Estimating Hysteresis Effects *

Francesco Furlanetto† Antoine Lepetit‡ Ørjan Robstad§
Juan Rubio-Ramírez¶ Pål Ulvedal‖

August 8, 2023

Abstract

In this paper we identify demand shocks that can have a permanent effect on output through hysteresis effects. We call these shocks permanent demand shocks. They are found to be quantitatively important in the United States, in particular in samples starting in the 1980s. Recessions driven by permanent demand shocks lead to a permanent decline in employment and investment, while output per worker is largely unaffected. We find strong evidence that hysteresis transmits through a rise in long-term unemployment and a decline in labor force participation and disproportionately affects the least productive workers.

Keywords: Hysteresis, Structural vector autoregressions, Sign restrictions, Long-run restrictions, Employment, Labor productivity, Local projections.

JEL codes: C32, E24, E32

---

*First Draft: September 2020. This working paper should not be reported as representing the views of Norges Bank, the Board of Governors of the Federal Reserve System, or the Federal Reserve Bank of Atlanta. The views expressed are those of the authors and do not necessarily reflect those of Norges Bank or of the Federal Reserve System. We would like to thank for useful comments the editor Francesco Bianchi, three anonymous referees, editor Emi Nakamura, one anonymous referee for the Norges Bank working paper series, our discussant Daniele Siena, Knut Are Aastveit, Jonas Arias, Drago Bergholt, Saroj Bhattarai, Olivier Blanchard, Paolo Bonomolo, Jeff Campbell, Fabio Canova, Yoosoon Chang, Larry Christiano, Todd Clark, Olivier Coibion, Thorsten Drautzburg, Martin Eichenbaum, Stefano Eusepi, Luca Fosso, Luca Gambetti, Cooper Howes, Michel Juillard, Kurt Lunsford, Christian Matthes, Elmar Mertens, Silvia Miranda Agrippino, Nicoló Maffei Faccioli, Florens Odendhal, Francesco Ravazzolo, Øistein Røisland, Aysegul Sahin, Samad Sarferaz, Tatevik Sekhposyan, Sanjay Singh, Anders Vredin, Hans Christian Wika as well as seminar participants at Boston College, the Federal Reserve Bank of Cleveland, Norges Bank, University of Texas Austin, Texas A&M University, Indiana University, Sveriges Riksbank, Friendly Faces Online seminar series, Sant'Anna School of Advanced Studies in Pisa, NYU Abu Dhabi, Swiss National Bank, Penn State University, Banco de México, 20th CEF Conference in Ottawa in June 2019, 15th Dynare Conference in Lausanne, 6th Conference on New Developments in Business Cycle Analysis at Norges Bank, 2021 SNDE online conference, 2021 IAAE online conference, 2021 Ventotene Workshop in Macroeconomics, 2021 Conference in honor of Fabio Canova in Hydra, 2021 Macroeconomic Dynamics Workshop in Milano, 2022 T2M conference in London, 2022 EABCN conference on Challenges in Empirical Macroeconomics since 2020 organized by Deutsche Bundesbank, 2022 Barcelona Summer Forum and 2022 ESEM conference in Milano.

†Norges Bank and BI Norwegian Business School. Address: Bankplassen 2, PB 1179 Sentrum, 0107 Oslo, Norway. E-mail: francesco.furlanetto@norges-bank.no.

‡Board of Governors of the Federal Reserve System. E-mail: antoine.lepetit@frb.gov.

§Norges Bank. E-mail: orjan.robstad@norges-bank.no.

¶Emory University, FEDEA, Federal Reserve Bank of Atlanta, and BBVA Research. E-mail: juan.rubio-ramirez@emory.edu

‖Nord University and Norges Bank. E-mail: pal.b.ulvedal@nord.no.

1 Introduction

Macroeconomists are used to decomposing output per capita into an upward stochastic trend, often thought of as determining potential output or productive capacity, and the transitory fluctuations around it, often interpreted as business cycles. According to the traditional view, unexpected changes to the trend are caused only by supply shocks, such as labor supply and total factor productivity (TFP) shocks, while the business cycle is mostly driven by shocks to the components of aggregate demand and monetary policy. Blanchard (2018) argues that the assumption that productive capacity is independent from demand shocks in general, and monetary policy in particular, has become the dominant paradigm in macroeconomics and is the basis of the inflation-targeting framework used by most central banks. The "independence assumption" is embedded in the standard toolkit of modern macroeconomic analysis. In fact, most dynamic stochastic general equilibrium (DSGE) models imply that demand shocks have either no or a small transitory effect on the trend (cf. Blanchard, 2018) and structural vector autoregressions (SVAR) are often identified assuming only one shock with permanent effects on output (cf. Blanchard and Quah, 1989). This shock is commonly interpreted as a supply shock.

One alternative (and minority) view, popularized by Blanchard and Summers (1986) in the 1980s, states that demand shocks (especially when causing large recessions) may have a permanent effect on potential output through hysteresis effects. Economic developments in Europe in the 1980s seemed to support the hysteresis view since unemployment was stabilizing at a higher level following each recession. However, the Great Moderation was interpreted by many economists as supportive of the traditional view, and research on hysteresis largely disappeared. The idea that recessions may have permanent effects on output has re-emerged in the aftermath of the Great Recession as estimates of potential output have been revised down continuously over several years. As of today, the debate is not closed. Supporters of the traditional view argue that the downward revisions mainly reflect lower pre-existing trends masked by the boom in the pre-Great Recession period (cf. Gordon, 2015; Fernald et al., 2017; Antolin-Diaz et al., 2017; Eo and Morley, 2022). In contrast, Summers (2014) interprets them as evidence of hysteresis and stated that *"any reasonable reader of the data has to recognize that the financial crisis has confirmed the doctrine of hysteresis more strongly than anyone could have anticipated."*

In order to wind up the debate, we use U.S. data on output per capita (output) growth, inflation, employment-to-population (employment) growth, and investment per capita (investment) growth for the period 1983:Q1-2019:Q4 to identify a SVAR that allows for two shocks with potentially permanent effects on output: a traditional supply shock and a more novel demand shock that we disentangle on the basis of the short-run co-movement between output growth and inflation as advocated by Summers (2015). We also allow for two transitory shocks; a demand and a supply shock with no permanent effects on either output or employment. In practice, we combine long-run zero and short-run sign restrictions to identify the four shocks using the methodology proposed by Arias et al. (2018). We focus our attention on the demand shock with potentially permanent effects, evaluate its importance for economic fluctuations, and analyze its transmission mechanism. The more important this shock is, the larger are the deviations from the independence assumption and the larger is the role for hysteresis effects.

Our main result is on *the relevance of hysteresis effects*. First, we find that demand shocks with potentially permanent effects do indeed lead to a permanent decrease in output. Thus, from then on, we simply call them permanent demand shocks. Second, we find that these shocks explain more than 50 percent of the fluctuations in long-run output growth. Such a relevant role for permanent demand shocks highlights that the traditional view is not supported by the data. Permanent demand shocks also have important negative permanent effects on employment and investment.

Our second result is related to *the transmission mechanism of hysteresis effects*. A permanent decline in output can be conveniently decomposed into an effect on employment and an effect on output per worker. Our impulse response functions (IRFs) show that hysteresis propagates almost exclusively through employment. Output per worker, which can be interpreted as a simple measure of labor productivity, is hardly affected at all, both in the short run and in the long run. Using local projection (LP) methods, we show that the permanent decline in employment is accompanied by an increase in long-term unemployment, a decline in participation, and an increase in applications (and awards) for disability insurance. These responses are consistent with standard hysteresis channels and compatible with the skill depreciation and reduced employability of long-term unemployed workers. When it comes to the neutral effect on labor productivity, the LP approach shows that it is likely the outcome of compensating effects. The share of employment in routine (see Jaimovich and Siu, 2020; Fernández-Villaverde et al., 2019), and arguably less productive, tasks decreases in response to a negative permanent demand shock. Along with an increase in capital intensity, this pushes up labor productivity, but is compensated by a protracted decrease in capital and labor utilization. Our results indicate that these forces roughly cancel each other, leaving output per worker largely unaffected after a permanent demand shock.[1] Notably, while invariant to permanent demand shocks, output per worker responds strongly to a permanent supply shock. Therefore, our results are consistent with supply shocks being the only drivers of labor productivity in the long run, as assumed in Galí (1999).

We contribute to the empirical literature on hysteresis (cf. Cerra et al., 2023, for a recent survey). Most studies restrict their attention to deep recessions and investigate

[1]Several recent papers use New Keynesian models with endogenous growth to examine the hypothesis that the slowdown in productivity following the Great Recession was to a large extent an endogenous response to the collapse in demand that caused the contraction in economic activity. See Benigno and Fornaro (2018), Guerron-Quintana and Jinnai (2019), Ikeda and Kurozumi (2019), Moran and Queralto (2018), Bianchi et al. (2019), Anzoategui et al. (2019), Garga and Singh (2021), and Schmöller and Spitzer (2021). Similar mechanisms are also present in agent-based models, as discussed in Dosi et al. (2018).

their impact on the economy's productive capacity. Cerra and Saxena (2008) find evidence of highly persistent effects on output. Since recessions are not necessarily all driven by demand shocks, Blanchard et al. (2015) focus on 22 recessions associated with intentional disinflations, mostly concentrated during the 1980s and early 1990s. These recessions are driven by large monetary policy shocks that reflect a surprise change in policy rather than the policy response to other shocks. They find that nearly two-thirds of these recessions are associated with a permanently lower output level and that a significant fraction of those are associated with permanently lower output growth. Our paper also connects with SVAR studies on labor market dynamics in the U.S. (cf. Galí and Hammour, 1992), in Scandinavian countries (cf. Jacobson et al., 1997), in Italy (Gambetti and Pistoresi, 2004), and in Spain (Dolado and Jimeno, 1997). None of these papers combine zero and sign restrictions to identify shocks. Notably, sign restrictions were introduced well after (cf. Faust, 1998; Canova and De Nicoló, 2002; Uhlig, 2005; Rubio-Ramirez et al., 2010) and the combination of sign and zero restrictions has become feasible only with the routines recently developed by Arias et al. (2018). In a recent paper, Benati and Lubik (2022) estimate a cointegrated SVAR for the U.S. over a long sample starting in 1954 and find some weak evidence of hysteresis. Finally, while we consider demand shocks with long-run effects on output, Maffei-Faccioli (2021) studies the impact of demand factors on output growth (an effect named by Ball, 2014, as super-hysteresis) in a SVAR with common trends and finds supportive evidence.

The paper proceeds as follows. Section 2 provides a brief description of our empirical set-up. Section 3 presents our main results. Section 4 discusses the channels of transmission of hysteresis. Section 5 examines the sensitivity of the results to the sample under consideration. Section 6 evaluates the robustness of the results in larger systems where additional shocks are identified. Section 7 proposes a Monte Carlo exercise to check whether hysteresis can be detected in a controlled model environment. Finally, Section 8 concludes.

2 The Model

The traditional approach to disentangle shocks with long-run effects on output from purely transitory disturbances relies on Blanchard and Quah (1989). Their decomposition uses data on GDP growth and unemployment and imposes that transitory shocks have a zero long-run effect on output. Despite not using data on inflation, the decomposition has been traditionally used to separate demand shocks (assumed to be transitory) from supply shocks (assumed to have permanent effects). In this paper, we argue that this traditional practice is not reliable and that both demand and supply shocks are likely to generate long-run effects. To illustrate our point, we estimate a bivariate VAR with 3 lags using data on inflation and output growth over two adjacent sample periods. Notably, we maintain the assumption that one shock has transitory effects on GDP while the second shock is allowed to have permanent effects. When estimated over the sample 1949:Q1-1982:Q4 (which largely overlaps with the original Blanchard-Quah sample), the model supports the traditional interpretation: the permanent shock behaves as a supply shock (thus generating a negative co-movement between output and prices) while the transitory shock transmits as a demand shock (generating a positive co-movement), as shown in Figure 1. However, when estimated over the sample 1983:Q1-2019:Q4, the very same model implies that the permanent shock propagates as a demand shock. This result suggests that the permanent component in the Blanchard-Quah decomposition is likely to commingle demand and supply shocks, both with permanent effects. Investigating the relevance and the propagation of these demand shocks with permanent effects is the goal of this paper. In the next step, we extend the original Blanchard-Quah methodology in order to isolate more than one transitory shock and more than one permanent shock.[2] We consider the standard reduced-form VAR model:

$$\mathbf{y}_t = \mathbf{C}_B + \sum_{i=1}^{P} \mathbf{B}_i \mathbf{y}_{t-i} + \mathbf{u}_t,$$

where $\mathbf{y}_t$ is an $N \times 1$ vector containing our N endogenous variables, $\mathbf{C}_B$ is an $N \times 1$ vector of constants, $\mathbf{B}_i$ for $i = 1, ..., P$ are $N \times N$ parameter matrices, with P the number of lags (3 in our specific case), and $\mathbf{u}_t$ the vector of innovations with $\mathbf{u}_t \sim N(0, \mathbf{\Sigma})$, where $\mathbf{\Sigma}$ is the $N \times N$ variance-covariance matrix. We rely on long-run zero and short-run sign restrictions to identify the shocks. We implement the restrictions using the conventional algorithm proposed by Arias et al. (2018).[3]

We combine the standard Minnesota prior with the sum-of-coefficients prior (Doan et al., 1984). This combination of priors will reduce overfitting for both the parameters that drive the temporary fluctuations and the underlying deterministic trend in the reduced-form VAR (initial conditions). Restraining the role of the initial conditions could be of great importance when trying to assess the long-run effects of demand and supply shocks. We follow Giannone et al. (2015) in the choice of hyperparameters. We use averages from 1949 to the beginning of our estimation sample as the dummy initial observation in the sum-of-coefficients prior, while the Minnesota prior is centered around the variables being independently and identically distributed. Because all our variables are considered in first differences, the prior has mean zero.[4]

[2]Figure A-1 in the Appendix replicates the original analysis by Blanchard and Quah (1989) in our framework and recovers their results when using their data.

[3]Arias et al. (2023) show that several of the criticisms in Baumeister and Hamilton (2015); Watson (2020) about the conventional algorithms have no basis. In particular, the paper formally shows that when the focus is on joint inference the uniform prior over the set of orthogonal matrices is not only sufficient but also necessary for inference based on a uniform joint prior distribution over the identified set for the vector of impulse responses.

[4]Because our data "should be" stationary, the sum-of-coefficients prior may have undesirable effects. In particular, this prior could have a bearing on the estimated persistence. Since the employment-to-population ratio is quite persistent, we find that the sum-of-coefficients priors are "active" in the sense of the data requiring relatively tight priors. We thank one referee for raising this point.

Figure 1: IRFs from Blanchard-Quah decomposition for different samples

(a) Estimated over the sample 1949:Q1 - 1982:Q4

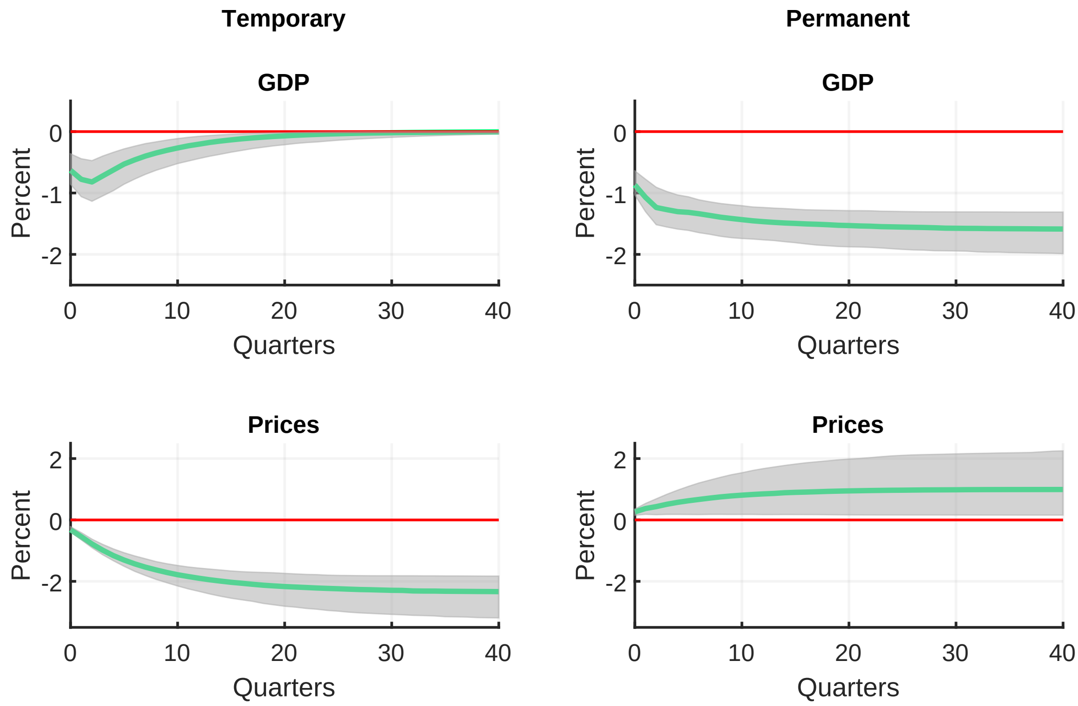

(b) Estimated over the sample 1983:Q1 - 2019:Q4

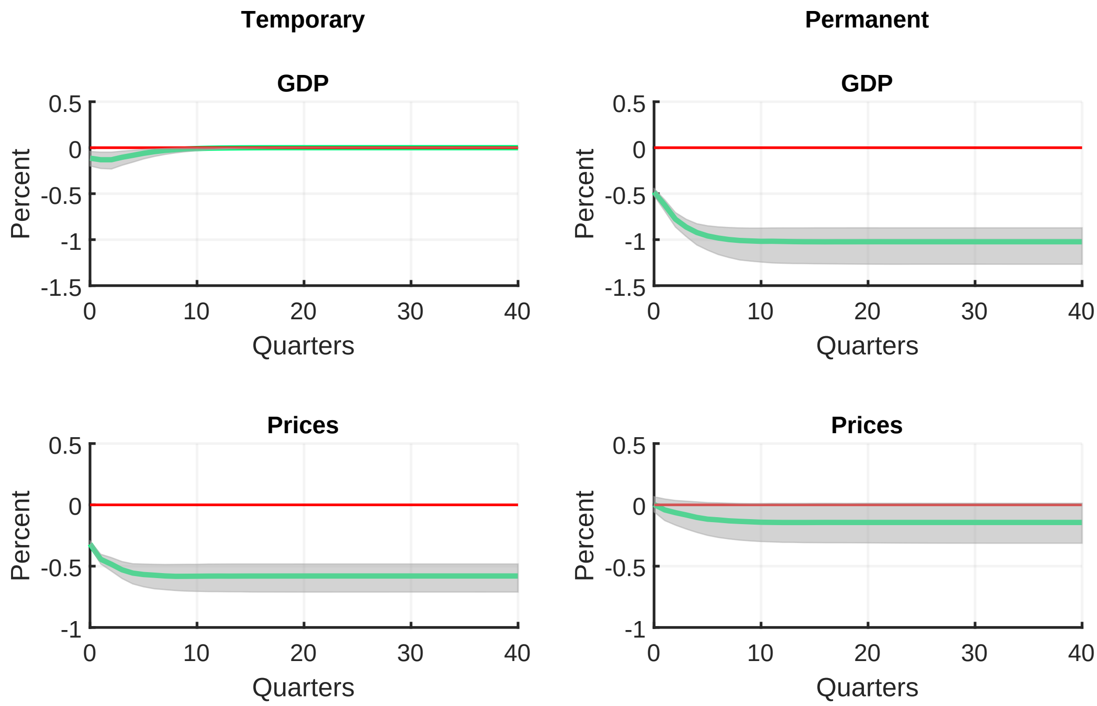

We use quarterly U.S. data on real GDP per capita (output), PCE deflator (prices), employment-to-population ratio (employment) and real investment per capita (investment) over the sample period 1983:Q1-2019:Q4. All variables enter our model in first differences. Since SVAR models identified with long-run restrictions are sensitive to trend breaks and low-frequency correlations (cf. Fernald, 2007), we have chosen to focus on a relatively homogeneous sample. We follow Galí (1999) and use data in first differences in order to allow for (without imposing) hysteresis effects on employment. A specification in levels would tilt the IRFs to converge back to zero, thus making hysteresis effects immaterial (at least in the long run). Nonetheless, we consider a specification with employment (instead of employment growth) in Section 5; our results survive. We consider three lags, in accordance with the Akaike information criterion (AIC). Two shocks are transitory, while the remaining two are allowed to have permanent effects. We note that data on (detrended) unemployment, rather than on the employment-to-population ratio, are used in Blanchard and Quah (1989). Our choice is based on the experience of the recovery from the Great Recession: while unemployment eventually recovered (albeit slowly) to its pre-Great Recession level, the employment-to-population ratio did not. Therefore, it seems more fruitful to search for hysteresis effects by looking at employment data (cf. Yagan, 2019).[5]

The identification assumptions are summarized in Table 1. We assume that there are two transitory shocks that have a zero long-run impact on output (as in Blanchard and Quah, 1989) and employment, thus implying that labor productivity is not affected in the long run by these two transitory shocks. While it is useful in sharpening identification, the long-run restriction on employment does not drive our results. We disentangle the two transitory shocks on the basis of the short-run contemporaneous co-movement between output growth and inflation: a transitory demand shock moves the two variables in the same direction, while a transitory supply shock moves them in the opposite direction. The long-run impact of the remaining two shocks is left unrestricted but the same contemporaneous sign restriction on the co-movement between output growth and inflation is used to classify them as demand or supply shocks. Thus, we identify a traditional supply shock with potentially long-run effects together with a more novel demand shock with also potentially long-run effects. It should be clear that we do not impose that these two shocks have any long-run effects; we only allow for such a possibility. As we will see below, they both have, in fact, permanent effects. For this reason, we will call them permanent demand and permanent supply shocks in the rest of the paper. All sign restrictions are imposed over the first two quarters. Finally, investment is left unrestricted, which allows us

[5]We discuss the results of a specification using data on unemployment in Section 6. In addition, all our results are confirmed when we use total hours worked per capita rather than the employment-to-population ratio in the SVAR, as reported in Figure A-2 in the Appendix. We believe that the discussion of hysteresis is most natural in terms of employment. However, it is important to cross-check our results since the literature using long-run restrictions to identify the effects of permanent shocks on the labor market has traditionally used total hours worked (cf. Galí, 1999 and Fernald, 2007 among others).

Table 1: Identifying restrictions

	Demand - perm.		Supply -perm.		Demand - temp.		Supply - temp.	
	Short-run	Long-run	Short-run	Long-run	Short-run	Long-run	Short-run	Long-run
Output	-		-		-	0	-	0
Prices	-		+		-		+	
Employment						0		0
Investment								

Note: Short-run restrictions are imposed on the first difference of the variable, while long-run restrictions are imposed on the level of the variable.

to investigate the strength of the investment channel emphasized by Benigno and Fornaro (2018) and in much of the endogenous growth literature.

A word of caution on the identification strategy needs to be added here. It should not be taken for granted that a negative permanent supply shock should lead to an increase in inflation on impact. It is possible to find parameterizations of the standard New Keynesian model under which a negative permanent technology shock leads to a decrease in inflation (cf. Galí et al., 2003, among others). However, the overwhelming majority of estimated New Keynesian models and SVAR models find a positive response of inflation to a contractionary permanent technology shock (cf. Christiano et al., 2003; Galí and Rabanal, 2004; Basu et al., 2006; Paciello, 2011; Altig et al., 2011, among others). More generally, since it is possible to construct knife-edge cases in which our restrictions are not satisfied on impact, we estimate a specification in which restrictions are imposed at horizon four; our results, which we present in Figure A-3 in the Appendix, are even stronger in such a case.

In the literature, hysteresis effects are often associated with recessions and not with booms. We note, however, that Ball et al. (1999) and, more recently, Aaronson et al. (2019) and Bluedorn and Leigh (2019) provide evidence of positive hysteresis where protracted expansions are associated with permanent increases in employment. In light of these results, our model's linear structure, although admittedly simple, seems to be a reasonable starting point to search for hysteresis effects.

3 Finding Hysteresis

In this Section, we present our results. In Figure 2, we plot cumulative IRFs to both the permanent demand and supply shocks. All IRFs plotted in the paper are in response to one standard deviation negative (contractionary) shocks. The solid line represents the posterior median at each horizon and the shaded area indicates the 16th and 84th percentiles obtained from the set of IRFs consistent with our identification assumptions. In Figure 3, we present the cumulative forecast error variance decomposition (FEVD) based on the point-wise median posterior estimate. Because we report cumulative IRFs and FEVDs, the figures show the (log) level response of the variables, not the growth rates.

Two main results stand out. First, permanent demand shocks are important drivers of output fluctuations: they account for more than 50 percent of output variations at all horizons. While permanent supply shocks are also important drivers of output, they play a more minor, yet not negligible, role. Second, we find that the output response to a permanent demand shock is surprisingly similar to the response to a permanent supply shock. Importantly, these similar dynamics are not the result of weakly identified shocks. The response of prices to the two shocks is substantially different despite being restricted only over the first two quarters. In addition, the decomposition of output between employment and output per worker reveals clear differences. The permanent demand shock propagates exclusively through employment and both IRFs and FEVD indicate that employment fluctuations in the long run are largely driven by permanent demand shocks. In contrast, the permanent supply shock propagates mainly through output per worker; both IRFs and FEVD indicate that output per worker is explained almost exclusively by permanent supply shocks in the long run. This result supports the identification scheme proposed by Galí (1999) to identify technology shocks in a SVAR where only one shock can have permanent effects on labor productivity.[6]

We also find that investment drops permanently in response to both permanent demand and supply shocks. While such a behavior of investment in response to permanent supply shocks is not surprising, standard theories of the business cycle predict that the drop in investment in response to demand shocks should only be temporary. However, the observed response of investment to permanent demand shocks is consistent with channels emphasized by Benigno and Fornaro (2018), where the interaction of endogenous growth and the zero lower bound on nominal interest rates can lead to prolonged periods of high unemployment and low growth, or Kozlowski et al. (2020), where the occurrence of extreme events such as the Great Recession generates persistent changes in beliefs and macroeconomic outcomes.

Summarizing, our results show that hysteresis effects are large. In particular, a negative permanent demand shock leads to a permanent decline in output and employment and this shock is an important driver of output fluctuations and the main driver of employment fluctuations, both at short and long horizons. The fact that hysteresis effects propagate mainly through employment while leaving labor productivity unchanged is a defining feature of our results.

We now briefly comment on transitory shocks whose cumulative IRFs are presented

[6]Unlike in Blanchard and Quah (1989) and Galí (1999), whose samples cover the 1950s to the 1970s, our permanent permanent supply shock generates co-movement between output and employment. Galí et al. (2003) also find that the response of hours to a permanent supply shock switches sign in the Volcker-Greenspan period and argue that it is related to a change in the conduct of monetary policy. Moreover, our permanent supply shock can capture forces other than technology such as labor supply or tax shocks, as in Mertens and Ravn (2011) and Cloyne et al. (2022).

Figure 2: IRFs to the permanent demand and supply shocks

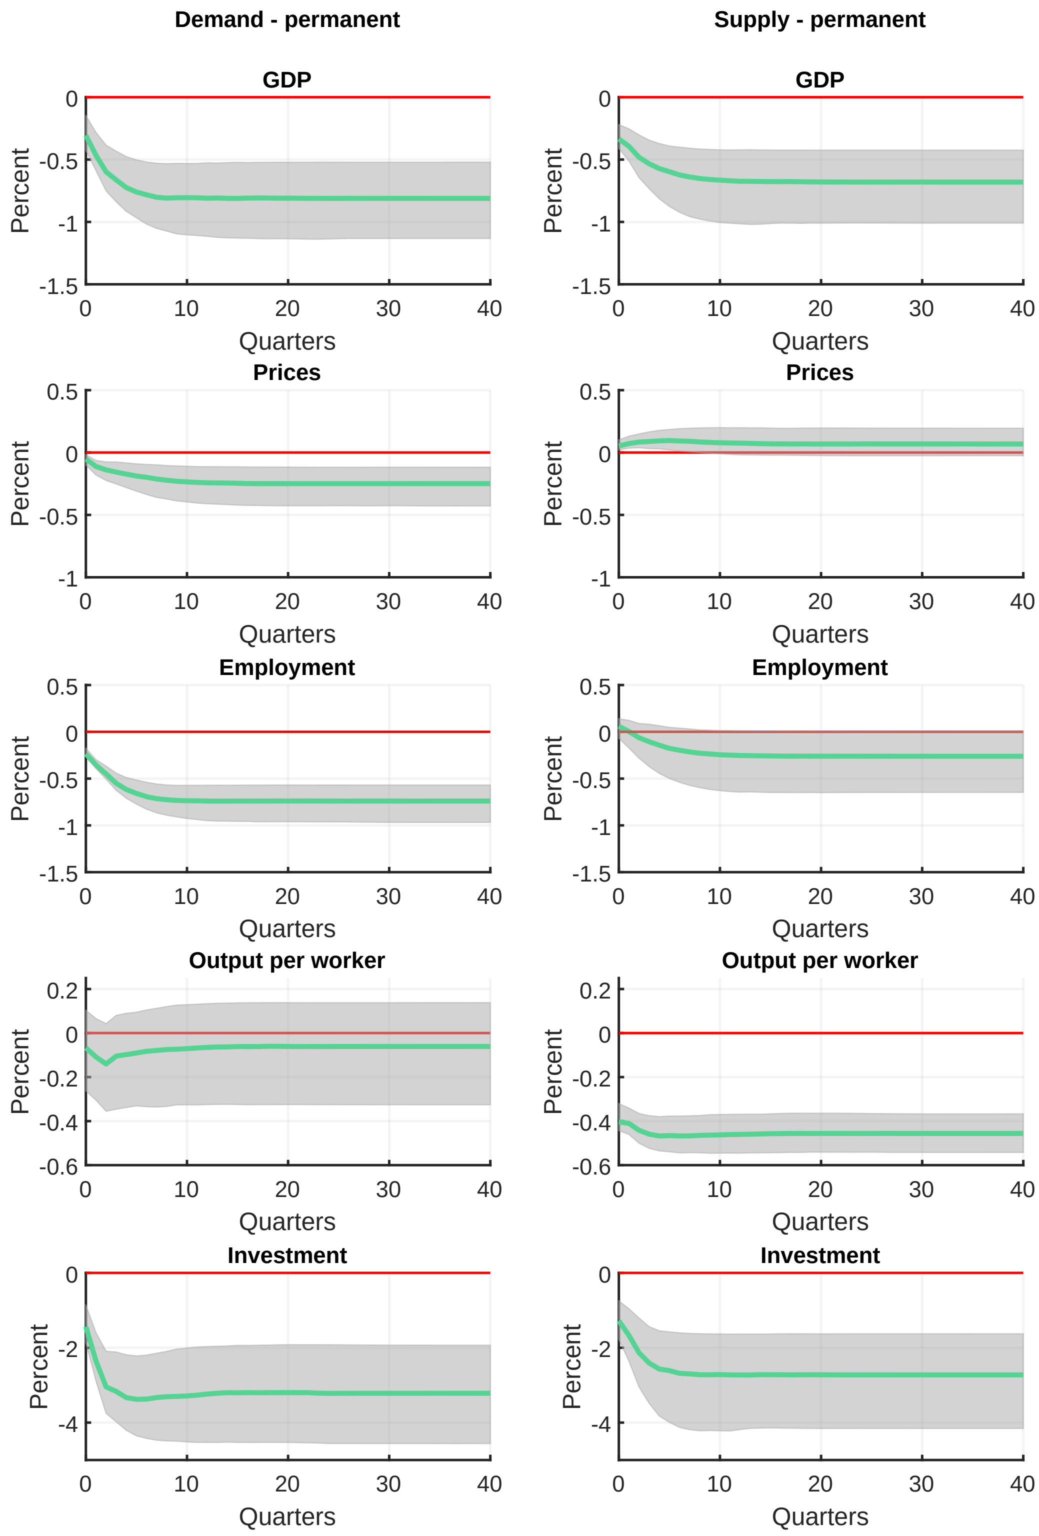

Figure 3: Forecast error variance decomposition

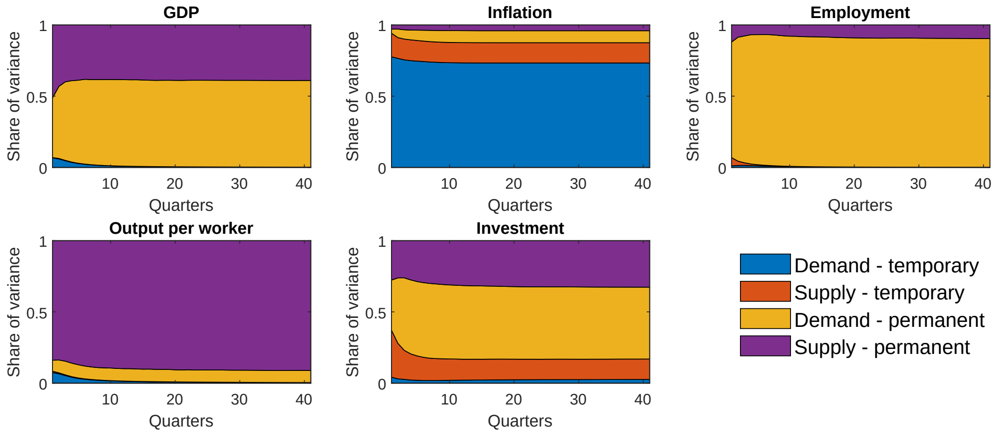

Figure 4: Historical decomposition of the growth rate in GDP per capita

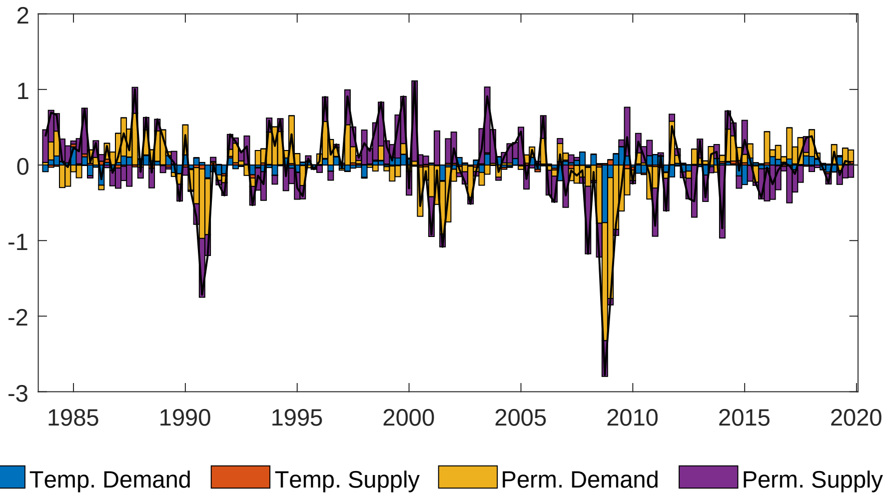

in Figure A-4 in the Appendix. The transitory supply shock marginally explains price and investment fluctuations. Transitory demand shocks induce a small fraction of output fluctuations but are the main drivers of fluctuations in prices at all horizons. This finding is consistent with evidence on the orthogonal response of prices to the business cycle since the 1990s (see, for example, Del Negro et al., 2020, among many others): inflation is driven by its own shock (our transitory demand shock) and this shock has little effect on real variables.[7] The limited role of transitory shocks is discussed further in Section 6.

In Figure 4 we present a historical decomposition for output growth (in deviation from its forecastable component) into the contribution of the four shocks. The conclusion is that permanent demand shocks are dominant in recessions. To reinforce that point, Figure 5 runs counterfactuals for output over the three recessions in our sample. The left column shows how output would have behaved in the absence of permanent demand shocks in

[7] Angeletos et al. (2020) similarly find that shocks explaining the bulk of fluctuations in real activity explain very little of movements in inflation, and vice versa.

Figure 5: Counterfactual path of output in recessions in the absence of permanent demand and supply shocks

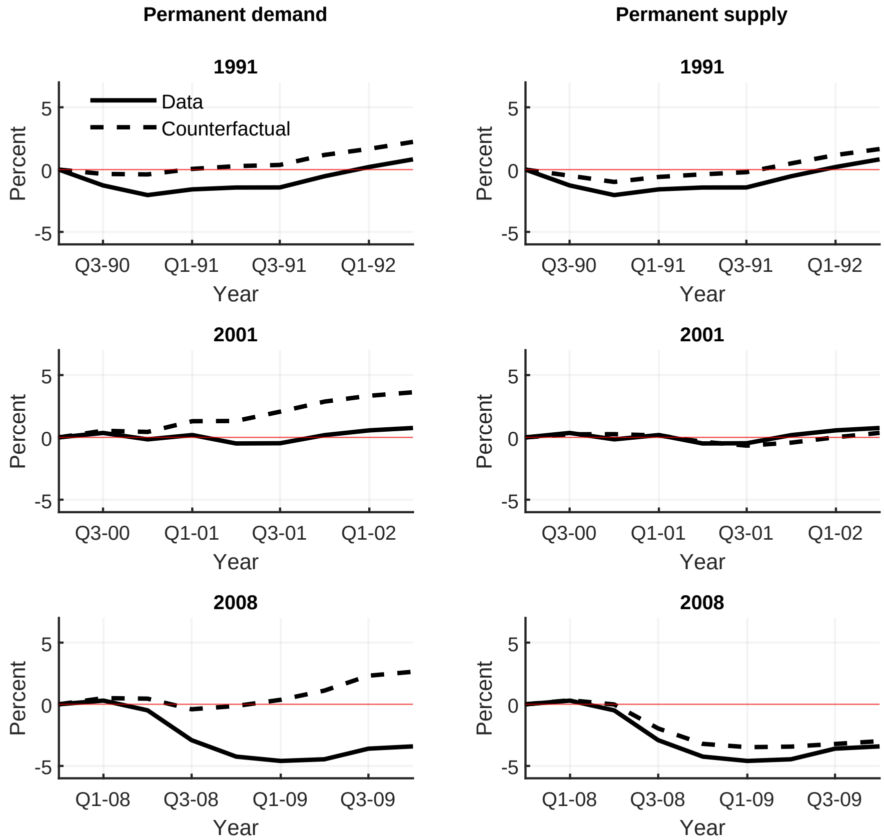

the recessions of 1991, 2001, and 2008. The right column performs the same exercise in the absence of permanent supply shocks. Clearly, permanent demand shocks are the most important driver of these three recessions.

It is also worthwhile to note that while permanent supply shocks are not important drivers of recessions in our sample, they contribute significantly to the sustained output growth of the late 1990s, in line with the high productivity growth in that period. Consistently with the results described above, transitory demand shocks explain a very small fraction of unexpected fluctuations in output.

Finally, we comment on the uncertainty surrounding our estimates. Being set-identified models, sign-restricted SVARs feature substantial uncertainty. This is partly the case also in our model and it is the price to pay for the simplicity of the identification scheme which relies on a minimum set of identification assumptions. Yet, uncertainty can be substantially reduced if sign restrictions are imposed over a longer horizon, if additional long-run restrictions are enforced or if narrative restrictions (see Antolín-Díaz and Rubio-Ramírez, 2018) are added on the top of the sign restrictions.

4 Understanding Hysteresis

We have documented that permanent demand shocks propagate almost entirely through employment, while they have a small effect on output per worker. We now investigate the more granular transmission channels of hysteresis by comparing the responses of other macroeconomic variables to permanent demand shocks. First, we aim to understand why hysteresis transmits through employment. Second, we will try to explain why permanent demand shocks leave output per worker almost unaffected in the long run. We use local projection (LP) methods to analyze the transmission mechanism. The LP approach was proposed by Jordà (2005) and further developed by Ramey and Zubairy (2018), Stock and Watson (2018), and Plagborg-Møller and Wolf (2021), among others, and consists in direct regressions of the variable of interest in period $t+h$ on a measure of an identified shock at time t, as well as on control variables. One could possibly argue that the additional variable of interest could be added to the SVAR. This approach would increment the number of parameters to be estimated and the uncertainty surrounding the IRFs. As argued by Montiel Olea and Plagborg-Møller (2021), once the shocks of interest are obtained, LP inference is arguably both simpler and more robust than standard autoregressive inference, whose validity is known to be sensitive to the persistence of the data and to the length of the horizon. Accordingly we consider the following regression for each horizon h:

$$y_{t+h} - y_{t-1} = \alpha_h + \sum_{s=1}^{3} \lambda_{h,s} \Delta y_{t-s} + \beta_h S_t + \varepsilon_{t+h} \tag{1}$$

where y_t is the macroeconomic variable of interest and S_t represents the time series for the shock. Because LP is less efficient than SVARs, we plot IRFs only up to horizon 20. In some cases, we conduct the same regression to trace the effects of permanent supply shocks for comparison. We follow a Bayesian approach whereby for each draw of the shock's distribution, we compute IRFs from Equation (1) using a noninformative normal-Wishart prior on the coefficients. The underlying idea is to calculate the IRFs of a large set of variables conditional on the distribution of shocks.[8]

4.1 Effects on Employment

Section 3 finds that hysteresis transmits mostly through employment. In order to understand the channels of propagation of hysteresis, Figure 6 shows the IRFs of five variables related to employment (the unemployment rate, the fraction of long-term unemployed, the participation rate, as well as applications and awards for disability insurance) to both permanent demand and supply shocks. These five variables are listed by Blanchard (2018) as the leading candidates to be considered when searching for the transmission channels

[8] Because we include lagged values of the variable of interest in the regression, we do not need to do autocorrelation adjustments to the posterior, which simplifies inference (Montiel Olea and Plagborg-Møller, 2021).

of hysteresis (see also Coibion et al., 2013). The left column of Figure 6 presents the responses to permanent demand shocks, while the right column plots the responses to permanent supply shocks.

The first obvious candidate is the unemployment rate. The IRF of unemployment to both permanent shocks is shown in the first row of the figure. Not surprisingly, the unemployment rate increases significantly in response to the demand shock, although we also observe a gradual decline after 2-3 years. The dynamics are similar for the permanent supply shock, but more muted. These responses indicate that the unemployment rate cannot be the main propagation channel of hysteresis. As Blanchard (2018) noted, if some workers become less employable or discouraged after a permanent demand shock, then the unemployment numbers will fail to fully recover the transmission channel of hysteresis. In the second row of Figure 6, we consider the ratio of long-term unemployment (unemployed for 27 weeks or more) to total unemployment: its response to permanent demand shocks is stronger and more persistent than its response to permanent supply shocks. Differences are even more striking when considering the participation rate: we find large and long lasting negative effects in response to permanent demand shocks and no effects (on average) in response to permanent supply shocks. This indicates that long-term unemployment and participation rates are important channels of the propagation of hysteresis. Finally, to reinforce the point that workers became less employable and discouraged after a permanent demand shock, in rows four and five, we consider applications and acceptances for disability insurance. As Blanchard (2018) puts it: *"Cyclical variations in applications for disability insurance can give information about the loss of morale among workers as a result of the state of the labor market. And once people are accepted and start receiving disability payments, terminations are rare. This implies that, to the extent that recessions lead to increases in disability insurance rolls, they have a hysteretic effect on the labor force."* Our results are consistent with disability insurance rolls having a hysteretic effect on the labor force. The figure shows that while applications and awards respond strongly to demand shocks, they hardly respond at all to supply shocks.

Finally, we would like to stress that the results in Figure 6 also show that the channels of propagation of hysteresis in recent years are substantially different from the ones discussed in the literature from the 1980s. In particular, Blanchard and Summers (1986) highlighted the role of insiders in wage formation and the potential impact this would have on the unemployment rate. Instead, our results are consistent with studies emphasizing the negative impact of recessions on morale and skills, leading to a rise in long-term unemployment and a decline in the participation rate.

Given that hysteresis transmits through employment, it is of interest to understand how it affects the employment levels of different gender and race groups. To do that we consider more disaggregated data based on gender and race, building on Aaronson et al. (2019). Figure 7 summarizes responses to permanent demand shocks. In the first row, we

Figure 6: IRFs to permanent demand and supply shocks on labor market variables

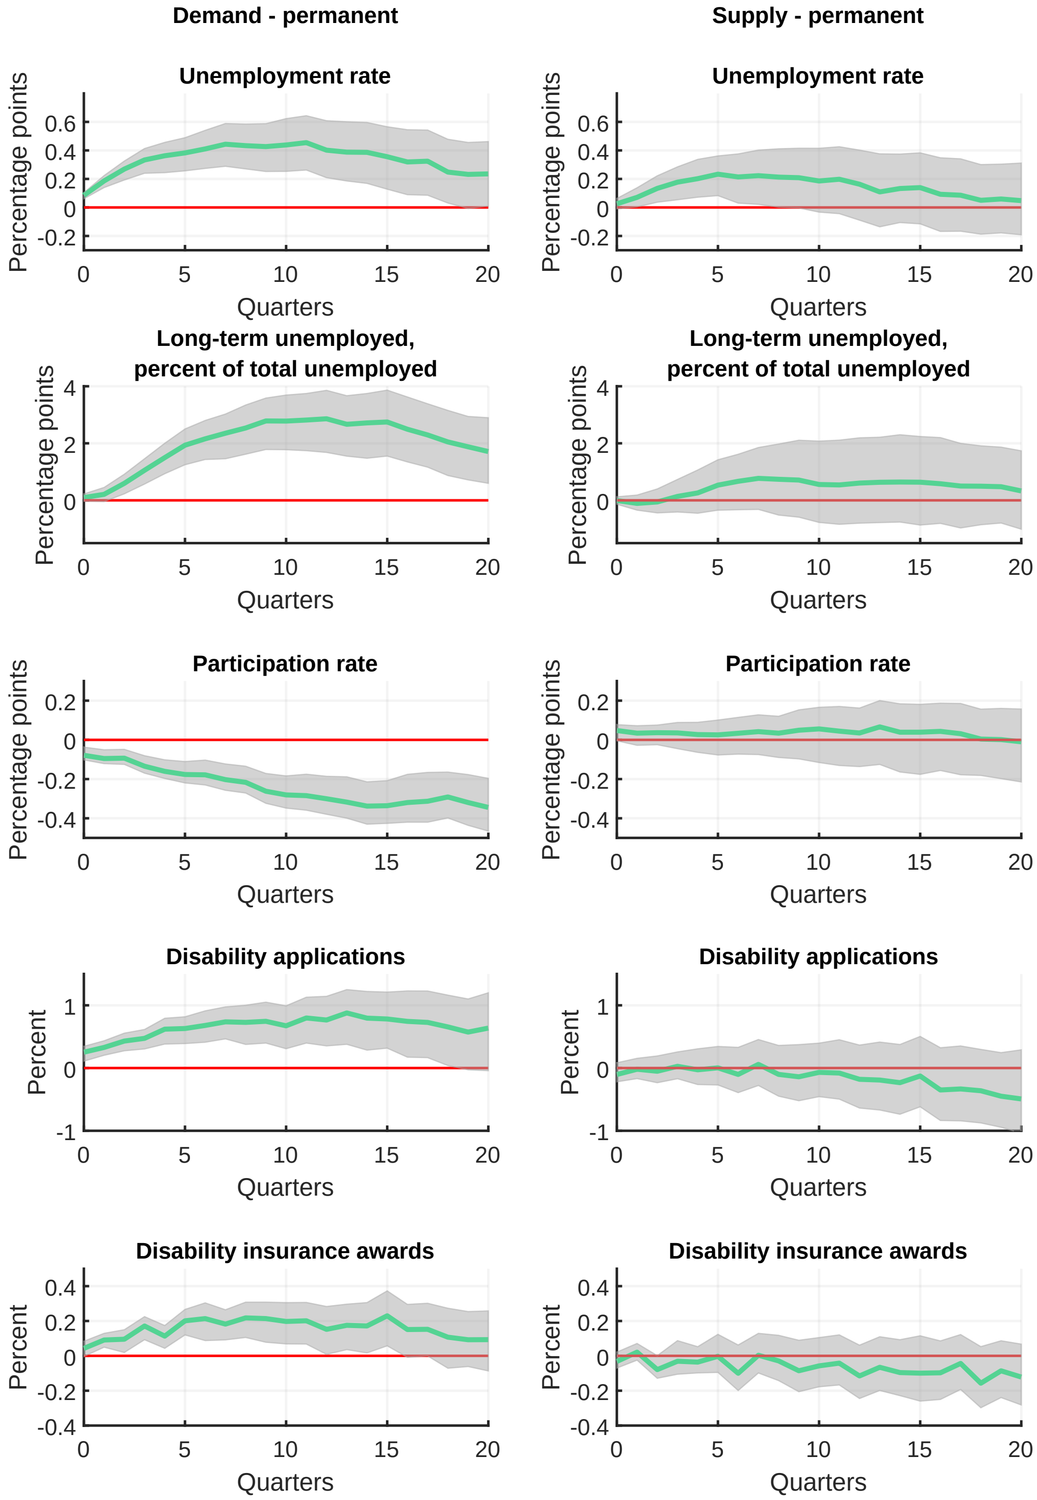

Figure 7: IRFs to permanent demand shocks on relative employment by gender and race

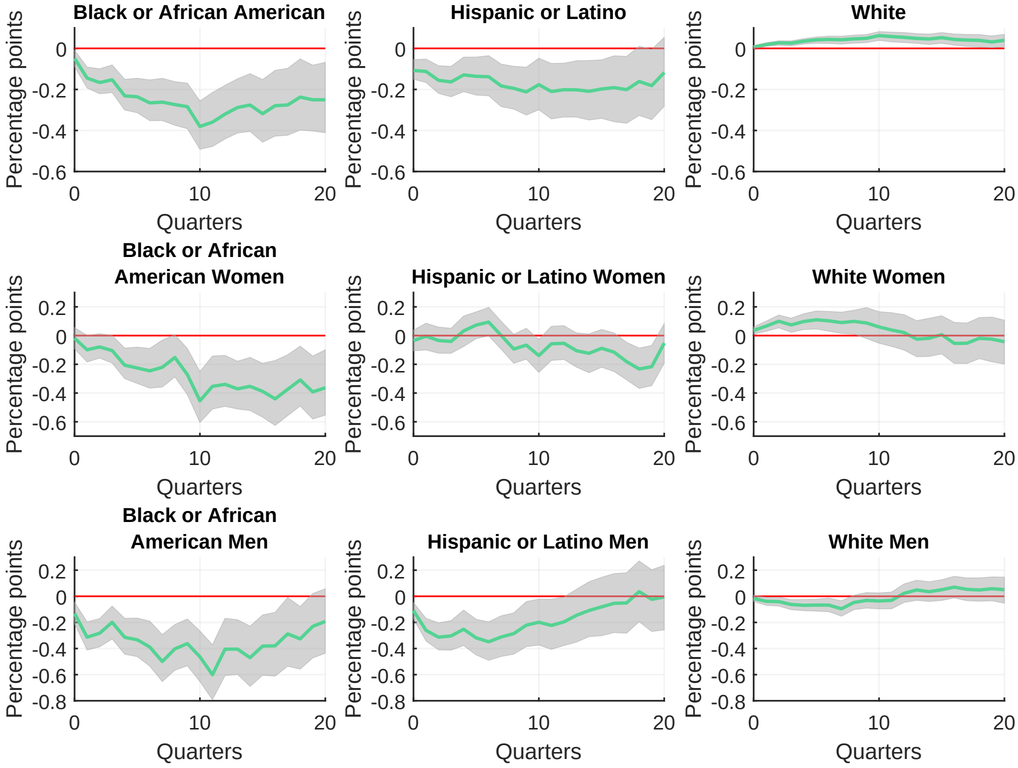

plot the IRFs of relative employment for black or African-American workers, Hispanic and Latino workers, and white American workers. Relative employment is calculated as the deviation of the employment-to-population ratio of a particular group from the aggregate employment-to-population ratio. Clearly, African-American and Hispanic workers are disproportionately affected by permanent demand shocks, while white Americans suffer a lower decline in employment than average. This confirms the results in Aaronson et al. (2019).

In the second and third rows, we consider the responses disaggregated by gender. As before, relative employment for each category is calculated in deviation from the employment-to-population ratio for the whole economy. We note that the employment rate for men falls more in the short run for all races. However, toward the end of the estimation horizon (20 quarters), the effects are roughly similar for men and women, conditional on race.

Figure 8: IRFs of labor productivity components to permanent demand shocks

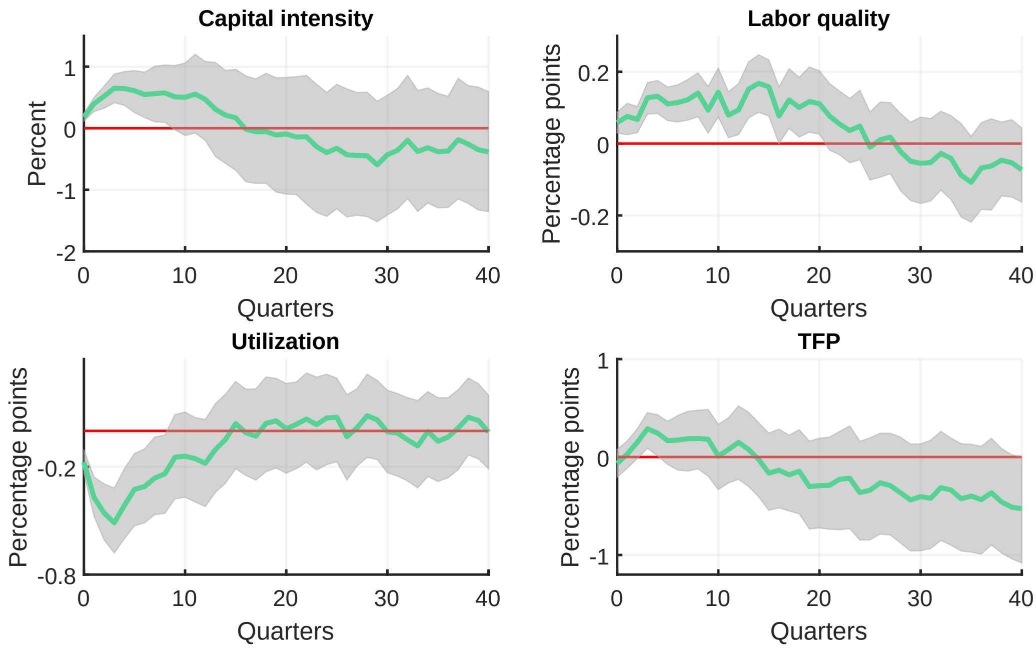

4.2 Effects on Labor Productivity

Our results show that strong hysteresis effects on employment are accompanied by limited ones on labor productivity.[9] At the same time, labor productivity is strongly affected by permanent supply shocks. In order to understand why labor productivity is not a channel of transmission of hysteresis, we decompose it following the methodology outlined in Fernald (2014). Labor productivity can be written as:

$$\Delta lnY_t - \Delta lnH_t = \alpha \left(\Delta lnK_t - \Delta lnH_t\right) + (1-\alpha)\, \Delta lnQ_t + \Delta lnU_t + \Delta lnA_t, \qquad (2)$$

where Y_t is output, H_t is total hours worked, K_t is capital, Q_t is labor quality, U_t is labor and capital utilization, A_t is utilization-adjusted total factor productivity, and α is the share of capital in total output. We call $(\Delta lnK_t - \Delta lnH_t)$ capital intensity. We use Fernald's up-to-date quarterly estimates of each of these variables to decompose movements in labor productivity into their underlying drivers.

While our measure of labor productivity is output per worker for the whole economy, Fernald's measure of labor productivity is output per hour worked in the U.S. business sector. Thus, our first step is to check whether the results reported in Section 3 survive with Fernald's measure of labor productivity. Figures A-5 and A-6 in the Appendix present results for the SVAR re-estimated using Fernald's measures of output and hours. Our previous results are confirmed and, as before, we find that labor productivity hardly moves in response to a permanent demand shock.

Next, we use the above decomposition to understand the drivers behind the responses

[9]Relatedly, Bhattarai et al. (2021) find that the 2006-2009 housing crisis had long-run scarring effects on employment but not on labor productivity.

of labor productivity to permanent demand shocks. The panels in Figure 8 plot the IRFs of the different terms on the right-hand side of Equation (2). The IRFs are obtained using the LP approach described above. In this case we show IRFs up to horizon 40 because we want to analyze the long-run effects of permanent demand shocks on TFP.

The Figure shows that the muted short-run response of labor productivity to a permanent demand shock is the result of two offsetting effects. On the one hand, capital intensity goes up as employment declines abruptly, while capital reacts only slowly to the decrease in investment. Along with an increase in labor quality and TFP, this tends to push labor productivity up. On the other hand, a large decline in utilization pushes labor productivity down. At longer horizons, TFP turns negative and the protracted slump in investment shown in Figure 2 results in a decrease in capital intensity. These long-run effects are compensated by the improvement in utilization.

A large literature (see Benigno and Fornaro, 2018; Moran and Queralto, 2018; Guerron-Quintana and Jinnai, 2019; Ikeda and Kurozumi, 2019; Bianchi et al., 2019; Anzoategui et al., 2019; Garga and Singh, 2021; Schmöller and Spitzer, 2021) has developed New Keynesian models with endogenous growth in which demand shocks can have long-lasting effects on output. A key transmission mechanism in this literature is that a contractionary demand shock results in a decline in productivity-enhancing investment, notably research and development (R&D), which later triggers a persistent slowdown in TFP. We find evidence in favor of this channel. The top-left panel of Figure 9 shows that our permanent demand shock leads to a protracted decline in R&D investment. In turn, this decline in R&D investment could potentially be responsible for the negative response of TFP at long horizons that we observe in Figure 8.

Moreover, the fact that labor quality increases significantly in response to permanent demand shocks suggests that demand-driven recessions are periods of intense restructuring in which the least productive units or workers are disproportionately affected. According to this view, firms grow fat during economic expansions before aggressively restructuring in recessions (Berger, 2012). This leaves them better equipped to serve demand with a smaller workforce in the ensuing expansion, thereby leading to jobless recoveries, in line with our evidence. The top-right panel of Figure 9 validates this interpretation. The employment share for skilled workers increases after a permanent demand shock, which means that employment for skilled workers declines less than for other workers. As an additional piece of evidence, we also consider the routine employment share, defined as the ratio of employment of workers performing routine tasks (as classified in Jaimovich and Siu, 2020) over total employment, on the bottom-left panel of Figure 9. We observe that permanent demand shocks persistently displace workers performing routine tasks. This result is in keeping with the fact that job polarization takes place mainly in recessions and generates jobless recoveries, as shown by Jaimovich and Siu (2020).

It is reasonable to think that similar composition effects are at play also on the firm

Figure 9: IRFs of selected variables to permanent demand shocks

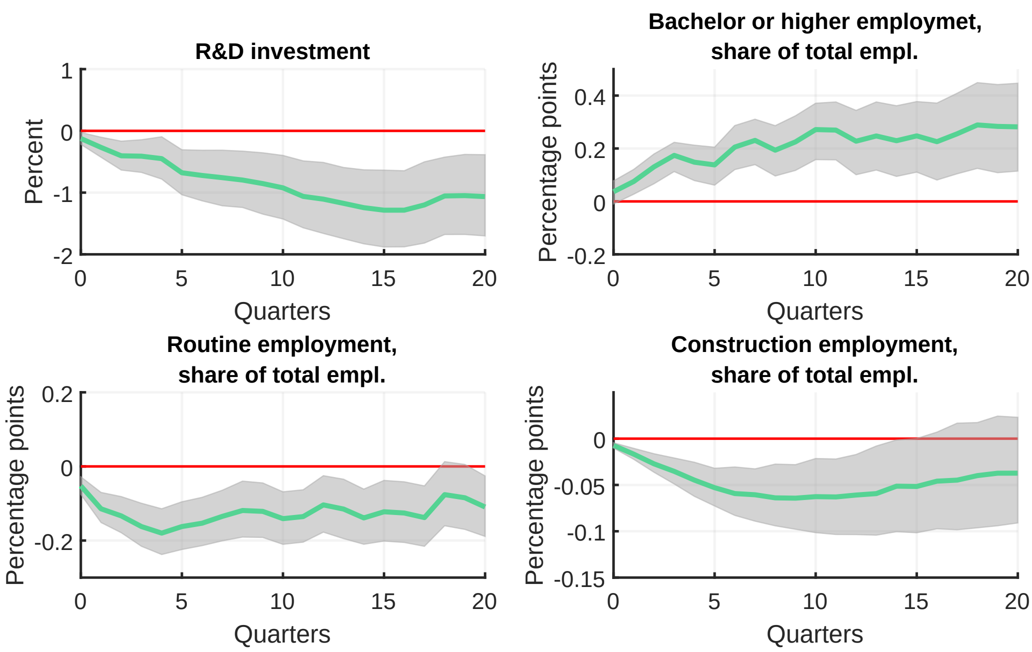

side, causing less efficient production units to become unprofitable and shut down. Our evidence is admittedly weaker in that dimension. However, one fact that is consistent with this narrative is the large and permanent effect of a negative demand shock on employment in the construction sector (see the bottom-right panel of Figure 9). Productivity in the construction sector is notoriously low and thus a shrinking level of economic activity in that sector will lead to an improvement in aggregate productivity.

All in all, we find that the limited effect of permanent demand shocks on labor productivity is the result of offsetting movements in labor and capital utilization, capital intensity, and labor quality. While utilization falls in the short run, both capital intensity and labor quality go up, thus explaining the neutral effect on the aggregate. We also uncover some evidence in favor of theories in which decreases in R&D investment in (demand-driven) recessions are followed by long-run declines in TFP growth.

5 The Great Recession and the Great Inflation

In this Section, we examine the sensitivity of our results to the sample under consideration. First, we check that our results are not driven by the presence of the Great Recession in the sample. Second, we discuss costs and benefits of estimating the model over a longer sample period including the Great Inflation of the 1970s.

5.1 The Role of the Great Recession

One obvious question of interest is whether the presence of the Great Recession, by far the largest recession in our data, is driving our results. To check this conjecture, we repeat our analysis over the period 1983:Q1-2007:Q4, thus excluding the Great Recession and its aftermath from the sample. The first line of Figure 10 shows the FEVD and the IRFs of output and employment to permanent demand shocks for this sample. Clearly, the role of permanent demand shocks is slightly lower in this period but still large. The IRFs show that the responses of output and employment to permanent demand shocks are marginally weaker than the ones reported in Figure 2 and the FEVD shows that the weight of permanent demand shocks diminishes when compared with the results in Figure 3. Hence, and not surprisingly, our model seems to suggest particularly strong hysteresis effects associated with the Great Recession. However, evidence of hysteresis is still present during the period associated with the Great Moderation. In Figure 11, we report selected results from the local projections analysis on the pre-Great Recession sample. Once again, all the key results on labor market variables (i.e. on unemployment, participation, disability applications and awards) are still present. Some of the impulse responses related to productivity are admittedly slightly weaker (this is the case for labor quality and TFP) but the broad message from these exercises is that hysteresis was present already in the pre-Great Recession period.

Another interesting question is whether our model sees the Great Recession as a large shock (or combination of shocks) or whether the propagation of shocks was different in such a peculiar episode. Since both the IRFs and the FEVD do not change substantially across the two samples, the SVAR seems to interpret the Great Recession mainly (although not exclusively) as a large shock. Still, one interesting question is whether the (rather limited) changes in the propagation of shocks are driven by a binding zero-lower-bound constraint on interest rates or weather they are related to other factors such as the financial origin of the shock or the presence of downward nominal wage rigidities. In order to isolate the role of both conventional and unconventional monetary policy, we estimate a version of our baseline model in which we introduce the shadow interest rate computed by Wu and Xia (2016) as an observable variable instead of investment. In Figure 12, we see that the impulse responses are almost identical to our baseline model over both samples (stopping in 2007:Q4 and stopping in 2019:Q4). This indicates that using the shadow rate instead of investment as the fourth variable in the SVAR does not change the propagation of shocks. Therefore, we now use the model with the shadow rate to investigate the role of monetary policy across samples and, in particular, during the zero-lower-bound period. The idea is to compare impulse responses from the model estimated over the full sample against a counterfactual in which monetary policy is unconstrained.[10] One way to do this is by assuming that the Federal Reserve follows the same policy

[10] We thank the editor for this suggestion.

Figure 10: IRFs and FEVD for output in alternative samples

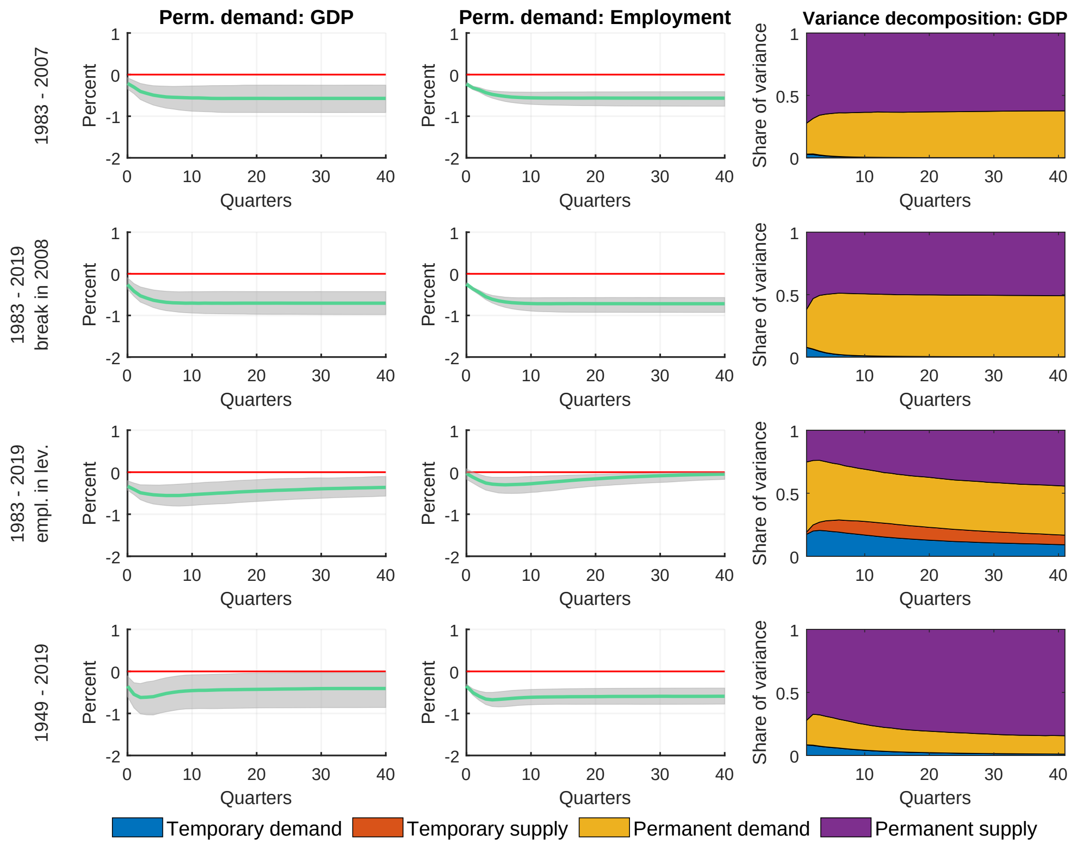

Figure 11: IRFs of selected variables to permanent demand shocks: sample 1983-2007

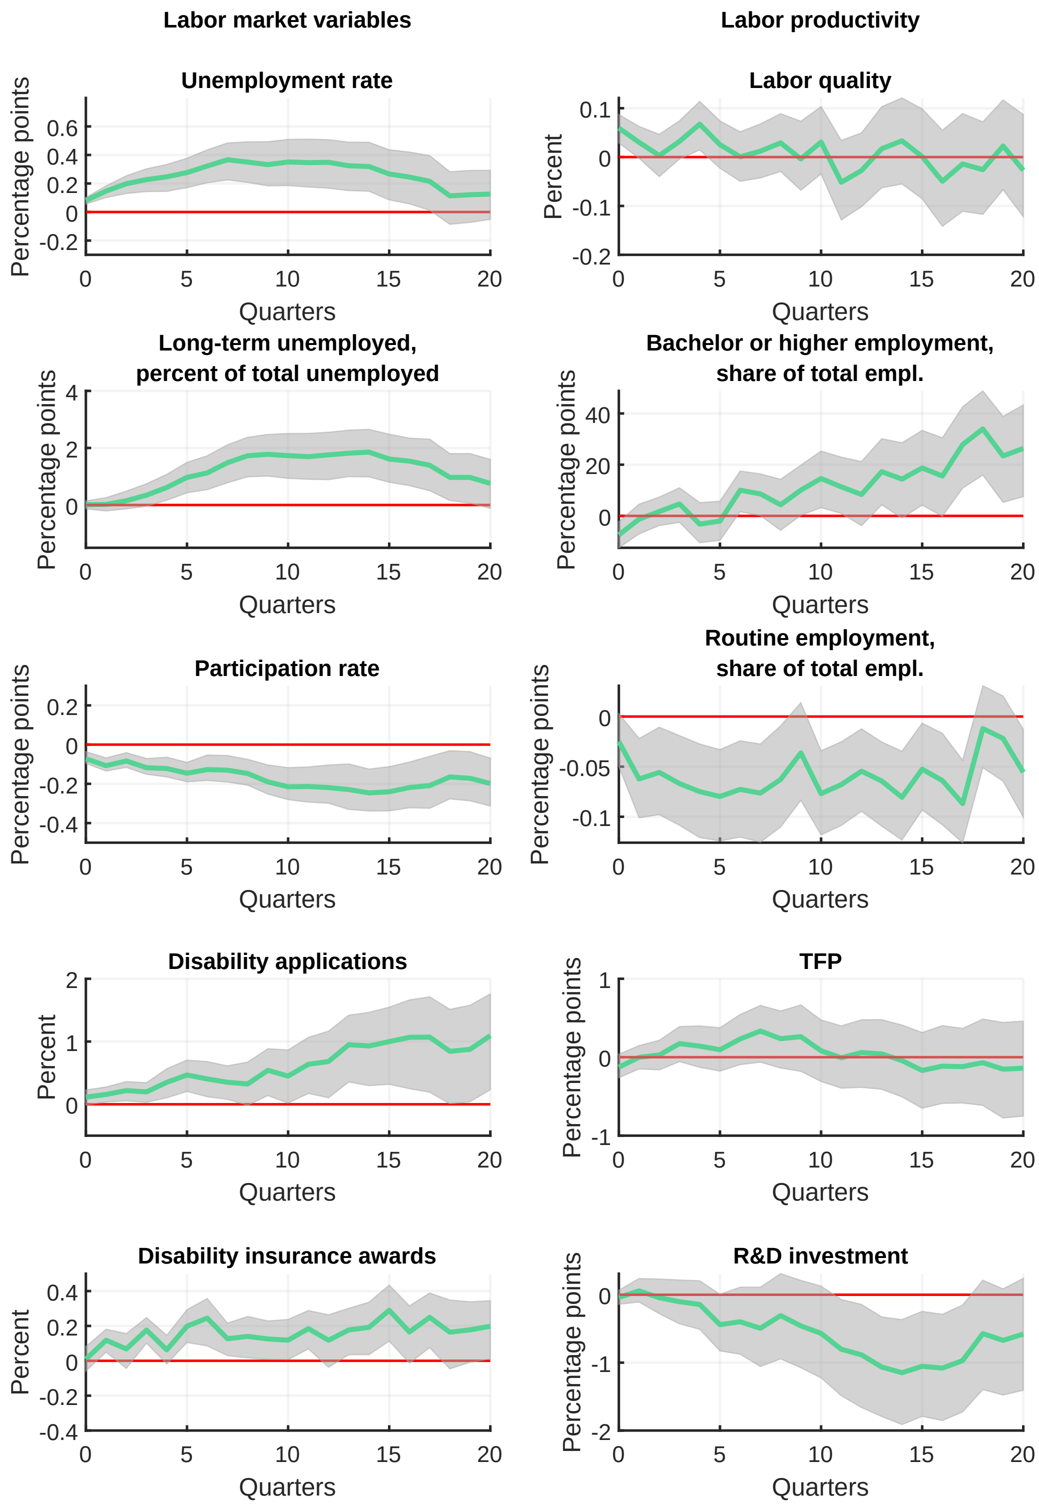

Figure 12: IRFs to the permanent demand and supply shock

Main model and a model with shadow rate estimated over different samples

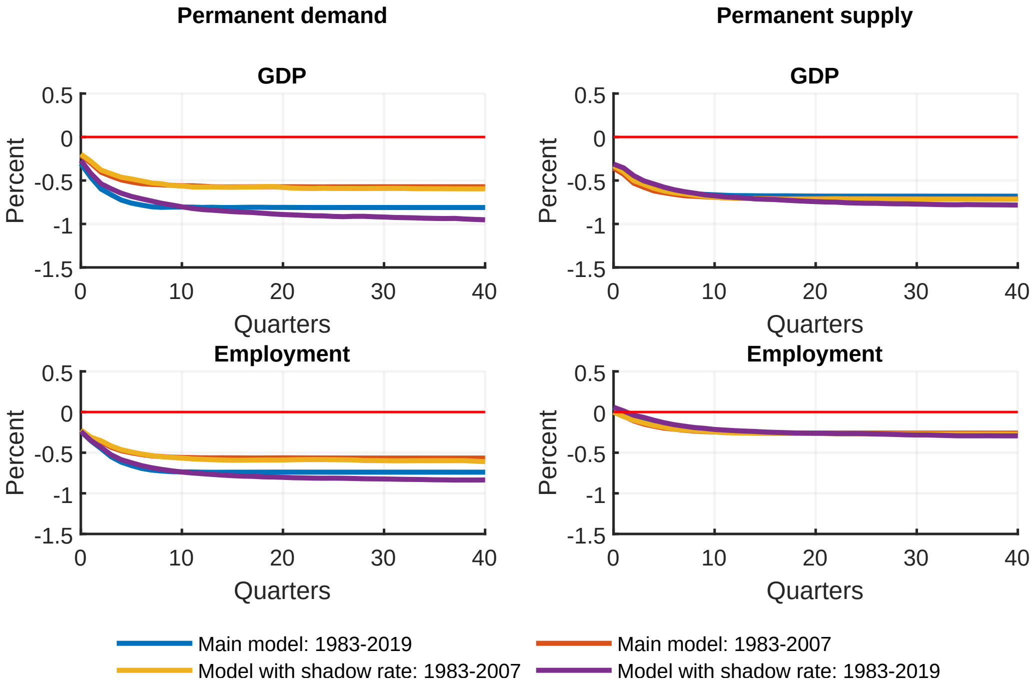

rule over the full sample as in the pre-Great Recession sample, and is not affected by the zero lower bound. In the counterfactual exercise, we do this by shifting the posterior parameter estimates in the interest rate equation so that the posterior mean is the same in the full sample as in the model estimated over the sample ending in 2007:Q4. For all other parameters, we keep the posterior from the full sample estimation. If the long-run responses to demand shocks are somewhat muted in the counterfactual where we change the interest rate equation, this could indicate that the zero lower bound has amplified the effects of demand shocks. We compare impulse responses estimated over the full sample (red lines) against impulse responses in the counterfactual with rescaled parameters in the interest rate equation (yellow lines) in Figure 13. We note that the two impulse responses for GDP and employment are almost indistinguishable, thus suggesting that differences in the propagation of shocks are not associated with changes in parameters in the interest rate equation. The interest rate itself responds more in the counterfactual exercise (yellow line) in the short-run but its response is less persistent over time when compared with the shadow rate (red line). All in all, parameters in the interest rate equation do not seem to explain the (rather limited) changes in the propagation of demand shocks over time. These results are in line with Bianchi and Melosi (2017), who find that uncertainty about fiscal policy rather than monetary policy is responsible for the comovements during the Great Recession.

Figure 13: IRFs to the permanent demand and supply shock

Model with shadow rate estimated over different samples

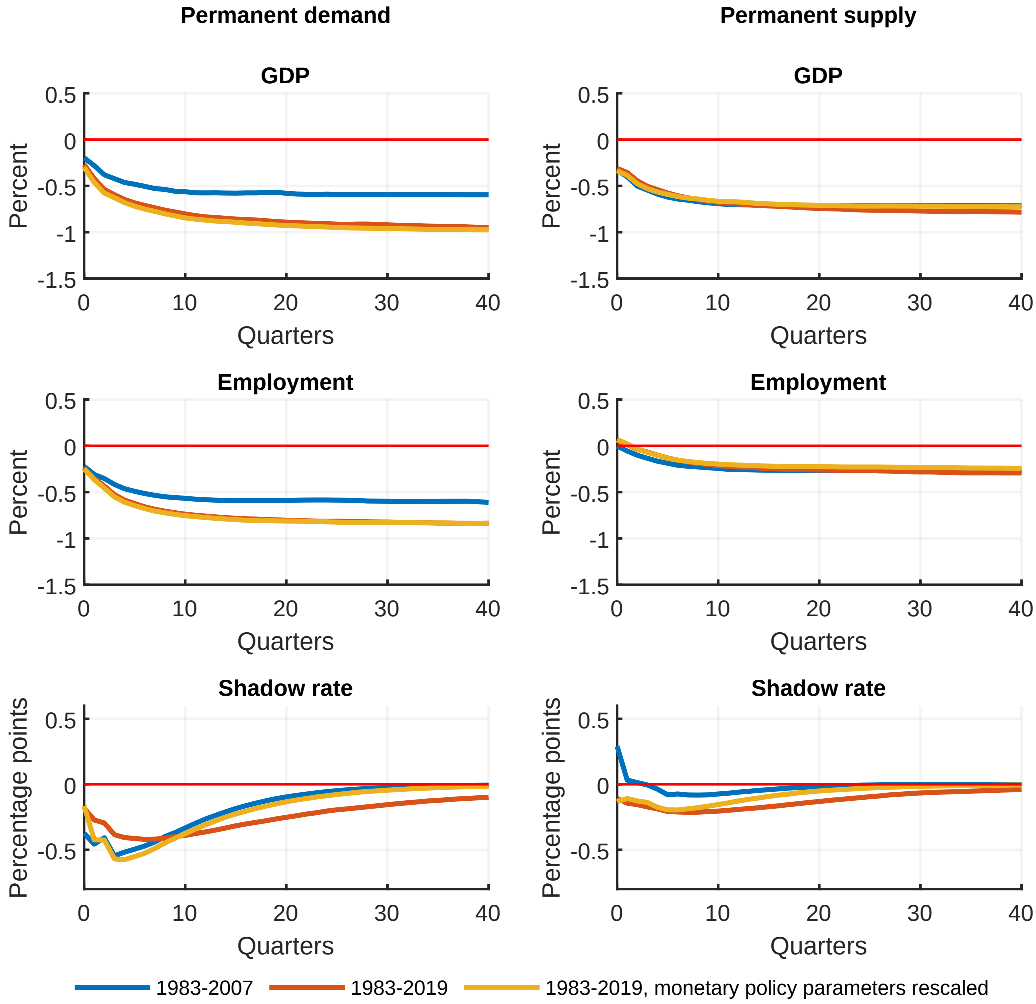

5.2 A Longer but Less Homogeneous Sample

Is our baseline sample sufficiently long to capture hysteresis effects? Ideally, a long sample is desirable when long-run restrictions are used (Faust and Leeper, 1997). However, the longer the sample, the more likely the presence of structural breaks, and it is well known that long-run restrictions are sensitive to structural breaks and low-frequency correlations (Fernald, 2007). Therefore, we face a trade-off between the length and the homogeneity of the sample. Our compromise is to use a sample starting in 1983, a relatively homogeneous sample period spanning 37 years of quarterly data. Coibion et al. (2018) discuss the very same trade-off and estimate Blanchard-Quah decompositions over samples of 30 years while Blanchard and Quah (1989) estimate their model over a sample of 37 years.

Given that VARs identified with long-run restrictions are sensitive to structural breaks, and although we are careful in choosing a relatively homogeneous sample, we follow the practical recommendation of Fernald (2007) and check whether our results are robust to various forms of detrending the data. First, we consider the following detrending methods:

(i) demeaning with a break imposed in 2008:Q1; *(ii)* quadratic detrending; *(iii)* detrending with a one-sided HP filter. The results for the first method is shown in the second line of Figure 10, while the results for the other two methods are reported in Figure A-7 in the Appendix. In all cases, we observe that our baseline results are confirmed. Second, we check that our results are robust to estimating a version of the model where the employment series is in levels rather than first differences. We perform this robustness check since studies that use long-run restrictions to estimate the effects of technology shocks find that whether hours enter in levels or first differences matters considerably for the results (Galí, 1999; Christiano et al., 2003; Basu et al., 2006; Fernald, 2007). In this case, all shocks have a transitory effect on employment as long as the system is stationary. The IRFs of output and employment to the permanent demand shock and the FEVD of output for this exercise are shown on the third line of Figure 10. As expected, the effects of the shock on employment are not permanent, although they are quite persistent. However, as in our baseline results, the shock has permanent effects on output and still explains a significant share of the forecast error variance of output at a 40-quarter horizon. Thus, these results show that our assumption that employment enters the SVAR in growth rates rather than levels is not crucial.

Keeping in mind our earlier warning that the heterogeneity of a longer sample may distort the results, we now examine how our results change when we extend our baseline sample further in the past. As a first exercise, we estimate our model over the longest possible post-war sample 1949:Q1-2019:Q4. Results are presented on the fourth line of Figure 10. Over this sample, permanent demand shocks have some long-run effects on output (the entire 68 percent confidence band is below the zero line) and rather large long-run effects on employment. Moreover, the role of permanent demand shocks in the variance decomposition of output is now substantially reduced but it is not negligible. These results seem to suggest that hysteresis is present even when using the full post-war sample. To examine this more closely, we conduct a second exercise where we estimate our model over an expanding window. Results are presented in Figure 14, where we show the contribution of permanent demand shocks to the FEVD of output at a 40-quarter horizon. Each point in the panel corresponds to a model estimated over a sample ending in 2019:Q4. The rightmost point corresponds to our baseline sample starting in 1983:Q1 while the leftmost point corresponds to the sample starting in 1949:Q1 that we just considered. We see that the large hysteresis effects on output that we detect in our baseline results progressively fade as the 1970s enter the sample. Samples starting in the 1950s or in the 1960s detect limited (although not negligible) hysteresis: in the long run most of the unexpected variation in output is explained by permanent supply shocks. Hence, the posterior probability that hysteresis effects are in fact present over the longest samples is low, although not zero. These results are consistent with Benati and Lubik (2022) who estimate a co-integrated SVAR over the sample 1954:Q1-2019:Q4 and find

Figure 14: Expanding window estimation

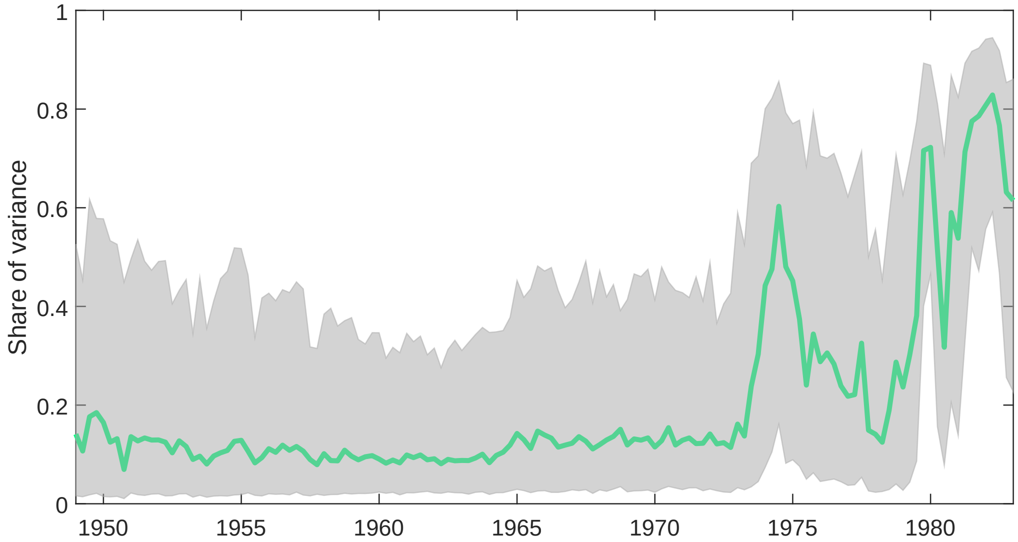

Note: Contribution of the permanent demand shock to the FEVD of GDP at a 40-quarter horizon. The estimation is conducted over an expanding window. The rightmost point in the chart corresponds to the sample 1983:Q1-2019:Q4. This initial sample is extended in the past as one moves left. The leftmost point in the chart corresponds to the sample 1949:Q1-2019:Q4.

limited evidence in favor of hysteresis.

How we can explain the sensitivity of our results to the sample under consideration? After all, Figure 14 shows that extending our baseline sample by only a few years reduces the importance of permanent demand shocks substantially. To gain intuition on this question, we rely on the analysis of a simple bi-variate system in output growth and inflation where we impose a long-restriction on output using the instrumental variables approach of Shapiro and Watson (1988). Our analysis follows closely the analytical discussion in Fernald (2007). Consider the following non-dynamic system

$$\Delta y_t = a\Delta^2 p_t + \varepsilon_t^P, \tag{3}$$
$$\Delta p_t = d\varepsilon_t^P + \varepsilon_t^T, \tag{4}$$

where Δy_t denotes output growth, Δp_t is inflation, Δp_t^2 is the change in inflation, ε_t^P is a shock with permanent effects on the level of output, and ε_t^T is a shock with temporary effects on the level of output. Since ε_t^P might affect the current change in inflation, in a first step one estimates a in equation (3) using past inflation as an instrument. In a second step, one can then estimate d in equation (4) with OLS. We are interested in the sign of d: if d is positive, permanent shocks lead to a positive contemporaneous co-movement between output growth and inflation that is characteristic of demand disturbances; if d is negative, permanent shocks lead to a negative contemporaneous co-movement between output growth and inflation that is characteristic of supply disturbances. The IV estimate of a is $\text{cov}\left(\Delta p_{-1}, \Delta y\right)/\text{cov}\left(\Delta p_{-1}, \Delta^2 p\right)$ and d equals $\text{cov}\left(\Delta p, \varepsilon^P\right)/\text{var}\left(\varepsilon^P\right)$. Given that the denominator of d is positive, we focus our attention on its numerator:

$$\text{cov}\left(\Delta p, \varepsilon^P\right) = \text{cov}\left(\Delta p, \Delta y - a\Delta^2 p\right) = \text{cov}\left(\Delta p, \Delta y\right) - \frac{\text{cov}\left(\Delta p_{-1}, \Delta y\right)}{\text{cov}\left(\Delta p_{-1}, \Delta^2 p\right)}\text{cov}\left(\Delta p, \Delta^2 p\right). \tag{5}$$

In the data, $\text{cov}\left(\Delta p, \Delta^2 p\right)$ is positive, in contrast to the negative $\text{cov}\left(\Delta p_{-1}, \Delta^2 p\right)$, and both are stable over time (these are the expected signs for a stationary autoregressive time series). In contrast, $\text{cov}\left(\Delta p, \Delta y\right)$ and $\text{cov}\left(\Delta p_{-1}, \Delta y\right)$ are unstable over time. Figure 15 plots the two covariances over an expanding window, where the rightmost point corresponds to the sample 2008:Q1-2019:Q4 and the leftmost point corresponds to the sample 1949:Q1-2019:Q4. We see that $\text{cov}\left(\Delta p, \Delta y\right)$ and $\text{cov}\left(\Delta p_{-1}, \Delta y\right)$ are firmly positive over our baseline sample, but that they turn negative as soon as the sample is extended in the 1970s. This helps explain our finding that permanent demand shocks are important over our baseline sample 1983:Q1-2019:Q4— in that sample d is positive and permanent shocks are more likely to look like demand shocks—but that their role shrinks as data from the 1970s enter the sample—in those samples d is negative and permanent shocks are more likely to look like supply shocks.[11]

We view these results as another practical illustration of the dangers of estimating a fixed-parameter SVAR over a long and heterogeneous sample, a point raised forcefully also by Bianchi et al. (2023). In the post-1980 period, regardless of the sample under consideration, the covariance between output and inflation is solidly positive and hysteresis is reliably detected by the SVAR. However, one needs to add only a few extra years of data, say by considering the sample 1978:Q1-2019:Q4, for this result to be turned on its head. Why? Because the extraordinary volatility in output growth and inflation observed during the years 1978 to 1982 dominates the smaller fluctuations observed between 1983 and 2019 in the computation of covariances. In other words, the model takes a lot of signal from the period 1978-1982 and little signal from the rest of the sample.

If one wishes to consider a long sample, using tools that can account for regime shifts, such as the transition from the high volatility environment of the 1970s to the calm of the Great Moderation, seems necessary. While these tools do exist—time-varying parameters and Markov-switching VARs (Bianchi and Melosi, 2017) being two examples, they cannot accommodate our identification scheme that relies on both sign and zero restrictions. Thus, we find it prudent to assign a greater weight to results obtained over homogeneous

[11] One may question the validity of using this toy model to interpret the results from the baseline SVAR. It turns out that the shocks extracted from the toy model and the VAR are highly positively correlated (a permanent shock in the SVAR is defined as the sum of permanent supply and demand shocks). This suggests that the main force at work in the toy model—the sign switch in $\text{cov}\left(\Delta p, \Delta y\right)$ and $\text{cov}\left(\Delta p_{-1}, \Delta y\right)$—is also driving the results in the baseline SVAR.

Figure 15: Evolution of the covariance between output and inflation over time.

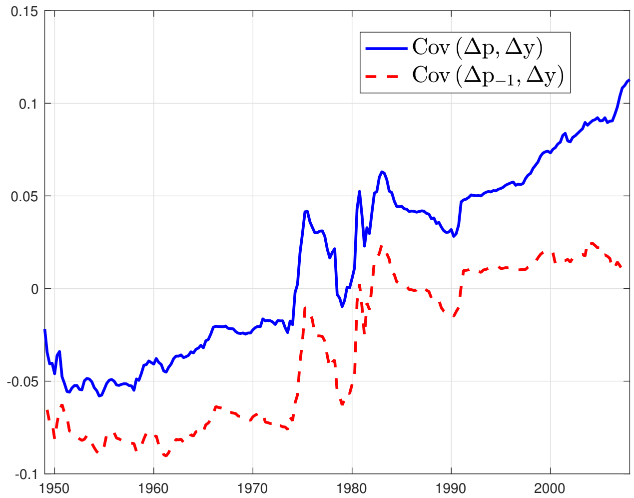

Note: Covariances between current output and current inflation (solid blue line) and current output and lagged inflation (red dashed line) over an expanding window. All samples end in 2019:Q4. The rightmost point in the chart corresponds to the sample 2008:Q1-2019:Q4. This initial sample is extended in the past as one moves left. The leftmost point in the chart corresponds to the sample 1949:Q1-2019:Q4.

samples.

6 Commingling with Other Shocks

In this Section, we evaluate the importance of permanent demand shocks in larger systems where we separately identify monetary shocks, labor supply shocks or sectoral supply shocks. We also briefly discuss the limited role of transitory shocks in our baseline results.

In a first extension, we consider two demand shocks with potentially permanent effects and we include the shadow rate computed by Wu and Xia (2016) as an additional variable in the system: a monetary and a non-monetary demand shock are distinguished on the basis of the short-run co-movement between inflation and interest rates. In Figure 16 we report the impulse responses of GDP to the two demand shocks and its FEVD. The two demand shocks explain jointly the lion's share of fluctuations with the non-monetary shock being the dominant driver. As shown by the IRFs, both shocks generate hysteresis effects. A long-run effect of monetary shocks on output is inconsistent with money neutrality but should be taken with caution being only marginally significant at the 68 percent level. More refined measures of monetary shocks should be used to investigate more in depth this surprising result. Jordà et al. (2020) propose a first important step in that direction in a local projection analysis using historical data over the period 1890-2015 (see also Bernanke and Mihov (1998) as an earlier attempt at testing the long-run money neutrality hypothesis).

Figure 16: IRFs and FEVD in the model with monetary shocks

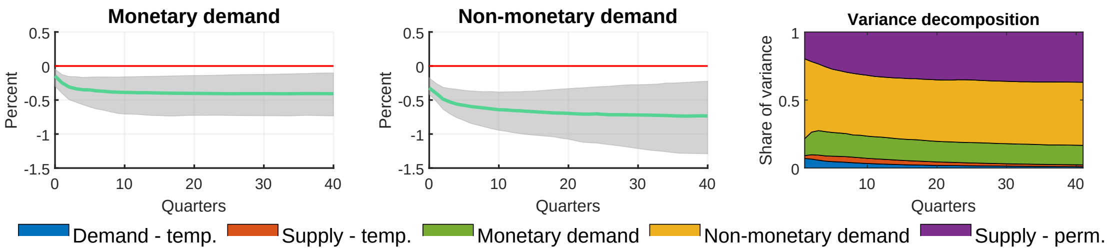

In a second experiment, we consider two permanent supply shocks and we include the real wage (measured as Average Hourly Earnings of Production and Nonsupervisory Employees deflated by the PCE deflator) as an additional variable in the system: a labor supply shock generates a negative short-run co-movement between real wages and GDP while a technology shock generates a positive co-movement. In addition, we assume that labor supply shocks have no long-run effects on labor productivity, a property satisfied by the neoclassical model. This exercise is particularly important in order to make sure that the long-run effects of permanent demand shocks on the employment-to-population ratio are not confounded with demographic changes. As shown in Figure 17, the permanent demand shock is still by far the main driver of output fluctuations with the labor supply shock absorbing explanatory power mainly from the technology shock. The labor supply shock is an important driver of employment but the permanent demand shock still generates important hysteresis effects in the labor market. In Figure A-8 in the Appendix, we present a second experiment to account for demographic factors in which we re-estimate our baseline model using the prime-age male employment-to-population ratio instead of the overall employment-to-population ratio. This change in variable helps controlling for labor supply factors such as women's increased labor force participation or changes in the length of education and in the retirement age. The role of permanent demand shocks is even more important in this specification.

In a third experiment, we evaluate the role of sectoral supply shocks. Theoretically, these shocks may lead to a positive co-movement between output growth and inflation in the aggregate in the presence of strong complementarities across sectors, as in Guerrieri et al. (2022). Thus, in principle, it is possible that our identification scheme labels these so-called "Keynesian" supply shocks as demand shocks. However, we think that there are at least two reasons to believe that Keynesian supply shocks are not important drivers of our identified permanent demand shock. First, these shocks are assumed to be transitory by Guerrieri et al. (2022), while our results show that permanent demand shocks (not transitory ones) are the important drivers of fluctuations in output and employment. Thus, to the extent that Keynesian supply shocks are captured in our analysis, they should be bundled up with transitory demand shocks, which have limited effects on real variables. Second, Baqaee and Farhi (2022) show that transitory sectoral supply shocks

Figure 17: IRFs and FEVDS in the model with labor supply and technology shocks

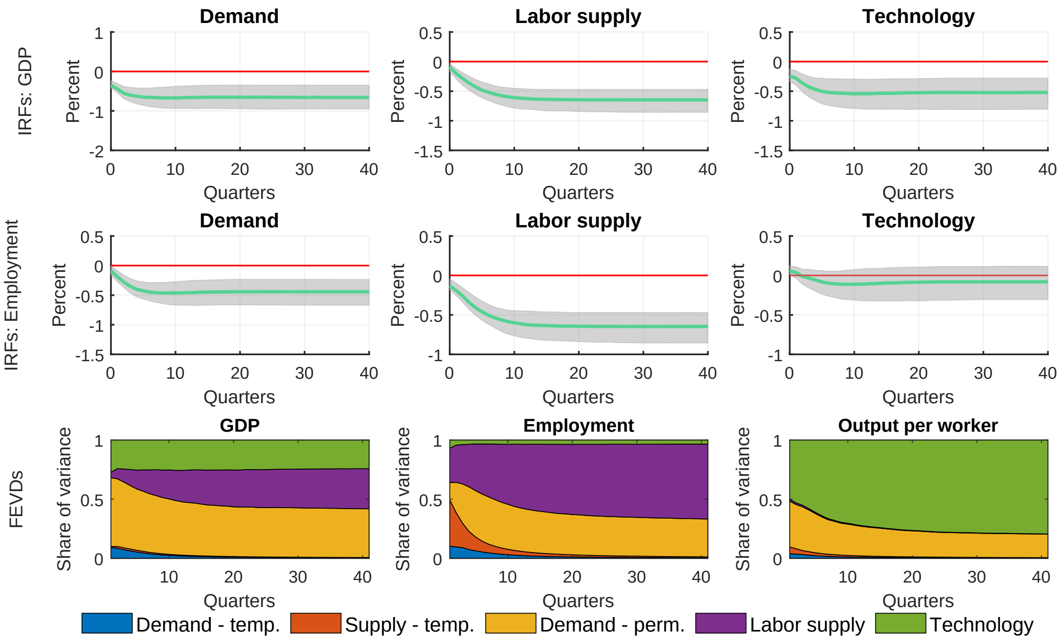

generate negative co-movement between output and inflation once input-output linkages are properly taken into account.[12] There are other kinds of sectoral supply shocks that may generate permanent effects, with the most natural candidates being shocks to the investment-specific technology. Therefore, we extend our model by including the relative price of investment as an observable. Following Fisher (2006), we identify an investment-specific technology shock as the only shock having a long-run effect on the relative price of investment. We also impose the innocuous assumption that the shock generates a negative co-movement between output (and investment) and the relative price of investment, while leaving inflation unrestricted. As shown in Figure 18, investment-specific shocks play a limited role in the model (except obviously for the relative price of investment) and permanent demand shocks retain an important explanatory power.

Finally, we address the limited role of transitory supply and demand shocks in our results. As shown for our baseline results in Figure 3, transitory shocks explain only a small fraction of fluctuations in output at all horizons. This result also holds for different samples, as seen in Figure 10, and may appear puzzling to the reader. We see two ways of accounting for this finding. First, it owes, in part, to our choice of variables. Indeed, swapping either employment or investment with unemployment, which is widely used to assess the cyclical position of the economy, results in a larger role for transitory shocks.

[12] Fornaro and Wolf (2020) show that *aggregate* supply shocks can also generate Keynesian effects under specific parameterizations in a Keynesian growth framework. However, these shocks propagate mainly through labor productivity, unlike our permanent demand shocks. For empirical evidence on the effects of Keynesian supply shocks, see Cesa Bianchi and Ferrero (2020).

Figure 18: Forecast error variance decomposition

(Model with the relative price of investment as an observable)

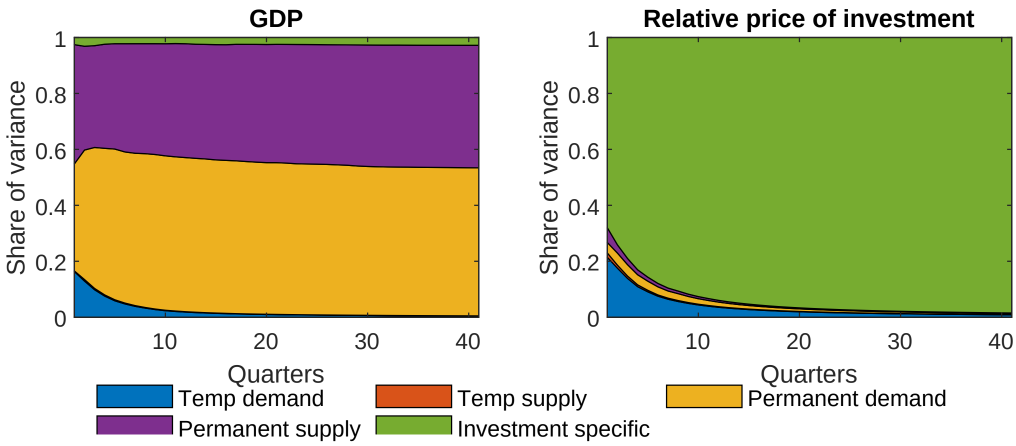

As shown in Figure A-9 in the Appendix, with unemployment in the SVAR transitory shocks account for about 25 percent of the FEVD of output and more than 50 percent of the FEVD of unemployment at short horizons. Second, our variance decomposition exercises do not show the contributions of shocks to business cycle fluctuations, but rather to deviations of variables from their deterministic trends. Consider the case of output. Because average output growth varies over time, the deviation of the output level from its deterministic trend mixes both a low-frequency component and a high-frequency component, and the size of the low-frequency component dwarfs that of the high-frequency component. Since the low-frequency component is necessarily explained by the permanent shocks, it is therefore not necessarily surprising that these shocks have a dominant role in the FEVD of output.

7 Monte Carlo Analysis

In this Section, we assess the performance of a SVAR model driven by permanent and transitory shocks in a controlled model environment. We aim to answer the following questions. First, if hysteresis is indeed present in the data, will the SVAR be able to detect it? Second, if hysteresis is absent from the data but fluctuations are persistent, is there a risk that we incorrectly detect the presence of hysteresis?

To answer these questions, we run a Monte Carlo experiment using the insider-outsider model of Galí (2022). This theoretical model is particularly suited for such an exercise since it allows for the possibility of hysteresis in the labor market, as in our empirical model. In the model, hysteresis in employment and output arises as the number of insiders in wage bargaining evolves endogenously in response to shocks. The model features three shocks: preference, technology, and price-markup. In Galí's original model, all three shocks have permanent effects on output and employment. To allow for the possibility that at least one shock has only temporary effects, we introduce a direct response of monetary policy to the price-markup shock. Therefore, the modified model features two supply shocks, one with temporary effects (price-markup shock) and one with permanent effects (technology shock), as well as a demand shock with permanent effects (preference shock). We scale the shock variances so that each shock's relative importance for output at a horizon of 40 quarters resembles what we obtain in our baseline empirical estimates. This modified model is used to simulate artificial data on which we estimate a SVAR model.

To focus on the evaluation of the identification scheme, we limit the sources of misspecification by estimating a 3-variable SVAR model that contains real wages, a key endogenous state variable in the theoretical model, alongside output and inflation. Note that the SVAR is still misspecified to the extent that it does not include all the endogenous state variables of the insider-outsider model. We identify the same shocks as in the data generating process: a temporary supply shock, a (potentially) permanent supply shock, and a (potentially) permanent demand shock. Sign restrictions on output and inflation are used to disentangle supply from demand and a zero long-run restriction on output is used to disentangle the temporary supply shock from the potentially permanent supply shock. For simplicity, we also eliminate small-sample uncertainty by estimating the SVAR on a large sample of 10,000 observations.[13]

What could possibly go wrong at this stage? A common concern with sign restrictions is that of identification uncertainty. Sign restrictions are weak information and can produce unreliable results if shocks have small variance or do not produce a mutually exclusive pattern of signs for the impulse responses (Canova and Paustian, 2011). A particular example is given in Wolf (2020), who shows that the surprising results of Uhlig (2005) regarding the effects of monetary policy can be explained by the fact that a linear combination of demand and supply shocks "masquerade" as monetary policy shocks in his framework. Although identification is strengthened by the presence of zero restrictions in our model, a similar issue could arise: for example, a combination of price-markup shocks (temporary supply) and preference shocks (permanent demand) could "masquerade" as technology shocks (permanent supply).

Figure 19 presents the results of our Monte Carlo exercise. We contrast the theoretical and estimated IRFs of output to preference and technology shocks in two distinct cases: a situation where there is hysteresis in the data generating process (upper panels), and a situation where there is no hysteresis (lower panels). In all panels, the red dashed lines are the theoretical IRFs from the insider-outsider model and the solid green lines and accompanying shaded areas are the median IRFs and the 68 percent confidence interval obtained from the SVAR.

We are now in a position to answer our two initial questions. Starting with the first

[13] In Figure A-10 in the Appendix, we consider the case of finite samples of the same length as our baseline sample. We find very similar results.

Figure 19: Monte Carlo exercise in population: Output IRFs

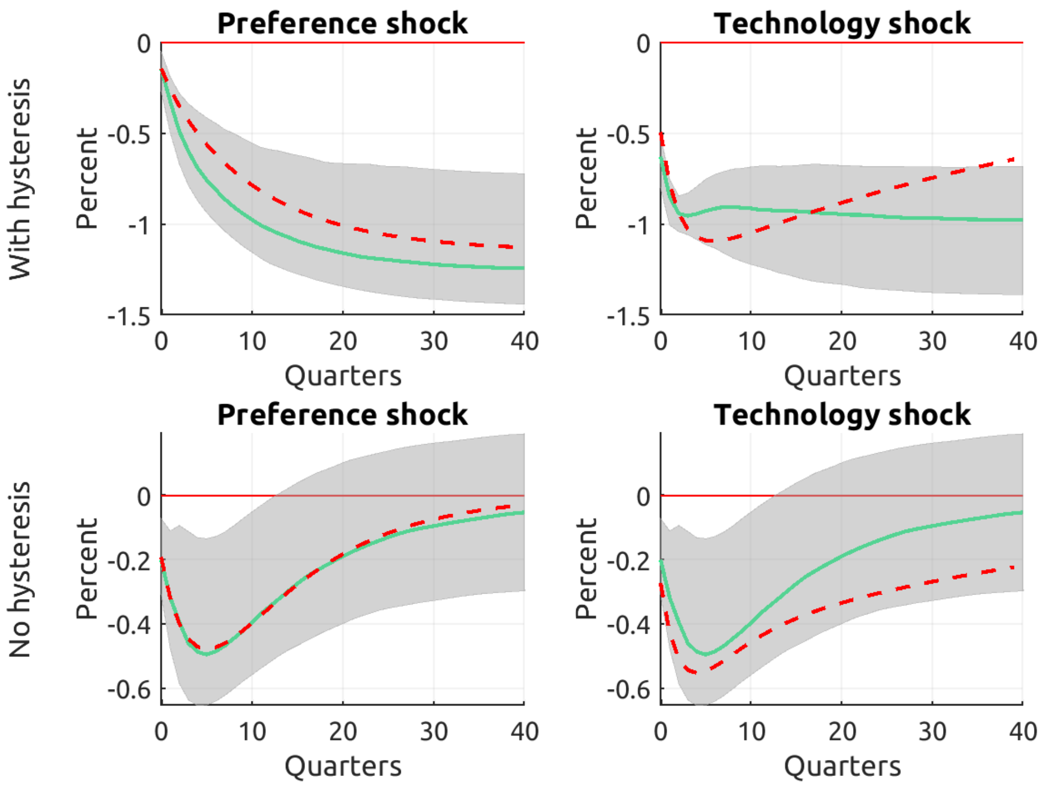

Note: Dashed red lines are the IRFs from the theoretical model. Solid green lines are the median IRFs from the SVAR. Shaded areas indicate the 16th and 84th percentiles. The first line reports IRFs of output to preference shock and technology shocks in the case *with* hysteresis. The second line reports IRFs of output to preference shocks and technology shocks in the case *without* hysteresis.

question, we see in the upper left panel that the median IRF of output to a preference shock estimated from the SVAR closely tracks the theoretical IRF and that the 68 percent confidence interval lies clearly below the zero line. Thus, we conclude that the SVAR will most likely be able to detect hysteresis if there is indeed hysteresis in the data. Turning to our second question, we should first note that the preference and technology shocks in the theoretical model have a persistence of 0.91 and 0.98 respectively, and we indeed see on the bottom line of Figure 19 that the theoretical IRFs of output to these two shocks return only sluggishly to zero. Yet, the SVAR does a fine job at capturing the main properties of the data generating process. Notably, the median IRF to preference shocks estimated from the SVAR is nearly identical to the theoretical IRF and the 68 percent confidence interval is centered around zero at a 40-quarter horizon. Thus, although the shocks are persistent, the SVAR evidence allows us to confidently conclude that there is no hysteresis in the data.[14]

8 Conclusion

In this paper we have challenged the independence assumption embedded in macroeconomic analysis. This assumption implies that output can be decomposed into a trend, in which surprises are driven only by supply shocks, and transitory fluctuations around the trend, mostly driven by demand shocks. We have shown that demand shocks may generate hysteresis effects. Recessions (and booms) driven by demand shocks may have permanent effects on output and employment. In particular, our results have shown that permanent demand shocks explain a significant share of the decline in real activity in the aftermath of the Great Recession. Hysteresis effects transmit through employment but do not affect output per worker. While our paper is purely empirical and does not provide normative implications, we believe it is important to have sound empirical evidence on the relevance of hysteresis effects to inform the policy discussion (cf. Garga and Singh, 2021; Galí, 2022).

It is also worth stressing that our simple analysis is only a first step toward estimating hysteresis effects. As shown by Benigno et al. (2015), non-linearities are potentially important in studying unemployment, labor productivity and their drivers. Introducing non-linearities in our set-up is certainly promising and desirable, although far from trivial insofar as the literature has not reached a consensus on how to integrate sign restrictions into non-linear models.

Another avenue for future research consists of disentangling further the origin of hysteresis effects. Bianchi et al. (2019) and Guerron-Quintana and Jinnai (2019) find an

[14] One may also wonder whether the SVAR model is able to recover the variance decomposition of the forecast errors from the theoretical model. Since the variance decomposition is based on the IRFs, and the SVAR does a good job at recovering the theoretical IRFs, the answer to this question is yes. This is true not only for output, but also inflation and wages.

important role for shocks related to investment (shocks to the marginal efficiency of investment and liquidity shocks, respectively), while supporting evidence on the long-run effects of monetary and fiscal shocks is provided in Jordà et al. (2020) and Fatás and Summers (2018).[15] Fernández-Villaverde et al. (2019) document the long-lasting effects of discount factor shocks in the presence of search complementarities. Farmer (2012) studies the persistent effects of confidence shocks on unemployment in Keynesian models. All of these shocks are bundled together in our analysis, and disentangling the different components would be worthwhile at the cost, however, of compromising the simplicity of our approach.

[15]In standard New Keynesian theory, demand shocks with long-lasting effects (like permanent changes in government purchases or permanent changes in the labor income tax rate when combined with hand-to-mouth behavior) are not usually seen as dominant drivers of the business cycle. Instead, the demand shocks driving the business cycle in these models (like financial shocks, preference or discount factor shocks) do not have long-run effects. It is in this sense that we find that our results present a challenge to standard business cycle theory and are important: it is the demand shocks that are driving the business cycle that are also found to have permanent effects.

References

Aaronson, S., Daly, M., Wascher, W., and Wilcox, D. (2019). Okun revisited: Who benefits most from a strong economy? *Brookings Papers on Economic Activity*, 2019:333–404.

Altig, D., Christiano, L., Eichenbaum, M., and Linde, J. (2011). Firm-specific capital, nominal rigidities and the business cycle. *Review of Economic Dynamics*, 14(2):225–247.

Angeletos, G.-M., Collard, F., and Dellas, H. (2020). Business cycle anatomy. *American Economic Review*, 110(10):3030–3070.

Antolin-Diaz, J., Drechsel, T., and Petrella, I. (2017). Tracking the slowdown in long-run GDP growth. *Review of Economics and Statistics*, 99(2):343–356.

Antolín-Díaz, J. and Rubio-Ramírez, J. F. (2018). Narrative sign restrictions for svars. *American Economic Review*, 108(10):2802–2829.

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.

Arias, J., Rubio-Ramirez, J., and Waggoner, D. (2018). Inference based on structural vector autoregressions identified with sign and zero restrictions: Theory and applications. *Econometrica*, 86(2):685–720.

Arias, J., Rubio-Ramirez, J., and Waggoner, D. (2023). Uniform priors for impulse responses.

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

Ball, L., Mankiw, G., and Nordhaus, W. (1999). Aggregate demand and long-run unemployment. *Brookings Papers on Economic Activity*, 1999(2):189–251.

Baqaee, D. and Farhi, E. (2022). Supply and demand in disaggregated keynesian economies with an application to the covid-19 crisis. *American Economic Review*, 112(5):1397–1436.

Basu, S., Fernald, J., and Kimball, M. (2006). Are technology improvements contractionary? *American Economic Review*, 96(5):1418–1448.

Baumeister, C. and Hamilton, J. D. (2015). Sign Restrictions, Structural Vector Autoregressions, and Useful Prior Information. *Econometrica*, 83(5):1963–1999.

Benati, L. and Lubik, T. (2022). Searching for hysteresis. Federal Reserve Bank of Richmond Working Paper 22-05.

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

Benigno, P., Ricci, L. A., and Surico, P. (2015). Unemployment and productivity in the long run: the role of macroeconomic volatility. *Review of Economics and Statistics*, 97(3):698–709.

Berger, D. (2012). Countercyclical restructuring and jobless recoveries. 2012 Meeting Papers 1179, Society for Economic Dynamics.

Bernanke, B. S. and Mihov, I. (1998). The liquidity effect and long-run neutrality. *Carnegie-Rochester Conference Series on Public Policy*, 49:149–194.

Bhattarai, S., Schwartzman, F., and Yang, C. (2021). Local scars of the US housing crisis. *Journal of Monetary Economics*, 119:40–57.

Bianchi, F., Kung, H., and Morales, G. (2019). Growth, slowdowns, and recoveries. *Journal of Monetary Economics*, 101:47–63.

Bianchi, F. and Melosi, L. (2017). Escaping the great recession. *American Economic Review*, 107(4):1030–58.

Bianchi, F., Nicolò, G., and Song, D. (2023). Inflation and real activity over the business cycle. National Bureau of Economic Research Working Paper 31075.

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

Blanchard, O., Cerutti, E., and Summers, L. (2015). Inflation and activity–two explorations and their monetary policy implications. NBER Working Paper 21726.

Blanchard, O. and Quah, D. (1989). The dynamic effects of aggregate demand and aggregate supply. *American Economic Review*, 79(4):655–673.

Blanchard, O. and Summers, L. (1986). Hysteresis and the European unemployment problem. *NBER Macroeconomics Annual*, 1:15–78.

Bluedorn, J. and Leigh, D. (2019). Hysteresis in labor markets? Evidence from professional long-term forecasts. Working Paper 19/114, International Monetary Fund.

Canova, F. and De Nicoló, G. (2002). Monetary disturbances matter for business fluctuations in the G-7. *Journal of Monetary Economics*, 49(6):1131–1159.

Canova, F. and Paustian, M. (2011). Business cycle measurement with some theory. *Journal of Monetary Economics*, 58(4):345–361.

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

Cerra, V. and Saxena, S. (2008). Growth dynamics: the myth of economic recovery. *American Economic Review*, 98(1):439–57.

Cesa Bianchi, A. and Ferrero, A. (2020). The transmission of Keynesian supply shocks. Manuscript.

Christiano, L. J., Eichenbaum, M., and Vigfusson, R. (2003). What happens after a technology shock? NBER Working Paper 9819.

Cloyne, J., Martinez, J., Mumtaz, H., and Surico, P. (2022). Short-term tax cuts, long-term stimulus. National Bureau of Economic Research Working Paper 30246.

Coibion, O., Gorodnichenko, Y., and Koustas, D. (2013). Amerisclerosis? The puzzle of rising US unemployment persistence. *Brookings Papers on Economic Activity*, 2013(2):193–260.

Coibion, O., Gorodnichenko, Y., and Ulate, M. (2018). The cyclical sensitivity in estimates of potential output. *Brookings Papers on Economic Activity*, pages 343–441.

Del Negro, M., Lenza, M., Primiceri, G., and Tambalotti, A. (2020). What's up with the Phillips curve? Brooking Papers on Economic Activity Conference Draft, Spring.

Doan, T., Litterman, R., and Sims, C. (1984). Forecasting and conditional projection using realistic prior distributions. *Econometric Reviews*, 3(1):1–100.

Dolado, J. and Jimeno, J. F. (1997). The causes of Spanish unemployment: A structural VAR approach. *European Economic Review*, 41(7):1281–1307.

Dosi, G., Pereira, M. C., Roventini, A., and Virgillito, M. E. (2018). Causes and consequences of hysteresis: aggregate demand, productivity, and employment. *Industrial and Corporate Change*, 27(6):1015–1044.

Eo, Y. and Morley, J. (2022). Why has the us economy stagnated since the great recession? *The Review of Economics and Statistics*, 104(2):246–258.

Farmer, R. E. (2012). Confidence, crashes and animal spirits. *The Economic Journal*, 122(559):155–172.

Fatás, A. and Summers, L. (2018). The permanent effects of fiscal consolidations. *Journal of International Economics*, 112:238–250.

Faust, J. (1998). The robustness of identified VAR conclusions about money. *Carnegie-Rochester Conference Series on Public Policy*, 49:207–244.

Faust, J. and Leeper, E. M. (1997). When do long-run identifying restrictions give reliable results? *Journal of Business & Economic Statistics*, 15(3):345–353.

Fernald, J. (2007). Trend breaks, long-run restrictions, and contractionary technology improvements. *Journal of Monetary Economics*, 54(8):2467–2485.

Fernald, J. (2014). A quarterly, utilization-adjusted series on total factor productivity. Working Paper 2012-19, Federal Reserve Bank of San Francisco.

Fernald, J., Stock, J., Hall, R., and Watson, M. (2017). The disappointing recovery of output after 2009. *Brookings Papers on Economic Activity*, 48:1–58.

Fernández-Villaverde, J., Mandelman, F., Yu, Y., and Zanetti, F. (2019). Search complementarities, aggregate fluctuations, and fiscal policy. NBER Working Paper 26210.

Fisher, J. D. (2006). The dynamic effects of neutral and investment-specific technology shocks. *Journal of Political Economy*, 114(3):413–451.

Fornaro, L. and Wolf, M. (2020). The scars of supply shocks. CEPR Discussion Paper 15423.

Galí, J. (1999). Technology, employment, and the business cycle: do technology shocks explain aggregate fluctuations? *American Economic Review*, 89(1):249–271.

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

Galí, J. and Hammour, M. (1992). Long run effects of business cycles. Working Paper Columbia - Graduate School of Business.

Galí, J., López-Salido, D., and Vallés, J. (2003). Technology shocks and monetary policy: assessing the Fed's performance. *Journal of Monetary Economics*, 50(4):723–743.

Galí, J. and Rabanal, P. (2004). Technology shocks and aggregate fluctuations: How well does the real business cycle model fit postwar us data? *NBER Macroeconomics Annual*, 19:225–288.

Gambetti, L. and Pistoresi, B. (2004). Policy matters. the long run effects of aggregate demand and mark-up shocks on the Italian unemployment rate. *Empirical Economics*, 29(2):209–226.

Garga, V. and Singh, S. (2021). Output hysteresis and optimal monetary policy. *Journal of Monetary Economics*, 117:871–886.

Giannone, D., Lenza, M., and Primiceri, G. E. (2015). Prior selection for vector autoregressions. *Review of Economics and Statistics*, 97(2):436–451.

Gordon, R. (2015). Secular stagnation: A supply-side view. *American Economic Review*, 105(5):54–59.

Guerrieri, V., Lorenzoni, G., Straub, L., and Werning, I. (2022). Macroeconomic implications of covid-19: Can negative supply shocks cause demand shortages? *American Economic Review*, 112(5):1437–74.

Guerron-Quintana, P. A. and Jinnai, R. (2019). Financial frictions, trends, and the great recession. *Quantitative Economics*, 10(2):735–773.

Ikeda, D. and Kurozumi, T. (2019). Slow post-financial crisis recovery and monetary policy. *American Economic Journal: Macroeconomics*, 11(4):82–112.

Jacobson, T., Vredin, A., and Warne, A. (1997). Common trends and hysteresis in Scandinavian unemployment. *European Economic Review*, 41(9):1781–1816.

Jaimovich, N. and Siu, H. (2020). Job polarization and jobless recoveries. *Review of Economics and Statistics*, 102(1):129–147.

Jordà, Ò. (2005). Estimation and Inference of Impulse Responses by Local Projections. *American Economic Review*, 95(1):161–182.

Jordà, Ò., Singh, S., and Taylor, A. (2020). The long-run effects of monetary policy. Working Paper 2020-09, Federal Reserve Bank of San Francisco.

Kozlowski, J., Veldkamp, L., and Venkateswaran, V. (2020). The tail that wags the economy: Beliefs and persistent stagnation. *Journal of Political Economy*, 128(8):2839–2879.

Maffei-Faccioli, N. (2021). Identifying the sources of the slowdown in growth: Demand vs. supply. Norges Bank Working Paper 9/2021.

Mertens, K. and Ravn, M. (2011). Technology-hours redux: Tax changes and the measurement of technology shocks. *NBER International Seminar on Macroeconomics*, 7(1):41–76.

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.

Moran, P. and Queralto, A. (2018). Innovation, productivity, and monetary policy. *Journal of Monetary Economics*, 93:24–41.

Paciello, L. (2011). Does inflation adjust faster to aggregate technology shocks than to monetary policy shocks? *Journal of Money, Credit and Banking*, 43(8):1663–1684.

Plagborg-Møller, M. and Wolf, C. K. (2021). Local projections and vars estimate the same impulse responses. *Econometrica*, 89(2):955–980.

Ramey, V. A. and Zubairy, S. (2018). Government Spending Multipliers in Good Times and in Bad: Evidence from US Historical Data. *Journal of Political Economy*, 126(2):850–901.

Rubio-Ramirez, J., Waggoner, D., and Zha, T. (2010). Structural vector autoregressions: Theory of identification and algorithms for inference. *Review of Economic Studies*, 77(2):665–696.

Schmöller, M. E. and Spitzer, M. (2021). Deep recessions, slowing productivity and missing (dis-) inflation in the euro area. *European Economic Review*, 134:103708.

Shapiro, M. D. and Watson, M. W. (1988). Sources of business cycle fluctuations. *NBER Macroeconomics annual*, 3:111–148.

Stock, J. H. and Watson, M. W. (2018). Identification and Estimation of Dynamic Causal Effects in Macroeconomics Using External Instruments. *The Economic Journal*, 128(610):917–948.

Summers, L. (2014). Fiscal policy and full employment. *Speech at Center on Budget and Policy Priorities Event on Full Employment.*

Summers, L. (2015). Demand side secular stagnation. *American Economic Review*, 105(5):60–65.

Uhlig, H. (2005). What are the effects of monetary policy on output? Results from an agnostic identification procedure. *Journal of Monetary Economics*, 52(2):381–419.

Watson, M. W. (2020). Comment. *NBER Macroeconomics Annual*, 34:182–193.

Wolf, C. K. (2020). Svar (mis)identification and the real effects of monetary policy shocks. *American Economic Journal:Macroeconomics*, 12(4):1–32.

Wu, J. C. and Xia, F. D. (2016). Measuring the macroeconomic impact of monetary policy at the zero lower bound. *Journal of Money, Credit and Banking*, 48(2-3):253–291.

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

For Online Publication: Appendix

Additional results

- Figure A-1 replicates the results in Blanchard and Quah (1989).
- Figure A-2 shows the IRFs and FEVDs of output and hours in an alternative to the baseline model where the employment-to-population ratio is replaced by a series of per capita hours worked.
- Figure A-3 shows the IRFs of output to permanent demand and supply shocks as well as the FEVD of output in an alternative to the baseline model where sign restrictions are imposed at horizon 4.
- Figure A-4 presents the IRFs to the two transitory shocks in the baseline model.
- Figures A-5 and A-6 shows the IRFs and FEVDs for the model estimated using business sector output and hours.
- Figure A-7 shows the IRFs of output and employment and the FEVD of output in two alternatives to the baseline model: 1) a model in which the data is quadratically detrended (first line); 2) a model in which the data is detrended with a one-sided HP filter (second line).
- Figure A-8 shows the IRFs and FEVDs of output and employment in an alternative to the baseline model where the employment-to-population ratio for the overall population is replaced by the prime-age male employment-to-population ratio.
- Figure A-9 shows the FEVD of output and unemployment in two models with unemployment. The first model swaps employment for unemployment (left column). The second model swaps investment for unemployment (right column).
- Figure A-10 shows the results of the Monte Carlo exercise in a case with small-sample uncertainty. We use the theoretical model to generate 50 samples of the same length as our baseline sample (284 observations with the first 136 observations used to set the dummy initial observation in the sum-of-coefficients prior). For each sample, we estimate the SVAR and keep 1,000 draws that satisfy our sign and zero restrictions. This leaves us with 50,000 posterior draws. Figure A-10 reports the median and 16th-84th confidence intervals.

Figure A-1: Replication of Blanchard and Quah

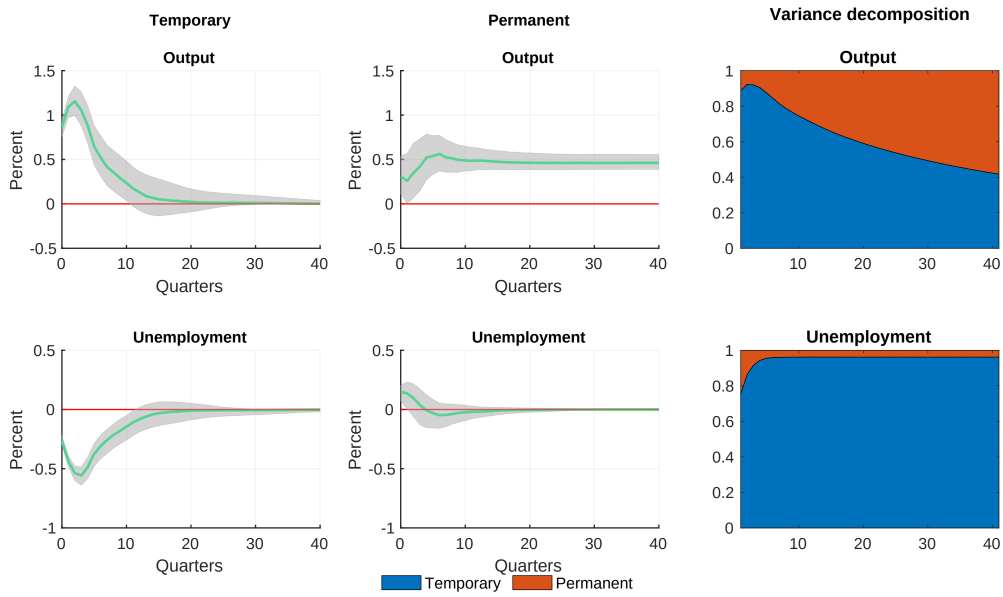

Figure A-2: IRFs and FEVDs of output and hours in the model with hours

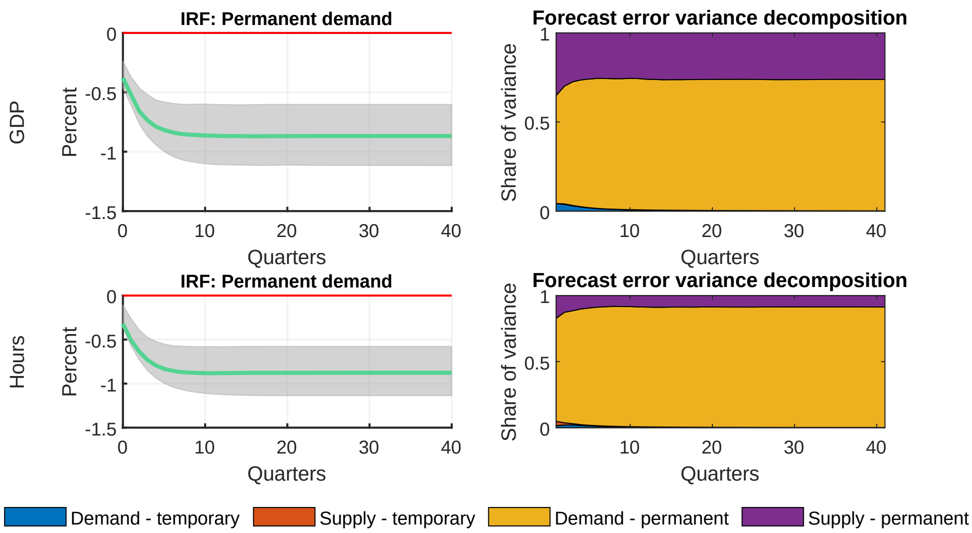

Figure A-3: IRFs and FEVD of output in the model with sign restrictions imposed at horizon 4

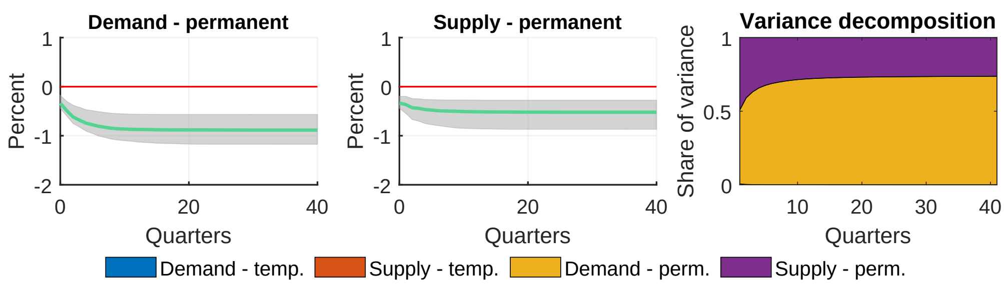

Figure A-4: IRFs to transitory demand and supply shocks

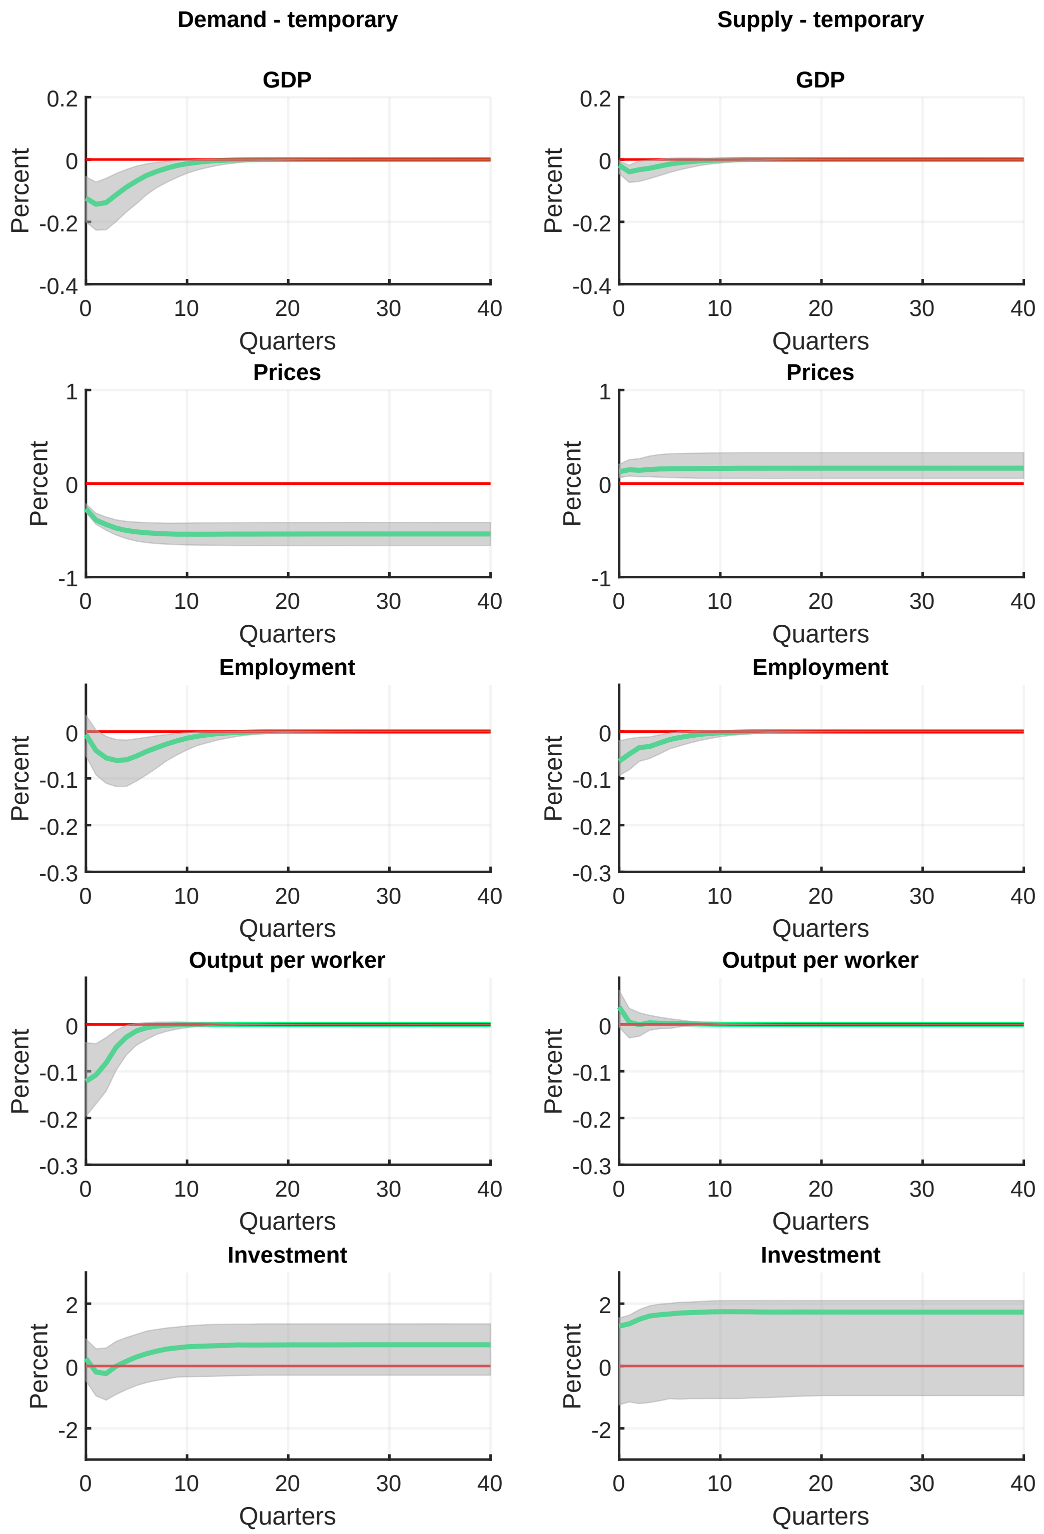

Figure A-5: IRFs to the permanent demand and supply shock

(Using business sector output and hours)

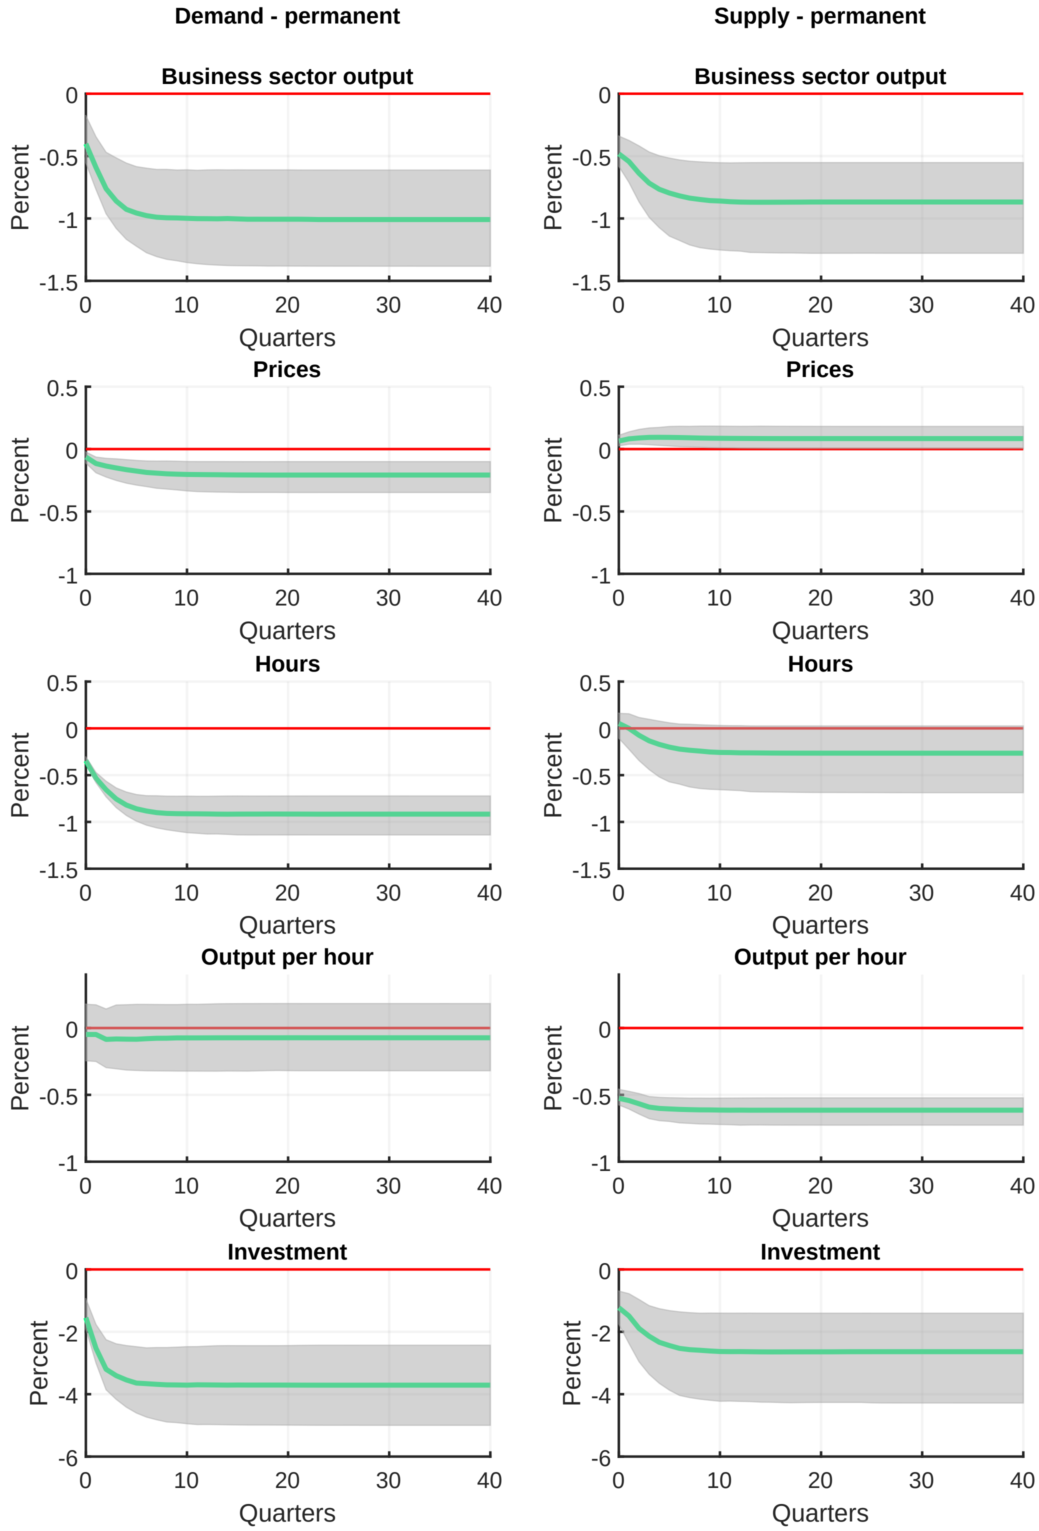

Figure A-6: Forecast error variance decomposition

(Using business sector output and hours)

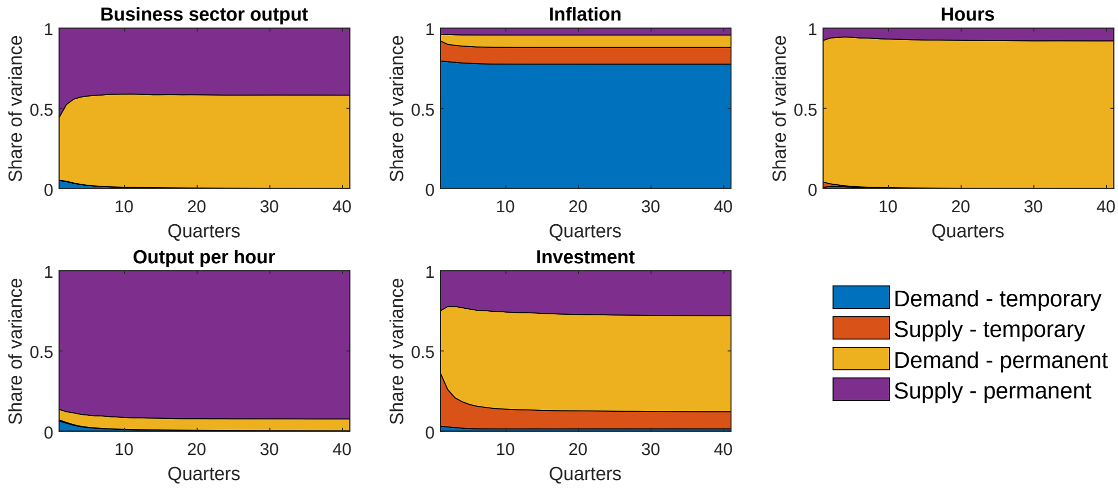

Figure A-7: IRFs and FEVDs for alternative detrending methods

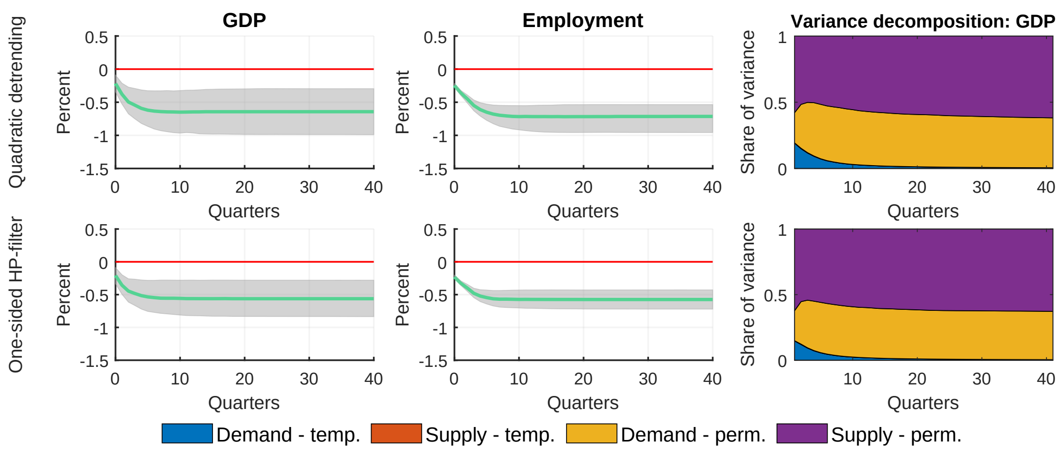

Figure A-8: IRFs and FEVDs of output and employment in the model with the prime-age male employment to population ratio

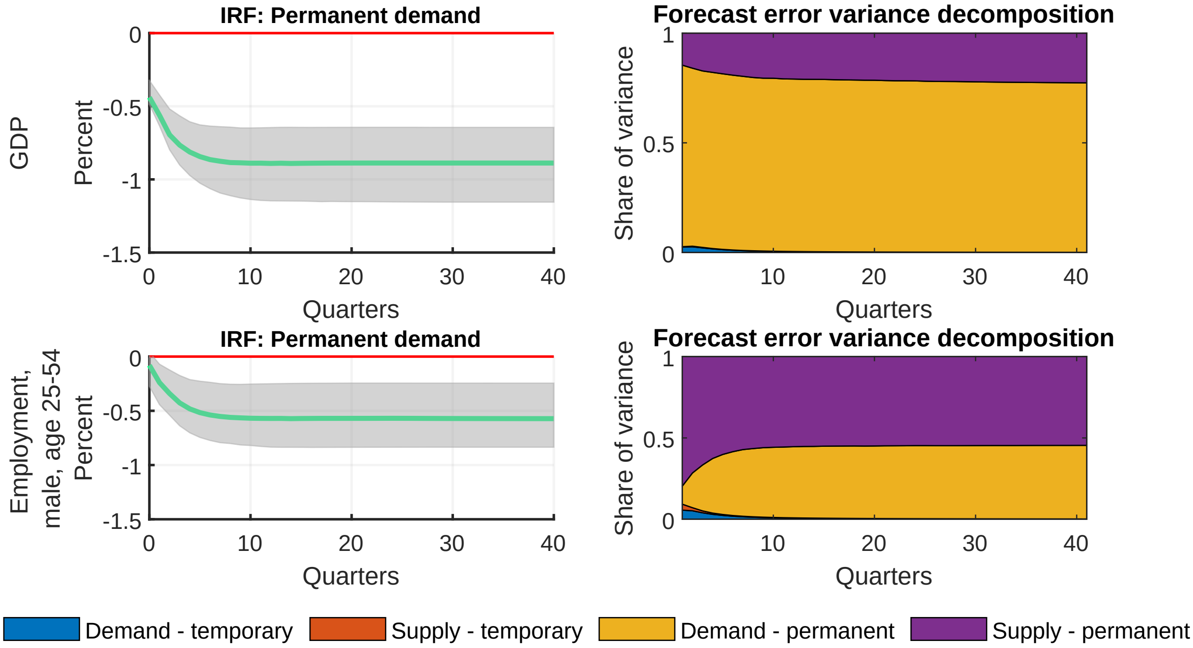

Figure A-9: FEVDs of output and unemployment in the models with unemployment

Figure A-10: Monte Carlo exercise with short samples: Output IRFs

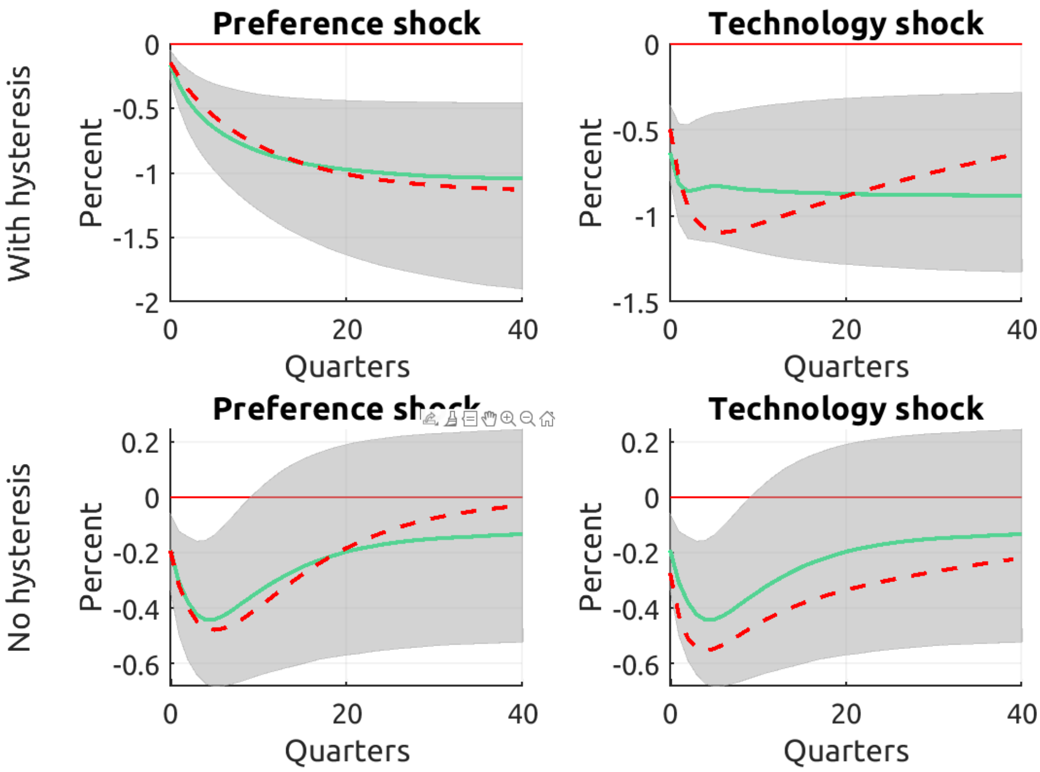

Note: Dashed red lines are the IRFs from the theoretical model. Solid green lines are the median IRFs from the SVAR. Shaded areas indicate the 16th and 84th percentiles. The first line reports IRFs of output to preference shock and technology shocks in the case *with* hysteresis. The second line reports IRFs of output to preference shocks and technology shocks in the case *without* hysteresis.