‘Missing disinflation’ and ‘missing inflation’ since the Great Recession: a VAR perspective*

Elena Bobeica† Marek Jarociński‡

This version: June 21, 2016

Abstract

In light of past data, inflation was puzzling neither during the ‘missing disinflation’ in the US and the euro area, nor during the euro area ‘missing inflation’. However, global and domestic factors played very different roles in these episodes. Economists and models that excessively focused on the global nature of inflation missed the contribution of deflationary domestic shocks and systematically over-predicted it during the euro area ‘missing inflation’ episode. These and other lessons follow from a suite of reduced-form and structural Bayesian vector autoregressions.

JEL Classification: E31, E32, F44

Keywords: Inflation Dynamics, International Transmission of Shocks, Bayesian Vector Autoregression, Conditional Forecast, Shock Identification

*This paper is part of the work of the Low Inflation Task Force of the ECB and the Eurosystem. We thank Fabio Canova, Matteo Ciccarelli, Luca Dedola, Thorsten Drautzburg, Giorgio Primiceri, Chiara Osbat, Mathias Trabandt and the members of the Eurosystem’s Workgroup on Econometric Modelling for their comments. The opinions in this paper are those of the authors and do not necessarily reflect the views of the European Central Bank and the Eurosystem.

†European Central Bank DG-Economics, email: elena.bobeica@ecb.int

‡European Central Bank DG-Research, email: marek.jarocinski@ecb.int

Non-technical summary

This paper answers two questions related to inflation developments since the start of the Great Recession. First, was inflation unusual during this time, given domestic and global developments? Second, what was the relative importance of global and domestic shocks? The period is particularly interesting as inflation in advanced economies has systematically surprised professional forecasters, raising questions about the reliability of existing models used for understanding inflation. In the US, in the aftermath of the financial crisis, inflation was expected to fall by more than it did, judging by the predictions of standard models such us reduced-form Phillips curves (see e.g. Ball and Mazumder 2011). This so-called 'missing disinflation' puzzle also characterizes other advanced economies where inflation proved to be rather resilient given the weakness in economic activity (see Friedrich 2016). Since 2012, particularly in the euro area, inflation has surprised in the opposite direction, falling to very low levels and giving birth to what we call the 'missing inflation' puzzle (see Constancio 2015). We analyze these two episodes, 'missing disinflation' and 'missing inflation', and compare the case of the US with that of the euro area.

We use vector autoregressive (VAR) models including both global and domestic variables as the tools to interpret the data. In a first step, we study the performance of out-of-sample conditional forecasts of inflation in a medium-scale VAR. We find that this flexible and data-rich model is able to account for inflation dynamics both in the US and in the euro area and, from the point of view of this model, no 'missing disinflation' puzzle exists. In fact, the observed fall in inflation in the first phase of the Great Recession appears to be large, not small, judging by the link between inflation and real activity prevailing over the past two decades. Global variables have a high explanatory power for the inflation swings recorded in this period. In the same econometric set-up, zooming in the post-2012 period for the euro area, we find again that the VAR is able to match inflation developments and we do not find a 'missing inflation' puzzle in the data. This time, the forecast based on domestic real variables matches the inflation path quite well, whereas global variables have a lower explanatory power.

These results are consistent with the conclusions we obtain based on smaller-scale structural VAR models, which disentangle the various shocks behind inflation dynamics in the Great Recession. In order to ensure robustness of results we employ three different identification schemes that use either restrictions on the timing of effects, or on the co-movement between global and domestic variables. Then we look at historical decompositions of inflation. We find that global shocks contributed positively to inflation during the US 'missing disinflation' episode. Domestic shocks have driven inflation down during the euro area 'missing inflation' episode.

Our results have important implications for modelling and policy analysis. We show that a sufficiently flexible and data-rich model allowing for both global and domestic factors matches well inflation developments since the Great Recession. The 'puzzles' seem to be an artefact of too small or too restricted models. We confirm the importance of the global nature of inflation (see e.g. Ciccarelli and Mojon 2010), but we also isolate an episode, the euro area 'missing inflation' period, where the link between domestic real activity and inflation appears crucial for explaining inflation. The slope of the Phillips curve implicit in our VAR appears to remain constant during and after the Great Recession. We find a remarkable similarity between the inflation forecast errors and the contribution of domestic shocks to inflation. Attuned to the global nature of inflation, professional forecasters may have missed the importance of domestic economic developments for euro area inflation.

1 Introduction

The dynamics of inflation in the Great Recession has puzzled economists. First, a 'missing disinflation' puzzle emerged when inflation in the advanced economies failed to fall as much as expected given the depth of the recession (see e.g. Hall 2011 on the US, and Friedrich 2016 on the rest of the advanced economies). Then came the 'missing inflation' puzzle, specific to the euro area, where inflation was unexpectedly low after 2012 (see e.g. Constancio 2015). These puzzles have led economists to question and reassess the relation between real activity and inflation, and reconsider the global nature of inflation. They did this using either small reduced-form models (e.g. Ball and Mazumder, 2011; Coibion and Gorodnichenko, 2015) or structural Dynamic Stochastic General Equilibrium (DSGE) models (e.g. Del Negro et al., 2015; Christiano et al., 2015; Gilchrist et al., 2015). In this paper we investigate these puzzles using reduced-form and structural Bayesian vector autoregressions (VARs) as the tools to interpret the data.

We use VARs including both global and domestic variables to answer two questions. First, was inflation unusual during the 'missing disinflation' and 'missing inflation' episodes, given domestic and global developments? Second, what were the roles of global and domestic shocks in these episodes? We answer these questions using out-of-sample conditional forecasts and in-sample historical decompositions, and draw lessons for economic modelling. We investigate both the euro area and the US. Our findings are as follows.

In light of a VAR estimated on the pre-crisis sample there was no 'missing disinflation' puzzle. Conditionally on the actual evolution of domestic real activity in the Great Recession our VARs predict a modest fall in inflation both in the euro area and in the US. The 'missing disinflation' puzzle appears in reduced form Phillips-curve type regressions (see e.g. Ball and Mazumder, 2011; Coibion and Gorodnichenko, 2015) and some tightly restricted DSGE models, but not in a VAR, a data-richer and less restricted model. The VAR has the opposite problem and predicts too little disinflation conditioning on domestic real activity. However, the forecast becomes very good when conditioned on the evolution of global variables instead, suggesting that global developments were important for inflation in this period.

On the second identified puzzle, the 'missing inflation' specific to the euro zone, a euro area VAR suggests that the disinflation was in line with the experience of the preceding two decades. Conditionally on the evolution of domestic real activity, the VAR correctly predicts a fall in inflation. This time, conditioning on global variables yields poor conditional forecasts of inflation. Hence, the 'missing inflation' is a puzzle only from the point of view of models that overestimate the global nature of inflation.

Historical decompositions also suggest that the two European disinflation episodes are of different nature: the first episode is driven more by global shocks, and the second one more by domestic shocks. In the US, global shocks contribute positively to inflation in the 'missing disinflation' episode. These lessons consistently follow from three alternative identification schemes, using either restrictions on the timing, or on the comovement of global and domestic variables to separate global from domestic shocks.

A striking picture emerges when we compare the errors of professional forecasters with the contributions of different shocks to inflation in the euro area: in the 'missing inflation' period the forecast errors and the contributions of domestic shocks look almost identical. A plausible story of 'missing inflation' is that economists, attuned to the global nature of inflation, missed the effect of domestic shocks.

One practical lesson is that models of inflation that are to be successful in the longer run need to incorporate both domestic and global factors. In the period since the Great Recession that we study, domestic factors are crucial in the 'missing inflation' episode, while global factors are crucial both before and after this episode.

We also find that the relation between real activity and inflation implied by our medium-sized VAR appears essentially unchanged during the Great Recession, which alleviates concerns about the time variation in the slope of the Phillips curve. It also justifies the use of a medium size VAR with constant coefficients, unlike in smaller models where time-variation in the coefficients appears important.

Our paper is related to a number of other works. A growing literature argues that inflation is largely a global phenomenon, as the globalisation process rendered inflation more sensitive to economic developments abroad (see e.g. Ciccarelli and Mojon, 2010; Mumtaz and Surico, 2012; Ferroni and Mojon, 2015). Another line of research argues that the Phillips curve has flattened and hence downplay the connection of inflation with domestic real activity fluctuations (see e.g. Simon et al., 2013; Blanchard et al., 2015). While we do not question these findings in general, we put them in a perspective by isolating a recent episode, the 'missing inflation' period in the euro area, when domestic real activity was crucial for explaining inflation dynamics.

Two closely related recent papers, Conti et al. (2015) and Ferroni and Mojon (2015), perform historical decompositions of euro area inflation into domestic and global shocks and also find a role for global shocks throughout the crisis period, and for domestic shocks in the 'missing inflation' period. We confirm the robustness of their conclusions in a structural VAR similar to Conti et al. (2015), with three identification schemes that use a similar logic but different detailed assumptions (our sign and zero restrictions follow Corsetti et al. 2014 and Baumeister and Benati 2013).

Our analysis leads to a similar conclusion as that of the DSGE papers that try to explain the 'missing disinflation' puzzle: the puzzle disappears in models that realistically account for the dynamics of a richer set of variables. In Del Negro et al. (2015) forward looking agents expect the central bank to prevent future marginal cost declines. In Christiano et al. (2015) interest rate spreads increase costs via the working capital channel, and, in addition, productivity falls. In Gilchrist et al. (2015) firms optimally raise prices when facing liquidity problems. In Bianchi and Melosi (2016) agents factor in an increasing probability of a switch to the fiscally led policy mix. All these complementary mechanisms imply that inflation does not fall in the Great Recession as much as a univariate regression of inflation on output gap would suggest.

Our results are also related to the explanation of the 'missing disinflation' proposed by Coibion and Gorodnichenko (2015). They argue that the growth of the price of oil between 2009 and 2011 affected household inflation expectations and, through them, inflation itself. We find that global shocks contributed positively to inflation in this period, and that conditional forecasts of inflation in this period are very good when conditioning on global variables, although the price of oil is not the most important of them.

There are two other papers that study VAR forecasts conditional on the actual developments in the Great Recession. Aastveit et al. (2014) use relatively small VARs and find that these conditional forecasts have large errors. Consequently, they advocate the use of time-varying models to account for the Great Recession. We also find that conditional forecasts in small VARs have large errors, but we show that these forecasts become very good when the VAR is sufficiently large. This finding is about out-of-sample conditional forecasts, which could in principle become more affected by estimation error in larger models, so it is not a forgone conclusion. Our finding is consistent with Bańbura et al. (2015), who also find, using a large VAR, that the conditional forecasts of the main macroeconomic variables match the actual outcomes well in the Great Recession, with the exception of financial variables.

In the rest of this paper, section 2 investigates the presence of the 'missing disinflation' and 'missing inflation' puzzles in the data using conditional forecasts, section 3 identifies structural shocks and studies their importance for inflation in the two episodes, and section 4 concludes.

2 Conditional forecasts of inflation

In this section we ask whether inflation dynamics was indeed unusual during the 'missing disinflation' and 'missing inflation' episodes, from the point of view of a Bayesian VAR estimated on the data preceding these episodes. We condition on the actual realizations of either domestic or global variables and compare the resulting conditional forecasts of inflation with the actual outcomes.

We use a medium-sized VAR, with 23 variables: a consumer price index, eight indicators of domestic real activity, eight global variables and six financial variables. When choosing the variables we try to focus on the most relevant ones according to the rank-

Table 1: VAR specification

	Variable	Transformation
Price index	€:Harmonized Index of Consumer Prices (HICP) /	log
	$:Consumer Price Index (CPI)	log
Domestic real activity variables	Real GDP	log
	Real consumption	log
	Real investment	log
	Total employment	log
	Unemployment rate	none
	Capacity utilization	none
	Consumer confidence	none
	Purchasing Managers' Index (PMI)	none
Global variables	Rest-of-the-world real GDP	log
	Price of oil	log
	Commodity prices	log
	Nominal Effective Exchange Rate (NEER)	log
	USD/EUR exchange rate	log
	Foreign* real GDP	log
	Foreign* consumer price index	log
	Foreign* short-term interest rate	none
Financial variables	Short-term interest rate (€:EONIA / $:Fed Funds rate)	none
	2-year government bond spread	none
	10-year government bond spread	none
	Mortgage bank lending spread	none
	Non-financial corporations bank lending spread	none
	Corporate bond spread	none
Prior hyper-parameters Sims and Zha (1998) notation	overall tightness $\lambda_1 = 0.05$	
	decay $\lambda_2 = 1$	
	'other weight' $\lambda_3 = 1$	
	standard deviation of the constant term $\lambda_4 = 2.5 \times 10^3$	
	weight on the no-cointegration (sum of coefficients) dummy observations $\mu_5 = 4$	
	weight on the one unit root dummy observation (initial dummy observation) $\mu_6 = 4$	

Notes: € denotes the euro area VAR, $ denotes the US VAR. *'Foreign' means 'US' in the euro area VAR and 'Euro area' in the US VAR. The spreads are calculated as the interest rate in question minus the short-term interest rate, except for the Corporate bond spread which is the average spread of a BBB 7 to 10-year corporate bond over government bond of the same maturity.

ing derived in Jarociński and Maćkowiak (2016), but with a few constraints: we use the same specification for the euro area and the US, we want to separate domestic from global variables (so we omit exports and imports which are difficult to classify), and for a-priori reasons we want the number of global variables to be equal to the number of domestic variables, so we end up including some global variables that do not appear relevant in Jarociński and Maćkowiak (2016). The variables are listed in the Table 1. The variables are observed at the quarterly frequency over the sample starting in 1990Q1 and the real variables, as well as the price indexes are seasonally adjusted (Appendix A contains details on the construction of some of the series and the data sources). As is the common practice in quarterly VARs, we include four lags of the endogenous variables. All the VARs include an intercept.

We estimate the VARs using the Bayesian approach in the tradition of Sims. More precisely, we include all variables in (log) levels and we use Bayesian priors of Litterman (1979) and Sims and Zha (1998). Since the VAR is relatively large, we follow the advice of Bańbura et al. (2010) and Giannone et al. (2015), and use tighter priors than is standard in smaller VAR. When setting the priors, we start with the prior hyperparameters of Sims and Zha (1998) and scale them by a factor of 4 (see Table 1). When the prior is too loose, the unconditional forecast implies exploding rate of inflation in the long run, so we tighten the prior until the unconditional forecast shows that inflation stabilizes in the long run. It turns out that we need a factor of 4 to achieve this. When computing the conditional forecasts we use the Gibbs sampler of Waggoner and Zha (1999) (see also Jarociński 2010 for some implementation details).

We start with attempts to replicate with a VAR the 'missing disinflation' puzzle in the US. We estimate the VAR using the data up to the peak before the Great Recession, dated by the NBER at 2007Q4. Then we forecast consumer prices conditionally on the actual path of the real activity variables in the Great Recession and afterwards. We plot the forecast with the 68% posterior uncertainty band, along with the actual data, in the first panel of Figure 1. The VAR is specified in terms of the level of the price index, but for the purpose of reporting we transform the levels into year-on-year growth rates.

The two vertical lines mark the beginning of 2009 and the end of 2011, delimiting the period of the 'missing disinflation' puzzle as defined e.g. in Coibion and Gorodnichenko (2015).

We can see that in a VAR the puzzle is not apparent: based on the joint dynamics of all the variables prior to the Great Recession and the real activity in the Great Recession, the VAR predicts a similar inflation as that actually observed, and hence the observed inflation does not appear puzzling. The conditional forecast in the first panel stands in contrast with the predictions of reduced-form Phillips curves, which, as shown e.g. in Coibion and Gorodnichenko (2015) tend to be much lower. For example, the green line with circles (labeled CG Phillips curve prediction) shows the prediction of their baseline Phillips curve specification, which is indeed much lower.[1] Why does the VAR predict a higher inflation than reduced-form Phillips curves?

Part of the reason lies in the inertia that the VAR captures. In particular, inflation was temporarily high just before the crisis, and this raises the whole forecast path. If, instead of starting the conditional forecast at the NBER-dated peak of real activity before the recession, we start in 2008Q3, when inflation peaked, we predict even higher inflation (see panel 2). By contrast, if we start in 2007Q1, before inflation picked up, the conditional forecast of inflation is lower and comes closer to replicating the 'missing disinflation' puzzle, but the puzzle is still far from obvious (see panel 3).

When instead of conditioning on real activity we condition on global variables the forecast becomes almost perfect irrespective of when we start the forecast (panels 4-6). After the first three plots one might have suspected that the conditional forecasts from this VAR are not very responsive to the conditioning assumptions, e.g. because the Bayesian prior is too restrictive and suppresses interdependencies between variables. However, panels 4-6 show that this is not the case.

In the next two panels we bring financial variables into the picture. The structural accounts of the Great Recession in Del Negro et al. (2015) and Christiano et al. (2015)

[1] We have replicated their Figure 1, Panel B, which is in terms of annualized quarterly inflation, and plotted it here in terms of annual inflation.

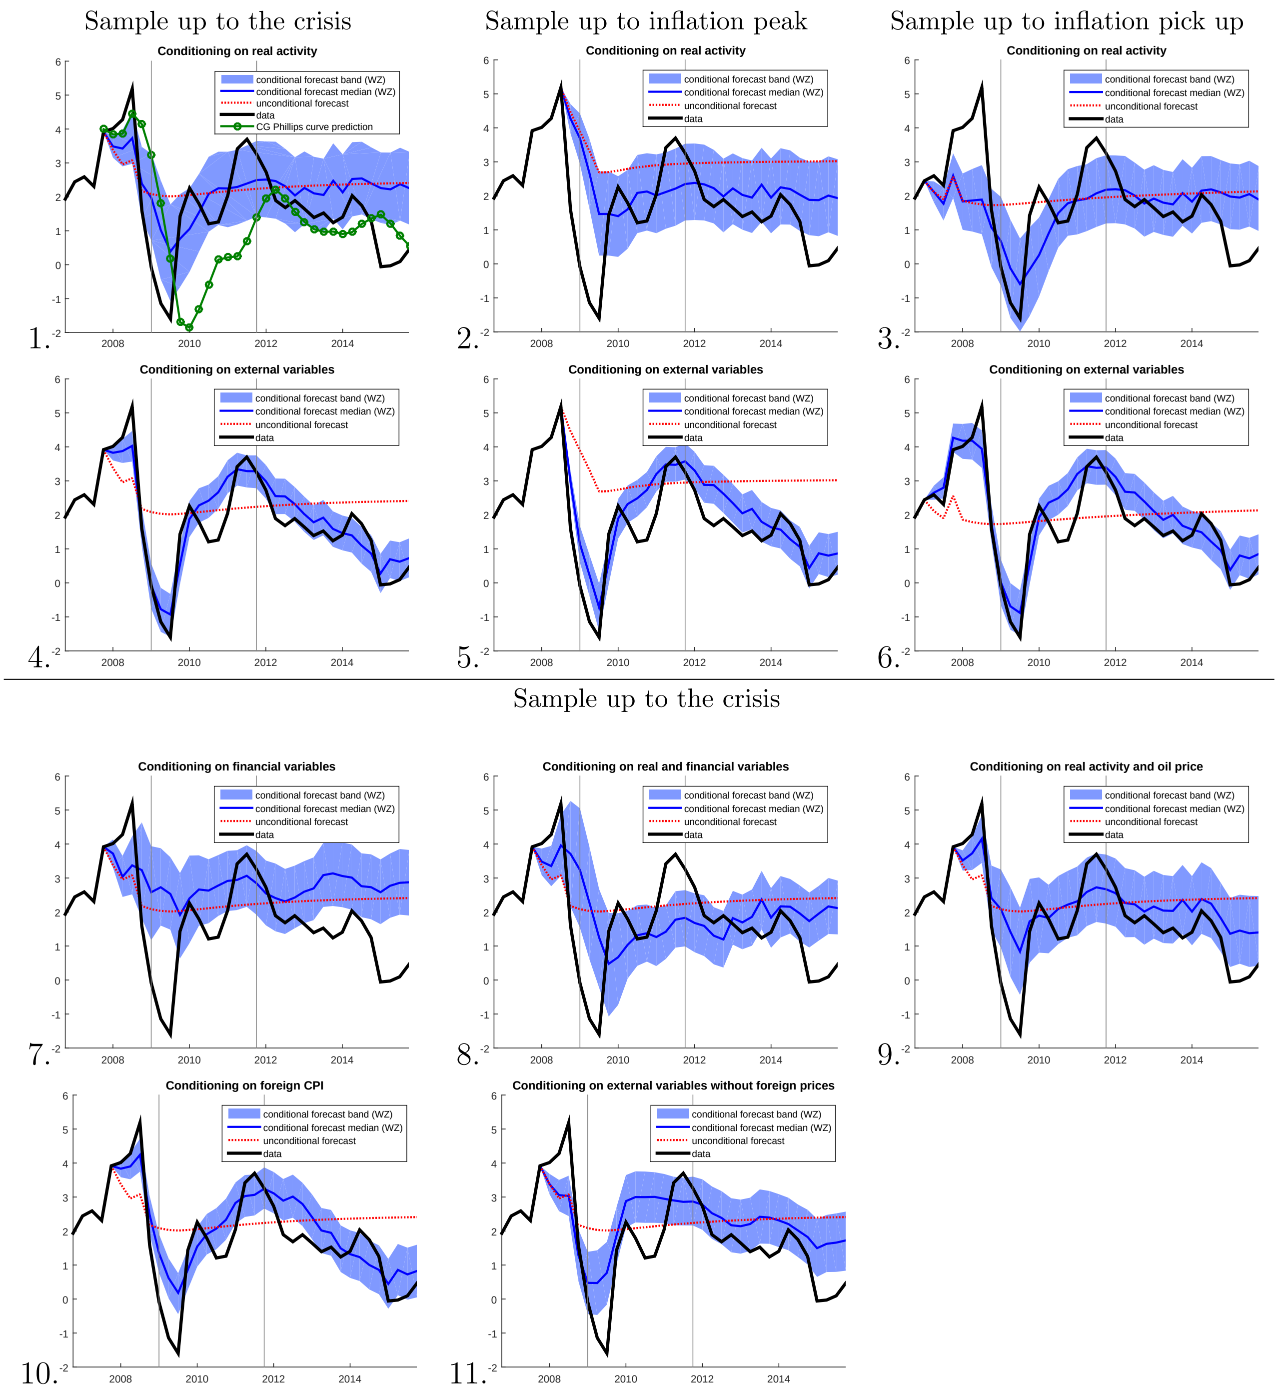

Figure 1: Conditional forecasts of US CPI (year-on-year growth)

suggest that interest rates and spreads, or their interaction with real activity, are important for inflation. However, the forecasts conditional on financial variables (panel 7) and on both real activity and financial variables (panel 8) do not match the actual inflation particularly well. This does not necessarily contradict the mentioned structural accounts of the Great Recession, because the Sims and Zha prior we use might be pushing the VAR parameters away from the cross-equation restrictions of these models. But the bottom line is that our VAR does not provide a 'smoking gun' evidence in favor of the structural accounts that rely on financial frictions.

Coibion and Gorodnichenko (2015) argue, instead, that the price of oil is crucial for understanding inflation in the Great Recession, so in panel 9 we condition on real activity and the price of oil. In this case we forecast a higher inflation than when conditioning on real activity alone. This improves the match of the forecast somewhat in 2011, but worsens it in 2009, so the VAR fails to produce a 'smoking gun' evidence in favor of the role of oil prices too.

The strongest message from the VAR is on the importance of the global nature of inflation. When we condition the forecast on only one variable, the euro area consumer prices, the resulting forecast becomes almost as good as in panel 4 (see panel 10). When we condition on the global variables excluding euro area consumer prices, the fit of the forecast is good but worsens compared with panel 4 (see panel 11).

We conclude that the VAR estimated for the US up to the Great Recession does not find the 'missing disinflation' puzzling in light of the observed real activity developments. The conditional forecast are, however, particularly close to the actual data when we condition on global variables. This confirms the global nature of inflation.

We now turn to the 'missing disinflation' and 'missing inflation' puzzles in the euro area. Figure 2 presents the conditional forecasts of euro area inflation generated with the euro area VAR. The first column focuses on the 'missing disinflation' puzzle, while the second and third columns focus on the 'missing inflation' puzzle. We start with the 'missing disinflation'.

Panel 1 shows that there is no 'missing disinflation' in the data according to the euro area VAR either. We estimate the VAR on the data up to the peak before the first recession, dated by the CEPR at 2008Q1 (and marked by a vertical line in the plot). We forecast inflation conditional on the actual domestic real activity during the recession

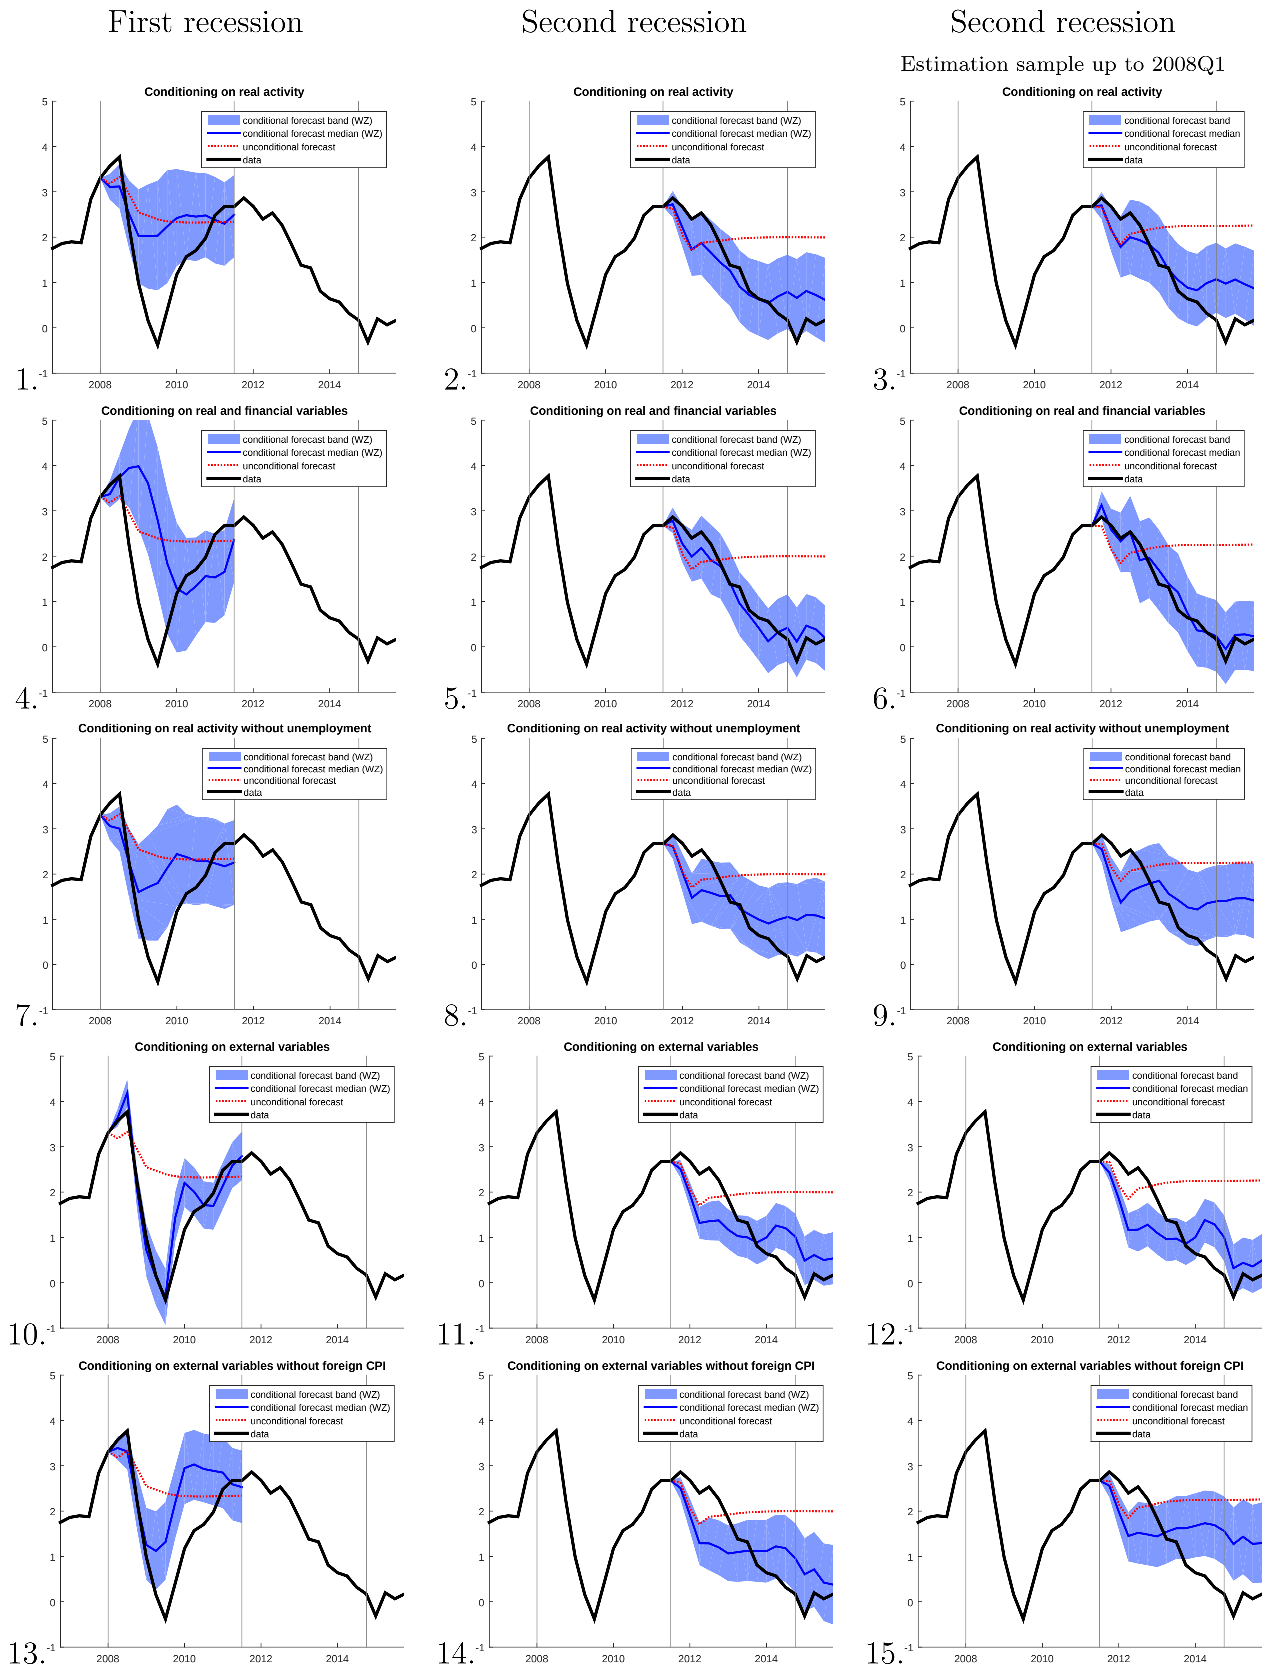

Figure 2: Conditional forecasts of euro area HICP (year-on-year growth)

and afterwards. We focus first on the period up to the peak before the second recession, dated by the CEPR at 2011Q3 (and marked by the second vertical line in the plot). Based on the the sample 1990Q1-2008Q1, a researcher who learns about the extent of the first euro area recession expects inflation to fall only by about 0.5 percentage point compared with his unconditional forecast. This is much less than the 4 percentage points fall actually observed. As in the US, adding the financial variables to the conditioning set yields a lower inflation forecast after some delay (panel 4).

Inflation again appears to be a global phenomenon: when in the same VAR we condition on the actual path of the external variables the forecast matches the actual outcomes very well (panel 10). Similarly to the US case, it turns out that the inflation of another major economy (here - US) is the crucial variable among the global ones. When we exclude the US inflation from the set of the conditions, the forecast conditional on the remaining variables gets farther from the actual inflation (panel 13). However, this forecast is still much better than the forecast conditional on real activity. We conclude that the VAR suggests a global nature of inflation in the first disinflation, and rather limited role for domestic real activity (and financial frictions).

The 'missing inflation' period is very different. We report the evidence for this in the second column in Figure 2. For the purpose of this paper we define the 'missing inflation' period as 2011Q3-2014Q4 (the end of this period is marked with the third vertical line). When we reestimate the VAR on the data up to 2011Q3, before the start of the second recession according to CEPR, the forecast based on the actual path of real activity variables in the 'missing inflation' period matches the outcomes very well (panel 2). The VAR predicts a steady fall in inflation, similar to what was observed, and hence suggests no 'missing inflation' puzzle in the actual data. The conditional forecast improves further after incorporating financial variables in the conditioning set (panel 5). Unemployment rate, which reached record high levels in the second recession, turns out to be the single most important variable among real activity indicators. When we omit the unemployment rate, the match between the conditional forecast and inflation becomes clearly worse (panel 8).

The forecast conditional on the external variables matches the outcomes poorly (panel 11). This forecast is too low initially, but then stays above 1% at the time when the actual inflation was falling gradually but steadily. When we exclude the US consumer prices from the conditioning set, the forecast becomes even poorer.

At the end of the sample, after 2014Q4, external variables again appear to be relevant for inflation. Recall that in his period the price of oil dropped substantially and the Chinese economy slowed down. These developments can be seen in the shape of the conditional forecasts based on external variables - they turn sharply downwards.

Comparing plots 1 and 2 one might suspect that the slope of the Phillips curve changed in the Great Recession. According to this interpretation, a VAR estimated on the pre-crisis sample would imply a weak relation between real activity and inflation, i.e. a flat Phillips curve. Then, after incorporating the data from the first recession, when inflation fell sharply, the VAR parameters would change to imply a steeper Phillips curve. However, the plots in the third column suggest that this is not the case, and there is no evidence of a change in the slope of the Phillips curve implied by the VAR. In the third column we conditionally forecast inflation starting in 2011Q1, but using the posterior of the VAR parameters based on the data up to 2008Q1.[2] Comparing panel 3 with panel 2 we see that the disinflation conditional on the second recession is roughly the same in the VAR estimated up to 2008Q1 and in the VAR estimated up to 2011Q3, suggesting that the relation between real activity and inflation remained roughly constant across the two VARs. The similarity of panels 5 and 6 suggests that the relation between real and financial variables, and inflation, also remained roughly constant.

What might have changed somewhat, although the evidence is weak, is the link between euro area inflation and global variables other than US inflation. If we focus on how much the conditional forecast deviates from the unconditional forecast (plotted with the dotted line), we can see that this deviation is smaller in the VAR estimated up to 2008Q1 (panel 15), and larger in the VAR estimated up to 2011Q3 (panel 14), which

[2] To keep the experiment clean, this time we sample directly from this posterior, without using the Gibbs sampler of Waggoner and Zha (1999) that would update it with the information contained in the assumed scenario.

might suggest that the link between inflation and global variables has become stronger. Such finding is to be expected if the data for 2008-2011 contain a disinflation driven more than usually by global variables.

Figure 3 presents the conditional forecasts based on the VAR estimated up to 2008Q1 and extended to the end of 2015. This figure is a concise illustration of the point that the first disinflation appears to be driven more by global variables (the two bottom plots fit the first disinflation better), and the second more by domestic variables (the two top plots fit the second disinflation better).

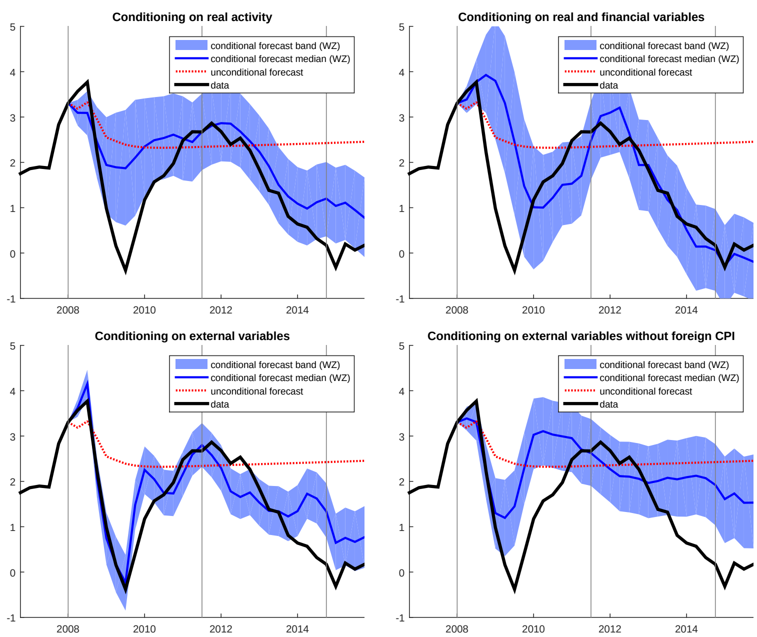

Figure 3: Conditional forecast of euro area HICP (year-on-year growth) extended to the end of 2015, VAR estimated on the sample up to 2008Q1.

To summarize, in light of VAR conditional forecasts, the ‘missing disinflation’ and ‘missing inflation’ episodes are very different from each other, but none of them appears to be puzzling. The first disinflation was not small, but rather large judging by the previous experience, but it can be explained by the dynamics of the global variables. The second disinflation was broadly in line with the past experience, and it can be explained by the weak domestic real activity. Adding the data on the Great Recession to the pre-crisis sample does not lead to a reassessment of the relation betweeen real activity and inflation. If anything, it leads to a somewhat stronger estimated relation between global variables and inflation.

The above discussion of the nature of the inflation fluctuations since the Great Recession, while suggestive, is reduced form and does not warrant conclusions on the types of shocks driving inflation in the two episodes. To identify the shocks we move on to structural VARs.

3 Structural approach: shock identification and historical decompositions

In this section we identify structural shocks and ask what their relative contributions are to the dynamics of prices since the start of the Great Recession. Since shock identification becomes more complex as the number of variables increases, in this section we limit the number of variables to ensure the tractability of the structural model. The variables we include are the price of oil, rest-of-the-world real GDP (or the share of domestic real GDP in the world real GDP), real GDP, consumer prices, short-term interest rate, 10-year bond spread, and the nominal effective exchange rate. In the priors we use the same hyperparameter values as Sims and Zha (1998). To study the robustness of the results we use three alternative identification schemes. Table 2 summarizes the identifying restrictions, which are all on the contemporaneous impulse responses.

I. Choleski identification

We start with the Choleski identification scheme. We order the variables so that the two global variables, the price of oil and the rest-of-the-world GDP, are first. We label the first two shocks as global, and the remaining ones as domestic. This identification relies on a timing restriction: global shocks affect all variables immediately, while domestic shocks affect global variables only with a delay. Studies of the global nature of inflation, such as Ciccarelli and Mojon (2010), use a similar logic. There is one more shock that we can label in this approach: the monetary policy shock. We place the short-term interest

Table 2: Sign and zero restrictions used to identify shocks in the VAR

Variable \ shock *I. Choleski*	Global shock	Global shock	Domestic shock	Domestic shock	Monetary policy	Spread shock	Exchange rate
Price of oil	+	0	0	0	0	0	0
Rest-of-the-world real GDP	•	+	0	0	0	0	0
Real GDP	•	•	+	0	0	0	0
Consumer prices	•	•	•	+	0	0	0
Short-term interest rate	•	•	•	•	+	0	0
Spread	•	•	•	•	•	+	0
Exchange rate	•	•	•	•	•	•	+

II. Corsetti et al. (2014) (CDL)	Oil supply	Global demand	Domestic demand	Domestic supply	Monetary policy	Spread shock	Exchange rate
Price of oil	+	+	•	• (≈0)	0	0	0
Share of world real GDP	•	−	+	+	0	0	0
Real GDP	−	+	+	+	0	0	0
Consumer prices	+	+	+	−	0	0	0
Short-term interest rate	0	•	•	•	+	0	0
Spread	•	•	•	•	•	+	0
Exchange rate	+ (•)	•	•	+	•	•	+

III. Corsetti et al. (2014) and Baumeister and Benati (2013) (CDL+BB)	Oil supply	Global demand	Domestic demand	Domestic supply	Monetary policy	Spread shock	Exchange rate
Price of oil	+	+	•	• (≈0)	•	•	0
Share of world real GDP	•	−	+	+	−	•	0
Real GDP	−	+	+	+	−	−	0
Consumer prices	+	+	+	−	−	−	0
Short-term interest rate	0	•	+	•	+	0	0
Spread	•	•	•	•	−	+	0
Exchange rate	+ (•)	•	•	+	+	+	+

Notes: • = unconstrained, + = positive sign, − = negative sign, 0 = zero restriction, ≈0 = magnitude restriction that centers the error band of the responses at zero. All restrictions are imposed on impact. The exchange rate is defined so that a + means an appreciation. In parentheses we show the restrictions used for the US VAR whenever they differed from the euro area VAR.

rate after real GDP and prices, but before the spread and the exchange rate, and the shock to the short-term interest rate is a recursively identified monetary policy shock (see Christiano et al. 1999 on the recursive identification of monetary policy shocks).

II. Corsetti et al. (2014) / Conti et al. (2015) identification

The second strategy relies on sign and zero restrictions that distinguish global shocks from domestic ones, and supply shocks from demand shocks. The key step is to distinguish domestic demand shocks from global demand shocks, and to achieve this end we follow Corsetti et al. (2014) (CDL). We first modify one variable: we replace the rest-of-the-world real GDP with the share of the domestic real GDP in the world real GDP. The identification focuses on the comovement of this share with the domestic real GDP. Consider a shock that increases real GDP. If the economy's share in world GDP also increases, this means that the shock had more effect on the domestic economy then on the rest of the world and we label the shock as domestic. By contrast, if the economy's share in world GDP falls, the shock had more effect on the rest of the world and we label the shock as global. CDL use the same logic (combined with some other restrictions) to identify productivity and global demand shocks affecting US manufacturing. A caveat is that one can imagine a shock that is global in nature but increases the demand for domestic products more than the demand for rest-of-the-world products. Such a shock would generate a positive comovement between domestic GDP and its world share. Hence, we would label this shock as domestic. This means that this identification might capture some foreign shocks, but we can reasonably hope that most of the shocks that we label domestic are indeed domestic.

Following this logic, we assume that a positive *global demand shock* increases the price of oil, consumer prices and real GDP, but reduces the domestic GDP share in the world. A positive *domestic demand shock* increases both real GDP and its world share, and consumer prices.

Next, we identify an *oil supply shock*. We assume that a negative oil supply shock that leads to an increase in oil prices has a negative impact on real activity and a positive one on inflation. This follows the lessons of the literature that identifies various types of oil related shocks by modelling the global crude oil market and then investigates the impact of the different shocks on the key macroeconomic variables (see for example Kilian 2009 and Baumeister and Peersman 2013 for the case of the US).

We assume that a positive *domestic supply shock* increases domestic output as well as its world share (following the CDL logic) and, in contrast to the demand shock, it reduces domestic prices. With only the above restrictions, oil supply shocks and domestic supply shocks are still not distinguished. To disentangle them, we require that the exchange rate appreciates after a positive domestic supply shock. Then we proceed differently in the US and in the euro area. In the euro area it turns out that the exchange rate appreciates also after a contractionary oil supply shock (we impose this as a restriction, but this restriction is hardly ever binding), and the two shocks are distinguished. In the US we don't restrict the response of the exchange rate to the oil supply shock, and instead impose a magnitude restriction that the response of the price of oil to a domestic supply shock is 'small' (in practice, less than 5% in magnitude). This distinguishes this shock from the oil supply shock.

III. Adding monetary policy identification of Baumeister and Benati (2013)

In the third identification scheme we keep the sign restrictions from the previous scheme and add further restrictions that identify monetary policy shocks in a different way. As we show, this modification changes the relative importance of domestic and global shocks.

Following Baumeister and Benati (2013) (BB) we assume that a contractionary *monetary policy shock* is an increase of the short term interest rate that has an immediate negative effect on output and prices, a negative effect on the bond spread and a positive effect on the exchange rate. Since bond yields do not respond one-to-one to the short rate, the spread shrinks after a short-term rate increase. Since the monetary policy shock is a special case of a domestic demand shock, domestic output and its world share move in the same direction, so the world share of output falls.

The reason we adopt the BB identification is that we think of it as representing an opposite extreme to the recursive identification. The recursive identification implies that monetary policy shocks have no contemporaneous effect on prices or output. It is well known that with this assumption researchers tend to find that monetary policy shocks contribute very little role to the economic fluctuations. By contrast, BB assume that monetary policy shocks affect prices and output immediately, with the expected sign, and in fact the policy effects turn out to be largest contemporaneously, and also in the long run contribute importantly to the economic fluctuations. So the recursive and the BB identification represent respectively the lower and the upper bound for the strength of the impact of monetary policy.

Following BB we also identify a *spread shock* that moves the spread, while leaving the short term interest rate unchanged. An contractionary spread shock increases the spread, appreciates the exchange rate, and reduces output and prices. The spread shock has been latent throughout the sample, but it is particularly useful to reflect nonstandard monetary policies when the short term interest rates are at their effective lower bound.

The *exchange rate shock* is common to all the approaches and moves the exchange rate but has no contemporaneous effect on all the other variables. We find that also in subsequent periods the shock has little effect on other variables. We need this additional shock because in a VAR the number of shocks is always equal to the number of variables, but we do not interpret this shock. The shock plays very little role in the dynamics of all the variables except exchange rate.

When implementing the above restrictions we randomly search for orthogonal matrices that rotate the Choleski factor and impose our sign restrictions while preserving the zero restrictions. Technical details on imposing sign and zero restrictions can be found in Arias et al. (2014). It takes only a few minutes to generate 1000 draws from the posterior of the structural VAR in the CDL identification, but this time extends to several hours for the CDL+BB identification.

In a closely related recent paper Conti et al. (2015) also study structural shocks affecting the 2012 disinflation in the euro area using similar sign restrictions. There are two main differences between our paper and theirs. First, our goal is to partition the total variability of inflation into the contributions of different shocks. By contrast, they are interested in the effects of a subset of shocks, and they control for a number of exogenous variables. Second, we differ in the details of the implementation, e.g. they impose sign restrictions on the response of world output and not on the euro area share in world output, and they do not impose any zero restrictions, while we found the zero restrictions to be helpful.

We report the impulse responses to all the shocks in the Appendix. They are mostly intuitive, except, perhaps, for the fact that the responses of the exchange rate are in many cases insignificant when not explicitly restricted.

In the rest of this section we focus on the historical decompositions in the Great Recession and around it.

Figure 4 reports historical decompositions of inflation (year-on-year change in consumer prices). We start the decomposition in 2006Q. The black line is the difference between the median unconditional forecast generated by our VAR and the actual series. Our VAR represents this difference as the effect of the shocks that we identified and bars show the contribution of each shock. For greater readability, we aggregate the contributions of the two global shocks (dark blue bars), the contributions of the domestic supply and demand shocks (light blue bars), and the contributions of monetary policy and spread shocks (yellow). The contributions of the exchange rate shock are very small and barely visible (dark red). The vertical lines represent the same dates as before: in the euro area the dates are 2008Q1 (the peak before the first CEPR recession), 2011Q3 (the peak before the second CEPR recession) and 2014Q4, and in the US the dates are 2009Q1 and 2011Q4 (the notional beginning and the end of the 'missing disinflation').

Let us focus first on the euro area 'missing disinflation' period. The dark blue bars in Figure 4, representing global shocks, account for most of the disinflation in each of the three identification schemes.

In the 'missing inflation' period the picture is different. We see that the dominant factor behind the disinflation is the switch of the contribution of domestic shocks from positive to negative. Global shocks play some role too, but their contributions are small.

Finally, the picture changes again around the end of 2014: the negative contributions of global shocks become important again.

Figure 4 outlines also a range for the possible contributions of monetary policy shocks. Recursively identified monetary policy shocks contribute very little to inflation dynamics.

The monetary policy shocks identified as in Baumeister and Benati (2013) imply that an excessively tight monetary policy (in the period when the zero lower bound for the interest rate was binding) contributed up to 0.5 of a percentage point to the disinflation. The contribution of monetary policy becomes slightly positive again at the end of the sample and the detailed results (reported in Appendix C) show that this is due to the spread shock, likely reflecting the quantitative easing that the ECB implemented in 2015.

To sum up, according to this analysis, both global and domestic shocks have driven euro area inflation in the crisis. In the 'missing disinflation' period global shocks have played a larger role, while in the 'missing inflation' period the role of domestic shocks was larger. These results put some perspective on the literature arguing in favour of inflation in advanced countries being largely a global phenomenon. Ciccarelli and Mojon (2010), using data up to the international financial crisis, find that global inflation explains more than two thirds of the variance of national inflation rates and that it also acts as an attractor for national inflation. Ferroni and Mojon (2015) revisit the relevance of global factors in the aftermath of the crisis and find that they still explain a dominant part of national inflation, albeit to a diminished extent. We confirm that global factors account for much of inflation dynamics, but we find that the 'missing inflation' episode in the euro area is mainly driven by domestic shocks.

Figure 5 compares the errors made by professional forecasters and contributions of selected types of shocks. The middle panel shows the errors of 1-year ahead forecasts reported in the Survey of Professional Forecasters (SPF), defined as the outcome minus the forecast. Forecasts tend to be clustered around similar values, and the errors of the official ECB forecasts, Eurozone Barometer and Consensus Economics are very similar. In the top and bottom panels we plot the contributions of, respectively, global and domestic shocks (including monetary policy), from the CDL+BB identification.

One-year ahead forecast errors and shock contributions are, of course, different objects, but they do carry some similarities. The contributions contain the effects of new shocks, that by definition could not have been anticipated a year earlier, and the propagation of all the past shocks. Forecast errors contain the effect of the same new

Euro area HICP | United States CPI

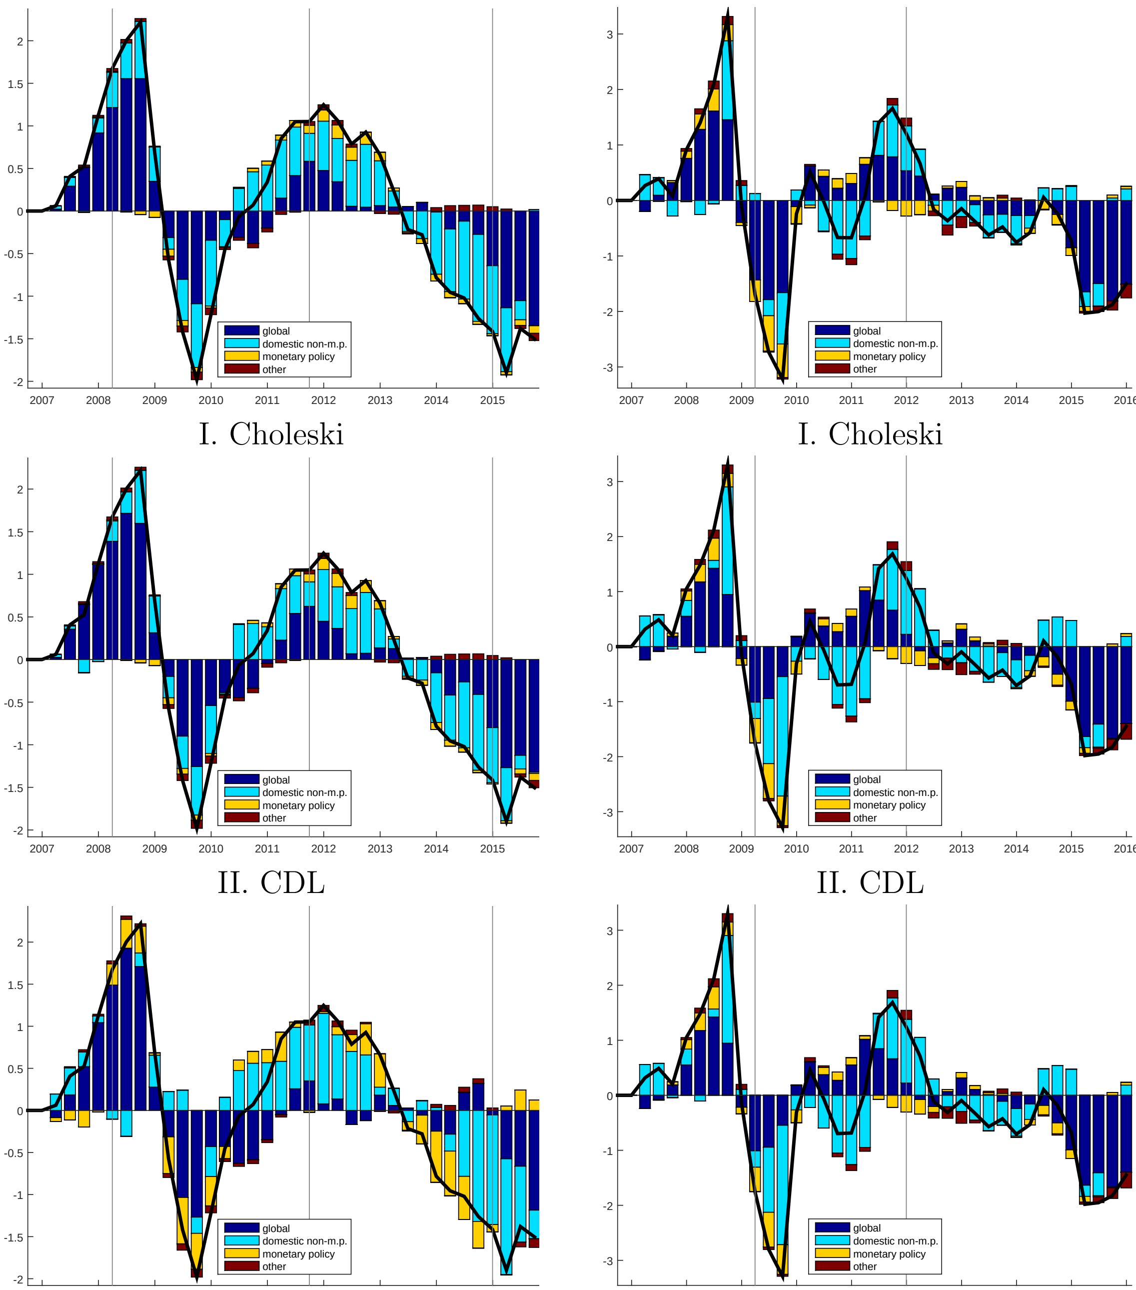

Note: The black line is the deviation of year-on-year inflation from the unconditional forecast, the bars show the contributions of different types of shocks to this deviation.

Figure 4: Historical decompositions of inflation in the euro area and in the US

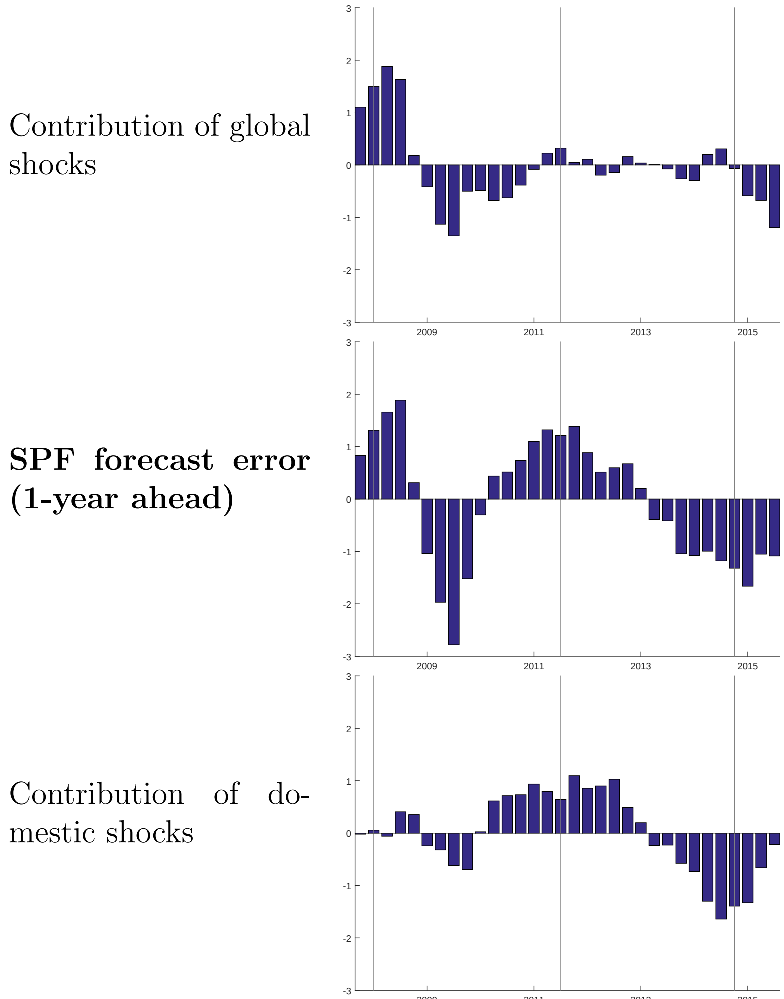

Figure 5: Errors of the SPF forecasts and the contributions of shocks

shocks and the part of the propagation of all the past shocks that is misspecified in the forecaster's model.

With this interpretation in mind, the comparison of the middle and the bottom panel is striking: in the 'missing inflation' period the forecast errors are very similar to the contributions of domestic shocks. This suggests a particular interpretation of the 'missing inflation' episode, namely that economists, attuned the global nature of inflation (which was indeed crucial in the preceding period), missed the effect of domestic shocks during this episode.

The key lesson from the US historical decompositions, reported in the right part of Figure 4, is that much of the 'missing disinflation' is explained by the positive contributions of global shocks in 2010-2011. Regarding other periods, the US results confirm the role of global shocks in the post-2014 disinflation, similarly as in the euro area. As for the short-lived deep disinflation in 2009, the alternative identification produce different results: the Choleski identification explains it with the global shocks (recall that the conditional forecasts of section 2 attributed it to global variables), while the CDL and CDL+BB identifications assign a larger role for domestic shocks in this episode. These results suggest that in 2009 the US shocks strongly and quickly spilled over to global variables. Finally, and similarly as in the euro area, monetary policy shocks identified recursively (Choleski and CDL identifications) explain a small fraction of the variation of inflation, while monetary policy shocks identified with the BB identification explain quite a lot of its variation.

4 Conclusions

In a medium size VAR, conditional forecasts of inflation in and after the Great Recession match the actual realizations very well. In light of this VAR, inflation appears to be puzzling neither during the 'missing disinflation' in the US and the euro area, nor during the euro area 'missing inflation' episode.

The two disinflation episodes in the euro area appear to be of different nature. Global variables, and global structural shocks that we identify in a smaller VAR, are key to understanding the 'missing disinflation'. Weak domestic real activity and deflationary domestic shocks are key to understanding the euro area 'missing inflation'. Economists and models that excessively focused on the global nature of inflation systematically overpredicted it during the 'missing inflation' episode. Our results suggest that the link between inflation and real activity has been relevant and remained unchanged during and after the Great Recession.

Our results have practical implications for economic modelling. First, an empirically successful model of inflation needs to account well for both global factors and domestic real activity. Second, our results alleviate concerns about the time variation in the slope of the Phillips curve. A medium size VAR with constant coefficients appears to be a reasonable empirical model for capturing the relation between inflation and the variables that are most relevant for it.

Appendix A Data and sources

Table A.1: Euro area data sources and definition

Variable	Source	Description
	Price index	
HICP	Eurostat, ECB	Harmonized index of consumer prices, seasonally adjusted
	Domestic real activity variables	
Real GDP	Eurostat	chain linked volume, calendar and seasonally adjusted
Real Consumption	Eurostat	chain linked volume, calendar and seasonally adjusted
Real Investment	Eurostat	chain linked volume, calendar and seasonally adjusted
Total employment	Eurostat	persons, calendar and seasonally adjusted data
Unemployment rate	Eurostat	standardised unemployment rate, seasonally adjusted
Capacity utilization	BIS	sector covered: manufacturing
Consumer confidence	European Commission	survey indicator, seasonally adjusted, balance of responses
Purchasing Managers' Index (PMI)	Markit	survey indicator, seasonally adjusted, units
	Global variables	
GDP of the rest of the world	national sources, IMF, OECD, authors' calculations	euro area real GDP index was extracted from the world GDP index using the euro area share in world GDP (expressed in PPS)
Price of oil in USD	Bloomberg	Brent crude oil price in US dollars
Price of non-energy commodities	OECD	prices of raw materials, Index Total excluding energy (US dollars)
Nominal effective exchange rate	ECB	nominal effective exchange rate vis–a–vis 19 trading partners
USD/EUR exchange rate	ECB	exchange rate against euro, spot (mid)
US real GDP	BIS	chain linked volume, calendar and seasonally adjusted
US consumer prices	BIS	consumer price index, seasonally adjusted
Federal funds rate	BIS	daily and monthly average
	Financial variables	
EONIA	ECB, AWM	historical close, average of observations through period. Available since 1999, back-linked using data from the Area Wide Model (AWM) database
2–year government bond spread	ECB	2–year government benchmark bond yield – EONIA
10–year government bond spread	ECB	10–year government benchmark bond yield – EONIA
Mortgage bank lending spread	national sources,	loans to households for house purchase, weighted average
	Eurostat, ECB	of the rates for the EA big 5 (weights based on nominal GDP) – EONIA
Non-financial corporations bank	ECB	bank interest rates - loans to corporations (other than revolving
lending spread		loans and overdrafts) – EONIA
Corporate bond spread	ECB	difference between 7 to 10–year corporate bond yield and 7 to 10–year government bond yield

Notes: The world GDP index was constructed as follows. Starting 1995, quarterly national accounts data from individual national sources were used for countries covering around 93 per cent of world GDP. For the remaining countries the annual data provided by the World Economic Outlook database (WEO) was interpolated. The interpolation does not affect the aggregate world GDP, as the share of these countries is relatively small. The aggregation of individual data has been done using the annual PPP weights of each country provided by IMF WEO and a fixed base index was constructed (2005=100). Before 1995, a large set of individual country data is either not available or not reliable. In order to back-link the world GDP index for the pre-95 period, the annual world output data provided by WEO was interpolated using the quarterly data for real output of OECD countries. This ensures that the annual growth rate of world GDP is in line with the IMF estimate and the infra-annual dynamics is given by developments in countries covering a large share of world GDP.

Table A.2: US data sources and definition

Variable	Source	Definition
	Price index	
CPI	BIS	Consumer price index, seasonally adjusted
	Real activity variables	
Real GDP	BIS	chain linked volume, calendar and seasonally adjusted
Real Consumption	BIS	chain linked volume, calendar and seasonally adjusted
Real Investment	BIS	chain linked volume, calendar and seasonally adjusted
Total employment	BIS	thousands of persons, seasonally adjusted
Unemployment rate	BIS	seasonally adjusted
Capacity utilization	BIS	sector covered: manufacturing.
Consumer confidence	Bloomberg	conference Board Consumer Confidence, seasonally adjusted
Purchasing Managers' Index (PMI)	Bloomberg	US Chicago Purchasing Managers Index, seasonally adjusted
	Global variables	
GDP of the rest of the world	national sources, IMF, OECD, authors' calculations	US GDP index was extracted from the world GDP index using the US share in world GDP (expressed in PPS)
Price of oil in USD	Bloomberg	Brent crude oil price in US dollars
Price of non-energy commodities	OECD	prices of raw materials, Index Total excluding energy (US dollars)
Nominal effective exchange rate	BIS	nominal effective exchange rate vis–a–vis 26 trading partners
USD/EUR exchange rate	ECB	exchange rate against euro, spot (mid)
Euro area real GDP	Eurostat	chain linked volume, calendar and seasonally adjusted
HICP	Eurostat, ECB	Harmonized index of consumer prices, seasonally adjusted
Federal funds rate	BIS	daily and monthly average
EONIA	ECB, AWM	historical close, average of observations through period. Available since 1999, back-linked using data from the Area Wide Model (AWM) database
	Financial variables	
Federal funds rate	BIS	dayly and monthly average
2–year government bond spread	BIS	US Treasury 2–year bond yield – Fed funds rate
10–year government bond spread	BIS	US Treasury 10–year bond yield – Fed funds rate
Mortgage bank lending spread	BIS	mortgage rate (30 years maturity) – Fed funds rate
Non-financial corporations bank lending spread	BIS	bank prime lending rate to NFCs – Fed funds rate
Corporate bond spread	ECB	difference between 7 to 10–year corporate bond yield and 7 to 10–year government bond yield

Figure A.1: Euro area data. The series marked (yoy) have been transformed to year-on-year changes for this plot, but we used levels for the estimation.

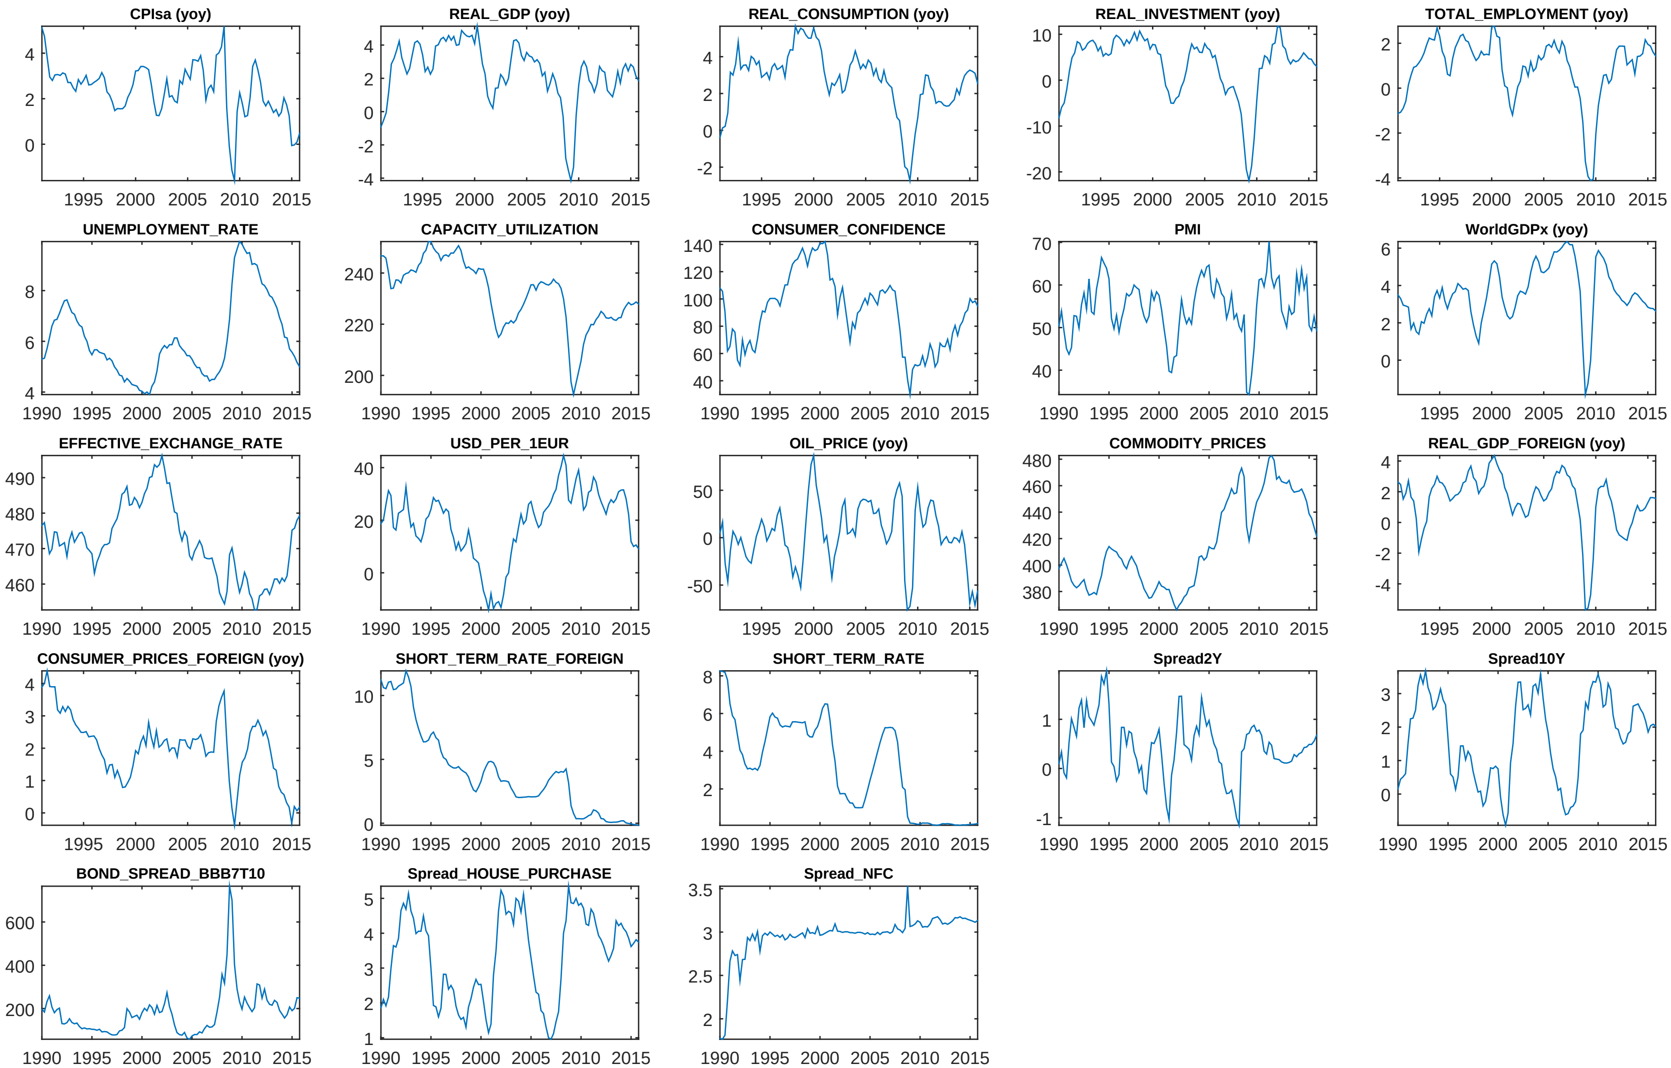

Figure A.2: US data. The series marked (yoy) have been transformed to year-on-year changes for this plot, but we used levels for the estimation.

Appendix B Impulse responses to all shocks

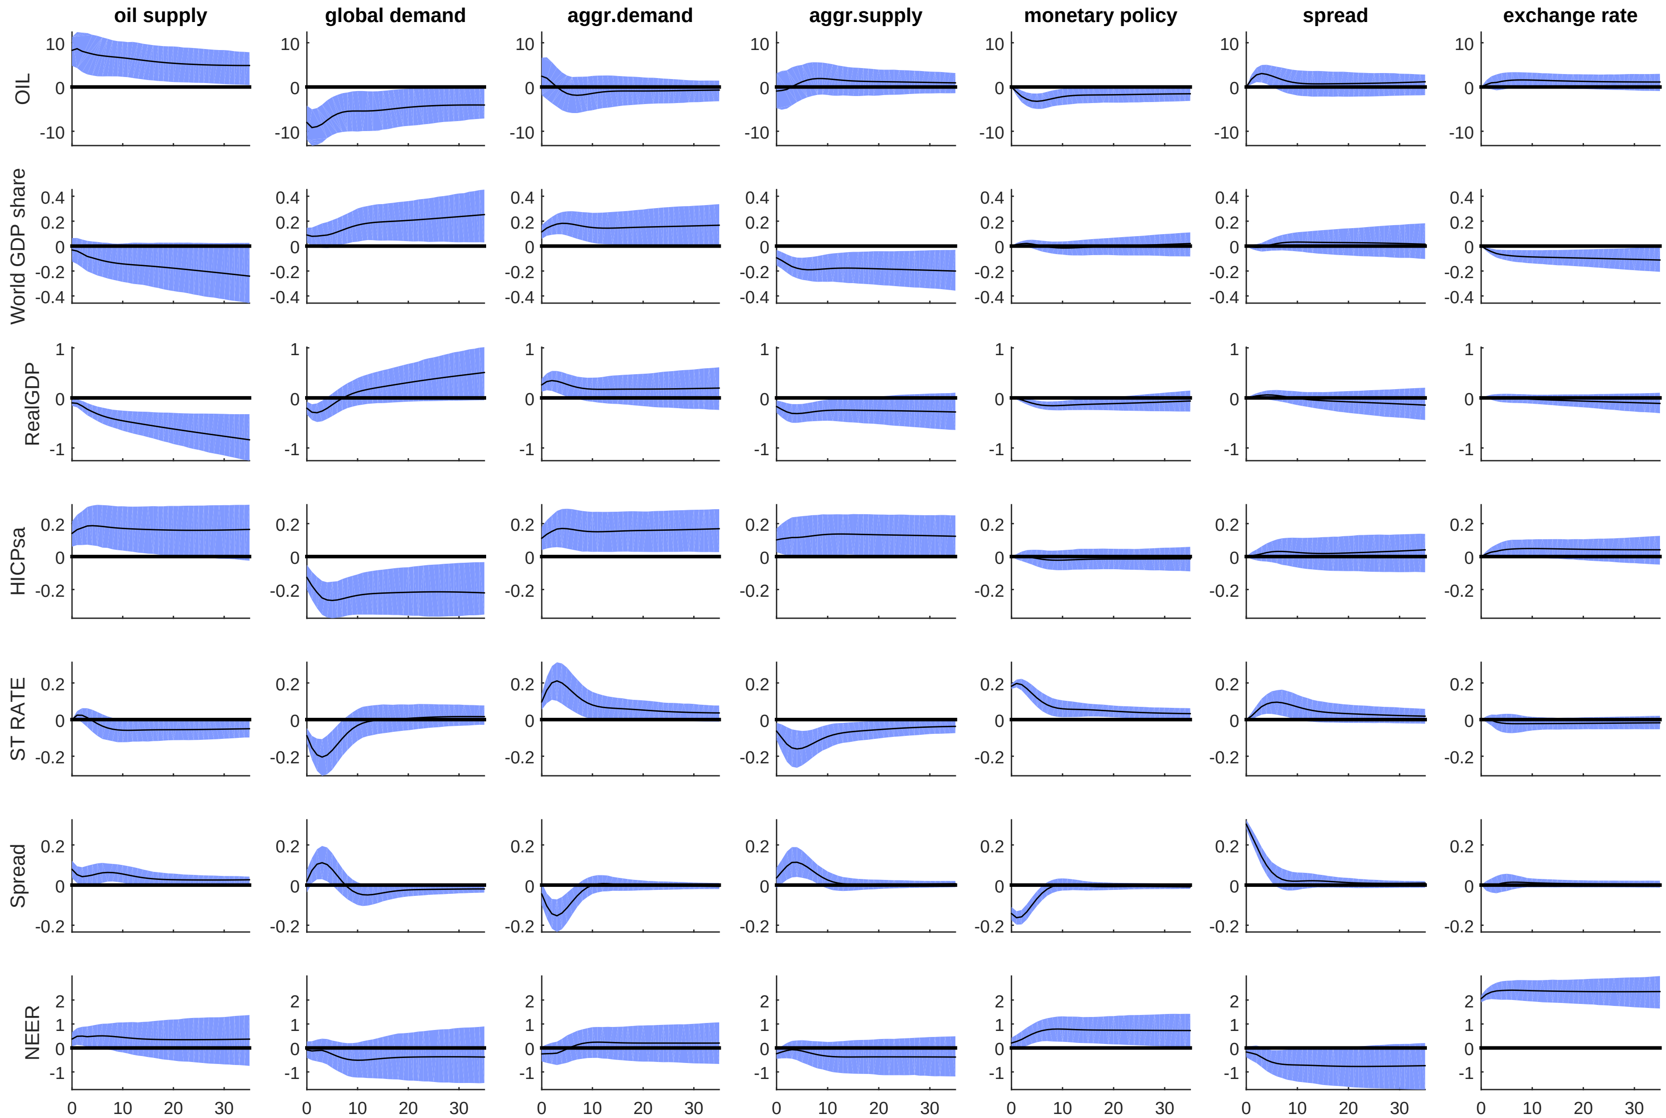

Figure B.1: Euro area impulse responses, identification II (CDL)

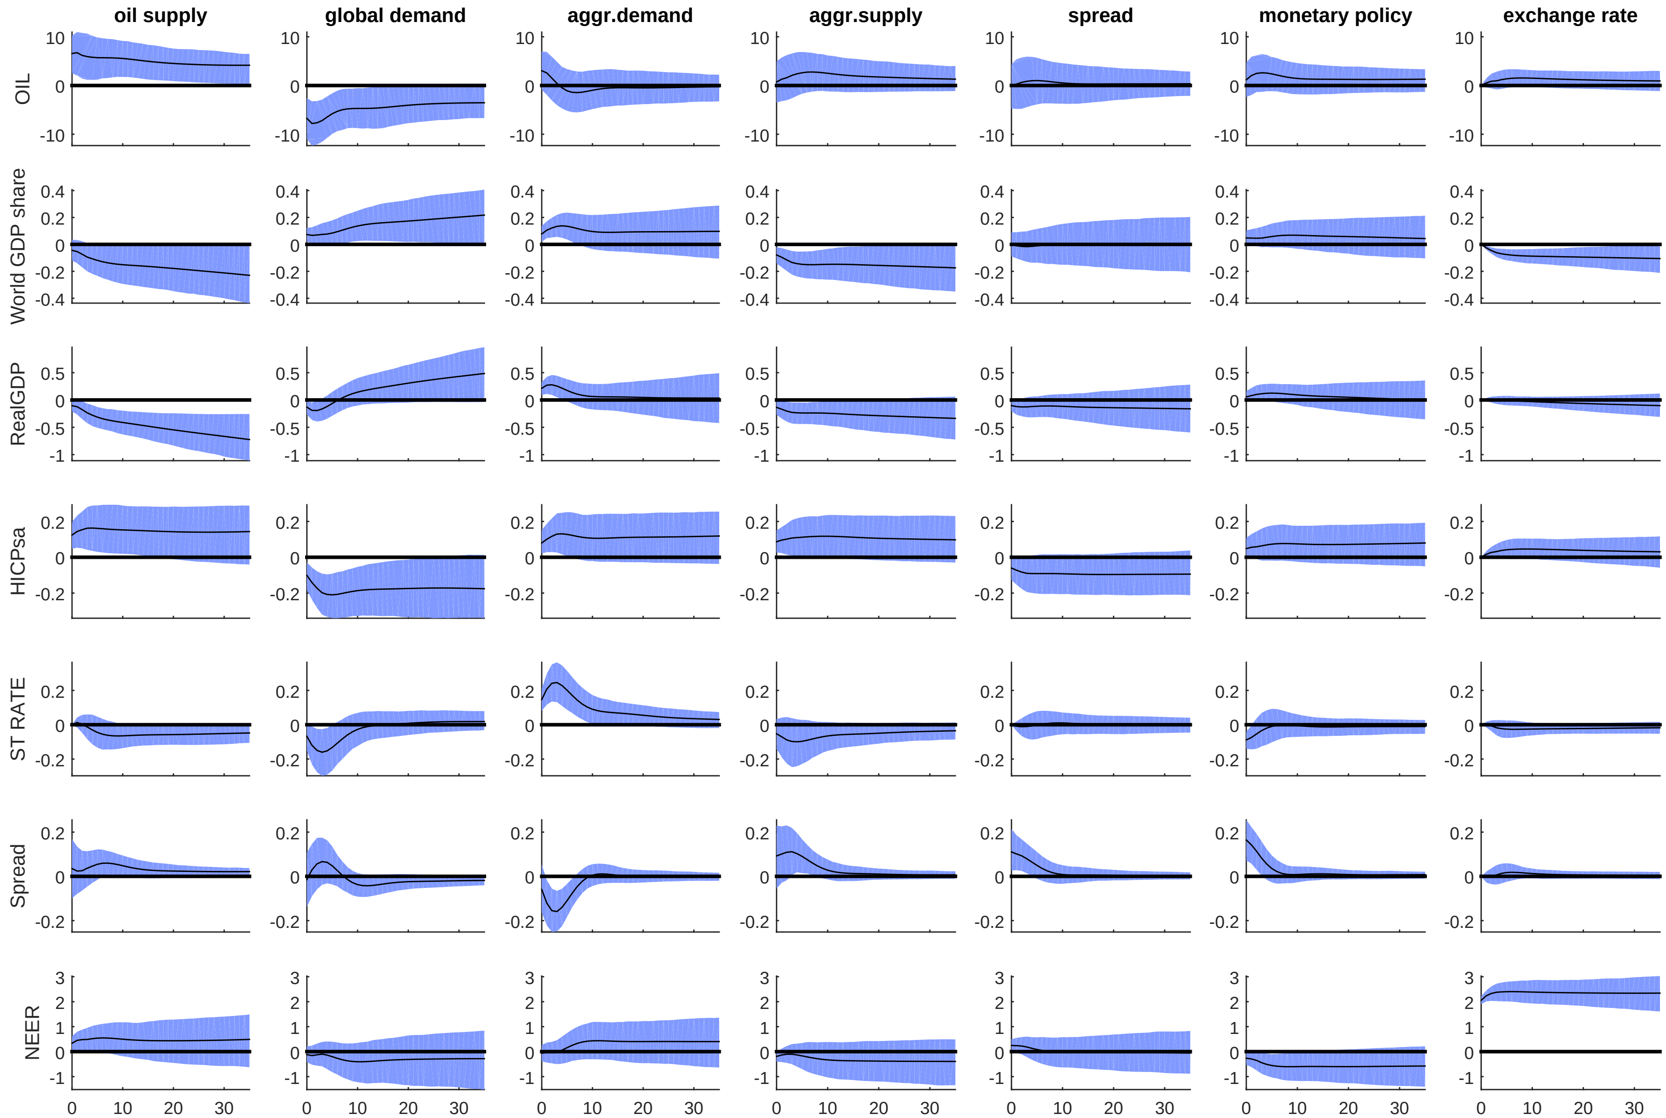

Figure B.2: Euro area impulse responses, identification III (CDL+BB)

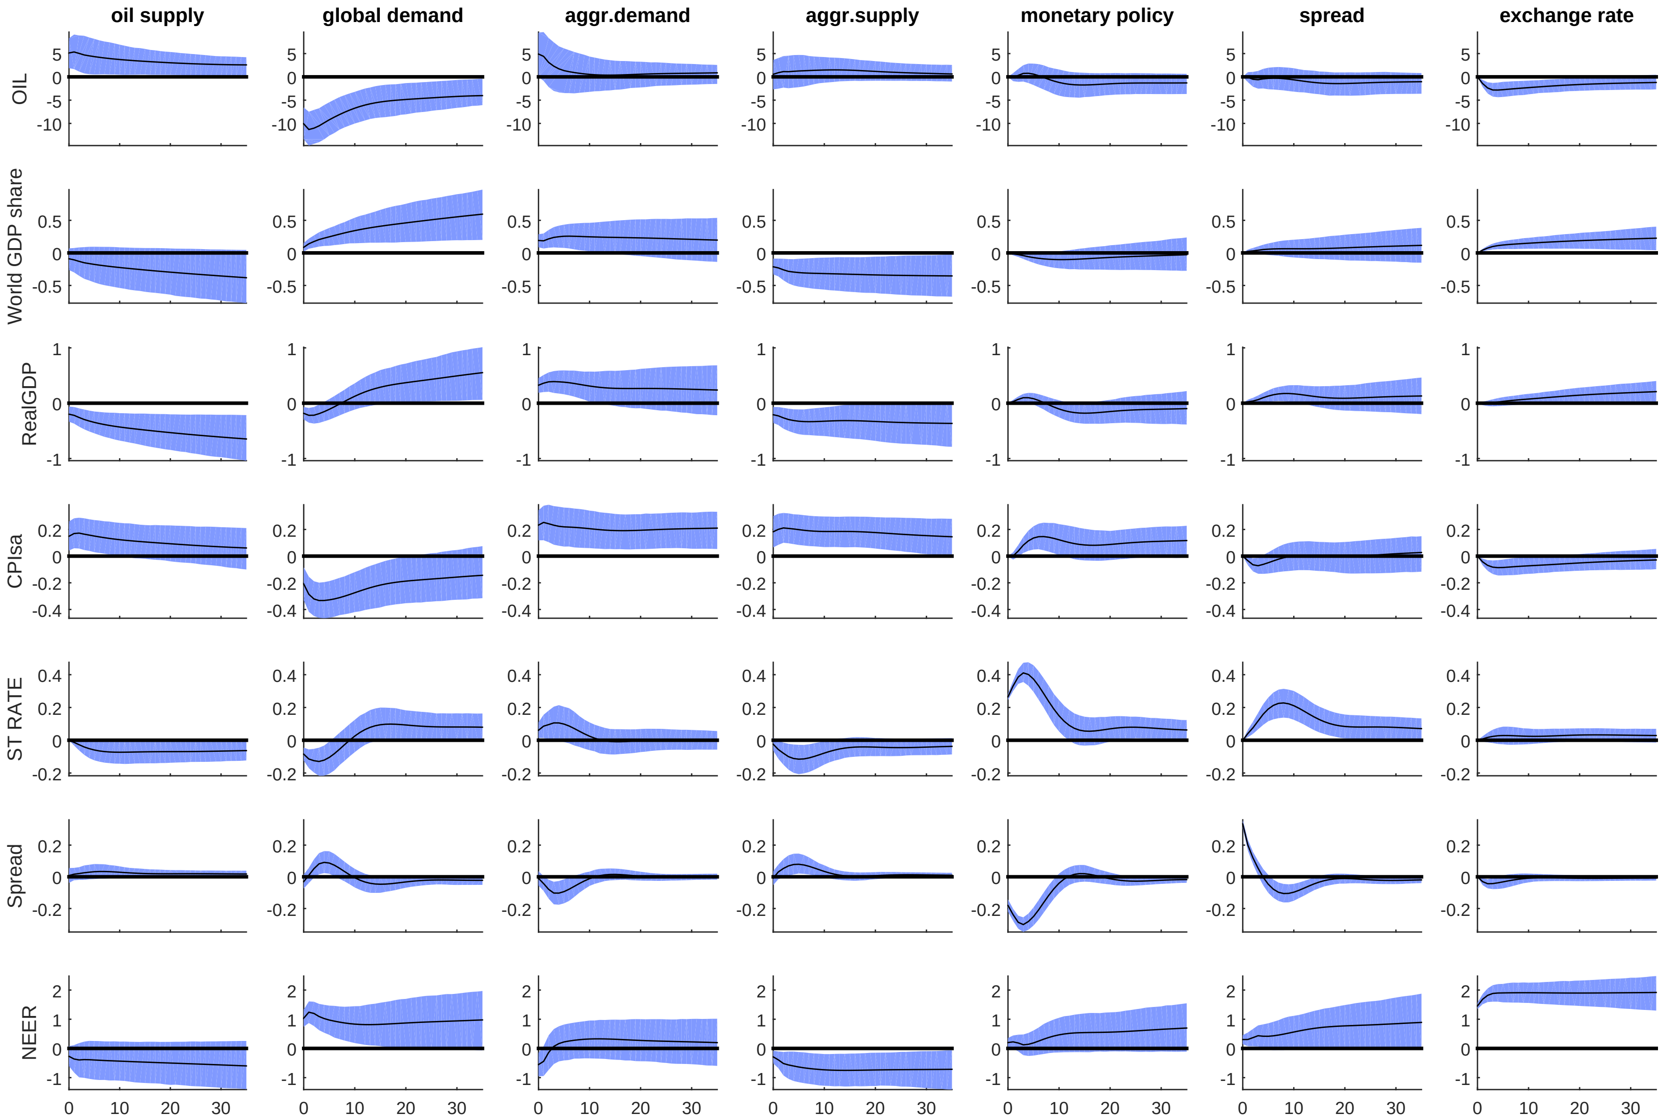

Figure B.3: US impulse responses, identification II (CDL)

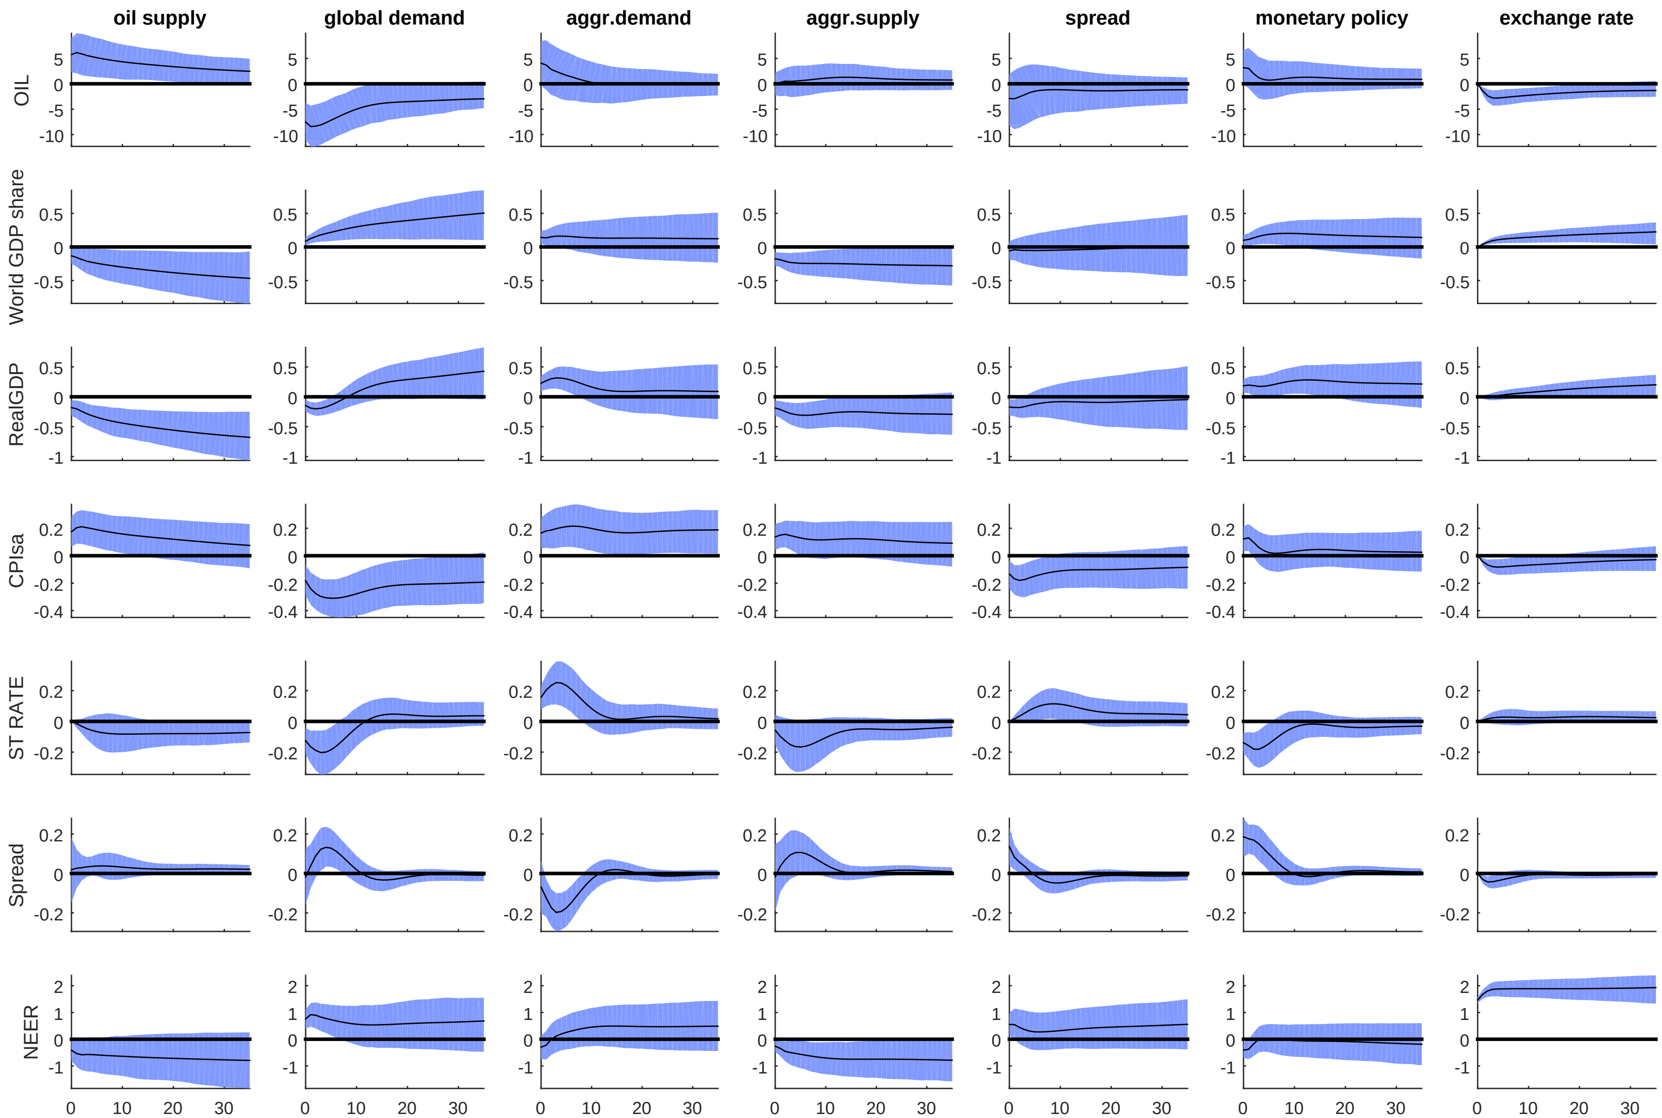

Figure B.4: US impulse responses, identification III (CDL+BB)

Appendix C Detailed historical decompositions

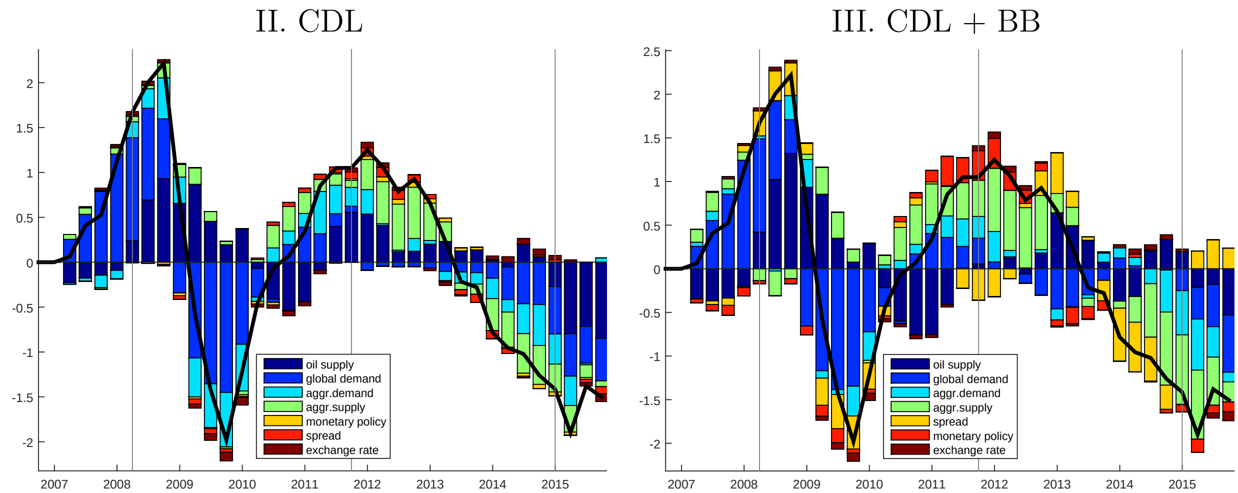

Figure C.1: Historical decompositions of euro area HICP: disaggregated effects of all the shocks

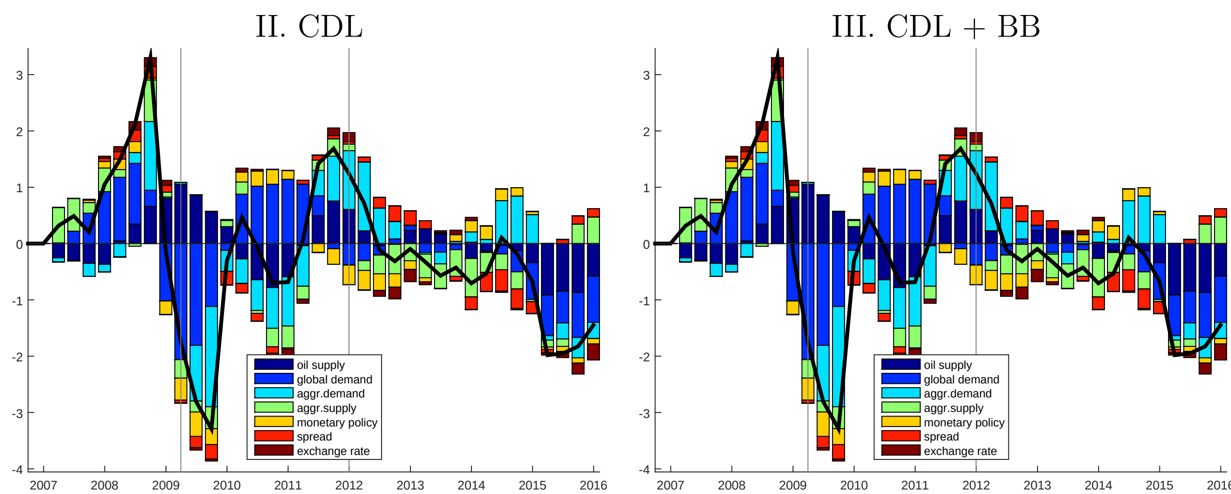

Figure C.2: Historical decompositions of US CPI: disaggregated effects of all the shocks

References

Aastveit, K. A., Carriero, A., Clark, T. E., and Marcellino, M. (2014). Have standard VARs remained stable since the crisis? Working paper no. 14-11, Federal Reserve Bank of Cleveland.

Arias, J. E., Rubio-Ramirez, J. F., and Waggoner, D. F. (2014). Inference based on SVARs identified with sign and zero restrictions: Theory and applications. International Finance Discussion Papers 1100, Board of Governors of the Federal Reserve System (U.S.).

Ball, L. and Mazumder, S. (2011). Inflation dynamics and the Great Recession. *Brookings Papers on Economic Activity*, 42(1 (Spring):337–405.

Bańbura, M., Giannone, D., and Lenza, M. (2015). Conditional forecasts and scenario analysis with vector autoregressions for large cross-sections. *International Journal of Forecasting*, 31(3):739 – 756.

Bańbura, M., Giannone, D., and Reichlin, L. (2010). Large Bayesian vector auto regressions. *Journal of Applied Econometrics*, 25(1):71–92.

Baumeister, C. and Benati, L. (2013). Unconventional monetary policy and the Great Recession: Estimating the macroeconomic effects of a spread compression at the zero lower bound. *International Journal of Central Banking*, 9(2):165–212.

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

Bianchi, F. and Melosi, L. (2016). Escaping the Great Recession. Unpublished manuscript, Cornell University and Federal Reserve Bank of Chicago.

Blanchard, O., Cerutti, E., and Summers, L. H. (2015). Inflation and activity: Two explorations and their monetary policy implications. Peterson Institute for International Economics, Working Paper WP15-19.

Christiano, L. J., Eichenbaum, M., and Evans, C. L. (1999). Monetary policy shocks: What have we learned and to what end? In Taylor, J. B. and Woodford, M., editors, *Handbook of Macroeconomics 1A*, pages 65–148. North-Holland, Amsterdam.

Christiano, L. J., Eichenbaum, M. S., and Trabandt, M. (2015). Understanding the Great Recession. *American Economic Journal: Macroeconomics*, 7(1):110–67.

Ciccarelli, M. and Mojon, B. (2010). Global inflation. *The Review of Economics and Statistics*, 92(3):524–535.

Coibion, O. and Gorodnichenko, Y. (2015). Is the Phillips curve alive and well after all? inflation expectations and the missing disinflation. *American Economic Journal: Macroeconomics*, 7(1):197–232.

Constancio, V. (2015). Understanding inflation dynamics and monetary policy. Speech at the Jackson Hole Economic Policy Symposium.

Conti, A. M., Neri, S., and Nobili, A. (2015). Why is inflation so low in the euro area? Temi di discussione (Economic working papers) 1019, Bank of Italy, Economic Research and International Relations Area.

Corsetti, G., Dedola, L., and Leduc, S. (2014). The international dimension of productivity and demand shocks in the US economy. *Journal of the European Economic Association*, 12(1):153–176.

Del Negro, M., Giannoni, M. P., and Schorfheide, F. (2015). Inflation in the Great Recession and New Keynesian models. *American Economic Journal: Macroeconomics*, 7(1):168–96.

Ferroni, F. and Mojon, B. (2015). Domestic and global inflation. Banque de France and Ecole Polytechnique, unpublished.

Friedrich, C. (2016). Global inflation dynamics in the post-crisis period: What explains the puzzles? *Economics Letters*, 142(2):31–34.

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

Gilchrist, S., Schoenle, R., Sim, J. W., and Zakrajsek, E. (2015). Inflation dynamics during the financial crisis. Finance and Economics Discussion Series 2015-12, Board of Governors of the Federal Reserve System (U.S.).

Hall, R. E. (2011). The long slump. *American Economic Review*, 101(2):431–69.

Jarociński, M. (2010). Conditional forecasts and uncertainty about forecast revisions in vector autoregressions. *Economics Letters*, 108(3):257–259.

Jarociński, M. and Maćkowiak, B. (2016). Granger-causal-priority and choice of variables in vector autoregressions. *Review of Economics and Statistics*. forthcoming.

Kilian, L. (2009). Not all oil price shocks are alike: Disentangling demand and supply shocks in the crude oil market. *American Economic Review*, 99(3):1053–69.

Litterman, R. B. (1979). Techniques of forecasting using vector autoregressions. Federal Reserve Bank of Minneapolis Working Paper number 115.

Mumtaz, H. and Surico, P. (2012). Evolving international inflation dynamics: World and country-specific factors. *Journal of the European Economic Association*, 10(4):716–734.

Simon, J., Matheson, T., and Sandri, D. (2013). The dog that didn't bark: has inflation been muzzled or was it just sleeping? *IMF World Economic Outlook*, pages 1–17. International Monetary Fund.

Sims, C. A. and Zha, T. (1998). Bayesian methods for dynamic multivariate models. *International Economic Review*, 39(4):949–68.

Waggoner, D. F. and Zha, T. (1999). Conditional forecasts in dynamic multivariate models. *The Review of Economics and Statistics*, 81(4):639–651.