Heterogeneity in Euro-Area Monetary Policy Transmission: Results from a Large Multi-Country BVAR Model[1]

Martin Mandler, Deutsche Bundesbank

Michael Scharnagl, Deutsche Bundesbank

Ute Volz, Deutsche Bundesbank

February 2016

Abstract

We study cross-country differences in monetary policy transmission across the large four euro-area countries (France, Germany, Italy and Spain) using a large Bayesian vector autoregressive model with endogenous prior selection. Drawing both on the posterior distributions of the cross-country differences in impulse responses as well as on a battery of other tests, we find real output to respond less negatively in Spain to monetary policy tightening than in the other three countries, while the decline in the price level is weaker in Germany. Bond yields rise more strongly and more persistently in France and Germany than in Italy and Spain.

Keywords: monetary policy, transmission mechanism, euro area, Bayesian vector autoregression

JEL-Classification: C11, C54, E52

[1] Contact address: Wilhelm-Epstein-Str. 14, D-60431 Frankfurt am Main. E-mail: martin.mandler@bundesbank.de, michael.scharnagl@bundesbank.de, ute.volz@bundesbank.de. The authors thank Carlo Altavilla, Jörg Breitung, Sandra Eickmeier, Björn Fischer, Michele Lenza, Harald Uhlig and seminar participants at the Deutsche Bundesbank for helpful comments and discussions. This paper represents the authors' personal opinions and does not necessarily reflect the views of the Deutsche Bundesbank or its staff.

1 Introduction

We analyse empirically whether the common monetary policy of the Eurosystem has heterogeneous effects on the four large member countries of the euro area. This issue is of high relevance for monetary policymakers because of its implications both for the effectiveness and for the welfare effects of the common monetary policy. The Eurosystem has defined price stability in terms of an average of national price developments weighted by the countries' relative household consumption expenditure shares. Hence, if inflation in larger countries responds very weakly to changes in the Eurosystem's monetary policy stance compared to the response in smaller ones, this implies that stronger changes in the policy stance would be required to bring average inflation back to target compared to a situation of identical responses across countries. Even in the latter case, pronounced differences in the responsiveness of output could nevertheless imply an asymmetrical distribution of the burden of adjusting to euro-area-wide inflationary disequilibria,[2] where welfare losses could be reduced by taking national information into account instead of just looking at euro-area-wide aggregates (e.g. De Grauwe, 2000; De Grauwe and Piskorski, 2001, Angelini, Del Giovane, Siviero and Terlizzese, 2002).[3]

From an optimal policy perspective, cross-country differences in the effects of monetary policy might also raise questions concerning the appropriateness of a central bank objective that is defined using a (weighted) price level average. If, for example, there is heterogeneity across the countries of a currency union that is due to a different degree of nominal rigidities, theory suggests assigning higher weights to countries with more pronounced rigidities (Benigno, 2004; Soares, 2008).

The most commonly used empirical methodology for studying monetary transmission without using structural dynamic stochastic general equilibrium models is based on vector autoregressive (VAR) models and the analysis of the effects of identified monetary policy shocks. Given the length of macroeconomic time series available for many countries, and in order to be able to include more than two or three time series per

---

[2] The welfare implications of such an observation, however, are not clear-cut. In simple New-Keynesian type models, the micro-founded welfare criterion for each country would depend on output gap volatility and on the volatility of inflation around the inflation target. Given that most VAR analyses indicate that monetary policy shocks make only a modest contribution to unexpected fluctuations in these variables, their volatility would depend on the (symmetric or asymmetric) shocks hitting each country and on the effectiveness of the common monetary policy in stabilising the economy in the face of these shocks. Such an analysis is beyond the purely empirical approach of this paper.

[3] In the study by De Grauwe and Piskorski (2001), the welfare losses from having the central bank only respond to union-wide aggregates instead of including national information in the reaction function turn out to be relatively small. However, Angelini et al. (2002) find substantially larger welfare losses.

country, researchers have often estimated VAR models for individual countries and then compared the impulse responses and other results across different countries.[4] This approach, however, either tends to ignore the transmission channels which work through the interactions and spillover effects among the countries (e.g. Berben, Locarno, Morgan and Vallés, 2004) or captures these interactions by including a small set of additional variables (e.g. trade balance, foreign output) in order to approximate these interactions.

Properly accounting for cross-country interactions is particularly relevant in a monetary union with highly integrated economies. Thus, the estimation of a multi-country model, which allows for the interaction of variables across countries, might be more appropriate for studying cross-country heterogeneity in monetary policy transmission. However, a consequence of this is that the inclusion of a reasonable number of variables for each country and an adequate number of lags quickly exhausts the available degrees of freedom. This curse of dimensionality is even more severe in the case of the euro area because only about 15 years of data are available for the common monetary policy regime.[5] Estimating the VAR model with Bayesian methods overcomes this limitation by augmenting the information in the data with prior information.

In this paper, we use a Bayesian VAR (BVAR) model for the four large euro-area countries (DE, ES, FR, and IT) to study cross-country heterogeneity in the effects of the Eurosystem's monetary policy. For the estimation of the BVAR we use an approach for the selection of priors recently proposed by Giannone, Lenza, and Primiceri (2015), which treats the hyperparameters controlling the priors as random variables for which relatively uninformative "hyperpriors" are assumed. Furthermore, the Bayesian estimation of the VAR and the impulse responses in a multi-country setting allows us to analyse potential differences in transmission across countries in a rigorous way. Instead of just a visual inspection and comparison of impulse responses between different countries, the estimation procedure provides simulated distributions of cross-country *differences* in impulse responses, allowing us to derive probabilistic statements about these differences, also taking estimation uncertainty into account. This rigorous approach to investigating differences in the effects of monetary policy shocks across countries is an important improvement on the previous literature, which almost exclusively relies on the comparison of point estimates or on "eyeballing" impulse response distributions.

4 See, for example, Mojon and Peersman (2001), and Dedola and Lippi (2005), Section 2. Some of these models include a small number of foreign variables as exogenous variables which, however, does not allow for feedback effects.

5 Adding pre-euro data might imply a structural break in the monetary policy regime.

Our analysis provides evidence of marked differences in the effects of monetary policy shocks across the four large euro-area countries. We find that the price level response in DE to an exogenous monetary policy shock is relatively weak compared to the other three countries. However, the output response in DE is fairly similar to or somewhat stronger than the one in FR and IT, while real output in ES responds much less. For the long-term bond yield we find stronger and more persistent reactions in DE and FR than in ES or IT.

2 Selected literature

Some of the previous studies of heterogeneity in monetary policy transmission in the euro area are based on pre-euro data. For example, Ehrmann (2000) compares the impulse responses to a monetary policy shock across 13 EU countries using single-country VAR models with four variables. His results indicate a much stronger reaction of output in DE than in ES and IT, and a weaker reaction in FR than in the other three countries. For inflation, the response is shown to be less pronounced in FR and IT than in DE and ES. However, he does not provide any evidence on the significance of the estimated differences. Mojon and Peersman (2001) also estimate individual VAR models for many euro-area countries using pre-euro data and impose identification assumptions based on how strong the national monetary policy was linked to the German one under the European Monetary System.[6] While their results indicate cross-country differences in the means of the impulse responses to a monetary policy shock, the single-country-plus-Germany setup does not allow this hypothesis to be rigorously tested.

Boivin, Giannoni and Mojon (2009) use a combination of pre-euro data and data since the introduction of the euro and estimate a factor-augmented VAR. The factor structure allows them to model all euro-area countries jointly and to recover the impulse responses to a monetary policy shock for each country from the common factors. Their results show qualitatively different reactions of long-term bond yields to a monetary policy shock in IT and ES compared to the other countries. Furthermore, while the reaction of output growth is shown to be similar across countries, they obtain marked differences in the responses of the components of GDP, with consumption and investment declining more strongly in IT and ES. Their study, however, does not provide any formal evidence on the significance of these differences and the probability bands around the impulse responses are fairly wide. Barigozzi, Conti and Luciani

[6] For other studies based on VARs for individual countries see, for example, Mihov and Scott (2001) or Rafiq and Mallick (2008).

(2014) estimate a structural dynamic factor model for several euro-area countries over the period from 1983 to 2007. Comparing the pre-euro and euro periods, they show that the monetary policy transmission mechanism has evolved towards more similar reactions, especially for output. Marked differences between countries remain for inflation and unemployment, which they attribute to country-specific labour market characteristics.

Using only data from the euro period, Cavallo and Ribba (2014) estimate a near-VAR, which includes a block of euro-area variables together with a single country's output and price level and assumes that while the individual country variables are affected by the euro-area variables, the euro-area variables evolve independently of the country block. Their estimates indicate relatively similar impulse responses of output and inflation to a euro-area restrictive monetary policy shock across countries.

Georgiadis (2015) studies cross-country heterogeneity in monetary transmission using a Global VAR (GVAR) model that includes 14 euro-area and 47 non-euro-area economies and is estimated on quarterly data from 1999 to 2009. Each country block for the euro-area economies contains output growth and inflation as endogenous variables. The common monetary policy is modelled as a block of its own in which the Eurosystem responds to euro-area aggregate output and inflation as well as to foreign interest rates, the exchange rate and oil prices. The impulse responses of real GDP to a euro-area monetary policy shock identified through sign restrictions differ markedly across euro-area countries, with DE showing the biggest (median) trough reaction followed by IT, ES and FR. Using correlation and regression analysis, Georgiadis explains these differences with different industry structures and differences in labour market rigidities.

Ciccarelli and Rebucci (2006) study heterogeneity in the reaction of real output in DE, ES, FR, and IT to a monetary policy shock. They use a time-varying-parameter VAR (TVP-VAR) with the four countries' real output as endogenous variables, some controls and a common monetary policy shock, which is identified as a shock to the German policy rate and constructed from a similar TVP-VAR in the four countries' monetary policy interest rates. They carefully try to control for the fact that policy rates in the other three countries could still deviate from the German rate over their sample period (1980-1997), and for possible exchange rate movements that might lead to a misidentification of the effects of the policy shock on output. While their focus is on the changes in the output effects of monetary policy shocks through time, they also present results on the cross-country differences of these effects, which indicate a smaller negative output response in FR following a monetary policy shock than in the other three countries.

Heterogeneity and time-variation in the effects of monetary policy shocks among EMU countries is also analysed by Ciccarelli, Maddaloni and Peydro (2013). They estimate a panel VAR model for two groups of euro-area countries: countries that came under stress in the financial and sovereign debt crises and those that did not. Because they include 15 variables per country, and due to the very short sample (2002Q4-2007Q4), they impose identical VAR coefficients and impact effects of shocks within each group and allow no cross-country correlations in the shock matrices, either within or between groups, as well as no direct spillover effects. According to their results, the output response in the non-stressed countries to a monetary policy shock was generally weaker and less persistent compared to the stressed countries, and the inflation response was weaker as well. Furthermore, while the strength of the output response has become stronger during the crises, their results show this change being more pronounced for the stressed countries as a result of changes in the amplification effects of the credit channel.

Much of the literature is, at least in part, based on data from the pre-euro period, when the future euro-area members still had individual monetary policies (although to differing extents). Therefore, the estimated differences in impulse responses to a monetary policy shock might either reflect differences in the way a country's economy reacts to monetary policy (transmission mechanism in a narrow sense) or differences in the country's monetary policy reaction function, which describes how the national monetary policy endogenously reacts to shock-induced movements in variables (e.g. Guiso, Kashyap, Panetta and Terlizzese, 2000). However, only the first element is relevant for any conclusions about heterogeneity after the introduction of the euro, as all countries within the euro area are subject to an identical monetary policy reaction function. This requires either focusing exclusively on data after the introduction of the euro, or carefully modelling the monetary policy reaction functions and the monetary policy shock, as in Mojon and Peersman (2001) and in Ciccarelli and Rebucci (2006).[7] In our paper, we follow the first strategy and use only data from 1999 onward.

[7] For countries such as Austria and the Netherlands, which tied their monetary policy very closely to the Bundesbank, this is less of an issue, but many other countries which would later join the euro area still retained some leeway in their monetary policy relative to Germany, despite being members of the European Exchange Rate Mechanism.

3 Econometric approach and data

3.1 Econometric approach

The estimation of the BVAR model follows the approach of Giannone, Lenza and Primiceri (GLP, 2015). Instead of specifying prior distributions for the parameters using fixed hyperparameters,[8] they treat these hyperparameters as random variables for which prior beliefs can be summarised by "hyperpriors".

The VAR model for n variables $y_t = \left(y_{1,t}, \ldots, y_{n,t}\right)'$ can be written as

$$y_t = c + A_1 y_{t-1} + \ldots + A_p y_{t-p} + \varepsilon_t \qquad \varepsilon_t \sim N\left(0, \Sigma\right). \tag{1}$$

c is a vector of intercepts, A_i is a $n \times n$ matrix of coefficients on lag i, p is the number of lags and ε_t is a vector of residuals with mean zero. For later reference, we collect the A_i matrices and c in the matrix Γ. The $n \times n$ covariance matrix of (reduced-form) shocks is represented by Σ.

GLP specify the prior in the spirit of the Minnesota approach, where all variables are assumed to follow independent random walks. The coefficients in A_i are assumed to be normally distributed with mean

$$E\left[A_{s,ij} \middle| \Sigma, \gamma\right] = \begin{cases} 1 & \text{if } i = j \text{ and } s = 1 \\ 0 & \text{otherwise} \end{cases}, \tag{2}$$

conditional on the vector of hyperparameters γ and on the covariance matrix Σ. Conditional on the hyperparameters, the coefficient on the first own lag of each variable in the VAR is assumed to have a mean of one, while all other coefficients have a prior mean of zero. The prior covariance of the coefficients in A_i is given by

$$\operatorname{cov}\left[A_{s,ij}, A_{r,hm} \middle| \Sigma, \lambda, \Psi\right] = \begin{cases} \lambda^2 \dfrac{1}{s^2} \dfrac{\Sigma_{ih}}{\Psi_{jj}} & \text{if } m = j \text{ and } r = s \\ 0 & \text{otherwise} \end{cases}, \tag{3}$$

with the hyperparameters λ and Ψ being elements of the vector γ. The lower the variance of the coefficient, the stronger it is pushed towards the prior mean (2).

[8] Litterman (1980) and Doan, Litterman and Sims (1984) determine the hyperparameters by minimising the forecast errors over a (non-overlapping) training sample. Bańbura, Giannone and Reichlin (2010) select the hyperparameters by a grid search and by matching the BVAR in-sample fit to that of a small VAR estimated by OLS.

Therefore, the higher the lag s, the more shrinkage is imposed. GLP use s^2 as the lag decay rate. The importance of the prior relative to the data is controlled by the overall shrinkage parameter λ. The larger λ is, the smaller the relative influence of the prior information. The term Σ_{ij}/Ψ_{jj} accounts for different scales of variables.

The prior on the covariance matrix is assumed to be inverse-Wishart:

$$\Sigma \sim IW(\Psi, d) \tag{4}$$

where $d = n + 2$ represents the degrees of freedom. The scale matrix Ψ is assumed to be diagonal, with the diagonal elements in Ψ, ψ_i being treated as hyperparameters.

GLP include additional priors commonly used in the literature by adding dummy observations to the data matrices. The "sum-of-coefficients" prior assumes that for each variable the sum of coefficients on own lags is equal to one and that the sums of coefficients on other variables are equal to zero. This introduces correlation among the coefficients and implies a unit root in each equation (no cointegration). The prior's tightness is controlled by a hyperparameter $\mu > 0$, with smaller values of μ implying a more informative prior. The "initial-dummy-observation" prior implies that the initial observations (average of first p observations of each variable) deliver a good forecast. Its tightness is controlled by another hyperparameter $\theta > 0$.

Thus, GLP use a hierarchical approach that decomposes the joint prior distribution of the VAR parameters as follows

$$p(\Gamma, \Sigma, \gamma) = p(\Gamma | \Sigma, \gamma) p(\Sigma | \gamma) p(\gamma), \tag{5}$$

where γ collects all the hyperparameters λ, μ, θ, and the elements in Ψ, $p(\Sigma|\gamma)$ is given by (4) and $p(\Gamma|\Sigma, \gamma)$ is given by (2) and (3). Conditional on the hyperparameters, the model retains a standard prior in the form of a normal-inverse Wishart, but the posterior distribution of the VAR parameters takes into account the uncertainty surrounding the hyperparameters.

The priors on the hyperparameters are chosen to be relatively uninformative, i.e. flat: the priors for λ, μ and θ are assumed to be Gamma distributions with mode equal to 0.2, 1 and 1. The joint posterior distribution of the VAR parameters in Γ and Σ can, conditional on the hyperparameters, be estimated by means of the Gibbs sampler, as in standard BVAR models.

The posterior of the hyperparameters is given by

$$p(\gamma|y) \propto p(y|\gamma)p(\gamma), \qquad (6)$$

with y representing the data. Since the hyperprior $p(\gamma)$ is known, the only unknown element in (6) is the marginal likelihood $p(y|\gamma)$. GLP show that a solution for the marginal likelihood can be obtained in closed form, which enables drawing directly from this distribution without any need for simulation.[9]

The Markov-Chain-Monte-Carlo (MCMC) algorithm for drawing from the posterior of the BVAR combines a Gibbs sampler for the VAR coefficients and covariances conditional on the hyperparameters with a Metropolis-Hastings step for generating draws of the hyperparameters using the closed form solution to (6). The simulation starts at the mode of the posterior of the hyperparameters, which is determined by numerical optimisation. [10]

3.2 Data and identification

In compiling our data set we follow the specification of Altavilla, Giannone and Lenza (2014) and include six country-specific variables:[11] real GDP (RGDP), HICP, M3 money stock, loans to the non-financial private sector (households and non-financial corporations) (LOANS), a short-term government bond yield (R1Y)[12] and a long-term government bond yield (R10Y). Two euro-area-wide variables are used: the EONIA as a proxy for the monetary policy stance and bond market implied volatility (VOLA). The model is estimated on quarterly data from 1999Q1 to 2014Q3.

Starting the estimation in 1999 and not in the pre-euro period has the important advantage that it avoids mixing up heterogeneity due to different effects of the common monetary policy shock on each country's economy, which can be attributed to different structural characteristics, and heterogeneity due to differences in monetary policy reaction functions across countries[13].

Identification of the monetary policy shock is achieved through a Choleski decomposition of the VAR covariance matrix with RGDP and HICP for each country ordered before the EONIA and all other variables ordered after it.[14] This assumption

[9] See Appendix A in GLP.
[10] Estimations and simulations were carried out in MATLAB using a modified and extended version of the code from Giorgio Primiceri's website
http://faculty.wcas.northwestern.edu/~gep575/GLPreplicationWeb.zip
[11] Appendix A provides details on the data.
[12] Two-year maturity in the case of DE because of data availability.
[13] See the discussion in Section 2.
[14] Since we only want to identify the monetary policy shock, the actual ordering within the groups of variables in front of/after the EONIA does not matter. See Christiano, Eichenbaum, and Evans (1999).

implies that output and prices do not respond to a monetary policy shock within the same quarter but that the other (financial) variables do. This is the identical assumption to the one made in Altavilla et al. (2014). As a robustness check, we also identify the monetary policy shock by a combination of zero and sign restrictions following the approach of Arias, Rubio-Ramírez and Waggoner (2014). The following results are based on 5,000 saved draws from the MCMC, which represent every 250th draw after discarding 1.75 million initial draws.[15]

4 Results

4.1 Effects of a monetary policy shock

Figures 1 - 5 present impulse responses following a monetary policy shock defined as an exogenous innovation in the EONIA of 25 basis points (bp) and identified using the Choleski decomposition. The graphs show the median (blue line) together with a 68% credible interval (grey shaded area) enclosed by the 16% and 84% percentiles of the posterior impulse response distribution. In interpreting the impulse response distributions, we base our assessment on whether the vast majority of the probability mass of the simulated posterior is located above or below zero. Impulse responses are grouped by country (Figures 1 - 4) while Figure 5 shows the two euro-area aggregate variables.

[15] For this benchmark specification our choice of step length in the Metropolis-Hastings step resulted in an overall acceptance ratio of 27.9%, which is in the recommended range between 20% and 30%. The very long chain in our simulations is necessary due to the very slow convergence of the MCMC algorithm. Basically, the Metropolis-Hastings step is not very efficient for some of the hyperparameters, which then translates into convergence problems in the Gibbs sampler for the VAR parameters, as well.

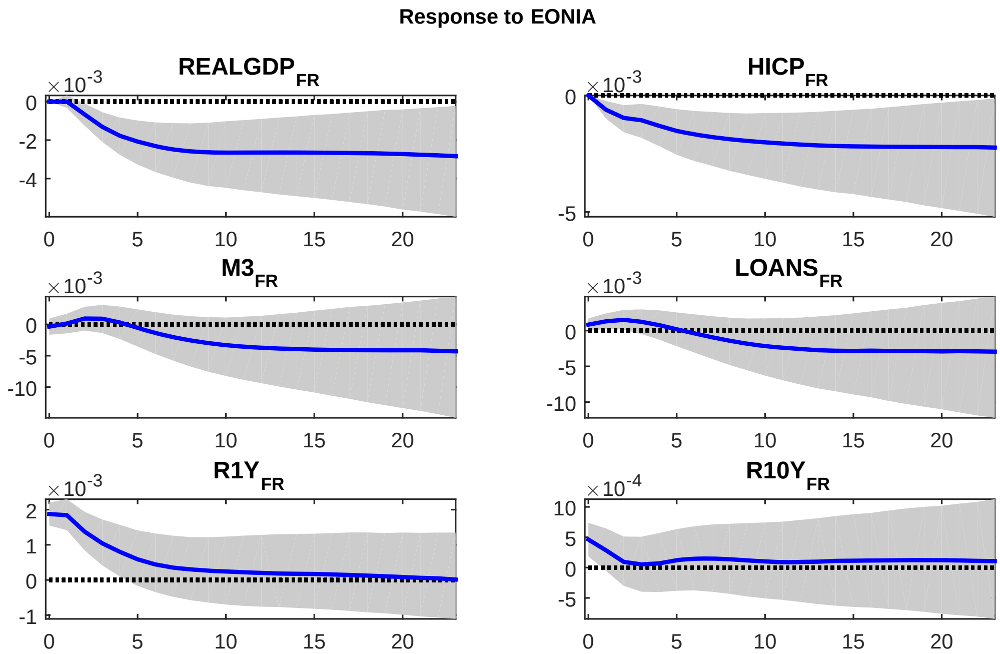
Response to EONIA
REALGDP FR
HICP FR
M3 FR
LOANS FR
R1Y FR
R10Y FR
×10-3
×10-4

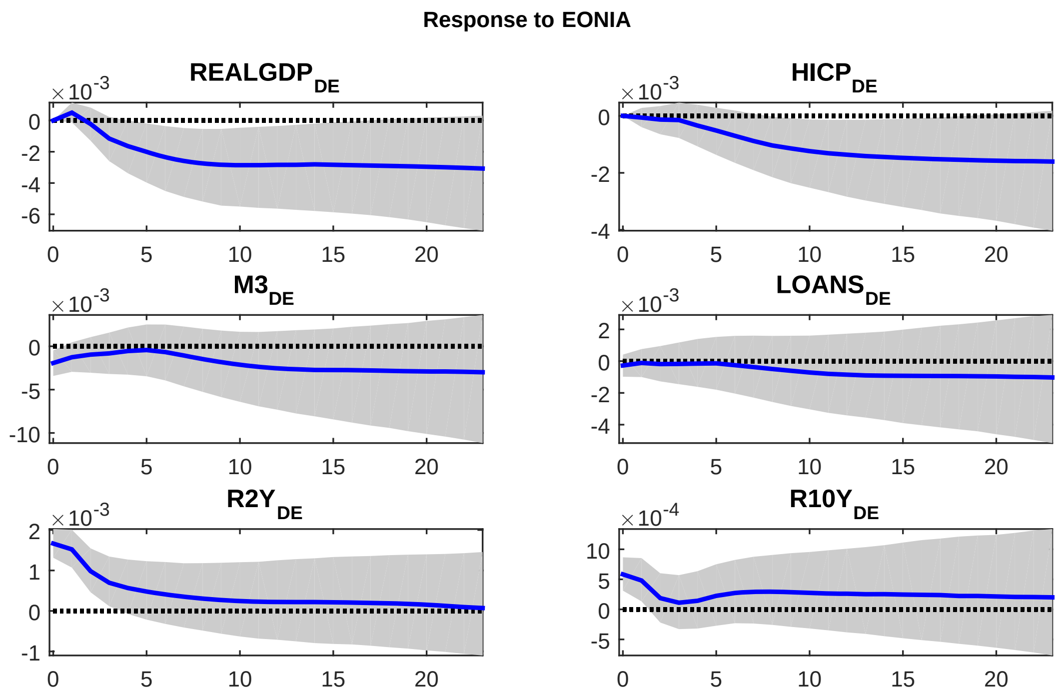
Response to EONIA
REALGDP DE
HICP DE
M3 DE
LOANS DE
R2Y DE
R10Y DE
×10-3
×10-4

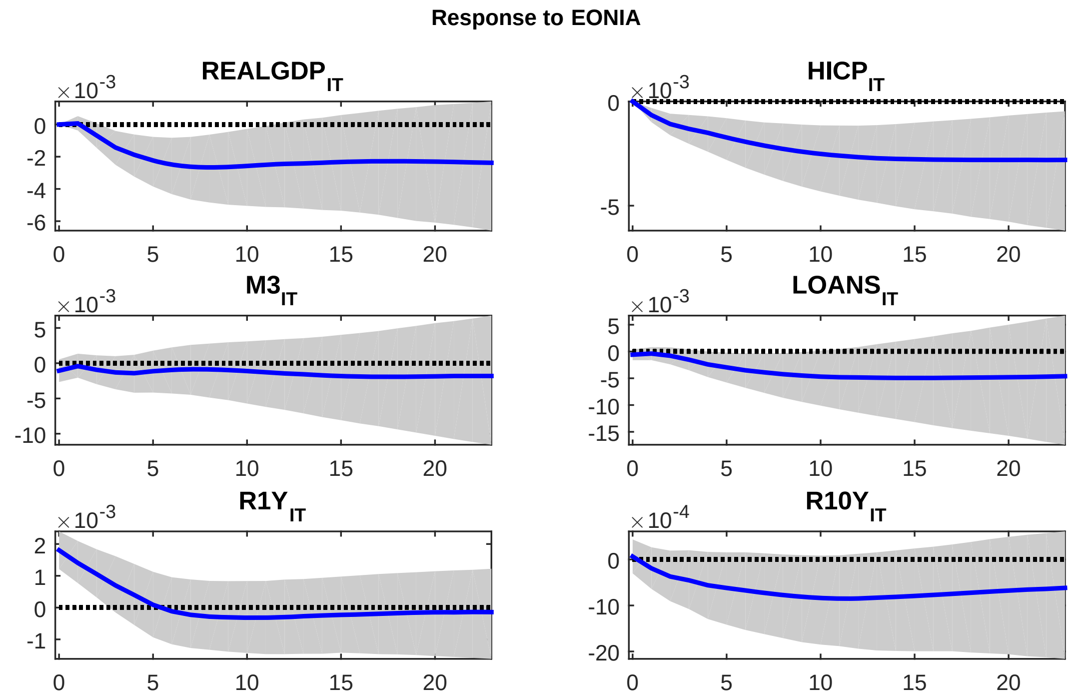
Response to EONIA
REALGDP IT
×10-3
HICP IT
M3 IT
LOANS IT
R1Y IT
R10Y IT
×10-4

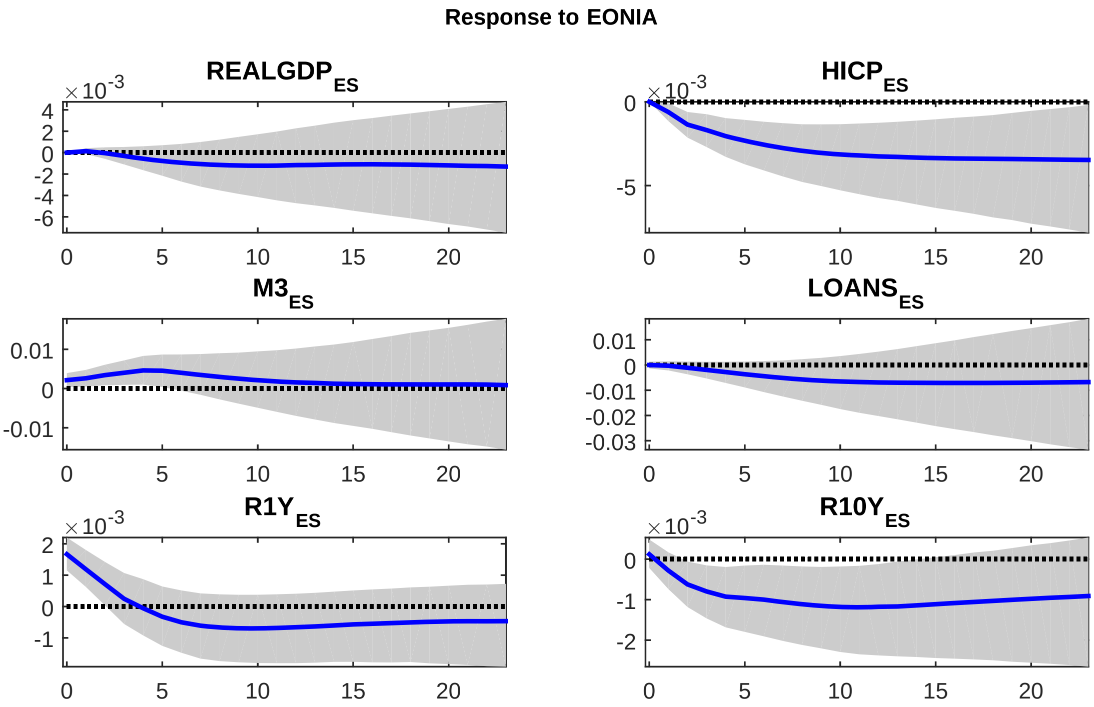
Response to EONIA
REALGDP ES
×10-3
HICP ES
M3 ES
LOANS ES
R1Y ES
R10Y ES

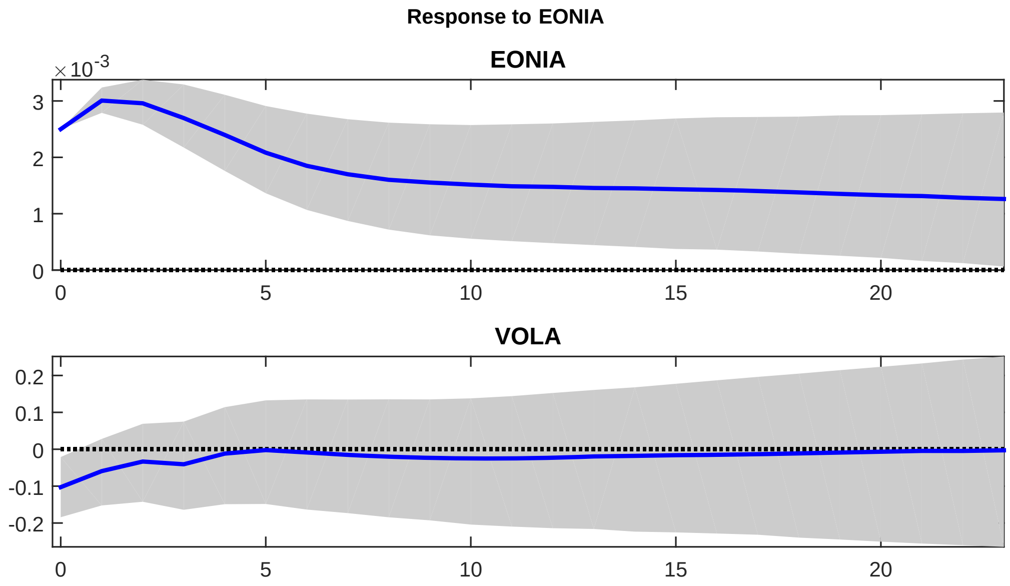

estimated much more imprecisely than in the other countries. Turning to bond yields, the short-term bond yield reacts positively in all countries. The long-term bond yield responds positively in DE and FR but returns quickly to zero. For IT and ES, after a few quarters, most of the mass of the posterior impulse response distribution is in the negative region. Overall, most of the responses in DE and FR, and to some extent also in IT, conform to the established stylised facts, whereas the responses in ES are counterintuitive for some variables.

Finally, turning to Figure 5, the bond market volatility indicator reacts negatively in the short-run to a contractionary monetary policy shock: a result similar to that shown in Bruno and Shin (2015) for the effects of a restrictive US monetary policy shock on the VIX. The EONIA response itself displays a high degree of persistence.

[16] The response in FR shows a temporary increase in loans. A disaggregated analysis that separates out loans to households and loans to firms indicates that while loans to households decline after the monetary policy shock, loans to firms tend to increase temporarily, particularly in DE and FR, as also shown in Giannone, Lenza and Reichlin (2012) for the euro area and Den Haan, Sumner and Yamashiro (2007) and Gertler and Gilchrist (1993) for the US. Giannone et al. (2012) show this effect to be mainly due to short-term loans to firms.

4.2 Are the effects of a monetary policy shock heterogeneous across countries?

In this section, we present results for various tests on the cross-country differences in the impulse responses to a monetary policy shock. Figures 6 - 11 show the median and a 68% credible interval of the posterior distribution of the difference between two countries' impulse responses to a 25bp restrictive monetary policy shock.[17] Negative values indicate that the impulse response of the variable in question in the first country lies below that in the second country, i.e. the variable falls more or rises less than in the second country. As before, we base our assessment on the location of the majority of the mass of the posterior distribution.

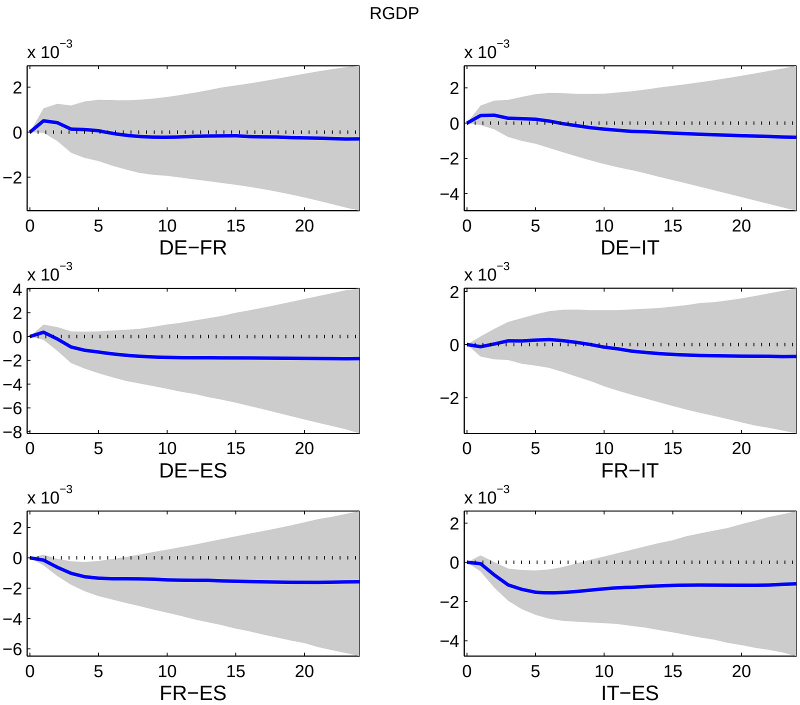

Figure 6: Cross-country differences between impulse responses to a 25bp monetary policy shock (median and 68% credible interval) – RGDP

[17] Technically, for each draw of the VAR parameters we compute the impulse responses for each variable and then the difference between the impulse responses for each pair of countries for a given variable. As this is just a function of the simulated parameter distribution of our BVAR model, the MCMC algorithm also approximates the posterior distribution of these impulse response differences.

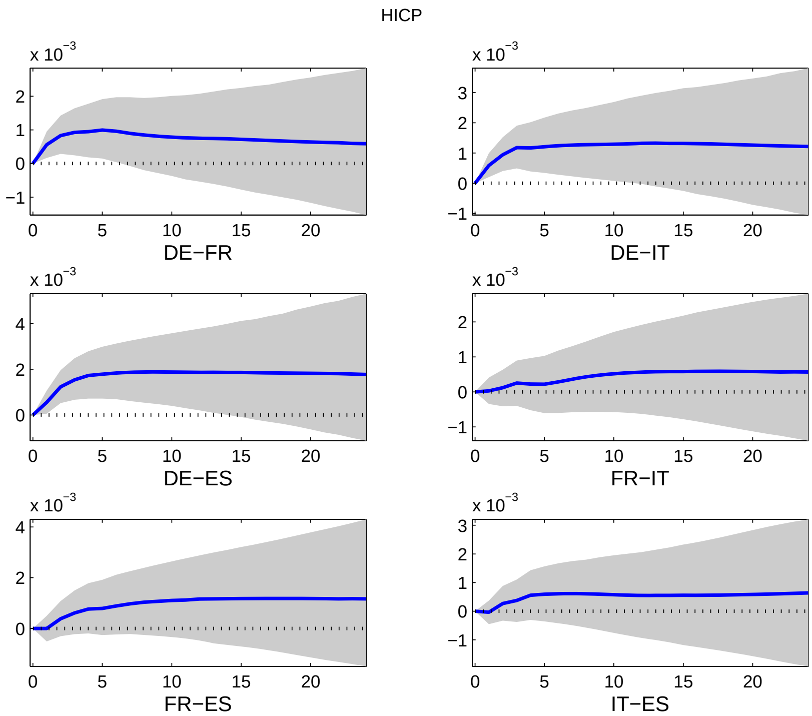

Figure 7: Cross-country differences between impulse responses to a 25bp monetary policy shock (median and 68% credible interval) – HICP

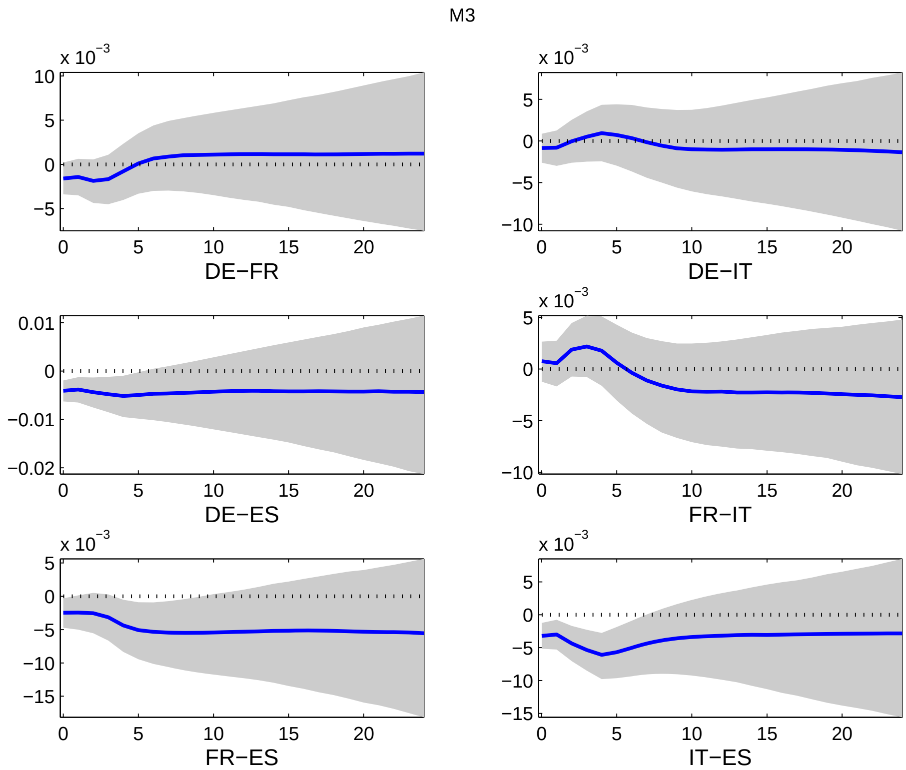

Figure 8: Cross-country differences between impulse responses to a 25bp monetary policy shock (median and 68% credible interval) – M3

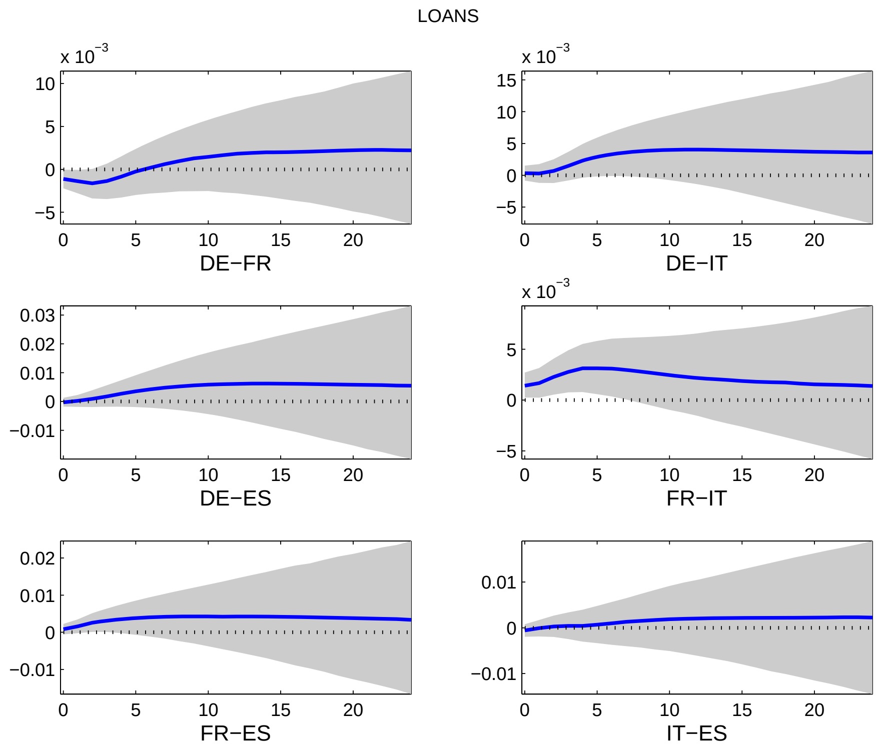

Figure 9: Cross-country differences between impulse responses to a 25bp monetary policy shock (median and 68% credible interval) – LOANS

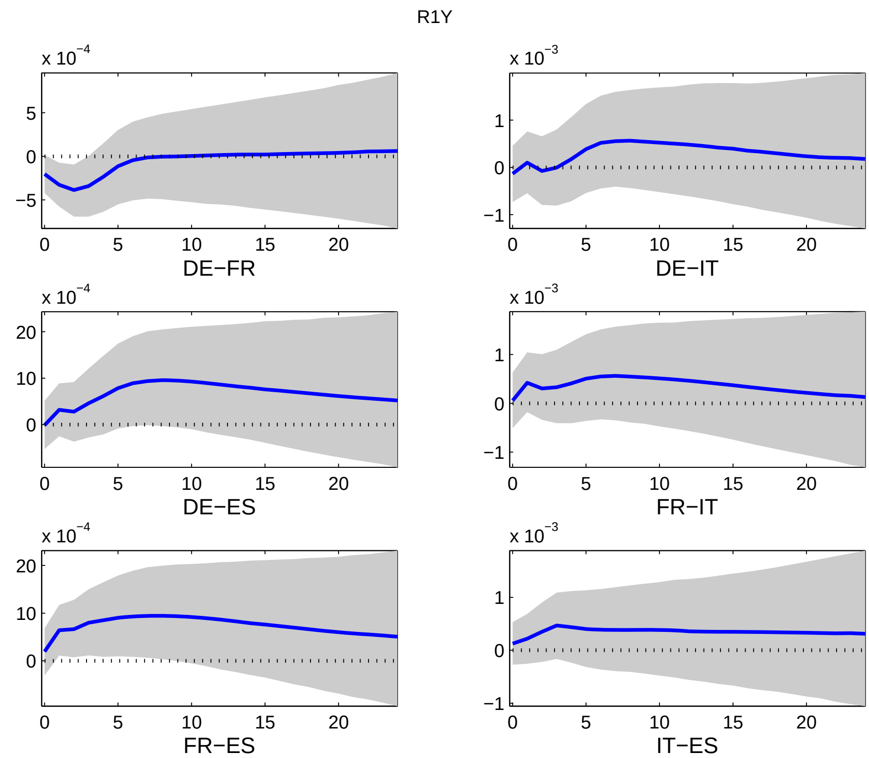

Figure 10: Cross-country differences between impulse responses to a 25bp monetary policy shock (median and 68% credible interval) – R1Y

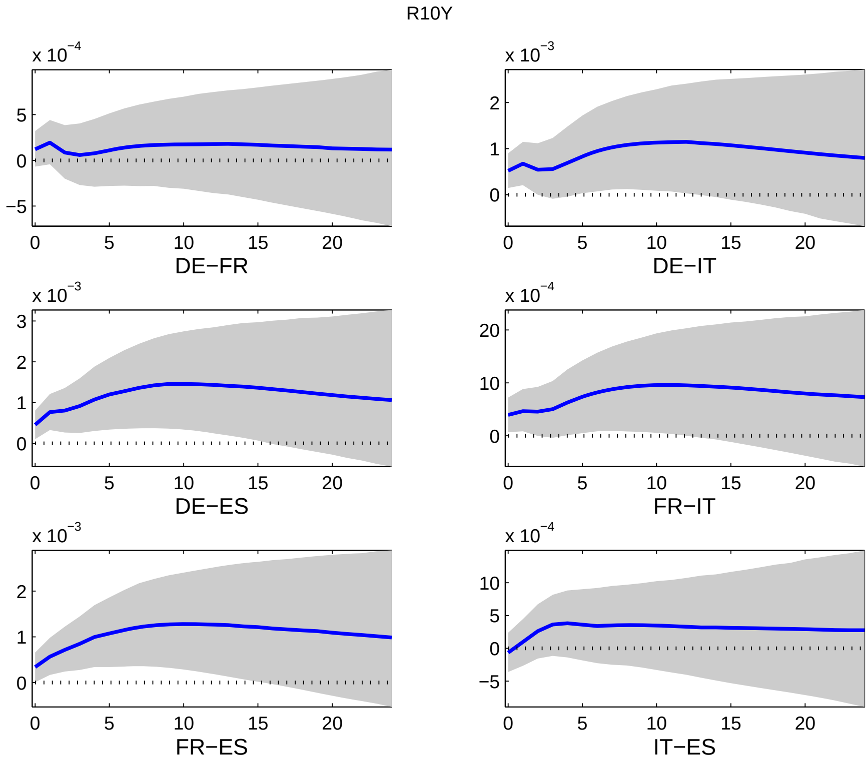

Figure 11: Cross-country differences between impulse responses to a 25bp monetary policy shock (median and 68% credible interval) – R10Y

Figure 6 shows that the real GDP response to a 25bp monetary policy shock is mainly different between ES and the other three countries, with stronger (i.e. more negative) responses in DE, FR and IT.

In contrast, for the HICP it is DE that behaves differently from the other countries (Figure 7), with the decline in the price level following a contractionary policy impulse being weaker than in the other three countries as of the quarter after the shock. There is also some evidence for the HICP falling less in FR than in ES.

For the other four variables, the estimated differences are less clear-cut. The cross-country difference in the impulse response of LOANS is shown to be positive between DE and ES/IT up to about ten to 15 quarters and between FR and ES/IT up to about ten quarters, while the difference between DE and FR is negative in the first few quarters following the shock. This reflects the negative reaction of bank lending in IT and ES (Figures 3 and 4), while the response is much weaker in DE and even positive in FR (Figures 1 and 2). Marked differences in the impulse responses of the short-term bond yield are estimated between DE and FR, DE and ES, and FR and ES, with a stronger and, relative to ES, more persistent increase in the yield in DE and FR and a smaller short-run increase in DE compared to FR. Finally, as could be expected from the impulse responses in Figures 1 – 4, for R10Y the distributions of the impulse response differences indicate pronounced deviations between DE and FR on the one hand, and IT and ES on the other, with stronger increases in the long-term bond yield in DE and FR compared to IT and ES. These differences also turn out to be very persistent.[18]

In addition to these tests on cross-country differences in the monetary transmission mechanism, we also performed tests of a classical or frequentist nature:

First, we tested for differences between the impulse response distributions for a given variable in two countries using a Kolmogorov-Smirnov test, as in Ciccarelli and Rebucci (2006).[19] The test almost universally rejects for any pair of countries that the impulse responses of real GDP, HICP, M3, loans to the non-financial private sector and the long-term bond yield are drawn from the same distributions at any horizon.[20] However, the Kolmogorov-Smirnov test is probably too strict for our purposes, which is reflected in the almost uniform rejection of the null hypothesis of identical distributions.[21] The reason for this is that it tests for the identity of the full cumulative distribution function of the variables. This includes the identity of all quantiles and all distributional statistics such as standard deviation, skewness, kurtosis etc. However, it is reasonable that a monetary policy maker might be more concerned about a subset of distributional characteristics, such as the median and a few quantiles, and not about the full distribution.

For this reason we performed two additional tests: a parametric test (t-test) on the identity of the means of the impulse response functions as in Ciccarelli and Rebucci (2006) and a non-parametric test of the identity of the median impulse response across countries.[22] The results on the differences in impulse responses provided by the one-sided versions of the tests are consistent, overall, with our results presented above. The

[18] Another way of looking at the differences in the reactions of the country-specific variables to a monetary policy shock is to use scatter plots similar to those in Cogley, Primiceri, and Sargent (2010). As it turns out, these scatter plots are more difficult to interpret than the figures based on the differences between the impulse responses across countries. The reason for this is that the high density of the dots and their overlap make it difficult to visually determine how substantial an asymmetry around the 45 degree line actually is. Therefore, we replaced these scatter plots with colour-coded contour plots of 3D histograms. Results using this approach are available upon request. The cross-country differences indicated by this exercise remain consistent with the results presented above.

[19] The Kolmogorov-Smirnov test is based on the differences between the empirical cumulative distribution functions of two random variables. For details, see, for example, Rinne (1997), pp. 550 ff.

[20] Results are available upon request.

[21] This is similar to the results in Ciccarelli and Rebucci (2006). In their Table 3, the Kolmogorov-Smirnov test rejects in the vast majority of cases, while the t-test only rejects for FR versus the other countries.

[22] The parametric test is a t-test based on the impulse responses at each horizon being normally distributed random variables (see, for example, Rinne (1997), pp. 536 ff.) and on the independence of the draws from both distributions. Neither of these assumptions might hold true in our exercise. For the test to identity the median, we compute for a given variable the difference between the impulse responses for two countries and then test whether the median of the distribution of impulse response differences is equal to zero using a Binomial distribution; see Rinne (1997), pp. 555 ff.

one-sided tests for comparing the means indicate the following rankings for the effects of a restrictive policy shock:

- RGDP (increasingly negative response): ES, IT, FR, DE,
- HICP (increasingly negative response): DE, FR, IT, ES,
- M3 (increasingly negative response): ES, IT, DE and FR
- LOANS (increasingly negative response): DE and FR, IT, ES,
- R1Y (increasingly positive response): ES, IT, DE and FR,
- R10Y (increasingly positive response): ES, IT, FR, DE.

The one-sided test for the median of the differences in impulse responses provides very similar conclusions.

Finally, as an additional indicator we follow the approach of Jarociński (2010, p. 845 ff.) and compute for each saved draw of VAR coefficients the output cost of a reduction in the price level for each country, which is defined as the average change in output up to horizon H divided by the deviation in the price level from the baseline at horizon H, i.e.

$$oc_H = \frac{\frac{1}{H}\sum_{h=1}^{H}\xi_h^{RGDP}}{\xi_H^{HICP}}, \tag{8}$$

where ξ_h^j, j = RGDP, and HICP denotes the impulse responses of variable j at horizon h to a monetary policy shock. This is the ratio of the average output loss to the reduction in the price level caused by a monetary policy shock. Figures 12 - 14 show the simulated distributions of this indicator for the four countries of interest at horizons of 4, 12, and 20 quarters.

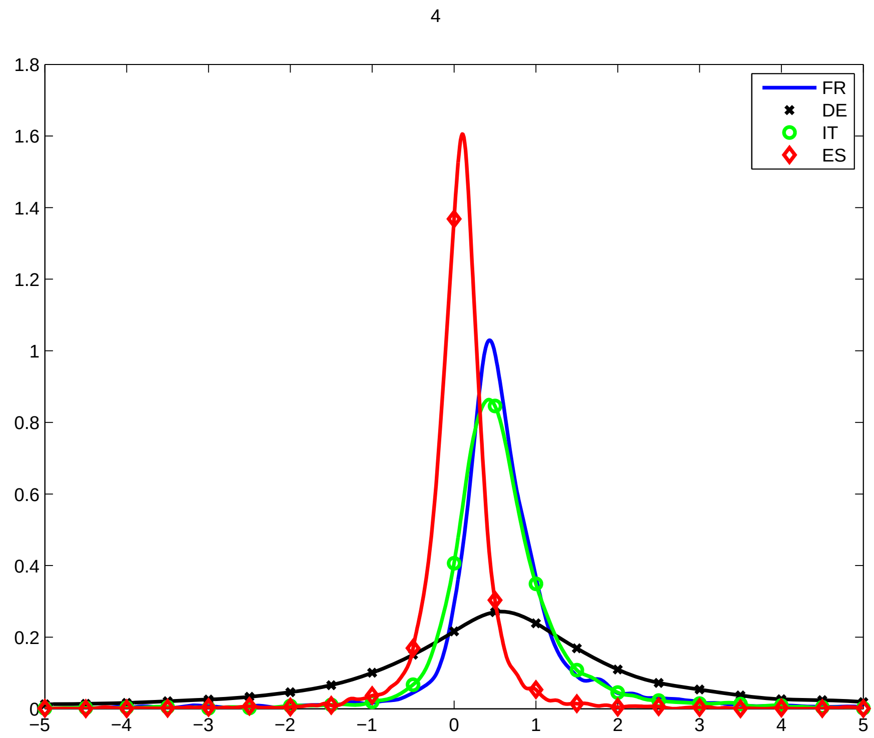

Figure 12: Posterior distribution of the output cost of price level reduction (four quarters horizon)

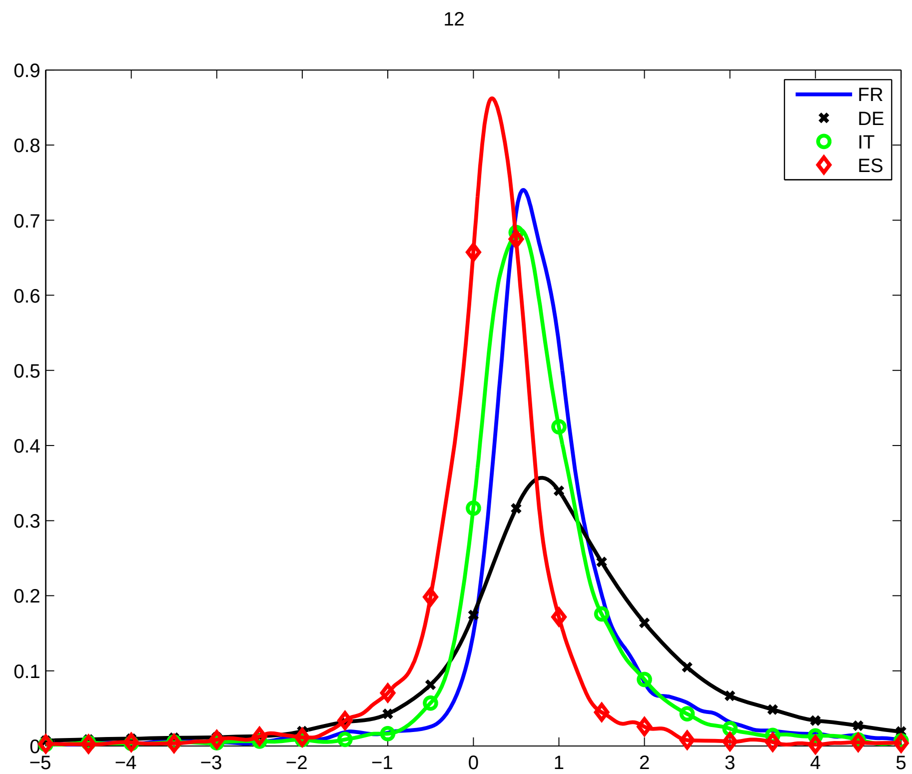

Figure 13: Posterior distribution of the output cost of price level reduction (twelve quarters horizon)

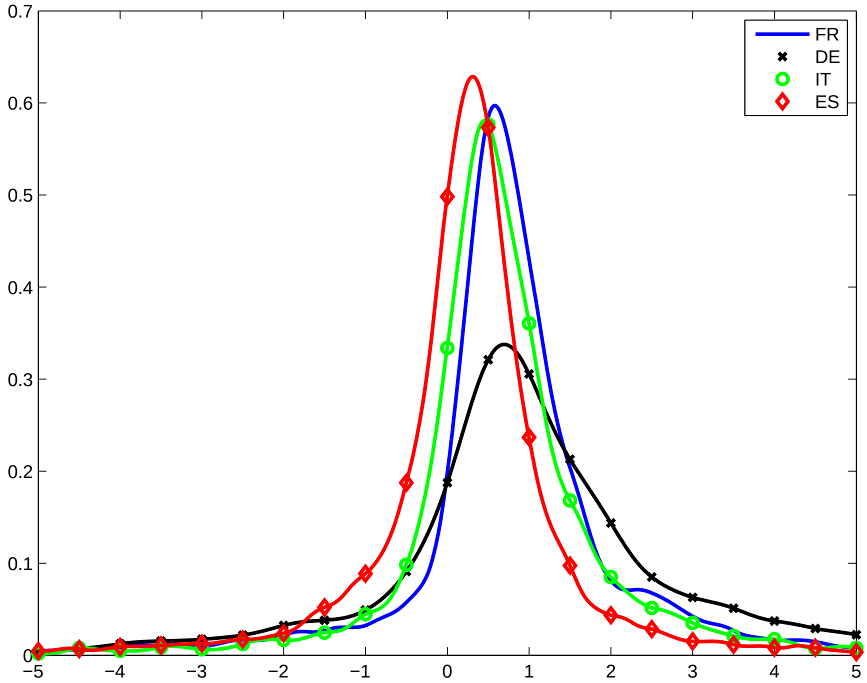

Figure 14: Posterior distribution of the output cost of price level reduction (20 quarters horizon)

The distribution for ES is furthest to the left, i.e. it implies the lowest average output loss. In fact, a substantial part of the distribution falls into the negative region, i.e. it shows the "wrong" sign. The estimated distributions of IT and FR are relatively close to each other. Overall, the distributions of ES, IT and FR converge as the horizon is extended. The distribution for DE is estimated very imprecisely at the short horizon and contracts when the horizon is lengthened. Its distribution displays positive skew, indicating substantial average output losses (in excess of the median) are more likely than small average losses (falling short of the median) for a given contraction in the price level.

4.3 Robustness checks

In addition to checking the effects of different prior specifications on our results,[23] we estimated various alternative versions of our BVAR models including fewer or different variables.[24] To reduce model size we estimated models without the bond market

[23] We estimated our model using fixed, i.e. exogenously given, hyperparameters following the specifications in Sims and Zha (1998) and Bańbura et al. (2010). This leads to narrower credible sets for the impulse responses, as the uncertainty about the hyperparameters is no longer included (Giannone et al., 2015). Our results concerning the cross-country comparisons of the impulse responses remain intact.

[24] We estimated these alternative models and derived the posterior distributions of the differences in impulse responses as in Figures 6 – 11. Results are not presented here in detail but are available upon request.

volatility indicator, with the spread between R10Y and R1Y instead of the individual variables and a "small" model containing only output, prices and the long-term bond yield for all countries together with the EONIA. We also replaced M3 by M1 or swapped real house prices, stock prices or the countries' real effective exchange rates for M3. In addition we also studied a model which included the euro-area real effective exchange rate as an additional 27th variable.[25]

For each of these models we computed the simulated distribution of the impulse responses to a monetary policy shock (identification via Choleski decomposition) and the resulting distribution of the cross-country differences.

Overall, our results, concerning both the impulse responses themselves and their differences, turn out to be quite robust across these specifications and are not strongly affected by dropping the volatility indicator, adding the euro-area real effective exchange rate, or by replacing M3 with M1, real house prices, stock prices or the real effective exchange rate. Reducing the model size by dropping M3, LOANS and R1Y or using the term spread instead of the two bond yields strengthens our results, particularly when including the spread.

In order to account for the possibility that the EONIA might not be an appropriate indicator of the monetary policy stance during the financial and sovereign debt crises because it does not reflect all the non-standard monetary policy measures employed by the European Central Bank, we re-estimated the model using the euro-area shadow short rate by Wu and Xia and obtained results very similar to the specification using the EONIA for all variables, with the cross-country differences for the bond yields becoming even more pronounced than in the baseline model.[26]

As an additional robustness check we identified the monetary policy shock through a combination of zero and sign restrictions using the algorithm of Arias et al. (2014).[27] Our results turned out to be largely robust with respect to this change, with the results concerning the impulse response differences in HICP and LOANS remaining intact and

[25] Except for a few select cases, we estimated the models using the Gibbs sampler only and kept the hyperparameters fixed at the mode of their posterior distribution in order to reduce estimation time. As our results for fixing the hyperparameters at their (model-specific) posterior mode indicate that the results remain close to the ones obtained from the full estimation, we chose this restricted estimation approach for most of the robustness tests in order to be able to check a wider range of models.

[26] The shadow short rate is a summary indicator of monetary policy that is derived from a term structure model. It is not restricted by the zero lower bound and also accounts for non-standard or unconventional monetary policy. See Wu and Xia (2016). The shadow short rate for the euro area is available at http://faculty.chicagobooth.edu/jing.wu/research/data/WX.html (downloaded on 8 December 2015).

[27] More details on the results from this exercise are available upon request. Our baseline identification imposed zero restrictions on the responses of output and prices as well as positive effects on the EONIA (normalisation) and on R1Y within the quarter in which the monetary policy shock occurs. Output and price responses were restricted to be non-positive in the following quarter.

the results for real GDP broadly robust, as well.[28] However, the cross-country differences in the impulse responses became somewhat less pronounced, which was to be expected as the sign restrictions – which are homogeneous across countries – can be interpreted as a prior against heterogeneity in the effects of monetary policy shocks on the restricted variables (at least as far as the sign of the effect is concerned).

5 Discussion and conclusions

Our results provide evidence of cross-country differences in the transmission of monetary policy shocks to macroeconomic aggregates in the euro area. We find that the response of the price level in DE is comparatively weak, as is the output response in ES. Bank lending turns out to respond more strongly in ES and IT than in DE and FR, and broad money growth tends to rise in ES, while tending to fall or to remain unresponsive in the other countries. The increase in the short-term bond yield is both more pronounced and more persistent in DE and FR than in ES. Finally, the reaction of the long-term bond yield indicates a separation into two groups, with DE/FR on one side and ES/IT on the other. Figure 15 summarises the overall evidence from the impulse responses (Figures 1 – 5), the Bayesian analysis of their differences (Figures 6 – 11) and the classical tests. The red (negative) and green (positive) arrows indicate the overall direction of the impulse responses, while the country abbreviations for each variable give the relative strength of these responses in the different countries. As already mentioned, it is important to remember that the mean/median comparisons produce a much finer ranking, since they ignore the large cross-country overlap of the posterior distributions of the impulse responses. Therefore, we suggest putting a higher weight on the results in Figures 6 – 11.

Our result of a stronger output response in Germany is consistent with the evidence in Ehrmann (2000) (using pre-euro data only) and Georgiadis (2015) (using data since the introduction of the euro). However, while these studies as well as the evidence in Ciccarelli and Rebucci (2006) indicate the output response to be weakest in FR, we find this to be the case for ES. Our results concerning the price level response also contrast with Ehrmann (2000), who shows FR and IT to display a weaker reaction than DE and ES. While the difference to Georgiadis (2015) can probably be attributed to the different empirical methodology, the other two analyses are also based on other sample periods.

Our results concerning the cross-country ranking of the impulse responses for bond yields, broad money and output are consistent with a reduction in money demand due to

[28] In principle, the ordering in Figure 15 remained intact.

increased opportunity cost and a stronger reduction in output being caused by the more pronounced increase in the longer-term bond yield, e.g. through a traditional interest-rate channel. The fact that bank lending actually contracts less for countries which experience stronger increases in bond yields and a more pronounced contraction in broad money and output is difficult to explain. One possible explanation supported by a disaggregated analysis of lending to firms and households is that bank lending to firms actually expands after a restrictive policy shock (Giannone et al., 2012, Gertler and Gilchrist 1995) and that this effect might partly compensate for the reduction in lending to households. Gertler and Gilchrist (1993) argue that the positive effect of a restrictive policy shock on bank lending reflects an increase in firms' demand for loans due to declining revenues caused by weakening economic activity. Den Haan et al. (2007) suggest as an explanation that banks rebalance their portfolios towards business loans following a monetary policy contraction. Reconciling the fact that prices respond less in countries with a stronger output response requires differently sloped aggregate supply curves (flatter in DE, steeper in ES).

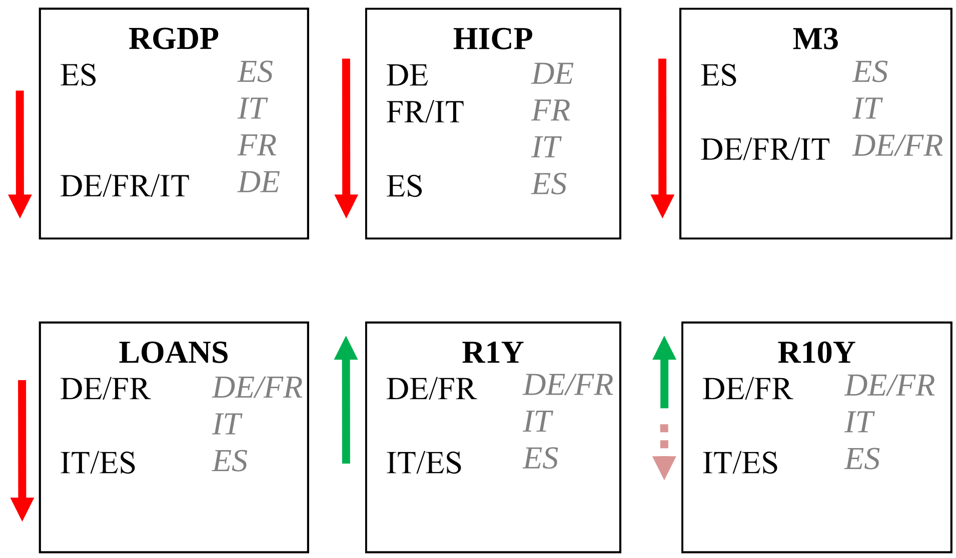

Note: Arrows indicate direction of impulse response to a contractionary monetary policy shock. In each block, the LHS column indicates the summary ranking based on the Bayesian credible intervals for the impulse response differences. The RHS column (in italics and grey) indicates the ranking based on the tests for the differences in mean and median.

Figure 15: Overall summary of relative effects of a monetary policy shock

It remains unclear as to whether our results provide a case against monetary policy responding to euro-area average or aggregate variables.[29] Although our results indicate some heterogeneity in the effects of monetary policy, monetary policy shocks are a very minor source of fluctuations in output and prices, which is a common result in the VAR literature. The relevant question therefore is whether the estimated differences in the effects of monetary policy have implications for the Eurosystem's ability to stabilise (output and) prices in the different countries and whether different effects of these stabilisation policies translate into different effects on welfare. This, however, not only depends on differences in monetary transmission, but also on the symmetry in the shocks affecting the economies and on possible differences in the effects of these shocks.[30]

References

Altavilla, Carlo, Domenico Giannone, and Michele Lenza (2014), The financial and macroeconomic effects of OMT announcement, ECB Working Paper Series, No. 1707.

Angelini, Paolo, Paolo Del Giovane, Stefano Siviero, and Daniele Terlizzese (2002), Monetary policy rules for the Euro Area: what role for national information?, Temi di Discussione, Number 467, Banca d'Italia.

Arias, Jonas, Juan Rubio-Ramírez, and Daniel Waggoner (2014), Inference based on SVARs identified with zero and sign restrictions: theory and applications, Dynare Working Paper Series, No. 30.

Bańbura, Marta, Domenico Giannone, and Lucrezia Reichlin (2010), Large Bayesian vector autoregressions, Journal of Applied Econometrics, 25(1), 71-92.

Barigozzi, Matteo, Antonio Conti and Matteo Luciani (2014), Do euro area countries respond asymmetrically to the common monetary policy?, Oxford Bulletin of Economics and Statistics, 76(5), 693-714.

[29] Apart from a well-founded welfare criterion, this would also require a model in which we could experiment with different policy reaction functions and try to derive an optimal (or restricted optimal) policy function. Since our model is not structural, the other equations will not be invariant to changes in the policy equation and therefore such experiments are not possible.

[30] For example, if all countries were hit by an inflationary shock but the size of the shock or its inflationary effects were exactly proportional to the effects of monetary policy on prices (smallest in DE, largest in ES), the heterogeneous responsiveness of prices to a monetary policy tightening in DE might not have any serious implications for the ability of the Eurosystem to stabilise its constituent economies.

Benigno, Pierpaolo (2004), Optimal monetary policy in a currency area, Journal of International Economics, 63(2), 293-320.

Berben, Robert-Paul, Alberto Locarno, Julian Morgan, and Javier Vallés (2004), Cross-country differences in monetary transmission, ECB Working Paper Series, No. 0400.

Boivin, Jean, Marc Giannoni, and Benoît Mojon (2009), How has the Euro changed monetary transmission?, in Daron Acemoglu, Kenneth Rogoff, and Michael Woodford (eds.), NBER Macroeconomics Annual 2008, 23, University of Chicago Press, Chicago, 77-125.

Bruno, Valentina and Hyun Song Shin (2015), Capital flows and the risk-taking channel of monetary policy, Journal of Monetary Economcis, 71, 119-132.

Cavallo, Antonella and Antonio Ribba (2014), Common macroeconomic shocks and business cycle fluctuations in euro area countries, International Review of Economics and Finance, 38(C), 377-392.

Christiano, Lawrence, Martin Eichenbaum, and Charles Evans (1999), Monetary policy shocks: what have we learned and to what end?, in John Taylor and Michael Woodford (eds.), *Handbook of Macroeconomics*, Vol. 1A, Elsevier, Amsterdam etc., 65-148.

Ciccarelli, Matteo and Alessandro Rebucci (2006), Has the transmission mechanism of European monetary policy changed?, European Economic Review, 50(3), 737-776.

Ciccarelli, Matteo, Angela Maddaloni, and José-Luis Peydró (2013), Heterogeneous transmission mechanism: monetary policy and financial fragility in the euro area, Economic Policy, 28(75), 459-512.

Cogley, Timothy, Giorgio Primiceri, and Thomas Sargent (2010), Inflation-gap persistence in the US, American Economic Journal: Macroeconomics, 2(1), 1875-1911.

Dedola, Luca and Francesco Lippi (2005), The monetary transmission mechanism: Evidence from the industries of five OECD countries, European Economic Review, 49(6), 1543-1569.

De Grauwe, Paul (2000), Monetary policy in the presence of asymmetries, Journal of Common Market Studies, 38(4), 593-612.

De Grauwe, Paul and Tomasz Piskorski (2001), Union-wide aggregates versus national data based monetary policies: does it matter for the Eurosystem, CEPR Discussion Paper Series, Discussion Paper No. 3036.

Den Haan, Wouter, Steven Sumner and Guy Yamashiro (2007), Bank loan portfolios and the monetary transmission mechanism, Journal of Monetary Economics, 54, 904-924.

Doan, Thomas, Robert Litterman and Chris Sims (1984), Forecasting and conditional projection using realistic prior distributions, Econometric Reviews, 3, 1-100.

Ehrmann, Michael (2000), Comparing monetary policy transmission across European countries, Weltwirtschaftliches Archiv, 136(1), 58-83.

Georgiadis, Georgios (2015), Examining asymmetries in the transmission of monetary policy in the euro area: Evidence from a mixed cross-section global VAR model, European Economic Review, 75, 195-215.

Gertler, Mark and Simon Gilchrist (1993), The role of credit imperfections in the monetary transmission mechanism: arguments and evidence, Scandinavian Journal of Economics, 95(1), 43-64.

Geweke, John (1992), Evaluating the accuracy of sampling-based approaches to the calculation of posterior moments, in: Jose, Bernardo, James Berger, Philip Dawid and Adrian Smith (eds), Bayesian Statistics 4, Oxford: Oxford University Press, 169-193.

Giannone, Domenico, Michele Lenza and Lucrezia Reichlin (2012), Money, credit, monetary policy and the business cycle in the Euro area, ECARES Working Paper 2012-008.

Giannone, Domenico, Michele Lenza and Giorgio Primiceri (2015), Prior selection for vector autoregressions, Review of Economics and Statistics, 97(2), 412-435.

Guiso, Luigi, Anil Kashyap, Fabio Panetta, and Daniele Terlizzese (2000), Will a common European monetary policy have asymmetric effects?, Temi di Discussione, Number 384, Banca d'Italia.

Jarociński, Marek (2010), Responses to monetary policy shocks in the East and West of Europe: a comparison, Journal of Applied Econometrics, 25(5), 833-868.

Litterman, Robert (1980), A Bayesian procedure for forecasting with vector autoregression, Working paper, Massachusetts Institute of Technology, Department of Economics.

Mihov, Ilian and Andrew Scott (2001), Monetary policy implementation and transmission in the European Monetary Union, Economic Policy, 16(33), 369-406.

Mojon, Benoît and Gert Peersman (2001), A VAR description of the effects of monetary policy in the individual countries of the euro area, European Central Bank Working Paper Series, No. 92.

Primiceri, Giorgio (2005), Time varying structural vector autoregressions and monetary policy, Review of Economic Studies, 72(3), 821-852.

Rafiq, M. Sohrab and Sushanta Mallick (2008), The effect of monetary policy on output in the EMU3: A sign restriction approach, Journal of Macroeconomics, 30(4), 1756-1791.

Rinne, Horst (1997), Taschenbuch der Statistik, 2nd ed., Verlag Harri Deutsch, Thun und Frankfurt am Main.

Sims, Christopher and Tao Zha (1998), Bayesian methods for dynamic multivariate models, International Economic Review, 29(4), 949-968.

Soares, Carla (2008), Heterogeneity in a monetary union and its impact on welfare, Economic Bulletin, Banco de Portugal, Autumn, 157-173.

Wu, Cynthia and Dora Xia (2016), Measuring the macroeconomic impact of monetary policy at the zero lower bound, Journal of Money, Credit, and Banking, forthcoming.

Appendix A: Data and data transformations

Country-specific series:

- real GDP (REALGDP or RGDP): working day and seasonally adjusted, chain linked volumes (rebased), source: ECB, ESA 2010 national accounts, main aggregates. Series extended backwards from 1995Q1 for ES and IT using OECD data based on ESA 2005.[31]
- HICP: Euro area (changing composition), overall index, monthly index, Eurostat, data neither seasonally nor working day adjusted.
- M3: national contributions: index of notional stocks, MFIs, central government and post office giro institutions reporting sector, all currencies combined, Euro area (changing composition) counterpart, non-MFIs excluding central government sector, denominated in euro, data neither seasonally nor working day adjusted, source: ECB Statistical Data Warehouse (SDW).
- loans (LOANS): Index of notional stocks (own construction) of MFI loans to private households and to non-financial corporations (loans of national MFIs to euro-area residents), from 1991 to 1996 outstanding amounts at the end of the period (stocks), from 1997 to 2009 index of notional stocks constructed from flows adjusted for reclassifications and revaluations, after 2009 constructed from flows adjusted for reclassifications, revaluations sales and securitisations, data neither seasonally nor working day adjusted, source: SDW.
- R1Y: Government bond yield (one year maturity, two years maturity for DE[32]), quarterly averages of daily data (DE, FR IT) or of monthly data (ES), sources: Bundesbank (DE, FR, IT), IMF IFS (ES).
- R10Y: Yield of 10-year government benchmark securities, quarterly averages of monthly data, source: OECD main economic indicators.
- real house prices (RHP): source: international house price database, Federal Reserve Bank of Dallas.
- stock prices (STOCK): source: IMF IFS.
- real effective exchange rate (REER): Indicator of price competitiveness compared to selected industrial countries, based on the total sales deflators. Bundesbank internal data series (ES, FR, IT).

[31] We use the seasonally adjusted series because the ESA 2005 data we use for extending the time series for ES and IT backwards is only available as seasonally adjusted data.
[32] The available IMF IFS series for one year government bond yields for Germany breaks off in 2007Q3.

Euro area series:

- EONIA: quarterly average of monthly values, before 1999 overnight interbank rate for DE, source: OECD main economic indicators and IMF IFS.
- bond market volatility index: euro area (changing composition), Eurex Generic 1st `RX` Future - implied bond volatility, end of period, source: SDW.
- real effective exchange rate (REEREA): ECB Real effective exch. rate producer prices deflated, euro area changing composition vis-à-vis the EER-12 group of trading partners (AU, CA, DK, HK, JP, NO, SG, KR, SE, CH, GB and US) against the euro, source: SDW.

Real GDP, HICP, M3, bank loans, real house prices and stock prices are transformed into 4xlog levels, while the interest rate variables and the bond market volatility are used in raw form, i.e. as decimal numbers.[33] The ordering we use for the Choleski decomposition in the baseline model is the same as in Altavilla et al. (2014): $RGDP_{FR}$, $HICP_{FR}$, $RGDP_{DE}$, $HICP_{DE}$, $RGDP_{IT}$, $HICP_{IT}$, $RGDP_{ES}$, $HICP_{ES}$, EONIA, $M3_{FR}$, $LOANS_{FR}$, $R1Y_{FR}$, $R10Y_{FR}$, $M3_{DE}$, $LOANS_{DE}$, $R2Y_{DE}$, $R10Y_{DE}$, $M3_{IT}$, $LOANS_{IT}$, $R1Y_{IT}$, $R10Y_{IT}$, $M3_{ES}$, $LOANS_{ES}$, $R1Y_{ES}$, $R10Y_{ES}$, VOLA.

[33] GLP argue that this transformation of variables results in an approximately identical scale, which is important since the hyperprior for Ψ is not scale invariant. It should be noted that this transformation also results in a scaling of the impulse responses, which we scale back to (simple) log levels before graphing.

Appendix B: Convergence diagnostics

Figures B1 and B2 present two statistics on the convergence of the MCMC algorithm for our baseline model. The first statistic (Figure B1) was proposed by Geweke (1992). It compares means of subsamples of the MCMC recursion. Let Ξ represent the full Markov chain after burn-in and $\xi^{(j)}$, j = 1,…,n denote the individual draws of the parameter vector. The first subsample consists of the first n_1 draws (10% of the chain after burn-in), the second subsample of the last n_2 draws (60% of the chain). The corresponding sample means are given by $\hat{\mu}_1$ and $\hat{\mu}_2$. The test statistic is

$$CD = \frac{\hat{\mu}_1 - \hat{\mu}_2}{\sqrt{\hat{\sigma}_1^2 / n_1 - \hat{\sigma}_2^2 / n_2}},$$

where

$$\hat{\mu}_1 = \frac{1}{n_1}\sum_{j=1}^{n_1} \xi^{(j)} \qquad \hat{\mu}_2 = \frac{1}{n_2}\sum_{j=n+1-n_2}^{n} \xi^{(j)}$$

The variances $\hat{\sigma}_1$ and $\hat{\sigma}_2$ are estimated by taking autocorrelation into account. The CD statistic converges in distribution to a standard normal distribution, i.e. a 5% critical value is equal to 1.96. At the 5% significance level, the test statistic in Figure B1 rejects convergence for about 6% of the VAR coefficients in Γ and about 6% of the elements of the VAR covariance matrix Σ.

Figure B2 displays the inefficiency factor. This statistic compares the standard error (se) of the sample mean $\overline{\xi}$ under possible correlation with the standard error under the hypothesis of independence ($\rho_j = 0$). ρ_j measures the correlation between two draws, j iterations apart.

$$IEF(\theta) = \frac{se^2(\overline{\xi})}{se^2(\overline{\xi}, \rho_j = 0)} = 1 + 2\sum_{j=1}^{n-1}\left(1 - \frac{j}{n}\right)\rho_j$$

In the ideal case of no correlation, the IEF statistic equals one. Values in the range of 1-5 indicate efficiency, but values in the range of 6-20 are still considered to be acceptable, especially for draws from the VAR covariance matrix (Primiceri, 2005). Using this metric there is only a possible convergence issue for one single element in Σ, specifically in the residual variance of the bond market volatility indicator. Thus, overall, both statistics indicate MCMC algorithm convergence.

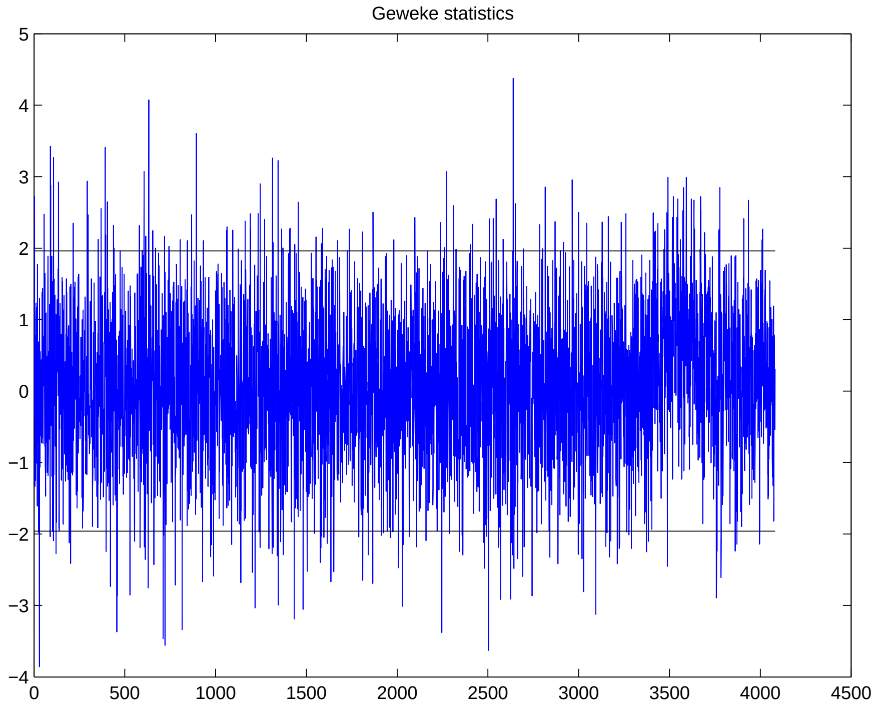

Figure B1: Geweke statistic

Note: The horizontal axis shows the VAR parameters with the first 3,406 parameters representing the VAR coefficients and the remaining 676 parameters representing the elements of the covariance matrix.

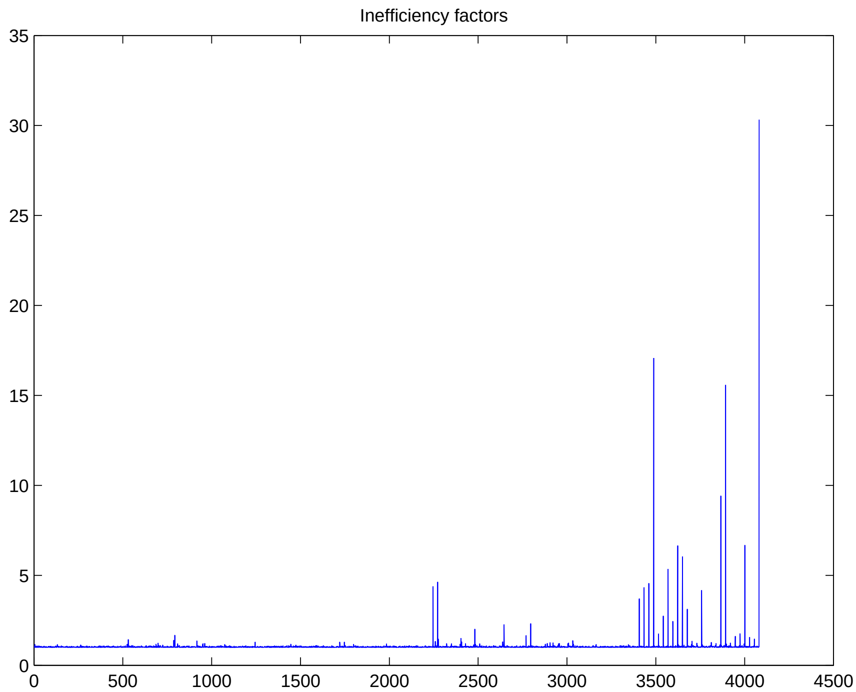

Figure B2: Inefficiency factors

Note: The horizontal axis shows the VAR parameters with the first 3,406 parameters representing the VAR coefficients and the remaining 676 parameters representing the elements of the covariance matrix.

Since our "objects of interest" are the impulse responses of the model variables following a monetary policy shock, we also tested the impulse response functions (impact quarter up to 24 quarters) for convergence, as it is theoretically possible that either instabilities in the VAR parameters might offset each other and might have only moderate effects on the impulse responses or that they might reinforce themselves and create more instabilities in the impulse responses. Our baseline estimation results, overall, in 2.3% non-converging impulse responses (from overall 650 = 25 (quarters) x 26 variables) according to the Geweke-statistic at the 5% level and a zero non-converging impulse response function according to the inefficiency factor using a critical value of 20.[34]

[34] The impulse responses with possible convergence issues almost all apply to the Spanish long-term bond yield.

Appendix C: Endogenous hyperparameter selection

Figure C1 shows graphs of the marginal posterior densities for three important hyperparameters: λ (overall shrinkage), μ (sum-of-coefficients prior), and θ (dummy-initial-observation prior). The posterior distributions turn out to be much tighter than the priors, indicating that the data is informative about the hyperparameters.[35] The posterior mode for λ is 0.7058 which is much higher than the values (for fixed hyperparameters) suggested in Sims and Zha (1998) (λ = 0.2) or in Bańbura et al. (2010) (λ = 0.1, for a model with 20 variables and about 40 years of data)[36]. In fact, the posterior assigns almost zero probabilities to these values. Thus, the posterior distribution of the λ parameter implies considerably less shrinkage than the conventional fixed hyperparameters. In contrast, the posterior distributions for θ and μ with modes of 0.5477 and 0.2406, respectively, imply values below those usually chosen in the literature, e.g. $\theta = \mu = 1$ in Sims and Zha (1998). The comparatively low values for these two parameters show that the posterior distribution of the hyperparameters puts more weight on the sum-of-coefficients priors and the dummy-initial-observation (common stochastic trend) priors than the standard choices in the literature using fixed hyperparameters.

[35] The graphs of the marginal posterior densities result from applying a normal kernel density estimator to the drawn distribution of each hyperparameter.
[36] The hyperparameters' values at the posterior mode are determined by numerical maximisation.

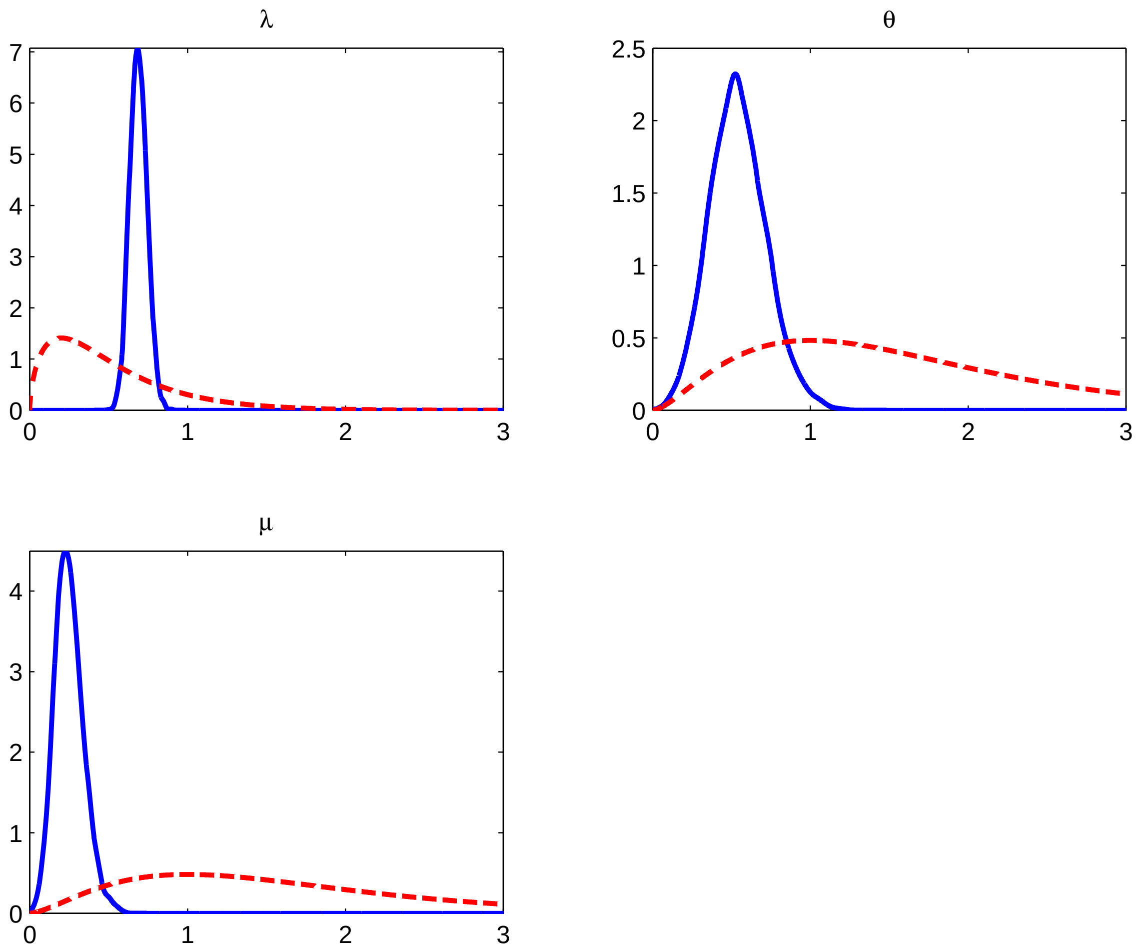

Figure C1: Prior (red, dashed) and simulated marginal posterior (blue, solid) densities of hyperparameters