Nonlinear Business-Cycle Anatomy*

Marco Brianti† Mario Forni‡ Luca Gambetti§ Antonio Granese¶

March 25, 2026

Abstract

Building on a frequency-domain identification within a nonlinear Structural Dynamic Factor Model, we study the nonlinear transmission of demand and supply shocks, the two shocks accounting for the bulk of fluctuations in U.S. macroeconomic variables. Supply shocks propagate symmetrically and are well described by linear dynamics. Demand shocks, by contrast, display strong sign asymmetries: contractionary shocks generate larger and more persistent declines in real activity, with limited adjustment of prices and nominal wages, an asymmetry amplified in booms. A New Keynesian model with downward nominal wage rigidity rationalizes these findings, highlighting the role of nominal rigidities as a source of nonlinearities.

JEL classification: C32, C51, E12, E32
Keywords: Frequency Domain, Nonlinear Time Series, Keynesian DSGE, Business Cycle, Economic Fluctuations

*We would like to thank Jacopo Bonchi, Vito Cormun, Luca Gemmi and seminar and conference participants at the Hitotsubashi University and Kyoto University for useful comments and suggestions.

†Università di Bologna; email: marco.brianti@unibo.it.

‡Università di Modena e Reggio Emilia, CEPR and RECent; email: mario.forni@unimore.it.

§Universitat Autonoma de Barcelona, Barcelona School of Economics, Università di Torino and Collegio Carlo Alberto; email: luca.gambetti@uab.es.

¶Università di Modena e Reggio Emilia and RECent; email: antonio.granese@unimore.it.

1 Introduction

The idea that business cycles exhibit nonlinear dynamics has a long standing tradition in macroeconomic research. Early contributions, such as Mitchell (1927), Keynes (1936), Burns and Mitchell (1946), and Friedman (1964) emphasize a fundamental asymmetry between expansions and recessions, with downturns typically unfolding more sharply and severely. Building on this insight, since the mid-1980s a large econometric literature has sought to test whether the statistical properties of macroeconomic time series align with this view. The evidence that emerges is broadly consistent with the earlier narrative: recessions tend to be sharper and deeper than expansions.[1] More recently, the literature has moved beyond reduced-form evidence to examine how macroeconomic fluctuations respond to identified structural shocks, with responses depending on the sign, magnitude, and the state of the economy. This work highlights pervasive asymmetries in the transmission of most macroeconomic disturbances.[2]

In parallel, a distinct literature has argued that the U.S. business cycle can be reliably described as arising from a small number of primitive driving forces. Originating from the pioneering work of Blanchard and Quah (1989), this line of research has progressively refined our understanding of the main sources of macroeconomic fluctuations. Angeletos et al. (2020) shows that a single macroeconomic shock can account for the bulk of business cycle fluctuations in the main macroeconomic variables.[3] Forni et al. (2025b), Granese (2025) and Avarucci et al. (2026) build on this line of work to show that two shocks, an

[1] See, among others, Neftci (1984), Falk (1986), DeLong and Summers (1986), Sichel (1989), Hamilton (1989), Diebold and Rudebusch (1990), Beaudry and Koop (1993), Tiao and Tsay (1994), Potter (1995), Kim and Nelson (1999), Morley and Piger (2012), and Dupraz et al. (2025).

[2] Evidence of such nonlinear responses has been documented for a range of structural shocks, including monetary policy shocks (Barnichon and Matthes, 2018; Tenreyro and Thwaites, 2016; Debortoli et al., 2025), financial shocks (Barnichon et al., 2022b; Forni et al., 2024), oil price shocks (Hamilton, 2003, 2011; Caravello and Martínez-Bruera, 2024; Forni et al., 2025a), government spending shocks (Auerbach and Gorodnichenko, 2012; Caggiano et al., 2015; Ramey and Zubairy, 2018; Barnichon et al., 2022a; Ben Zeev et al., 2023), and uncertainty shocks (Caggiano et al., 2022; Andreasen et al., 2024).

[3] Angeletos et al. (2020) uses a frequency-domain version of the maxshare approach of Faust (1998) and Uhlig (2004). See also Basu et al. (2021), Brianti and Cormun (2024), and Chahrour and Jurado (2022) for papers using a similar approach.

aggregate demand and an aggregate supply, provide a reliable and comprehensive characterization of U.S. business cycle fluctuations.[4] These analyses, however, are linear, and therefore silent on whether the main drivers of the cycle are themselves responsible for the nonlinear dynamics documented above.

The main contribution of this paper is to bring together these two strands of the literature by filling this gap, tracing the nonlinear dynamics of the business cycle back to its main drivers. In this sense, we extend the analysis of business cycle anatomy to a nonlinear framework.

To do so, we develop a nonlinear extension of the Structural Dynamic Factor Model (Stock and Watson, 2005; Forni et al., 2009), in which the factors—estimated from a dataset of 118 quarterly U.S. time series covering the period 1961-I to 2019-IV—depend on the structural shock of interest (either demand or supply) as well as its nonlinear functions. The large informational content of the DFM allows the identification of aggregate demand and supply shocks from the Wold representation of the factors, ignoring any nonlinear terms. Identification is implemented using frequency-domain techniques as in Forni et al. (2025b). Identification of the demand (supply) shock is achieved by maximizing (minimizing) the spectral covariance between inflation and real economic activity variables over business-cycle frequencies. Once the shocks are estimated, they are used, together with their nonlinear functions, as exogenous variables in a VARX for the factors.[5]

Guided by the literature discussed above, we first focus on sign asymmetry, testing whether a contractionary shock generates effects of comparable magnitude and persistence as an otherwise identical expansionary one. This is perhaps the most natural form of nonlinearity to consider, as it speaks directly to the early observation that recessions are sharper and more severe than expansions. This form of nonlinearity is in practice captured by augmenting the VARX specification with the absolute value of the shock of interest.

[4] Similar evidence has been documented for the euro area by Brignone and Mazzali (2025).
[5] The method is an extension to the Structural Dynamic Factor Model of the approach developed in Forni et al. (2024) and Debortoli et al. (2025) for Structral VARs.

We then enrich the model by allowing sign asymmetry to interact with the state of the economy, introducing a second source of nonlinearity. The propagation of a shock may thus differ not only across signs, but also across expansionary and recessionary periods. In other words, the difference between contractionary and expansionary shocks may itself depend on whether the economy is in a boom or a recession.

Our results point to a sharp contrast between supply and demand shocks. Supply shocks are an important driver of macroeconomic fluctuations, but their transmission is well described by linear dynamics: responses to positive and negative supply disturbances are largely symmetric. By contrast, demand shocks display strong and statistically significant asymmetries across a broad set of macroeconomic variables, leading us to reject the null of linear propagation. Positive demand shocks generate transitory expansions in real economic activity, broadly consistent with standard business-cycle theory. Negative demand shocks produce contractions that are substantially deeper and more persistent than the expansions triggered by an equally-sized positive shock. The asymmetry, however, cuts in opposite directions for quantities and prices: while real activity responds more strongly to contractionary shocks, prices and nominal wages respond less, exhibiting limited downward adjustment. Monetary policy also displays asymmetries consistent with those of prices and quantities: because prices adjust little following adverse shocks, inflation-based policy rules provide only modest accommodation when real activity contracts, allegedly amplifying the persistence of downturns. A stronger response of monetary policy to real activity, for example, through a rule that reacts more aggressively to output gaps during recessions, would likely mitigate both the depth and the persistence of these effects. Finally, sign asymmetries survive after controlling for the state of the economy. However, the difference between contractionary and expansionary shocks appears to be more pronounced when the economy is in a boom than in a recession.

To rationalize these findings, we develop a standard New Keynesian model augmented with downward nominal wage rigidity (DNWR), a natural departure from the textbook

framework in light of the limited downward adjustment of prices and nominal wages observed in the data. The key mechanism is simple: since nominal wages cannot fall, real wage adjustment depends on inflation reducing the real value of wages. Following expansionary shocks, real wages are under upward pressure and the constraint is naturally slack, leaving the economy's dynamics unaffected. The friction becomes relevant only in downturns, where real wages may need to fall but the DNWR constraint prevents the necessary nominal adjustment. In this context, inflation plays the "greasing–the–wheel" role emphasized by Olivera (1964) and Tobin (1972): by lowering the real value of wages, it substitutes for the nominal wage cuts that firms would otherwise wish to implement. Crucially, however, the degree of "grease" is not invariant to the source of the disturbance. An adverse demand shock lowers both output and prices, tightening the real-wage lower bound precisely when it is most needed, amplifying the contraction in quantities while dampening the response of prices. A negative technology shock, by contrast, generates inflationary pressure that loosens the real-wage bound, leaving the constraint slack and producing near-symmetric responses. The model thus replicates the key empirical pattern through a single, transparent mechanism.

We further validate the empirical strategy via a Monte Carlo exercise in which the model serves as the data-generating process, confirming that the econometric estimator reliably recovers the structural shocks and their nonlinear impulse responses.

2 Econometric methodology

In this section we present our empirical model and show how to estimate the nonlinear effects of the shocks. In what follows, we focus on a generic shock of interest, denoted by u_t^i.

2.1 A Nonlinear Structural Dynamic Factor Model

Let x_t be a n-dimensional vector of observable macroeconomic variables, transformed to obtain stationarity. The vector x_t is assumed to be part of an infinite-dimensional panel of

time series. Each variable x_{it}, $i = 1, \ldots, n$, is decomposed into the sum of two unobservable and mutually orthogonal components: a common component χ_{it} and an idiosyncratic component ξ_{it} (see Forni et al., 2009, for further details). The idiosyncratic components capture variable-specific or sector-specific shocks, as well as measurement errors. In contrast, the common components account for most of the co-movement among macroeconomic variables, as they are driven by the same $r < n$ latent static factors, $F_t = (F_{1t}, \ldots, F_{rt})'$. Formally, the observable vector x_t admits the representation

$$x_t = \chi_t + \xi_t = \Lambda F_t + \xi_t \tag{1}$$

where Λ is an $n \times r$ matrix of factor loadings. Equation (1) provides the static factor representation, in which the latent factors affect the common components contemporaneously. The dynamic nature of the model is introduced through a moving-average representation of the static factors, augmented by a nonlinear transformation of a structural shock of interest.

Structural representation. The r-dimensional vector F_t admits the following structural representation:

$$F_t = \nu + \Gamma(L)v_t + \gamma(L)u_t^i + \beta(L)g(u_t^i) \tag{2}$$

where ν is a vector of constants, u_t^i denotes the structural shock of interest and v_t is an m-dimensional vector containing other structural shocks. The term $g(u_t^i)$ is a contemporaneous nonlinear transformation of the shock of interest. In principle, $g(u_t^i)$ may be vector-valued and include multiple nonlinear transformations and may also depend on other shocks or lagged variables (see Debortoli et al., 2025).

We assume that the elements of $[v_t' \; u_t^i]'$ are mutually independent at all leads and lags, with zero mean and identity covariance matrix. The serial independence assumption implies that all structural shocks, including u_t^i, are uncorrelated with past values of both $g(u_t^i)$ and F_t. The polynomial $\gamma(L) = \gamma_0 + \gamma_1 L + \gamma_2 L^2 + \ldots$ collects the IRFs to the shock of in-

terest, while $\beta(L) = \beta_0 + \beta_1 L + \beta_2 L^2 + \ldots$ collects the IRFs associated with the nonlinear transformation of the shock. Finally, $\Gamma(L) = \Gamma_0 + \Gamma_1 L + \Gamma_2 L^2 + \ldots$ is an $r \times m$ matrix of IRFs associated with the remaining structural shocks. Combining equations (1) and (2), the common component can be written as

$$\chi_t \quad = \quad \nu^* + \Gamma^*(L) v_t + \gamma^*(L) u_t^i + \beta^*(L) g(u_t^i),$$

where $\nu^* = \Lambda\nu$, $\Gamma^*(L) = \Lambda\Gamma(L)$, $\gamma^*(L) = \Lambda\gamma(L)$ and $\beta^*(L) = \Lambda\beta(L)$. Our object of interest is the effect of the structural shock u_t^i on the common component χ_t. The total effect of a shock of interest $u_t^i = \bar{u}^i$ can be obtained by combining the two terms $\gamma^*(L)$ and $\beta^*(L)$ as:

$$IRF(u_t^i = \bar{u}^i) = \gamma^*(L)\bar{u}^i + \beta^*(L) g(\bar{u}^i). \tag{3}$$

VARX Representation. Turning back to equation (2), we assume that the term $\Gamma(L)v_t$, which is not necessarily invertible, has an invertible Wold representation, $\Gamma(L)v_t = \Phi(L)\epsilon_t$. By inverting $\Phi(L)$, vector F_t also admits a VARX representation:

$$C(L)F_t = \delta + C(L)\gamma(L)u_t^i + C(L)\beta(L)g(u_t^i) + \epsilon_t$$

where $C(L) = \Phi(L)^{-1}$, and rearranging

$$F_t = \delta + \tilde{C}(L)F_{t-1} + \tilde{\gamma}(L)u_t^i + \tilde{\beta}(L)g(u_t^i) + \epsilon_t \tag{4}$$

where $\delta = C(1)\nu$, $\tilde{C}(L) = C_1 L + \ldots + C_p L^p$, $\tilde{\gamma}(L) = C(L)\gamma(L)$, and $\tilde{\beta}(L) = C(L)\beta(L)$.

VAR Representation. Since the factors are stationary, F_t admits a Wold representation. We assume that this representation is invertible, so that F_t can be written as a (possibly infinite-order) VAR:

$$A(L)F_t = \mu + e_t \tag{5}$$

where $\mu = A(1)\,\mathbb{E}(F_t)$, the vector $e_t \sim WN(0, \Sigma_e)$, $A(L)$ is an $r \times r$, stable polynomial matrix, with $\mathbb{E}(\cdot)$ denoting the unconditional expected value operator. Equation (5) rules out the presence of unit roots in the Wold representation of the factors. We test this condition empirically and find no evidence of unit roots.[6] This result supports the existence of a stable VAR representation for the factors and validates the conventional DFM specification adopted in the analysis.

Relationship between the VAR and VARX representations. To clarify the relationship between equations (4) and (5), consider projecting the term $\tilde{\gamma}(L)u_t^i + \tilde{\beta}(L)g(u_t^i)$ onto the space spanned by a constant and the past history of F_t:

$$\tilde{\gamma}(L)u_t^i + \tilde{\beta}(L)g(u_t^i) = \theta + B(L)F_{t-1} + \omega_t.$$

Substituting this projection into equation (4) yields the VAR representation

$$F_t = (\delta + \theta) + [\tilde{C}(L) + B(L)]F_{t-1} + \epsilon_t + \omega_t.$$

Hence, in equation (5) we have

$$A(L) = I - \tilde{C}(L) - B(L), \qquad \mu = \delta + \theta, \qquad e_t = \epsilon_t + \omega_t.$$

While the structural shock can still be recovered by imposing appropriate identification restrictions, the associated nonlinear dynamics are absorbed into the reduced-form residuals and therefore remain hidden in the VAR specification. This highlights the importance of estimating the VARX model, which allows the nonlinear effects to be explicitly identified and tested.

[6] This condition can be assessed empirically by verifying that the spectral density matrix of F_t has full rank at frequency zero as suggested in Barigozzi et al. (2021). To this end, we apply the Dynamic eigenvalue Difference Ratio (DDR) estimator of Avarucci et al. (2026) over the full frequency range $[0, \pi]$ and at zero frequency. In both cases, the estimator detects two dynamic shocks, with no evidence of rank reduction at frequency zero. This conclusion is further confirmed by the criterion of Hallin and Liška (2007), applied at zero frequency and across all frequencies.

2.2 Identification

Our main identifying assumption is that the VAR in equation (5) is informationally sufficient for the shock of interest. In other words, u_t^i is fundamental, that is, it can be recovered as a linear combination of the VAR residuals:

$$u_t^i = w_i' e_t, \tag{6}$$

where w_i is the vector combining the reduced-form residuals into the structural shock of interest, which also corresponds to the linear part of the impact effect of u_t^i. While this assumption is demanding in a low-dimensional SVAR framework, it is considerably weaker in the rich information environment of a high-dimensional factor model (Bernanke et al., 2005; Forni et al., 2009; Forni and Gambetti, 2010; Forni et al., 2025c).

Identification of u_t^i and the associated IRFs requires selecting the vector w_i. To do so, we use the two-steps frequency-domain approach developed by Forni et al. (2025b). In the first step, the procedure extracts two reduced-form shocks that account for the largest share of variance in key macroeconomic aggregates at frequencies longer than 6 quarters, thereby filtering out short-run fluctuations. The target set includes the variances of the main real and nominal variables (GDP, consumption, investment, the unemployment rate, inflation, and the federal funds rate).[7] In the second step, these reduced-form shocks are then orthogonally rotated to recover economically interpretable structural shocks. The demand shock is identified as the disturbance that maximizes the covariance between inflation and GDP growth at business-cycle frequencies, while the supply shock minimizes

[7] In this application, the target set differs slightly from that used in Forni et al. (2025b), reflecting the specific focus of this study on business-cycle fluctuations. Importantly, the results are not driven by this choice. As a robustness check, the analysis is replicated using alternative identification targets, namely, the benchmark specification in Forni et al. (2025b) and a single-variable target based on GDP, as in Angeletos et al. (2020). Across all specifications, the impulse response functions display the same patterns, indicating that the findings are robust to the choice of target.

this covariance. By construction, the two shocks turn out to orthogonal (see Forni et al., 2025b for details).

2.3 Estimation

Model estimation unfolds as follows. First, we estimate the factors F_t using the first r principal components and the factor loadings Λ as the corresponding eigenvectors. Second, we estimate a VAR for the factors using OLS and identify the demand and supply shock as discussed above. Third, we include the estimated structural shock and its nonlinear function as exogenous regressors in equation (4) and estimate the coefficients $C(L)$, $\tilde{\gamma}(L)$, and $\tilde{\beta}(L)$ using OLS.[8] The impulse response functions (IRFs) associated with the linear and nonlinear components are then obtained from the equations $\gamma^*(L) = \Lambda C(L)^{-1}\tilde{\gamma}(L)$ and $\beta^*(L) = \Lambda C(L)^{-1}\tilde{\beta}(L)$ using the corresponding estimates. Finally, the total effect of a shock $\bar{u}^i$ is obtained using equation (3).[9]

3 Nonlinear Propagation of Business-Cycle Shocks

In this section, we investigate whether the propagation of the main business-cycle shocks exhibits nonlinearities. Section 3.1 introduces the data and model specification, while Section 3.2 tests for sign asymmetries in impulse response functions. Section 3.3 then quantifies their importance using frequency-domain variance and historical decompositions. Section 3.4 turns to joint sign- and state-dependence, and Section 3.5 assesses the robustness of the results to alternative specifications.

3.1 Data and Model Specification

For the empirical analysis, we use the quarterly macroeconomic dataset developed in Granese (2025), which is also used in Forni et al. (2025b). The dataset consists of 118

[8] To keep the model tractable and for the sake of parsimony we estimate two separate VARX with the same endogenous variables for the demand shock and for the supply shock.

[9] The confidence bands are obtained by treating u_t^i as an observed variable, since considering the uncertainty related to the first step in the context of the factor model would be complicated.

U.S. time series covering the period 1961:I-2019:IV. All series are transformed to reach stationarity; as argued above, this does not entail cointegration problems. The complete list of variables and transformations is provided in Appendix B.

We set a value r of the static factors, using the criterion of Bai and Ng (2002), with the penalty adjustment proposed by Alessi et al. (2010), finding a number of static factors $\hat{r} = 11$. The shock of interest is estimated using a VAR(4) on the estimated factors. The VARX is estimated using four lags as well.

In the robustness section, we take into account the uncertainty in estimating the number of static factors, and repeat the analysis with different specifications of $\hat{r}$. We also repeat the analysis using 3, 5, and 6 lags for both the VAR and the VARX.

3.2 Sign Asymmetries

We investigate whether the propagation of the main business-cycle shocks depends on their sign. To do so we set $g(u_t^i) = |u_t^i|$, the absolute value of the shock. Figures 1 and 2 show the impulse responses to supply and demand shocks, respectively. For each shock, the first column reports responses to a positive shock, the second to a negative shock (multiplied by minus one for ease of comparison), and the third their difference. The latter captures potential sign asymmetries: if the responses to shocks of opposite sign coincide, the difference is zero, whereas deviations from zero indicate sign dependence. Solid lines report point estimates, while the shaded areas denote the 68% and 90% bootstrap confidence bands.

Supply shocks. Figure 1 reports the sign-dependent impulse responses of output, prices, the unemployment rate, and the federal funds rate to a supply shock. Overall, the responses are consistent with the conventional interpretation of supply shocks: real activity expands after positive shocks and contracts after adverse ones, while prices move in the opposite direction (see Forni et al., 2025b, for linear evidence). However, it is particularly noteworthy that the responses are highly symmetric. For output and prices, the differ-

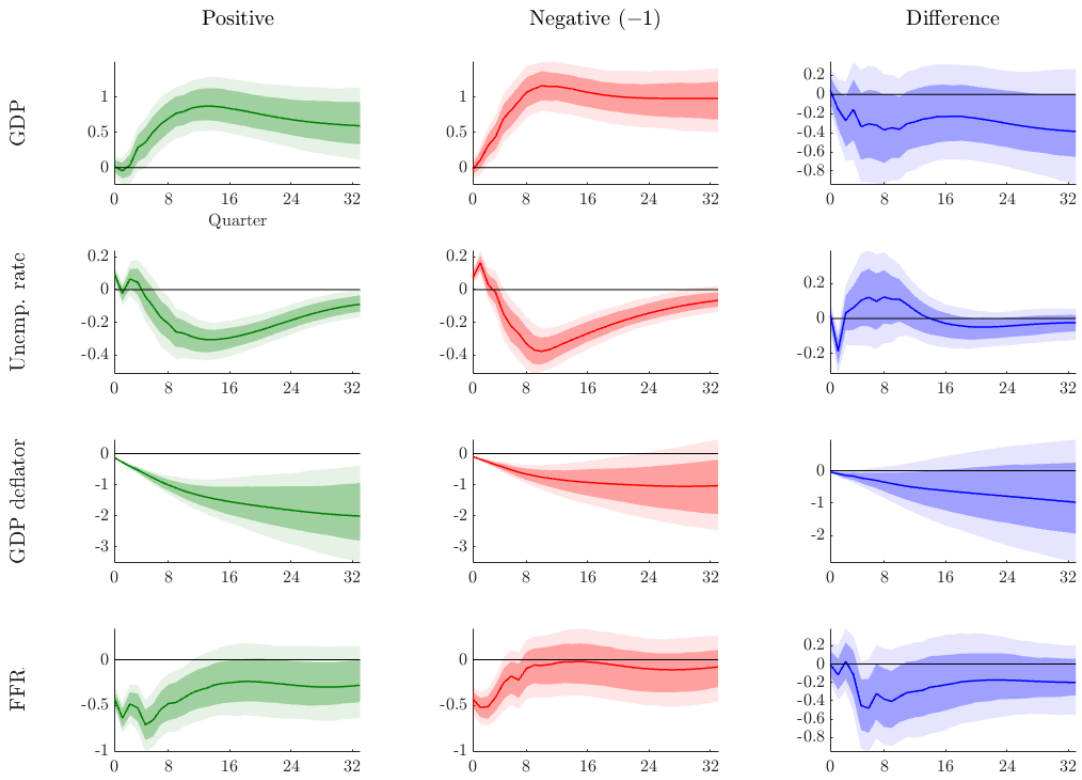


Figure 1: Sign-dependent impulse responses to supply shocks

Notes. Sign-dependent impulse responses to a supply shock. The first column reports responses to a positive shock, the second column reports responses to a negative shock (multiplied by minus one to ease comparison), and the third column reports the difference between the two responses. Solid lines denote point estimates, while shaded areas represent 68% and 90% bootstrap confidence bands.

ence between positive and negative shocks is never statistically significant across horizons. The unemployment rate exhibits a small and short-lived asymmetry on impact that quickly fades, while for the FFR the difference becomes briefly significant at the 68% level around the sixth quarter. Taken together, these results imply that we cannot reject the null of linearity for supply disturbances.[10]

Demand shocks. The picture changes for demand shocks. Figure 2 reports the sign-dependent impulse responses of output, prices, the unemployment rate, and the federal

[10] Figure A1 in the Appendix reports the responses of additional real variables, including consumption and investment, together with other labor-market indicators. In all cases, the difference between positive and negative shocks is not statistically significant, confirming the absence of sign dependence for supply disturbances.

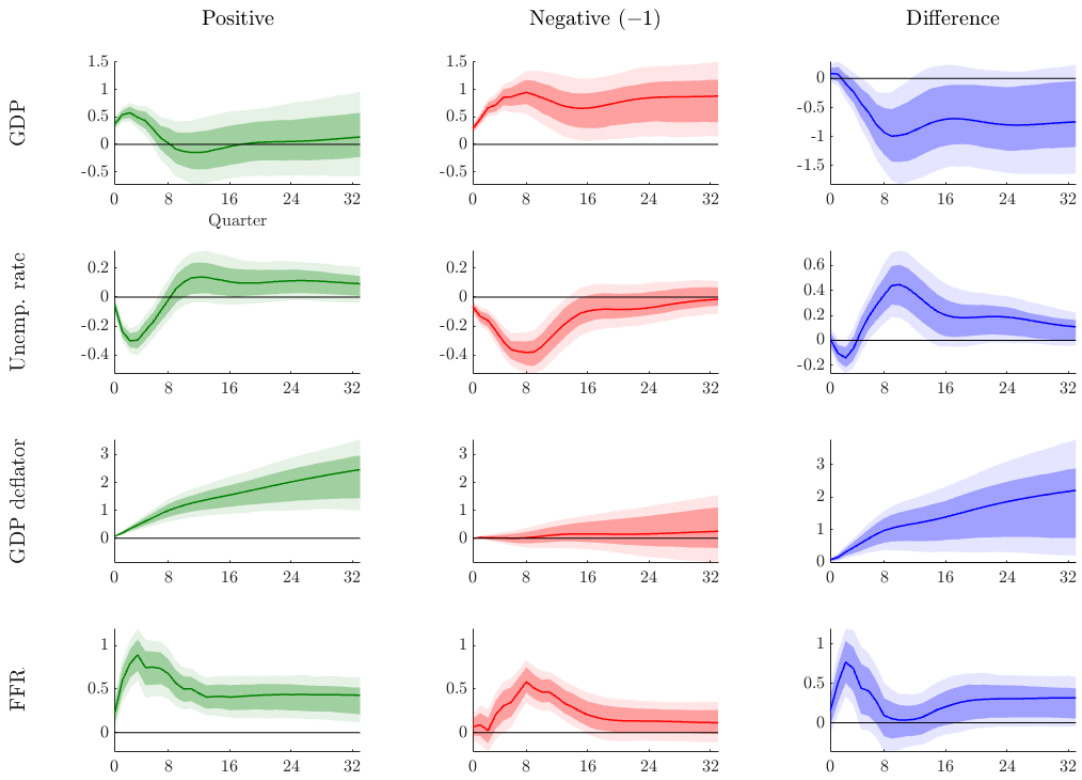


Figure 2: Sign-dependent impulse responses to demand shocks

Notes. Sign-dependent impulse responses to a demand shock. The first column reports responses to a positive shock, the second column reports responses to a negative shock (multiplied by minus one to ease comparison), and the third column reports the difference between the two responses. Solid lines denote point estimates, while shaded areas represent 68% and 90% bootstrap confidence bands.

funds rate. In this case, the null of linearity is clearly rejected, as the propagation mechanism depends on the sign of the shock. The difference between positive and negative shocks remains statistically significant for all the variables across several horizons. For real activity, contractionary shocks generate substantially larger and more persistent effects than expansionary ones. Positive shocks generate conventional, short-lived responses: GDP rises on impact, reaches a peak of approximately 0.5% after a few quarters, and gradually returns to baseline within two years. Negative shocks, by contrast, lead to deep and persistent contractions. Following a negative shock, GDP reaches a trough close to -1% after about two years and remains statistically different from zero at long horizons.

The dynamics of the unemployment rate are mirrored: it rises sharply and persistently after contractionary shocks, while its decline following positive shocks is smaller and less persistent. This pattern, which also emerges in the responses of consumption, investment, hours worked, and the employment ratio (Figure A2), echoes the hysteresis dynamics documented by Furlanetto et al. (2025) and Carnevale and Di Francesco (2025), although in our setting persistence is driven exclusively by adverse demand shocks.

Nominal variables display the opposite asymmetry. Expansionary demand shocks generate sustained inflationary pressures, whereas prices show little response to contractionary shocks. A similar pattern emerges for nominal wages (Figure A2), which rise persistently following positive shocks but exhibit little and statistically insignificant adjustment after negative ones, consistent with the widespread view that nominal variables are subject to downward rigidities. The federal funds rate moves accordingly, reacting more strongly to positive shocks than to negative ones. This result points to a more muted monetary policy response to contractionary demand shocks. While this may partly reflect the effective lower bound, it is also consistent with a policy rule that places greater weight on inflation stabilization than on real activity. As price adjustments are muted following adverse shocks, a rule focused mainly on inflation stability provides limited accommodation to real activity, even in the presence of sizable output losses.

Summary. The evidence points to a central role of downward nominal rigidities in shaping the asymmetric transmission of demand shocks. Following positive demand shocks, adjustment operates through both quantities and prices: output expands, prices and wages rise, and monetary policy tightens accordingly. Following negative shocks, by contrast, prices and nominal wages exhibit little downward adjustment, leaving quantities to absorb most of the response. That is, firms reduce hours and employment rather than labor costs, unemployment rises persistently, and real activity experiences a deeper and more prolonged contraction. Because inflation exhibits a muted response, a policy rule focused primarily on inflation stabilization provides limited accommodation, so that the policy

rate declines only modestly despite substantial output losses. It follows that a stronger response of monetary policy to real activity would better cushion the adverse effects of a contractionary demand shock (Debortoli et al., 2025).

In contrast, supply shocks remain largely symmetric and are well described by linear dynamics. Nonlinear features of U.S. business-cycle fluctuations therefore originate primarily from demand disturbances.

3.3 Quantifying Sign Asymmetries

Variance decomposition. To quantify the importance of sign-dependent nonlinearities, we apply a frequency-domain variance decomposition that separates each shock into linear and nonlinear components. Using the estimated impulse responses, we estimate the share of macroeconomic variance explained by each component over the business-cycle band ($2\pi/32 \leq \omega \leq 2\pi/6$) and the long-run band ($0 \leq \omega \leq 2\pi/80$).[11] Table 1 reports the results.

Panel A shows that the nonlinear component of supply shocks explains a negligible share of variance at both business-cycle and long-run frequencies. Across all variables, the nonlinear contribution does not exceed 1–2 percent of total variance, indicating that sign-dependent nonlinearities play essentially no role in the transmission of supply shocks. At the same time, supply shocks account for a substantial share of macroeconomic fluctuations. At long-run frequencies, they explain more than half of the variance of GDP, consumption, investment, hours, and unemployment. At business-cycle frequencies, they account for roughly one third of the variance of GDP, investment, hours, and unemployment, and close to half of consumption variance, consistent with the linear analysis of Forni et al. (2025b). Overall, while supply shocks are a major driver of macroeconomic fluctuations, their propagation is effectively linear.

Panel B presents the corresponding decomposition for demand shocks. Three findings stand out. First, unlike supply, demand shocks display a sizable nonlinear component. At

[11] We assume zero covariance between the linear and nonlinear components, as their empirical correlation is negligible. The explained variance in the numerator is computed using the VARX-based impulse responses, while total variance in the denominator is obtained from the DFM impulse responses.

Table 1: Frequency-domain variance decomposition: demand and supply

	Business cycle			Long run		
	Linear	Nonlinear	Total	Linear	Nonlinear	Total
Panel A: Supply shock						
GDP	28.2	0.9	29.1	62.9	0.6	63.5
Consumption	48.9	0.5	49.4	74.8	0.3	75.1
Investment	32.3	1.1	33.4	57.5	0.9	58.4
GDP deflator	58.2	1.0	59.2	44.2	1.3	45.5
Hours	29.8	0.8	30.6	53.5	0.5	54.0
Unemp. rate	32.5	1.5	34.0	67.1	0.4	67.5
FFR	39.7	2.0	41.7	22.1	2.0	24.1
Panel B: Demand shock						
GDP	43.4	6.1	49.5	11.8	5.3	17.1
Consumption	28.1	13.2	41.3	12.3	7.3	19.6
Investment	40.7	8.0	48.7	15.3	6.1	21.4
GDP deflator	22.3	10.7	33.0	22.9	6.6	29.5
Hours	36.0	6.6	42.6	17.2	7.5	24.7
Unemp. rate	33.7	8.2	41.9	19.2	7.4	26.6
FFR	34.4	7.9	42.3	43.0	3.3	46.3

Notes. Frequency-domain variance decomposition of supply (Panel A) and demand (Panel B) shocks. Entries report the percentage share of variance explained by the linear component, the nonlinear component, and their sum (Total) over the business-cycle band ($2\pi/32 \leq \omega \leq 2\pi/6$) and the long-run band ($0 \leq \omega \leq 2\pi/80$). The nonlinear component captures sign-dependent effects.

business-cycle frequencies, the nonlinear component explains 6.1% and 13.2% of the variance of GDP and consumption, respectively, with notable effects also on prices (10.7%) and labor-market variables (6.6%–8.2%). These contributions persists at long-run frequencies, ranging from 5.3% for GDP to 7.5% for hours worked. Second, once the linear and nonlinear components are combined, demand shocks emerge as the main drivers of business-cycle fluctuations, explaining roughly half of the variance of GDP and investment and more than 40% of that of hours worked and unemployment. Third, allowing for sign-dependent nonlinearities reveals a substantially larger long-run role for demand shocks than suggested by the linear model in Forni et al. (2025b), with the total long-run contribution of demand reaching about 17% for GDP growth and about 20% for consumption and investment.

Historical decomposition. We next examine the time-series contribution of demand shocks through a historical decomposition of the unemployment rate. Figure 3 reports

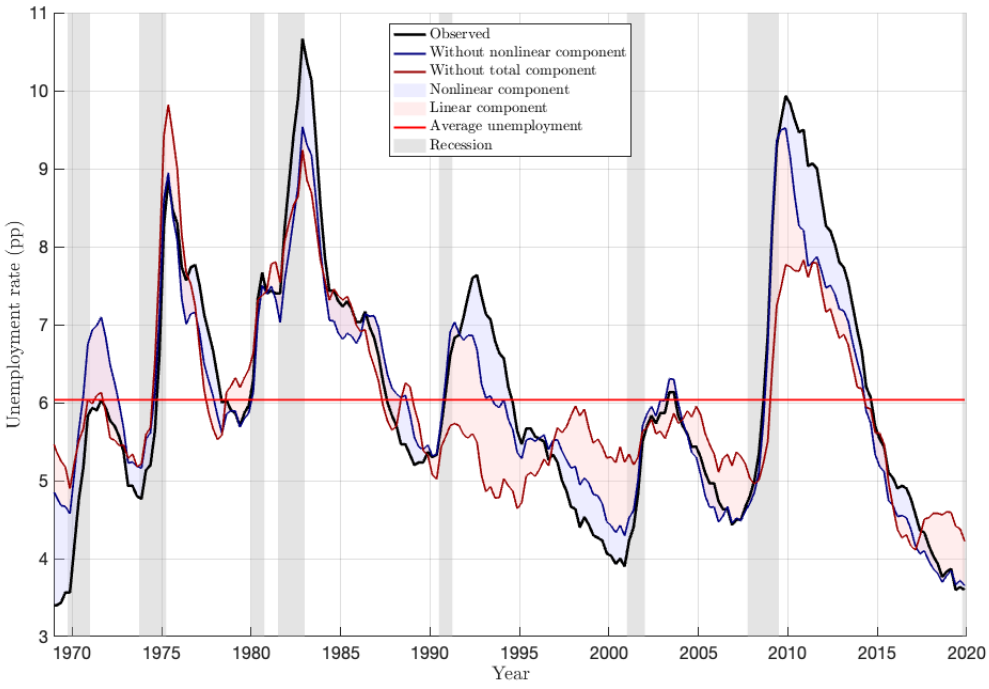


Figure 3: Historical decomposition of the unemployment rate

Notes. Historical decomposition of the unemployment rate. The black line denotes observed unemployment. The blue line removes only the nonlinear component of the demand shock, while the dark red line removes the entire demand contribution. The blue shaded area represents the nonlinear demand component, defined as the gap between observed unemployment and the counterfactual excluding only the nonlinear part. The red shaded area represents the remaining linear demand contribution, defined as the gap between the two counterfactual paths. The horizontal red line marks the 6 percent unemployment threshold separating low- and high-unemployment regimes.

the results. Based on the estimated impulse responses, we trace out the period-by-period contribution of the identified demand shock series and decompose it into a linear and nonlinear component. This procedure yields two counterfactual paths for unemployment: one that removes only the nonlinear demand component, shown by the blue line, and one that removes the full demand contribution, shown by the dark red line. The latter, therefore, captures unemployment dynamics driven by all other shocks, mainly supply disturbances. The blue shaded area measures the nonlinear demand contribution, defined as the gap between observed unemployment and the counterfactual that excludes only the nonlinear component. The red shaded area captures the remaining linear demand contribution,

defined as the gap between the two counterfactual paths. The horizontal red line indicates the 6 percent unemployment threshold that separates low- and high-unemployment regimes.

Nonlinear demand effects are economically relevant mainly during major recessions, particularly the early 1980s downturns, the Gulf War recession, and the Great Recession. In these episodes, they contribute not only to the rise in unemployment but also to its slow recovery, reinforcing persistence. By contrast, during the oil shock episodes of the 1970s, the dark red counterfactual lies above observed unemployment, indicating that supply disturbances were the primary drivers of the increase in unemployment and demand shocks were partly offsetting these effects: without them, unemployment would have risen even more.

3.4 Adding state dependence

We next study whether the propagation of the main business-cycle shocks also depends on the state of economy. Specifically, we assess whether the sign asymmetries for demand shocks and the sign symmetries for supply shocks continue to hold when conditioning on whether the economy is in an expansion or a recession.

To study joint sign- and state-dependence, we let $g(u_t^i)$ be a vector-valued function,

$$g(u_t^i) = \left[\, |u_t^i| \quad s_{t-1}u_t^i \quad s_{t-1}|u_t^i| \,\right]',$$

where s_t is an indicator of labor-market tightness, equal to 1 when the unemployment rate is below 6 percent and 0 otherwise. The absolute-value term $|u_t^i|$ captures sign asymmetry, while the interaction terms $s_{t-1}u_t^i$ and $s_{t-1}|u_t^i|$ allow the transmission of both positive and negative shock to vary with the state of the economy in the previous period, thus distinguishing between low- and high-unemployment regimes.

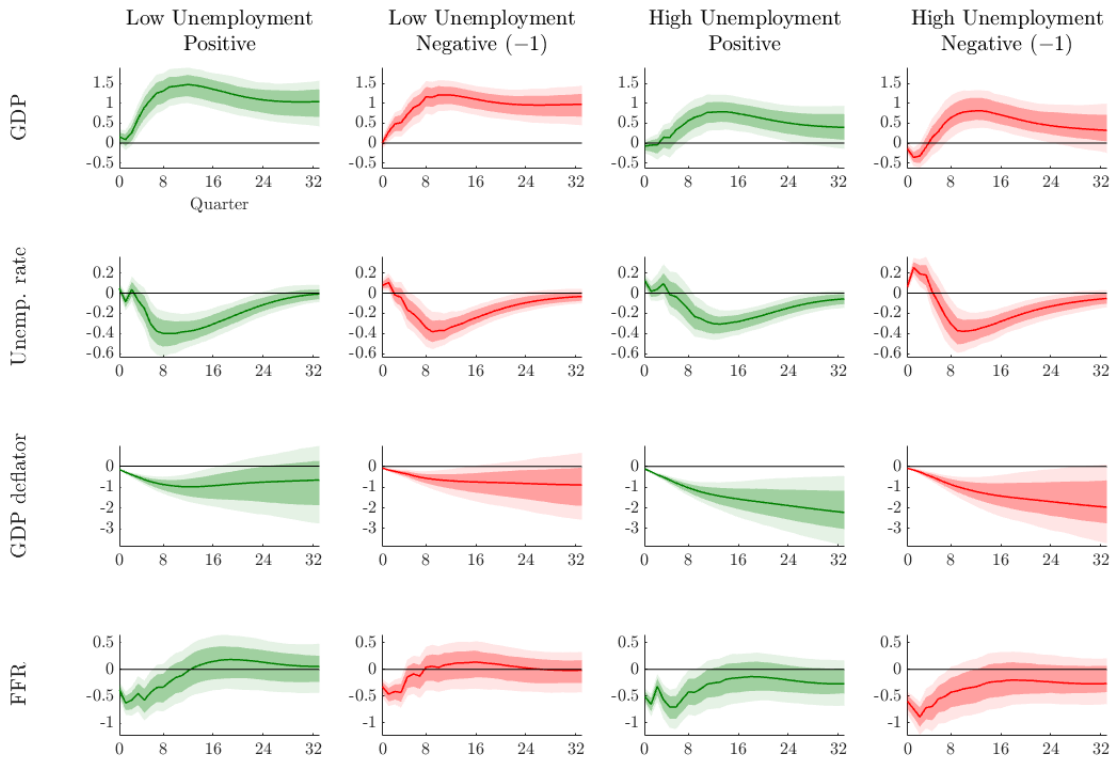


Figure 4: State-dependent impulse responses Supply

Notes. State-dependent impulse responses to supply shocks. The first two columns report responses in low-unemployment states (unemployment rate below 6 percent), while the last two columns report responses in high-unemployment states (unemployment rate above 6 percent). Within each regime, responses to positive shocks are shown in green solid lines, while responses to negative shocks (multiplied by minus one to ease comparison) are shown in red solid lines. Shaded areas represent 68% and 90% bootstrap confidence bands.

In this case, the impulse response to a shock $u_t^i = \bar{u}^i$ is a straightforward extension of equation (3):

$$IRF(\bar{u}^i) = \gamma^*(L)\,\bar{u}^i + \beta_1^*(L)\,|\bar{u}^i| + \beta_2^*(L)\,s_{t-1}\bar{u}^i + \beta_3^*(L)\,s_{t-1}|\bar{u}^i|.$$

Supply. Figure 4 reports the sign- and state-dependent impulse responses of output, prices, the unemployment rate, and the federal funds rate to positive and negative supply shocks. Consistent with the absence of sign asymmetries documented above, the responses to positive and negative supply disturbances remain largely symmetric across

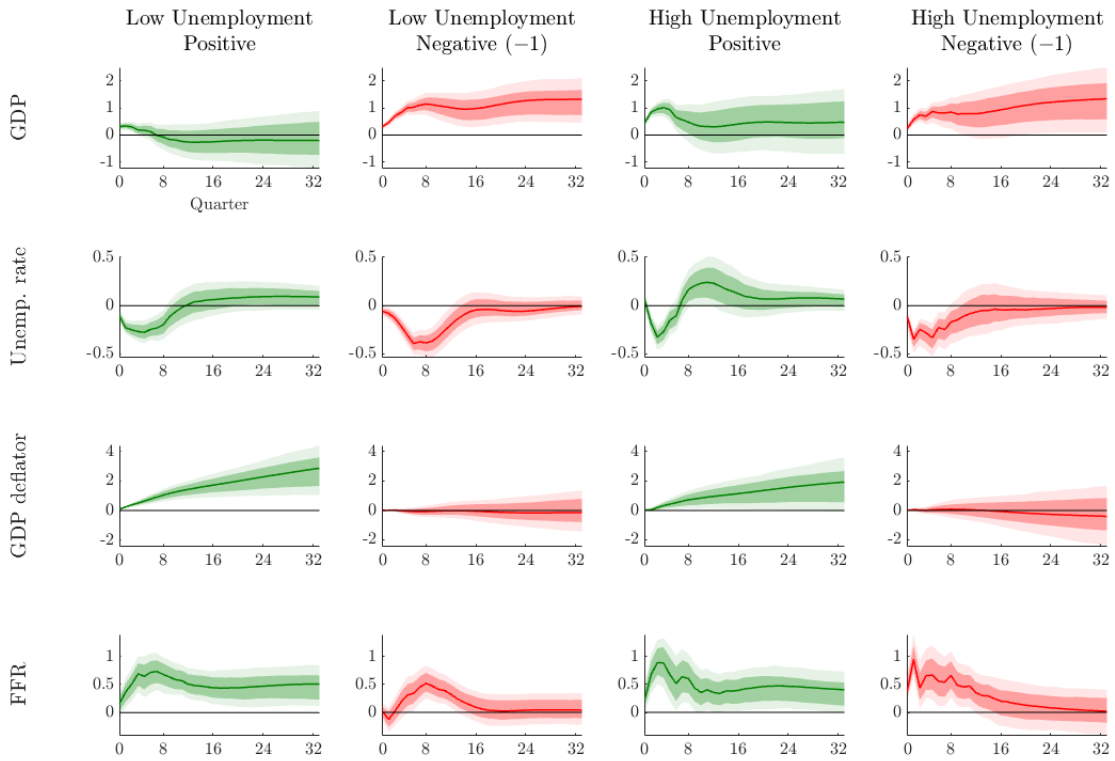


Figure 5: State-dependent impulse responses Demand

Notes. State-dependent impulse responses to demand shocks. The first two columns report responses in low-unemployment states (unemployment rate below 6 percent), while the last two columns report responses in high-unemployment states (unemployment rate above 6 percent). Within each regime, responses to positive shocks are shown in green solid lines, while responses to negative shocks (multiplied by minus one to ease comparison) are shown in red solid lines. Shaded areas represent 68% and 90% bootstrap confidence bands.

labor-market states. State dependence is limited. The responses of GDP, consumption and investment are slightly smaller in high-unemployment economy, for both positive and negative shocks.[12] By contrast, price responses are somewhat stronger in high-unemployment states, with the GDP deflator adjusting more than in low-unemployment regimes. The federal funds rate moves consistently with these price dynamics.

Demand. Figure 5 reports the corresponding sign- and state-dependent impulse responses for demand shocks. The sign asymmetries documented above remain robust once we con-

[12] Additional responses for consumption, investment, and labor-market variables are reported in Figure A2 in the Appendix.

dition on the state of the economy. At the same time, conditioning on the state reveals additional differences in the propagation of expansionary shocks. Following positive demand shocks, GDP increases on impact in both regimes, peaks within the first few quarters, and gradually fades after roughly two to three years. The expansion is nevertheless stronger when the shock occurs in a high-unemployment economy, where the peak response is about twice as large as in low-unemployment states. Investment exhibits a similar pattern (see Figure A1 in the Appendix A). On the nominal side, prices increase in both regimes, though more strongly when unemployment is low. Thus, while expansionary shocks have only temporary real effects, they stimulate activity more in slack labor markets and generate stronger price pressures when labor markets are tight. By contrast, the role of state dependence is considerably weaker following negative demand shocks.

3.5 Robustness

We assess the robustness of our main results under a set of alternative specifications. Figures 6 and 7 report the GDP impulse responses obtained from three robustness exercises. First, we vary the number of static factors in the DFM, considering $r = 10, 12,$ and 13 instead of the baseline $r = 11$. Thinner dotted lines correspond to lower values of r, thicker ones to higher values, and the solid line denotes the baseline specification. Second, we change the lag length of the VAR between three and six lags. As before, thinner dotted lines correspond to lower lag orders and thicker ones to higher orders, with the baseline shown by the solid line. Third, we modify the identification target used in the first step of the frequency-domain procedure. The thinnest dotted line corresponds to a single-variable target based solely on GDP, as in Angeletos et al. (2020), while the thickest line corresponds to the benchmark target set in Forni et al. (2025b). In all cases, the alternative specifications yield impulse responses that are very similar to the baseline and remain within the confidence bands. Taken together, these exercises show that the nonlinear GDP

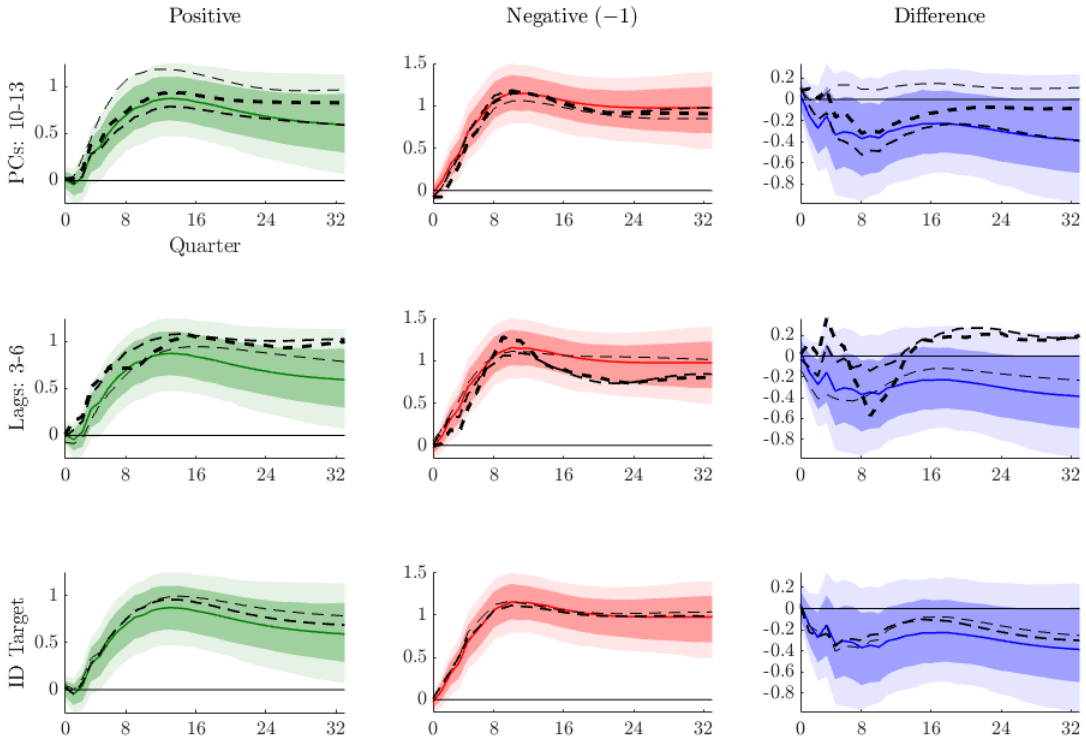


Figure 6: Robustness Supply

Notes. GDP impulse responses to supply shocks under alternative specifications. The solid line denotes the baseline model. Dotted lines with different thicknesses represent three robustness exercises. Thinner dotted lines correspond to lower values of the parameter considered, while thicker dotted lines correspond to higher values. Shaded areas represent the 68% and 90% bootstrap confidence bands from the baseline specification.

response to the main business-cycle shocks is highly robust to alternative choices of the number of factors, lag length, and identification target.

4 A simple NK model with wage rigidities

The model presented in this section serves two purposes. First, it provides a simple qualitative explanation for the empirical findings, rationalizing why demand shocks display strong sign asymmetry while supply shocks remain approximately symmetric. Second, it is used to assess, via Monte Carlo simulation, whether the econometric estimator can recover structural shocks and their nonlinear impulse responses.

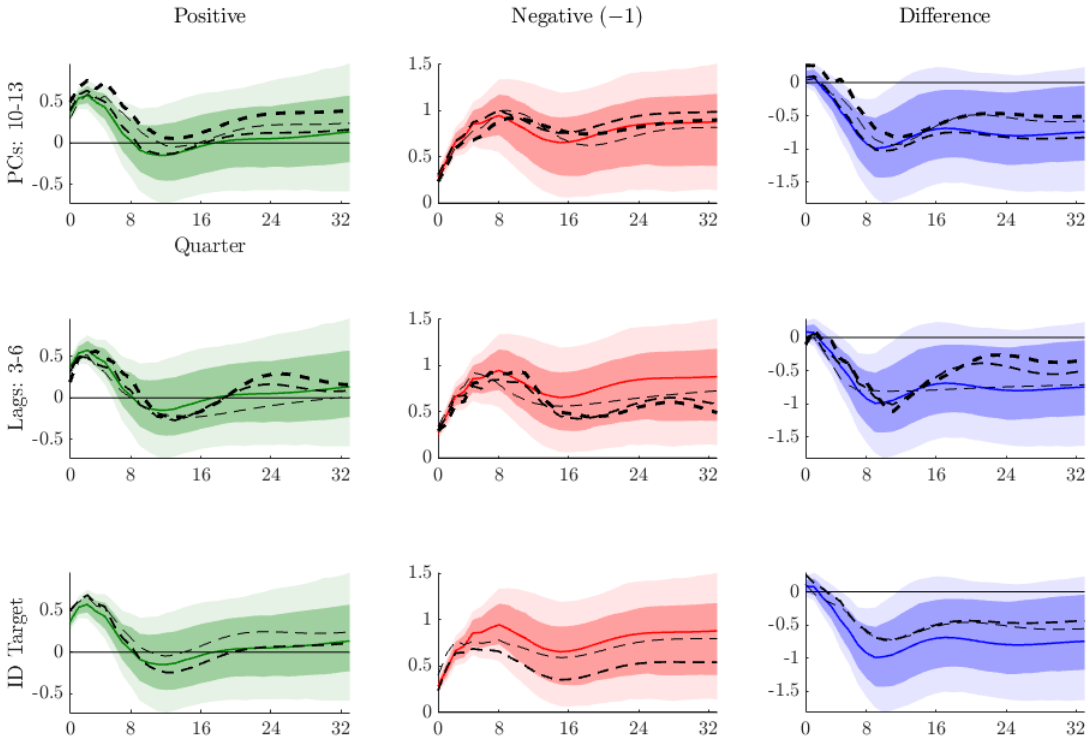


Figure 7: Robustness Demand

Notes. GDP impulse responses to demand shocks under alternative specifications. The solid line denotes the baseline model. Dotted lines with different thicknesses represent three robustness exercises. Thinner dotted lines correspond to lower values of the parameter considered, while thicker dotted lines correspond to higher values. Shaded areas represent the 68% and 90% bootstrap confidence bands from the baseline specification.

Section 4.1 outlines a standard New Keynesian model with downward nominal wage rigidity. Section 4.2 develops the underlying intuition. Section 4.3 presents numerical simulations supporting the mechanism. Section 4.4 conducts a Monte Carlo exercise using the model as the data-generating process to evaluate the performance of the empirical estimator.

4.1 General setup

We consider the standard New Keynesian (NK) framework of Galí (2015). The economy is populated by a continuum of identical households that consume a composite final

good, supply labor, and hold risk-free bonds; a continuum of monopolistically competitive intermediate-good firms that use labor as the sole variable input and set prices according to a Calvo (1983) mechanism; and a monetary authority that sets the nominal interest rate according to a Taylor-type rule.

Consistent with the empirical evidence, the model features two sources of shocks: demand and supply. Following the conventional modeling prescription, we characterize demand shocks as preference shocks, a_t. A positive preference shock lowers households' effective discount factor, thus increasing their incentive to shift consumption toward the present. Supply shocks are instead characterized as technology shocks, z_t, whereby a positive technology shock allows firms to produce a larger quantity of intermediate goods for a given level of labor input.

A key departure from the textbook NK framework is the introduction of downward nominal wage rigidity (DNWR). This assumption implies that nominal wages cannot decrease, that is,

$$W_t = \max\left\{W_t^*,\ W_{t-1}\right\}, \tag{7}$$

where W_t^* is the optimal wage that would prevail in the absence of the constraint.[13]

When the constraint binds, the realized wage exceeds the market-clearing wage, so that firms demand fewer hours than households are willing to supply. We therefore define unemployment as the difference between desired labor supply and actual hours worked. By construction, unemployment is zero when the constraint is not binding, and it increases with its tightness.

Remaining details of the model are reported in Sections C.1-C.6 of the Online Appendix.

[13] See Schmitt-Grohé and Uribe (2016) for an early theoretical representation of the DNWR constraint. Additional arguments and empirical evidence in favor of downward wage rigidity can be found in Malthus (1967), Bewley (1999), Fallick et al. (2016), and Barattieri et al. (2014).

4.2 Intuition

To understand why demand-type shocks exhibit strong sign asymmetries while supply-type shocks remain approximately symmetric, it is useful to examine their propagation mechanisms and how they interact with the DNWR constraint.

To shed light on the mechanism, it is useful to express the constraint in real terms:

$$w_t \geq \frac{w_{t-1}}{\Pi_t}$$

where $w_t \equiv W_t/P_t$ and $\Pi_t \equiv P_t/P_{t-1}$ denotes the real wage and inflation, respectively. Since nominal wages cannot decrease, the constraint implies that real wages can fall only insofar as inflation erodes their purchasing power. When inflation is high, real wages may adjust downward without requiring nominal wage cuts. Conversely, low inflation restricts downward real wage adjustment, making the constraint more likely to bind. This formulation highlights that the DNWR effectively imposes a lower bound on real wages, with inflation governing the extent to which real wage can adjust.

To build intuition, we approximate the model's dynamics by treating the DNWR constraint as nonbinding in states where firms would otherwise choose to raise real wages. This simplification abstracts from cases in which inflation is sufficiently negative to activate the constraint even when real wages are under upward pressure—a scenario that is unlikely given our calibration of steady-state annual inflation at 2 percent (see Section 4.3).[14] Under this approximation, the model's conditional responses simplify considerably: since both expansionary preference and technology shocks place upward pressure on real wages, the DNWR constraint is slack in both cases and can therefore be set aside.

In light of the fact that the DNWR constraint is not likely to bind following expansionary shocks, asymmetry across positive and negative disturbances arises when the constraint

[14] This is confirmed by the numerical impulse responses in Figure 8: unemployment—the indicator of whether the constraint binds—remains approximately zero following expansionary preference and technology shocks.

binds in response to contractionary shocks. We therefore trace the asymmetry of preference shocks, and the symmetry of technology shocks, to differences in the likelihood that adverse shocks trigger the DNWR constraint.[15]

The real-wage formulation of the DNWR clarifies why the constraint is more likely to bind following adverse demand shocks than after adverse technology shocks. Both types of adverse shocks put downward pressure on wages, increasing the likelihood that the DNWR constraint binds. However, they differ in how they affect inflation and thus the effective tightness of the real-wage lower bound. An adverse demand shock puts downward pressure on prices, making the real-wage lower bound tighter and increasing the likelihood that the DNWR constraint binds. By contrast, an adverse technology shock raises inflation, allowing real wages to adjust downward even when nominal wages are constrained.

This intuition speaks directly to the Olivera (1964)–Tobin (1972) argument that positive inflation "greases the wheels" of the labor market by relieving downward nominal wage rigidity. Our analysis reveals, however, that the degree of "grease" inflation provides is not invariant to the source of the disturbance. Adverse technology shocks are self-relieving in the sense that they endogenously generate the inflation needed to erode the real-wage lower bound, whereas adverse demand shocks are self-tightening, simultaneously compressing real activity and withdrawing the inflationary margin that would otherwise facilitate real-wage adjustment. The extent to which steady-state inflation insulates the economy from the costs of nominal wage rigidity thus depends critically on the composition of the shocks driving the business cycle.[16]

[15] Note that we implicitly treat DNWR as the primary source of nonlinearity. While the model is nonlinear across many dimensions, the occasionally binding DNWR constraint accounts for the dominant share of nonlinear behavior. Figure C2 in the Online Appendix provides further numerical evidence supporting this point.

[16] A more detailed discussion of the intuition, along with a graphical illustration, can be found in Section C.7 of the Online Appendix.

4.3 Numerical simulations

The model is calibrated at a quarterly frequency using standard parameter values from the New Keynesian literature. One important deviation from the conventional calibration is that steady-state inflation is set to match a 2 percent annual inflation target. This departure from the common assumption of zero steady-state inflation is not only more consistent with observed monetary policy practice, but it also affects the tightness of the DNWR constraint. In particular, higher steady-state inflation reduces the likelihood that the constraint binds, making the implied degree of downward nominal wage rigidity more realistic. In addition, the persistence and standard deviations of the preference and technology shocks are calibrated so that positive realizations of each shock generate identical output dynamics. This normalization ensures that differences in asymmetric responses—across shocks and across shock signs—reflect the nature of the underlying disturbance rather than differences in magnitude or persistence.

The model is solved using a global projection method based on Chebyshev polynomials. Exogenous preference and technology shocks are discretized on a finite grid, while policy functions for the endogenous state variables are approximated using Chebyshev polynomials over their respective domains. To study potentially asymmetries to positive and negative realizations of preference and technology shocks, we compute generalized impulse response functions (GIRFs) following Koop et al. (1996). Additional details on the calibration, the solution method and the definition and estimation of the GIRFs are reported to Sections C.8-C.10 of the Online Appendix.

The model is intentionally parsimonious and is not designed to quantitatively replicate the empirical impulse responses. The absence of state variables in the production function and of standard real rigidities limits the model's ability to generate the degree of persistence observed empirically. We therefore interpret the model's success as replicating the qualitative pattern of amplification and attenuation across shock types and signs, rather than matching the full dynamics of the empirical responses. We conjecture that enrich-

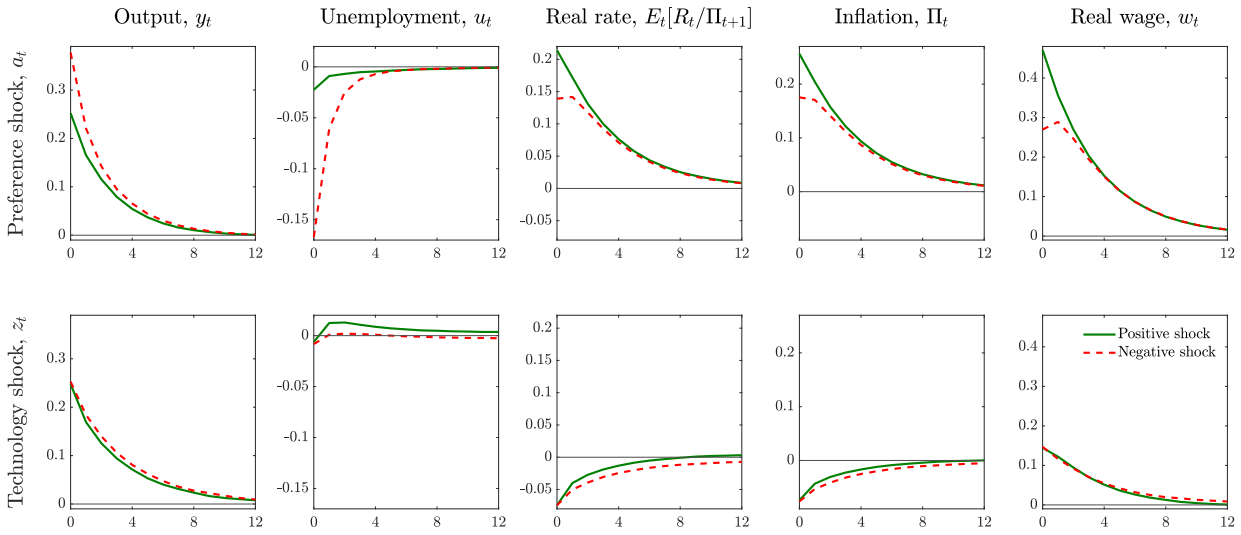


Figure 8: Model-implied impulse responses to preference and technology shocks

Notes. Impulse responses of output y_t, unemployment u_t, the real interest rate $\mathbb{E}_t[R_t/\Pi_{t+1}]$, inflation Π_t, and the real wage w_t to two–standard-deviation preference shocks (first row) and technology shocks (second row). Responses to positive shocks are shown by solid green lines, while responses to negative shocks are multiplied by minus one to ease comparison and are plotted using red dotted lines. The figure reports unconditional generalized impulse response functions (GIRFs) à la Koop et al. (1996), computed by Monte Carlo simulation as the average difference between two economies that share the same history and sequence of future stochastic realizations, but differ in the realization of a one-time structural shock at impact. Responses of output and the real wage are expressed as percentage deviations, while responses of unemployment, the real interest rate, and inflation are reported in percentage-point deviations. Additional details on the calibration, the solution method and the definition and estimation of the GIRFs are reported to Sections C.8-C.10 of the Online Appendix.

ing the model with these features would propagate the short-run amplification mechanism into the medium run, bringing the model's dynamics closer to their empirical counterparts. With this interpretive lens in mind, Figure 8 reports impulse responses to positive and negative preference shocks (first row) and to positive and negative technology shocks (second row). The responses to negative realizations are multiplied by minus one to ease comparison.

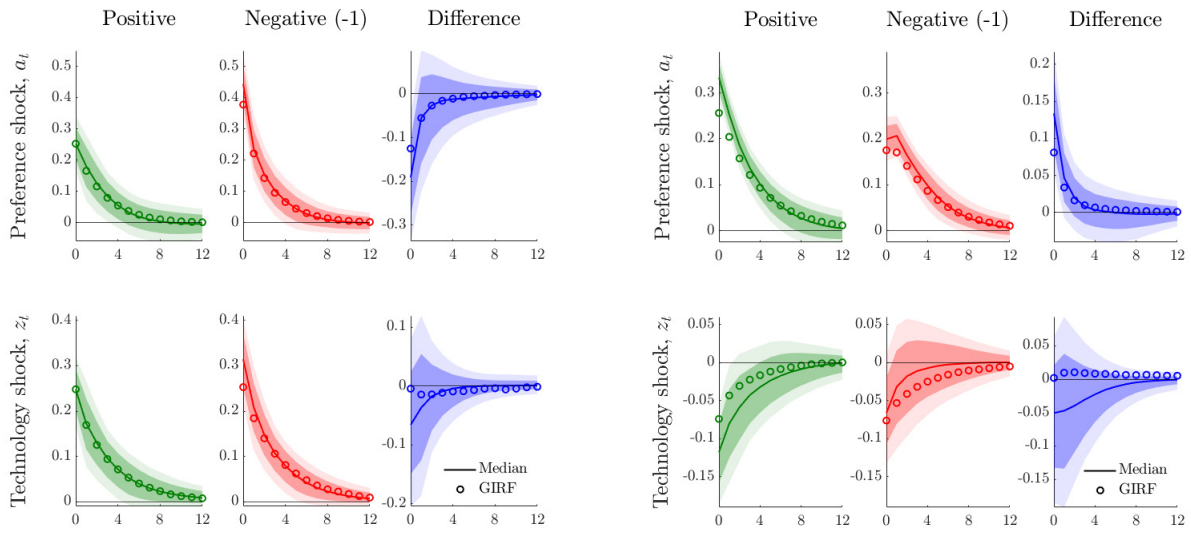
The impulse responses confirm the mechanism outlined in the previous section. In the case of preference shocks, negative realizations generate amplified effects on quantities and a more muted response of prices. The behavior of unemployment suggests that this asymmetry is largely driven by the downward nominal wage rigidity (DNWR) constraint, as unemployment reflects the degree to which the constraint is binding. Following nega-

tive preference shocks, unemployment rises sharply, indicating that the DNWR constraint binds, whereas its response to positive shocks is negligible, consistent with the constraint being slack or only weakly binding. When the constraint binds after adverse shocks, real wage adjustment is limited, which amplifies the contraction in output through a larger decline in hours worked and dampens the response of inflation via a more muted reduction in marginal costs.[17]

Interestingly, the monetary authority's response—measured by movements in the real interest rate—is weaker following negative preference shocks, even though the associated downturn is more severe. This pattern is counterfactual, as one would typically expect a stronger policy response to larger contractions. Put differently, for the same adverse preference shock, monetary policy responds more aggressively in an economy without the DNWR constraint than in an otherwise identical economy where the constraint is present. This outcome, however, is a direct consequence of the Taylor rule assigning greater weight to inflation stability rather than to output stability. Because the DNWR constraint substantially dampens the response of prices—despite much larger movements in real activity—the implied adjustment of the policy rate is correspondingly weaker. This asymmetry suggests that, in the presence of DNWR, a monetary policy framework more strongly oriented toward output stabilization may be warranted. This policy recommendation is in line with Debortoli et al. (2025).

In the case of technology shocks, by contrast, negative and positive realizations generate largely symmetric responses in both prices and quantities. Unsurprisingly, unemployment responds only negligibly to these shocks, implying that the DNWR constraint never binds and therefore does not introduce asymmetries into the economy.

[17] Impulse responses of hours worked, nominal policy rate, nominal wage inflation, real marginal costs, and exogenous shock processes are reported in Figure C3 of the Online Appendix.



(a) Output, y_t (b) Inflation, Π_t

Figure 9: Monte Carlo simulation

Notes. Impulse responses of output y_t (panel a) and inflation Π_t (panel b) to two–standard-deviation preference shocks (first row) and technology shocks (second row). Responses to positive shocks are plotted in green, responses to negative shocks are in red and multiplied by minus one to ease comparison, and their differences in blue. The circles denote the unconditional generalized impulse response functions (GIRFs) à la Koop et al. (1996), also reported in Figure 8. The solid lines and shaded areas represent the median and the 68% and 90% percentile bands obtained from a Monte Carlo simulation, in which we repeatedly simulate artificial datasets matching the empirical dataset in both time dimension and number of variables, supplementing the model's original variables with additional series constructed as linear combinations thereof, and apply the same econometric procedure as in the empirical analysis. For each shock type, the estimated impulse responses are normalized such that the median impact response to a positive shock matches the corresponding GIRF. Responses of output are expressed as percentage deviations, while responses of inflation are reported in percentage-point deviations. Additional details on the calibration, the solution method, and the Monte Carlo simulation are provided in Sections C.8, C.9, and C.11 of the Online Appendix.

4.4 Monte Carlo simulation

The objective of this section is to conduct a Monte Carlo simulation to assess whether the econometric estimator used in the empirical analysis is able to recover the structural shocks and, most importantly, their nonlinear impulse responses. To this end, we use the theoretical model as the data-generating process and repeatedly simulate artificial datasets that match the empirical dataset in both time dimension and number of variables, supplementing the model's original variables with additional series constructed as linear combinations thereof. For each simulated dataset, we then apply the same econometric procedure as in

the empirical analysis and store the estimated impulse responses. Additional details on the procedure are reported in Section C.11 of the Online Appendix.

Figure 9 displays the results of the Monte Carlo exercise for output and inflation, reporting the median and confidence intervals of the estimated impulse responses alongside the GIRFs from Figure 8.

The econometric procedure performs remarkably well in recovering the model-implied nonlinear impulse responses. Across all specifications, the GIRFs usually lie within the 68% percentile band and track the median estimated response closely, indicating that the estimator successfully captures the key nonlinear features of the data-generating process. In particular, the estimator correctly replicates the approximate symmetry of technology shocks across positive and negative realizations, as well as the asymmetric pattern of preference shocks, whereby adverse shocks generate amplified responses in quantities and a more muted adjustment in prices relative to their expansionary counterparts.

5 Conclusions

This paper studies the propagation of economic fluctuations using a nonlinear extension of the Structural Dynamic Factor Model. Our results uncover a sharp contrast between the two main drivers of U.S. business-cycle fluctuations. Supply shocks propagate symmetrically and are well captured by linear dynamics. Demand shocks, by contrast, exhibit pronounced and economically meaningful sign asymmetries: adverse shocks lead to larger and more persistent contractions, with limited downward adjustment in prices and nominal wages, leaving real activity to absorb most of the adjustment. These asymmetries are stronger during booms and are quantitatively important, accounting for a substantial share of both business-cycle and long-run variation in output, consumption, investment, and labor market variables. Historical decompositions further show that nonlinear demand effects are concentrated in major recessions, contributing both to the depth of downturns and the sluggishness of subsequent recoveries.

A simple New Keynesian model with downward nominal wage rigidity rationalizes these findings through a clear and intuitive mechanism. The constraint is largely irrelevant after expansionary shocks, when real wages are naturally pushed upward. It becomes binding in downturns, where inflation—by eroding the real value of wages—acts as a substitute for the nominal wage cuts firms would otherwise implement. Crucially, this adjustment operates asymmetrically across shocks. Supply shocks generate inflationary pressure that keeps the real-wage floor slack, whereas demand shocks simultaneously reduce output and eliminate the inflationary margin needed for real-wage adjustment, thereby amplifying the contraction. A Monte Carlo exercise confirms that the econometric estimator can reliably recover these nonlinear dynamics from the data.

Taken together, these results point to downward nominal wage rigidity as a key amplification mechanism for contractionary demand shocks and suggest that monetary policy frameworks with a stronger emphasis on output stabilization may be better suited to mitigating the depth and persistence of demand-driven downturns.

References

ALESSI, L., M. BARIGOZZI, AND M. CAPASSO (2010): "Improved penalization for determining the number of factors in approximate factor models," *Statistics & Probability Letters*, 80, 1806–1813.

ANDREASEN, M. M., G. CAGGIANO, E. CASTELNUOVO, AND G. PELLEGRINO (2024): "Does risk matter more in recessions than in expansions? Implications for monetary policy," *Journal of Monetary Economics*, 143, 103533.

ANGELETOS, G. M., F. COLLARD, AND H. DELLAS (2020): "Business cycle anatomy," *American Economic Review*, 110, 3030–3070.

AUERBACH, A. J. AND Y. GORODNICHENKO (2012): "Measuring the Output Responses to Fiscal Policy," *American Economic Journal: Economic Policy*, 4, 1–27.

AVARUCCI, M., M. CAVICCHIOLI, M. FORNI, AND P. ZAFFARONI (2026): "Frequency-band estimation of the number of factors," *Journal of the American Statistical Association*, 1–13.

BAI, J. AND S. NG (2002): "Determining the number of factors in approximate factor models," *Econometrica*, 70, 191–221.

BARATTIERI, A., S. BASU, AND P. GOTTSCHALK (2014): "Some evidence on the importance of sticky wages," *American Economic Journal: Macroeconomics*, 6, 70–101.

BARIGOZZI, M., M. LIPPI, AND M. LUCIANI (2021): "Large-dimensional dynamic factor models: Estimation of impulse–response functions with I (1) cointegrated factors," *Journal of Econometrics*, 221, 455–482.

BARNICHON, R., D. DEBORTOLI, AND C. MATTHES (2022a): "Understanding the Size of the Government Spending Multiplier: It's in the Sign," *The Review of Economic Studies*, 89, 87–117.

BARNICHON, R. AND C. MATTHES (2018): "Functional Approximation of Impulse Responses," *Journal of Monetary Economics*, 99, 41–55.

BARNICHON, R., C. MATTHES, AND A. ZIEGENBEIN (2022b): "Are the Effects of Financial Market Disruptions Big or Small?" *Review of Economics and Statistics*, 104, 557–570.

BASU, S., G. CANDIAN, R. CHAHROUR, AND R. VALCHEV (2021): "Risky Business Cycles," NBER Working Paper 28693, National Bureau of Economic Research, revised September 2024.

BEAUDRY, P. AND G. KOOP (1993): "Do Recessions Permanently Change Output?" *Journal of Monetary Economics*, 31, 149–163.

BEN ZEEV, N., V. A. RAMEY, AND S. ZUBAIRY (2023): "Do Government Spending Multipliers Depend on the Sign of the Shock?" *AEA Papers and Proceedings*, 113, 382–387.

BERNANKE, B. S., J. BOIVIN, AND P. ELIASZ (2005): "Measuring the effects of monetary policy: a factor-augmented vector autoregressive (FAVAR) approach," *The Quarterly journal of economics*, 120, 387–422.

BEWLEY, T. F. (1999): *Why wages don't fall during a recession*, Harvard university press.

BLANCHARD, O. J. AND D. QUAH (1989): "The Dynamic Effects of Aggregate Demand and Supply Disturbance," *American Economic Review*, 79, 655–673.

BRIANTI, M. AND V. CORMUN (2024): "Expectation-driven boom-bust cycles," *Journal of Monetary Economics*, 146, 103575.

BRIGNONE, D. AND M. MAZZALI (2025): "It is all about demand and supply: A dualistic view of the euro area business cycle," *Bank of England Staff Working Paper Series*.

BURNS, A. F. AND W. C. MITCHELL (1946): *Measuring Business Cycles*, New York: National Bureau of Economic Research.

CAGGIANO, G., E. CASTELNUOVO, V. COLOMBO, AND G. NODARI (2015): "Estimating fiscal multipliers: News from a non-linear world," *The Economic Journal*, 125, 746–776.

CAGGIANO, G., E. CASTELNUOVO, AND G. NODARI (2022): "Uncertainty and monetary policy in good and bad times: A replication of the vector autoregressive investigation by Bloom (2009)," *Journal of Applied Econometrics*, 37, 210–217.

CALVO, G. A. (1983): "Staggered prices in a utility-maximizing framework," *Journal of monetary Economics*, 12, 383–398.

CARAVELLO, T. E. AND P. MARTÍNEZ-BRUERA (2024): "Disentangling Sign and Size Nonlinearities," *SSRN Electronic Journal*, working paper.

CARNEVALE, O. P. AND D. DI FRANCESCO (2025): "Are Hysteresis Effects Nonlinear?" *Available at SSRN 5520619*.

CHAHROUR, R. AND K. JURADO (2022): "Recoverability and expectations-driven fluctuations," *The Review of Economic Studies*, 89, 214–239.

DEBORTOLI, D., M. FORNI, L. GAMBETTI, AND L. SALA (2025): "Nonlinear Monetary Policy Tradeoffs," *The Economic Journal*, ueaf132.

DELONG, J. B. AND L. H. SUMMERS (1986): "Are Business Cycles Symmetrical?" in *The American Business Cycle: Continuity and Change*, ed. by R. J. Gordon, Chicago: University of Chicago Press, 166–179.

DIEBOLD, F. X. AND G. D. RUDEBUSCH (1990): "A Nonparametric Investigation of Duration Dependence in the American Business Cycle," *Journal of Political Economy*, 98, 1275–1296.

DUPRAZ, S., E. NAKAMURA, AND J. STEINSSON (2025): "A plucking model of business cycles," *Journal of Monetary Economics*, 152, 103766.

FALK, B. (1986): "Further Evidence on the Asymmetric Behavior of Economic Time Series over the Business Cycle," *Journal of Political Economy*, 94, 1096–1109.

FALLICK, B., M. K. LETTAU, AND W. L. WASCHER (2016): "Downward nominal wage rigidity in the United States during and after the Great Recession," Tech. rep., Federal Reserve Bank of Cleveland.

FAUST, J. (1998): "The Robustness of Identified VAR Conclusions about Money," *Carnegie-Rochester Conference Series on Public Policy*, 49, 207–244.

FORNI, M., A. FRANCONI, L. GAMBETTI, AND L. SALA (2025a): "Asymmetric Transmission of Oil Supply News," *Quantitative Economics*, 16, 947–979.

FORNI, M. AND L. GAMBETTI (2010): "The dynamic effects of monetary policy: A structural factor model approach," *Journal of Monetary Economics*, 57, 203–216.

FORNI, M., L. GAMBETTI, A. GRANESE, L. SALA, AND S. SOCCORSI (2025b): "An american macroeconomic picture: Supply and demand shocks in the frequency domain," *American Economic Journal: Macroeconomics*, 17, 311–341.

FORNI, M., L. GAMBETTI, M. LIPPI, AND L. SALA (2025c): "Common components structural VARs," *Journal of Business & Economic Statistics*, 1–24.

FORNI, M., L. GAMBETTI, N. MAFFEI-FACCIOLI, AND L. SALA (2024): "Nonlinear transmission of financial shocks: Some new evidence," *Journal of Money, Credit and Banking*, 56, 5–33.

FORNI, M., D. GIANNONE, M. LIPPI, AND L. REICHLIN (2009): "Opening the black box: Structural factor models with large cross sections," *Econometric Theory*, 25, 1319–1347.

FRIEDMAN, M. (1964): "The "Plucking Model" of Business Fluctuations," in *Monetary Studies of the National Bureau of Economic Research*, ed. by M. Friedman, Chicago: University of Chicago Press, 7–27, introduces the plucking model of business cycles.

FURLANETTO, F., A. LEPETIT, Ø. ROBSTAD, J. RUBIO-RAMÍREZ, AND P. ULVEDAL (2025): "Estimating hysteresis effects," *American Economic Journal: Macroeconomics*, 17, 35–70.

GALÍ, J. (2015): "Monetary policy, inflation, and the business cycle: an introduction to the new keynesian framework and its applications second edition," *Economics Books*.

GRANESE, A. (2025): "Two main business cycle shocks are better than one," *Journal of Economic Dynamics and Control*, 105231.

HALLIN, M. AND R. LIŠKA (2007): "Determining the number of factors in the general dynamic factor model," *Journal of the American Statistical Association*, 102, 603–617.

HAMILTON, J. D. (1989): "A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle," *Econometrica*, 57, 357–384.

——— (2003): "What Is an Oil Shock?" *Journal of Econometrics*, 113, 363–398.

——— (2011): "Nonlinearities and the Macroeconomic Effects of Oil Prices," *Macroeconomic Dynamics*, 15, 364–378.

KEYNES, J. M. (1936): *The General Theory of Employment, Interest and Money*, London: Macmillan.

KIM, C.-J. AND C. R. NELSON (1999): *State-Space Models with Regime Switching: Classical and Gibbs-Sampling Approaches with Applications*, Cambridge, MA: MIT Press.

KOOP, G., M. H. PESARAN, AND S. M. POTTER (1996): "Impulse response analysis in nonlinear multivariate models," *Journal of econometrics*, 74, 119–147.

MALTHUS, T. R. (1967): *Essay on Principle of Population 1798*, Readex Microprint.

MCCRACKEN, M. AND S. NG (2020): "FRED-QD: A quarterly database for macroeconomic research," Tech. rep., National Bureau of Economic Research.

MITCHELL, W. C. (1927): *Business Cycles: The Problem and Its Setting*, New York: National Bureau of Economic Research.

MORLEY, J. AND J. PIGER (2012): "The asymmetric business cycle," *Review of Economics and Statistics*, 94, 208–221.

NEFTCI, S. N. (1984): "Are Economic Time Series Asymmetric over the Business Cycle?" *Journal of Political Economy*, 92, 307–328.

OLIVERA, J. H. (1964): "On Structural Inflation and Latin-American'Structuralism'," *Oxford economic papers*, 16, 321–332.

POTTER, S. M. (1995): "Forecasting U.S. Recessions with Leading Indicators via Nonlinear Models," *Journal of Forecasting*, 14, 187–203.

RAMEY, V. A. AND S. ZUBAIRY (2018): "Government Spending Multipliers in Good Times and in Bad: Evidence from US Historical Data," *Journal of Political Economy*, 126, 850–901.

SCHMITT-GROHÉ, S. AND M. URIBE (2016): "Downward nominal wage rigidity, currency pegs, and involuntary unemployment," *Journal of Political Economy*, 124, 1466–1514.

SICHEL, D. E. (1989): "Business Cycle Asymmetry: A Deeper Look," *Journal of Monetary Economics*, 24, 85–103.

STOCK, J. H. AND M. W. WATSON (2005): "Implications of dynamic factor models for VAR analysis," .

TENREYRO, S. AND G. THWAITES (2016): "Pushing on a string: US monetary policy is less powerful in recessions," *American Economic Journal: Macroeconomics*, 8, 43–74.

TIAO, G. C. AND R. S. TSAY (1994): "Modeling Multiple Time Series with Applications," *Journal of Time Series Analysis*, 15, 489–518.

TOBIN, J. (1972): "Inflation and Unemployment." *American Economic Review*, 62.

UHLIG, H. (2004): "Do Technology Shocks Lead to a Fall in Total Hours Worked?" *Journal of the European Economic Association*, 2, 361–371.

In-Text Appendix

Nonlinear Busyness Cycle Anatomy

by Marco Brianti, Mario Forni, Luca Gambetti, Antonio Granese

A Additional Results

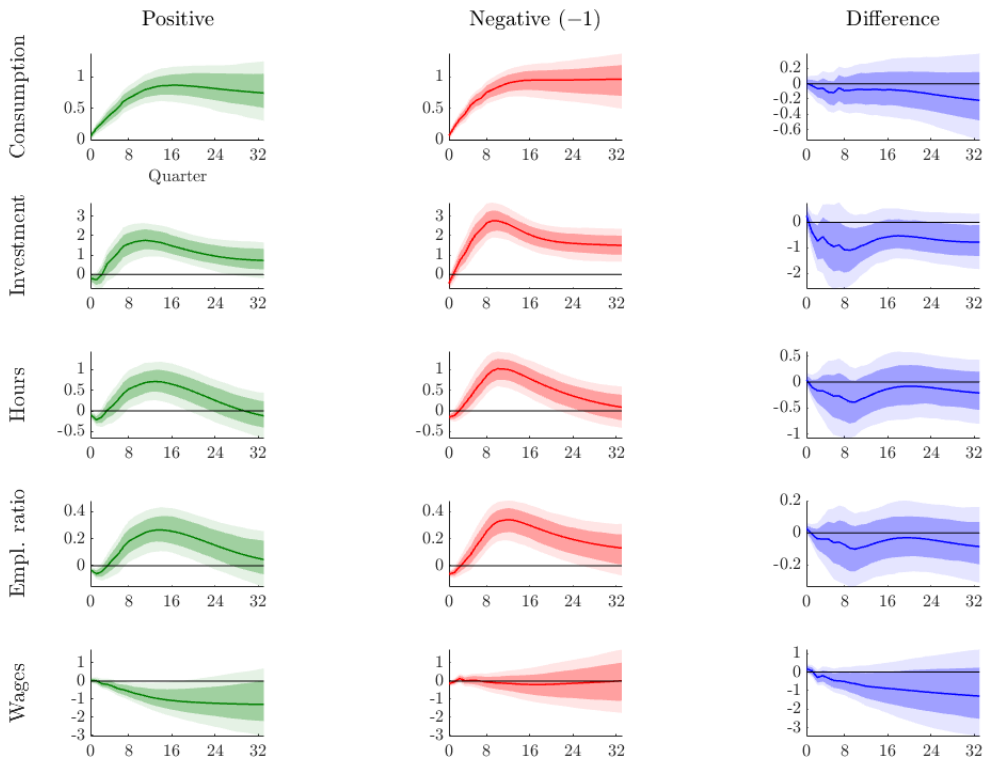


Figure A1: Supply

Notes. Sign-dependent impulse responses to a supply shock. The first column reports responses to a positive shock, the second column reports responses to a negative shock (multiplied by minus one to ease comparison), and the third column reports the difference between the two responses. Solid lines denote point estimates, while shaded areas represent 68% and 90% bootstrap confidence bands.

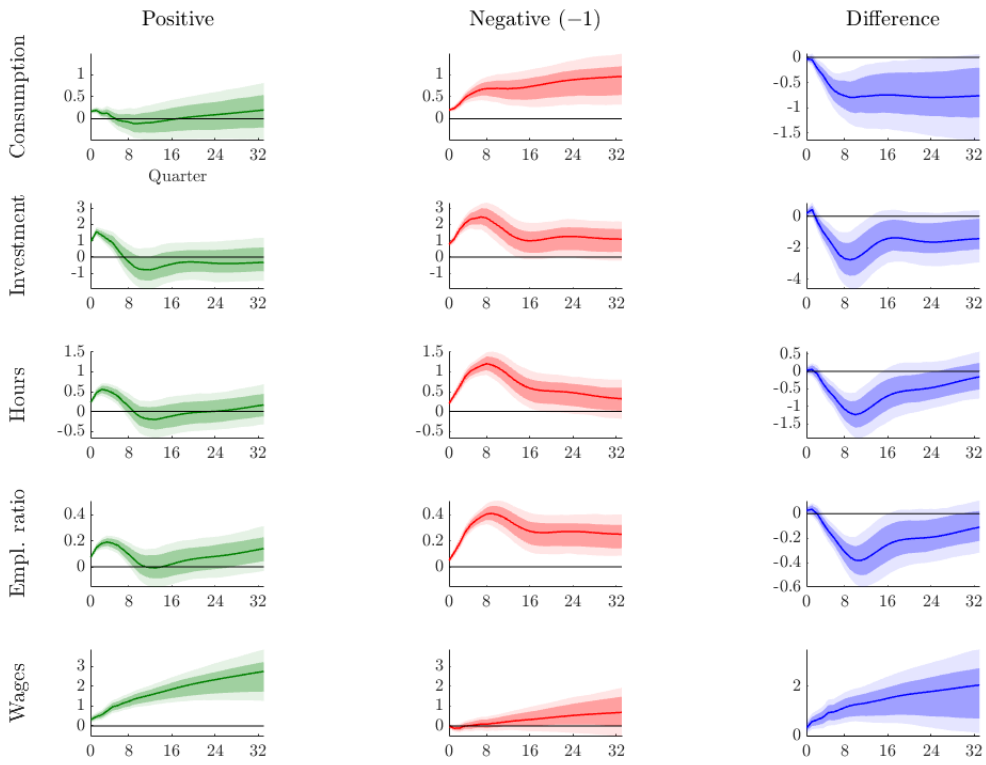


Figure A2: Demand

Notes. Sign-dependent impulse responses to a demand shock. The first column reports responses to a positive shock, the second column reports responses to a negative shock (multiplied by minus one to ease comparison), and the third column reports the difference between the two responses. Solid lines denote point estimates, while shaded areas represent 68% and 90% bootstrap confidence bands.

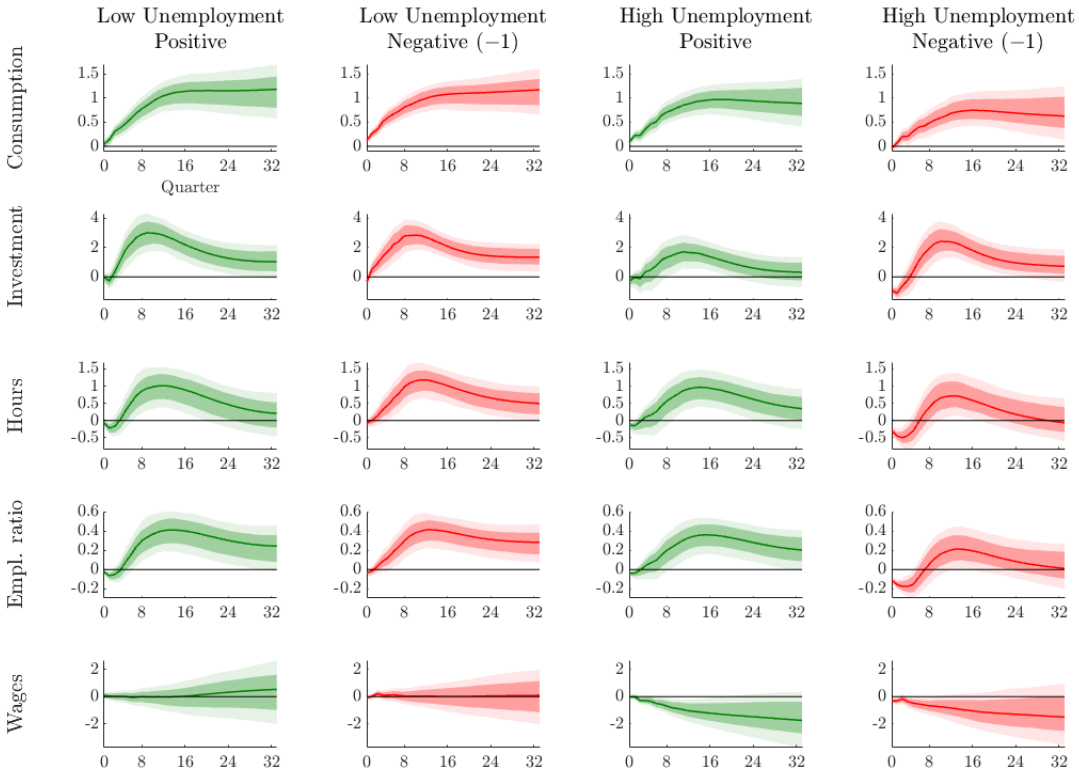


Figure A3: State-dependent impulse responses supply shock

Notes. State-dependent impulse responses to supply shocks. The first two columns report responses in low-unemployment states (unemployment rate below 6 percent), while the last two columns report responses in high-unemployment states (unemployment rate above 6 percent). Within each regime, responses to positive shocks are shown in green solid lines, while responses to negative shocks (multiplied by minus one to ease comparison) are shown in red solid lines. Shaded areas represent 68% and 90% bootstrap confidence bands.

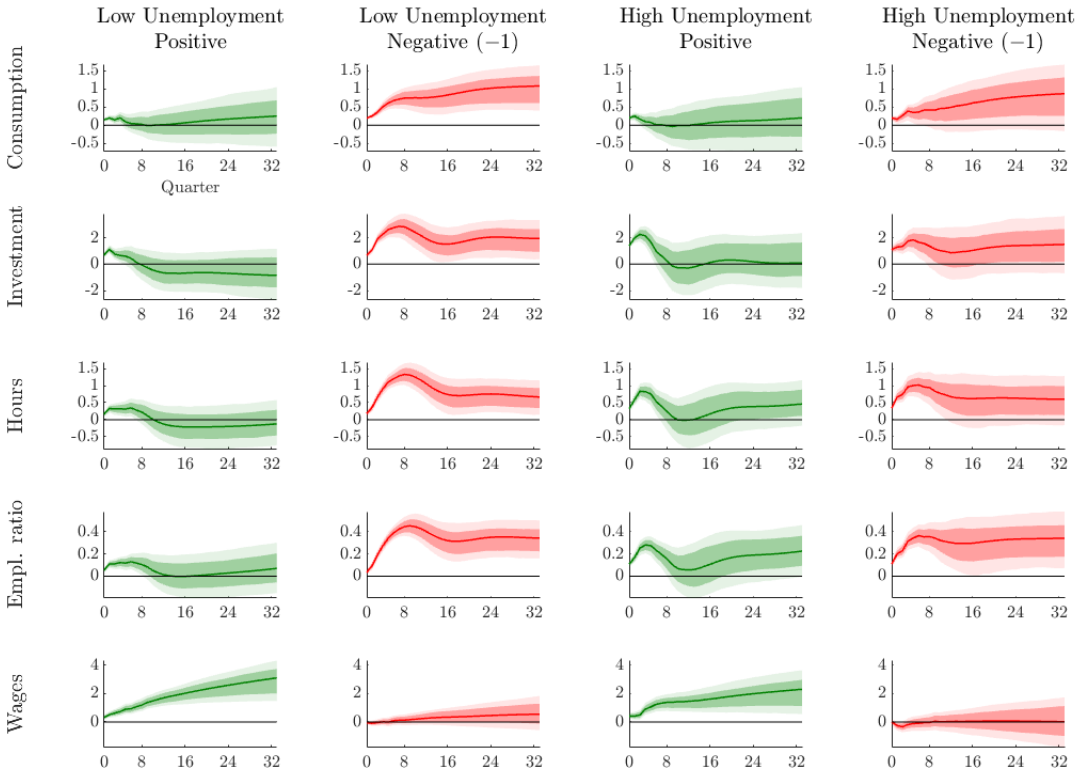


Figure A4: State-dependent impulse responses demand shock

Notes. State-dependent impulse responses to demand shocks. The first two columns report responses in low-unemployment states (unemployment rate below 6 percent), while the last two columns report responses in high-unemployment states (unemployment rate above 6 percent). Within each regime, responses to positive shocks are shown in green solid lines, while responses to negative shocks (multiplied by minus one to ease comparison) are shown in red solid lines. Shaded areas represent 68% and 90% bootstrap confidence bands.

B Data

For the description of each variable see McCracken and Ng (2020). For variables not in the FRED-QD dataset, refer to the Mnemonic and note. Treatment codes: 1 = no treatment; 2 = first difference, Δx_t; 4 = $\log(x_t)$; 5 = log of the first difference, $\Delta \log(x_t)$.

ID	FRED-QD ID	Mnemonic	Treatment Code	Note
1	1	GDPC1/CNP16OV	5	
2	2	PCECC96/CNP16OV	5	
3	3	PCDGx/CNP16OV	5	
4	4	PCESVx/CNP16OV	5	
5	5	PCNDx/CNP16OV	5	
6	6	GPDIC1/CNP16OV	5	
7	7	FPIx/CNP16OV	5	
8	8	Y033RC1Q027SBEAx/CNP16OV	5	
9	9	PNFIx/CNP16OV	5	
10	10	PRFIx/CNP16OV	5	
11	11	A014RE1Q156NBEA	1	
12	12	GCEC1/CNP16OV	5	
13	13	A823RL1Q225SBEA	1	
14	14	FGRECPTx/CNP16OV	5	
15	15	SLCEx/CNP16OV	5	
16	16	EXPGSC1/CNP16OV	5	
17	17	IMPGSC1/CNP16OV	5	
18	18	DPIC96/CNP16OV	5	
19	19	OUTNFB/CNP16OV	5	
20	20	OUTBS/CNP16OV	5	
21		(PCESVx+PCNDx)/CNP16OV	5	
22		(PCDGx+FPIx)/CNP16OV	5	
23	22	INDPRO/CNP16OV	5	
24	23	IPFINAL/CNP16OV	5	
25	24	IPCONGD/CNP16OV	5	
26	25	IPMAT/CNP16OV	5	
27	28	IPDCONGD/CNP16OV	5	
28	30	IPNCONGD/CNP16OV	5	
29	31	IPBUSEQ/CNP16OV	5	
30	35	PAYEMS/CNP16OV	2	
31	36	USPRIV/CNP16OV	2	

Continued on next page

Continued from previous page

ID	FRED-QD ID	Mnemonic	Treatment Code	Note
32	38	SRVPRD/CNP16OV	2	
33	39	USGOOD/CNP16OV	2	
34	51	USGOVT/CNP16OV	2	
35	57	CE16OV/CNP16OV (EMRATIO)	2	
36	58	CIVPART	2	
37	59	UNRATE	1	
38	60	UNRATESTx	1	
39	61	UNRATELTx	1	
40	62	LNS14000012	1	
41	63	LNS14000025	1	
42	64	LNS14000026	1	
43	74	HOABS/CNP16OV	5	
44	76	HOANBS/CNP16OV	5	
45	77	AWHMAN	1	
46	79	AWOTMAN	1	
47	81	HOUST/CNP16OV	5	
48	95	PCECTPI	5	
49	96	PCEPILFE	5	
50		GDPDEF	5	GDP: Implicit Price Deflator
51	97	GDPCTPI	5	
52	98	GPDICTPI	5	
53	120	CPIAUCSL	5	
54	121	CPILFESL	5	
55	122	WPSFD49207	5	
56	123	PPIACO	5	
57	124	WPSFD49502	5	
58	126	PPIIDC	5	
59	129	WPU0561	5	
60	130	OILPRICEx	5	
61	135	COMPRNFB	5	
62	138	OPHNFB	5	
63	139	OPHPBS	5	
64	140	ULCBS	5	
65	142	ULCNFB	5	
66	143	UNLPNBS	5	
67		dtfp	1	Fernald's TFP growth
68		dtfp_util	1	Fernald's TFP growth CU adjusted
69		dtfp_I	1	Fernald's TFP growth – Investment

Continued on next page

Continued from previous page

ID	FRED-QD ID	Mnemonic	Treatment Code	Note
70		dtfp_C	1	Fernald's TFP growth – Consumption
71		dtfp_I_util	1	Fernald's TFP growth CU – Investment
72		dtfp_C_util	1	Fernald's TFP growth CU – Consumption
73	144	FEDFUNDS	1	
74	145	TB3MS	1	
75	146	TB6MS	1	
76	147	GS1	1	
77	148	GS10	1	
78	150	AAA	1	
79	151	BAA	1	
80	152	BAA10YM	1	
81	156	GS10TB3Mx	1	
82		BAA-AAA	1	
83		GS10-FEDFUNDS	1	
84		GS1-FEDFUNDS	1	
85		BAA-FEDFUNDS	1	
86	158	BOGMBASEREALx/CNP16OV	5	
87	160	M1REAL/CNP16OV	5	
88	161	M2REAL/CNP16OV	5	
89	163	BUSLOANSx/CNP16OV	5	
90	164	CONSUMERx/CNP16OV	5	
91	166	REALLNx/CNP16OV	5	
92	168	TOTALSLx/CNP16OV	5	
93	188	UMCSENTx	1	
94		Business Condition 12 Months	1	Michigan Consumer Survey
95		Business Condition 5 Years	1	Michigan Consumer Survey
96		Current Index	1	Michigan Consumer Survey
97		Expected Index	1	Michigan Consumer Survey
98		News Index: Relative	1	Michigan Consumer Survey
99	197	UEMPMEAN	1	
100	201	GS5	1	
101	210	CUSR0000SAC	5	
102	211	CUSR0000SAD	5	
103	212	CUSR0000SAS	5	
104	213	CPIULFSL	5	
105	245	S&P 500	5	
106	246	S&P: indust	5	
107		S&P 500/GDPDEF	5	

Continued on next page

Continued from previous page

ID	FRED-QD ID	Mnemonic	Treatment Code	Note
108		S&P: indust/GDPDEF	5	
109		JLN Macro Unc 1-month	1	Jurado, Ludvigson and Ng Uncertainty
110		JLN Macro Unc 3-month	1	Jurado, Ludvigson and Ng Uncertainty
111		JLN Macro Unc 12-month	1	Jurado, Ludvigson and Ng Uncertainty
112		DPCCRC1Q027SBEAx/CNP16OV	5	Real PCE Excluding food and energy
113		DFXARC1M027SBEAx/CNP16OV	5	Real PCE: Food
114		DNRGRC1Q027SBEAx/CNP16OV	5	Real PCE: Energy goods and services
115		UEMP27OV/UNEMPLOY	1	
116	136	RCPHBS	5	
117	132	CES2000000008x	5	
118	133	CES3000000008x	5	

Online Appendix

Nonlinear Busyness Cycle Anatomy

by Marco Brianti, Mario Forni, Luca Gambetti, Antonio Granese

C An NK model with downward nominal wage rigidities

In this section we outline the structure and derivation of the New Keynesian model used in the main text.

The economy consists of a continuum of identical households that choose final-good consumption c_t, hours worked h_t, and holdings of risk-free bonds B_{t+1}; a representative final-good producer that aggregates intermediate varieties into the composite good y_t; a continuum of intermediate-good firms $j \in [0,1]$ that take production and pricing decisions to maximize the discounted value of profits; and a monetary authority that sets the nominal risk-free interest rate i_t.

Section C.1 formulates the household problem, while Sections C.2 and C.3 describe the behavior of final- and intermediate-good producers. Section C.4 aggregates individual decisions and Section C.5 closes the model. Section C.6 presents the flexible-price counterpart. Finally, Section C.8 discusses the calibration, Sections C.9 and C.10 describe the numerical solution and simulation strategy, and Section C.12 reports the simulation results and additional findings.

C.1 Household

A representative household chooses consumption c_t, hours worked h_t, and next period's nominal bond holdings B_{t+1} to maximize the expected discounted value of utility over an infinite horizon. The household derives utility from consuming the composite good and experiences disutility from supplying labor. Each period, it earns labor income at the nominal wage W_t, receives nominal profits X_t from the firms it owns, and pays lump–sum taxes $P_t T_t$. To reallocate resources across time, the household saves in one–period nominal bonds that pay a gross nominal return $1 + i_t$, where i_t is the net risk-free rate set by the monetary authority.

Formally, the household solves:

$$\begin{aligned} \max_{\{c_t, h_t, B_{t+1}\}_{t=0}^{\infty}} \quad & \mathbb{E}_0 \left\{ \sum_{t=0}^{\infty} \beta^t e^{a_t} \left[\frac{c_t^{1-\gamma}}{1-\gamma} - \chi \frac{h_t^{1+\varphi}}{1+\varphi} \right] \right\} \\ \text{subject to} \quad & P_t c_t + B_{t+1} + P_t T_t \leq W_t h_t + (1 + i_{t-1}) B_t + X_t, \\ \text{for} \quad & t = 0, 1, \ldots, \infty, \end{aligned}$$

where $\beta \in (0,1)$ is the subjective discount factor, $\gamma > 0$ denotes the relative risk aversion, $\chi > 0$ scales the disutility of labor with respect to the utility of consumption, $\varphi > 0$ is the inverse Frisch elasticity of labor supply, and a_t is a preference shock to the household's period utility.[18]

For simplicity, one can also express the budget constraint in real terms (that is, in terms of the consumption good) by dividing it by the aggregate price index P_t. That is,

$$c_t + b_{t+1} + T_t \leq w_t h_t + \frac{(1 + i_{t-1})}{\Pi_t} b_t + x_t, \tag{C.1}$$

where $b_{t+1} = B_{t+1}/P_t$ denotes amount of bonds expressed in real terms, $w_t = W_t/P_t$ represents the real wage, $\Pi_t = P_t/P_{t-1}$ denotes the gross inflation rate, and $x_t = X_t/P_t$ are the real profits.

First order conditions yield the labor supply:

$$w_t = \chi h_t^{\varphi} c_t^{\gamma};$$

and the Euler equation for risk-free bonds:

$$1 = (1 + i_t)\, \mathbb{E}_t \left[\Lambda_{t,t+1} \frac{1}{\Pi_{t+1}} \right] \tag{C.2}$$

[18] In this framework, the preference shifter e^{a_t} acts as an aggregate *demand shock*: a rise in a_t raises the marginal utility of consumption, increasing households' effective desire to consume at given prices and interest rates, and thereby shifting aggregate demand and output.

where

$$\Lambda_{t,t+1} = \beta e^{(a_{t+1}-a_t)} \left(\frac{c_{t+1}}{c_t} \right)^{-\gamma} \tag{C.3}$$

is the stochastic discount factor.

Downward nominal wage rigidities We assume that nominal wages cannot fall below a lower bound. Formally,

$$w_t = \max \left\{ \chi h_t^{\varphi} c_t^{\gamma},\ \rho_w \frac{w_{t-1}}{\Pi_t} \right\}, \tag{C.4}$$

where the first term corresponds to the household's notional (flexible) labor–supply wage and the second term captures the downward rigidity through the parameter $\rho_w \in (0, 1]$. When the rigidity binds, the actual wage exceeds the household's desired wage, implying that firms demand less labor than households are willing to supply. We therefore define unemployment as

$$U_t \equiv \left(\frac{w_t}{\chi\, c_t^{\gamma}} \right)^{\frac{1}{\varphi}} - h_t, \tag{C.5}$$

which by construction satisfies $U_t \geq 0$, with equality when the wage constraint is not binding.

C.2 Final-good producer

The final good is produced by a representative competitive firm that aggregates a continuum of differentiated intermediate varieties indexed by $j \in [0, 1]$. This representative firm combines these inputs into a single homogeneous composite good—whose output is denoted by y_t—with the following production function:

$$y_t = \left[\int_0^1 y_t(j)^{\frac{\varepsilon-1}{\varepsilon}} dj \right]^{\frac{\varepsilon}{\varepsilon-1}},$$

where $y_t(j)$, produced by an intermediate-good firm, is the quantity of the jth intermediate input and $\varepsilon > 1$ is the elasticity of substitution between different intermediate goods.

The representative final–good producer chooses the quantity of each differentiated variety to minimize the cost of producing the aggregate good. Taking as given the prices of intermediate inputs $P_t(j)$ and the required level of final output y_t, the firm solves

$$\begin{aligned} \min_{\{y_t(j)\}} \quad & \int_0^1 P_t(j)\, y_t(j)\, dj \\ \text{subject to} \quad & \left[\int_0^1 y_t(j)^{\frac{\varepsilon-1}{\varepsilon}} dj\right]^{\frac{\varepsilon}{\varepsilon-1}} \geq y_t. \end{aligned}$$

The associated first–order conditions imply the standard demand schedule for each intermediate input:

$$y_t(j) = \left(\frac{P_t(j)}{P_t}\right)^{-\varepsilon} y_t,$$

where the aggregate price index P_t is given by

$$P_t = \left[\int_0^1 P_t(j)^{\,1-\varepsilon} dj\right]^{\frac{1}{1-\varepsilon}}.$$

Using this index, the total cost of acquiring the intermediate inputs can be expressed compactly as

$$\int_0^1 P_t(j)\, y_t(j)\, dj = P_t\, y_t.$$

C.3 Intermediate-good producer

There is a continuum of monopolistically competitive intermediate–good firms indexed by $j \in [0, 1]$. Each firm produces a differentiated variety $y_t(j)$ and faces the downward–sloping demand function

$$y_t(j) = p_t(j)^{-\varepsilon}\, y_t,$$

where $p_t(j) = P_t(j)/P_t$ is the relative price of variety j. Production uses labor as the sole input and exhibits constant returns to scale:

$$y_t(j) = e^{z_t} h_t(j),$$

where e^{z_t} is an aggregate technology level common to all firms, and $h_t(j)$ denotes labor hired by firm j.

Nominal profits for firm j are given by

$$Q_t(j) = P_t(j)y_t(j) - (1-\tau)W_t h_t(j),$$

where W_t is the nominal wage and $\tau \in [0,1)$ is a labor subsidy rate. The subsidy reduces the firm's effective labor cost to $(1-\tau)W_t$. Dividing by the aggregate price level P_t yields real profits:

$$q_t(j) = p_t(j)\, y_t(j) - (1-\tau)\, w_t\, h_t(j),$$

where $q_t(j) = Q_t(j)/P_t$.

The firm problem can be written as:

$$\begin{aligned} \max_{\{p_t(j),h_t(j)\}_{t=0}^{\infty}} \quad & \mathbb{E}_0 \left\{ \sum_{t=0}^{\infty} \Lambda_{0,t} \left[p_t(j)^{1-\varepsilon} y_t - (1-\tau) w_t h_t(j) \right] \right\} \\ \text{subject to} \quad & p_t(j)^{-\varepsilon} y_t \leq e^{z_t} h_t(j). \end{aligned}$$

Under flexible prices, the optimal pricing condition is the standard constant–markup rule,

$$p_t(j) = \mathcal{M}(1-\tau)\frac{w_t}{e^{z_t}},$$

where w_t/e^{z_t} is the real marginal cost and the desired markup $\mathcal{M} = \varepsilon/(\varepsilon-1)$. Thus, in a flexible–price equilibrium, all firms choose the same price and set it as a constant markup over marginal cost.

Instead of flexible prices, we assume in what follows a price-setting environment as in Calvo (1983). Any firm can reset its price in any given period only with probability $1-\theta$, independently of the time elapsed since its last adjustment.

The nominal profits in period $t+k$ of a firm that last re-optimized its price $P_t^*(j)$ in period t are given by:

$$Q_{t+k|t}(j) = P_t^*(j) y_{t+k|t}(j) - (1-\tau) W_{t+k} h_{t+k|t}(j),$$

where $y_{t+k|t}(j)$ and $h_{t+k|t}(j)$ indicate the output produced and the amount of hours employed in period $t+k$ by a firm j that last re-optimized its price in period t.

These can also be expressed in real terms by dividing over P_{t+k}:

$$q_{t+k|t}(j) = \frac{p_t^*(j)}{\Pi_{t,t+k}} y_{t+k|t}(j) - (1-\tau) w_{t+k} h_{t+k|t}(j)$$

where $p_t^*(j) = P_t^*(j)/P_t$ and $\Pi_{t,t+k} \equiv P_{t+k}/P_t$.

The firm re-optimizing in period t will choose the relative price $p_t^*(j)$ that maximizes the current market value of the profits generated while that price remains effective. That is,

$$\begin{aligned} \max_{\{p_t^*(j), h_{t+k|t}(j)\}_{k=0}^{\infty}} \quad & \sum_{k=0}^{\infty} \theta^k \, \mathbb{E}_t \left\{ \Lambda_{t,t+k} \left[\frac{p_t^*(j)^{1-\varepsilon}}{\Pi_{t,t+k}^{1-\varepsilon}} y_{t+k} - (1-\tau) w_{t+k} h_{t+k|t}(j) \right] \right\} \\ \text{subject to} \quad & \frac{p_t^*(j)^{-\varepsilon}}{\Pi_{t,t+k}^{-\varepsilon}} y_{t+k} \leq e^{z_{t+k}} h_{t+k|t}(j). \end{aligned}$$

The Lagrangian is

$$\begin{aligned} \mathcal{L} = \sum_{k=0}^{\infty} \theta^k \, \mathbb{E}_t \Bigg\{ \Lambda_{t,t+k} \Bigg[& \left(\frac{p_t^*(j)^{1-\varepsilon}}{\Pi_{t,t+k}^{1-\varepsilon}} y_{t+k} - (1-\tau) w_{t+k} h_{t+k|t}(j) \right) \\ & + \lambda_{t+k|t}(j) \left(e^{z_{t+k}} h_{t+k|t}(j) - \frac{p_t^*(j)^{-\varepsilon}}{\Pi_{t,t+k}^{-\varepsilon}} y_{t+k} \right) \Bigg] \Bigg\}, \end{aligned}$$

where $\lambda_{t+k|t}(j)$ is the marginal cost at period $t+k$ of the firm j last re-optimizing at period t. The first order condition with respect to $h_{t+k|t}(j)$ is

$$\lambda_{t+k|t}(j) = (1-\tau) \frac{w_{t+k}}{e^{z_{t+k}}} \quad \Longrightarrow \quad \lambda_{t+k} = (1-\tau) \frac{w_{t+k}}{e^{z_{t+k}}},$$

which implies that the marginal cost is common across all firms and does not depend on the timing of the last reset. Moreover, the first order condition with respect to $p_t^*(j)$ is

$$\sum_{k=0}^{\infty} \theta^k \mathbb{E}_t \left\{ \Lambda_{t,t+k} \left[\frac{(1-\varepsilon) p_t^*(j)^{-\varepsilon}}{\Pi_{t,t+k}^{1-\varepsilon}} + \lambda_{t+k} \frac{\varepsilon p_t^*(j)^{-\varepsilon-1}}{\Pi_{t,t+k}^{-\varepsilon}} \right] y_{t+k} \right\} = 0,$$

which is

$$p_t^* = \mathcal{M}(1-\tau) \frac{\sum_{k=0}^{\infty} \theta^k \mathbb{E}_t \left\{ \Lambda_{t,t+k} \Pi_{t,t+k}^{\varepsilon} y_{t+k} w_{t+k} / e^{z_{t+k}} \right\}}{\sum_{k=0}^{\infty} \theta^k \mathbb{E}_t \left\{ \Lambda_{t,t+k} \Pi_{t,t+k}^{\varepsilon-1} y_{t+k} \right\}}, \tag{C.6}$$

where $p_t^*(j) = p_t^*$ as all the resetting firms choose identical prices in equilibrium.

C.4 Aggregation

This section derives the aggregate quantities implied by the optimal decisions of households and firms.

Inflation Let $\Theta_t \subset [0,1]$ denote the measure of firms that do not reset their price in period t and therefore continue to charge their prior price $P_{t-1}(j)$. It follows that the aggregate price level can be written as

$$\begin{aligned} P_t^{1-\varepsilon} &= \int_{\Theta_t} P_{t-1}(j)^{1-\varepsilon} \, dj + (1-\theta)(P_t^*)^{1-\varepsilon} \\ &= \theta (P_{t-1})^{1-\varepsilon} + (1-\theta)(P_t^*)^{1-\varepsilon}. \end{aligned}$$

Dividing both sides by P_t,

$$1 = \theta \Pi_t^{\varepsilon-1} + (1-\theta) \left(p_t^*\right)^{1-\varepsilon},$$

which implies,

$$\Pi_t = \left[\frac{1 - (1-\theta)(p_t^*)^{1-\varepsilon}}{\theta} \right]^{\frac{1}{\varepsilon-1}}. \tag{C.7}$$

Production Because only a subset of firms adjusts its price in period t, adjusting and non–adjusting firms generally charge different prices. Given a common demand schedule, these price differences imply different quantities sold, and therefore different production levels. Since all firms operate the same technology, the resulting variation in output is reflected entirely in the amount of labor they employ.

Thus, we define total hours h_t as

$$\begin{aligned} h_t &\equiv \int_0^1 h_t(j)dj \\ &= \frac{1}{e^{z_t}} \int_0^1 y_t(j)dj = \frac{y_t}{e^{z_t}} \int_0^1 p_t(j)^{-\varepsilon} dj. \end{aligned}$$

By defining price dispersion as

$$D_t = \int_0^1 p_t(j)^{-\varepsilon} dj \geq 1,$$

the aggregate production function is

$$y_t = \frac{e^{z_t}}{D_t} h_t. \tag{C.8}$$

Moreover, D_t takes the following recursive form:

$$\begin{aligned} D_t &\equiv \int_0^1 p_t(j)^{-\varepsilon} dj \\ &= \int_{\Theta_t} \left(\frac{P_{t-1}(j)}{P_t} \right)^{-\varepsilon} dj + \int_{\Theta_t^c} (p_t^*)^{-\varepsilon} dj \\ &= \Pi_t^{\varepsilon} \int_{\Theta_t} p_{t-1}(j)^{-\varepsilon} dj + (1-\theta)(p_t^*)^{-\varepsilon} \\ &= \theta \Pi_t^{\varepsilon} D_{t-1} + (1-\theta)(p_t^*)^{-\varepsilon}. \end{aligned} \tag{C.9}$$

Finally, Aggregate profits are defined as:

$$\begin{aligned} q_t &\equiv \int_0^1 p_t(j)y_t(j) - (1-\tau)w_t h_t(j)dj \\ &= y_t - (1-\tau)w_t h_t. \end{aligned}$$

Demand Substitute in the real budget constraint of the household (C.1), the bond clearing condition $b_t = b_{t+1} = 0$, the above definition of aggregate profits, and the government budget constraint $T_t = \tau w_t h_t$ to get:

$$y_t = c_t. \tag{C.10}$$

C.5 Closing the model

Monetary authority The nominal risk-free interest rate i_t is determined by the monetary authority according to a conventional Taylor-type rule:

$$1 + i_t = (1 + i_{ss}) \left(\frac{\Pi_t}{\bar{\Pi}}\right)^{\rho_\pi} \left(\frac{x_t}{\bar{x}}\right)^{\rho_x} m_t,$$

where $\bar{\Pi}$ denotes the central bank's inflation target. The variable

$$x_t \equiv \frac{y_t}{y_t^f}$$

is the output gap, defined as actual output relative to the flexible-price allocation y_t^f. The parameter $\bar{x}$ represents the target output gap, and m_t captures exogenous, unanticipated deviations from the policy rule.

Exogenous shocks The process for exogenous preferences a_t is

$$a_t = \rho_a a_{t-1} + \sigma_a \varepsilon_t^a, \qquad \varepsilon_t^a \sim \mathcal{N}(0,1), \;\; \rho_a \in (0,1), \tag{C.11}$$

where ε_t^a can be interpreted as a preference shock. Moreover, the process for technology is

$$z_t = \rho_z z_{t-1} + \sigma_z \varepsilon_t^z, \qquad \varepsilon_t^z \sim \mathcal{N}(0,1), \tag{C.12}$$

where ε_t^z can be interpreted as a technology shock.

C.6 Flexible price model

To close the model, it is useful to characterize the equilibrium allocation under flexible prices. In this environment, all intermediate–good firms can adjust their price every period and therefore set the optimal constant markup:

$$p_t^f = \mathcal{M}(1-\tau)\,\frac{w_t^f}{e^{z_t}},$$

and since all firms choose the same price, the relative price satisfies $p_t^f = 1$, implying $D_t^f = 1$. Aggregate production therefore reduces to

$$y_t^f = e^{z_t} h_t^f.$$

Combining the household's labor supply condition,

$$w_t^f = \chi \left(h_t^f\right)^{\varphi} \left(c_t^f\right)^{\gamma},$$

with market clearing $c_t^f = y_t^f$ yields the intratemporal condition determining efficient labor supply.

Combining all the previous equations, the flexible–price level of output is

$$y_t^f = [\chi \mathcal{M}(1-\tau)]^{-\frac{1}{\gamma+\varphi}} \left(e^{z_t}\right)^{\frac{1+\varphi}{\gamma+\varphi}}. \tag{C.13}$$

C.7 Graphical intuition

Figure C1 provides a stylized graphical summary of the previous discussion. The left panel illustrates the effects of an adverse preference shock, a_t, while the right panel shows the impact of a negative technology shock, z_t. In both cases, the focus is on the labor market, with equilibrium real wages, w_t, and hours worked, h_t, determined by labor supply, $L^S(a_t, z_t)$, and labor demand, $L^D(a_t, z_t)$.

An adverse preference shock leads households to desire lower consumption today relative to the future. With sticky prices, weaker demand translates into lower production and thus lower labor demand. At the same time, weaker preferences for current consumption increase labor supply through a wealth effect. As shown by the green arrows on the left panel, the net result—absent the DNWR constraint—is a mild decline in hours worked and a large fall in the real wage. This pronounced decline in real wages reduces marginal costs, inducing firms that can reoptimize prices to cut them and thereby generating a decline in inflation. Lower inflation raises the effective lower bound on real wages, thereby tightening the DNWR constraint, denoted by the horizontal blue line in the figure. When the constraint binds—that is, when the horizontal blue line lies above the equilibrium given by the intersection of the two green lines—real wages remain mechanically higher. This amplifies the decline in hours worked (and thus output), as indicated by the blue arrows, while attenuating the fall in marginal costs and, consequently, in prices and inflation.[19]

A negative technology shock reduces the marginal product of labor, leading to a decline in labor demand. At the same time, labor supply increases through a wealth effect, as lower productivity reduces households' lifetime income. As shown by the red arrows on the left panel, the net result is a mild increase in hours worked and a large fall in the real wage.[20] Although real wages fall, marginal costs increase because the associated

[19] This intuition follows the same mechanism emphasized by Dupraz et al. (2025) and Debortoli et al. (2025).

[20] In principle, the labor supply response reflects the interaction of wealth and substitution effects, which makes the direction of the shift ambiguous. However, under conventional calibrations and consistent with empirical evidence, the wealth component dominates, so that a negative technology shock typically raises hours worked while lowering real wages.

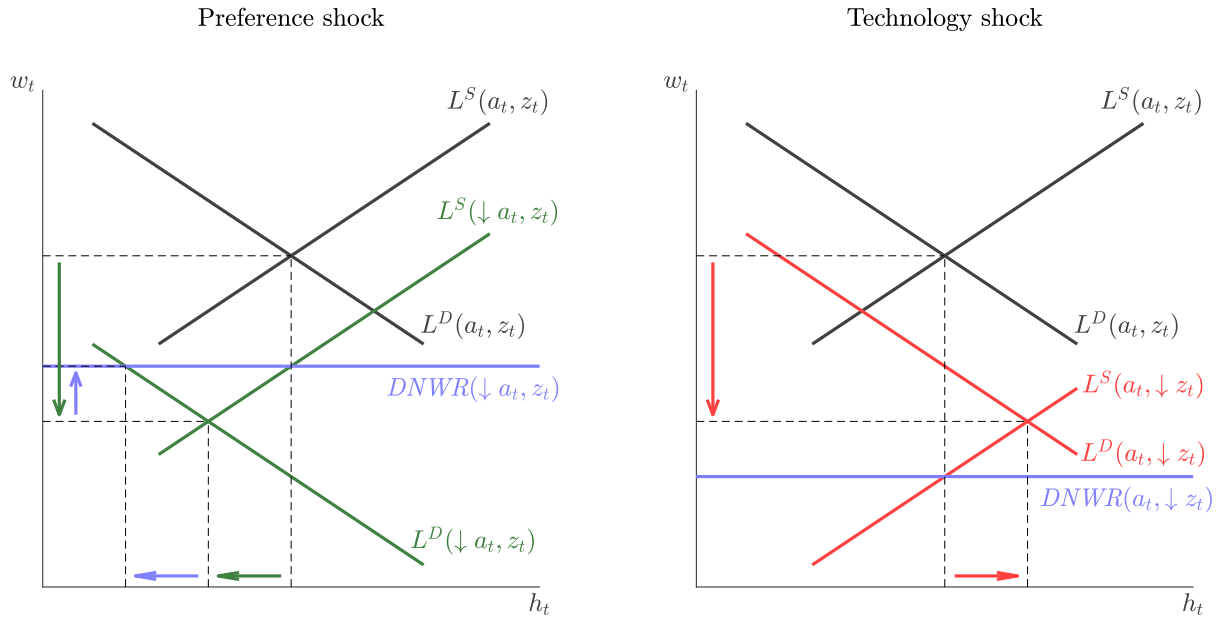


Figure C1: Model intuition

Notes. Stylized graphical summary of the model-implied labor market dynamics, with hours worked on the x-axis and the real wage w_t on the y-axis. Upward-sloping curves represent labor supply $L^S(a_t, z_t)$, downward-sloping curves represent labor demand $L^D(a_t, z_t)$, and horizontal blue lines represent the downward nominal wage rigidity constraint $DNWR(a_t, z_t)$; all curves are functions of preference a_t and technology z_t. Black curves correspond to the initial state, green curves illustrate the shifts induced by an adverse preference shock (a decline in a_t), and red curves illustrate the shifts induced by a negative technology shock (a decline in z_t).

drop in labor productivity remains the dominant force. This induces firms that can reset their prices to raise them, thereby leading to an acceleration in inflation. Higher inflation effectively loosens the DNWR constraint, making it slack—that is, the horizontal blue line lies below the equilibrium given by the intersection of the two red lines—or, if binding, less tight than under demand shocks. As a result, the asymmetry between positive and negative technology shocks is expected to be negligible relative to the asymmetry between expansionary and adverse preference shocks.

C.8 Calibration

The model is calibrated at a quarterly frequency using standard values from the New Keynesian literature. The subjective discount factor is set to $\beta = 0.99$, implying a steady-state real interest rate of approximately 4 percent at an annual rate. Preferences feature loga-

rithmic utility over consumption, corresponding to $\gamma = 1$, and an inverse Frisch elasticity of labor supply equal to $\varphi = 1$, implying a unit Frisch elasticity. The elasticity of substitution across differentiated intermediate goods is $\varepsilon = 6$, which implies a desired gross markup of $\mathcal{M} = \varepsilon/(\varepsilon - 1) = 1.2$. The Calvo price stickiness parameter is set to $\theta = 0.75$, corresponding to an average price duration of four quarters. Monopolistic distortions are offset in steady state by a labor subsidy $\tau = 1/\varepsilon$, ensuring that the flexible-price allocation is efficient. The monetary authority follows a Taylor-type rule with a response to inflation of $\rho_\pi = 1.5$, satisfying the Taylor principle, and a response to the output gap of $\rho_x = 0.5/4$. The steady-state gross inflation target is set to $\bar{\Pi} = 1.00496$, corresponding to a 2 percent annual inflation rate, while the target output gap $\bar{x} = .9993$ equals its steady state value, x_{ss}. Downward nominal wage rigidities are governed by the parameter ρ_w, which is set to one, implying that nominal wages are prevented from declining in nominal terms. The disutility-of-labor parameter $\chi = 9.0063$ in order to have steady-state hours worked, h_{ss}, equal to a third. We choose the persistence of preference and technology shocks to be identical, setting $\rho_a = \rho_z = 0.75$. The standard deviations are calibrated to $\sigma_a = 0.0065$ and $\sigma_z = 0.0018$ so that, on impact, a positive technology shock and a positive preference shock each induce an output response of approximately one quarter of a percentage point. By construction, this calibration aligns both the persistence and the impact magnitude of the two shocks, so that any differences in asymmetric responses across shocks and across shock signs can be attributed to the nature of the underlying disturbance rather than to differences in magnitude or persistence.

C.9 Numerical solution

We solve the model using a global projection method based on Chebyshev polynomials and collocation. The state vector combines exogenous processes and predetermined endogenous variables. The two exogenous states, (a_t, z_t), are discretized using a two-dimensional Tauchen method, yielding grids $\{a^i\}_{i=1}^{N_a}$ and $\{z^j\}_{j=1}^{N_z}$ and an associated Markov transition

matrix P over the $N_a \times N_z$ joint states. The endogenous predetermined states are price dispersion D_{t-1} and the lagged real wage w_{t-1}, so that the full state vector is

$$s_t \equiv (a_t, z_t, D_{t-1}, w_{t-1}).$$

Given s_t, the equilibrium is characterized by policy rules for hours worked and inflation,

$$h_t = \mathcal{H}(s_t), \qquad \Pi_t = \mathcal{P}(s_t),$$

and all remaining endogenous variables are recovered using market-clearing conditions, equilibrium identities, and optimality conditions.

Chebyshev approximation. For each discrete exogenous state (a^i, z^j), the policy functions $\mathcal{H}$ and $\mathcal{P}$ are approximated over the endogenous state space (D_{t-1}, w_{t-1}) on compact domains $D_{t-1} \in [D_{\min}, D_{\max}]$ and $w_{t-1} \in [w_{\min}, w_{\max}]$ using tensor-product Chebyshev polynomials. Let $\tilde{D}$ and $\tilde{w}$ denote the affine transformations mapping (D_{t-1}, w_{t-1}) into the interval $[-1, 1]$. For polynomial degrees n_D and n_w, the approximations take the form

$$\mathcal{H}(a^i, z^j, D_{t-1}, w_{t-1}) \approx \sum_{m=0}^{n_D} \sum_{\ell=0}^{n_w} \alpha^h_{m,\ell}(i,j)\, T_m(\tilde{D})\, T_\ell(\tilde{w}),$$

$$\mathcal{P}(a^i, z^j, D_{t-1}, w_{t-1}) \approx \sum_{m=0}^{n_D} \sum_{\ell=0}^{n_w} \alpha^\pi_{m,\ell}(i,j)\, T_m(\tilde{D})\, T_\ell(\tilde{w}),$$

where $T_m(\cdot)$ and $T_l(\cdot)$ denotes the Chebyshev polynomial of order m and l, respectively, while $\alpha^h_{m,l}(i,j)$ and $\alpha^\pi_{m,l}(i,j)$ their associated optimal coefficients.

Collocation. Collocation nodes are chosen at the Chebyshev zeros in each endogenous state dimension and mapped back to the domain $[D_{\min}, D_{\max}] \times [w_{\min}, w_{\max}]$. At each collocation node and for each exogenous state (a^i, z^j), equilibrium is imposed by setting to zero the residuals of (i) the household Euler equation for the risk-free bond and (ii) the optimal price-setting condition. Expectations are computed using the Markov transition matrix P

together with the model-implied law of motion for the endogenous state variables. The resulting system of nonlinear equations in the polynomial coefficients $\{\alpha^h_{m,\ell}(i,j), \alpha^\pi_{m,\ell}(i,j)\}$ is solved using a Newton method.

Implementation details. In the baseline calibration, the exogenous states are discretized using $N_a = N_z = 11$ grid points with bounds of ± 3 unconditional standard deviations around their means. The endogenous state domains are set to $D_{t-1} \in [1.0003, 1.002]$ and $w_{t-1} \in [0.97, 1.03]$. Policy functions are approximated using polynomial degrees $n_D = 2$ and $n_w = 7$. Accuracy is assessed by evaluating equilibrium residuals on fine grids with 200 points in each endogenous state dimension.[21]

C.10 Sign-dependent impulse response functions

We are interested in impulse responses to both *positive* and *negative* realizations of preference and technology shocks. To allow the dynamic responses of endogenous variables to differ across the sign of the shock, we compute generalized impulse response functions (Koop et al., 1996).

Generalized Impulse Response Functions. Following Koop et al. (1996), the generalized impulse response functions (GIRFs) for variable y at horizon h is defined as the difference between two conditional expectations that share the same history up to $t-1$, but differ in the realization of the time-t structural shock. Let Ω_{t-1} summarize the information set generated by the history of states up to $t-1$; let ε^j_t denote the time-t structural innovation of type j, with $j \in \{a, z\}$ indexing preference and technology shocks, respectively; and let δ denote the shock magnitude. The GIRF of variable y at horizon h, conditional on a given history Ω_{t-1}, is defined as

$$\text{GIRF}^j_y(h, \delta \mid \Omega_{t-1}) \equiv \mathbb{E}_t\left[y_{t+h} \,\middle|\, \Omega_{t-1}, \varepsilon^j_t = \delta\right] - \mathbb{E}_t\left[y_{t+h} \,\middle|\, \Omega_{t-1}, \varepsilon^j_t = 0\right], \tag{C.14}$$

[21] Increasing n_D to three or more leaves the impulse responses essentially identical.

where $h = 0, 1, \ldots, H$ and $\mathbb{E}_t[\cdot]$ denotes expectations conditional on the time-t information set. Note that if y_{t+h} is expressed in logs, then multiplying (C.14) by 100 yields the percentage deviation of y_{t+h} from its counterfactual path in the absence of the shock.

The object of interest is the unconditional GIRF that integrates (C.14) over histories drawn from the ergodic distribution:

$$\text{GIRF}_y^j(h, \delta) \equiv \mathbb{E}_{\Omega_{t-1}}\left[\text{GIRF}_y^j(h, \delta \mid \Omega_{t-1})\right], \tag{C.15}$$

where $\mathbb{E}_{\Omega_{t-1}}[\cdot]$ denotes the expectation taken with respect to the distribution of histories Ω_{t-1}.

Numerical estimation. We approximate (C.15) by Monte Carlo simulation. The procedure is as follows:

1. From $t - \tau$ to $t - 1$, simulate a single path of the exogenous Markov state using the transition matrix P. The treated and untreated economies are identical over this period.

2. At time t, both economies evolve according to the Markov transition implied by P. In addition, in the treated economy we add an exogenous innovation $\varepsilon_t^j = \delta$, which shifts the state at time t relative to the untreated economy.[22]

3. From time $t + 1$ to time $t + H$, both economies are simulated forward using the same sequence of random transitions of the exogenous Markov process. Thus, because the treatment changes the state at time t, the same transition can generate different subsequent states in the treated and untreated economies, preserving the correct stochastic evolution of the exogenous process in both cases.

[22] The magnitude of the treatment δ is set to two standard deviations of the corresponding shock. This choice ensures that, across Monte Carlo replications, the sign of the treatment is preserved with high probability despite the presence of other contemporaneous stochastic shocks.

4. Given the simulated evolution of the exogenous states, we compute—in both the treated and untreated economies—the response of the model's endogenous variables under the assumption that, at time $t - \tau - 1$, the endogenous state variables are at their steady-state values.[23]

5. For each exogenous and endogenous variable, we compute the difference between the treated and untreated economies at each horizon. For variables whose interpretation is more meaningful in percentage terms, we divide this difference by the level of the variable in the untreated economy at time t, which is kept fixed across horizons. As a result, the responses can be interpreted as percentage deviations from the untreated economy at the time of impact.

6. Steps 1–5 are repeated 20,000 times. The unconditional generalized impulse responses are then computed as the average of the differences between the treated and untreated economies.

C.11 Monte Carlo simulation

The objective of the Monte Carlo simulation is to assess whether the econometric estimator used in the empirical analysis is able to recover the structural shocks and, most importantly, their nonlinear impulse responses. To this end, we use the theoretical model as the data-generating process and apply the same econometric procedure employed in the empirical section to the simulated data. In this way, we evaluate whether our empirical approach is able to correctly capture the nonlinearities implied by the model. The procedure is as follows:

1. Simulate an economy with $T = 230$ observations by generating a sequence of random transitions of the exogenous Markov process.

[23] As usual, we assume that agents cannot anticipate subsequent transitions of the exogenous processes.

2. Construct the matrix $\mathcal{X}^{originals}$ of size $(T \times 17)$ containing the following time series: the natural logarithm of price dispersion $\log D_t$, the natural logarithm of hours worked $\log h_t$, price inflation Π_t, the natural logarithm of output $\log y_t$, the natural logarithm of consumption $\log c_t$, the policy rate R_t, the natural logarithm of the real wage $\log w_t$, real marginal cost λ_t, the real interest rate $\mathbb{E}_t\left[R_t/\Pi_{t+1}\right]$, unemployment u_t, wage inflation Π_t^w, preferences a_t, technology z_t, the first lag of preferences a_{t-1}, the first lag of technology z_{t-1}, the second lag of preferences a_{t-2}, and the second lag of technology z_{t-2}.

3. Construct the matrix $\mathcal{X}^{synthetic}$ of size $(T \times 101)$ by generating additional variables as linear combinations of the variables in $\mathcal{X}^{originals}$. Each new variable is defined as a weighted sum of these variables, where the weights are independently drawn from a uniform distribution $U(-1, 1)$.

4. Form the matrix

$$\mathcal{X} = \left[\mathcal{X}^{originals}, \mathcal{X}^{synthetic}\right],$$

   which has dimension $(T \times 118)$.

5. Generate the matrix $\tilde{\mathcal{X}}$ by adding measurement error to each variable in $\mathcal{X}$, where the standard deviation of the measurement error equals 5% of the standard deviation of the corresponding variable.

6. Apply the same econometric procedure used in the empirical analysis to $\tilde{\mathcal{X}}$ and store the estimated impulse responses.

7. Repeat steps 1–6 for 5,000 Monte Carlo replications and compute the median responses together with the 68% and 90% confidence bands.

Additional details. The dimension of $\tilde{\mathcal{X}}$ is chosen to match the size of the dataset used in the empirical analysis. In addition, the same set of weights is used across all simulated

economies. Finally, differently from the parametrization adopted in the empirical analysis, we impose nine factors and one lag in the Dynamic Factor Model.

C.12 Other results

This section presents additional results that complement the analysis in the main text. While informative, these results are sufficiently peripheral to be reported here. In particular, Figure C2 reports the impulse responses implied by an otherwise identical version of the model in which the DNWR constraint is absent. Moreover, Figure C3 displays impulse responses of additional endogenous variables in the baseline model that are not shown in the main text. Finally, Figure C4 reports additional results related to the Monte Carlo simulations.

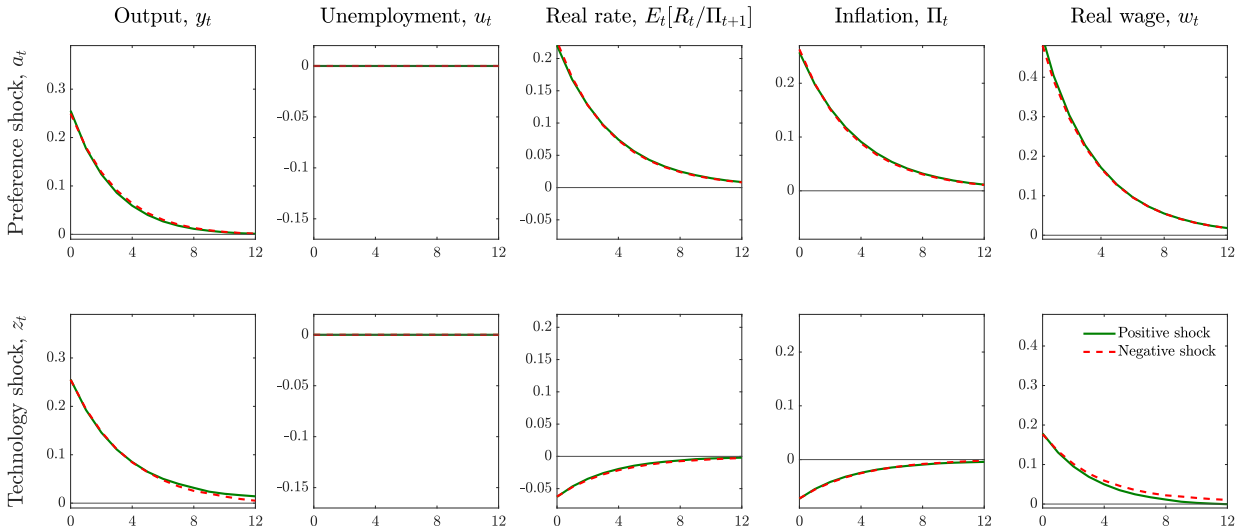


Figure C2: Model-implied impulse responses without DNWR

Notes. Impulse responses of output y_t, unemployment u_t, the real interest rate $\mathbb{E}_t[R_t/\Pi_{t+1}]$, inflation Π_t, and the real wage w_t to two–standard-deviation preference shocks (first row) and technology shocks (second row) implied by an otherwise identical model without downward nominal wage rigidities. Responses to positive shocks are shown by solid green lines, while responses to negative shocks are plotted in absolute value using red dotted lines. The figure reports unconditional generalized impulse response functions (GIRFs) à la Koop et al. (1996), computed by Monte Carlo simulation as the average difference between two economies that share the same history and sequence of future stochastic realizations, but differ in the realization of a one-time structural shock at impact. Responses of output and the real wage are expressed as percentage deviations, while responses of unemployment, the real interest rate, and inflation are reported in percentage-point deviations. Additional details on the calibration, the solution method and the definition and estimation of the GIRFs are reported to Sections C.8-C.10 of this document.

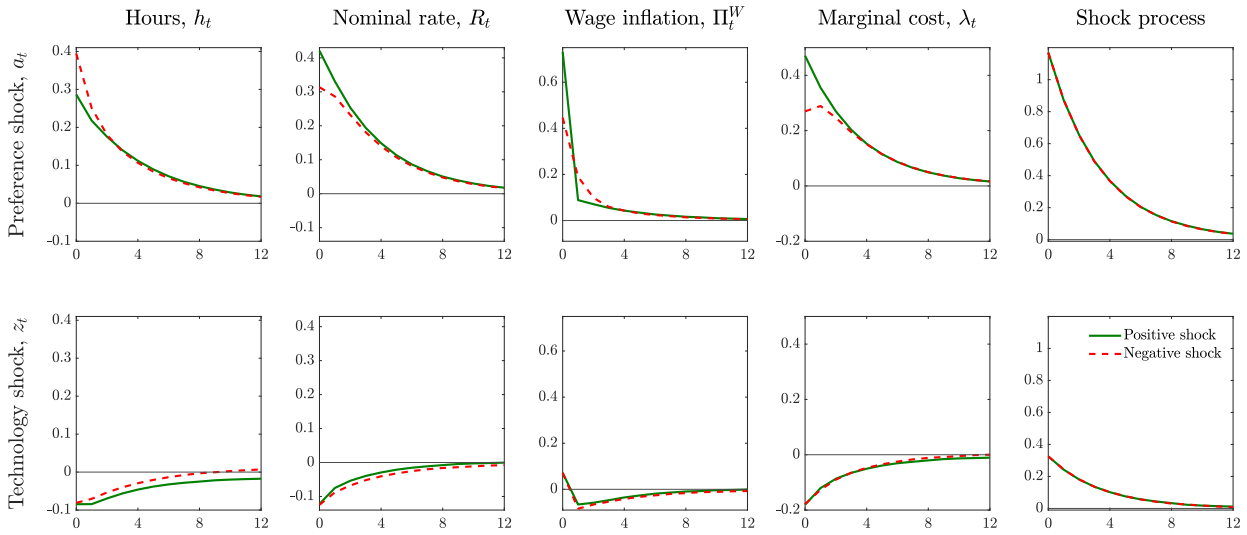


Figure C3: Model-implied impulse responses of additional variables

Notes. Impulse responses of hours h_t, nominal rate R_t, nominal wage inflation Π_{t+1}^W, marginal cost λ_t, and the shock processes a_t/z_t to two–standard-deviation preference shocks (first row) and technology shocks (second row). Responses to positive shocks are shown by solid green lines, while responses to negative shocks are plotted in absolute value using red dotted lines. The figure reports unconditional generalized impulse response functions (GIRFs) à la Koop et al. (1996), computed by Monte Carlo simulation as the average difference between two economies that share the same history and sequence of future stochastic realizations, but differ in the realization of a one-time structural shock at impact. Responses of hours, marginal cost and shock processes are expressed as percentage deviations, while responses of the nominal rate and nominal wage inflation are reported in percentage-point deviations. Additional details on the calibration, the solution method and the definition and estimation of the GIRFs are reported to Sections C.8-C.10 of this document.

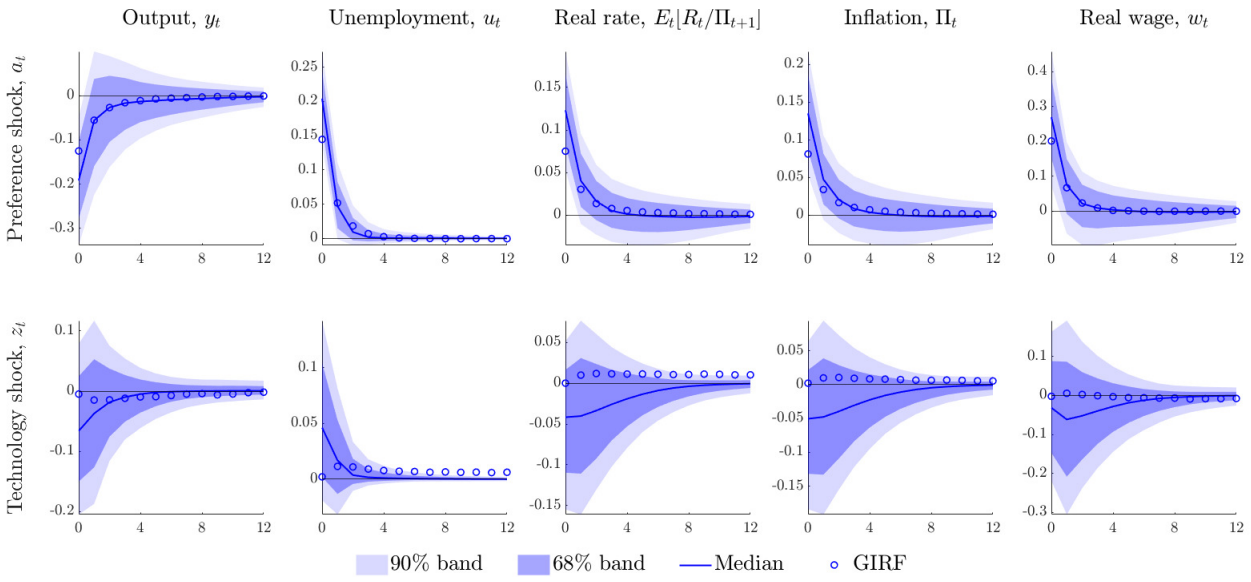


Figure C4: Monte Carlo simulation

Notes. Differences between the impulse responses to two–standard-deviation positive and negative preference shocks (first row) and to two–standard-deviation positive and negative technology shocks (second row) for output y_t, unemployment u_t, the real interest rate $\mathbb{E}_t[R_t/\Pi_{t+1}]$, inflation Π_t, and the real wage w_t. The circles denote the differences implied by the unconditional generalized impulse response functions (GIRFs) à la Koop et al. (1996). The solid lines and shaded areas represent the median difference and the corresponding 68% and 90% percentile bands obtained from the Monte Carlo simulations. Responses of output and the real wage are expressed as percentage deviations, while responses of unemployment, the real interest rate, and inflation are reported in percentage-point deviations. Additional details on the calibration, the solution method, and the Monte Carlo simulation are provided in Sections C.8, C.9, and C.11.