YES YOUNG ECONOMISTS SEMINAR

HRVATSKA NARODNA BANKA
EUROSUSTAV

19th YOUNG ECONOMISTS' SEMINAR
TO THE 32nd DUBROVNIK ECONOMIC CONFERENCE

May 28 – 29, 2026, Dubrovnik, Croatia

Marthe Mareels

The Long-Run Effects of Temporary Oil Supply Disruptions in European Countries

Draft version

Please do not quote

The Long-Run Effects of Temporary Oil Supply Disruptions in European Countries

Marthe Mareels*

April 2026

Abstract

This paper provides evidence that supply-driven oil price increases depress GDP for more than a decade after oil prices have returned to their pre-shock level. These hysteresis effects are driven by two channels: a contraction in investment, and a steady decline in labour force participation. In contrast, oil price decreases do not produce sustained economic gains, indicating asymmetric effects. Comparing these dynamics to those following persistent demand shocks, I find that these hysteresis effects are caused by similar mechanisms.

*Ghent University (Email: marthe.mareels@ugent.be). I would like to thank Luca Alfieri, Arthur Apostel, Tobias Broer, Jan De Loecker, Monika De Ridder, Sandra Eickmeier, Gerdie Everaert, Francesco Furlanetto, Freddy Heylen, Lukas Lehner, Mariano Montoya, Lorenzo Mori, Marvin Nöller, Gert Peersman, Michele Piffer, Christophe Van Langenhove, Gianluca Violante and workshop participants at Ghent University, the 18th UNITO - Collegio Carlo Alberto PhD Workshop in Economics, the 6th ERMEES European Macroeconomics Workshop, the 19th RGS doctoral conference, and the 21th Doctorissimes conference for helpful comments and suggestions.

1 Introduction

Hysteresis refers to the idea that to the idea that transitory economic disturbances can have very long-lived or even permanent effects, which challenges the conventional separation between trend and cycle. Existing research on this topic is however mainly focused on demand-shock hysteresis. Empirical evidence on long-run effects of temporary supply shocks remains scarce, even though there is reason to expect these shocks to affect the economy persistently (Blanchard, 2018; Fornaro and Wolf, 2023; Coibion et al., 2017).

The demand-shock hysteresis literature highlights two main mechanisms behind such persistence. First, a labour market channel: prolonged unemployment erodes workers' skills and discourages job search, leading to lasting detachment from the labour market (Pissarides, 1992; Huckfeldt, 2022). Second, a capital formation channel: because investment is pro-cyclical, temporary downturns reduce capital accumulation and productivity, generating scarring effects that outlast the initial shock (Dixit, 1992; Anzoategui et al., 2019).

Supply shocks may reinforce these mechanisms and thereby amplify hysteresis effects. Unlike negative demand shocks, which can be counteracted by monetary easing, an adverse supply shock limits the central bank's ability to stimulate the economy, as it coincides with rising inflation. At the same time, unfavourable supply shocks increase firms' input costs and may compress their markups. The combination of higher interest rates and lower profit margins may tighten firms' financial constraints, discouraging investment and hence resulting in a slower recovery (Ma and Zimmermann, 2023; Beqiraj et al., 2025). Yet despite these theoretical reasons to expect stronger or more persistent effects, direct empirical evidence on supply-shock hysteresis remains limited.

This paper studies whether temporary supply shocks can have persistent effects on economic activity, and through which channels these effects operate. I focus on European countries in a panel local projections setting, using an externally identified oil supply shock which causes oil prices to rise sharply in the short run before returning to their pre-shock level after one year. This shock can hence be seen as a *temporary supply disruption* for non-oil producing countries, and makes them particularly well-suited for studying the possibility of hysteresis effects: if output remains depressed after oil prices have reverted, this provides evidence for supply-shock hysteresis.

I find that following a temporary oil price increase, real GDP decreases and stays depressed for over a decade, long after oil prices have returned to their pre-shock level. Investment falls sharply and fails to recover, pointing to a capital formation channel through which productivity is durably reduced. These results align with the predictions of the endogenous

growth theory, supporting the idea that temporary supply shocks can have persistent effects through depressed investment and productivity (Fornaro and Wolf, 2023). At the same time, the unemployment rate increases and the labour force participation rate decreases steadily, which points to an additional role for labour market hysteresis.

I compare these dynamics to hysteresis effects following demand shocks. I find that degree of persistence and the relative importance of the transmission channels are broadly similar across the two shock types, despite theoretical predictions that supply shocks should generate stronger persistence. While monetary policy responds less accommodative to supply shocks, profit margins prove more resilient, suggesting that firms pass through a substantial share of cost increases to prices. This limits the squeeze on margins and offsets the tighter interest rate environment: although supply shocks constrain monetary easing, the sharper margin compression under demand shocks leads to a comparable tightening of firms' financial conditions. As a result, investment responses, and hence persistence, are similar across the two shock types.

The effects are asymmetric. I find no evidence for positive hysteresis: temporary favourable shocks do not yield long-run gains. The persistent effects are entirely driven by adverse oil supply shocks, consistent with Jordà et al. (2024) and Aikman et al. (2022), but in contrast to studies that document lasting gains from expansions (Girardi et al., 2020; Bluedorn and Leigh, 2019; Ball and Onken, 2022; Carnevale and Di Francesco, 2025). Labour market rigidities prevent cost adjustment on the downside, while the combination of asymmetric monetary policy and profit margin dynamics suppresses the investment response on the upside, ruling out positive hysteresis in European economies.

The rest of this paper is structured as follows. Section 2 discusses the existing literature. Section 3 outlines the empirical strategy and data used. Section 4 sets out the main results, Section 5 compares the results with demand-shock hysteresis, and Section 6 tests asymmetries from the sign of the shock. Section 7 concludes.

2 Literature

2.1 Transmission Mechanisms

The literature has proposed two main mechanisms to explain hysteresis effects: (i) a labour market channel, and (ii) a capital accumulation channel.[1]

[1] For a comprehensive overview of the hysteresis literature development, see Cerra et al. (2023).

Labour market channel A large literature highlights the labour market as a key transmission channel for hysteresis. A traditional strand emphasizes institutional features, such as labour unions, employment protection, and unemployment benefits, that can make wage-setting insufficiently responsive to labour market slack, thereby sustaining elevated wages and unemployment following adverse shocks (Blanchard and Summers, 1986; Galí, 2022; Blanchard and Wolfers, 2000; Campos et al., 2025; Schoefer, 2025). A complementary mechanism operates through duration and signalling effects. Prolonged unemployment erodes workers' skills and human capital, while discouragement leads people to exit the labour force entirely, reducing productive capacity (Pissarides, 1992; Topel, 1991; Clark and Summers, 1982). Additionally, job displacement may signal lower ability or motivation to employers, further reducing re-employment prospects. Empirical evidence at both the micro and macro level supports the importance of these channels: persistent output losses are often accompanied by lasting declines in labour force participation and shifts in the natural rate of unemployment, indicating that cyclical downturns can have durable effects on labour market outcomes (Kroft et al., 2013; Davis and Von Wachter, 2011; Huckfeldt, 2022; Ball et al., 1999; Ball, 2009; Ball and Onken, 2022; Furlanetto et al., 2025).

Capital accumulation channel A more recent channel is proposed by the endogenous growth literature. Because investment is pro-cyclical, a temporary downturn reduces capital formation through lower physical investment, R&D, and learning-by-doing, leading to a sustained shortfall in productive capacity that outlasts the initial shock (Dixit, 1992; Fatas, 2000; Benigno and Fornaro, 2018; Anzoategui et al., 2019). Jordà et al. (2024), using monetary policy shocks for a panel of pegged countries, and Cloyne et al. (2022), for temporary tax cuts in the U.S., document long-run effects that operate through the capital accumulation channel, leading to persistent productivity losses

In a recent theoretical contribution to this literature, Fornaro and Wolf (2023) argue that the idea of hysteresis should not be limited to demand shocks, as temporary supply shocks may also have persistent economic effects. In their framework, a transitory negative supply shock reduces firms' expected profitability and raises their cost of funds, leading to lower investment. Since investment drives productivity growth in their endogenous growth model, the result is a permanent shift in the level of output. This contrasts with the New Keynesian framework where temporary supply shocks have only transitory effects on real activity. They also show that this scarring effect depresses aggregate demand through a negative wealth effect on households, and that it can prolong inflationary pressures by sustaining firms' marginal costs.

Building on this, there are reasons to expect supply shocks to amplify the capital accumu-

lation channel relative to demand disturbances. Unlike negative demand shocks, which can be offset by monetary easing, adverse supply shocks coincide with rising inflation, constraining the central bank's ability to stimulate the economy. At the same time, higher input costs may compress firms' profit margins, weakening internal financing. The combination of tighter monetary conditions and reduced profitability can amplify financial constraints, leading firms to cut back investment more sharply and persistently. These mechanisms suggest that temporary supply shocks may generate stronger and more persistent investment shortfalls, and thereby possibly reinforce hysteresis effects.

2.2 Asymmetries in Hysteresis

Hysteresis is traditionally viewed negatively, as it is associated with long recessions and persistent economic slack. Some however argue that the mechanisms causing hysteresis could also prevail in the opposite case, i.e. a positive shock to the economy could lead to a lower unemployment rate or heightened investments for an extended period, which may cause persistent positive effects (Girardi et al., 2020; Cloyne et al., 2022; Carnevale and Di Francesco, 2025). For example, temporary positive shocks could generate a "hot economy" where firms invest more or where the unemployed get drawn into the labour market, accumulating human capital.

3 Data & Methodology

I construct a panel dataset of European countries covering the period 1975Q1 - 2019Q4 with quarterly data from the OECD Economic Outlook. For the baseline results I restrict the sample to the countries that were OECD members in 1975. I exclude the oil-producing countries in my panel. This choice is motivated by two key considerations. First, as shown by Känzig (2021) and Mori and Peersman (2024), oil supply news shocks have only a temporary impact on oil prices, while global oil production adjusts sluggishly and persistently. As a result, oil-producing countries may experience such shocks as persistent supply shocks, whereas non-producing countries face them as temporary cost-push shocks, raising input prices in the short run. Second, an increase in oil prices may be good news for oil producers, resulting in a favourable shock for this industry. As a baseline rule I exclude the countries where the maximum oil rents-to-GDP ratio exceeds 1% in at least one year during the estimation period (1975–2019). This leaves me with a panel of 15 countries.[2] The results are robust to this threshold.

I estimate impulse responses using Local Projections (LP), which offer two key advan-

[2] The European countries in the baseline panel are Austria, Belgium, Switzerland, Germany, Spain, Finland, France, Greece, Ireland, Iceland, Italy, Luxembourg, Netherlands, Portugal and Sweden. Countries that are excluded because of their oil production are the UK, Norway and Denmark.

tages for this analysis. First, they produce more robust inference at long horizons than VARs, which are known to be sensitive to the persistence of the data (Montiel Olea and Plagborg-Møller, 2021). Second, they provide a flexible framework for investigating potential nonlinearities in the impulse responses. Specifically, I estimate the following Local Projections specification:

$$y_{i,t+h} - y_{i,t-1} = \alpha_i^h + \beta^h S_t + \theta^h \Delta X_{i,t} + \varepsilon_{i,t+h} \tag{1}$$

for horizon $h = 0, 1, 2, ..., H$ and country i at time t.

$y_{i,t+h}$ is the variable of interest, log real GDP, and $X_{i,t}$ a set of lagged controls, including the outcome variable as well as lagged values of other macro variables.[3] I calculate heteroskedasticity and autocorrelation robust Driscoll and Kraay (1998) standard errors.[4]

S_t is a temporary cost-push shock. This type of shock is particularly well-suited for examining supply-shock hysteresis as it is a temporary supply shock that does not affect economic fundamentals directly. Following Fornaro and Wolf (2023) I look at energy price shocks, identified here as oil price shocks in the spirit of Känzig (2021) and augmented with financial variables as suggested in Mori and Peersman (2024). By measuring the changes in oil futures prices from OPEC announcements in a sufficiently tight window, Känzig (2021) isolates the impact of news about future oil supply. Using this as an external instrument in a VAR model, he identifies a structural oil supply news shock. Mori and Peersman (2024) however show that these structural shocks are Granger-caused by financial variables, and address this omitted variable bias by incorporating a set of financial indicators into the VAR model. For shock S_t I use the identified monthly structural shock series aggregated to the quarterly frequency, rather than the raw surprise series. This avoids the power problem that arises when aggregating monthly surprises to the quarterly frequency (Bilbiie and Känzig, 2023), and allows me to estimate impulse responses starting in 1975, whereas the surprise series only starts in 1983.

4 Evidence for Supply-Shock Hysteresis

The figures in this section display cumulative impulse responses to a one standard deviation increase in the oil supply news shock corresponding to an 8% oil price increase

[3] The control variables are the exchange rate, log GDP deflator, the unemployment rate, log real gross fixed capital formation, log labour productivity, the labour force participation rate, log oil prices and log oil production. I hold the sample length constant for all h.

[4] Canova (2024) documents that if the units of a panel are dynamically heterogeneous, the local projection estimator is inconsistent. I test for this using his proposed Hausman test for each horizon, and I cannot reject the null of dynamic homogeneity. Additionally, Figure A.2.3 reports the IRFs for the individual countries in my panel.

after one quarter, estimated over a 10-year (40-quarter) horizon. As the responses are cumulative, they represent the change in variables in (log) levels relative to the baseline. The solid line shows the point estimate, with dark and light shaded areas indicating 68 and 90 percent confidence bands, respectively. All baseline results use 4 lags of the dependent and explanatory variables.

4.1 Long-Run Effects on Economic Activity

A necessary condition for identifying hysteresis is that the shock itself is transitory. That is, I want to investigate whether there is a persistent response of economic activity after the price shock has dissipated. The first panel of Figure 1 plots the impulse response of the oil price to a one standard deviation increase in the oil supply news shock. The response shows that the oil price returns to its pre-shock baseline within approximately four quarters, confirming the temporary nature of the shock. This provides a suitable setting to test for hysteresis: if output remains depressed after the shock has faded, it would indicate existence of hysteresis effects.

Following an oil price increase, output falls on impact and contracts by 1% after 5 years (Figure 1). The response remains persistently below the baseline throughout the entire horizon – long after the oil price has returned to its initial level. After 10 years, real GDP is estimated to be 1% lower than would have been the case in absence of an oil price shock. Notably, following the initial decline, GDP neither continues to fall nor rebounds, but instead stabilizes at a lower level. This indicates that the *growth rate* of GDP has returned to its pre-shock path (Figure A.5.1) but without the temporary acceleration in growth that would compensate for the earlier loss, as predicted by standard New Keynesian models. The finding that GDP does not continue to decrease also implies that there is no evidence of *superhysteresis*, in which not only the level but also the long-run growth rate of output is permanently affected (see, for example, Ball (2014)).

To assess how much of the variation in real GDP is driven by oil supply shocks, I report the Forecast Error Variance Decomposition in Figure A.5.2, estimated following Gorodnichenko and Lee (2020). This shows that the oil price shock accounts for around 10% of the long-run variation, which is an economically meaningful share of the forecast error variance of real GDP, and supports the interpretation that oil supply shocks are a quantitatively relevant source of the persistent output dynamics documented above.

Figure 1: Impulse responses to an oil supply news shock

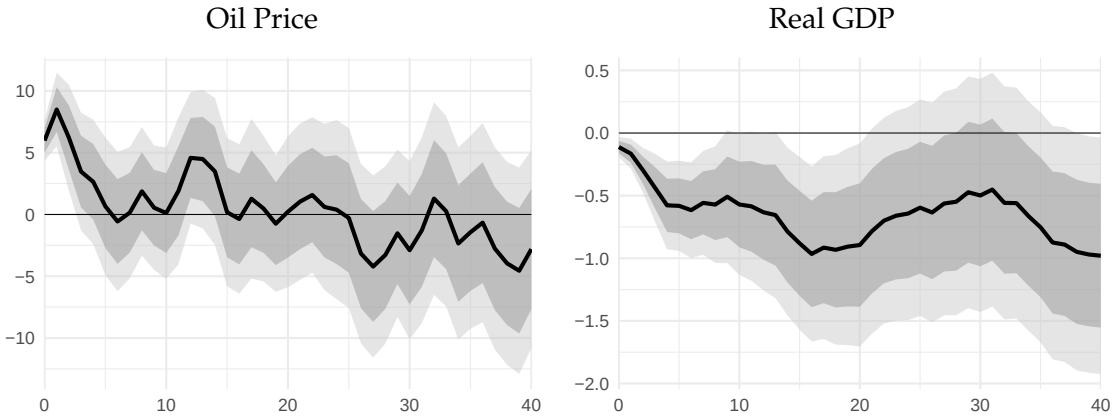

Note: The figure shows impulse responses to an oil supply news shock. Results are estimated over the period 1975Q1-2019Q4 using 4 lags of the dependent variable and control variables. The solid line is the point estimate and the dark and light shaded areas are 68 and 90 percent confidence bands, respectively.

4.2 Transmission Mechanisms

To understand how transitory price shocks can cause a persistent fall in real GDP, I compare the responses of other macroeconomic variables to the oil price shocks (Figure 2).

A first observation is that the unemployment rate increases immediately after the oil price shock, and stays elevated over the horizon. This is accompanied by a lagged, but continuous, decline in the labour force participation rate, indicating that people leave the labour market and do not return over time. These patterns are consistent with the classic hysteresis effects documented in response to demand-driven recessions, where prolonged slack leads to labour market detachment (Yagan, 2019; Furlanetto et al., 2025).

The results also point to an important role for the capital formation channel. Investment falls by roughly 1.5 percent relative to the baseline and remains persistently depressed. Notably, this decline exceeds the contraction in real GDP, implying a sustained fall in the investment-to-GDP ratio. Firms therefore actively scale back capital accumulation beyond what lower output alone would imply, pointing to a structural shift in investment behaviour rather than a purely cyclical response.

This aligns with the mechanism in Fornaro and Wolf (2023), where a temporary adverse supply shock compresses expected profitability and raises firms' cost of funds, reducing incentives to invest. In their endogenous growth framework, this investment shortfall translates into a permanent reduction in productive capacity. Consistent with this, I find that gross operating profits — a measure of economy-wide profit volume, capturing the margin between output prices and labour costs scaled by real GDP — decline significantly

Figure 2: Impulse responses to an oil supply news shock

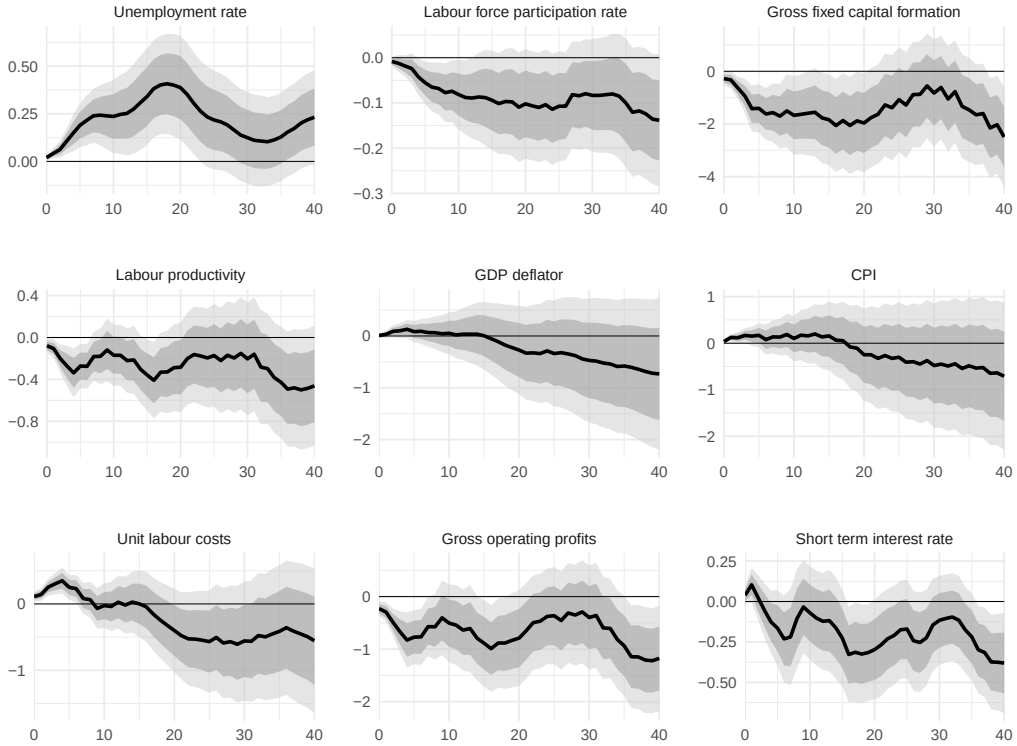

Note: Impulse response function of the unemployment rate, the labour force participation rate (of total population aged 15-74), gross fixed capital formation, labour productivity (output per worker), the GDP deflator, CPI, unit labour costs, gross operating profits and the short term interest rate to an oil price shock, with 68 and 90% CI. Results are estimated over the period 1975Q1-2019Q4. Horizon of the responses is quarterly. Additional estimation results are shown in Figure A.5.3.

and persistently, and that labour productivity contracts by approximately 0.3 percent within the first year.[5] While point estimates remain consistently negative thereafter, the wide error bands call for caution in interpreting these longer-run productivity effects.

The response of prices is more ambiguous. The CPI increases in the short run, but the responses are not statistically significant at longer horizons and gradually decline over the estimated horizon. Similarly, the GDP deflator, which reflects the prices of domestically produced goods, and hence do not include the oil prices, increases in the short run, suggesting that firms are able to pass through at least part of the cost increase to output prices, limiting the squeeze on profit margins. However, it does not remain significantly

[5] Gross operating profits are constructed as real GDP multiplied by a proxy for the profit margin, namely the difference between the GDP deflator and unit labour costs. This difference serves as a measure of the aggregate profit margin: the GDP deflator captures the price level of domestically produced output, while unit labour costs capture the labour cost per unit of output, so their difference approximates the margin firms retain after covering labour costs. Multiplying by real GDP scales this margin to the economy-wide level, yielding a measure of total profit volume. Figure A.5.3 in Appendix A.5 shows the IRFs for the components of this measure.

elevated at longer horizons, which may partly reflect demand-side dampening.

Notably, monetary policy responds with an easing stance following the shock. Rather than tightening in response to the short-run inflationary pressure, policymakers appear to look through the supply disruption and prioritise supporting activity. This is consistent with the broader empirical literature on monetary responses to oil shocks (Kilian and Lewis, 2011; De Winne and Peersman, 2016; Castelnuovo et al., 2024).[6] Since monetary tightening was one of the mechanisms through which supply shocks were expected to generate stronger hysteresis than demand shocks, its absence raises the question of whether the two shock types ultimately produce a different degree of persistence — a question I address directly in Section 5.

4.3 The Investment and Productivity Channel: Evidence from Annual Data

The quarterly results point to a sustained decline in capital formation, raising the question of how this translates into lasting damage to the economy's productive capacity. I gather data from the Penn World Tables on total factor productivity (TFP) and the capital stock, as well as data on registered patents (residential and non-residential) from the World Bank. The patent data is reported from 1980 onwards. However, both of these dataseries are reported at the annual frequency. In order to examine the responses of these investment-related variables, I aggregate the oil price shocks to the annual level to then estimate the effects on the annual frequency. Figure 3 shows the impulse responses for capital formation, the real capital stock, patents registered, total factor productivity and labour productivity.

Figure 3: Impulse responses to an oil supply news shock

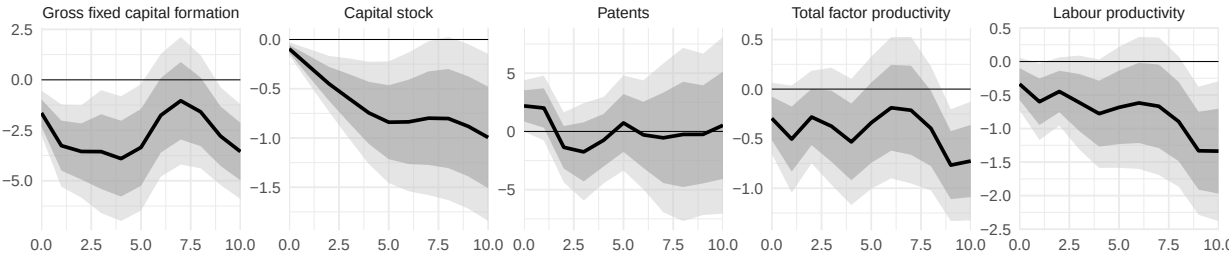

Note: Impulse responses to an oil supply news shock. Results are estimated over the period 1980-2019 using 1 lag of the dependent variable and control variables. The solid line is the point estimate and the dark and light shaded areas are 68 and 90 percent confidence bands, respectively. Horizon of the responses is annual.

[6] The lower interest rates may also help explain why the sustained inflationary pressure predicted by Fornaro and Wolf (2023) is not observed in the data, as in their framework the monetary tightening amplifies the price increase caused by the fall in productivity.

The capital stock decreases strongly, with -1% after 10 years with regards to the baseline scenario, in line with the earlier results on capital formation. Patent registrations do not react to the oil price shock, suggesting that the slowdown in investment does not extend to innovative activity. Both total factor productivity and labour productivity exhibit negative responses, reinforcing the evidence that temporary supply shocks can leave lasting scars on the productivity of the economy.

4.4 Robustness

I show the results using more lags in appendix A.1, in which case the conclusions become even stronger. I further show that the results are robust to including late OECD joiners to the baseline panel. To account for structural changes in oil market dynamics, I perform a robustness check in which the estimation period is restricted to the post-Volcker disinflation era, in line with the findings of Baumeister and Peersman (2013) who document decreasing oil price elasticities over time. These results, along with those from alternative estimation windows, are gathered in Appendix A.2.

Additional robustness tests include using the Känzig (2021) and the Baumeister and Hamilton (2019) oil supply shocks rather than the Mori and Peersman (2024) oil supply news shock. The conclusions are robust to these alternative specifications (see Appendix A.3).

5 Comparing Supply- and Demand-Shock Hysteresis

The baseline results establish that temporary supply shocks generate persistent output losses through both capital formation and labour market channels. A natural follow-up question is whether these hysteresis effects differ in degree or mechanism from those produced by demand shocks. Theory suggests they might: adverse supply shocks constrain monetary easing while potentially compressing firms' profit margins through higher input costs. Both forces tighten firms' financial conditions and could amplify the investment contraction relative to what a demand shock of similar size would produce.

However, the results in Section 4.2 showed that monetary policy did not tighten following the oil supply shock, raising the question of whether these theoretical amplification mechanisms translate into a meaningfully different degree of persistence in practice. To answer this, I benchmark the supply-shock results against those from an identified demand shock that also causes hysteresis. I first test whether the overall degree of persistence differs, and then decomposing the transmission mechanisms to understand why.

I test whether the relative persistence of output differs between supply and demand shocks,

or in other words, whether supply-induced hysteresis is more persistent than demand-induced hysteresis. To benchmark my supply-shock results against those from demand disturbances, I use the permanent demand shock of Furlanetto et al. (2025). This shock is identified using long-run and sign restrictions and is characterized by a non-temporary effect on output in the US data. Given the size of the US economy and its central role in global trade and financial markets, this shock plausibly represents an (external) demand impulse that is common across European countries, making it a natural benchmark for the European panel studied here.

I re-estimate Specification (1), now including both the oil supply news shock S_t^S and the permanent demand shock S_t^D on the right-hand side. Because of the availability of the demand shock, IRFs are estimated over the period 1985Q1-2019Q4. The resulting impulse response functions are shown in Figure 4.

Figure 4: Impulse responses of real GDP to an oil supply news shock and a permanent demand shock

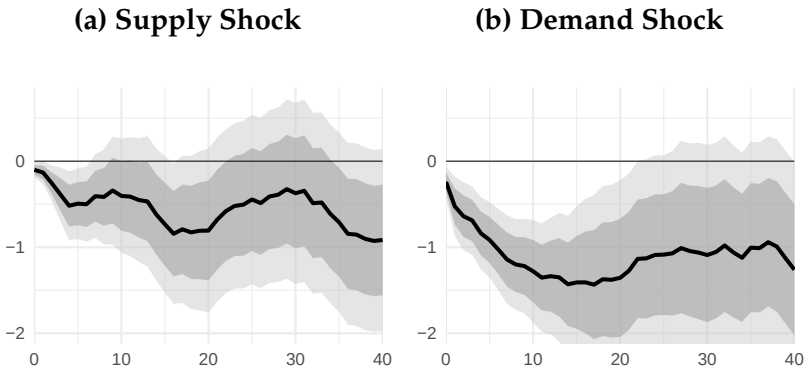

Note: Impulse responses to an oil supply news shock (panel a) and a permanent demand shock (panel b). Results are estimated over the period 1985Q1-2019Q3 using 4 lags of the dependent variable and control variables. The solid line is the point estimate and the dark and light shaded areas are 68 and 90 percent confidence bands, respectively. Horizon of the responses is quarterly.

A first noteworthy finding is that the identified permanent US demand shock also produces persistent effects in European economies. The IRFs show that, while the initial output response is larger following a demand shock, the long-horizon responses converge. This suggests that supply shocks may display relatively stronger persistence conditional on their initial impact. However, formal tests reported in Appendix A.5.4 find no statistically significant difference in persistence between the two shock types. I therefore conclude that temporary supply shocks do not generate a meaningfully different degree of hysteresis in output than demand shocks.

Next, in order to uncover more about the relative importance of different transmission mechanisms, I rescale the impulse responses of key transmission variables by the absolute value of the peak GDP response for each shock type. This approach isolates the relative importance of each transmission mechanism conditional on the size of the business-cycle output decline.

To better understand this similarity in the degree of hysteresis across shock types, I rescale the impulse responses of key transmission variables by the absolute value of the peak GDP response for each shock type. This normalization allows me to compare the mechanisms conditional on the overall size of the initial output contraction.

Figure 5: Impulse responses to demand and supply shocks

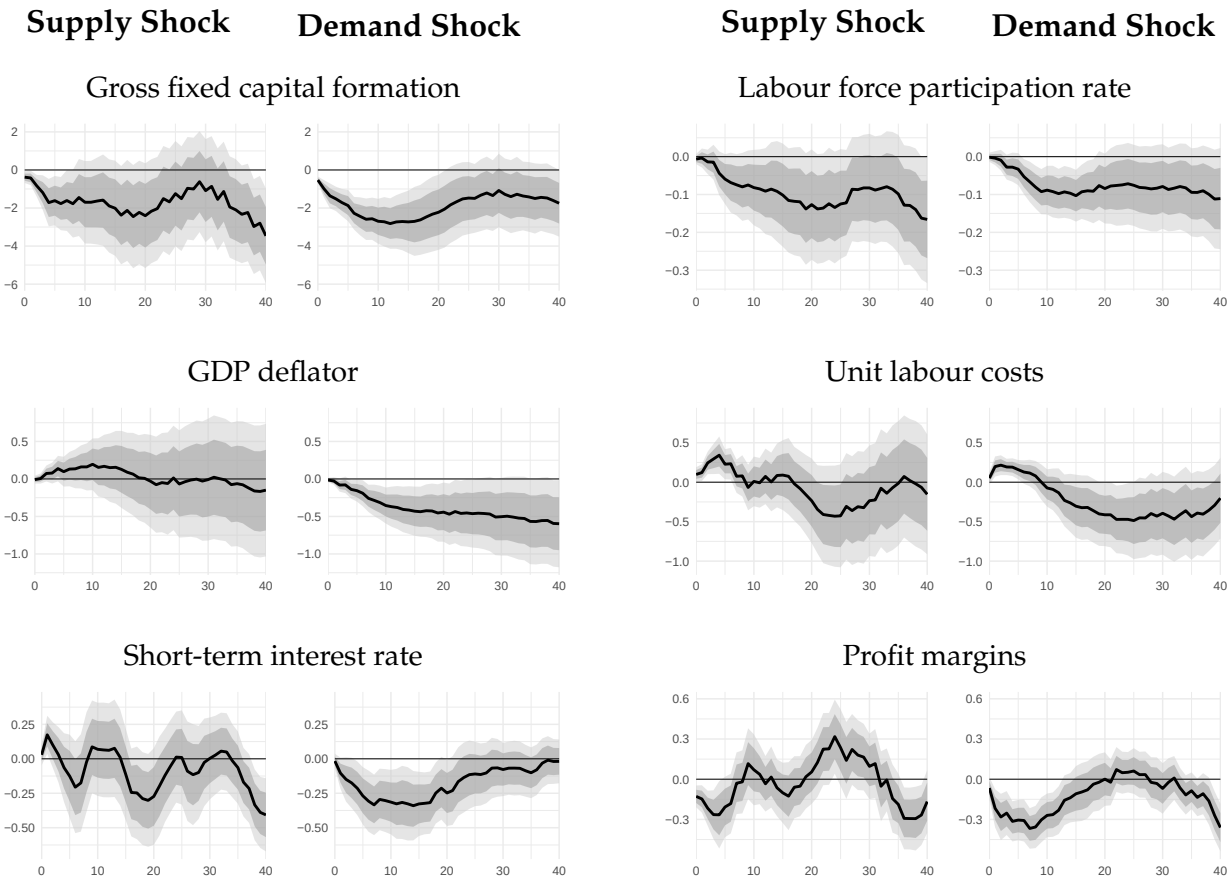

Note: Shaded areas denote 68% and 90% confidence intervals. In each variable pair, the left panel is the supply shock and the right panel is the demand shock. Dependent variables are rescaled with the peak response of real GDP to the respective shock. Results are estimated over the period 1985Q1-2019Q4. All responses are shown at a quarterly frequency.

Both for investment as for the labour force participation rate, the normalized responses are broadly comparable across the two shock types. Both shocks generate a persistent decline in employment of similar magnitude relative to the output loss. In Appendix A.5.4 I estimate the difference between the response functions, and the difference between the two

is centred around zero throughout the horizon. This suggests that, per unit of GDP lost at its peak response, the employment or investment adjustment does not depend on whether the recession was triggered by a supply or demand disturbance. Taken together, these results suggest that supply shocks causing hysteresis move through similar mechanisms as demand shocks causing hysteresis, and that there is no important difference in the severity of the scarring between the two shocks.

The similarity of the investment responses across supply and demand shocks may appear surprising, given that a more constrained monetary policy response following supply shocks would be expected to depress investment more strongly. Two partially offsetting mechanisms help reconcile this theoretical prediction with the empirical findings (Figure 5). First, consistent with the view that monetary policy is more constrained following adverse supply disturbances, the short-term interest rate falls more sharply after a demand shock than a supply shock. Importantly, however, rates do not tighten following supply shocks as theory would predict, limiting the extent to which monetary policy amplifies the investment contraction. Second, however, profit margins – proxied by the difference between the GDP deflator and unit labour costs – compress more after a demand shock than a supply shock. Although unit labour costs respond similarly across both shock types, the demand shock generates a sharp decline in output prices, squeezing margins to a greater degree. This counteracts the theoretical amplification mechanism: while supply shocks constrain monetary easing, the sharper margin compression under demand shocks broadly equalises the overall tightening of firms' financial conditions. Together, these two forces explain why the normalised investment responses, and by extension the degree of hysteresis, do not differ significantly across shock types.

6 One-Sided Hysteresis: Asymmetric Effects of Oil Price Shocks

The baseline results in previous sections implicitly assume that the economic response to oil price shocks is symmetric across the sign of the shock. This would mean that both oil price increasing as oil price decreasing shocks have persistent output effects. In this section, I relax this restriction and explicitly test whether the hysteresis mechanisms documented above operate symmetrically.

To investigate this question empirically, I explicitly distinguish between positive and negative price shocks to assess whether hysteresis mechanisms operate symmetrically. More specifically, I estimate following specification:

$$y_{i,t+h} - y_{i,t-1} = \alpha_i^h + \beta_p^h S_t^p + \beta_n^h S_t^n + \theta^h \Delta X_{i,t} + \varepsilon_{i,t+h} \tag{2}$$

for horizon $h = 0, 1, 2, ..., H$ and country i.

Structural shock $S_{i,t}$ is divided in its positive and negative counterpart, S_t^p and S_t^n respectively. For the positive (negative) shock series, I take the value of S_t if $S_t > (<)\ 0$ and 0 if otherwise. S_t^p then represents an oil supply news shock that increases oil prices. Following Kilian and Vigfusson (2011), both shock series enter the specification simultaneously to avoid the estimation biases that arise from censored regressions, which have been shown to produce inconsistent and exaggerated estimates of the effects of positive shocks when negative values are omitted.

To facilitate comparison, I scale both impulse responses so that they represent the effect of an oil price increase. The response to a negative shock is therefore multiplied by -1. The impulse responses for the positive and negative oil supply news shocks both indicate a temporary effect on oil prices, and can hence be interpreted as a temporary cost-push shock (Figure 6a). One caveat is worth noting: the negative shock exhibits a sign reversal at longer horizons, though this is not statistically distinguishable from zero. This should be kept in mind when interpreting the longer-run responses.

I formally test whether the oil price responses from the positive and negative shocks are significantly different from each other, using the method proposed by Forni et al. (2024). I estimate following equation,

$$y_{i,t+h} - y_{i,t-1} = \alpha^h + \beta^h S_t + \gamma^h |\boldsymbol{S}_t| + \theta^h \Delta X_t + \varepsilon_{t+h} \tag{3}$$

for horizon $h = 0, 1, 2, ..., H$

In this specification, S_t represents the oil supply news shock and captures the symmetric (linear) effect on oil prices. The inclusion of $|S_t|$ allows the response to vary depending on the sign of the shock. In other words, it tests for asymmetries between positive and negative shocks. A statistically significant γ^h would indicate that positive (S_t^p) and negative (S_t^n) shocks have asymmetric effects on oil prices. The resulting impulse responses (Figure 6b) do not point to meaningful asymmetries.

The left-side panel of Figure 7 shows the effect on the dependent variables for a shock that increases the oil price (positive shock) and a shock that decreases the oil price (negative shock). Recalling that both IRFs are expressed in terms of an equivalent oil price increase, a divergence between the two lines directly reveals asymmetry in the economic response.

Both shocks affect real GDP symmetrically in the short run, leading to a 0.5% change in real GDP in the first year. This is consistent with the findings of Kilian and Vigfusson (2011) that do not find evidence for short-term asymmetric responses of the US economy to oil

Figure 6: Testing the transitory and symmetric nature of the oil supply news shocks

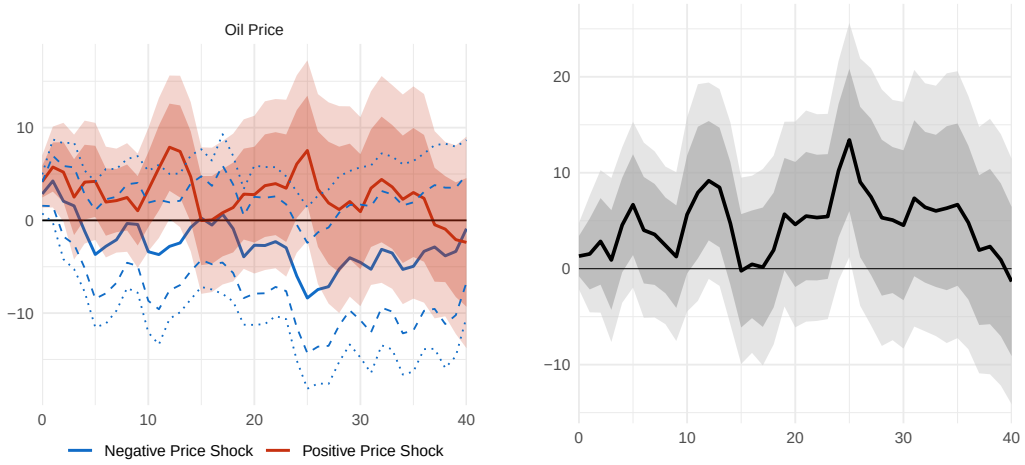

(a) Positive and negative price shocks

(b) Asymmetry test (γ^h)

Note: The first panel shows the impulse responses of a positive (red) and negative (blue) oil supply news shock. Both impulse responses are scaled to represent the effect of a one standard deviation oil price increase. The second panel shows the impulse responses for the absolute value of an oil supply news shock, ie γ^h in specification (3). This tests whether positive and negative shocks affect the oil price asymmetrically, where significance would mean asymmetric effects.

price increases and decreases. However, after a favourable (price decreasing or negative) shock, real GDP returns to the baseline after 3 years, while the effect from an unfavourable (price increasing or positive) shock remains persistent over the full estimation horizon. The right-side panel of Figure 7 shows the results of the asymmetry tests (γ^h in Specification (3)). We can see that the responses become significantly different from each other at the 68% confidence level after 10 quarters. This implies that the asymmetry in the responses does not lie in the initial response to the oil price shock, but rather that there is an asymmetry in the hysteresis mechanisms at play.

Both the response of investments as well as the labour force participation rate are larger and more persistent following an oil price increase than for an oil price decrease. Similarly, labour productivity reacts clearly to price-increasing shocks, but shows no significant longer-run response to price-decreasing shocks. Examining the other macroeconomic variables, the responses seem to be symmetric in the short-run, whereas only the positive shocks cause significant long-term responses.

The asymmetry in long-run output effects is reflected in several underlying variables. On the labour market, nominal wages rise quickly and significantly following a favourable shock, but do not fall following an adverse shock, pointing to downward nominal wage rigidity. Additionally, hours worked per worker fall significantly after an oil price increase, whereas they do not rise after a price decrease. These asymmetric responses in wages and

Figure 7: Impulse responses testing asymmetry

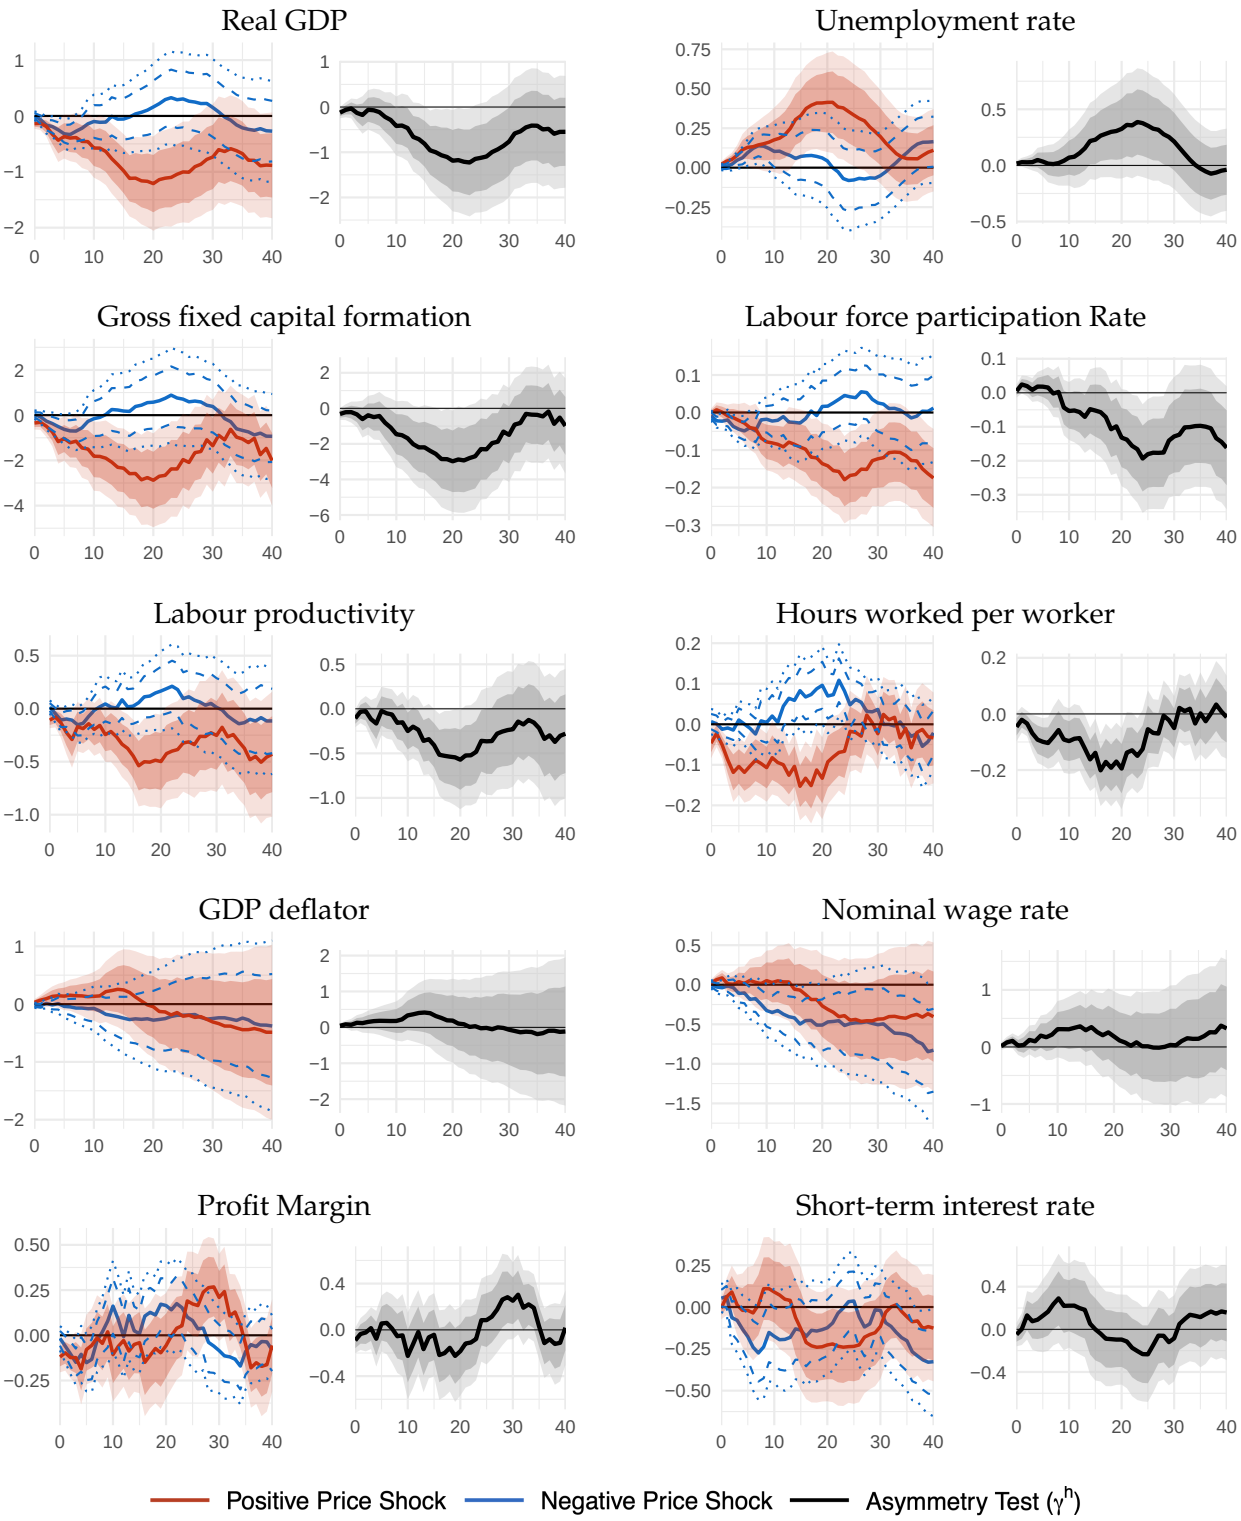

Note: The first panels show the impulse response function for positive (red) and negative (blue) oil price shocks. Both impulse responses are scaled to represent the effect of a one standard deviation oil price increase. The second panels show the impulse responses for the absolute value of the oil supply news shock, i.e. γ^h in equation 3.

hours worked point to labour market rigidities as an important mechanism preventing positive hysteresis in European economies.

Two additional asymmetries show the important role for the asymmetric investment response. First, monetary policy responds asymmetrically: a favourable shock triggers a tightening, while an adverse shock is not followed by an easing. This limits the extent to which firms can benefit from cheaper financing following a positive supply shock, dampening the investment response on the upside. Second, profit margins compress more sharply following an adverse shock than they increase following a favourable one. The downwardly rigid wages documented above maintain high labour costs for firms even as output falls, squeezing margins and discouraging investment. Following a favourable shock, by contrast, the rapid rise in wages offsets the gains from lower input costs, again limiting margin improvement and preventing the positive feedback loops that would generate lasting gains in investment and productivity. Additionally, it has been documented that external finance constraints typically bind more tightly following output contractions than expansions (Balke, 2000; Manaresi and Pierri, 2024), reinforcing the asymmetry: firms face tighter funding conditions precisely when margins are most compressed, further discouraging investment on the downside while leaving the upside relatively unaffected..

Together, these forces explain why adverse supply shocks leave lasting scars while favourable shocks do not. Labour market rigidities prevent cost adjustment on the downside, while the combination of asymmetric monetary policy and profit margin dynamics suppresses the investment response on the upside, ruling out positive hysteresis in European economies.

The asymmetry documented here also has broader theoretical implications, since standard macroeconomic models often impose linearity, implying that the direction of a shock should not affect the persistence of its impact. However, my results add to a growing body of empirical research challenges this assumption, suggesting that the economy may respond asymmetrically to positive and negative disturbances (Forni et al., 2024; Jordà et al., 2024).

7 Conclusion

While previous empirical work has focused on demand-shock hysteresis, this paper provides new empirical evidence that temporary supply disruptions can have persistent effects on macroeconomic outcomes. Using panel local projections with an externally identified shock, I document that temporary increases in oil prices lead to sustained reductions in output, investment, productivity and labour market participation, long after

prices have reverted to their pre-shock levels. These findings stand in contrast to the predictions of standard New Keynesian models, which posit that temporary supply shocks only have transitory effects on real activity.

The results identify two key channels through which supply-shock hysteresis materializes. First, a *capital formation channel*: investment falls sharply after the shock and remains subdued, leading to a persistent decline in the capital stock, and contributing to persistently lower labour productivity and total factor productivity. Second, a *labour market channel*: the unemployment rate rises and the participation rate declines persistently, pointing to labour market detachment consistent with classic hysteresis theories.

I explore asymmetries in the effects of these shocks. The evidence suggests that hysteresis is primarily driven by adverse shocks; favourable supply shocks do not generate similar long-lasting gains. This asymmetry underscores the risks associated with even temporary adverse supply disruptions and the limits of high-pressure economies in generating persistent benefits.

Overall, this paper challenges the conventional wisdom that transitory oil supply shocks cannot affect long-run economic activity, and hence highlights the need for policymakers to consider the lasting consequences of transitory supply disturbances. The relevance of these findings extends beyond the historical sample. The sharp rise in energy prices following Russia's invasion of Ukraine in 2022, and more recently the renewed geopolitical tensions in the Middle East stemming from the conflict in Iran, serve as stark reminders that temporary oil supply disruptions remain a live concern for European policymakers. Although oil prices have partially retreated from their peaks, the results of this paper suggest that the damage may already be accumulating through subdued investment and rising labour market detachment. My findings suggest that the standard toolkit of demand-side stabilization may be insufficient in mitigating the long-run costs of such shocks, pointing instead to the need for policies that directly support investment and labour market reattachment in the aftermath of supply-driven recessions.

References

Aikman, D., Drehmann, M., Juselius, M., and Xing, X. (2022). The Scarring Effects of Deep Contractions. *Bank of Finland Research Discussion Paper, No. 12*.

Anzoategui, D., Comin, D., Gertler, M., and Martinez, J. (2019). Endogenous technology adoption and r&d as sources of business cycle persistence. *American Economic Journal: Macroeconomics*, 11(3):67–110.

Balke, N. S. (2000). Credit and economic activity: credit regimes and nonlinear propagation of shocks. *Review of Economics and Statistics*, 82(2):344–349.

Ball, L. (2014). Long-term damage from the great recession in oecd countries. *European Journal of Economics and Economic Policies*, 11(2):149–160.

Ball, L., Mankiw, N. G., and Nordhaus, W. D. (1999). Aggregate demand and long-run unemployment. *Brookings papers on economic activity*, 1999(2):189–251.

Ball, L. and Onken, J. (2022). Hysteresis in unemployment: evidence from oecd estimates of the natural rate. *International Finance*, 25(3):268–284.

Ball, L. M. (2009). Hysteresis in unemployment: Old and new evidence. Working Paper 14818, National Bureau of Economic Research.

Baumeister, C. and Hamilton, J. D. (2019). Structural interpretation of vector autoregressions with incomplete identification: Revisiting the role of oil supply and demand shocks. *American Economic Review*, 109(5):1873–1910.

Baumeister, C. and Peersman, G. (2013). Time-varying effects of oil supply shocks on the us economy. *American Economic Journal: Macroeconomics*, 5(4):1–28.

Benigno, G. and Fornaro, L. (2018). Stagnation traps. *The Review of Economic Studies*, 85(3):1425–1470.

Beqiraj, E., Cao, Q., Minetti, R., and Tarquini, G. (2025). Persistent slumps: Innovation and the credit channel of monetary policy. *European Economic Review*, 172:104946.

Bilbiie, F. O. and Känzig, D. R. (2023). Greed? profits, inflation, and aggregate demand. Technical report, National Bureau of Economic Research.

Blanchard, O. (2018). Should we reject the natural rate hypothesis? *Journal of Economic Perspectives*, 32(1):97–120.

Blanchard, O. and Wolfers, J. (2000). The role of shocks and institutions in the rise of european unemployment. *Monetary Policy and Unemployment*.

Blanchard, O. J. and Summers, L. H. (1986). Hysteresis and the european unemployment problem. *NBER Chapters*, pages 15–90.

Bluedorn, J. C. and Leigh, D. (2019). Hysteresis in labor markets? evidence from professional long-term forecasts. International Monetary Fund.

Campos, N. F., Eichenauer, V. Z., and Sturm, J.-E. (2025). Are unemployment differentials among advanced economies still explained by the shocking institutions hypothesis? *Economic Systems*.

Canova, F. (2024). Should we trust cross-sectional multiplier estimates? *Journal of Applied Econometrics*, 39(4):589–606.

Carnevale, O. P. and Di Francesco, D. (2025). Are hysteresis effects nonlinear? *Available at SSRN 5520619*.

Castelnuovo, E., Mori, L., and Peersman, G. (2024). Commodity price shocks and global cycles: Monetary policy matters. Cesifo working paper, CESifo.

Cerra, V., Fatás, A., and Saxena, S. C. (2023). Hysteresis and business cycles. *Journal of Economic Literature*, 61(1):181–225.

Clark, K. B. and Summers, L. H. (1982). Labour force participation: timing and persistence. *The review of economic studies*, 49(5):825–844.

Cloyne, J., Martinez, J., Mumtaz, H., and Surico, P. (2022). Short-term tax cuts, long-term stimulus. Technical report, National Bureau of Economic Research.

Coibion, O., Gorodnichenko, Y., and Ulate, M. (2017). The cyclical sensitivity in estimates of potential output. Technical report, National Bureau of Economic Research.

Davis, S. J. and Von Wachter, T. M. (2011). Recessions and the cost of job loss. Technical report, National Bureau of Economic Research.

De Winne, J. and Peersman, G. (2016). Macroeconomic effects of disruptions in global food commodity markets: Evidence for the united states. *Brookings Papers on Economic Activity*, 2016(2):183–286.

Dixit, A. (1992). Investment and Hysteresis. *Journal of Economic Perspectives*, 6(1):107–132.

Driscoll, J. C. and Kraay, A. C. (1998). Consistent covariance matrix estimation with spatially dependent panel data. *Review of economics and statistics*, 80(4):549–560.

Fatas, A. (2000). Do business cycles cast long shadows? short-run persistence and economic growth. *Journal of economic growth*, 5:147–162.

Fornaro, L. and Wolf, M. (2023). The scars of supply shocks: Implications for monetary policy. *Journal of Monetary Economics*.

Forni, M., Gambetti, L., Maffei-Faccioli, N., and Sala, L. (2024). Nonlinear transmission of financial shocks: Some new evidence. *Journal of Money, Credit and Banking*, 56(1):5–33.

Furlanetto, F., Lepetit, A., Robstad, Ø., Rubio-Ramírez, J., and Ulvedal, P. (2025). Estimating hysteresis effects. *American Economic Journal: Macroeconomics*, 17(1):35–70.

Galí, J. (2022). Insider–outsider labor markets, hysteresis, and monetary policy. *Journal of Money, Credit and Banking*, 54(S1):53–88.

Girardi, D., Paternesi Meloni, W., and Stirati, A. (2020). Reverse hysteresis? persistent effects of autonomous demand expansions. *Cambridge Journal of Economics*, 44(4):835–869.

Gorodnichenko, Y. and Lee, B. (2020). Forecast error variance decompositions with local projections. *Journal of Business & Economic Statistics*, 38(4):921–933.

Huckfeldt, C. (2022). Understanding the Scarring Effect of Recessions. *American Economic Review*, 112(4):1273–1310.

Jordà, Ò., Singh, S. R., and Taylor, A. M. (2024). The long-run effects of monetary policy. *Review of Economics and Statistics*, pages 1–49.

Känzig, D. R. (2021). The macroeconomic effects of oil supply news: Evidence from opec announcements. *American Economic Review*, 111(4):1092–1125.

Kilian, L. and Lewis, L. T. (2011). Does the Fed Respond to Oil Price Shocks? *The Economic Journal*, 121(555):1047–1072. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1468-0297.2011.02437.x.

Kilian, L. and Vigfusson, R. J. (2011). Are the responses of the us economy asymmetric in energy price increases and decreases? *Quantitative Economics*, 2(3):419–453.

Kroft, K., Lange, F., and Notowidigdo, M. J. (2013). Duration dependence and labor market conditions: Evidence from a field experiment. *The Quarterly journal of economics*, 128(3):1123–1167.

Ma, Y. and Zimmermann, K. (2023). Monetary policy and innovation. Technical report, National Bureau of Economic Research.

Manaresi, F. and Pierri, N. (2024). The asymmetric effect of credit supply on firm-level productivity growth. *Journal of Money, Credit and Banking*, 56(4):677–704.

Montiel Olea, J. L. and Plagborg-Møller, M. (2021). Local projection inference is simpler and more robust than you think. *Econometrica*, 89(4):1789–1823.

Mori, L. and Peersman, G. (2024). Estimating the macroeconomic effects of oil supply news. Technical report, CESifo Working Paper.

Pissarides, C. A. (1992). Loss of skill during unemployment and the persistence of employment shocks. *The quarterly journal of economics*, 107(4):1371–1391.

Schoefer, B. (2025). Eurosclerosis at 40: Labor market institutions, dynamism, and european competitiveness. Technical report, National Bureau of Economic Research.

Topel, R. (1991). Specific capital, mobility, and wages: Wages rise with job seniority. *Journal of political Economy*, 99(1):145–176.

Yagan, D. (2019). Employment hysteresis from the great recession. *Journal of Political Economy*, 127(5):2505–2558.

A Appendix

A.1 Regressions with more lags of control variables

Figure A.1.1: 5 lags of dependent and control variables

Note: Impulse responses for an oil supply news shock, with 68 and 90% CI. Results are estimated over the period 1975Q1-2019Q4, using 5 lags of the dependent variable and control variables. Horizon of the responses is quarterly.

Figure A.1.2: 8 lags of dependent and control variables

Note: Impulse responses for an oil supply news shock, with 68 and 90% CI. Results are estimated over the period 1975Q1-2019Q4, using 8 lags of the dependent variable and control variables. Horizon of the responses is quarterly.

A.2 Alternative Samples

Figure A.2.1: Impulse responses to an oil supply news shock for alternative estimation periods

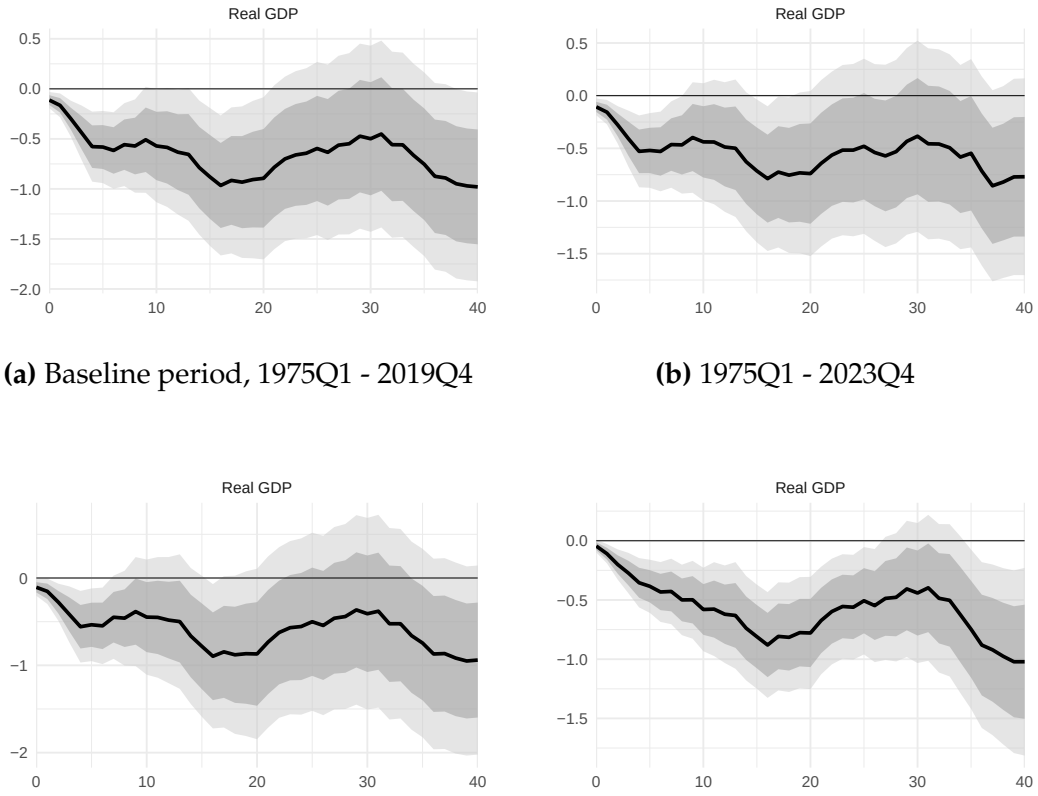

(a) Baseline period, 1975Q1 - 2019Q4

(b) 1975Q1 - 2023Q4

(c) 1985Q1 - 2019Q4

(d) 1970Q1 - 2015Q4

Note: The panel (a) shows the baseline results. Panel (b)-(d) show the impulse responses for the results estimated over alternative sample periods.

Figure A.2.2: Impulse responses to an oil supply news shock for alternative panels

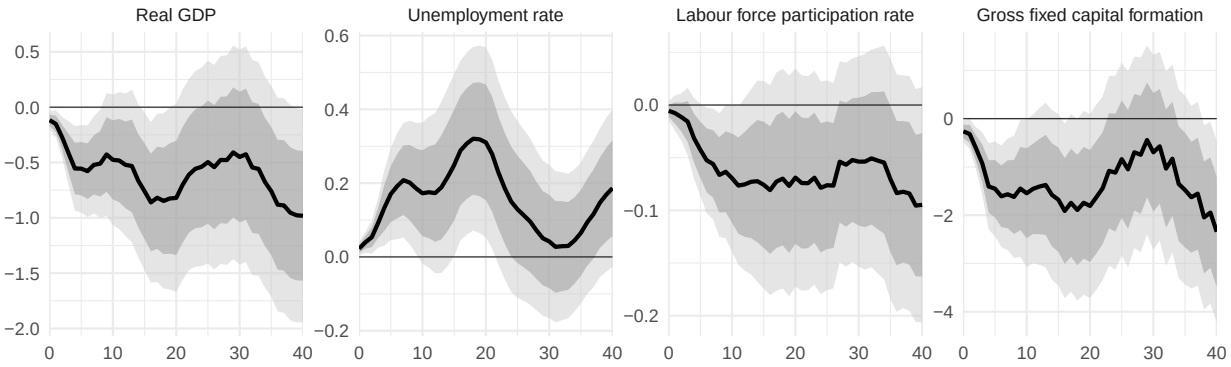

(a) Including European OECD member countries that joined after 1975

Note: Impulse responses are estimated over the baseline period 1975Q1 - 2019Q4 for alternative samples of countries.

Figure A.2.3: Impulse responses of real GDP to an oil supply news shock for individual countries in the baseline sample

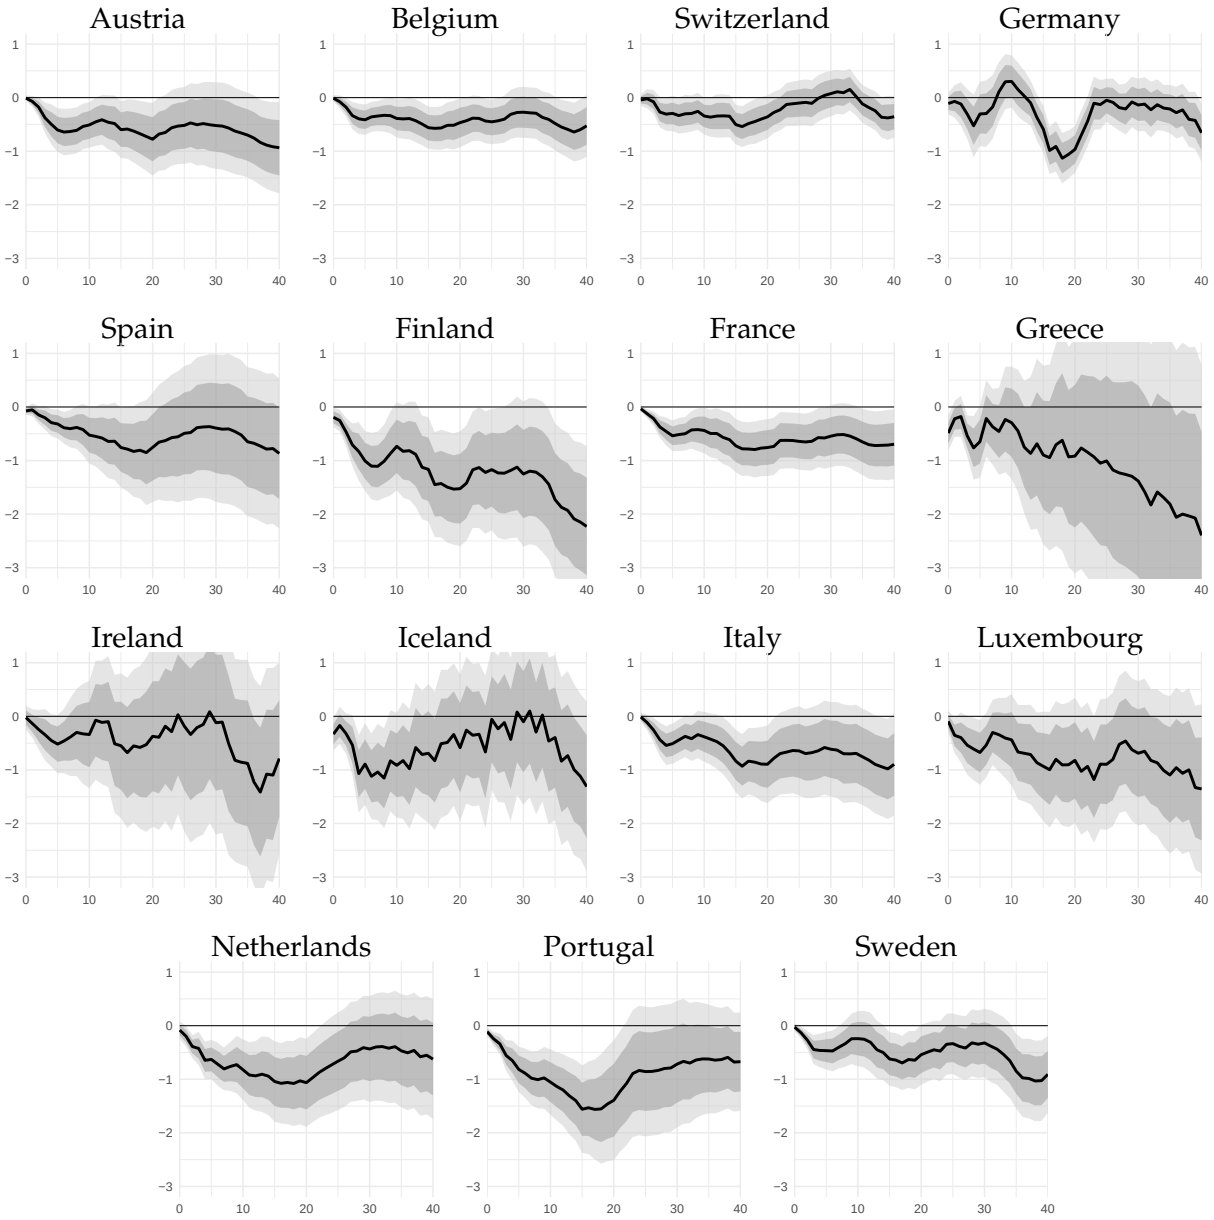

Note: Impulse responses are estimated over the baseline period 1975Q1 - 2019Q4 for each individual country separately. Due to data constraints, the results for Germany are estimated from its reunfication in 1990 onwards. I include 4 lags of the dependent variable and control variables (oil prices, oil production, the exchange rate, the GDP deflator and the unemployment rate).

Figure A.2.4: Impulse responses to an oil supply news shock for individual oil-producing countries

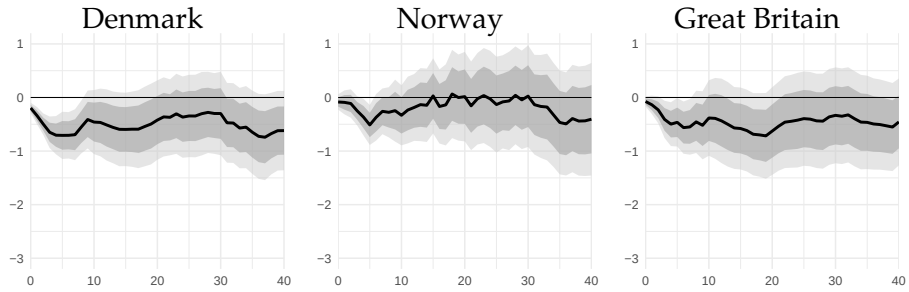

Note: Impulse responses are estimated over the baseline period 1975Q1 - 2019Q4 for each individual country separately. I include 4 lags of the dependent variable and control variables (oil prices, oil production, the exchange rate, the GDP deflator and the unemployment rate).

A.3 Alternative shock identification

Figure A.3.1: Impulse responses to an oil shock as identified by Känzig (2021), for estimation period 1975Q1-2019Q4

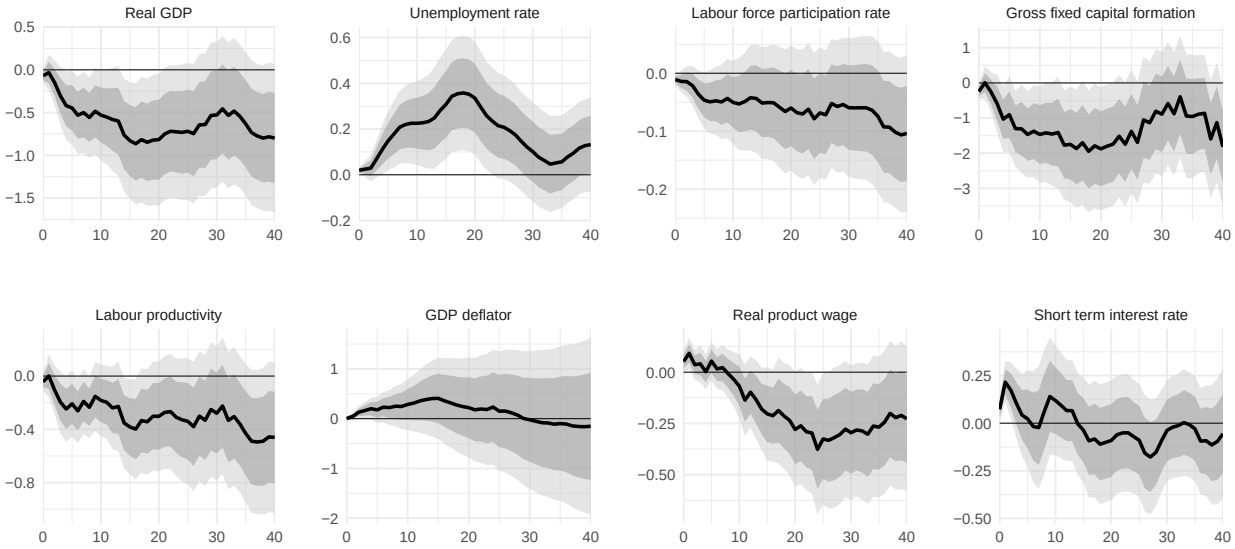

Note: Impulse response function for a Känzig (2021) oil shock, with 68 and 90% CI. Results are estimated over the period 1975Q1-2019Q4 using 4 lags of the dependent variable and control variables. Horizon of the responses is quarterly.

Figure A.3.2: Impulse responses to an oil shock as identified by Baumeister and Hamilton (2019), for estimation period 1985Q1-2019Q4

Note: Impulse response function for a Baumeister and Hamilton (2019) oil shock, with 68 and 90% CI. Results are estimated over the period 1985Q1-2019Q4 using 4 lags of the dependent variable and control variables. Horizon of the responses is quarterly.

A.4 Regressions with annual data

Figure A.4.1: Impulse responses to an oil supply news shock, annual

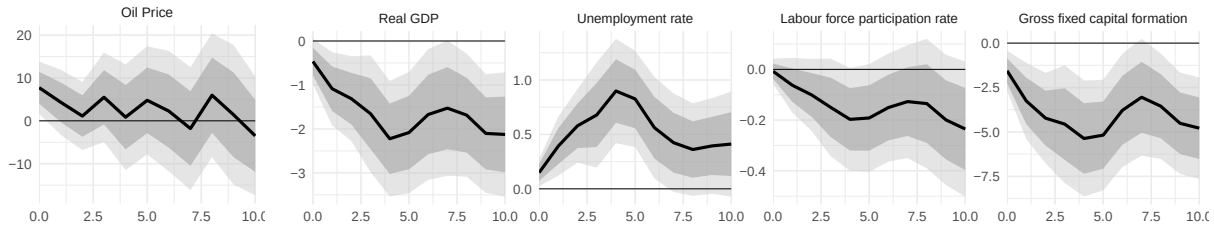

Notes: Impulse responses to an oil supply news shock. Results are estimated over the period 1980Q1-2019Q4 using 1 lag of the dependent variable and control variables. The solid line is the point estimate and the dark and light shaded areas are 68 and 90 percent confidence bands, respectively. Horizon of the responses is annual.

A.5 Additional Results

A.5.1 Testing for Superhysteresis

Figure A.5.1: Impulse response of real GDP growth to an oil supply news shock

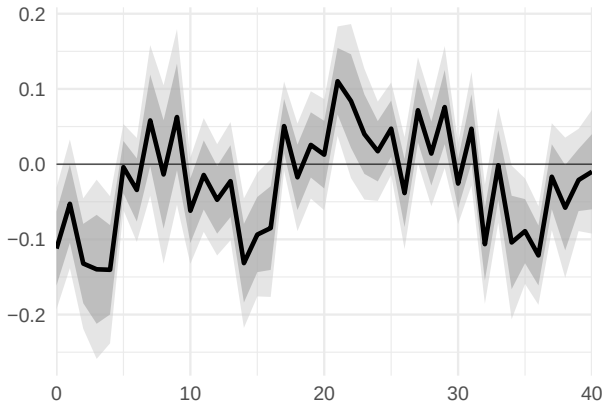

Note: Impulse responses to an oil supply news shock. Results are estimated over the period 1975Q1-2019Q4 using 4 lags of the dependent variable and control variables. The solid line is the point estimate and the dark and light shaded areas are 68 and 90 percent confidence bands, respectively.

A.5.2 Forecast Error Variance Decompositions

Figure A.5.2: Forecast Error Variance Decomposition for Real GDP

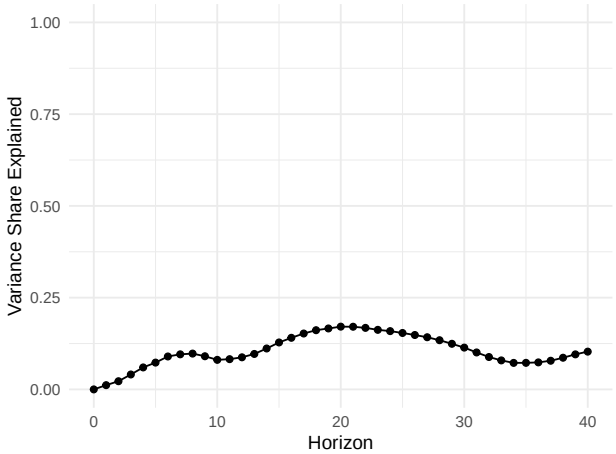

Note: The table reports the share of forecast error variance explained by the identified shock. Values are estimated following the method of Gorodnichenko and Lee (2020) and reported on a scale of 0 to 1.

A.5.3 Additional results from baseline specification

Figure A.5.3: Impulse responses to an oil supply news shock

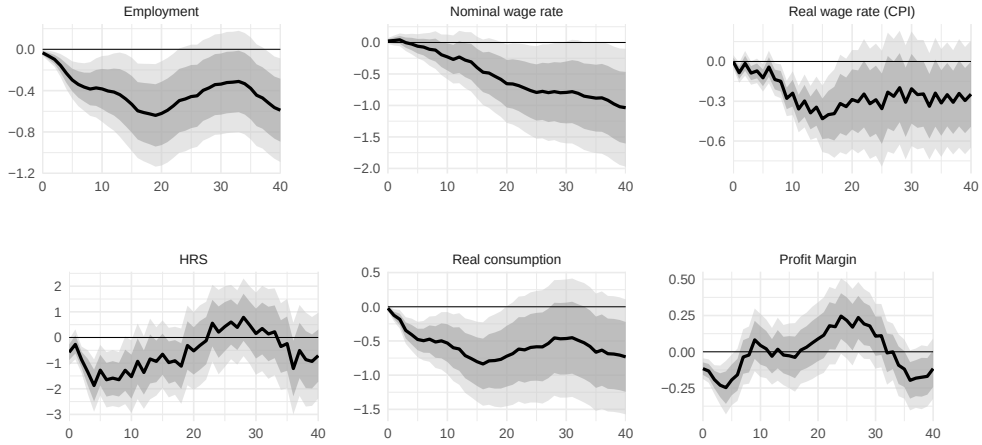

Note: Shaded areas denote 68% and 90% confidence intervals. Results are estimated according to specification (1). All responses are shown at a quarterly frequency.

A.5.4 Supply versus demand-driven hysteresis

To test whether the persistence is significantly different for supply and demand shocks, while controlling for their business-cycle effects, I estimate following regression:

$$\frac{y_{(i,t+h)} - y_{i,t-1}}{y_{(i,t+peak)} - y_{i,t-1}} = \alpha_{i,h} + \beta_S^h S_t^S + \beta_D^h S_t^D + \theta_h \Delta X_{i,t} + u_{i,t+h} \tag{4}$$

with $y_{(t+peak)}$ the value of real GDP at quarter $t + peak$, with $peak$ defined as the horizon at which output reaches its maximum business-cycle response following the shock in the baseline Local Projections specification (Specification (1)).

Intuitively, it captures the fraction of the initial output loss that persists over time. A larger value of the ratio indicates that a larger share of the initial decline remains, implying a stronger *degree of hysteresis* following the shock.

In other words, I test whether demand shocks generate a different degree of persistence in output than supply shocks. The results are reported in the second panel of Figure A.5.4. The results show that the estimated degrees of hysteresis for supply shocks are larger than for demand shocks, but they are not statistically significant. This suggests that, in this specification, supply shocks do not generate stronger hysteresis in output than demand shocks.

Figure A.5.4: Difference of IRFs (supply - demand) for Real GDP

(a) Difference **(b) Normalized Difference**

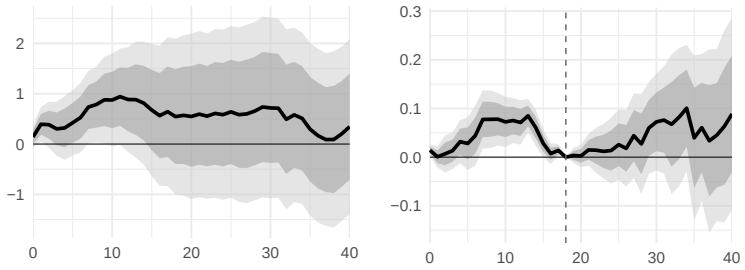

Note: Panel a plots the difference between the supply and demand shock coefficients at each horizon, with confidence bands constructed from the variance-covariance matrix as $\hat{V}_h^{S,S} + \hat{V}_h^{D,D} - 2\hat{V}_h^{S,D}$. Panel b plots the normalized difference, as in Specification (4).

Figure A.5.5: Difference of impulse responses to demand and supply shocks

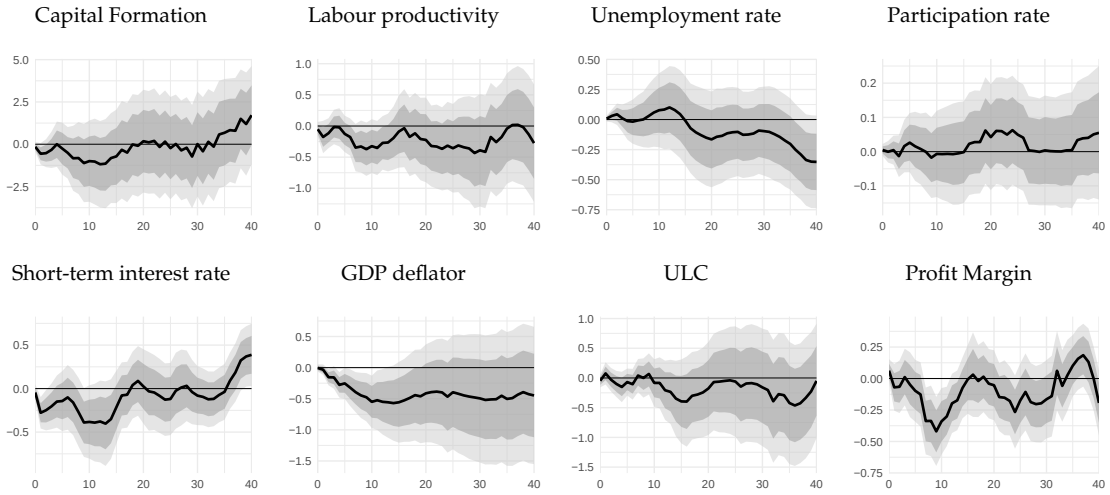

Note: Shaded areas denote 68% and 90% confidence intervals. Dependent variables are rescaled with the peak response of real GDP to the respective shock. Results are estimated over the period 1985Q1-2019Q4. All responses are shown at a quarterly frequency.

A.6 References

Baumeister, C. and Hamilton, J. D. (2019). Structural interpretation of vector autoregressions with incomplete identification: Revisiting the role of oil supply and demand shocks. *American Economic Review*, 109(5):1873–1910.

Gorodnichenko, Y. and Lee, B. (2020). Forecast error variance decompositions with local projections. *Journal of Business & Economic Statistics*, 38(4):921–933.

Känzig, D. R. (2021). The macroeconomic effects of oil supply news: Evidence from opec

announcements. *American Economic Review*, 111(4):1092–1125.