimal reset price increases and inflation spikes. In response to the price surge, Ricardian households partially smooth consumption: the expected path of tighter monetary policy and future tax liabilities induces them to save a fraction of the windfall.

Inflation broadens the distribution of relative prices, pushing the price-dispersion wedge v_p upward. The associated fall in effective productivity offsets part of the fiscal stimulus but does not overturn the positive output effect. The Taylor rule reacts to the inflationary pressure by increasing the nominal policy rate, reinforcing the consumption-smoothing objective of Ricardian households.

Government expenditure rises immediately, whereas taxation adjusts only gradually: the primary deficit therefore widens. Although stronger output initially lowers the debt-to-GDP ratio through the denominator, higher real interest payments and the persistent spending impulse dominate after a few quarters, so debt peaks above the steady state and persistently decays.

Considering the shock enters the resource constraint directly, the level of trend inflation, $\bar{\pi}$, leaves the initial output response largely unchanged. It does, however, shape the nominal adjustment: a higher $\bar{\pi}$ flattens the Phillips curve, so the on-impact inflation spike is smaller, but price dispersion rises by more and decays more slowly. Consequently, inflation remains above its low-inflation counterpart for several quarters, while the policy-rate response is attenuated. The public-debt dynamics are broadly unaffected, so the debt-to-GDP trajectory differs little across targets.

Although the output multipliers for a government spending shock are less sensitive to adjustments

Figure 8: Impulse responses to a 1 pp government-spending shock under alternative trend-inflation targets

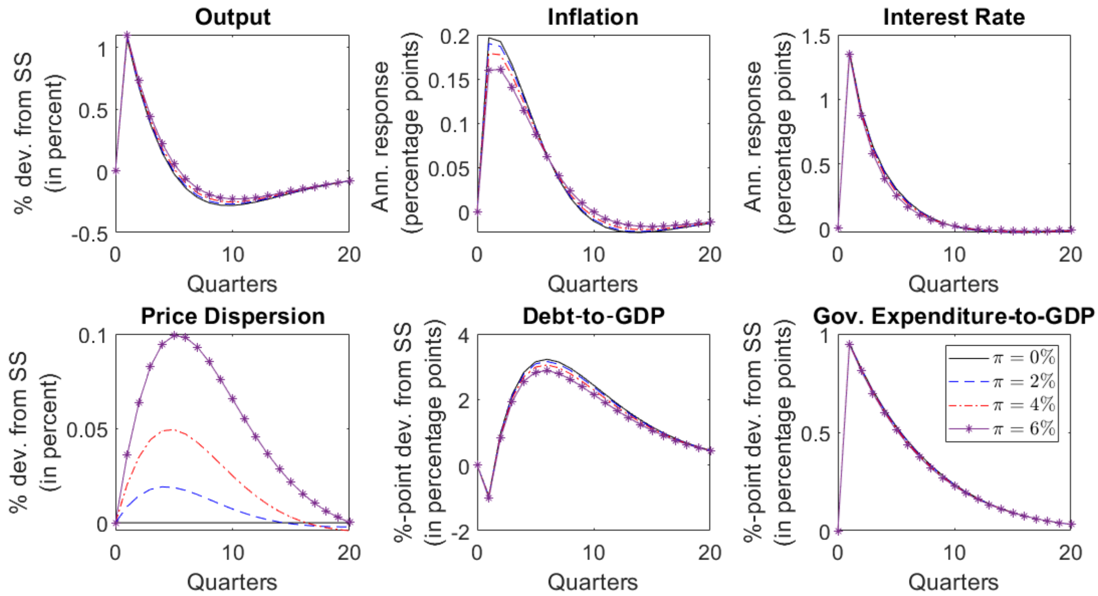

Notes: First-order IRFs over 20 quarters after a one-percentage-point annualised shock to government spending (quarterly s.d. 0.05, calibrated to raise G/Y by about 1 pp on impact). Output and price dispersion are shown as percentage deviations from steady state; inflation, the nominal rate, and the real rate as annualised percentage-point deviations; debt-to-GDP and government-expenditure-to-GDP as percentage-point changes. Lines correspond to trend-inflation targets $\bar{\pi} \in \{0, 2, 4, 6\}\%$ (black solid, blue dashed, red dash-dot, purple circles).

Source: Author's Dynare simulations.

in $\bar{\pi}$ we are interested in understanding the transmission mechanisms. The value this brings is in understanding how nominal variables—i_t, π_t, and $v_{p,t}$—are sensitive in ways that feed into debt-service costs through the policy-rate path and the persistence of dispersion. This highlights the value of analysing fiscal rules jointly with monetary policy when trend inflation is elevated, a question pertinent to future work on disinflation policy.

4.4 Phillips curve under trend inflation and the sacrifice ratio

This section studies how trend inflation shapes (i) the short-run steepness of the Phillips curve and (ii) the implied sacrifice ratio.

Generalised New Keynesian Phillips Curve with explicit coefficients For simplicity of notation, we let Π denote gross steady-state trend inflation. We provide detailed derivations in the Technical Appendix. The log-linearised Phillips curve in marginal-cost form is

$$\hat{\pi}_t = \tilde{\beta}_\pi(\Pi)\,\mathbb{E}_t[\hat{\pi}_{t+1}] + \kappa_{mc}(\Pi)\,\widehat{mc}_t + \gamma(\Pi)\,\hat{\pi}_{t-1}, \tag{12}$$

with coefficients

$$\theta_1(\Pi) \equiv \phi_p \beta \Pi^{\varepsilon_p - \zeta_p \varepsilon_p}, \tag{13}$$

$$s_0(\Pi) \equiv \frac{1-\phi_p}{(1-\phi_p)+\phi_p \Pi^{\zeta_p(1-\varepsilon_p)}}, \qquad s_1(\Pi) \equiv \frac{\phi_p \Pi^{\zeta_p(1-\varepsilon_p)}}{(1-\phi_p)+\phi_p \Pi^{\zeta_p(1-\varepsilon_p)}}, \tag{14}$$

$$\tilde{\beta}_\pi(\Pi) \equiv \frac{\theta_1(\Pi)}{1+\theta_1(\Pi)\zeta_p s_1(\Pi)}, \qquad \gamma(\Pi) \equiv \frac{\zeta_p s_1(\Pi)}{1+\theta_1(\Pi)\zeta_p s_1(\Pi)}, \tag{15}$$

$$\kappa_{mc}(\Pi) \equiv \frac{s_0(\Pi)}{1+\theta_1(\Pi)\zeta_p s_1(\Pi)}. \tag{16}$$

Figure 9 plots the structural slope $\kappa_{mc}(\Pi)$—the contemporaneous mapping from marginal cost to inflation, holding fixed the forward-looking and indexation terms. Our object $\kappa_{mc}(\Pi)$ in (12) varies with two distinct margins: (i) price stickiness ϕ_p and (ii) trend inflation Π (through indexation ζ_p). Holding Π fixed, the usual Calvo comparative static obtains:

$$\frac{\partial \kappa_{mc}(\Pi)}{\partial \phi_p} < 0,$$

and so greater price stickiness implies a flatter Phillips slope, consistent with Jacome et al. (2025). In our explicit coefficients, both the numerator $s_0(\Pi)$ falls and the denominator $1+\theta_1(\Pi)\zeta_p s_1(\Pi)$ rises with ϕ_p, reinforcing $\partial\kappa_{mc}/\partial\phi_p < 0$. By contrast, in Figure 9 we hold ϕ_p fixed and vary the trend inflation Π. With partial indexation ($\zeta_p > 0$) we obtain

$$\frac{\partial \kappa_{mc}(\Pi)}{\partial \Pi} > 0,$$

so that $\kappa_{mc}(\Pi)$ *rises* with trend inflation and tends toward the no-indexation bound $\kappa_{mc} = 1-\phi_p$ as Π increases. When there is no indexation ($\zeta_p = 0$), $\kappa_{mc}(\Pi) \equiv 1-\phi_p$ is constant in Π.

Reduced-form impulse response function slope To obtain a reduced-form summary comparable to empirical estimates, we summarise the joint $(\widehat{mc}_t, \hat{\pi}_t)$ impulse-response path from a monetary-policy shock by the slope of a line fitted through the origin over the first H quarters shown as:

$$\beta_0^{\mathrm{RF}}(\Pi; H) = \arg\min_b \sum_{t=0}^{H} \left(\hat{\pi}_t - b\,\widehat{mc}_t\right)^2 = \frac{\sum_{t=0}^{H} \widehat{mc}_t \hat{\pi}_t}{\sum_{t=0}^{H} \widehat{mc}_t^2}. \tag{17}$$

Given the structural relation is anchored at the origin (zero intercept in $(\widehat{mc}_t, \hat{\pi}_t)$ space), the no-constant slope β_0^{RF} is the empirical counterpart to the theoretical κ-tangent.

Figure 9: Inflation–marginal cost across trend-inflation targets

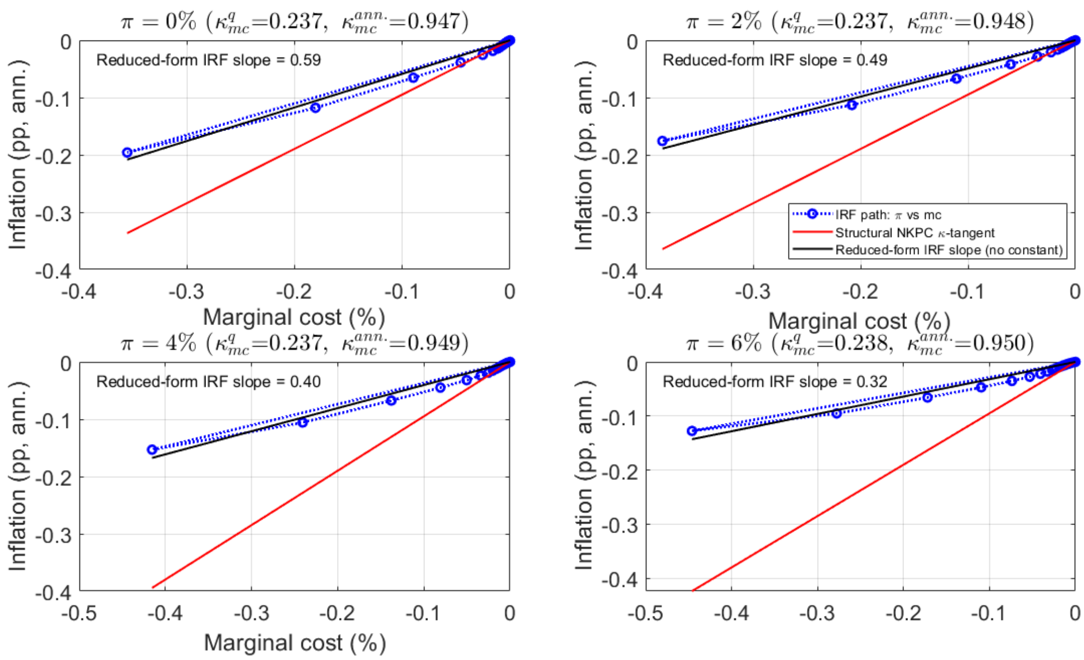

Notes: First-order IRFs over 20 quarters after a one-percentage-point *annualised* monetary-policy shock (innovation scaled to deliver a 1 pp impact). Blue dotted line with markers: IRF path of inflation vs marginal cost (π_t in pp, annualised; mc_t in %). Red line: structural NKPC κ-tangent (annualised) implied by the explicit $\kappa_{mc}(\bar{\Pi})$. Black line: reduced-form IRF slope (no constant) fitted to the path. Panels compare steady-state trend inflation $\bar{\pi} \in \{0, 2, 4, 6\}$%. Calibration: $\phi_p = 0.75$, $\varepsilon_p = 11$, $\beta = 0.9884$, $\zeta_p = 0.1$ (quarterly). The red line is scaled by 4 to match annualised units for inflation.
Source: Author's Dynare simulations.

Gap form for interpretation We do not construct a flexible-price counterfactual in this paper. For interpretation, we define an *activity gap* as the log deviation of output from the trend-inflation steady state, $x_t \equiv \hat{y}_t$ where $\hat{y}_t$ is the deviation of output from its steady state. The model's cost mapping implies

$$\widehat{mc}_t = \Psi_{mc}(\Pi)\, x_t, \qquad \kappa_y(\Pi) \equiv \kappa_{mc}(\Pi)\, \Psi_{mc}(\Pi),$$

so that (12) can be written as

$$\hat{\pi}_t = \tilde{\beta}_\pi(\Pi)\, \mathbb{E}_t[\hat{\pi}_{t+1}] + \kappa_y(\Pi)\, x_t + \gamma(\Pi)\, \hat{\pi}_{t-1}. \tag{18}$$

We use (18) only to interpret slopes. The activity gap series used below comes directly from the impulse response functions.

Sacrifice ratios Over a horizon H, we report two measures constructed from the monetary policy shock impulse response functions of the activity gap x_t and annualised inflation $\hat{\pi}_t^{\text{ann}}$ (pp). The factor 1/4 converts quarter sums into years:

$$\mathrm{SR}^{\text{shock}}(\Pi; H) = -\frac{1}{4}\frac{\sum_{h=0}^{H} x_{t+h}}{\hat{\pi}^{\text{ann}}_{t+H} - \hat{\pi}^{\text{ann}}_{t}} \quad \text{(percent-years per pp)}, \tag{19}$$

$$\mathrm{SR}^{1\text{pp}}(\Pi; H) = \frac{1}{4}\sum_{h=0}^{H}\left|\frac{x_{t+h}}{\max_{0\le s\le H}\left|\hat{\pi}^{\text{ann}}_{t+s} - \hat{\pi}^{\text{ann}}_{t}\right|}\right| \quad \text{(normalised to a 1-pp inflation move)}. \tag{20}$$

Figure 10: Sacrifice ratios and Phillips-curve steepness across trend-inflation targets

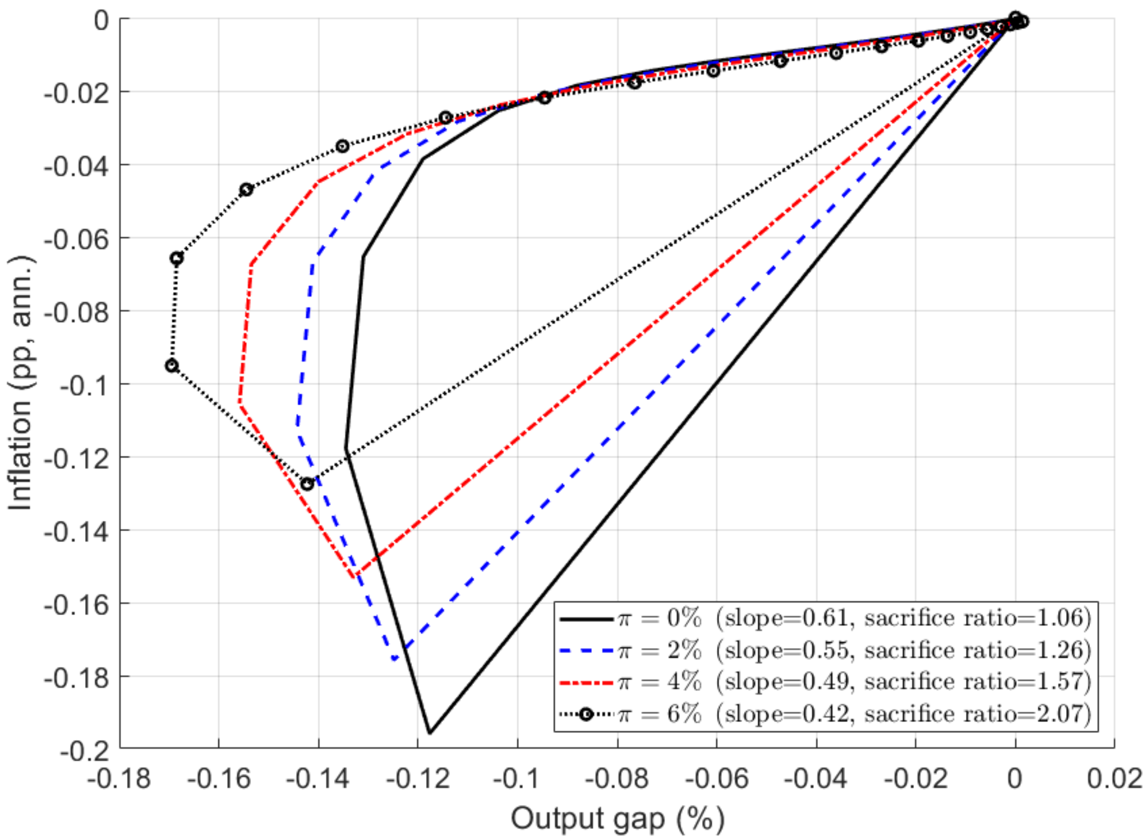

Notes: Figure plots the joint impulse response function path of annualised inflation (pp, ann.) against the activity gap x_t (output gap, %) over the first $H = 8$ quarters following a contractionary monetary-policy shock. Lines correspond to trend-inflation targets $\bar{\pi} \in \{0, 2, 4, 6\}$%. The text labels report (i) the reduced-form IRF slope — fitted line of $\hat{\pi}_t$ on x_t through the origin over $t = 0, \ldots, H$—and (ii) the sacrifice ratio $\mathrm{SR} = -(1/4)\sum_{t=0}^{H} x_t/(\hat{\pi}_H - \hat{\pi}_0)$ in years per percentage point.

Source: Author's Dynare simulations.

Two robust patterns emerge. First, the reduced-form short-run Phillips curve flattens with higher trend inflation: $\beta_0^{\text{RF}}(\Pi; H)$ declines monotonically as $\bar{\pi}$ rises (as found by Ascari and Sbordone, 2014). Second, the sacrifice ratio rises with trend inflation. In our calibration, a 1 percentage point disinflation at $\bar{\pi} = 0\%$ has a cumulative cost equivalent to keeping output a little over 1% below potential output for one year or 0.5% over two years. At $\bar{\pi} = 6\%$ the cumulative cost doubles to a little over 2% for one year to achieve a 1 percentage point disinflation. Intuitively, even though $\kappa_{mc}(\Pi)$ increases with Π under partial indexation, a larger $\tilde{\beta}_\pi(\Pi)$ places more weight on expected future inflation this makes $\hat{\pi}_t$ adjust less on impact to a given $\widehat{mc}_t$, requiring a larger output contraction to achieve a given disinflation.

5 Sensitivity analysis and policy implications

Having shown that price dispersion becomes a first-order distortion under positive trend inflation, we now assess how this modifies the classic New-Keynesian policy trade-off. In the textbook model, an optimal rule that closes the output gap automatically stabilises inflation–known as the 'divine coincidence' (Blanchard and Galí, 2007). Ascari and Sbordone (2014) demonstrate that this property breaks down once steady-state inflation is positive: changes in $\bar{\pi}$ shift the reference point around which policy is designed, so eliminating the output gap no longer guarantees price stability.

5.1 The labour market

Trend inflation interacts with wage-setting frictions as meaningfully as with price stickiness. This subsection documents three channels. First, we measure the steady-state employment cost that arises from wage-dispersion wedges when adjusting key wage-setting parameters. Second, we link those wedges to effective productivity, labour input, and consumption-equivalent welfare for Ricardian and non-Ricardian households. Third, we show how wage rigidity changes the economy's short-run response to a monetary-policy shock under low versus higher trend inflation. Taken together, the results highlight why wage dynamics deserve attention in any policy debate on trend inflation.

Costs of Wage Dispersion Positive trend inflation distorts not only goods prices but also wages when households face Calvo-style adjustment frictions. Figure 11 illustrates how three key wage-setting parameters—the adjustment probability (ϕ_w), the elasticity of substitution across labour types (ε_w) and the labour disutility (η)—shape steady-state labour under alternative targets $\bar{\pi} \in \{3, 5\}\%$.

With wage rigidity, positive trend inflation $\bar{\pi}$ creates a dispersion wedge across individual wages. Holding steady–state output fixed, this wedge raises the effective hours required to produce the same quantity of output. We present this as an employment cost the percentage-point increase in steady–state hours relative to a baseline (first point in each panel). Formally, for parameter z and baseline z_0, Employment cost$(z) \equiv 100 \times \big(N(z) - N(z_0)\big)$, where $N = v_w(\bar{\pi}, \phi_w, \varepsilon_w, \eta).N^d$ and $N^d = Y.v_p(\bar{\pi}, \phi_p, \varepsilon_p, \zeta_p)$.

Panel A shows how lowering wage stickiness through the Calvo parameter ϕ_w reduces the share

Figure 11: The cost of price and wage dispersion

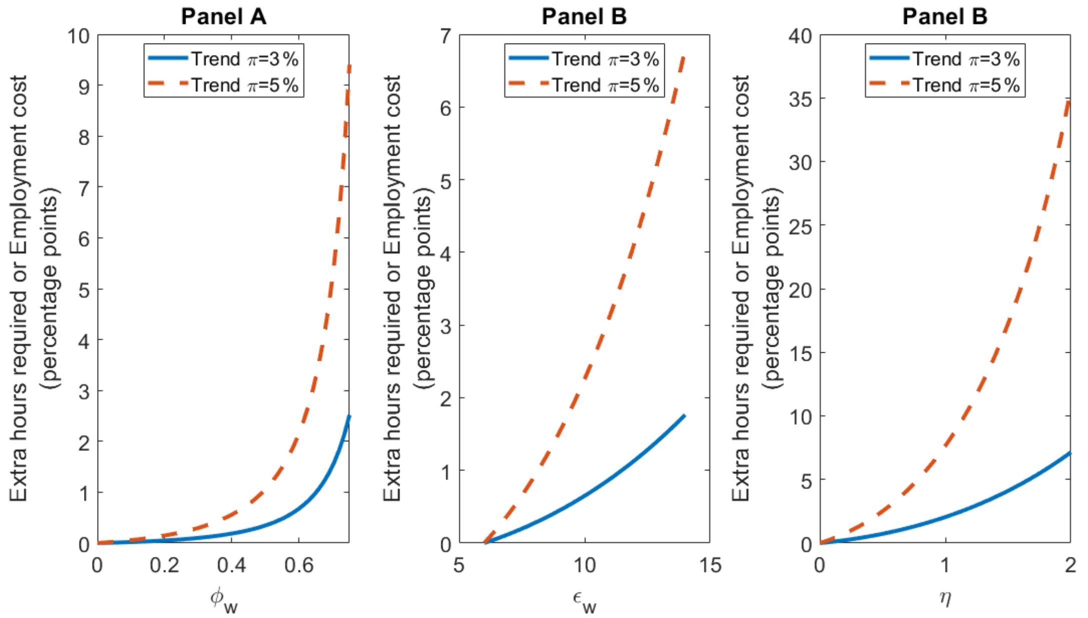

Notes: Employment cost equals the increase in steady–state hours (percentage points) required to produce the same output as the baseline (first point in each panel). *Panel A* varies the Calvo wage parameter $\phi_w \in [0, 0.75]$; *Panel B* varies labour substitutability $\varepsilon_w \in [6, 14]$; *Panel C* varies labour disutility $\eta \in [0, 2]$. Each panel reports two targets, $\bar{\pi} \in [3\%, 5\%]$. All values are computed from the model's steady state with price and wage rigidity, holding the output target fixed and allowing hours to adjust.
Source: Author's Dynare simulations.

of workers who cannot reset wages, narrowing wage dispersion. This narrower wedge reduces the labour hours needed to produce the same level of output. Conversely, higher ϕ_w raises the employment cost, with a markedly steeper rise at $\bar{\pi} = 5\%$ than at $\bar{\pi} = 3\%$.

In Panel B we show that a higher ε_w makes labour types closer substitutes, shrinking the desired wage mark-ups. However, under positive trend inflation and partial indexation, the wage-staggering term $1 - \phi_w(1 + \bar{\pi})^{(1-\zeta_w)(\varepsilon_w(1+\eta))}$ from equation 7 tightens, raising v_w and implying more labour hours are needed for the same output. This amplification is stronger at higher $\bar{\pi}$.

A higher η does not raise the optimal reset wage w, but magnifies the wage-dispersion wedge via the exponents of v^w in equation 7. With positive trend inflation and partial indexation, the denominator term $1 - \phi_w(1 + \bar{\pi})^{(1-\zeta_w)\varepsilon_w(1+\eta)}$ shrinks more rapidly as η increases, so v^w rises. Therefore, Panel C shows that holding steady-state output fixed, firms require more labour hours as η increases, with the effect more pronounced at $\bar{\pi} = 5\%$ than at $\bar{\pi} = 3\%$.

These patterns confirm that labour-market frictions magnify the real resource costs of positive trend inflation and should be acknowledged in policy assessments.

Price– and wage–dispersion wedges Figure 12 shows how rising trend inflation $\bar{\pi}$ widens price and wage dispersion, lowering effective productivity and raising the hours needed to produce a given output.

Panel A shows how as $\bar{\pi}$ increases, price dispersion v_p rises and aggregate effective productivity $\tilde{A} = A/v_p$ falls, while total factor productivity A is flat. The productivity loss is therefore a dispersion effect.

In Panel B, the focus shifts to wage dispersion and the employment wedge. Total hours satisfy $N = v_w N^d$. If $v_w = 1$, all workers are paid the same (no wage dispersion), so $N^d = N$ and there is no inefficiency. When $v_w > 1$, wage dispersion widens, meaning not all firms employ workers at the optimal wage and the aggregate labour input the economy gets per hour worked is less efficient. With higher $\bar{\pi}$, wage dispersion v_w rises steeply. Labour demand in efficiency units N^d also rises, but remains well below total hours supplied N. The gap $N - N^d$ is exactly the wage-dispersion wedge. The gap reflects an efficiency wedge: wage dispersion distorts the allocation of work across households, so firms obtain fewer effective labour services per hour supplied.

Figure 12: Trend inflation and steady-state dispersions under price and wage rigidity

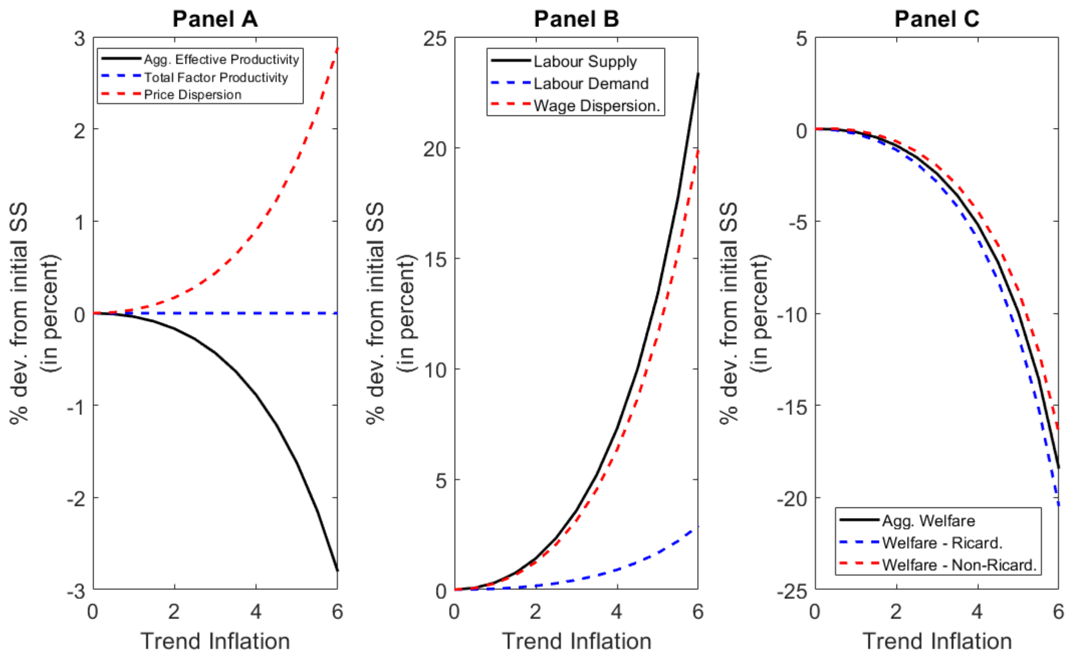

Notes: Percent deviations of key variables from the zero-loss baseline steady state (at $\bar{\pi} = 0\%$) as trend inflation $\bar{\pi}$ varies. *Panel A* plots aggregate effective productivity $\tilde{A}$ (solid black), total factor productivity A (dashed blue), and price dispersion v_p (dash-dot red) for $\bar{\pi} \in \{0, 0.5, \ldots, 6\}\%$. *Panel B* shows labour supply N (solid black), labour demand N_d (dashed blue), and wage dispersion v_w (dash-dot red) over the same range of $\bar{\pi}$. *Panel C* plots overall welfare W (solid black), Ricardian welfare W_R (dashed blue), and non-Ricardian welfare W_{NR} (dash-dot red) against $\bar{\pi}$. All values are obtained from Dynare steady-state simulations of the New Keynesian model with Calvo price stickiness $\phi_p = 0.75$, price indexation $\zeta_p = 0.1$, Calvo wage stickiness $\phi_w = 0.75$, and no wage indexation $\zeta_w = 0.1$.

Source: Author's Dynare simulations.

Panel C reports welfare. Aggregate welfare falls monotonically with $\bar{\pi}$. Losses are larger for Ricardians than for non-Ricardians.

Welfare variation Figure 13 reports consumption–equivalent welfare as trend inflation $\bar{\pi}$ rises and the share of rule–of–thumb (non-Ricardian) households varies, $\omega = \omega_w \in \{0, 0.4, 0.8\}$. Throughout the figure, we adopt a transfer–residual fiscal closure: transfers adjust endogenously to satisfy the government budget constraint given fixed fiscal instruments. This makes the fiscal trade-off transparent—higher $\bar{\pi}$ raises price and wage dispersion, yielding an efficiency loss but also reshuffles resources across households via the household transfer rule.

Figure 13: Consumption-equivalent welfare variation under non-Ricardian households

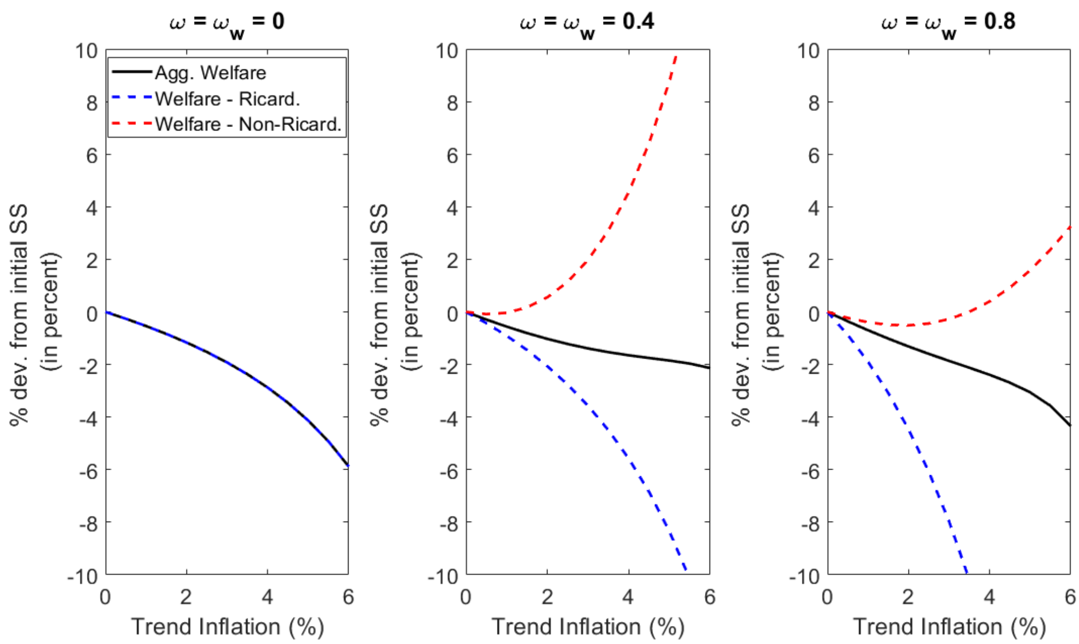

Notes: Percent deviations of steady-state consumption-equivalent welfare from the zero-loss baseline steady state (at $\bar{\pi} = 0\%$) as trend inflation $\bar{\pi}$ varies over $\{0, 0.5, \ldots, 8\}\%$. *Panels A–C* plot consumption-equivalent welfare, W (solid black), W_R (dashed blue) and W_{NR} (dash-dot red), for $\omega \in 0, 0.4, 0.8$—corresponding to no non-Ricardian households, a medium share, and a dominant share—while varying trend inflation $\bar{\pi}$ over $0, 0.5, \ldots, 8\%$. Closure: transfers adjust to balance the government budget constraint with debt, spending and tax revenue targets fixed. All values are from Dynare steady-state simulations of the New Keynesian model with $\phi_p = 0.75$, $\zeta_p = 0.1$, $\phi_w = 0.75$, and $\zeta_w = 0.1$.
Source: Author's Dynare simulations.

All Ricardian ($\omega_w = 0$, left panel). We find that with fully optimising households, the dispersion wedge dominates and aggregate welfare W falls monotonically as $\bar{\pi}$ rises.

Mixed economy ($\omega_w = 0.4$, middle panel). Endogenous transfers partly shield non-Ricardians, so non-Ricardian transfers W_{NR} increases with $\bar{\pi}$ while Ricardian transfers W_R declines more steeply. Aggregate welfare W falls less than in the $\omega = 0$ case as redistribution offsets some efficiency loss.

Non-Ricardians dominant ($\omega_w = 0.8$, right panel). When rule-of-thumb households form the majority, the redistributive channel is strong: W_{NR} rises with $\bar{\pi}$ and W_R falls sharply. The aggregate path W reflects the balance between dispersion costs and transfers, yielding a markedly flatter and potentially non-monotonic profile relative to $\omega = 0$.

Welfare profiles are sensitive to how the government budget constraint is closed. Under alternative closures—such as holding transfers fixed and letting debt adjust, or varying tax rates—the same dispersion mechanism operates, yet shifting the fiscal adjustment across instruments can produce different aggregate welfare profiles. We adopt the transfer-residual closure because it most directly exposes the fiscal trade-off in a heterogeneous-household setting. Results under other closures are available on request.

Monetary policy and labour market rigidities Figure 14 compares the transmission of a one-percentage-point monetary policy tightening under two types of nominal rigidity: price stickiness only (dashed lines, $\phi_w = 0$) and joint price and wage stickiness (solid lines, $\phi_p = \phi_w = 0.75$). Each case is shown for low ($\bar{\pi} = 0\%$, black) and moderate ($\bar{\pi} = 4\%$, blue) trend inflation.

Figure 14: Monetary-policy shock under price & wage versus price-only rigidity

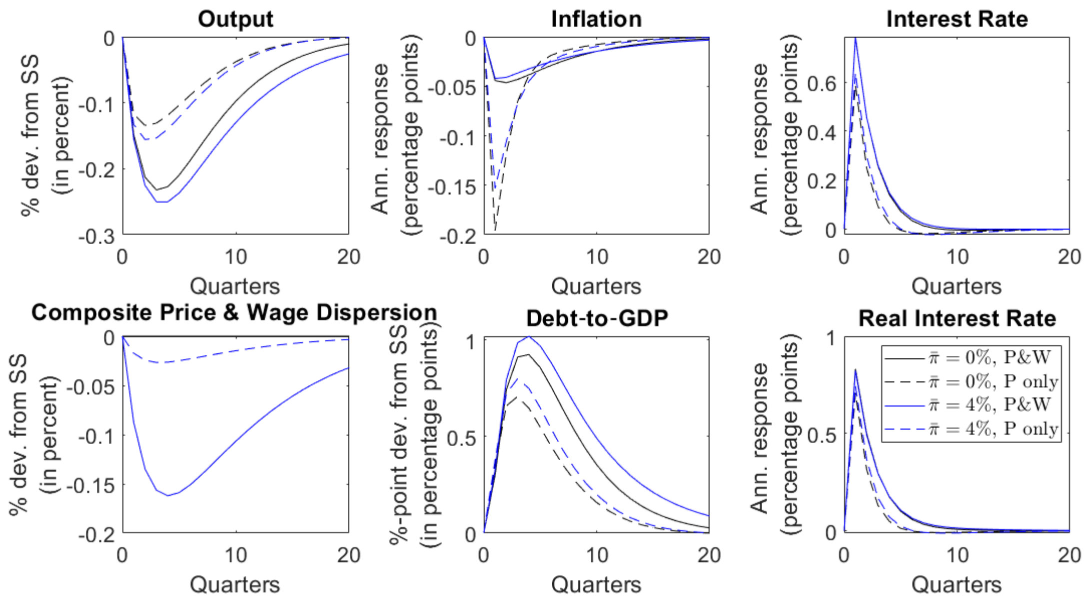

Notes: Impulse responses over 20 quarters to a one-percentage-point (annualised) monetary-policy tightening (quarterly s.d. 0.0025). Solid lines show simulations with Calvo price stickiness $\phi_p = 0.75$ and wage stickiness $\phi_w = 0.75$ (price & wage rigidity); dashed lines set $\phi_w = 0$ (price rigidity only). Black traces correspond to trend-inflation $\bar{\pi} = 0\%$, blue to $\bar{\pi} = 4\%$. Legend: P&W = price and wage rigidity | P only = price rigidity only.
Source: Author's Dynare simulations.

When wages are flexible, output falls by less and recovers faster. Without wage stickiness, firms can adjust relative wages to reflect weaker demand, allowing labour to reallocate smoothly and limiting the misallocation channel identified earlier.

Inflation, by contrast, drops more sharply in the price-only case. With flexible wages, prices fully adjust to the monetary policy shock, so nominal and real interest rates revert more quickly to baseline. When wages are sticky, part of the adjustment is absorbed through slower wage responses, which dampens the disinflation, resulting in a more persistent output contraction.

This difference in adjustment speed carries through the model. The composite price–wage dispersion term declines more rapidly when only prices are rigid, whereas with both frictions, dispersion response is comparatively dampened, reflecting the added wage-dispersion inefficiency. Consequently, the debt-to-GDP ratio rises more persistently in the joint-rigidity case, as weaker output and slower disinflation delay fiscal stabilisation.

Figure 15: Monetary-policy shock with increasing labour market rigidity

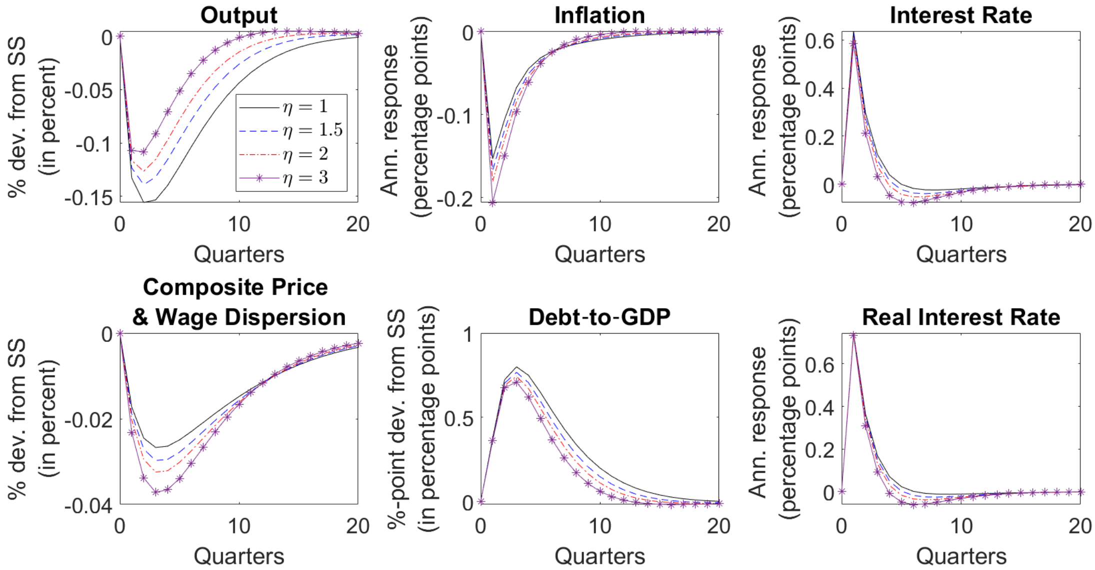

Notes: Impulse responses over 20 quarters to a one-percentage-point (annualised) monetary-policy tightening (quarterly s.d. 0.0025). Simulations show results for higher inverse Frisch elasticity values: $\eta = \{1, 1.5, 2, 3\}$.
Source: Author's Dynare simulations.

Overall, wage rigidity reduces the efficiency of the labour market, weakens the policy transmission to inflation, and amplifies the real costs of monetary tightening—especially under higher trend inflation.

Figure 15 presents impulse responses to a one-percentage-point monetary-policy tightening, simulated at a trend inflation rate of $\bar{\pi} = 4\%$, for progressively higher inverse Frisch elasticity values, $\eta = 1, 1.5, 2, 3$. Higher η reflects stronger labour disutility—a steeper labour–leisure trade-off and a weaker willingness of households to vary hours in response to real-wage changes.

Under positive trend inflation, this real rigidity amplifies the policy transmission to prices while cushioning output. As η increases, households reduce hours more sharply when real wages fall, tightening labour supply and intensifying the initial disinflation. Inflation therefore contracts more on impact, while output falls slightly less and recovers more quickly. The smaller and shorter-lived output contraction limits the rise in the debt-to-GDP ratio, and the composite price–wage dispersion term declines more rapidly as weaker labour adjustment dampens relative-price distortions.

Overall, at $\bar{\pi} = 4\%$, greater labour disutility strengthens the disinflationary impact of monetary policy but softens its output cost. Real rigidities therefore redistribute the adjustment burden from quantities to prices, reinforcing the earlier finding that labour-market characteristics shape the effectiveness of policy in economies with positive trend inflation.

5.2 Bayesian estimation

To explore the empirical relevance of our simulation findings, we estimate the model on South African data (2000Q1–2024Q2) using Bayesian methods under three alternative steady-state inflation assumptions: (i) zero trend inflation ($\bar{\pi} = 0$), (ii) a positive constant trend inflation (set at $\bar{\pi} = 6\%$), and (iii) a positive time-varying trend inflation. We focus on three key observables–real GDP, GDP-deflator inflation, and the 3-month Treasury bill rate–so that the model's information set aligns with the variables entering the Taylor rule. The baseline specification includes four structural shocks: technology, price-elasticity ('markup'), preference, and monetary policy; in the time-varying $\bar{\pi}$ specification we add an exogenous shock to the central bank's inflation target. We also estimate a variant that allows persistence in the monetary policy shock process (Model 4)[5]. This suite of estimations enables us to assess how a positive trend inflation rate, whether fixed or evolving, affects the model's fit and the inferred policy parameters in the Taylor rule.

Table 2 presents the prior distributions and posterior estimates of the structural parameters and shock processes for each model variant. Posterior estimates are obtained via a Metropolis-Hastings Markov Chain Monte Carlo (MCMC) procedure[6]. We run three chains of 200,000 draws each, discarding the first 50% of draws as burn-in, and adjust the Metropolis-Hastings jump scale parameter to achieve an acceptance rate of between a quarter and a third. For consistency across the different specifications, we fix the same prior means and variances for all parameters (as reported in Table 2, columns 2–4). The estimation achieves good convergence diagnostics, and the log marginal data densities (reported at the bottom of Table 2) indicate a better fit for the models that incorporate a positive time-varying trend inflation component.

Bringing the data to the model yields several insights about key parameters. The estimated Calvo price stickiness ϕ_p is generally lower once a positive trend inflation is introduced. For example, under a zero-inflation steady state ϕ_p falls between 0.71 and 0.85, whereas with $\bar{\pi} > 0$ it falls to within 0.65 and 0.76 (Table 2). This confirms that moderate trend inflation shortens effective price contract durations – firms reset prices more frequently when trend inflation is higher. In contrast, the degree of price indexation ζ_p remains low (around 0.13–0.4) and stable across all specifications. This consistently low indexation implies that backward-looking price-setting plays a limited role

[5]Model 4 specifies an AR(1) process for monetary policy disturbances, capturing the idea that policy decisions or signals in one quarter can carry over into subsequent periods even without new actions.

[6]While the posterior estimates reported here are locally identified, they are obtained using a standard Metropolis–Hastings MCMC routine and do not reflect a full exploration of the parameter space spanning both determinacy and indeterminacy regions. In practice, the combination of the model's non-linear structure, limited sample length, and the tendency of our estimation to favour highly persistent policy reaction functions and a relatively high degree of nominal price rigidity (including indexation) can lead to an over-reliance on shock and persistence mechanisms to match the data, complicating precise comparison of posterior parameter estimates (see Haque (2022)). Preliminary experiments estimating a small-scale linearised generalised New Keynesian model with trend inflation, using both US and South African data and a sequential Monte Carlo approach that jointly samples determinate and indeterminate solutions, appear more promising. A full implementation of this estimation strategy is left to future work focusing on disinflation policy.

in inflation dynamics, regardless of steady-state inflation assumptions. These results underscore how allowing for positive trend inflation in the model affects nominal rigidity estimates: it reduces the implied price stickiness but does not require additional indexation, meaning that trend inflation rather than indexation accounts for most of the inflation persistence in the data.

Turning to the monetary policy parameters, we find that accounting for trend inflation significantly alters the estimated Taylor rule coefficients. In particular, the inflation-response coefficient ϕ_π is higher once a positive $\bar{\pi}$ is internalised. Under the conventional zero-trend inflation model, the posterior mean of ϕ_π is around 1.8, whereas with a 6% trend or an estimated time-varying trend, ϕ_π rises into the 2.0–2.4 range at the posterior mean (see Table 2). The 90% high posterior density interval for ϕ_π also shifts upward (spanning roughly 1.2 up to 2.9 across specifications), indicating a stronger stance against inflation: the SARB reacted more aggressively to inflation than one would estimate under the assumption of $\bar{\pi} = 0$. This more forceful policy response to inflation leaves correspondingly less room for output stabilisation in the policy rule. Indeed, the output weight ϕ_y in the Taylor rule tends to be slightly lower in the positive $\bar{\pi}$ estimations (the posterior mean of ϕ_y falls from about 0.44 to 0.25 in our best-fit specification with time-varying trend inflation), consistent with a strong relative shift in focus toward price stability when trend inflation is high.

The interest rate smoothing coefficient ρ_i remains high in all estimations, but its interpretation changes once trend inflation is included. In the model without trend inflation, we estimate an extremely persistent policy rate (posterior $\rho_i \approx 0.95$ - 0.96, near the upper bound of the prior). When the model instead allows for a positive trend inflation, the posterior mean is essentially unchanged (in one variant, it fell to roughly 0.75, though in our full sample estimations and using the Taylor rule specification (11) it remains around 0.95; see Table 2).[7] It is clear, however, that even with trend inflation, the South African policy rate is estimated to be quite inertial, reflecting genuine interest-rate smoothing behaviour alongside positive trend-inflation targeting.

Differences in the estimated shock processes further highlight the role of trend inflation. The model that assumes $\bar{\pi} = 0$ tends to attribute relatively large disturbances to certain structural shocks in order to match the data's variability. For instance, the standard deviation of the technology shock required to fit output and inflation in the zero-inflation model is about 6.1 (in percentage-point terms), whereas in the specifications that allow for a positive inflation target this falls to roughly 4.2–4.9 (Table 2, ϵ_a-shock). A similar pattern holds for other shocks. By allowing an explicit shock to the inflation target (with its own estimated persistence $\rho_{\pi^T} \approx 0.75$) in the time-varying $\bar{\pi}$ model, the estimation can capture low-frequency movements in inflation without resorting as much to large, persistent supply or demand shocks. This reallocation of explanatory power indi-

[7] We find that while the Taylor rule is generally well identified based on tests using Dynare's `identification` toolbox, the estimated parameters and identified shocks are *very sensitive* to the specification of the Taylor rule. We intend to explore this in more detail in forthcoming research and will include our alternative Taylor rule estimations in an online Appendix.

cates that omitting trend inflation can lead to misidentification: the model without $\bar{\pi}$ tries to soak up trend-driven inflationary dynamics by overstating shock magnitudes (although the persistence of shock processes remains stable across estimations). Including a trend inflation process yields more moderate shock estimates and a clearer separation between trend-driven inflation fluctuations and other supply or demand-side disturbances.[8]

Overall, these Bayesian estimation results conform to our earlier hypothesis that ignoring trend inflation in policy analysis paints an incomplete picture. Estimating the Taylor rule without a trend inflation term would understate the stance against inflation. In a high–trend-inflation environment, policymakers face a steeper output-inflation trade-off: the central bank must respond more decisively to inflation deviations and accept a smaller role for output stabilisation. This finding is consistent with a higher sacrifice ratio – more output loss is required to bring inflation back to target when trend inflation is elevated. Moreover, maintaining a higher $\bar{\pi}$ comes with continuously greater price dispersion in the economy, which erodes aggregate productivity and welfare. Although a full analysis of disinflation policies is beyond the scope of this paper (and is deferred to future work), our estimation evidence already highlights the costs associated with a higher trend inflation regime. The next section explores these dynamics further by comparing impulse responses across the different model specifications and running counterfactuals to illustrate the endogenous interactions between trend inflation and price dispersion.

[8] Interestingly, we find that price elasticity of demand (price markup) shocks are not well identified and contribute negligibly to the dynamics of variables. Instead, if we specify the markup shock in an ad hoc wedge in the Phillips curve (or optimal reset price in the nonlinear setup) markup shocks become important. This observation suggests that the new Keynesian Phillips curve is generally not well specified to capture the transmission of shocks from marginal costs to inflation deviations. We intend to explore this further in future research.

Table 2: Posterior estimates

	Prior distribution			NoPiT (Mod 1)		PiT (Mod 2)		PiTshockNoAR (Mod 3)		PiTshock (Mod 4)	
Parameter	Dist.	Prior mean	Prior SD	Post.	90% HPD	Post.	90% HPD	Post.	90% HPD	Post.	90% HPD
Households											
σ (inverse EIS)	invg	2.00	0.25	1.91	[1.55, 2.25]	1.85	[1.52, 2.16]	1.86	[1.52, 2.19]	1.86	[1.53, 2.18]
ω (non-Ricardians)	beta	0.10	0.05	0.02	[0.00, 0.03]	0.02	[0.00, 0.03]	0.02	[0.00, 0.03]	0.03	[0.00, 0.06]
Firms											
ϕ_p (Calvo)	beta	0.75	0.10	0.78	[0.71, 0.85]	0.71	[0.65, 0.76]	0.71	[0.66, 0.76]	0.75	[0.71, 0.80]
ζ_p (indexation)	beta	0.50	0.10	0.24	[0.13, 0.36]	0.27	[0.14, 0.40]	0.26	[0.13, 0.39]	0.22	[0.11, 0.32]
Monetary policy											
ρ_i (smoothing)	beta	0.75	0.10	0.96	[0.95, 0.97]	0.96	[0.95, 0.97]	0.96	[0.95, 0.97]	0.95	[0.93, 0.97]
ϕ_π (inflation weight)	norm	1.50	0.50	1.82	[1.23, 2.39]	2.02	[1.45, 2.56]	2.12	[1.52, 2.67]	2.36	[1.76, 2.93]
ϕ_y (output weight)	beta	0.50	0.10	0.44	[0.28, 0.59]	0.43	[0.27, 0.58]	0.42	[0.27, 0.57]	0.25	[0.12, 0.38]
Shock persistence parameter											
ρ_a (technology)	beta	0.75	0.10	0.82	[0.72, 0.93]	0.85	[0.77, 0.94]	0.86	[0.79, 0.94]	0.84	[0.72, 0.96]
ρ_p (price markup)	beta	0.75	0.10	0.75	[0.59, 0.91]	0.75	[0.59, 0.92]	0.75	[0.60, 0.91]	0.75	[0.59, 0.92]
ρ_v (preference)	beta	0.75	0.10	0.28	[0.18, 0.38]	0.27	[0.17, 0.36]	0.26	[0.17, 0.36]	0.29	[0.14, 0.39]
ρ_{rp} (risk premium)	beta	0.75	0.10	0.71	[0.56, 0.86]	0.69	[0.54, 0.84]	0.68	[0.54, 0.84]	0.69	[0.55, 0.85]
ρ_{π^T} (inflation target)	beta	0.75	0.05	—	—	—	—	0.75	[0.67, 0.82]	0.77	[0.68, 0.90]
$\rho_{i,\text{shock}}$ (MP shock)	beta	0.75	0.10	—	—	—	—	—	—	0.68	[0.58, 0.78]
Std. dev. of shocks (×100)											
ϵ_a (technology)	invg	1.00	Inf	6.10	[2.90, 9.53]	4.93	[2.90, 6.95]	4.60	[2.62, 6.49]	4.16	[1.87, 7.61]
ϵ_p (price markup)	invg	1.00	Inf	0.85	[0.24, 1.48]	0.95	[0.22, 1.83]	0.96	[0.23, 1.83]	0.76	[0.25, 1.26]
ϵ_v (preference)	invg	1.00	Inf	21.8	[17.1, 26.2]	21.0	[16.6, 25.2]	21.2	[16.8, 25.6]	19.8	[15.3, 28.3]
ϵ_i (monetary policy)	invg	1.00	Inf	0.16	[0.14, 0.18]	0.16	[0.14, 0.19]	0.15	[0.13, 0.18]	0.18	[0.14, 0.22]
ϵ_{rp} (risk premium)	invg	1.00	Inf	0.62	[0.23, 1.02]	0.62	[0.24, 1.00]	0.63	[0.25, 1.05]	0.58	[0.24, 0.94]
ϵ_{π^T} (inflation target)	invg	1.00	Inf	—	—	—	—	0.98	[0.29, 1.65]	0.69	[0.26, 1.30]
Log data density (MHM)				-395.11		-398.70		-397.48		-379.73	

Note: Values in brackets are 90% HPD intervals. All shock standard deviations have been multiplied by 100 and are expressed in percentage points. EIS = elasticity of intertemporal substitution.

Trend inflation and monetary policy dynamics Figures 16 and 17 present impulse response functions following an estimated monetary policy shock under the alternative steady-state trend inflation scenarios described in Table 2. The shaded areas indicate 90 percent highest posterior density intervals.

Figure 16 contrasts the baseline monetary policy responses under zero-trend inflation (Model 1) and fixed positive trend inflation at 6 percent (Model 2). Under positive trend inflation, a tightening of monetary policy–approximately a 50 basis point increase in annualised terms–raises the real interest rate more sharply on impact, resulting in a quicker and deeper disinflationary response. By contrast, in the zero-trend scenario, the nominal interest rate increase translates into a smaller immediate rise in the real interest rate, but with greater persistence. Consequently, the policy-induced decline in inflation is muted, more gradual, and prolonged.

Importantly, the weaker disinflationary response under zero-trend inflation carries a higher real cost. Output falls more substantially and recovers more slowly than in the positive-trend scenario. This divergence arises primarily from the endogenous role of price dispersion under positive trend inflation: as trend inflation increases, price dispersion becomes a first-order state variable, significantly influencing the monetary transmission mechanism. Under higher steady-state inflation, the contractionary shock reduces the price dispersion wedge notably—an adjustment mechanism absent from conventional zero-inflation models. Allowing for trend inflation therefore alters the balance between stabilising prices and real activity, underscoring a critical policy trade-off that conventional New Keynesian frameworks without trend inflation may underestimate.

Figure 17 explores how alternative Taylor rule specifications affect the transmission of monetary policy shocks. Both Models 3 and 4 incorporate a time-varying inflation target, which aligns with trend inflation in the steady state. Model 4 differs by allowing persistence in the monetary policy shock itself. The comparison highlights that introducing a time-varying inflation target (Model 3) does not drastically alter the baseline dynamics observed in Model 2, suggesting that the mere presence of target variability does not substantially change the monetary policy transmission channels.

However, when persistence is introduced into the monetary policy shock process (Model 4), the impact on the economy is amplified. Although the nominal interest rate moves minimally on impact, the persistence of monetary policy disturbances implies expectations-driven dynamics that sharply amplify the real interest rate response and the resulting fall in inflation. The contraction in output is moderated, however, by the substantial reduction in price dispersion that accompanies the shock. This dampening effect highlights the critical role of the price dispersion channel, which mitigates output losses associated with disinflationary monetary policy in environments characterised by positive trend inflation.

Figures 18 through 21 provide complementary counterfactual analyses to further illustrate these

Figure 16: Impulse responses for the estimated baseline models `Mod.1` and `Mod.2` with and without trend inflation

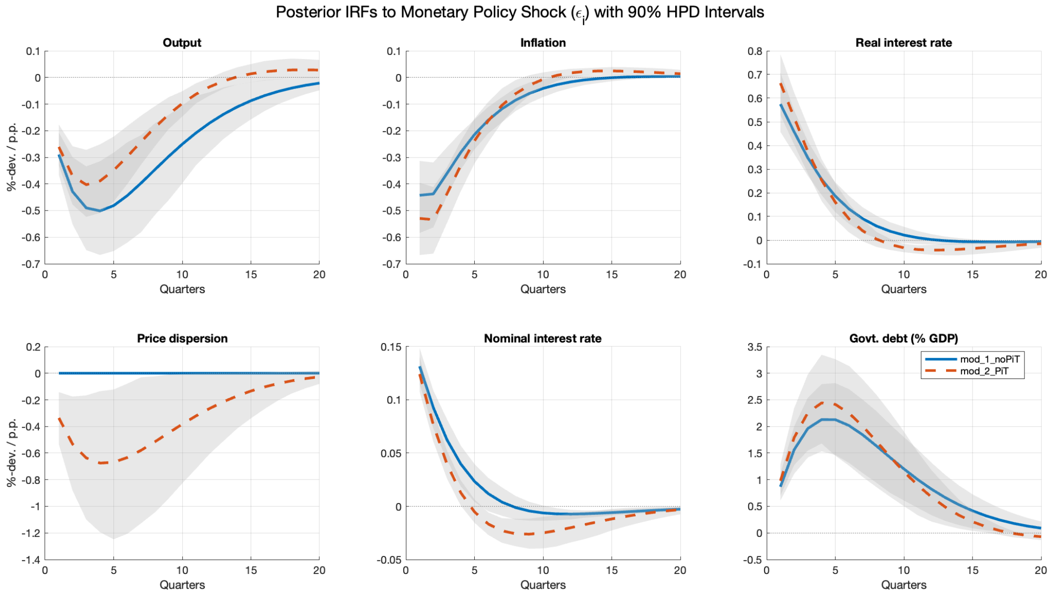

Figure 17: Impulse responses for the estimated baseline models `Mod.3` and `Mod.4` with and without trend inflation.

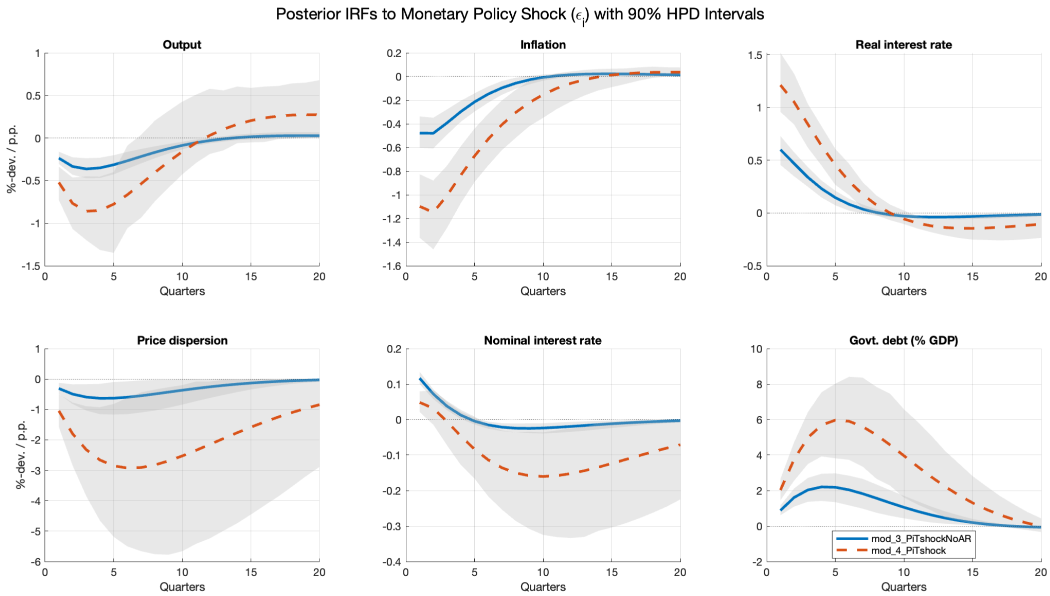

Figure 18: Impulse responses in the no-trend-inflation model (`Mod_1_noPiT`), where π^{trend} is restricted to zero during estimation compared to the counterfactual of adding positive trend inflation to a calibrated version of the estimated model.

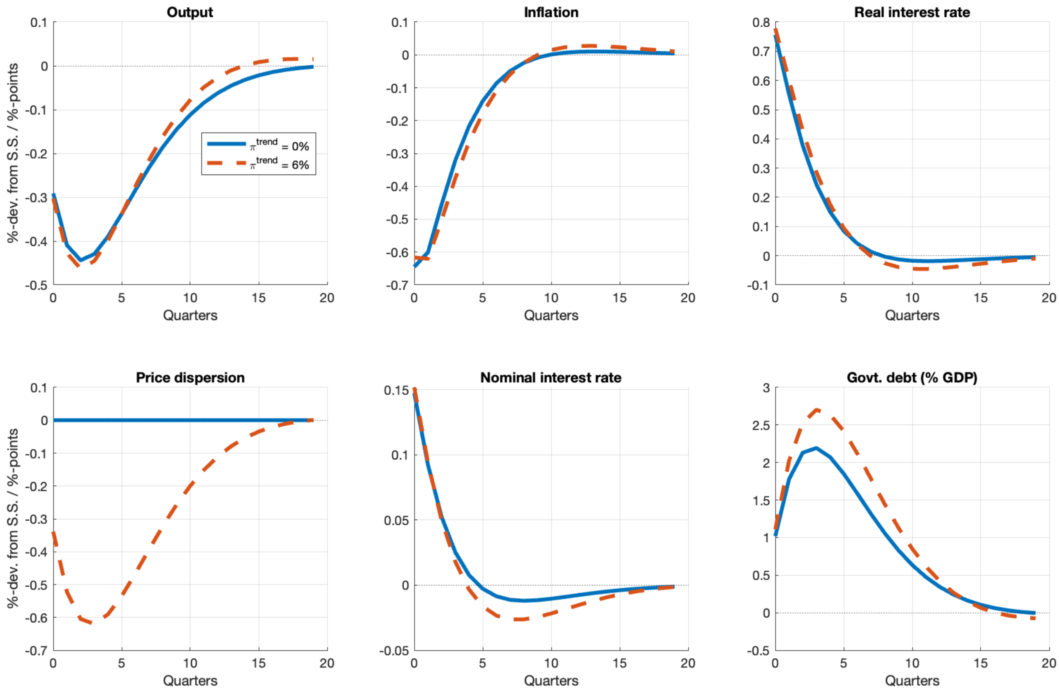

Figure 19: Impulse responses to a monetary policy shock in the baseline model with estimated trend inflation (`Mod_2_PiT`). The figure compares the dynamic responses under two trend inflation scenarios: $\pi^{\text{trend}} = 0\%$ and 6%.

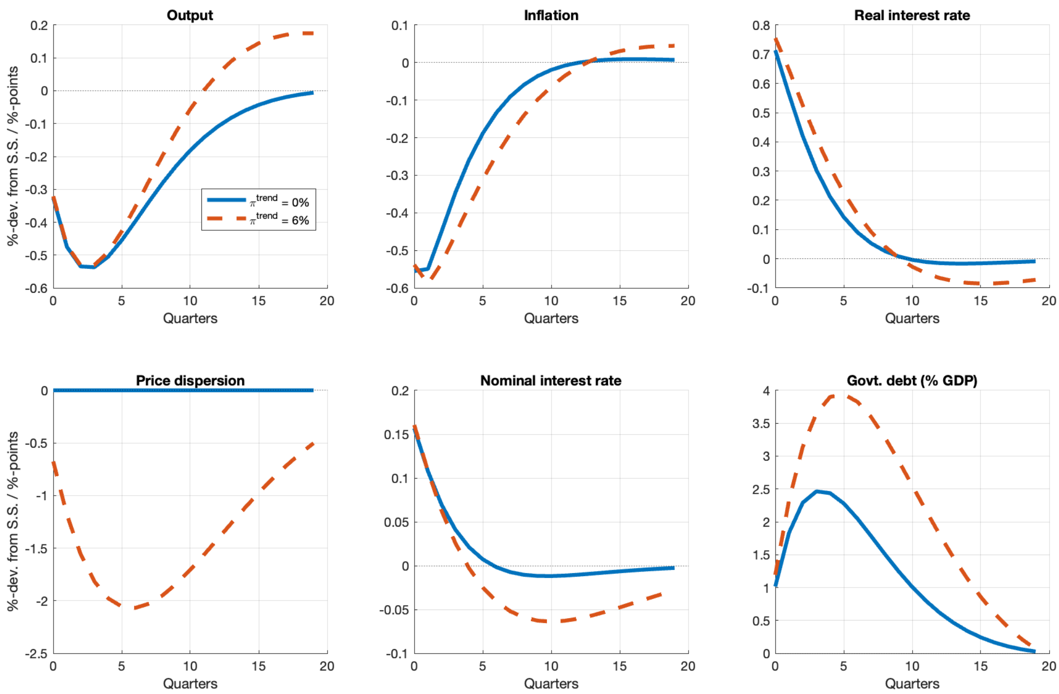

Figure 20: Impulse responses in the no-trend-inflation model (`Mod_3_PiTshockNoAR`), where π^{trend} is restricted to zero during estimation.

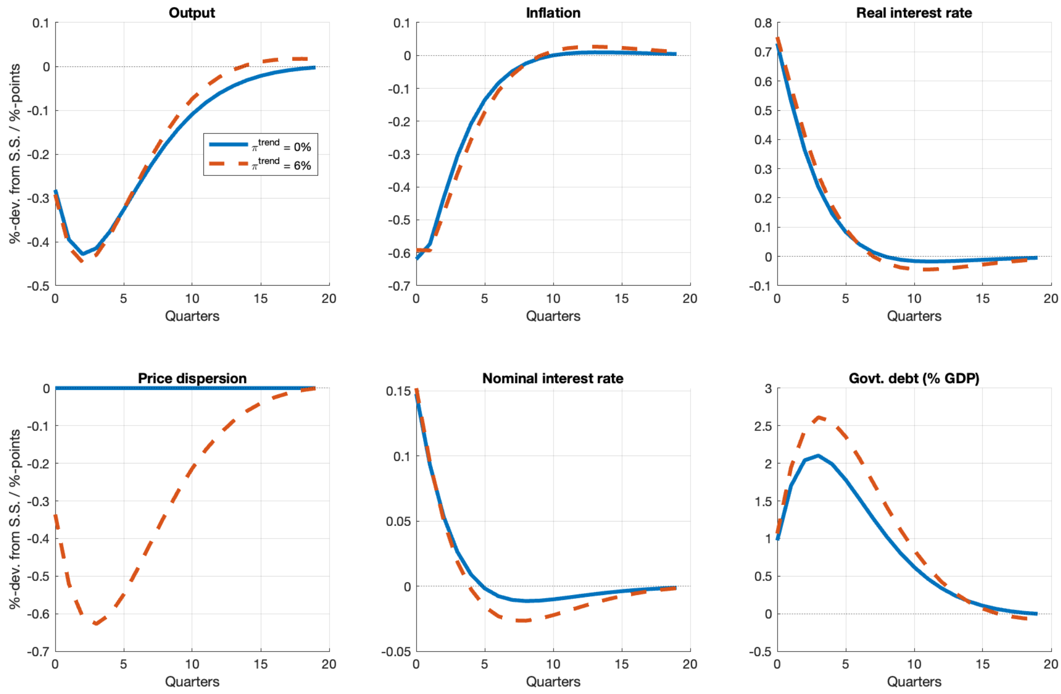

Figure 21: Impulse responses in the no-trend-inflation model (`Mod_4_PiTshock`), where π^{trend} is restricted to zero during estimation.

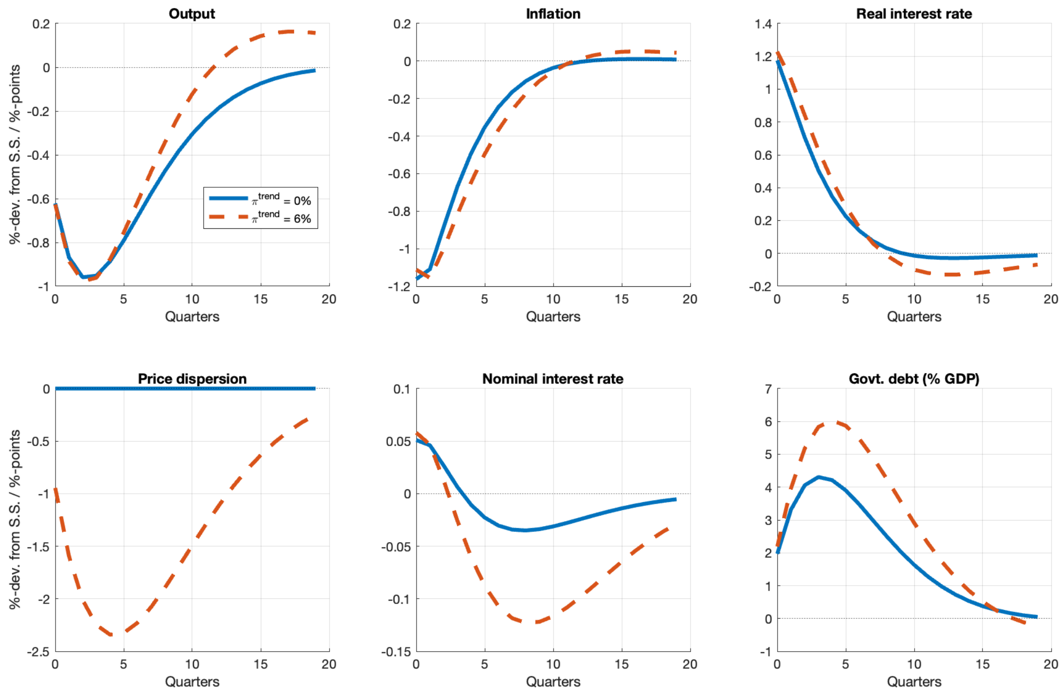

Figure 22: Time-varying inflation target based on Model 4.

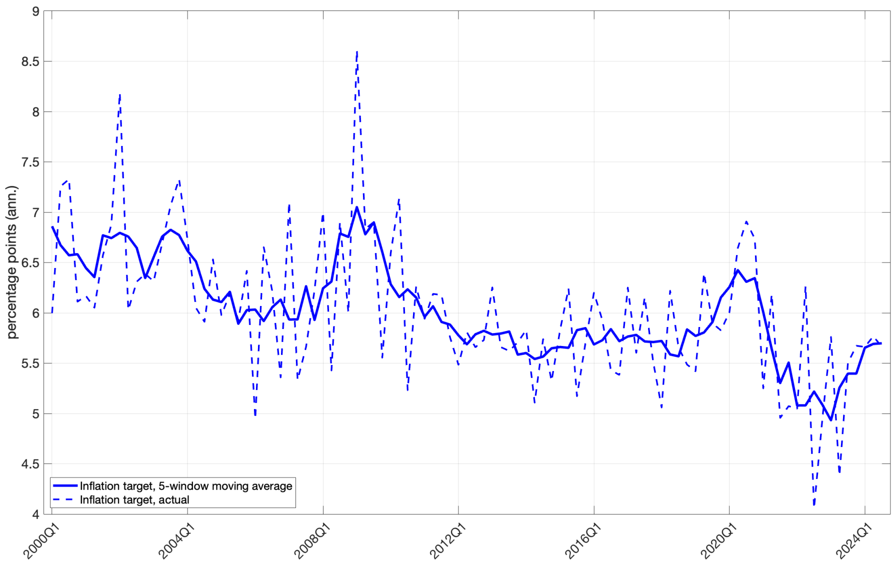

dynamics. Figure 18 contrasts the baseline zero-trend-inflation estimation (solid line) against a counterfactual scenario with positive trend inflation (dashed line). Conversely, Figures 19–21 depict the estimated scenarios with positive trend inflation ($\bar{\pi}$ = 6%) alongside their zero-trend-inflation counterfactuals. These comparisons consistently demonstrate that neglecting positive trend inflation results in underestimating the aggressiveness required of monetary policy and overstating the persistence needed to achieve disinflation.

Finally, Figure 22 plots the estimated time-varying inflation target implied by Model 4, which is broadly consistent with empirical findings in the literature (see, e.g., du Rand et al., 2023, and citations therein). At the onset of inflation targeting in South Africa (2000Q1), trend inflation comfortably exceeded the upper bound of the target band (6 percent). By the mid-2000s, trend inflation became anchored around this upper bound. In the wake of the 2008 global financial crisis, the SARB then appears to have tolerated higher inflation as the economy was also heading into the peak of a credit cycle as well as a price surge in commodities (see, e.g., Hollander et al., 2019; Hollander and Havemann, 2021). The ensuing recession in 2009 saw the implicit inflation target steadily declining, stabilising just below the upper bound around 2011 and remaining in the 5.5 to 6 percent range for much of the following decade. A modest upward drift started in 2019 but reversed sharply during the COVID-19 pandemic in 2020, falling briefly to around 5 percent. Only towards the end of the sample period (2024Q2) did trend inflation again reach pre-pandemic levels. Given recent inflation outcomes below 4 percent since mid-2024, this recent reversal may reflect mean-reversion dynamics, possibly indicating a shift towards a preference for lower inflation within the SARB's target band. However, a detailed assessment of this recent disinflationary stance is beyond this paper's scope and will be examined in forthcoming research.

In summary, incorporating positive trend inflation explicitly into a standard two-agent New Keynesian DSGE framework reveals important policy lessons that standard zero-trend models overlook:

- Elevated macroeconomic costs: higher steady-state inflation consistently generates greater price dispersion, which depresses steady-state output and welfare over the long run.
- Understated policy stance: Taylor-rule estimations ignoring trend inflation can underestimate the required aggressiveness of monetary tightening needed to stabilise inflation.
- Flatter effective Phillips curve: ignoring trend inflation can understate the sacrifice ratio and with a flatter Phillips curve at higher inflation, larger output gaps are needed to achieve the same disinflation.

For South Africa, explicitly modelling positive trend inflation is essential for accurately assessing monetary policy responses and their real-side implications—especially since trend inflation has been anchored for a long time near the upper bound of the 3–6 percent target band and now the lower bound of 3% is the official target (with a 1 percent tolerance band). Neglecting this dimension risks significant misjudgment of both the appropriate monetary policy stance and its associated real economic consequences, potentially forcing the SARB to forego greater output-stabilisation opportunities than implied by conventional New Keynesian frameworks.

6 Conclusion

This paper investigates the macroeconomic implications of positive trend inflation within a medium-scale New Keynesian fiscal DSGE model tailored to South Africa. Our analysis underscores the critical role of price dispersion as an economic channel that substantially magnifies the real costs of trend inflation. Even modest levels of positive trend inflation significantly increase resource misallocation, flatten the Phillips curve, and reduce the economy's flexibility to absorb shocks, thereby raising welfare losses and amplifying the sacrifice ratio associated with disinflation policies.

Bayesian estimations on South African quarterly data reveal that ignoring trend inflation in standard monetary policy formulations leads to systematic misinterpretations of both the required policy stance and its efficacy. Specifically, policy inertia is overstated, and the aggressiveness necessary to stabilise inflation is systematically underestimated. Allowing trend inflation to drift upwards with a target band exacerbates macroeconomic volatility, reduces steady-state output, and amplifies welfare costs.

From a policy perspective, these findings carry clear implications for central banks. Anchoring trend inflation closer to the lower end of a target band significantly reduces price dispersion, restores responsiveness of inflation to economic slack (a steeper Phillips curve), and improves the macroeconomic welfare trade-off. Conversely, tolerating higher trend inflation within an existing target band magnifies real resource misallocation, undermines policy effectiveness, and imposes larger welfare costs on the economy.

Future research could further enhance policy insights by incorporating explicit estimation of structural parameters and shocks in the labour market and examining in more detail the interaction of trend inflation with government debt dynamics. In forthcoming research, we examine the macroeconomic implications of moving from a 3–6% target band to a point target of 3% with a ± 1 percentage point tolerance, with emphasis on short- and long-run output–inflation trade-offs.

References

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

Bhatnagar, A. (2023). Monetary policy with non-ricardian households. *The Quarterly Review of Economics and Finance 89*, 12–26.

Blanchard, O. and J. Galí (2007). Real wage rigidities and the new keynesian model. *Journal of money, credit and banking 39*, 35–65.

Damjanovic, T. and C. Nolan (2010). Relative price distortions and inflation persistence. *The Economic Journal 120*(547), 1080–1099.

du Rand, G., H. Hollander, and D. van Lill (2023). A deep learning approach to estimation of the phillips curve in south africa. WIDER Working Paper Series wp-2023-79, World Institute for Development Economic Research (UNU-WIDER).

Eo, Y., L. Uzeda, and B. Wong (2023). Understanding trend inflation through the lens of the goods and services sectors. *Journal of Applied Econometrics 38*(5), 751–766.

Erceg, C. J., D. W. Henderson, and A. T. Levin (2000). Optimal monetary policy with staggered wage and price contracts. *Journal of monetary Economics 46*(2), 281–313.

Foresto, A., M. Reid, and J. Rakgalakane (2025). State dependence of the phillips curve what does this mean for monetary policy. Technical report.

Galí, J., J. D. López-Salido, and J. Vallés (2007). Understanding the effects of government spending on consumption. *Journal of the european economic association 5*(1), 227–270.

Garín, J., R. Lester, and E. Sims (2016). On the desirability of nominal gdp targeting. *Journal of Economic Dynamics and Control 69*, 21–44.

Haque, Q. (2022). Monetary policy, inflation target, and the great moderation: An empirical investigation. *International Journal of Central Banking*.

Havemann, R. and H. Hollander (2024). Fiscal policy in times of fiscal stress (or what to do when r>g). *Journal of Policy Modeling 46*(5), 1020–1054.

Hollander, H., R. Gupta, and M. E. Wohar (2019). The impact of oil shocks in a small open economy new-keynesian dynamic stochastic general equilibrium model for an oil-importing country: The case of south africa. *Emerging Markets Finance and Trade 55*(7), 1593–1618.

Hollander, H. and R. Havemann (2021). South africa's 2003–2013 credit boom and bust: Lessons for macroprudential policy. *Economic History of Developing Regions 36*(2), 339–365.

Hollander, H., R. Havemann, and D. Steenkamp (2024). The macroeconomics of establishing a basic income grant in south africa. *South African Journal of Economics 92*(1), 57–68.

Horn, A. J., L. Martin, J. H. Pretorius, and D. Steenkamp (2025, March). Gaps in the South African Inflation Targeting Debate. MPRA Paper 124010, University Library of Munich, Germany.

Jacome, L. I., N. E. Magud, S. Pienknagura, and M. Uribe (2025). The legacy of high inflation on monetary policy rules. Technical report, National Bureau of Economic Research.

Kemp, J. H. and H. Hollander (2020). A medium-sized, open-economy, fiscal DSGE model of South Africa. WIDER Working Paper Series wp-2020-92, World Institute for Development Economic Research (UNU-WIDER).

King, R. G. and A. L. Wolman (1996). Inflation targeting in a st. louis model of the 21st century.

Mohimont, J. (2022). Welfare effects of business cycles and monetary policies in a small open emerging economy. *Journal of Economic Dynamics and Control 136*, 104316.

Walsh, C. E. (2009). Inflation targeting: what have we learned? *International Finance 12*(2), 195–233.