

**Analysing and forecasting price dynamics across
euro area countries and sectors: A panel VAR approach**

by

Stéphane DÉES
Jochen GÜNTNER*)

Working Paper No. 1410

September 2014

Analysing and forecasting price dynamics across euro area countries and sectors: A panel VAR approach*

Stéphane Déés[†]

Jochen Güntner

European Central Bank

Johannes Kepler University Linz

September 1, 2014

Abstract

This paper uses a panel VAR (PVAR) approach to estimating, analysing and forecasting price dynamics in four different sectors – industry, services, construction, and agriculture – across the four largest euro area economies – Germany, France, Italy and Spain – and the euro area as a whole. By modelling prices together with real activity, employment and wages, we can disentangle the role of unit labour costs and profit margins as the factors affecting price pressures on the supply side. In out-of-sample forecast exercises, the PVAR model fares comparatively well against common alternatives, although short-horizon forecast errors tend to be large when we consider only the period of the recent financial crisis. The second part of the paper focuses on Spain, for which prediction errors during the crisis are particularly large. Given that its economy faced dramatic sectoral changes due to the burst of a housing bubble, we use the PVAR model for studying the transmission of shocks originating from the Spanish construction sector to other sectors. In a multi-country extension of the model, we also allow for spillovers to the other euro area countries in our sample.

Keywords: Cost pressures, forecasting, impulse response analysis, panel VAR models.

JEL Classification: C33, C53, E31, E37.

*The views expressed in this paper are those of the authors and do not necessarily reflect the views of the European Central Bank (ECB) or the Eurosystem. We thank Thomas Westermann, Elke Hahn, an anonymous referee of the ECB Working Paper Series and colleagues from the ECB for helpful comments.

[†]Corresponding author: Stéphane Déés, European Central Bank, Kaiserstrasse 29, 60311 Frankfurt am Main, Germany. Telephone: (+49)(0)69 1344 8784. Email: stephane.dees@ecb.int.

1 Introduction

While forecasting price inflation is fundamental for private sector and policy decision-making, it has always been a challenging exercise. Inflation forecasting can either be purely subjective or require more or less sophisticated techniques. Faust and Wright (2013) review the state of the art in inflation forecasting and state “an explosion in the number and variety of methods in recent years”, ranging from traditional time series approaches to more structural economic models, such as single Phillips Curve equations or complete Dynamic Stochastic General Equilibrium models. Using an extensive set of possible predictors has also gained popularity in recent years – e.g. methods based on factor-augmented vector autoregression (FAVAR) models, as proposed by Bernanke et al. (2005) –, as have methods based on financial market indicators, which extract forward-looking information about expected future inflation.

Owing to instability in inflation forecasting, however, it has often proved difficult to outperform even very simple models, such as a random walk (see, e.g., Atkeson and Ohanian, 2001). The stability of inflation dynamics during the Great Recession has also called into question the usefulness of fundamentals-based approaches, such as Phillips-Curve equations, in predicting inflation dynamics (see Ball and Mazumder, 2011; Bassetto et al., 2013).

Rather than from continuously increasing the degree of complexity of forecasting techniques, forecast accuracy might benefit from the informational content of disaggregated data. Since the disaggregated contain at least as much information as the aggregated time series, increasing the information set on which forecasts are based could improve the accuracy of out-of-sample (OOS) forecasts, at least theoretically. Since inflation and other macroeconomic variables represent contemporaneous aggregates, it seems plausible that the use of disaggregated data facilitates an increase in forecast accuracy (compare Luetkepohl, 1984b). Further, disaggregated information can be helpful in retrieving common drivers of the aggregated series. Finally, forecast errors of disaggregated components might partially cancel out (compare Theil, 1954).

In this paper, price dynamics are analysed from the supply side of the economy. We consider disaggregation along the sectoral dimension for the four largest euro area economies – Germany, France, Italy and Spain – and the euro area as a whole. The aim is to provide a model for estimating, analysing and forecasting price dynamics in four different sectors – industry, services, construction and agriculture.

The supply-side sectoral approach in this paper has several important advantages. First, we

can test whether combining disaggregated forecasts based on sectoral data is superior to methods based on economy-wide time series.¹ Second, a disaggregated approach provides information about sector-level price dynamics in the euro area, which allows for a close monitoring of disaggregated prices. This close-up perspective of price dynamics has received increasing attention since the global financial crisis, given the role of relative prices in the build-up of macroeconomic imbalances. Third, as our approach aims at modelling prices together with real activity, wages and employment, it also enables us to disentangle the factors affecting price dynamics from the cost side by accounting for pressures coming from both labour costs and profit margins.² Figure 1 reveals that there is nontrivial heterogeneity across countries regarding the sources of price pressures. In the decade preceding the crisis, for instance, unit labour costs rose strongly in Spain, while they were comparatively subdued in Germany. Finally, encompassing the supply side yields forecasts of additional variables, such as unit labour costs and profit margins, which are crucial for assessing future developments in the competitiveness between sectors or countries and for anticipating the investment decisions of firms in the near future, respectively.

Our approach relies on estimating VAR models that can account for the potential static and dynamic interdependencies between the variables and sectors of interest as well as for the contemporaneous and lagged influence of exogenous driving forces, such as fluctuations in world demand or oil prices. Given the size and complexity of the system, we face two main issues. On the one hand, estimating separate sector-specific models would be relatively parsimonious in terms of the number of coefficients, while it ignores any interdependencies between sectors. On the other hand, a large-scale VAR model of the entire economy would quickly run into the curse of dimensionality. With $N = 4$ sectors, $K = 4$ variables, $p = 2$ lags, and an intercept, we would have to estimate $N \cdot K \cdot p + 1 = 33$ parameters per equation, even when ignoring the possible influence of exogenous variables. Since our sector-level data is available from 1995Q1 only, there is little hope of obtaining precise coefficient estimates based on 70 observations per variable; in particular, if the model is supposed to serve in constructing forecasts.

As a consequence, a suitable shrinkage method is required in order to reduce the parameter

¹Similar “bottom-up” approaches from the supply side in a data-rich environment have been used for forecasting real GDP growth by Drechsel and Scheufele (2012) for Germany, Barhoumi et al. (2012) for France and Hahn and Skudelny (2008) for the euro area. Beck et al. (2011) also study euro area inflation from both a sectoral and a country perspective and find that the sectoral as well as the country-specific component of inflation help explain euro area inflation dynamics.

²Maurin et al. (2011) model profit dynamics in the four largest euro area countries (Germany, France, Italy and Spain) and the euro area as a whole, considering three main sectors (manufacturing, construction and services) in each economy, based on a vector autoregression (VAR) approach.




Figure 1: GDP deflators, unit labour costs and profit margins (*left column*) and sectoral value added deflators (*right column*)

space of the model, while preserving the possibility of static and dynamic interdependencies between different sectors. Due to the fact that the number of observation units ($N = 4$ sectors) is small relative to the number of observation periods, a panel vector autoregression (PVAR)




Figure 1 continued: GDP deflators, unit labour costs and profit margins (*left column*) and sectoral value added deflators (*right column*)

approach seems well suited for the task at hand. Given the limited data availability, we abstain from estimating a time-varying parameter model, as proposed in Canova and Ciccarelli (2009), although we might thus miss some of the variation in the interdependencies between variables, sectors and countries, e.g. due to structural changes.³

Forecasting economy-wide variables based on the PVAR approach requires contemporaneous aggregation of the respective sector-level forecasts. Tiao and Guttman (1980), Kohn (1982), Luetkepohl (1984a), and others show that aggregating forecasts is generally preferable to forecasting the aggregates directly, if the data-generating process (DGP) is known in terms of its

³See, e.g., Canova et al. (2012) for recent evidence of variations in European and national real business cycles over time, based on aggregate macroeconomic time series.

order and coefficients.⁴ In practice, however, the DGP is rarely known and a trade-off arises with respect to forecast accuracy. Luetkepohl (1984b) shows that, even if the order and coefficients are consistently estimated, the information gain from using the disaggregated time series might be more than offset by the higher specification and estimation errors of a less parsimoniously parameterized process, especially at long forecast horizons. As the sample size increases, the MSPE component due to specification and estimation uncertainty becomes sufficiently small, and the forecast based on the disaggregated multivariate process is again more accurate than directly forecasting the aggregate. However, these conclusions are based on asymptotic theory. Given that both asymptotic and small-sample simulation results depend on the DGP of the multiple time series,⁵ relative forecast accuracy ultimately remains an empirical question.

Forecasting euro area-wide time series entails at least two dimensions of contemporaneous aggregation. Forecasts of macroeconomic variables can be aggregated across countries (see, e.g., Marcellino et al., 2003) and across subcomponents (see, e.g., Hubrich, 2005). Using monthly data from 1992.1 to 2001.12, Hubrich (2005) finds that direct forecasts of euro area HICP are often more accurate than aggregating component forecasts, indicating higher estimation and specification error at horizons above 6 months. In particular, contemporaneous aggregation seems to increase rather than reduce bias, if unexpected events, such as the surge in unprocessed food and energy prices in 2000, affect components in the same direction. Benalal et al. (2004) explore both dimensions simultaneously, selecting the best model in terms of OOS MSPE from a wide class of uni- and multivariate alternatives for five components and overall HICP for the euro area and its four largest member countries. Regarding the aggregation of HICP components, indirect forecasts perform better for the euro area at short horizons, while direct forecasts are favourable at longer horizons and at the country level. Both Hubrich (2005) and Benalal et al. (2004) find that, for euro area “core inflation”, i.e. HICP excluding non-processed food and energy prices, empirical evidence is more favourable for aggregating component forecasts. Aggregating country-specific forecasts, in turn, is generally less accurate than forecasting euro area inflation, although the differences are small in magnitude and based on a “synthetic” euro area consisting of only its four largest members.

⁴If the disaggregated time series are approximately uncorrelated and have similar stochastic structures, there is no information gain from using a multivariate model of the disaggregated variables (see, e.g., Luetkepohl, 1984b, 2006), and the mean squared prediction errors (MSPEs) will be identical. Luetkepohl (2006) also considers aggregating univariate forecasts of the individual components, which is, however, weakly inferior to aggregating forecasts based on a multivariate model, if the process is known, while it may be more or less accurate than forecasting the aggregate time series directly.

⁵See Luetkepohl (1984b; 1987) and Hendry and Hubrich (2011).

The rest of the paper is structured as follows. Section 2 contains a general introduction to PVAR models, based on Canova and Ciccarelli (2013). Section 3 discusses model specification and estimation for the euro area and its four largest member countries. Section 4 presents the forecasting performance of the PVAR models. Section 5 illustrates how the model facilitates the analysis of cross-sector linkages within an economy by studying the effects of shocks originating from the Spanish construction sector. A model extension to account for spillovers between euro area countries is also proposed. Section 6 concludes.

2 The PVAR Model

Suppose there is a cross-section of N macroeconomic observation units (e.g. countries, regions, sectors,...), which are inherently linked to each other, and that for each unit i , a set of K macroeconomic variables of interest is observed over time.

One possibility to simultaneously account for the interdependencies between the variables *within one unit* as well as *between units* is by estimating the following large-scale VAR(p) model:

$$Y_t = \nu + A_1 Y_{t-1} + \dots + A_p Y_{t-p} + e_t, \quad (1)$$

where Y_t is an $(N \cdot K \times 1)$ vector of endogenous variables, ν is an $(N \cdot K \times 1)$ vector of intercepts, A_j , $j = 1, \dots, p$ are $(N \cdot K \times N \cdot K)$ matrices of slope coefficients, and $e_t \sim iid(0, \Sigma_e)$ is an $(N \cdot K \times 1)$ vector of possibly contemporaneously correlated reduced-form disturbances.⁶ Note that the VAR representation in (1) does not exploit the panel structure of the data, i.e. the fact that the $N \cdot K$ variables under consideration correspond to only K distinct variables observed for each of the N units.

A panel VAR has the same structure as a standard VAR model, i.e. each endogenous variable is assumed to depend on lagged values of itself and of all other endogenous variables. However, the representation also accounts for the cross-sectional dimension in the data. Let $y_{i,t}$ denote the $(K \times 1)$ vector of endogenous variables for unit i and $Y_t = (y'_{1,t}, y'_{2,t}, \dots, y'_{N,t})'$ denote the $(N \cdot K \times 1)$ vector of stacked $y_{i,t}$, $i = 1, \dots, N$. We can then write the PVAR model equation by equation as

$$y_{i,t} = \nu_i + A_{1,i} Y_{t-1} + \dots + A_{p,i} Y_{t-p} + e_{i,t}, \quad i = 1, \dots, N \quad (2)$$

⁶Note that the coefficient matrices A_j must be absolutely summable for a moving average representation of (1) to exist. This can be ensured, e.g., by taking first differences of the endogenous variables.

where ν_i is a $(K \times 1)$ vector of intercepts, $A_{j,i}$, $j = 1, \dots, p$, $i = 1, \dots, N$, are $(K \times N \cdot K)$ matrices of slope coefficients, and $e_{i,t}$ is a $(K \times 1)$ vector of possibly contemporaneously correlated reduced-form disturbances.

Suppose the variables in Y_t might also depend on an $(M \times 1)$ vector of *weakly exogenous* variables, such as world demand and oil prices, which are assumed to be independent of contemporaneous or lagged fluctuations in Y_t . If we assume that these variables follow a VAR(p^x), the panel VAR with exogenous driving forces (PVARX) can be written as

$$y_{i,t} = \nu_{1i} + \sum_{l=1}^p A_{l,i} Y_{t-l} + \sum_{l'=0}^q B_{l',i} X_{t-l'} + e_{1i,t} \quad (3)$$

$$X_t = \nu_2 + \sum_{l=1}^{p^x} C_l X_{t-l} + e_{2,t}, \quad (4)$$

where $B_{l',i}$, $l' = 0, \dots, q$ are $(N \cdot K \times M)$ matrices of exogenous coefficients and $e_{1i,t}$ and $e_{2,t}$ are assumed to be uncorrelated. Note that the vector of weakly exogenous variables is the *same for all* units i and that the latter might depend on the former contemporaneously, i.e. with lag 0.

Following the terminology in Canova and Ciccarelli (2013), the PVARX representation in (3) and (4) can account for (i) “dynamic interdependencies”, since p lags of all endogenous variables of all units enter the model for unit i ; (ii) “static interdependencies”, since the $e_{1i,t}$ are generally correlated across units i ; (iii) “cross-sectional heterogeneity”, since the intercept, the slope coefficients, and the variance of $e_{1i,t}$ are generally unit-specific.

In this regard, the PVARX is very similar to the large-scale VAR model in (1), augmented by a set of exogenous variables. As a consequence, *unrestricted estimation* of the model in (3) and (4) faces exactly the same curse of dimensionality. Including an intercept, each equation contains $G = N \cdot K \cdot p + M \cdot (q + 1) + 1$ unknown coefficients, while the total number of unknown coefficients amounts to $N \cdot K \cdot G$.⁷

This problem could be solved by selectively modelling the dynamic interdependencies between some units, while imposing zero restrictions on others, or by grouping units and assuming that the interdependencies only exist within but not across groups (compare Canova and Ciccarelli, 2012). Instead, we proceed by exploiting the panel structure of the data. Canova and Ciccarelli (2004, 2009) propose the use of *cross-sectional* shrinkage methods in order to deal with the curse of dimensionality.

⁷For $N = 4$ units, $K = 4$ endogenous variables, $M = 4$ exogenous variables, and lag order $p = q = 2$, e.g., this corresponds to $G = 45$ coefficients per equation and a total of $N \cdot K \cdot G = 720$ coefficients.

In the following, we neglect the law of motion of the exogenous variables in (4), effectively assuming that they are *strictly exogenous*. This is possible even in forecasting exercises, where we can condition on available projections of world demand, oil prices, and others.

Following Canova and Ciccarelli (2013), we start by writing (3) in simultaneous equations format:

$$Y_t = Z_t \cdot \delta + e_t, \quad (5)$$

where Y_t and e_t have been defined before, $Z_t = I_{N \cdot K} \otimes (I, Y'_{t-1}, \dots, Y'_{t-p}, X'_t, \dots, X'_{t-q})$, $\delta = (\delta'_1, \dots, \delta'_N)'$, and δ'_i are $(K \cdot G \times 1)$ vectors containing stacked the *rows* of the coefficient matrices

$$[\nu_{1,i}, A_{1,i}, \dots, A_{p,i}, B_{0,i}, \dots, B_{q,i}].$$

Allowing all coefficients to vary between cross-sectional units prevents a meaningful unrestricted estimation of the $(N \cdot K \cdot G \times 1)$ coefficient vector δ .

Suppose that we are not interested in modelling all the details of δ but rather in robust parameter estimates for impulse response analysis and forecasting. Assume further that δ can be factorized as a linear combination of a *lower-dimensional vector* θ , e.g.

$$\delta = \Xi_1 \theta_1 + \Xi_2 \theta_2 + \Xi_3 \theta_3 + \Xi_4 \theta_4 + \dots + u_t, \quad (6)$$

where the Ξ_i , $i = 1, 2, \dots$, are matrices of dimension $(N \cdot K \cdot G \times d_i)$, respectively, and θ_i are the corresponding *mutually orthogonal* factors, which determine the entries in δ . Here, θ_1 might capture unit-specific components ($d_1 = N$), θ_2 endogenous variable-specific components ($d_2 = K$), θ_3 endogenous lag-specific components ($d_3 = p$), and θ_4 exogenous variable-specific components ($d_4 = M$), while u_t absorbs any remaining idiosyncratic noise in the unrestricted coefficient vector.

The obvious advantage of factoring δ as in (6) is a substantial reduction in the dimensionality of the parameter space. In the above example, we must now estimate $N + K + p + M$ instead of $N \cdot K \cdot G$ unrestricted coefficients. In other words (compare Canova and Ciccarelli, 2012), the factorization transforms a large-scale PVARX model into a parsimonious seemingly unrelated regressions (SUR) model, such that we can rewrite (5) with the help of (6) as

$$Y_t = \sum_{i=1}^r \mathcal{Z}_{i,t} \theta_i + v_t, \quad (7)$$

where $\mathcal{Z}_{i,t} = Z_t \Xi_i$, $i = 1, 2, \dots$, captures, e.g., unit-specific, endogenous variable-specific, endogenous lag-specific, and exogenous variable-specific information in the data, and $v_t = e_t + Z_t u_t$. By construction, $\mathcal{Z}_{i,t}$ has a slow moving average structure that captures low frequency movements in the data, which is a convenient feature in OOS forecasting.

Economically, equation (7) decomposes the fluctuations in the endogenous variables in Y_t into mutually orthogonal components. In the above example, one can think of $\mathcal{Z}_{1,t}\theta_1$, $\mathcal{Z}_{2,t}\theta_2$, $\mathcal{Z}_{3,t}\theta_3$, and $\mathcal{Z}_{4,t}\theta_4$, as unit-specific, endogenous variable-specific, endogenous lag-specific, and exogenous variable-specific indicators, respectively (compare Canova and Ciccarelli, 2012).

3 Modelling the Largest Euro Area Countries and Sectors

3.1 Data

The PVAR model is composed of a set of VARX models for the main economic sectors of the euro area and the four largest euro area countries, i.e. the industrial, construction, services and agricultural sectors. Each of the sectoral models includes four endogenous variables: the basic price GDP deflator (or value added deflator), real value added, compensation and employment (measured by the number of employees) for the respective sector. Implicit in this set of variables are unit labour costs and profit margins as the two cost components of the GDP deflator. The sectoral data in quarterly frequency are available from Eurostat from the year 2000 onwards based on the Nace 2 classification for economic sectors and are backdated for the purpose of this analysis to the year 1995 on the basis of the previously available Nace 1 classification.

The economy-wide data and forecasts are derived from the sectoral models by aggregating the data and forecasts of the sectoral variables. This amounts to a simple sum for employment, whereas, due to chain linking, the total real value added data are no longer perfectly additive. In the latter case, the aggregates are obtained by weighting the growth rates of the chain-linked series with the shares obtained from the series on value added “at basic prices in previous year prices”, which is available from 2001Q1 onwards. For aggregating prior observations as well as the forecasts, weights based on the chain-linked series are applied, which, in any case, proved to provide similar results to those based on the data in previous year prices for the period starting in 2001Q1.

3.2 Model Specification

We are interested in the analysis of four sectors with four variables per sector, implying $N = 4$ and $K = 4$. Moreover, the model includes $M = 4$ exogenous variables: *world demand*, *oil prices*, *short-term interest rates*, and the *effective exchange rate*. All corresponding time series are taken from the Eurosystem macroeconomic projections.⁸

Estimation of the PVARX model requires choosing several parameter values such as the lag order of endogenous and exogenous variables. Although standard lag order-selection criteria can be applied, we set $p = q = 2$ rather arbitrarily in the following illustration. Note that $q = 2$ implies that the model accounts for the influence of the exogenous variables *contemporaneously* as well as at lags of one and two quarters.

The model's specification crucially depends on how the unrestricted coefficient vector δ is factorized.⁹ In line with Canova and Ciccarelli (2009), we assume that the factorization in (6) is exact, i.e. $u_t = 0$, and hence $v_t = e_t$ in (7). A convenient implication of this assumption is that we can estimate θ and thus δ consistently by multivariate least squares (MLS). In contrast to Canova and Ciccarelli (2013), we abstract from the possibility of time-varying coefficients, due to the relatively short sample period and because we are primarily interested in the forecasting properties of the tool.

Throughout Section 4, we use four factors and an equation-specific intercept in order to shrink the parameter space of δ :¹⁰

1. The $(N \times 1)$ vector θ_1 captures *sector-specific* components in the *endogenous* variables.
2. The $(K \times 1)$ vector θ_2 captures *variable-specific* components in the *endogenous* variables.
3. The $(M \times 1)$ vector θ_3 captures *variable-specific* components in the *exogenous* variables.
4. The $(q \times 1)$ vector θ_4 captures *lag-specific* components in the *exogenous* variables.¹¹

The PVAR procedure also permits combining the information in endogenous and exogenous variables, e.g. in a (1×1) vector/scalar of *common* components. However, the corresponding $(N \cdot K \cdot G \times 1)$ dimensional regressor Ξ_{common} will often be a linear combination of Ξ_i , $i = 1, \dots, 4$,

⁸For more details on the Eurosystem staff macroeconomic projection exercises, the interested reader is referred to <http://www.ecb.europa.eu/pub/pdf/other/staffprojectionsguideen.pdf>

⁹A large variety of factors can be chosen, each capturing the information in a certain subset of endogenous and exogenous variables, respectively. See, e.g., Canova and Ciccarelli (2012, p. 20) for an example of how the Ξ_i and $\mathcal{Z}_{i,t}$ are constructed.

¹⁰The intercept is not factorized for obvious reasons. Alternatively, we could demean all time series.

¹¹We drop the contemporaneous lag category of exogenous variables to avoid collinearity between θ_3 and θ_4 .

inducing collinearity between the factors. As a consequence, there might be *no unique solution* to the least squares minimization problem.

Stacking the T observation of Y_t , Z_t , and e_t in the $(T \cdot N \cdot K \times 1)$ vectors Y , Z , and e , respectively, we can rewrite equation (7) as

$$Y = Z\theta + e. \quad (8)$$

It is now straightforward to obtain the MLS estimate $\hat{\theta} = (Z'Z)^{-1}(Z'Y)$ and to transform it into $\hat{\delta} = \Xi \cdot \hat{\theta}$, which in turn allows us to compute $\hat{e}_t = Y_t - Z_t\hat{\delta}$, $t = 1, \dots, T$, and $\hat{\Sigma}_e = \frac{\hat{e}_t\hat{e}_t'}{T - (N + K + M + q + N \cdot K)}$.

To facilitate the analysis of impulse response functions, forecasting, and inference, we convert $\hat{\delta}$ back to $[\hat{\nu}, \hat{A}_1, \dots, \hat{A}_p, \hat{B}_0, \dots, \hat{B}_q]$ and construct the associated companion matrix

$$\hat{\mathbf{A}} = \begin{bmatrix} \hat{A}_1 & \hat{A}_2 & \dots & \hat{A}_{p-1} & \hat{A}_p \\ I_{N \cdot K} & 0 & \dots & 0 & 0 \\ 0 & I_{N \cdot K} & \dots & 0 & 0 \\ \vdots & \vdots & \ddots & 0 & \vdots \\ 0 & 0 & \dots & I_{N \cdot K} & 0 \end{bmatrix}. \quad (9)$$

3.3 Model Selection and Estimation

Given the different possible ways to shrink the parameter space of δ , the benchmark model above could be challenged by many alternative specifications of the PVARX model. Obvious degrees of freedom are the lag order p and q of endogenous and exogenous variables, respectively, which are of secondary interest here. Instead, we focus on the implications of alternative factorizations of the unrestricted coefficient vector δ for the model's in-sample fit.

More precisely, we compare the *in-sample fit* of our benchmark specification of θ with several alternative specifications. For this purpose, we use the maximum log likelihood (MLL), the Akaike information criterion (AIC), and the Schwarz information criterion (SIC).¹² Note that the optimal model must *maximise* the MLL and *minimise* the information criteria.

¹²Assuming normality of the error terms, the conditional ML estimator for VAR models coincides with the multivariate LS estimator (see, e.g., L pp. 87), while

$$\tilde{\Sigma}_e = \frac{T - (\#parameters)}{T} \cdot \hat{\Sigma}_e. \quad (10)$$

In the benchmark PVARX model, e.g., $\#parameters = \underbrace{N}_{\theta_1} + \underbrace{K}_{\theta_2} + \underbrace{M}_{\theta_3} + \underbrace{q}_{\theta_4}$.

Table 1: Model selection criteria for alternative factorizations of the parameter vector δ

		Benchmark	Model A	Model B	Model C	Model D	Model E
Euro Area	MLL	-477.32 ³	-474.57 ²	-558.08	-492.61	-493.55	-469.53 ¹
	AIC	-8.86 ³	-8.91 ²	-7.01	-8.53	-8.51	-8.91 ¹
	SIC	-7.89 ²	-7.91 ¹	-6.56	-7.70	-7.67	-7.75 ³
Germany	MLL	-532.72 ²	-533.14 ³	-582.92	-538.45	-547.97	-523.66 ¹
	AIC	-6.80 ²	-6.76 ³	-5.80	-6.75	-6.47	-6.89 ¹
	SIC	-5.82 ¹	-5.74 ²	-5.34	-5.90	-5.62	-5.72 ³
France	MLL	-369.21 ²	-355.56 ¹	-448.68	-391.46	-387.52	-371.11 ³
	AIC	-11.61 ²	-11.98 ¹	-9.74	-11.07	-11.19	-11.38 ³
	SIC	-10.63 ²	-10.97 ¹	-9.29	-10.22	-10.34	-10.20 ³
Italy	MLL	-421.81 ²	-421.94 ³	-456.03	-429.00	-431.72	-412.27 ¹
	AIC	-10.06 ²	-10.03 ³	-9.53	-9.97	-9.89	-10.17 ¹
	SIC	-9.08 ¹	-9.02 ²	-9.07	-9.12	-9.04	-8.99 ³
Spain	MLL	-466.78 ²	-468.44 ³	-509.31	-488.15	-480.43	-456.69 ¹
	AIC	-8.74 ²	-8.66 ³	-7.96	-8.23	-8.46	-8.86 ¹
	SIC	-7.76 ¹	-7.65 ³	-7.50	-7.38	-7.61	-7.69 ²
# parameters	θ	30	31	14	26	26	36
# parameters	δ	720	720	720	720	720	720

Notes: Each entry reports the value of the in-sample selection criterion in row, based on the respective model in column. Superscript indices rank models according to their in-sample fit.

Table 1 reports the MLL, AIC, and SIC for our benchmark model and five alternative factorizations of δ . Model A *adds a common component* of endogenous variables to the benchmark specification. Model B is identical to the benchmark specification *without an intercept*. Model C is identical to the benchmark specification *without the sector-specific endogenous components*. Model D is identical to the benchmark specification *without the variable-specific endogenous components*. Model E is identical to the benchmark specification *with variable- & lag-specific exogenous components*, i.e. θ_3 and θ_4 are replaced by a single $(M \cdot q \times 1)$ vector.

The results in Table 1 suggest that it is difficult to beat the benchmark factorization in terms of its in-sample fit. Only Model E, which replaces the separate variable-specific and lag-specific exogenous factors by a common variable- and lag-specific exogenous factor outperforms the benchmark model according to two of the three criteria for all countries except France. The corresponding values of the MLL and AIC indicate that modelling joint variable- and lag-specific components reduces the unexplained variance in the sample, while the SIC penalizes for the fact that the specification thus becomes less parsimonious.¹³ Model A, which includes a common endogenous factor, is in the ballpark, whereas Models B, C, and D perform significantly worse than our benchmark model according to all three criteria.

3.4 Component Indicators

As mentioned in Section 2, we can interpret the $\mathcal{Z}_{i,t}\theta_i$, $i = 1, 2, \dots$, as indicators for the relative importance of the corresponding components for fluctuations in the endogenous variables, since the time series have been standardised prior to the estimation.

Figures 2 to 4 plot the $N = 4$ unit-specific indicators of the endogenous variables, $\mathcal{Z}_{1,t}\theta_1$, the $K = 4$ variable-specific indicators of the endogenous variables, $\mathcal{Z}_{2,t}\theta_2$, and the $M = 4$ variable-specific indicators of the exogenous variables $\mathcal{Z}_{3,t}\theta_3$, respectively, for the euro area as a whole.¹⁴ Note that the $q = 2$ lag-specific indicators of the exogenous variables are not shown, because they have little economic interpretation.

In Figure 2, all sector-specific indicators, except the one for agriculture, display a pronounced drop in late 2008. The drop in the construction-related indicator is more gradual and persistent, reflecting substantial cross-country heterogeneity in this sector within the euro area. The variable-specific indicators in Figure 3 illustrate important labour market adjustments during

¹³Recall that the net change in the number of parameters is $M \cdot (q + 1) - M - q > 0 \quad \forall M, q$.

¹⁴Similar indicators are available for each of the four euro area countries upon request.




Figure 2: Time series of euro area *sector-specific* component indicators of endogenous variables




Figure 3: Time series of *endogenous variable-specific* component indicators for the euro area

the crisis as well as the unprecedented drop in real VA in 2009 and the renewed downturn in 2012. Note that the pronounced drop in real activity was not reflected in a similar reduction in price and wage levels, as measured by the indicators of VAD and compensation per employee, pointing to the presence of downward rigidities in price and wage setting.

Finally, world demand was the dominant *exogenous* driver of fluctuations in Y_t both before and during the crisis. Figure 4 also suggests the role of the crisis-related drop in oil prices




Figure 4: Time series of *exogenous variable-specific* component indicators for the euro area

together with the changes in interest rates as important factors in driving the fluctuations in the endogenous variables.

4 Out-Of-Sample Forecast Performance

The in-sample fit of a particular specification of the PVAR model is not necessarily indicative of its out-of-sample forecast performance. It is well-known that model selection based on in-sample criteria might be subject to “over-fitting”, i.e., a less parsimonious specification is preferred, although this might worsen the model’s out-of-sample forecast accuracy. For this reason, we continue by conducting a recursive *pseudo out-of-sample* (POOS) forecast exercise.¹⁵

First, the model is estimated on the initial estimation period, i.e. the first R quarters of the sample. Then, the h -quarter-ahead forecast of the endogenous variables is constructed based on the estimated PVAR coefficients, the lagged observations of the endogenous and the *realized* contemporaneous observations – hence the term “pseudo” – of the exogenous variables. As in Benalal et al. (2004), the conditional forecast for period $t + h$ is formed recursively based on an expanding estimation window $[1, \dots, t]$ until the end of the sample in 2012Q3. Assuming an independent multivariate white noise process for the error term ν_t in equation (7), the

¹⁵A genuine out-of-sample exercise based on real-time data would have been more satisfactory. While including some of the series of interest for the euro area as a whole, however, the ECB Real Time Database does not satisfy our data requirements to a sufficient extent.

conditional expectation given past observations is an optimal h -quarter-ahead forecast of Y_t .¹⁶ For stationary processes, the forecast will be unbiased, and forecast intervals of bounded length can be constructed using asymptotic or bootstrap methods.

Subtracting the point forecast for period $t + h$ from the realized data for the same period, we can compute the mean squared prediction error (MSPE) for each sector-specific variable. By aggregating the sector-specific POOS forecasts for period $t + h$ and subtracting the result from the realized aggregate data, we obtain the corresponding economy-wide MSPE of the four endogenous variables.

Table 2 reports the corresponding MSPE for aggregate total real VA and the VAD for a forecast horizon of $h = 1$ and an initial estimation period of $R = 30$ quarters.¹⁷ As we are interested in the cost components of the price dynamics, we also compute the implied forecasts for unit labour costs (ULC) and profit margins (PMA) based on the forecasts for the total number of employees, compensation per employee and real value added.¹⁸

The ranking of models according to their MSPE implies some striking differences relative to the ranking according to the in-sample criteria. In particular, Model B, which is identical to the benchmark model without the equation-specific intercepts, performs best in terms of the VA and VAD MSPE, although it was outperformed by most alternative models in terms of its in-sample fit. Accordingly, including an equation-specific intercept, which equals the average growth rate of the corresponding variable, deteriorates the POOS forecast performance of the PVAR model *in the very short run*.

While the results in Table 2 reveal that the PVAR model under consideration is prone to over-fitting, this should be taken with a grain of salt. On the one hand, due to the short sample period 1995Q2–2012Q3, the Great Recession accounts for a significant part of the evaluation period and dominates thus the reported MSPEs, while it is not clear whether putting too much weight on an exceptional event of this kind is desirable. On the other hand, the alternatives with an equation-specific intercept generally outperform Model B at longer forecast horizons, i.e. for $h > 1$, as the economy returns to its long-run equilibrium growth path.

In line with evidence in Benalal et al. (2004), OOS forecasts for Italy are less accurate than for the three other countries and the euro area as a whole, regardless of the forecast horizon.

¹⁶If the shocks are not independent but uncorrelated, the conditional expectation remains the best *linear* forecast, but may not be the best in a larger class containing nonlinear functions (compare Luetkepohl, 2006).

¹⁷For the sake of brevity, the results for alternative forecast horizons and the sector-specific MSPEs are only available upon request.

¹⁸ $ULC = \frac{\text{Number of Employees} \cdot \text{Compensation per employee}}{\text{Real Value Added}}$ and $PMA = VAD - ULC$.

Table 2: Mean squared prediction errors of economy-wide aggregate variables for alternative factorizations of the parameter vector δ

		Benchmark	Model A	Model B	Model C	Model D	Model E
Euro Area	real VA	0.4327 ³	0.4396	0.3864 ¹	0.4182 ²	0.4468	0.4371
	VAD	0.0378 ²	0.0400 ³	0.0420	0.0366 ¹	0.0464	0.0401
	ULC	0.2771	0.2723 ²	0.2512 ¹	0.2757 ³	0.2891	0.2794
	PMA	0.2487	0.2498	0.2154 ¹	0.2395 ²	0.2411 ³	0.2472
Germany	real VA	1.0838 ²	1.1187	1.1018	1.0842 ³	1.0567 ¹	1.0931
	VAD	0.1540	0.1586	0.1500 ³	0.1352 ¹	0.1751	0.1376 ²
	ULC	0.8426 ³	0.8697	0.7546 ¹	0.8329 ²	0.8865	0.8490
	PMA	1.0153	1.0526	1.0337	0.9789 ²	1.0142 ³	0.9645 ¹
France	real VA	0.2543 ³	0.2537 ²	0.2716	0.2595	0.2756	0.2461 ¹
	VAD	0.0542 ³	0.0494 ¹	0.0546	0.0502 ²	0.0690	0.0571
	ULC	0.1441	0.1450 ²	0.1431 ¹	0.1410 ³	0.1594	0.1456
	PMA	0.1140 ³	0.1230	0.1099 ²	0.1193	0.1073 ¹	0.1107
Italy	real VA	0.6953 ³	0.7014	0.5873 ¹	0.7044	0.6373 ²	0.7813
	VAD	0.5815 ³	0.5900	0.6170	0.5783 ²	0.6356	0.5104 ¹
	ULC	2.1832 ³	2.1931	2.3036	2.1831 ²	2.1581 ¹	2.3694
	PMA	1.2447	1.2471	1.1886 ²	1.2326 ³	1.1130 ¹	1.3142
Spain	real VA	0.1947 ³	0.1944 ²	0.1826	0.2101	0.2580	0.1917 ¹
	VAD	0.3998	0.4155	0.3671 ¹	0.4810	0.3679 ²	0.3762 ³
	ULC	0.7402	0.7458	0.6542 ²	0.8607	0.7031 ³	0.6316 ¹
	PMA	1.1816	1.2082	1.1043 ²	1.2370	0.9548 ¹	1.1556 ³
# parameters	θ	30	31	14	26	26	36
# parameters	δ	720	720	720	720	720	720

Notes: Each entry reports the MSPE from a recursive pseudo OOS forecast exercise with initial estimation period 1995Q1-2002Q3, based on the respective PVAR specification.

Superscript indices rank forecasts according to their MSPE.

We also conducted another POOS forecast exercise, in which we compare the benchmark PVAR model with three popular alternatives, including a random walk with drift (RW), a univariate autoregressive (AR) and a multivariate autoregressive (VAR) process.¹⁹ Like the PVAR model, the AR model forecasts the four endogenous variables at the sector level before aggregating the univariate forecasts, whereas the VAR model forecasts the economy-wide aggregates directly, using a single multivariate process.

Similarly to Table 2, Table 3 reports the corresponding MSPE for aggregate total real VA, the VAD, ULC, and profit margins, respectively, for a forecast horizon of $h = 1$, for the euro area and the four countries. In this horse-race, the PVAR model ranks first in 9 and second in 8 out of 20 cases, while it never ranks last. When considering all the different horizons up to 12 quarters ahead, the PVAR model ranks first in 35% of the cases and second in 31% of the cases. Among the variables, it is interesting to mention that the PVAR model ranks first in 65% of the cases for the VAD. Contrary to the results of Atkeson and Ohanian (2001), the random walk does not outperform the PVAR model, except for Germany.

While the performance of the PVAR model in this POOS forecast exercise is satisfactory, it is important to emphasize that the presence of the financial crisis and the Great Recession in the evaluation period may distort the results. To take into account the role of this exceptional event, Table 4 reports the MSPEs for the crisis relative to those for the benchmark evaluation period for all countries and the euro area as a whole. Each number corresponds to the ratio between the MSPE for variable k and horizon h based on the benchmark model evaluated over 2008Q2-2012Q3 and the corresponding MSPE evaluated over 2002Q4-2012Q3. Accordingly, a ratio larger than one indicates a deterioration in forecast accuracy during the Great Recession relative to the benchmark evaluation period.

The results show that the PVAR model generally performs worse during the crisis, especially at shorter forecast horizons. At horizons of 6 and 12 quarters however, the MSPE ratios for the euro area, Germany, France, and Italy decrease below one, indicating more precise forecasts, on average. In contrast, those for Spain remain above unity, with only a single exception, and tend to increase rather than decrease.

A possible shortcoming of our model is that it assumes constant parameters throughout the estimation and forecast periods. Although the PVAR coefficients are re-estimated in each round

¹⁹See, e.g., Gardner (1985) and Marcellino et al. (2003) for the satisfactory performance of univariate models in OOS forecasts.

Table 3: Mean squared prediction errors of economy-wide aggregate variables for alternative forecasting models

		Benchmark PVAR	Random Walk	AR	Aggr. VAR
Euro Area	real VA	0.4327 ³	0.5692	0.4227 ²	0.2959 ¹
	VAD	0.0378 ¹	0.0446 ³	0.0431 ²	0.0764
	ULC	0.2771 ¹	0.2871 ²	0.3968	0.2931 ³
	PMA	0.2487 ²	0.2307 ¹	0.3325	0.3063 ³
Germany	real VA	1.0838 ²	1.0393 ¹	1.0986 ³	1.1752
	VAD	0.1540 ³	0.0847 ¹	0.1024 ²	0.3209
	ULC	0.8426 ²	0.8670 ¹	1.1070	0.8423 ¹
	PMA	1.0153 ²	0.8704 ¹	1.1815	1.1098 ³
France	real VA	0.2543 ²	0.3722	0.2818 ³	0.1429 ¹
	VAD	0.0542 ¹	0.0690 ³	0.0622 ²	0.1623
	ULC	0.1441 ¹	0.1510 ²	0.1971	0.1656 ³
	PMA	0.1140 ¹	0.1283 ²	0.1695	0.1375 ³
Italy	real VA	0.6954 ³	0.7606	0.6559 ²	0.3603 ¹
	VAD	0.5815 ¹	0.6586 ²	0.7595 ³	0.9790
	ULC	2.1832 ²	2.0396 ¹	2.2311 ³	4.1452
	PMA	1.2447 ²	1.1201 ¹	1.2513 ³	2.5487
Spain	real VA	0.1947 ¹	0.6128	0.3379 ²	0.3494 ³
	VAD	0.3998 ¹	0.5682 ³	0.4279 ²	0.6236
	ULC	0.7402 ¹	1.1250	0.8998 ²	0.9518 ³
	PMA	1.1816 ²	0.9755 ¹	1.4856	1.4464 ³

Notes: Each entry reports the MSPE from a recursive pseudo OOS forecast exercise with initial estimation period 1995Q1-2002Q3, based on the respective model. Superscript indices rank forecasts according to their MSPE.

Table 4: Mean squared prediction errors of economy-wide aggregate variables for the crisis relative to those for the benchmark evaluation period

	Horizon	1	3	6	12
Euro Area	real VA	1.87	1.81	0.25	0.32
	VAD	1.09	1.04	1.12	1.09
	ULC	1.67	1.58	0.58	0.24
	PMA	1.77	1.63	0.46	0.31
Germany	real VA	1.82	1.96	0.48	0.14
	VAD	1.01	1.68	1.75	0.62
	ULC	1.91	1.84	0.31	0.46
	PMA	1.75	1.77	0.44	0.40
France	real VA	1.74	1.69	0.74	1.66
	VAD	0.91	0.90	0.85	3.01
	ULC	1.06	1.03	0.67	2.11
	PMA	1.07	0.84	0.50	0.48
Italy	real VA	1.88	1.82	0.31	0.27
	VAD	0.77	0.61	0.50	0.29
	ULC	0.72	0.68	0.73	0.41
	PMA	0.94	0.84	0.88	0.45
Spain	real VA	1.68	1.70	1.25	2.05
	VAD	1.85	1.95	2.27	1.68
	ULC	1.24	1.41	1.46	1.89
	PMA	1.50	1.49	1.59	0.93

Notes: Each entry reports the MSPE from a recursive pseudo OOS forecast exercise with initial estimation period 1995Q1-2008Q1 relative to the MSPE from a recursive pseudo OOS forecast exercise with initial estimation period 1995Q1-2002Q3. Both exercises are based on the benchmark specification of the PVAR model.

of the recursive forecast, changes over time are generally small. Alternatively, the estimation period can be designed as a rolling window, i.e., in each round, the first quarter of the previous estimation period is dropped, while a new last quarter is added, thus rolling the estimation window forward. On the one hand, this allows for more variation in the coefficient estimates used for OOS forecasts over time. On the other hand, it foregoes the efficiency gains from expanding the estimation period. The trade-off between a lower bias in coefficient estimates accommodating variations over time and higher estimation uncertainty relative to the recursive OOS exercise will be reflected in the relative MSPE of both methods. Tables A.1 and A.2 in Appendix A show that the MSPEs are generally larger than those in Table 2 and 3, respectively, indicating a non-trivial role for efficiency gains in reducing the OOS forecast errors.²⁰

Both Table 4 and Table A.3 in the appendix reveal that short-horizon forecast errors tend to be larger during the recent crisis period, in particular for Spain. Following the global financial crisis, the fourth largest euro area economy confronted a severe economic recession and dramatic sectoral adjustments at the same time, affecting especially the Spanish construction sector. The disaggregated approach in this paper could therefore add to our understanding of the role played by cross-sector linkages in the recent economic developments in Spain.

5 Cross-Sector Linkages and the Spanish Construction Sector

This section uses the PVARX model to quantify the role of the Spanish construction sector in generating forecast errors during the crisis. Moreover, we employ impulse response functions to investigate how reduced-form shocks originating from the construction sector affect the other sectors and the economy as a whole. We finally provide an extension of the PVARX model in the country dimension in order to assess whether a shock in the Spanish construction sector is transmitted to the other euro area economies.

Economic developments in Spain have been largely affected by the housing boom-bust cycle, which was long and severe. This may explain why external, common factors are not sufficient to account for the larger forecast errors observed in the previous section in the case of Spain. House prices almost tripled between 1997 and early 2008, while the construction of housing more than doubled from its 1995 level. The share of investment in construction increased from 15% of GDP in 1995 to 22% of GDP in 2006-2007, which represented a significant diversion of

²⁰Note that the efficiency-gains interpretation is also consistent with a relatively lower MSPEs of the random walk forecast, as the latter is parsimoniously parameterised, i.e. only the drift must be estimated.




Figure 5: Time series of Spanish *sector-specific* component indicators of endogenous variables




Figure 6: Time series of *endogenous variable-specific* component indicators for Spain

productive resources from the tradable sector to the non-tradable construction sector (European Commission, 2013).

While a timid adjustment had already started in early 2007, the financial crisis triggered a stronger correction in the Spanish housing market. The fall in house prices and the contraction in residential investment led to severe declines in construction value added and employment. The weight of construction in employment had increased from 9% to almost 14% from 1995 to

2007, before declining sharply to less than 6% in 2013. Between 2008 and 2012, employment in the Spanish construction sector dropped by 1.5 million, which represents almost half of the fall in the total number of employed persons. The developments in the housing market and in the construction sector might therefore have played an important role in the Spanish recession and the sharp increase in unemployment.

Figure 5 plots the sector-specific indicators of the PVARX model for Spain, which reflect the housing boom before the financial crisis and the subsequent triple-dip recession in construction. The latter might have rescinded the gentle signs of recovery in industry and services in 2010 and 2011. The variable-specific indicators in Figure 6 illustrate the turmoil in the Spanish labour market and the unprecedented drop in value added in 2009. Following a short-lived recovery in 2010, the variable indicators of value added and the number of employees have been pointing downwards again. Surprisingly, the pronounced drop in real economic activity did not lead to a similar reduction in Spanish price and wage levels, as measured by the indicators of VAD and compensation per employee, respectively, which remained relatively flat throughout the crisis.

5.1 The Role of the Construction Sector in Forecast Errors during the Crisis

First, we investigate the ability of the PVAR model to forecast the Great Recession in Spain conditional on contemporaneous and lagged observations of the exogenous variables, taking the position of a forecaster at the start of the crisis and recursively predicting the path of the endogenous variables. Due to the fact that we condition our forecast on realized data for world demand, oil prices, etc., this represents a *pseudo out-of-sample* (POOS) forecast exercise.

Figure 7 plots the POOS forecasts of the sector-specific endogenous variables against the realized data for 2008Q3–2012Q3, i.e., the estimation period ends in 2008Q2 and recursive one-quarter-ahead conditional forecasts are made until the end of the full sample. A common weakness and critique of multivariate econometric models is that they fail to forecast the Great Recession OOS, even when conditioning on contemporaneous and past observations of the exogenous variables. The reason is that the strong comovement between the endogenous variables and exogenous driving forces, such as world demand and oil prices, e.g., only emerged at the start of the crisis. It is therefore not surprising that models estimated on the relatively calm sample period 1995Q2–2008Q2 fail to forecast the unprecedented downturn in industrial VA after the bankruptcy of Lehman Brothers. This is also reflected by a substantial change in parameter values, when estimating the PVAR model with and without the crisis period included.

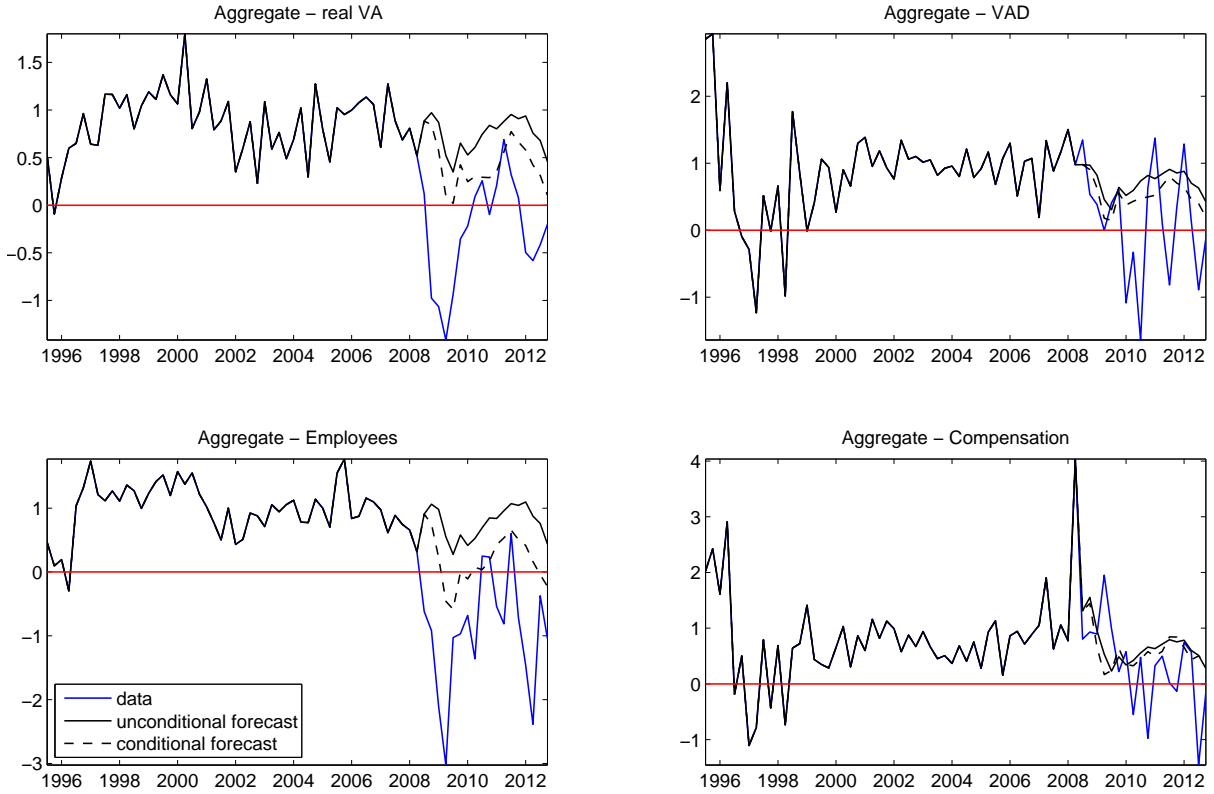


Figure 7: Unconditional and conditional POOS forecasts of economy-wide Spanish variables for 2008Q3–2012Q3 (Point forecasts)

Given the developments in the Spanish construction sector during the crisis, we perform another forecast exercise, in which we condition the forecasts on realized data of the *endogenous variables in the construction sector*. Although this does not exactly quantify the contribution of construction-related shocks to economy-wide forecast errors, it gives an idea about the importance of forecast errors in the construction sector. Figure 7 illustrates that, once we condition on the realized values of the latter variables, the ability of the PVARX model to track real activity and employment dynamics in Spain during the crisis improves substantially. In the case of employment, in particular, the forecast errors decrease by almost one half. Given that the construction sector accounted for no more than 15% of total employment during this period, our results indicate strong spillover effects from the former to the rest of the economy.

5.2 Impulse Response Functions

The previous analysis shows that conditioning the forecast on realised data for the construction sector allows to reduce the forecast errors at the aggregate level. We now turn to an impulse response analysis based on the benchmark PVAR model for the Spanish economy, in order to quantify the spillover effects from a particular sector to the whole economy.

Figure 8 plots the impulse response functions to a negative reduced-form innovation in real value added in the construction sector. The dashed and dotted lines indicate approximate 68 and 95% confidence intervals based on 1,000 replications of a recursive-design wild bootstrap that accounts for potential heteroscedasticity of unknown form in the error terms (see Gonçalves and Kilian, 2004). Note that all impulse responses are in terms of the variables' standard deviations. A direct consequence of the structure imposed by factorising δ are the broadly similar impulse response functions of variables belonging neither to the same sector nor to the same category as the disturbed variable.

Due to the fact that we consider a reduced-form innovation, only value added in construction responds on impact. Other variables in the same sector and the same variable in other sectors respond with a lag of one period, while the remaining variables respond with a lag of two periods. Note also that the point estimates of the impulse response functions are generally negative and almost always statistically significant in the medium run. Real value added decreases also in the other sectors. The maximum impact is reached after 3 quarters and has a magnitude of around one tenth of the initial decline in the construction sector.

As a consequence, the other variables also adjust downwards. In particular, in all sectors but agriculture, prices (i.e. VAD) remain significantly below baseline over the horizon considered. To get a better understanding of how cost pressures affect price responses, Figure 8 also shows the impacts of the decrease in construction VA on the cost components, i.e. unit labour costs and profit margins.²¹ Due to the decline in real activity, unit labour costs in the construction sector jump on impact. This is fully compensated by a fall in profit margins, as prices remain unchanged at the time of the shock.

The same patterns are also reflected in the other sectors, with some delay. To avoid strong price adjustments, firms tend to absorb higher labour cost, arising from the immediate fall in productivity with relatively sticky wages, by cutting their profit margins. It is interesting to note that, because downward wage adjustments are stronger in the services sector than in industry, unit labour costs fall short of the baseline after one year in the former, whereas they remain elevated throughout the simulation horizon in the latter. As a result, profit margins slightly increase above baseline after a few quarters in the services sector, whereas they remain below baseline in industry.

²¹Similarly to the forecast evaluation, we compute and plot impulse response functions for unit labour costs ($ULC = \frac{\text{Number of Employees} \cdot \text{Compensation per Employee}}{\text{Real Value Added}}$) and profit margins ($PMA = VAD - ULC$).




Figure 8: Impulse response functions of Spanish variables to a one-standard-deviation shock to *total real VA in Spanish construction* (Point estimates with one- and two-standard error confidence intervals)

As a caveat, it must be pointed out that there is very little economic interpretation to the reduced-form innovations in Figures 8. Instead, the purpose of this exercise is to show that the PVARX model allows for a transmission of shocks between sectors. In order to account for static interdependencies, we would have to identify *structural* innovations by imposing identifying restrictions on the contemporaneous covariance matrix $\hat{\Sigma}_e$ of the reduced-form innovations $e_{i,t}$, as in Canova et al. (2012). Maurin et al. (2011) identify structural shocks in a VAR including the same variables based on a recursive decomposition with employment ordered first, reacting instantaneously only to idiosyncratic shocks and with lag to all other shocks, followed by wages, prices and output. We abstain from this approach here, because (i) any ordering is questionable and (ii) finding a plausible extension to a multi-sector VAR model, where there is no theoretical motivation for ordering either the variables or the sectors, is even more difficult.

5.3 Assessing Spillovers Between Countries: A Multi-Country Extension

The PVAR model presented above captures a cross-section of $N = 4$ sectors and $K = 4$ variables. Our applications to Spanish sector-level data can be enriched by adding a *country dimension* to the model in order to allow for dynamic dependencies between the different euro area countries.

5.3.1 A Multi-Country PVAR Model

In equation (1), Y_t now corresponds to a $(C \cdot N \cdot K \times 1)$ vector of endogenous variables, ν to a $(C \cdot N \cdot K \times 1)$ vector of intercepts, A_j , $j = 1, \dots, p$ are $(C \cdot N \cdot K \times C \cdot N \cdot K)$ matrices of slope coefficients, and $e_t \sim iid(0, \Sigma_e)$ is a $(C \cdot N \cdot K \times 1)$ vector of possibly contemporaneously correlated reduced-form disturbances.

Letting $y_{c,i,t}$ denote the $(K \times 1)$ vector of endogenous variables for country c and sector i in period t and $Y_t = (y'_{1,1,t}, y'_{1,2,t}, \dots, y'_{2,1,t}, y'_{2,2,t}, \dots, y'_{C,N,t})'$ denote the $(C \cdot N \cdot K \times 1)$ vector of stacked $y_{c,i,t}$, $c = 1, \dots, C$, $i = 1, \dots, N$, we can write the PVARX model in (3) and (4) for C countries as

$$y_{c,i,t} = \nu_{1ci} + \sum_{l=1}^p A_{l,c,i} Y_{t-l} + \sum_{l'=0}^q B_{l',c,i} X_{t-l'} + e_{1ci,t} \quad (11)$$

$$X_t = \nu_2 + \sum_{l=1}^{p^x} C_l X_{t-l} + e_{2,t}, \quad (12)$$

where $B_{l',c,i}$, $l' = 0, \dots, q$ are $(K \times M)$ matrices of exogenous coefficients, and $e_{1ci,t}$ and $e_{2,t}$ are assumed to be uncorrelated.

The specification of the multi-country PVAR model is identical to the one for Spain, except that we include a $(C \times 1)$ vector θ_5 of *country-specific* components in the *endogenous* variables in addition to $\theta_1, \dots, \theta_4$ from Section 3.

Before considering the impulse responses to an innovation in the Spanish construction sector, we want to investigate how this multi-country version of the PVARX model performs relative to the country-specific models in terms of its out-of-sample forecast accuracy. For this purpose, Table 5 reports the MSPEs of the multi-country PVARX relative to those of the individual country models. In general, the relative MSPEs are larger than unity, indicating larger forecast errors of the multi-country PVARX model, on average over the evaluation period. Accordingly, the country-specific approach seems to be superior from the point of view of forecast accuracy. Nevertheless, the multi-country PVARX model enables us to quantify the spillover effects of country- and sector-specific disturbances to the other sectors within the same country as well as to sectors in the country's euro area neighbours. Moreover, the extension allows us to illustrate the informational content in the country-specific endogenous variables.

Table 5: Mean squared prediction errors of economy-wide aggregate variables for the multi-country PVAR model relative to those for the country-specific models

	Horizon	1	3	6	12
Germany	real VA	0.82	1.32	1.72	4.82
	VAD	1.04	1.79	2.93	6.47
	ULC	0.95	1.17	1.16	1.52
	PMA	0.93	1.26	1.33	3.08
France	real VA	1.20	1.30	1.47	2.76
	VAD	0.88	0.94	1.56	2.01
	ULC	1.07	0.96	1.02	0.48
	PMA	1.09	1.36	2.04	4.18
Italy	real VA	0.81	1.42	1.82	4.15
	VAD	1.21	1.64	2.40	11.41
	ULC	1.10	1.18	1.24	4.02
	PMA	0.98	1.03	1.11	1.27
Spain	real VA	1.31	1.02	0.92	1.70
	VAD	1.09	0.97	1.09	2.01
	ULC	0.88	0.91	0.81	0.78
	PMA	0.84	0.85	0.87	1.14

Notes: Each entry reports the MSPE from a recursive pseudo OOS forecast exercise for the multi-country PVAR model relative to results for the benchmark specification of the country-specific models, all with initial estimation period 1995Q1-2008Q2.

5.3.2 Illustrative Results

Figure 9 plots the country-specific component indicators for France, Germany, Italy, and Spain, based on the benchmark specification of the multi-country PVAR model. The graph illustrates several interesting developments before and during the Great recession. First, it reflects Germany's relatively low and Spain's comparatively high economic growth rates before the crisis as well as the German catch-up after 2005. Second, we can clearly see the increased comovement of country-specific indicators starting during 2008-2009 as well as the diverse paths of recovery in these four countries. Third, we detect the renewed sharp decline in the Spanish and, to a lesser extent, in the Italian economy in 2012, which largely undid the previous mild recoveries. Thus, the figure is again highly informative about the evolution of economic activity in the largest euro area countries during the sample period.

Figure 10 plots the impulse response functions of all $C \cdot N \cdot K = 64$ endogenous variables for horizons 1,...,15 to a reduced-form innovation in VA in the *Spanish construction sector*. The purpose of this exercise is to illustrate that shocks are transmitted between the different sectors within a country as well as between sectors in different countries, at least within the euro area. Hence, the multi-country PVARX model captures the dynamic interdependencies in both the sector and the country dimension. Note that, although the spillover effects are quantitatively small, they are almost always statistically significant (at an approximate 68 and 95% confidence level) for an extended period. By construction, there are parallels between the impulse response




Figure 9: Time series of *country-specific* component indicators of endogenous variables

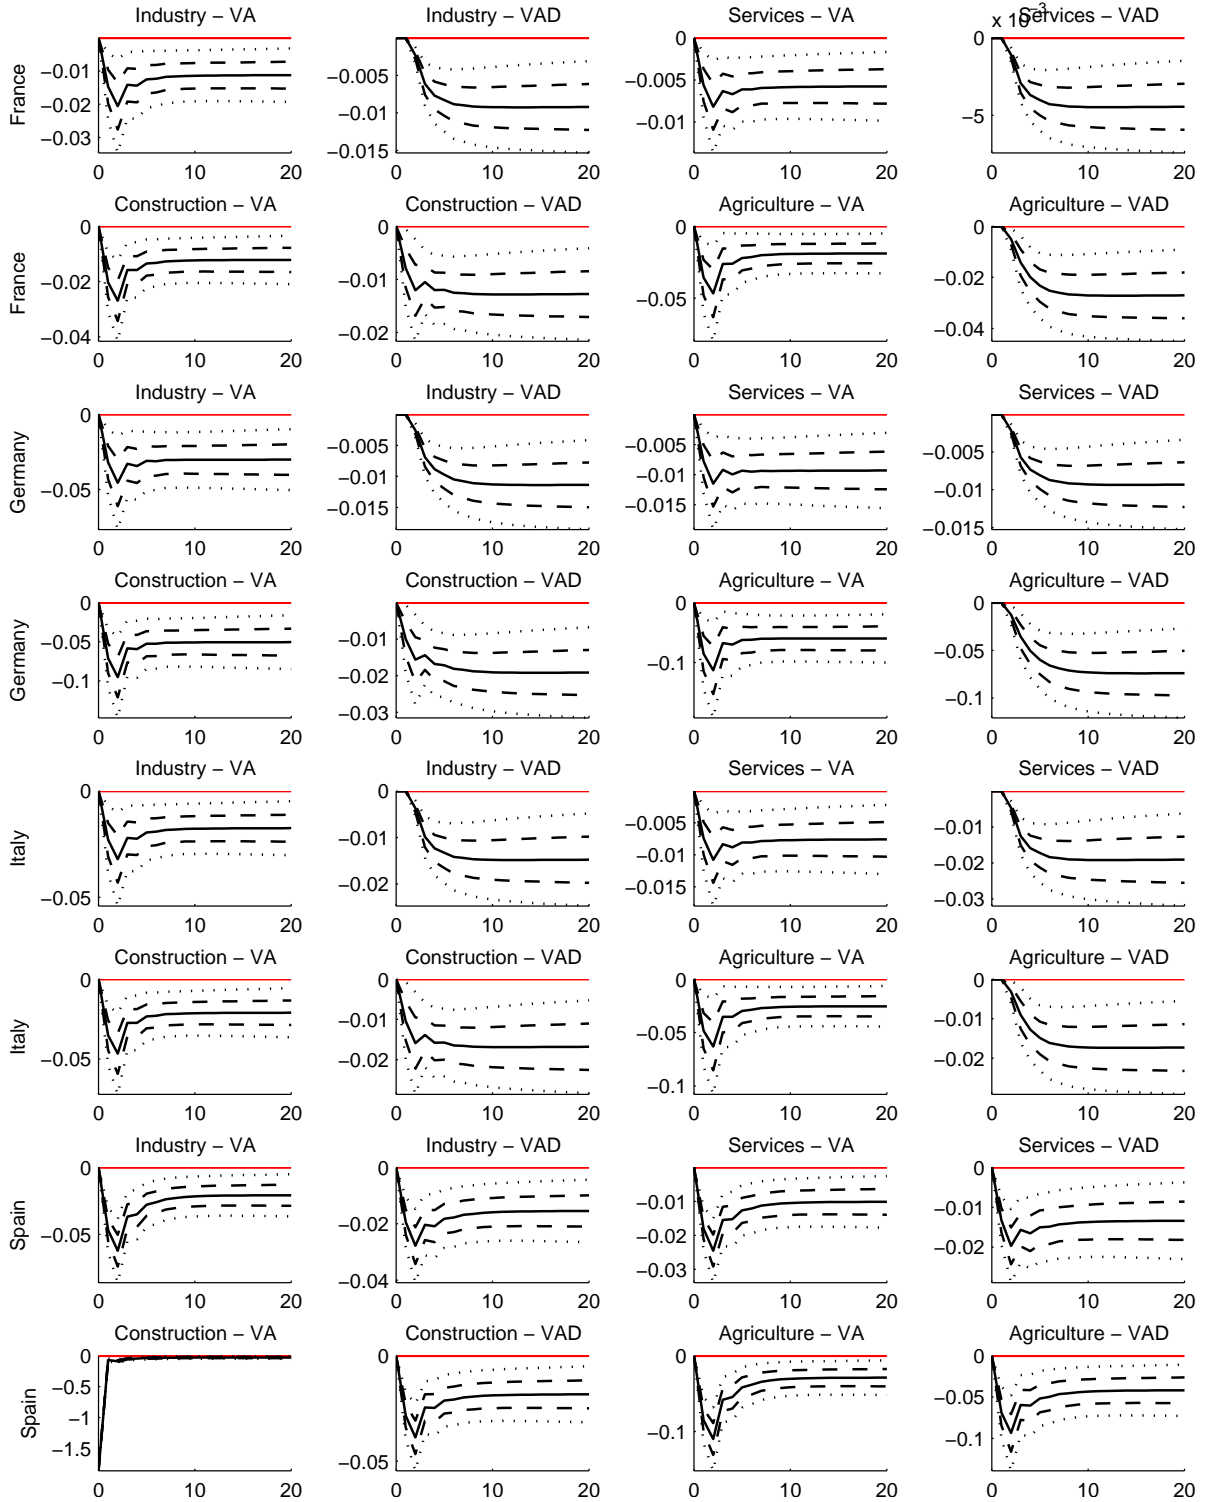


Figure 10: Impulse response functions of selected variables to a one-standard-deviation shock to *total real VA in Spanish construction* based on the multi-country PVAR model (Point estimates with one- and two-standard error confidence intervals)

functions of the same variables or variables in the same sector *across countries*. In comparison with the country-specific model, the multi-country PVARX model leads to more persistent responses. In particular, the impact on prices seems to be permanent in most cases.

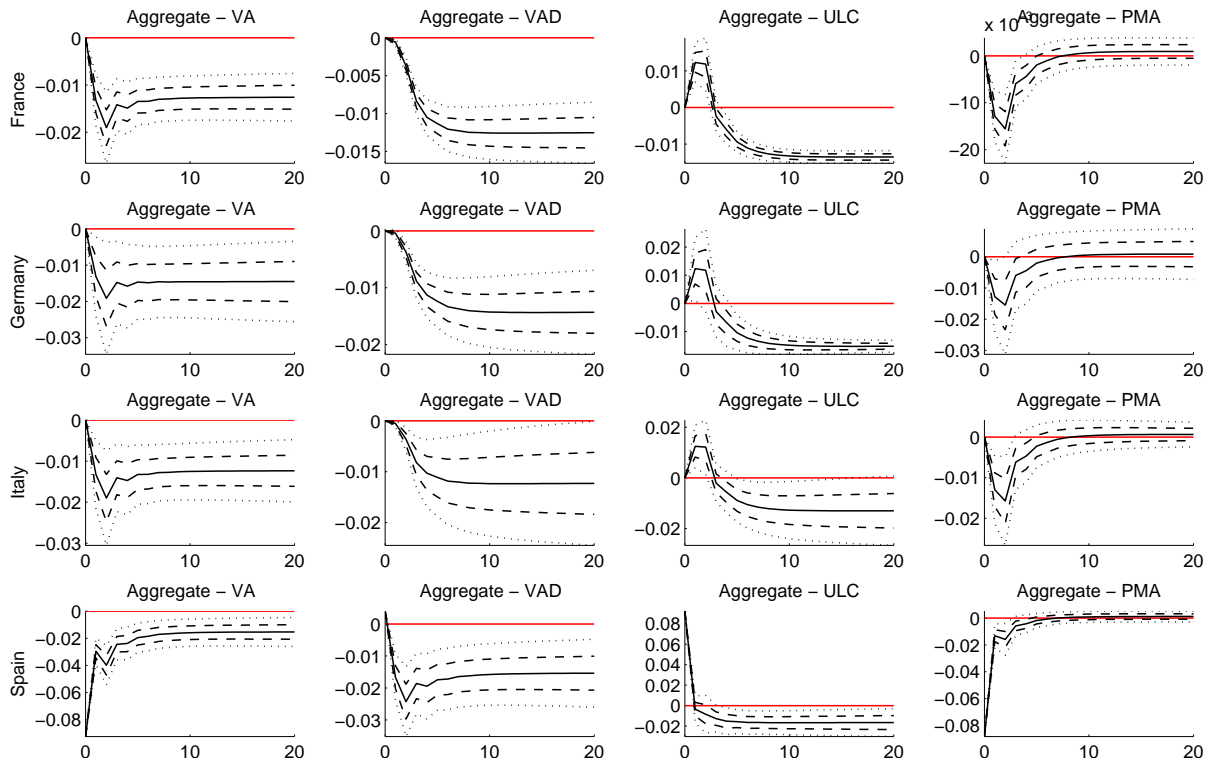


Figure 11: Impulse response functions of selected variables to a one-standard-deviation shock to *total real VA in Spanish construction* based on the multi-country PVAR model (Point estimates with one and two-standard error confidence intervals)

At the aggregate level (see Figure 11), a negative one-standard-deviation innovation in total real VA in the Spanish construction sector, i.e. a decline by around 2%, leads to an immediate decrease in economy-wide total real VA of around 0.1%. The VAD is affected permanently, with a maximum impact after 2 quarters. As for the country-specific model, the price response induces a decline of profit margins in the short term, which more than offsets the increase in ULC coming from the reduction in value added. In the medium term, profit margins are gradually restored, as real activity returns to the baseline. In contrast, the shock affects ULC permanently, as employment remains below the baseline. The response of ULC also explains the persistent effect on VAD.

The impulse responses of Spanish variables are transmitted to the other euro area countries through cross-country linkages, which in turn amplify the direct effects on the Spanish economy. Indeed, when comparing the results from the country-specific model in Figure 8 with those in Figure 11, the responses of total real VA and VAD are larger and more persistent for the multi-country PVARX model.

As mentioned before, our model cannot account for possible changes in the transmission of shocks between variables, sectors and countries over time. Using a time-varying PVAR model of

six macroeconomic time series for ten European countries, however, Canova et al. (2012) show that the transmission of German and U.S. real GDP shocks to the growth rates of the other countries is largely constant in terms of its sign and shape between 1998Q3 and 2002Q1, and thus over a crucial part of our sample period, while there is some variation in the magnitude of the responses. Unfortunately, a robustness check similar to the rolling-window OOS forecast exercise is not available here.

6 Concluding Remarks

In this paper we propose the use of a PVAR approach to analysing and forecasting price dynamics from a sectoral perspective. Focusing on sector-level data enables us to take into account cost-push factors from the supply side. The PVAR models are estimated for four economic sectors – industry, services, construction and agriculture – in the euro area and its four largest member countries. By modelling prices together with real activity, employment and wages, we are able to decompose price dynamics into unit labour costs and profit margins.

Our modelling strategy proves to be more accurate than simple time-series approaches in out-of-sample forecasting. The disaggregated approach also performs well relative to direct forecasts of the aggregate variables. While the forecast accuracy of the PVAR model is satisfactory overall, forecast errors tend to be larger during the financial crisis period. Among the euro area countries considered, the mean squared prediction errors are particularly large for Spain. Given that the Spanish economy was confronted with severe structural changes due to the burst of the housing bubble, we further employ the PVAR model to illustrate the transmission of shocks originating from the construction sector to the rest of the economy and, in a multi-country extension, to the other euro area economies.

We find the PVAR approach to be useful both for forecasting purposes and as an analytical tool. An important advantage is that it can easily be extended to include additional sectors or countries as well as more disaggregated data. Given a longer sample period, a time-varying parameter version, as proposed in Canova and Ciccarelli (2009), could account for structural changes in the interdependencies within and between sectors. It is furthermore important to note that the OOS forecast exercise in this paper is not equivalent to a *real-time* forecasting experiment. Besides the fact that the OOS forecasts are conditional on realized rather than predicted observations of the exogenous variables, the PVARX and all other candidate models

are estimated using *ex-post revised* rather than real-time data vintages. Given that the GDP and VA deflator are subject to more important revisions than HICP, the results presented here might be less general than those in Hubrich (2005) and Benalal et al. (2004). A real-time PVAR analysis is complicated by the fact that the data for the euro area are in changing composition and that base years change across data vintages, to name just two. As a consequence, we leave these extensions for future research.

References

Atkeson, Andrew and Lee E. Ohanian (2001), “Are Phillips curves useful for forecasting inflation?”, *Quarterly Review*, Federal Reserve Bank of Minneapolis, 25(1), 2-11.

Ball, Laurence M. and Sandeep Mazumder (2011), “Inflation dynamics and the Great Recession”, *Brookings Papers on Economic Activity* 42(1), 337–405.

Barhoumi, Karim, Olivier Darné, Laurent Ferrara and Bertrand Pluyaud (2012), “Monthly GDP forecasting using bridge models: Application for the French economy”, *Bulletin of Economic Research* 64(s1), s53-s70.

Bassetto, Marco, Todd Messer and Christine Ostrowski (2013), “Forecasting inflation and the Great Recession”, Federal Reserve Bank of Chicago *Economic Perspectives* 2013Q3, 79-106.

Benalal, Nicholai, Juan Luis Diaz del Hoyo, Bettina Landau, Moreno Roma and Frauke Skudelny (2004) “To aggregate or not to aggregate? Euro area inflation forecasting”, *ECB Working Paper Series* 374.

Canova, Fabio and Matteo Ciccarelli (2004), “Forecasting and turning point predictions in a Bayesian panel VAR model”, *Journal of Econometrics* 120(2), 327-359.

Canova, Fabio and Matteo Ciccarelli (2009), “Estimating multicountry VAR models”, *International Economic Review* 50(3), 929-959.

Canova, Fabio and Matteo Ciccarelli (2013), “Panel Vector Autoregression Models”, *ECB Working Paper Series* 1507.

Canova, Fabio, Matteo Ciccarelli and Eva Ortega (2012), “Do institutional changes affect business cycles? Evidence from Europe”, *Journal of Economic Dynamics and Control* 36(10), 1520-1533.

Drechsel, Katja and Rolf Scheufele (2012), “Bottom-up or Direct? Forecasting German GDP in a Data-rich Environment”, *Swiss National Bank Working Paper* 2012-16.

European Commission (2013), “Macroeconomic Imbalances – Spain 2013”, *European Economy Occasional Paper* 134.

Faust, Jon and Jonathan H. Wright (2013), “Forecasting Inflation”, in *Handbook of Economic Forecasting* 2A, eds. G. Elliott and A. Timmermann, Amsterdam, Netherlands: North-Holland.

Gardner, Everette S. (1985), “Exponential smoothing: The state of the art”, *Journal of Forecasting* 4(1), 1-28.

Gonçalves, Silvia and Lutz Kilian (2004), “Bootstrapping autoregressions with conditional heteroskedasticity of unknown form”, *Journal of Econometrics* 123(1), 89-120.

Hahn, Elke and Frauke Skudelny (2008), “Early estimates of euro area real GDP growth – a bottom up approach from the production side”, *ECB Working Paper Series* 975.

Hendry, David F. and Kirstin Hubrich (2011), “Combining Disaggregate Forecasts or Combining Disaggregate Information to Forecast an Aggregate”, *Journal of Business & Economic Statistics* 29(2), 216-227.

Hubrich, Kirstin (2005), “Forecasting euro area inflation: Does aggregating forecasts by HICP component improve forecast Accuracy?”, *International Journal of Forecasting* 21(1), 119-136.

Kohn, Robert (1982), “When is an aggregate of a time series efficiently forecast by its past?”, *Journal of Econometrics* 18(3), 337-349.

Luetkepohl, Helmut (1984a), “Linear transformations of vector ARMA processes”, *Journal of Econometrics* 26(3), 283-293.

Luetkepohl, Helmut (1984b), “Forecasting Contemporaneously Aggregated Vector ARMA Processes”, *Journal of Business & Economic Statistics* 2(3), 201-214.

Luetkepohl, Helmut (1987), “Forecasting Aggregated Vector ARMA Processes”, Berlin, Germany: Springer-Verlag.

Luetkepohl, Helmut (2005), “New Introduction to Multiple Time Series Analysis”, New York, NY: Springer-Verlag.

Luetkepohl, Helmut (2006), “Forecasting With VARMA Processes”, in *Handbook of Economic Forecasting*, eds. G. Elliott, C.W.J. Granger and A. Timmermann, Amsterdam, Netherlands: Elsevier.

Marcellino, Massimiliano, James H. Stock and Mark W. Watson (2003), “Macroeconomic forecasting in the Euro area: Country specific versus area-wide information”, *European Economic Review* 47(1), 1-18.

Maurin, Laurent, Moreno Roma and Igor Vetlov (2011), “Profit dynamics across the largest euro area countries and sectors”, *ECB Working Paper Series* 1410.

Theil, Henri (1954), “Linear Aggregation of Economic Relations”, Amsterdam, Netherlands: North-Holland Publishing.

Tiao, G.C. and Irwin Guttman (1980), “Forecasting contemporaneous aggregates of multiple time series”, *Journal of Econometrics* 12(2), 219-230.

Appendix A. Rolling-Window OOS Forecast Performance

Table A.1: Relative MSPEs of economy-wide aggregate variables for alternative factorizations of the parameter vector δ

		Benchmark	Model A	Model B	Model C	Model D	Model E
Euro Area	real VA	0.9741	0.9465	0.9482	0.9751	0.8683	0.8224
	VAD	1.1560	1.1329	1.1357	1.1071	1.1079	1.3541
	ULC	1.0347	1.0478	1.0079	1.0305	1.0041	1.0190
	PMA	0.9455	0.9294	0.9444	0.9621	0.9805	0.8526
Germany	real VA	1.0686	1.0660	0.9257	1.3084	1.0283	0.8818
	VAD	1.0072	0.9524	1.0516	3.5716	0.9687	1.1372
	ULC	1.0314	1.0500	0.9711	1.5473	0.9950	1.0190
	PMA	1.1045	1.1091	0.8943	1.9537	1.0374	0.9683
France	real VA	1.0666	0.9822	1.0428	1.0077	0.9804	1.1829
	VAD	1.0663	1.3626	1.1614	1.1570	0.9857	0.9982
	ULC	0.9420	0.9828	0.9391	1.0283	0.9550	0.9297
	PMA	1.1343	1.0605	1.2045	1.1100	1.0044	1.1501
Italy	real VA	1.0095	0.9839	1.0046	0.8848	0.9911	0.9307
	VAD	1.1012	1.1234	1.1188	1.0968	1.0698	1.0972
	ULC	1.0952	1.1151	1.0684	1.1299	1.0105	1.0946
	PMA	1.0679	1.0631	1.0017	1.0830	0.9732	1.0651
Spain	real VA	1.1103	1.1490	1.1118	0.9882	0.91040	1.0538
	VAD	1.1318	1.3240	1.0034	1.1312	1.1409	1.1814
	ULC	1.1685	1.2533	0.9995	1.2273	1.1333	1.2189
	PMA	1.0812	1.0638	0.9592	1.1270	1.0210	1.0756
# parameters	θ	30	31	14	26	26	36
# parameters	δ	720	720	720	720	720	720

Notes: Each entry reports the MSPE from a pseudo OOS forecast exercise with a rolling estimation window of 30 quarters relative to the recursive equivalent with initial estimation period 1995Q1-2002Q3.

Table A.2: Relative MSPEs of economy-wide aggregate variables for alternative forecasting models

		Benchmark PVAR	Random Walk	AR	Aggr. VAR
Euro Area	real VA	0.9742	0.8940	1.0921	1.3404
	VAD	1.1560	0.9751	1.0910	4.5111
	ULC	1.0347	0.9847	1.1851	2.9342
	PMA	0.9455	1.0046	1.1147	1.1646
Germany	real VA	1.0686	0.9845	1.2194	1.2071
	VAD	1.0072	0.8653	0.8582	1.5045
	ULC	1.0314	0.9994	1.2938	1.5301
	PMA	1.1045	1.0532	1.3895	1.7234
France	real VA	1.0666	0.8946	1.0377	2.3005
	VAD	1.0663	0.9952	1.0483	1.2454
	ULC	0.9420	0.9612	0.9532	1.1408
	PMA	1.1344	1.0016	0.9273	3.0104
Italy	real VA	1.0095	0.9139	1.0047	1.9569
	VAD	1.1012	0.7469	0.6176	1.8584
	ULC	1.0953	0.9769	0.9191	2.9524
	PMA	1.0679	0.9678	1.1144	2.8134
Spain	real VA	1.1103	0.7861	0.9171	2.6526
	VAD	1.1318	0.9715	1.1140	2.6470
	ULC	1.1685	0.9123	1.1848	3.9147
	PMA	1.0812	1.0164	1.1439	3.2988

Notes: Each entry reports the MSPE from a pseudo OOS forecast exercise with a rolling estimation window of 30 quarters relative to the recursive equivalent with initial estimation period 1995Q1-2002Q3.

Table A.3: Relative MSPEs of economy-wide aggregate variables for the crisis relative to those for the benchmark evaluation period

	Horizon	1	3	6	12
Euro Area	real VA	1.87	1.63	0.24	0.37
	VAD	0.96	0.75	0.81	0.84
	ULC	1.57	1.34	0.51	0.23
	PMA	1.84	1.51	0.39	0.31
Germany	real VA	1.73	1.39	0.39	0.11
	VAD	0.77	1.75	2.18	0.46
	ULC	1.84	1.66	0.23	0.43
	PMA	1.62	1.53	0.40	0.38
France	real VA	1.74	1.69	0.74	1.66
	VAD	0.91	0.90	0.85	3.01
	ULC	1.06	1.03	0.67	2.11
	PMA	1.07	0.84	0.50	0.48
Italy	real VA	1.77	1.69	0.34	0.27
	VAD	0.74	0.56	0.48	0.32
	ULC	0.71	0.61	0.58	0.35
	PMA	0.90	0.66	0.72	0.34
Spain	real VA	1.53	1.77	1.89	3.80
	VAD	1.71	1.67	1.40	1.40
	ULC	1.24	1.65	1.91	6.62
	PMA	1.47	1.71	1.94	5.48

Notes: Each entry reports the MSPE from a rolling-window pseudo OOS forecast exercise with evaluation period starting in 2008Q3 relative to the MSPE for an evaluation period starting in 2002Q4. Both exercises are based on the benchmark specification of the PVAR model.