DEUTSCHE BUNDESBANK
EUROSYSTEM

Discussion Paper

Deutsche Bundesbank
No 06/2020

Partial pooling with cross-country priors: An application to house price shocks

Markus Roth

Deutsche Bundesbank, Wilhelm-Epstein-Straße 14, 60431 Frankfurt am Main,
Postfach 10 06 02, 60006 Frankfurt am Main

Tel +49 69 9566-0

Please address all orders in writing to: Deutsche Bundesbank,
Press and Public Relations Division, at the above address or via fax +49 69 9566-3077

Internet http://www.bundesbank.de

ISBN 978–3–95729–670–2 (Printversion)
ISBN 978–3–95729–671–9 (Internetversion)

Non-technical summary

Research Question

In the paper the impact of house price bust shocks on economic activity and consumption is analyzed. To this end statistical models of the aggregate housing sectors are estimated for euro area countries in order to quantify the damage to be expected from sudden falls in house prices. As a byproduct, the approach delivers a dissimilarity measure of countries' real estate markets.

Contribution

In a novel statistical approach, cross-country information is employed in order to mitigate the small-sample problem. An estimator for countries' real estate market dissimilarities is proposed which, on the one hand, considers country specificities, and on the other hand, allows for general cross-country economic mechanisms. The model parameters are determined based on data information which contain consistent time series for the countries under consideration. This yields a stabilization of estimates.

Results

The model predicts negative effects of house price bust shocks on economic activity and consumption which vary in size across countries. That is, while the general economic transmission mechanisms of the shocks are similar, the severity of their respective damage differs. It appears that in a low-interest rate environment the effect of shocks is larger.

Nichttechnische Zusammenfassung

Fragestellung

In dem Papier wird der Einfluss von Hauspreisschocks auf die Wirtschaftstätigkeit und den Konsum analysiert. Dazu werden statistische Modelle der aggregierten Immobiliensektoren für Länder der Eurozone geschätzt, um damit den zu erwartenden Schaden eines plötzlichen Verfalls der Immobilienpreise zu quantifizieren. Als Nebenprodukt liefert der Ansatz ein Maß für die Unterschiedlichkeit der Immobilienmärkte einzelner Länder.

Beitrag

In einem neuartigen statistischen Ansatz werden länderübergreifende Informationen verwendet, um das Problem kleiner Stichproben zu mildern. Es wird ein Schätzverfahren für die Unterschiedlichkeit der Immobilienmärkte in den betrachteten Ländern vorgeschlagen, welches auf der einen Seite Besonderheiten der Länder berücksichtigt und auf der anderen Seite generelle länderübergreifende ökonomische Mechanismen abbildet. Die Modellparameter werden mit Hilfe von Informationen eines Datensatzes bestimmt, welcher konsistente Zeitreihen für die betrachteten Länder beinhaltet. Das führt dazu, dass die Schätzergebnisse stabilisiert werden.

Ergebnisse

Das Modell sagt negative Effekte von Hauspreisschocks auf die Wirtschaftstätigkeit und den Konsum voraus, wobei sich die Größenordnung der Effekte über die Länder hinweg unterscheidet. Das heißt, während sich die generellen ökonomischen Transmissionsmechanismen der Schocks im Querschnitt der Länder ähneln, ist die Schwere des zugehörigen Schadens unterschiedlich. Auch zeigt sich, dass der Einfluss der Schocks bei niedrigen Zinsen größer ist.

Bundesbank Discussion Paper No 06/2020

Partial pooling with cross-country priors: An application to house price shocks*

Markus Roth
Deutsche Bundesbank

Abstract

A structural Bayesian vector autoregression model predicts that – when accompanied by a decline in consumer confidence – a one-percent decrease in house prices is associated with a contraction of economic activity by 0.2 to 1.2 percent after one year. Results point to important second-round effects and additional exercises highlight the amplifying role of (*i*) the mortgage rate and (*ii*) consumers' expectations. A novel econometric approach exploits information available from the cross section. Shrinkage towards a cross-country average model helps to compensate for small country samples and reduces estimation uncertainty. As a by-product, the method delivers measures of cross-country heterogeneity.

Keywords: Bayesian model averaging, dummy observations, house price shocks

JEL classification: C11, C33, E44.

*Contact address: Directorate General Financial Stability, Deutsche Bundesbank, Wilhelm-Epstein-Strasse 14, 60431 Frankfurt am Main, Germany. Phone: +49 69 9566 8561. E-Mail: markus.roth@bundesbank.de. I thank Nikolay Hristov, Martin Mandler, Philipp Marek, Frieder Mokinski, Michael Richter, Frank Schorfheide, Harald Uhlig, Edgar Vogel and an anonymous referee for helpful comments and discussions. Suggestions by Michael Dear are gratefully acknowledged. Remaining errors are my sole responsibility. This paper represents the author's personal opinions and does not necessarily reflect the views of the Deutsche Bundesbank.

1 Introduction

The US housing bubble bust in 2007 and its consequences for the economy underlined the important role of house prices[1] for the real sector (see e.g. Piazzesi and Schneider, 2016; Lambertini, Mendicino, and Punzi, 2017). In a speech, Praet (2019) emphasizes the role of house prices for financial stability in the euro area.

> "*The effects of residential real estate [RRE] bubbles on financial stability are also well known. More than two-thirds of the past 46 systemic banking crises were preceded by boom-bust patterns in house prices. And recessions coinciding with house price busts have been found to yield a cumulative loss in GDP that is around three times greater than in recessions without such busts.*"

Since housing market shocks, potentially initiated by sudden reversals in households' confidence in future macroeconomic conditions, helped to ignite the recent financial crisis, (see e.g. Lambertini, Mendicino, and Punzi, 2013), house price developments are being given particular attention by most macroprudential authorities. Especially in a low-interest rate environment, their growth rates may accelerate and decouple from fundamentals. With respect to regulation, the European Systemic Risk Board (see ESRB, 2016) identifies the residential real estate market to be an important segment, writing that

> "*vulnerabilities in the residential real estate [RRE] market is a key responsibility of macroprudential authorities*".

Similarly, Bernanke (2010) notes that

> "*[a]lthough a number of developments helped trigger the crisis, the most prominent one was the prospect of significant losses on residential mortgage loans to subprime borrowers that became apparent shortly after house prices began to decline.*"

He considers a fall in house prices to be a potential trigger of financial crises. In this respect, interest regularly centers on early-warning tools, i.e. assessments of the appropriateness of price or mortgage developments, where "abnormal" dynamics imply warning signals which potentially trigger the use of regulatory instruments. However, a macroprudential authority might also be interested in quantifying the benefits of regulation in terms of damage from sudden house price drops avoided. The analysis focuses on the losses caused by, as opposed to the likelihood of, a sudden deterioration in the housing market and proposes a model with which to estimate the impact of house price shocks on economic activity as measured by gross domestic product (GDP) and consumption. Considering the household sector at an aggregate level illustrates the transmission channel of adverse shocks to house prices.

Housing markets are usually expected to feature regional specificities which call for a country-based assessment. However, general economic mechanisms should not depend on national peculiarities. For instance, the ESRB (2016) highlights the importance of national specificities that coexist with market commonalities. Consequently, I estimate

[1] Instead of using the comprehensive term "residential real estate prices", the term "house prices" is used for convenience.

a dynamic multivariate model at the country level which incorporates information from other countries in order to obtain an econometrically robust model and economically meaningful results. The underlying assumption is that, while unspecified, there exist common economic mechanisms across countries but their relative importance is unknown. Even though similar estimation strategies are available, the procedure is novel to the literature. It combines the advantages of Bayesian hierarchical approaches with a dummy observations prior distribution informed by the cross section.

This paper makes a technical contribution by describing a novel way to consider panel data-based information. In contrast to previous approaches, variation in the cross section is fully processed. The estimation procedure offers a balanced way to account for observations of foreign countries. It outlines a simple and convenient method to form sensible prior distributions if data at the country level is scarce but observations are available along the cross section. The setup allows me to estimate extensive models imposing a well-chosen prior distribution tightness. It documents the degree of cross-country heterogeneity, which mainly shows up in the magnitude of effects. Qualitatively, countries' impulse response functions feature similar patterns. The effects on GDP and consumption are quite substantial, ranging from -0.2 to about -1.2 percent after one year given a particular one percent drop in house prices on impact. Second-round mechanisms seem important because the initial reactions are considerably smaller. The approach shares many similarities with existing methods. It contains average cross-country information in the form of dummy observations which are weighted with a tightness parameter. A hierarchical Bayesian model averaging (BMA) approach is used to estimate the prior tightness parameter. The posterior distribution of this parameter may further be used to illustrate heterogeneity in the sample.

Similar approaches have been surveyed by Canova and Ciccarelli (2013), who conclude that "*partial pooling*" may impose an appropriate degree of homogeneity. Estimation strategies that follow their suggestion identify improvements with respect to out-of-sample prediction or estimation precision.[2] The model frameworks probably most closely related to the present study are those of Jarociński (2010) and Koop and Korobilis (2016). The former comprises shrinkage of the posterior slope coefficients to a cross-country common mean. The latter allows for cross-sectional homogeneity[3] as well as interdependencies. Previous approaches typically feature prior distributions with diagonal covariance matrices on the slope coefficients; they might overemphasize unreasonable combinations of the slope coefficients, as argued by Del Negro and Schorfheide (2011). Similarly, it is common not to impose beliefs on the residual covariance matrix. In contrast, the present approach features a prior distribution fully considering the coefficient covariance matrix as well as the residual covariance matrix. The proposed method is easily implemented and offers a plausible and transparent treatment of cross-country heterogeneity.

Most macro models covering the housing sector focus on the effects of monetary policy on house prices or mortgage credit. In this respect, only a few studies investigate empirically the impact of shocks to house prices on the economy. Ahamada and Sanchez (2013) show for US data that a one-percent decrease in house prices has a substantial impact on

[2] See for instance Zellner and Hong (1989), Rickman (1995), Canova (2005), Jarociński (2010), Calza, Monacelli, and Stracca (2013), or Koop and Korobilis (2016).

[3] Frey and Mokinski (2016) propose a related concept in a non-panel context. They show that imposing prior beliefs about the similarity of coefficients may improve forecasting performance.

consumption and GDP of about -0.5 to -1 percent after ten quarters. They identify intensified effects in subsamples following housing market deregulation. Disney, Gathergood, and Henley (2010) estimate a microeconometric model of household savings and infer a marginal propensity to consume with respect to housing wealth shocks of 0.01. In a euro area panel analysis, Nocera and Roma (2017) find that positive housing demand shocks increase GDP by 0.09 percent.[4] Gustafsson, Stockhammar, and Österholm (2016) employ a conditional forecasting exercise and show that a fall in house prices has an adverse effect on the economy, especially if the shock is combined with an international recession. Calza et al. (2013) show that contractionary monetary policy shocks have more pronounced effects on residential investment and house prices if mortgage markets are more flexible. They employ a "stochastic pooling" approach in order to account for cross-country similarities as well as specificities in the impulse response functions. Since the econometric setup of studies is hardly comparable, the literature documents a large variety of results, where macroeconometric approaches seem to identify more severe effects. The present study confirms those findings and attributes them to the expectations and the interest rate.

The remainder of this paper is organized as follows. Section 2 lays out the empirical strategy and discusses identifying assumptions. Section 3 summarizes the data. Results are discussed in section 4, and section 5 concludes.

2 Model

This section provides an overview of the empirical technique with respect to the implementation of the prior distribution, the model averaging procedure, and the structural shock identification. The approach involves a standard vector autoregression (VAR) model with a hierarchical cross-sectional prior specification[5] which is used to estimate the impact of house price adjustments on the real economy. The prior distribution is used as a device for partial pooling, i.e. for merging observations across countries to an appropriate degree. In fact, the model brings together existing building blocks which are well-understood but have not yet been combined in this particular way in one unified model framework.

2.1 Reduced form

Interest centers on a VAR model with respect to a particular country indexed by c

$$\boldsymbol{y}_{c,t} = \sum_{j=1}^{p} \boldsymbol{B}_{c,j} \boldsymbol{y}_{c,t-j} + \boldsymbol{e}_{c,t}, \quad \text{for } t = 1, \ldots, T_c \tag{1}$$

where $\boldsymbol{y}_{c,t}$ is an $n \times 1$ vector containing the endogenous model variables, the $n \times n$ matrices $\boldsymbol{B}_{c,j}$ contain the slope parameters of the model, and $\boldsymbol{e}_{c,t} \sim \mathcal{N}(\boldsymbol{0}, \boldsymbol{\Sigma}_c)$ is a $n \times 1$ vector of reduced-form residuals.[6] All variables considered enter the model in first differences. Time

[4]The authors use a mean-group average estimate according to Pesaran and Smith (1995).

[5]For a typical hierarchical approach in the context of Bayesian VARs see Giannone, Lenza, and Primiceri (2015).

[6]Note that the model does not allow for interdependencies across countries while the prior distribution proposed below accounts for potential cross-sectional homogeneity. I argue that, with respect to the

series $\{\boldsymbol{y}_{c,t}\}_{t=1}^{T}$ are standardized ex ante country-by-country, which allows the omission of an intercept term in the model.[7] Equation (1), together with the distributional assumption on the reduced form error vector $\boldsymbol{e}_{c,t}$, defines a likelihood function and describes my assumption about the data generating process (DGP).

Given the panel structure of the data, depending on the appropriateness of priori homogeneity assumptions there are two strategies to estimate a VAR model. First, one may assume a country-level perspective and estimate the model for each of the $c = 1, \dots, C$ countries separately, ignoring information in the cross section. Second, observations are pooled across countries carrying the assumption of homogeneous models. The former approach has the advantage of providing results for each cross-sectional unit. If the number of coefficients is large, the number of observations may be insufficient in order to estimate the model properly. While it is generally possible to obtain coefficients that minimize the in-sample fit, they may feature poor out-of-sample performance. The so-called problem of "overfitting" delivers misleading results in small samples. Pooling the observations sacrifices cross-sectional specific results for an extended dataset. Attention centers on inference at a country level, yet I seek to overcome the overfitting problem in a more careful way using a sensible prior distribution.

One representation of the VAR model (1) likelihood function is

$$f\left(\boldsymbol{Y}_c|\boldsymbol{\beta}_c, \boldsymbol{\Sigma}_c\right) = (2\pi)^{-\frac{nT_c}{2}} \det\left(\boldsymbol{\Sigma}_c\right)^{-\frac{T_c}{2}} \exp\left\{-\frac{1}{2}\text{tr}\left[\boldsymbol{\Sigma}_c^{-1}\left(\boldsymbol{Y}_c - \boldsymbol{X}_c\boldsymbol{B}_c\right)'\left(\boldsymbol{Y}_c - \boldsymbol{X}_c\boldsymbol{B}_c\right)\right]\right\}, \tag{2}$$

where $\boldsymbol{B}_c \equiv [\boldsymbol{B}_{c,1}, \cdots, \boldsymbol{B}_{c,p}]'$ is $(n \cdot p) \times n$, $\boldsymbol{Y}_c \equiv [\boldsymbol{y}_{c,1}, \cdots, \boldsymbol{y}_{c,T}]'$, $\boldsymbol{x}_{c,t} \equiv [\boldsymbol{y}'_{c,t-1}, \cdots, \boldsymbol{y}'_{c,t-p}]$, $\boldsymbol{X}_c \equiv [\boldsymbol{x}'_{c,1}, \cdots, \boldsymbol{x}'_{c,T}]'$, and $\boldsymbol{\beta}_c \equiv \text{vec}(\boldsymbol{B}_c)$. The Bayesian strategy allows me to specify a prior distribution which contains information about the model before analyzing the data in $\boldsymbol{Y}_c$. A well-chosen prior addresses the problem of overfitting. In the context of VAR models, most frequently the so-called Minnesota prior (see Litterman, 1979; Doan, Litterman, and Sims, 1984; Sims and Zha, 1998), which in its basic representation, shrinks the posterior distribution to a simple random walk model, is used. Alternatively, I propose the likelihood function of the cross-country pooled observations to serve as a reasonable device to form a priori beliefs.

Consider a panel model that pools the observations with respect to all countries except country c. The corresponding (non-country-c) likelihood function reads

$$\begin{aligned} f(\boldsymbol{Y}_c^*|\boldsymbol{\beta}_c^*, \boldsymbol{\Sigma}_c^*) = (2\pi)^{-\frac{nT_c^*}{2}} \det\left(\boldsymbol{\Sigma}_c^*\right)^{-\frac{T_c^*}{2}} \cdots \\ \cdots \times \exp\left\{-\frac{1}{2}\text{tr}\left[\left(\boldsymbol{\Sigma}_c^*\right)^{-1}\left(\boldsymbol{Y}_c^* - \boldsymbol{X}_c^*\boldsymbol{B}_c^*\right)'\left(\boldsymbol{Y}_c^* - \boldsymbol{X}_c^*\boldsymbol{B}_c^*\right)\right]\right\}, \end{aligned} \tag{3}$$

where $\boldsymbol{Y}_c^* \equiv [\boldsymbol{Y}'_1, \cdots, \boldsymbol{Y}'_{c-1}, \boldsymbol{Y}'_{c+1}, \cdots, \boldsymbol{Y}'_C]'$, $\boldsymbol{X}_c^* \equiv [\boldsymbol{X}'_1, \cdots, \boldsymbol{X}'_{c-1}, \boldsymbol{X}'_{c+1}, \cdots, \boldsymbol{X}'_C]'$, and $T_c^* \equiv T_1 + \cdots + T_{c-1} + T_{c+1} + \cdots + T_C$. Merging both likelihood functions (2) and (3) results in a panel VAR model where observations are pooled across countries and country specificities are washed out. As the focus is on inference at a country level, "full" pooling is not a sensible option. However, one can still interpret the likelihood in equation (3)

economic application, potential homogeneity between housing markets should be more relevant.

[7] Variables are rescaled according to their standard deviations in order to ensure meaningful magnitudes of impulse response functions.

as a priori information which may be appropriately discounted by a tightness parameter λ_c.[8] In other words, I propose a prior distribution for $\boldsymbol{B}_c$ and $\boldsymbol{\Sigma}_c$ that fulfills

$$p(\boldsymbol{\beta}_c, \boldsymbol{\Sigma}_c|\lambda_c) \propto \left[f(\boldsymbol{Y}^*_c|\boldsymbol{\beta}_c, \boldsymbol{\Sigma}_c)\right]^{\lambda_c}. \tag{4}$$

Note that the likelihood function on the right hand side corresponds to the one in equation (3), where $\boldsymbol{\beta}^*_c$ and $\boldsymbol{\Sigma}^*_c$ have been replaced by $\boldsymbol{\beta}_c$ and $\boldsymbol{\Sigma}_c$ respectively. In order to avoid double-counting of data, the prior distribution does not contain any observations with respect to country c. This ensures disjoint information sets of prior distribution and likelihood.

It has to be noted that the framework presented here still generates inference that is specific to the cross-sectional unit. The country-c model description in equation (1) or its respective likelihood function in equation (2) define the researcher's belief about the DGP. The introduction of a particular prior distribution does not question this belief. Rather it serves as a device to account for possibly insufficient sample sizes at the country level. Unspecified similarities between countries' models may prove useful in this context.

Following the mixed estimation strategy of Theil and Goldberger (1961), the prior distribution is implemented via dummy observations, i.e. by appending discounted observations to the sample of country c.[9] This procedure describes a partial pooling approach because merging country-c data with a non-country-c data is limited by λ_c. Section B of the appendix shows that the dummy observations prior may be expressed in terms of a conjugate normal-inverse-Wishart distribution to the likelihood function (2).[10] As a consequence, the corresponding posterior distribution inherits the well-established properties derived with respect to the conjugate prior distribution, e.g. it belongs to the same distributional family, sampling can be done in an efficient way, and a closed-form solution for the marginal data density exists.

The posterior distribution originates from the product of the likelihood (2) and the prior distribution (4). Following Del Negro and Schorfheide (2011), it is of a normal-inverse-Wishart form

$$\left(\boldsymbol{\beta}_c, \boldsymbol{\Sigma}_c\right)|\boldsymbol{Y}_c, \lambda_c \sim \mathcal{NIW}\left[\overline{\boldsymbol{\beta}}_c, \left(\overline{\boldsymbol{X}}'_c\overline{\boldsymbol{X}}_c\right)^{-1}, \overline{\boldsymbol{S}}_c, \overline{T}_c - n\cdot p\right], \tag{5}$$

with

$$\overline{T}_c \equiv T_c + \underline{T}_c, \quad \overline{\boldsymbol{Y}}_c \equiv \begin{bmatrix}\underline{\boldsymbol{Y}}_c\\ \boldsymbol{Y}_c\end{bmatrix}, \quad \overline{\boldsymbol{X}}_c \equiv \begin{bmatrix}\underline{\boldsymbol{X}}_c\\ \boldsymbol{X}_c\end{bmatrix},$$

$$\overline{\boldsymbol{B}}_c \equiv \left(\overline{\boldsymbol{X}}'_c\overline{\boldsymbol{X}}_c\right)^{-1}\left(\overline{\boldsymbol{X}}'_c\overline{\boldsymbol{Y}}_c\right), \quad \overline{\boldsymbol{\beta}}_c \equiv \text{vec}\left(\overline{\boldsymbol{B}}_c\right), \quad \overline{\boldsymbol{S}}_c \equiv \left(\overline{\boldsymbol{Y}}_c - \overline{\boldsymbol{X}}_c\overline{\boldsymbol{B}}_c\right)'\left(\overline{\boldsymbol{Y}}_c - \overline{\boldsymbol{X}}_c\overline{\boldsymbol{B}}_c\right),$$

and

$$\underline{T}_c \equiv T^*_c\lambda_c, \quad \underline{\boldsymbol{Y}}_c \equiv \boldsymbol{Y}^*_c\sqrt{\lambda_c}, \quad \underline{\boldsymbol{X}}_c \equiv \boldsymbol{X}^*_c\sqrt{\lambda_c}.$$

[8]Before estimating the model I standardize the variables to have mean zero and standard deviation one. It follows that the relative volatilities of the variables do not appear with respect to the specification of the prior distribution tightness parameter. Still, the standard deviations are considered with respect to impulse response analyses.

[9]In the context of multivariate time series models this technique has been applied by Sims and Zha (1998), amongst others.

[10]A similar representation of this prior distribution has been used by, e.g., Koop and Korobilis (2010).

The particular specification used in the present paper follows the exposition by Del Negro and Schorfheide (2003) and Del Negro and Schorfheide (2004). The authors construct a dummy observations prior distribution weighted by a tightness parameter. Making reference to Zellner (1971), they show that the corresponding posterior distribution is of a normal-inverse-Wishart form. Del Negro and Schorfheide (2011) argue that implementing priors via dummy observations helps to formulate reasonable a priori beliefs about the covariance matrix of the slope coefficients.[11] Section D shows that the slope coefficient covariance matrix is generally non-diagonal with respect to stationary VARs and a Monte Carlo exercise demonstrates the benefit not setting the respective off-diagonal elements to zero.

The posterior reduces to the likelihood of the country-specific VAR model if $\lambda_c = 0$ and it reduces to the likelihood of the pooled model if $\lambda_c = 1$.[12] Any convex combination in the range $[0, 1]$ constitutes a reasonable "intermediate" value. The Bayesian model averaging approach outlined in section 2.2 offers the possibility of estimating the prior tightness parameter λ_c within those bounds.

The specification of the prior distribution shares some similarities with the model of Jarociński (2010) as it shrinks the posterior distribution towards a cross-sectional average. However, the posterior distribution in equation (5) and the respective parameter definitions formulate a more general approach as (i) the prior covariance matrix of $\boldsymbol{\beta}_c$ is not restricted to be diagonal and (ii) the prior distribution is informative with respect to the residual covariance matrix $\boldsymbol{\Sigma}_c$ as well. A possibly minor advantage results from the dummy observations approach. Implementation of prior information becomes rather simple because the researcher only has to pool country c's observations with λ_c-discounted observations of the remaining countries instead of specifying each diagonal element of the prior covariance matrix of $\boldsymbol{\beta}_c$ separately.

The prior distribution shrinks the country-c model towards a discounted non-country-c average model. The non-country-c model is discounted by assuming a common tightness parameter λ_c for the remaining $C - 1$ countries. Alternatively, the researcher might want to define a separate tightness parameter for each country entering the prior distribution. This would allow the role to be discriminated for each country separately. While conceptually straightforward, this approach adds additional layers of complexity to the procedure because the distribution of the tightness parameters becomes multidimensional. Such an approach would represent a separate research question and is therefore not pursued any further at this point.

2.2 Bayesian model averaging

In terms of the BMA concept, the choice of a particular value for the hyperparameter λ_c defines an econometric model. Probabilities are assigned to each model, i.e. to each value

[11] The proposed dummy observations prior is a "data-based" quantity because it is constructed from actual observations. In this respect it shares similarities with the so-called g-prior proposed by Zellner (1986). Setting the slope covariance proportional to the observed $\boldsymbol{X}_c'\boldsymbol{X}_c$ matrix, off-diagonal elements are generally non-zero.

[12] Note that for too loose priors, i.e. sufficiently small values of λ_c, the prior distribution is not well-defined. For this reason, the prior on λ_c assigns zero probability to that range. Section 2.2 provides a more formal description and section F of the appendix demonstrates results with an alternative prior specification.

of λ_c, considered. Model statistics which depend on the choice of λ_c are weighted according to their respective model probabilities. A prior distribution on the hyperparameter is combined with the marginal data density[13] in order to obtain a posterior distribution over the models indicated by λ_c. The posterior distribution constitutes the models' weights in the BMA procedure. The hierarchical approach described above allows the researcher to treat hyperparameters as objects to be estimated.

In brief, the marginal data density measure is interpreted in terms of the likelihood and is combined with a (hyperprior) distribution over λ_c. The posterior distribution fulfills

$$p\left(\lambda_c|\boldsymbol{Y}_c\right) \propto p\left(\boldsymbol{Y}_c|\lambda_c\right) \times p(\lambda_c).$$

Conveniently, a closed-form representation of the marginal data density corresponding to the posterior (5) is available and documented by Del Negro and Schorfheide (2011); it reads

$$p\left(\boldsymbol{Y}_c|\lambda_c\right) = (2\pi)^{-nT_c} \frac{\det\left(\overline{\boldsymbol{X}}_c'\overline{\boldsymbol{X}}_c\right)^{-\frac{n}{2}} \det\left(\overline{\boldsymbol{S}}_c\right)^{-\frac{\overline{T}_c-k}{2}}}{\det\left(\underline{\boldsymbol{X}}_c'\underline{\boldsymbol{X}}_c\right)^{-\frac{n}{2}} \det\left(\underline{\boldsymbol{S}}_c\right)^{-\frac{\underline{T}_c-k}{2}}} \times \frac{2^{\frac{n(\overline{T}_c-k)}{2}} \prod_{i=1}^{n} \Gamma\left(\frac{\overline{T}_c-k+1-i}{2}\right)}{2^{\frac{n(\underline{T}_c-k)}{2}} \prod_{i=1}^{n} \Gamma\left(\frac{\underline{T}_c-k+1-i}{2}\right)}.$$

The existence of an exact closed-form solution for the marginal data density makes the hierarchical BMA approach tractable and accurate. It remains to specify the hyperprior distribution. As the two polar cases $\lambda_c = 0$ and $\lambda_c = 1$ limit the hyperparameter in a reasonable way, it seems sensible to choose the prior of λ_c to follow a Beta distribution

$$\lambda_c \sim Beta\left(\alpha, \beta\right).$$

The support of the Beta distribution features $\lambda_c \in [0, 1]$ which implicitly bounds the posterior draws of the tightness parameter to this range.[14] In order to impose a relatively uninformative prior, $\alpha = \beta = 1$ is chosen for the analysis. In fact, this specification reduces the Beta distribution to a uniform distribution bounded between zero and one.

The posterior distribution of the tightness parameter $p(\lambda_c|\boldsymbol{Y}_c)$ is used to weight a particular posterior statistic of interest $S(\boldsymbol{B}_c, \boldsymbol{\Sigma}_c)$. Inference is based on the average of the weighted statistic

$$p(S[\boldsymbol{B}_c, \boldsymbol{\Sigma}_c]|\boldsymbol{Y}_c) = \sum_{\lambda_c} p\left(\lambda_c|\boldsymbol{Y}_c\right) \times p\left(S[\boldsymbol{B}_c, \boldsymbol{\Sigma}_c]|\boldsymbol{Y}_c, \lambda_c\right). \quad (6)$$

In particular, a fixed number of draws from the posterior distribution of the VAR model depending on the tightness parameter are generated. In the following, based on the draws, the statistics of interest, e.g. percentiles of impulse response functions, are calculated. The final estimate of a particular percentile of the impulse response function is computed as

[13]The marginal data density is a common measure of fit in Bayesian analyses. Jarociński and Maćkowiak (2013) provide a discussion on the appropriateness of the marginal data density and alternative measures. A later version of this study (see Jarociński and Maćkowiak, 2017) omits the respective paragraph.

[14]In order to yield a well-defined prior distribution it is required that $\lambda_c \geq (p+1)n/T_c^*$. As a consequence, the indicator function $\mathbb{I}(\lambda_c \geq [p+1]n/T_c^*)$ which equals one for "large enough" values of λ_c and zero otherwise is multiplied to the probability density function of the Beta distribution to yield $p(\lambda_c)$.

the (posterior model probability) weighted average over a predefined set of possible values for λ_c. BMA describes a sensible treatment because it considers uncertainty with respect to the "correctness" of the model.[15]

2.3 Structural form

The reduced-form model of equation (1) is not sufficient to draw clear economic conclusions. I am more interested instead in the effects of structural shocks and the corresponding structural model

$$\boldsymbol{A}_{c,0}\boldsymbol{y}_{c,t} = \sum_{j=1}^{p} \boldsymbol{A}_{c,j}\boldsymbol{y}_{c,t-j} + \boldsymbol{u}_{c,t}, \quad \text{for } t = 1, \ldots, T_c, \tag{7}$$

where $\boldsymbol{A}_{c,j} \equiv \boldsymbol{A}_{c,0}\boldsymbol{B}_{c,j}$ and $\boldsymbol{u}_{c,t} \equiv \boldsymbol{A}_{c,0}\boldsymbol{e}_{c,t}$ with $\boldsymbol{u}_{c,t} \sim \mathcal{N}(\boldsymbol{0}_{n,1}, \boldsymbol{\mathcal{I}}_n)$.[16] In order to identify a structural house price bust shock, a combination of sign and zero restrictions is imposed on the impulse response functions.

The exceptional importance of expectations with respect to house price developments has been shown by empirical and theoretical analyses.[17] Similarly, the ESRB (2016) states that boom-bust housing cycles may be particularly fueled by excess optimism. For instance, Lambertini et al. (2017) argue that, in the run-up to the crisis, optimistic expectations led to an increased demand for housing, feeding into a rise in house prices. Borrowing capacity improved along with mortgage financing, thus ultimately causing the house price boom to become self-reinforcing. Similar forces may operate along downturns.

The above considerations are summarized by identifying assumptions that include a sudden drop in confidence in future economic developments accompanied by a decline in house prices. The corresponding innovation is denoted a "house price bust" shock. Table 1 depicts the identifying assumptions of the baseline shock (S1) as well as of alternative versions (S2)-(S4).[18] Other variables in the system remain unaffected on impact, which orthogonalizes the shock from standard economic driving forces like aggregate demand or aggregate supply shocks. In fact, the identified house price shock may equally well be called a particular housing demand shock that originates from households' expectations. The fact that the baseline identification scheme (S1) contains the assumption of a decline in households' economic outlook defines a very particular type of a housing demand shock which is expected to have more severe consequences for the economy compared to an ("ordinary") shock that does not contain this additional assumption (S3). The shock is said to capture a decline in prices which may have the potential to cause a deep recession

[15] There are two alternatives to the model averaging approach. First, the researcher selects the model at the posterior mode of λ_c ("Bayesian Model Selection"). If inference is based on the posterior mode, it rests upon a generally incorrect point estimate of λ_c. Second, the researcher samples the pair $(S[\boldsymbol{B}_c, \boldsymbol{\Sigma}_c], \lambda_c)$ in order to obtain draws from the respective joint distribution ("MC^3"). Koop (2003, ch. 11) mentions that MC^3 is useful if the number of models considered is too huge which is not the case with respect to the present application.

[16] $\boldsymbol{0}_{n,k}$ denotes an $n \times k$ matrix of zeros and $\boldsymbol{\mathcal{I}}_n$ denotes the $n \times n$ identity matrix.

[17] See for instance Piazzesi and Schneider (2009), Lambertini et al. (2013), Towbin and Weber (2016), and Lambertini et al. (2017).

[18] Note that, while generally possible, I do not identify the shocks jointly. Rather, depending on the type of analysis, I select one of the three alternatives and estimate the respective structural model.

	Shock	$g_{c,t}$	$r_{c,t}$	$x_{c,t}$	$p_{c,t}$	$z_{c,t}$
(S1)	House price bust (baseline)	0	0	↓	↓	0
(S2)	House price bust (flexible rate)	0	↓	↓	↓	0
(S3)	House price (w/o outlook)	0	0	0	↓	0
(S4)	House price (agnostic)			unrestricted		

Table 1: The following abbreviations are used: $g_{c,t}$: gross domestic product, $r_{c,t}$: mortgage rate, $x_{c,t}$: economic outlook, $p_{c,t}$: real residential property prices, $z_{c,t}$: consumption.

or a even a crisis. Since it depends on households' expectations with respect to future economic conditions, the shock may also be labeled a "confidence shock" which should be related to actual future economic activity.

The identifying assumptions on the structural shock represent very specific beliefs about the house price shock. In order to define an agnostic benchmark, (S4) represents a house price decline without any identifying assumptions. In particular, this specification resembles a shock in a recursively identified VAR model, where house prices are ordered first. Equivalently, it is akin to the so-called "generalized" impulse response functions approach proposed by Pesaran and Shin (1998). This measure has the advantage of not requiring arguably subjective identifying assumptions. Given the model, the impulse responses of all variables are computed from the data alone and do not contain further beliefs of the researcher. The flipside of this agnostic measure is that it produces responses of shocks without a proper economic interpretation. Even though a meaningful assessment is rarely appropriate, I use (S4) as a benchmark in order to evaluate the responses of the baseline identification scheme.

From a technical perspective, sign and zero restrictions are imposed on $\boldsymbol{A}_{c,0}$ (see e.g. Uhlig, 2005, 2017). Given draws from the reduced-form VAR, Arias, Rubio-Ramirez, and Waggoner (2014) [19] propose an efficient way to obtain rotations of the model that satisfy a set of sign and zero restrictions. In section 2.1 it was shown that the posterior distribution of the reduced-form model (1) is standard. Sampling follows Del Negro and Schorfheide (2011) who describe an algorithm with respect to the matricvariate-normal-inverse-Wishart distribution.[20]

Algorithm (Reduced-form sampling and identifying restrictions)**.**

1. *Draw* $\boldsymbol{\Sigma}_c^{(s)} \sim \mathcal{IW}\left(\overline{\boldsymbol{S}}_c, \overline{T}_c - k\right)$.
2. *Draw* $\boldsymbol{\beta}_c^{(s)} | \boldsymbol{\Sigma}_c \sim \mathcal{N}\left(\overline{\boldsymbol{\beta}}_c, \boldsymbol{\Sigma}_c^{(s)} \otimes \left(\overline{\boldsymbol{X}}_c' \overline{\boldsymbol{X}}_c\right)^{-1}\right)$.
3. *Draw a proposal for* $\boldsymbol{A}_{c,0}^{(s)}$ *according to Arias et al. (2014).*

[19] Arias, Rubio-Ramirez, and Waggoner (2018) clarify that the algorithm in the earlier version of their paper is not invariant to the ordering of structural shocks. Since I identify only one structural shock for each model this disadvantage should not apply at this point.

[20] Equivalently, the distributions in terms of the normal as opposed to the matricvariate-normal distribution may be formalized. In fact, draws from the normal distribution can be rearranged to serve as draws from the respective matricvariate-normal distribution.

4. *Generate impulse response functions based on zero restrictions and a candidate* $\{\boldsymbol{\Sigma}_c^{(s)}, \boldsymbol{\beta}_c^{(s)}, \boldsymbol{A}_{c,0}^{(s)}\}$. *If sign restrictions are fulfilled, keep the candidate and iterate* s, *otherwise discard the draw and generate another candidate.*

5. *Return to 1. until the required number of draws has been obtained.*[21]

3 Data

The model considered is estimated on quarterly data for a set of 16[22] euro area countries. The maximum time span ranges from 2003:Q1 to 2019:Q1 but is not available for all countries in the sample. It is important to note that the sample size varies over countries as the number of observations is relevant with respect to the discussion of the prior tightness.[23] Table 2 in the results section provides an overview.

Importantly, the definition of the time series should be as consistent as possible across countries. Variables considered for the model are expected to illustrate the effects of sudden house price shocks at a macroeconomic level. Obviously, a description of the household sector is crucial with respect to the transmission of house price shocks to GDP. Consequently, variable selection particularly focuses on the behavior of this sector at an aggregate level. The system is generally conformable with the empirical macro literature (see e.g. Calza et al., 2013 or Ahamada and Sanchez, 2013).

The analyses employ publicly available series of real residential property prices from the Bank for International Settlements (BIS) Residential Property Price database. While the respective original data sources may differ and the definitions of the time series do not match perfectly across countries, they constitute a relatively homogeneous collection of house price series. With respect to the time series properties, a notable degree of cross-sectional heterogeneity is identified. The column entitled "$100 \times \sigma_c$" of table 2 in section 4 covers the standard deviation as a measure of the series' volatility. It shows considerable heterogeneity, ranging from 0.97 in Germany to 6.14 in Lithuania. The role of the consumer outlook has been discussed in section 2.3. The respective data is taken from the OECD and covers consumer expectations over the future tendency of the general economic situation. As mentioned in section 2.3, expectations play a crucial role with respect to sudden price drops. The indicator represents the balance of answers to a survey and is generally bounded between -100 and 100. Finally, the object of study is the impact of house price bust shocks on output or consumption. In the analysis they serve as a measure of damage to the economy if sudden price drops occur. Data on real GDP and real individual consumption expenditures is taken from the European Central Bank (ECB) Statistical Data Warehouse. The mortgage rate is taken from the Monetary Financial Institutions (MFI) Interest Rate Statistics of the ECB which is provided by the ECB Statistical Data Warehouse as well.

[21] For each value of λ_c 1,000 draws from the posterior distribution are generated. This amounts to a total of 100,000 draws for country c to be considered for the model averaging procedure.

[22] The selection of euro area countries is based on the availability of time series data.

[23] In general it is not necessary for the prior to cover the same sample as the likelihood. With respect to Bayesian analyses it is common practice in applied research (see e.g. Primiceri, 2005 or Jarociński, 2010) to employ a pre-sample analysis in order to construct or calibrate the prior distribution. This method has also been mentioned by Koop (2003, ch. 12) and Canova (2007, ch. 10).

A detailed description of the time series may be obtained from the respective data sources. Data identifiers are provided in section A of the appendix. Except for the mortgage rate and the consumer outlook indicator, natural logarithms are applied to the variables before calculating first differences. The consumer outlook series are divided by 100 in order to rescale their magnitude.

4 Results

4.1 Prior tightness

The tightness parameter λ_c is an important component of the model because it sets the degree of partial pooling with the cross section. The hierarchical approach generates as a by-product the corresponding posterior distribution. The respective prior distribution reflects the belief that λ_c should be located in the range between zero and one. The lower bound represents a very loose prior on $\boldsymbol{B}_c$ and $\boldsymbol{\Sigma}_c$ that does not include any observations of other countries, and the upper bound represents a tight prior which fully incorporates other countries' observations. The posterior results with respect to λ_c either contain valuable information about the informativeness of cross-country mechanics for each country-specific VAR model or reflect a degree of cross-country heterogeneity. In fact, it is not possible to disentangle whether country c's data is relatively informative such that cross-sectional information is not required or whether the country-c model differs strongly from the (non-country-c) pooled model.

Figure 1 illustrates the mode and two credible sets of the posterior distribution of λ_c.[24] The boxplots document considerable differences across countries. In fact, there are some countries for which the model devotes relatively small weights to average cross-country information like Finland, France, Germany, Latvia, and Luxembourg with mode estimates below 0.2.[25] At the other extreme, the VAR models of Estonia, Greece, Portugal, and Spain require very tight priors, which shifts their respective posterior results towards the pooled VAR model. In fact, the upper bound of the hyperprior distribution seems to be binding for three countries in the sample. Still, values of $\lambda_c > 1$ are hardly interpretable and, while econometrically possible, relaxation of the restrictive bounds would appear to be unreasonable. The posterior distribution of λ_c seems to vary with the degree the respective countries' housing market has been hit by the global financial crisis. Stylized facts of Rodrigues and Lourenço (2015) show that countries like Spain, Greece, the Netherlands, or Portugal experienced declines in terms of the price-to-rent ratio or the price-to-income ratio whereas for Austria, Belgium, Germany, Finland, and France those figures remained fairly stable. Ireland[26] marks an exceptionally low estimate of λ_c while

[24] With respect to the lag length, the analyses generally assume the baseline specification $p = 4$ to cover at least one year. An additional exercise in section 4.4 demonstrates the impact of p on the tightness parameter.

[25] It has to be noted that credible intervals cannot reach zero because I assign zero probability to models where λ_c is too small, i.e. where the prior distribution is not well-defined.

[26] It is well-known that the time series of Irish GDP growth rates exhibits an unusual outlier value. Notes of the European Commission (see https://ec.europa.eu/eurostat/documents/24987/6390465/Irish_GDP_communication.pdf) and the OECD (see http://www.oecd.org/sdd/na/Irish-GDP-up-in-2015-OECD.pdf) state that the reason for this data anomaly follows from a relocation with respect to multinational corporations. Results from a robustness exercise, where the exceptional

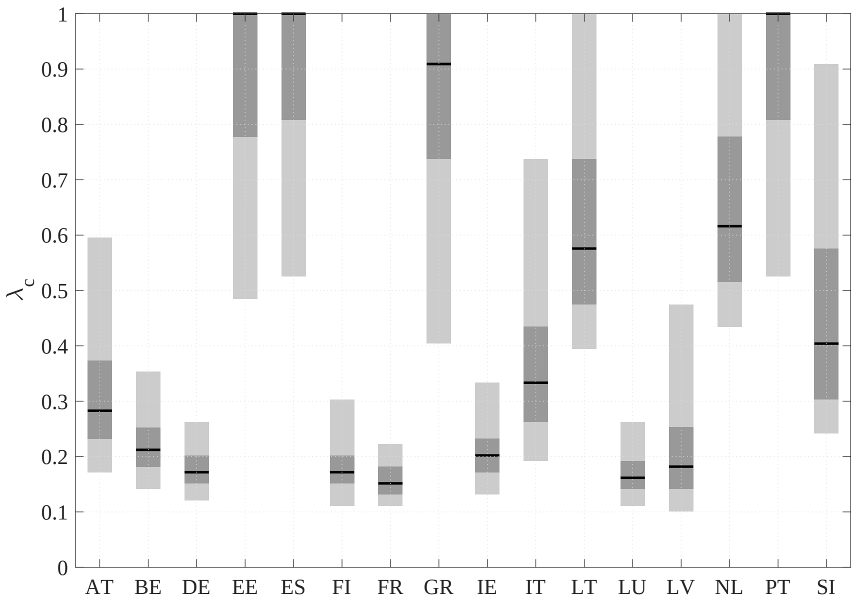

Figure 1: Posterior distribution of the tightness parameters λ_c across countries. The gray shaded areas denote the 68 percent and the 90 percent Bayesian (highest probability density) credible intervals, respectively.

the country has experienced considerable declines in both ratios.

In a similar context, Calza et al. (2013) focus on mortgages and distinguish country groups along a set of mortgage market indicators. As results depicted in figure 1 hardly align with their heterogeneity pattern, λ_c does not (or at least not exclusively) seem to capture differences in the mortgage market. My results resemble the ones of Koop and Korobilis (2016) in the sense that findings with respect to model homogeneity cannot be related to "simple narratives" that divide the euro zone into core and periphery regions. The ESRB (2016) confirms that it is difficult to relate country-specific results to certain parameters. One reason may be that interaction of different aspects is important. Moreover, as hinted above, λ_c does not only capture the degree of heterogeneity,[27] it also reflects the relative informativeness of country-c data compared to the rest of the sample. In other words, if the estimate of λ_c is small, country c's model either is very different from the other countries' or its data is sufficiently informative.

Table 2 provides further description regarding the prior importance. The column entitled T_c lists the number of actual observations for the respective country while column $\overline{T}_c^m$ contrasts the number of artificial dummy observations[28] used for the construction of the prior distribution

$$\overline{T}_c^m = \lambda_c^m T_c^*.$$

growth rate is replaced by the sample mean, are virtually indistinguishable from the baseline analysis. The respective figures are available upon request.

[27] Note that the unqualified interpretation of λ_c in terms of a measure of homogeneity is not correct because samples are not fully congruent across countries. The corresponding bias in interpretation should be more relevant if the number of observations was considerably lower than average (see table 2). It is still valid to interpret the tightness parameter in terms of the prior distribution informativeness.

[28] In contrast to actual observations, I define the number of artificial dummy observations as a continuous variable. The point estimate is calculated as the product of the mode of λ_c and the actual sample size of the dummy observations T_c^*.

	$100 \times \sigma_c$	T_c	$\overline{T}_c^m$	$\overline{T}_c^m / T_c$	λ_c^m
AT	2.22	60	216.08	3.60	0.28
BE	1.37	60	162.06	2.70	0.21
DE	0.97	60	131.19	2.19	0.17
EE	5.87	52	772.00	14.85	1.00
ES	2.54	60	764.00	12.73	1.00
FI	1.31	60	131.19	2.19	0.17
FR	1.76	60	115.76	1.93	0.15
GR	2.23	23	728.18	31.66	0.91
IE	3.46	60	154.34	2.57	0.20
IT	1.12	60	254.67	4.24	0.33
LT	6.14	53	443.91	8.38	0.58
LU	1.50	43	126.22	2.94	0.16
LV	2.98	29	144.55	4.98	0.18
NL	1.58	60	470.75	7.85	0.62
PT	2.04	40	784.00	19.60	1.00
SI	2.46	44	315.15	7.16	0.40

Table 2: Descriptive statistics of the data and the cross-country VAR models. σ_c is the standard deviation of the house price growth series $p_{c,t}$, T_c denotes the sample length of country c, and $\overline{T}_c^m$ denotes the effective number of prior dummy observations at the mode λ_c^m.

The superscript "m" indicates the posterior mode estimate. Apparently, even countries with relatively low estimates of λ_c feature quite a substantial degrees of cross-country prior information. Consider the case of Finland, which receives a relatively loose prior tightness parameter but still contains more than twice as many observations from the cross section compared to country c's sample size. One might explain this finding by a stable relationship between house prices and the macro variables across countries. It seems that the general economic mechanism prevails in the sample. While it is not straightforward to compare the numerical estimate of λ_c with the corresponding measures of heterogeneity in Jarociński (2010), his estimates seem to support considerable degrees of partial pooling as well. Unfortunately, Canova (2005), who employs a similar approach, does not cover a detailed analysis with respect to the consequences of heterogeneity on his results.

In order to illustrate the impact of the prior tightness on impulse response functions, figure 2 depicts impulse response functions for each value of λ_c considered in the range $[0, 1]$. The graphs show cumulated impulse response functions of GDP growth for the case of Austria and Italy. The countries serve as representative examples.[29] Darker lines indicate $\lambda_c \to 1$, i.e. responses that feature a tighter prior. In order not to overload the graph, the focus lies on the median impulse response function over the posterior draws. Note that a fully loose prior distribution ($\lambda_c = 0$) generates impulse response functions of a single-country maximum likelihood VAR model. In fact, for Austria the prior tightness has a marked effect on the dynamics as the reaction of output becomes to some extent

[29]The two countries are chosen in order to show distinctive patterns with respect to the impact of the tightness parameter on the simulation results. Selection is meant to avoid any suggestive motives. Results with respect to other countries are available upon request.

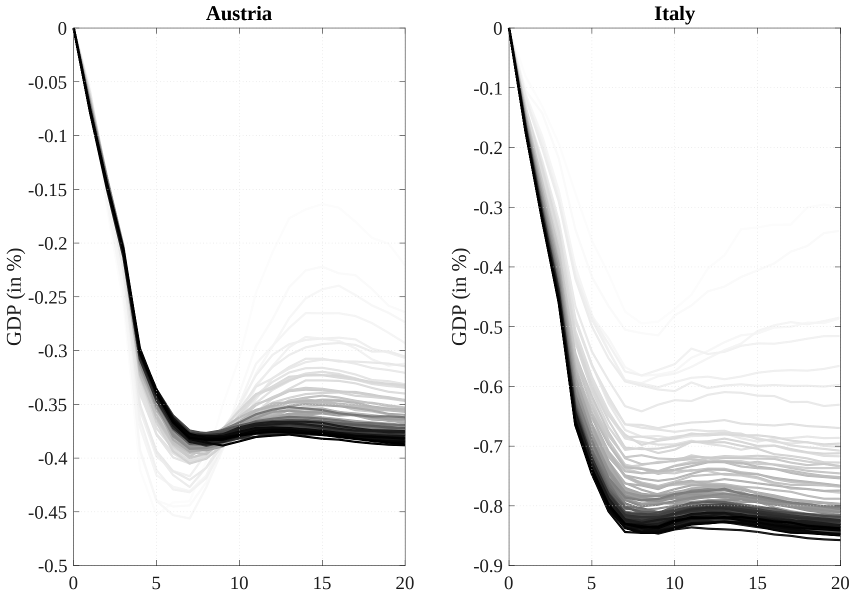

Figure 2: Median impulse response functions for $\lambda_c \in [0, 1]$ for the examples of Austria and Italy. Light gray lines indicate small values (closer to single country model) and darker lines indicate greater values (closer to pooled model) of λ_c.

more gradual. It seems that tighter priors induce stability in the response dynamics, which enhances the reliability of the results. As the tightness parameter approaches its maximum, the model converges to the pooled VAR rescaled with respect to the standard deviations of Austria. With respect to the second example, Italy, the prior works in another direction and amplifies the estimated effects. Still, responses seem to converge to more stable dynamics. Figure 2 should be regarded as an illustrative example for the general impact of the tightness parameter. While the quantitative effects might be diverse across countries, the feature of stabilizing the responses prevails.

A sensible prior distribution is designed to reduce estimation uncertainty. In order to evaluate the usefulness of my approach, I attempt to quantify the degree to which cross-sectional correlations improve upon the precision of impulse response function estimates. Relatively wide credible intervals indicate dispersed posterior distributions and vice versa. I construct a precision measure by computing the distance between predefined percentiles of the posterior distributions of the respective impulse response functions. In fact, I compute the 16 to 84 ($q = 68$) as well as the 5 to 95 ($q = 90$) interpercentile ranges IPR_q of impulse response functions with respect to output and consumption $h = 5$ quarters after the shock

$$\Delta(y_{j,t+h}) \equiv 100 \times \left[IPR_q\left(y_{j,t+h}|\lambda_c = \lambda_c^m\right) - IPR_q\left(y_{,jt+h}|\lambda_c = 0\right)\right].$$

Figures 3 and 4 depict the percentage point difference Δ between one standard deviation shock responses at the posterior mode λ_c^m and a model estimated without cross-sectional prior information, i.e. where $\lambda_c = 0$.[30] Negative values indicate that credible intervals of the model with prior information are tighter than the respective model estimated on country-specific data only. Results strongly indicate that cross-country information helps

[30]The model with $\lambda_c = 0$ essentially reflects the respective maximum likelihood estimate.

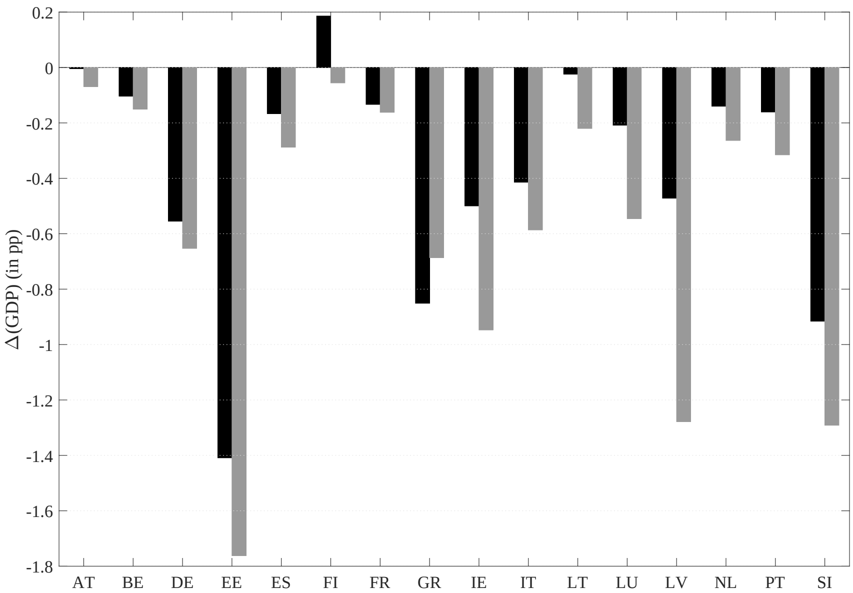

Figure 3: Difference Δ between interpercentile ranges of posterior impulse response functions 4 quarters after the shock with respect to GDP. Black bars correspond to the 68 percent and gray bars to the 90 percent Bayesian credible intervals, respectively.

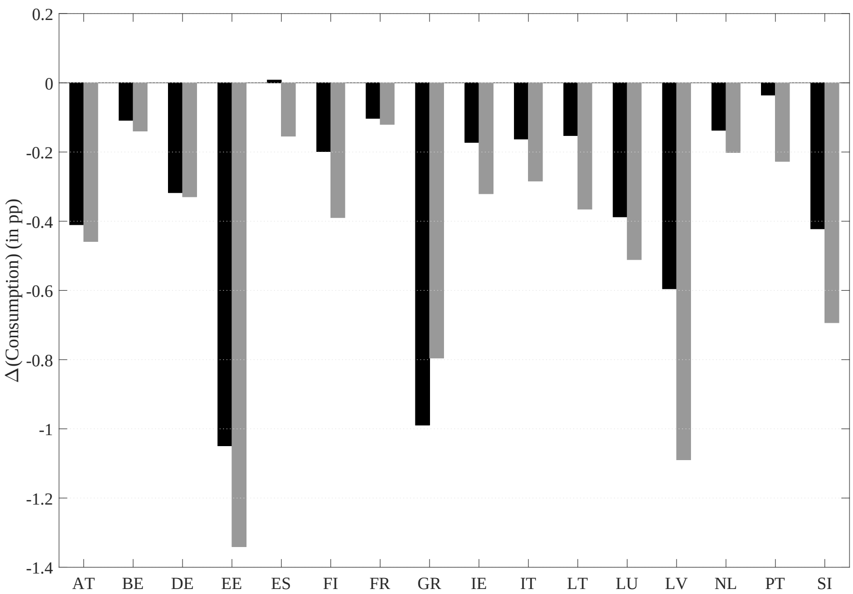

Figure 4: Difference Δ between interpercentile ranges of posterior impulse response functions 4 quarters after the shock with respect to GDP. Black bars correspond to the 68 percent and gray bars to the 90 percent Bayesian credible intervals, respectively.

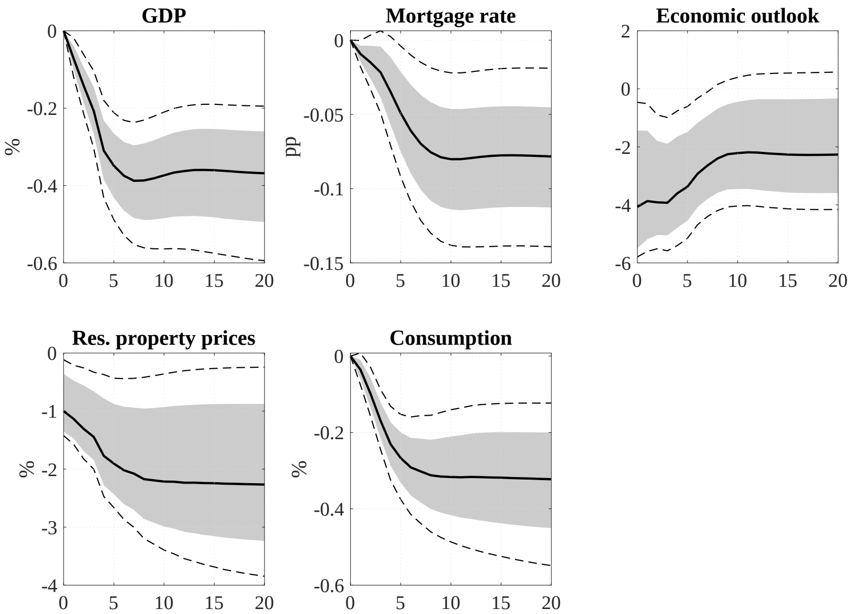

Figure 5: Quarterly cumulative impulse response functions for Austria. Solid lines depict medians, gray-shaded areas denote 68 percent and dashed lines 90 percent Bayesian credible intervals, respectively.

to improve precision with respect to estimates of impulse response functions. Accordingly, the vast majority of countries exhibit tighter credible intervals. Improvements are more pronounced with respect to the 90 percent credible interval.

Figures 3 and 4 refer to one particular horizon only. Inspection of the complete patterns of Δ[31] shows that the advantage with respect to the precision tends to improve with horizon h.

4.2 Adverse shocks to house prices

The impact of house price bust shocks is investigated with an impulse response analysis. Weighted averages of median response functions are reported for the baseline house price bust shock (S1) defined in table 1. In particular, a predetermined grid that defines possible values of λ_c is used and 1,000 draws from the structural VAR impulse response functions are generated for each λ_c. According to equation (6), percentiles of responses are calculated and averages across λ_c are computed using weights obtained from the respective posterior distribution $p(\lambda_c|\boldsymbol{Y}_c)$ as discussed in section 4.1. The shock hits the economy in period 0. The shapes of the impulse response functions are very similar across countries. Therefore for a qualitative discussion, it shall suffice to analyze a particular country example in figure 5. It depicts responses of all variables in the VAR system with respect to the model of Austria. Since the model is estimated in first differences, quarterly impulse responses of all model variables are cumulated over the horizon considered. The house price bust shock is rescaled such that the median drop in house prices on impact is

[31]Results are not depicted but are available upon request.

one percent.[32]

The identified shock has a negative impact on output and consumption. Results show that it induces a gradual decay of real GDP over the first year. The responses subsequently persist at the new and lower level. Theoretically, a drop in house prices should have a direct effect on household wealth and probably on household lifetime wealth as well. The latter crucially determines the potential of households to obtain credit to finance their expenditures. In addition, it increases lenders' losses in the case of a default which may activate a precautionary savings motive. Consequently, a weakened wealth situation should reduce aggregate consumption and housing investment (ESRB, 2016). Aggregate demand contracts. Accordingly, consumption shrinks for about 8 quarters and remains at the new lower level. A decrease in house prices reduces households' nominal housing wealth which presumably limits consumption possibilities. Mortgage rates show a tendency to decrease subsequently. The economic outlook worsens according to the restriction and improves along the first two years. In the following, the response, surprisingly, remains at a level below zero. It may not fully revert to zero due to its positive relationship to GDP. Results are robust and quite homogeneous across countries. Credible sets of GDP and consumption responses are almost exclusively different from zero, which emphasizes the unambiguity of the effects. During the dynamic simulation experiment, house prices subsequently tend to fall further, which may in turn accelerate the responses of the other variables. This finding is important because it documents considerable second-round effects in this context. Furthermore, the house price bust shock features a decline in consumers' economic outlook. Section 4.4 shows that this particular sign restriction intensifies the amplitude of the impulse response functions.

While the general pattern of the responses is very similar across countries, quantitative differences are to be noticed. Figure 6 condenses information from impulse response functions with respect to GDP. It shows percentiles of cumulated responses after one year.[33] Results illustrate heterogeneity with respect to a one-percent house price shock. Responses of output are similarly sharp for about 2/3 of the countries considered. The response with respect to Finland, Germany, Greece, Ireland, and Luxembourg are exceptionally different as they feature a decline in GDP of more than one percent after one year. The majority of countries exhibit more moderate effects of between about -0.2 and -0.6 percent. GDP reactions are distinctly negative as none of the credible sets considered contains zero.

The magnitude of the GDP reaction is relatively pronounced compared to what one would expect with respect to an "ordinary" housing demand shock.[34] The identified house price bust shock, as noted in section 2.3, contains a decline in households' expectations on future economic developments. In this respect, the responses do capture an effect on consumers' confidence as well, which may have an additive negative impact on future

[32] Due to the linear nature of the model, effects of positive house price shocks are easily obtained by reversing the impulse response functions.

[33] Section C in the appendix illustrates the effects after one quarter. One quantifies a relevant amplification effect over time but the cross-country comparison yields results comparable to those after one year.

[34] Berger, Guerrieri, Lorenzoni, and Vavra (2018) note that, compared to the effects predicted by theoretical models, empirical estimates of the elasticity of consumption to changes in house price are quite large. Case, Quigley, and Shiller (2013) find that theoretical models that account for the relevant frictions predict much higher elasticities.

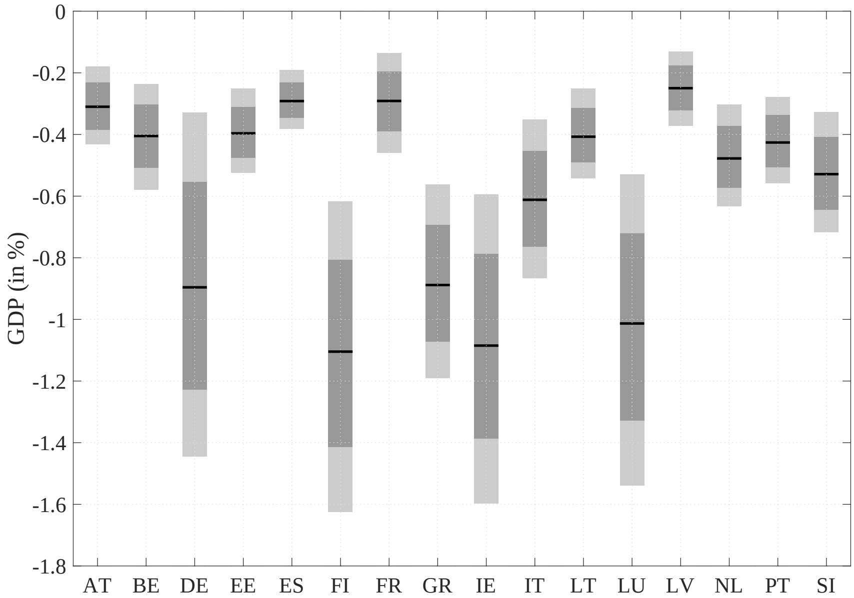

Figure 6: Impulse response functions of GDP 4 quarters after the shock. Solid lines depict medians and the gray shaded areas denote 68 percent and 90 percent Bayesian credible intervals, respectively.

economic activity. Section 4.4 sheds some light on the role of consumers' expectations.

Comparing the estimation uncertainty in figure 6 with the posterior distributions of λ_c in figure 1 suggests conclusions with respect to the informativeness of country-specific data. Consider Finland, Germany, and Luxembourg, which exhibit loose prior distributions (small values of λ_c) that may either indicate that those countries' models are relatively different from the "average" model or that their data is highly informative. Considering their respective impulse response functions' credible sets, which are relatively wide, rejects the latter hypothesis and rather indicates less informative country-specific data in those cases.

A similar conclusion may be drawn when considering the responses of consumption to the house price shock, which are depicted in figure 7. While the general effect on consumption is negative across all countries, the quantitative effect differs along similar groups of countries. Again one observes quite strong drops for Finland and Greece in excess of -0.75 percent. The effect seems to be about three to four times higher than for countries with more moderate responses ranging between about -0.25 and -0.5 percent.

Instead of simulating one-standard-deviation shocks, the above analyses consider impulse response functions rescaled to a negative one-percent decrease in house prices. The transformation makes the shocks comparable across countries but disregards the likelihood of the shocks occurring. Assuming a normal distribution, one-standard deviation shocks appear with a probability of about 15.87 percent. Since negative one-percent shocks may be more or less likely than the standard deviation shock, figure 8 contrasts the percentile of a negative standard deviation with the respective percentiles for the shocks (S1), (S3), and a negative one-percent quarterly growth rate of house prices. This illustration shows that the house price bust shock is quite unlikely to occur for some countries. For example, the German negative one-percent house price bust shock induces large declines in economic activity but seems to be rather rare. Furthermore, the figure confirms that

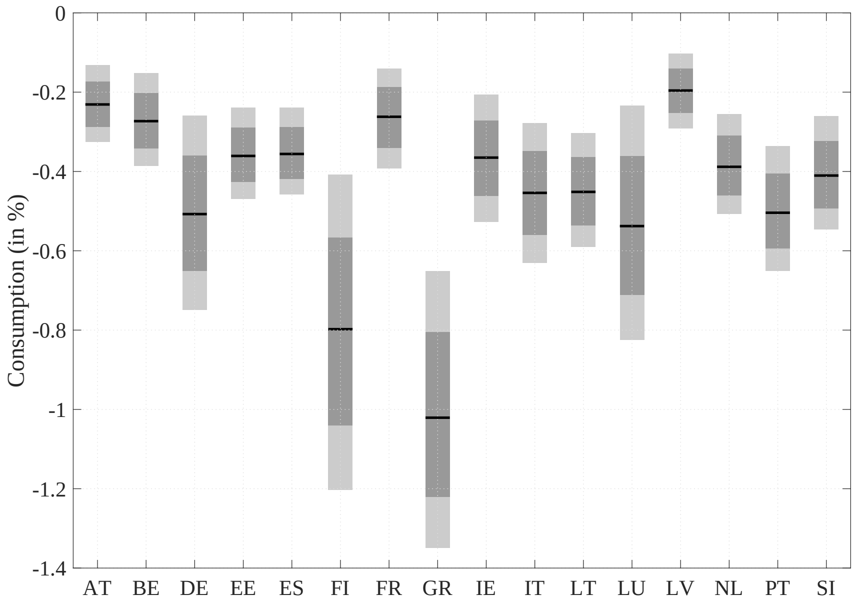

Figure 7: Impulse response functions of consumption 4 quarters after the shock. Solid lines depict medians and the gray shaded areas denote 68 percent and 90 percent Bayesian credible intervals, respectively.

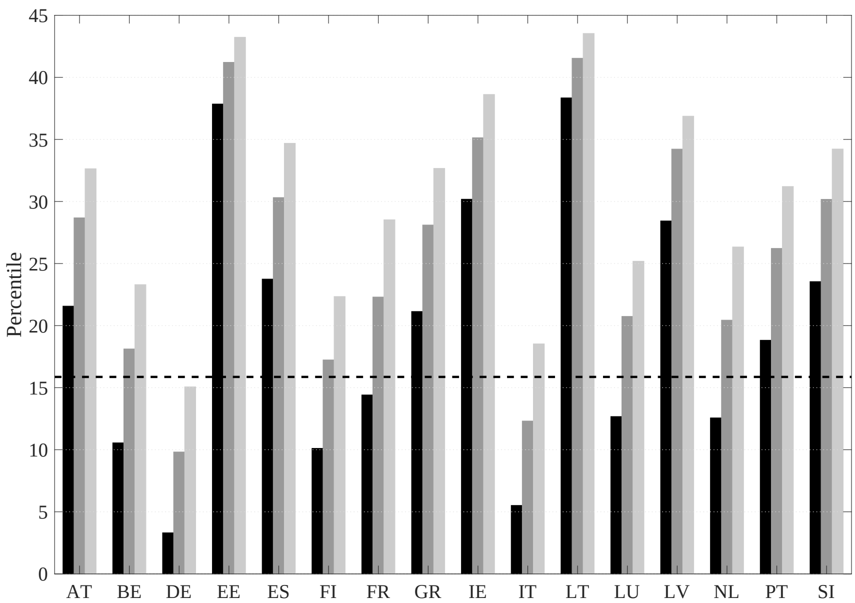

Figure 8: Black bars indicate the percentiles of a negative one-percent shock with respect to the baseline identification scheme (S1), dark gray bars indicate respective percentiles for the identification scheme (S3) and the light gray bars indicate the percentile of a negative one-percent quarterly growth in house prices assuming normally distributed growth rates. The dashed line represents one standard deviation.

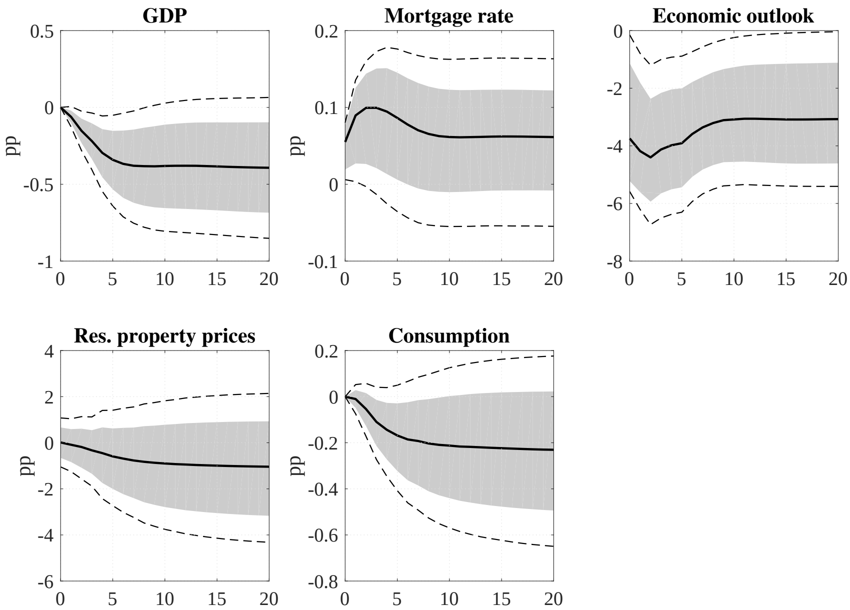

Figure 9: Amplification effect for Austria. Solid lines depict medians, gray shaded areas denote 68 percent and dashed lines 90 percent Bayesian credible intervals, respectively.

the identification scheme (S1; black bars) which contains a decline in consumer economic outlook is generally less likely than the shock with a zero restriction on that variable (S3; dark gray bars). Both shocks are also less likely than an almost unrestricted[35] fall in house prices (light gray bars).

4.3 Restriction on the mortgage rate

It stands to reason that the mortgage rate has a particular effect on house price reactions. Episodes of low interest rates constitute easy financing conditions and allow borrowers to afford more expensive homes which, in turn, may drive up house prices. Conversely, a monetary authority can use the short-term interest rate as a policy instrument in order to dampen economic downturns. In a situation where the nominal interest rates and the respective policy instrument are close to the zero lower bound, the potential to stabilize the economy becomes limited. In this respect, one may argue that the zero restriction on the mortgage rate is crucial for the magnitude of the output and consumption responses. This leads to a quantification of an amplification effect that may arise from the fact that interest rates are not allowed to be adjusted in response to the shock. Corresponding impulse response functions feature the same identifying assumptions except for the mortgage rate which is assumed to carry the effects of accommodative monetary policy. As an example, figure 9 contains the amplification effect for the case of the Austrian economy. In particular, draws of the baseline shock (S1; see table 1) responses are subtracted from the corresponding responses of the "flexible rate" shock (S2).[36] The current experiment tries to capture at least the effect of initially low mortgage rates and probably defines

[35] The only underlying model assumption is that house price growth rates are normally distributed with variance σ_c^2.

[36] Note that the responses are obtained from two independent simulations. The efficiency of this exercise may be improved by jointly identifying both shocks for one model.

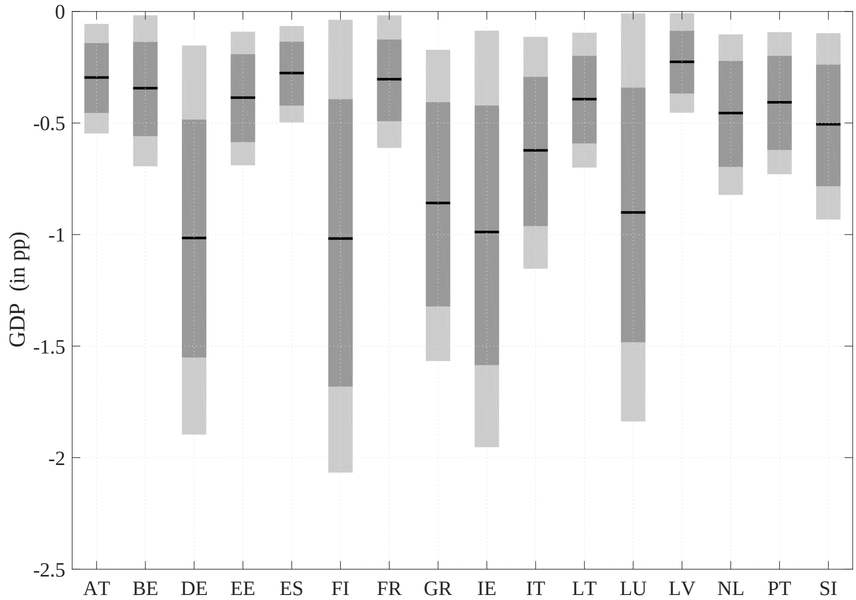

Figure 10: Amplification effect with respect to output 4 quarters after the shock. Solid lines depict medians and the gray shaded areas denote 68 percent and 90 percent Bayesian credible intervals, respectively.

a lower bound of the true effect.[37] Subsequently, the same BMA approach as for the impulse response functions of section 4.2 is applied to the differences between impulse response functions. With respect to the dynamics of output and GDP, simulations feature a gradual decline over about two years followed by a persistent behavior remaining at a lower level. The effect on house prices does not appear to vary distinctly between (S1) and (S2).

In order to illustrate potential cross-country differences, analogously to the impulse response analysis, the amplification effects with respect to the responses of output and consumption one year after the shock are depicted for each country in the sample. Figure 10 contains the effect with respect to the response of output. Across countries, results show a quite homogeneous pattern of consistently negative point estimates with respect to the amplification effect. Quantitatively, the point estimate of the GDP response is lower if mortgage rates are restricted to zero on impact. The impact ranges from about -0.25 percentage point to about -1 percentage point. A closer look at the results reveals the importance of the restriction on the mortgage rate. Accommodative rates seem to fully terminate the adverse effects. The experiment shows that house price shocks are particularly amplified if mortgage rates are restricted to adjust accordingly.

Considering the responses of consumption, figure 11 shows a similar picture. Quantitatively, the amplification is in most cases less than half the size as in the corresponding exercise with respect to GDP. However, the effects with respect to consumption are less clear because some uncertainty bands partially cover the zero line.

[37] In section 4.4 the zero lower bound issue is approximated by conditional forecasts.

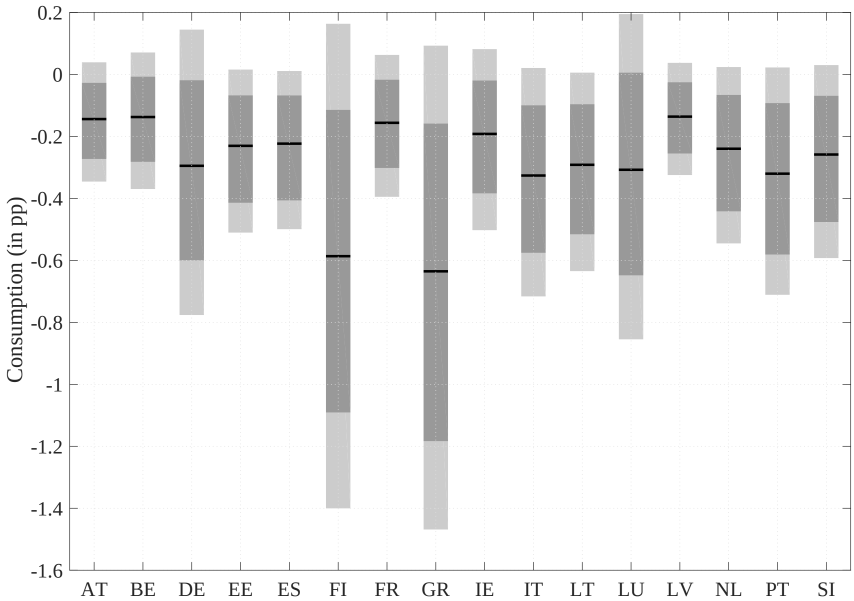

Figure 11: Amplification effect with respect to consumption 4 quarters after the shock. Solid lines depict medians and the gray shaded areas denote 68 percent and 90 percent Bayesian credible intervals, respectively.

4.4 Additional exercises

Prior tightness and lag length

Generally, the BMA approach can be used to address the problem of choosing an appropriate number of lags p in the VAR model. In particular, one may add a lag length dimension to the set of models considered which increases the level of computational complexity. While it is not the focus of this paper to determine the optimal lag length, one may be interested in the effect of p on the tightness parameter λ_c. Ceteris paribus, a positive impact may be expected because, as p increases, so does the number of parameters to be estimated. Larger models should require more observations or tighter prior distributions. Figure 12 plots the posterior distribution of λ_c for an otherwise unchanged VAR model with a lag length of $p = 6$. The results confirm the hypothesis. In terms of the lag length, more complex models yield a tighter prior distribution but the effect differs in magnitude. Comparing the posterior mode estimates, λ_c^m is roughly between 0.05 and 0.2 higher if 6 instead of 4 lags are considered. There are some outliers; Latvia, where the mode estimate remains virtually unchanged and Lithuania, the Netherlands, and Slovenia, where the mode estimate is more than 0.4 higher for the 6-lag model. The mode estimate with respect to Estonia, Portugal, and Spain still represent corner solutions.

Prior tightness and the number of variables

The general conclusion of the previous section – that more complex models require tighter prior distributions – should also carry over to other extensions, e.g. for models with more endogenous variables. Though it might be difficult to identify the effect of model complexity in this setup, a small model with three instead of six variables is estimated. The VAR model employs the baseline specification but contains only GDP, consumption and

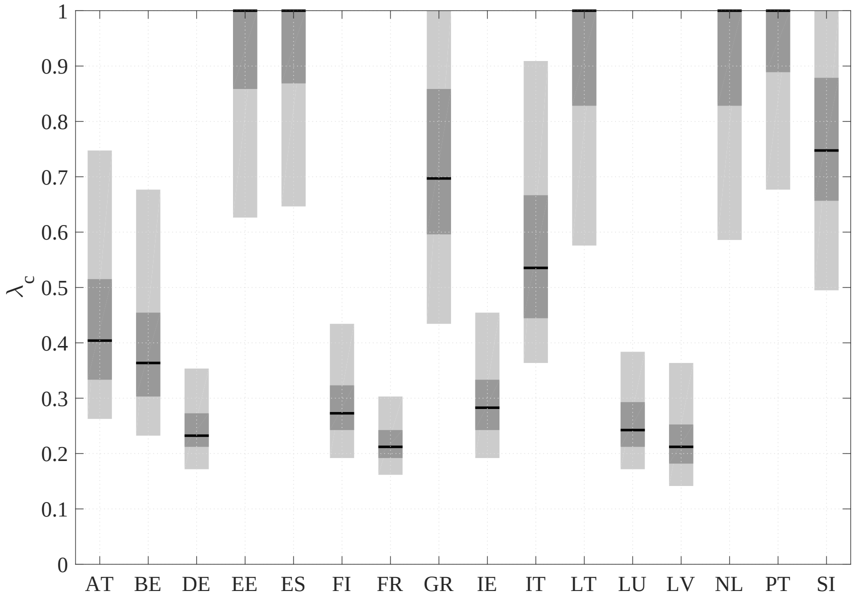

Figure 12: Posterior distribution of the tightness parameter λ_c for a model with $p = 6$ lags. The gray shaded areas denote the 68 percent and the 90 percent Bayesian (highest probability density) credible intervals, respectively.

house prices as endogenous variables. Again, interest centers on the posterior distribution of λ_c which is depicted in figure 13. The point estimate, represented by the mode of posterior distribution, consistently indicates a lower prior tightness parameter if the number of variables in the model is reduced. Some countries feature posterior distributions where the mass concentrates close to zero. While this experiment cannot constitute a proper ceteris paribus analysis, it may still be interpreted in terms of the above hypothesis of a negative relationship between prior distribution tightness and model complexity. In fact, as this result is in line with general Bayesian wisdom, it strengthens the plausibility of the current hierarchical approach. Similarly, section E of the appendix shows that the tightness parameter seems to be a negative function of the sample size. This finding illustrates that if less information (shorter sample) is available from the likelihood, more information is required from the prior distribution.

The role of expectations

One may argue that results are crucially affected by the negative sign restriction on business expectations. Figures 14 and 15 show impulse responses of output and consumption for a housing price bust shock without an initial reaction of households' expectations. The respective restrictions are labeled (S3) in table 1. Qualitatively, the results are akin to the baseline analysis in section 4.2. Quantitatively, the impact of the shock is more muted, which does not come as a surprise since business confidence is usually positively correlated with the business cycle. Compared to the baseline identification (S1), the precision of the estimates seems to improve slightly as the error bands narrow a little bit. Generally, the results do not question the baseline identification scheme as the qualitative results are similar.

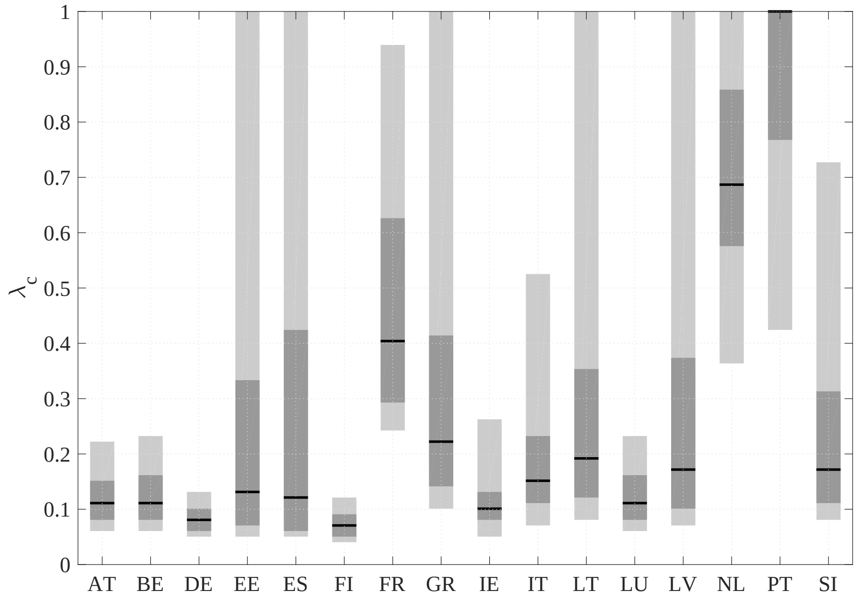

Figure 13: Posterior distribution of the tightness parameter λ_c across countries for a model with a reduced number of variables. The gray shaded areas denote the 68 percent and the 90 percent Bayesian (highest probability density) credible intervals, respectively.

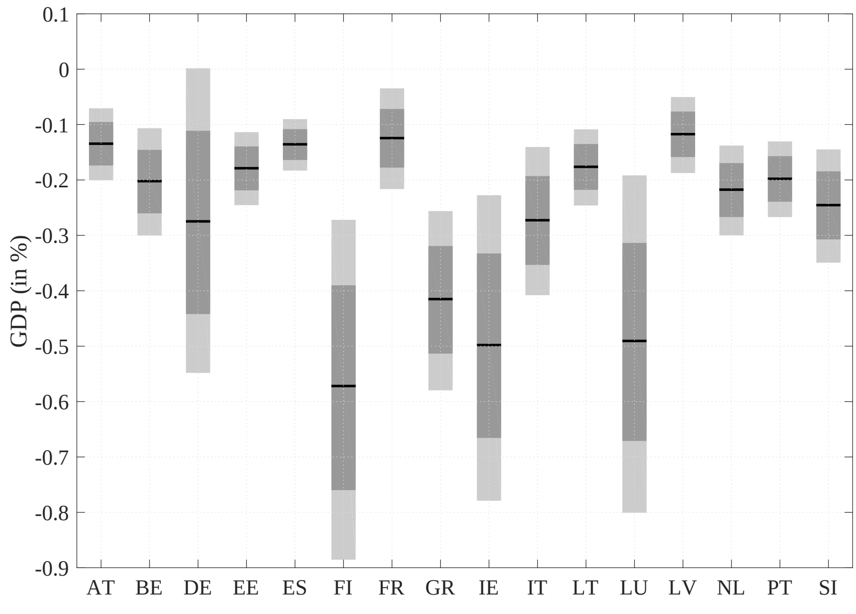

Figure 14: Impulse response functions of output 4 quarters after the shock. Alternative identification with business outlook unaffected on impact. Solid lines depict medians and the gray shaded areas denote 68 percent and 90 percent Bayesian credible intervals, respectively.

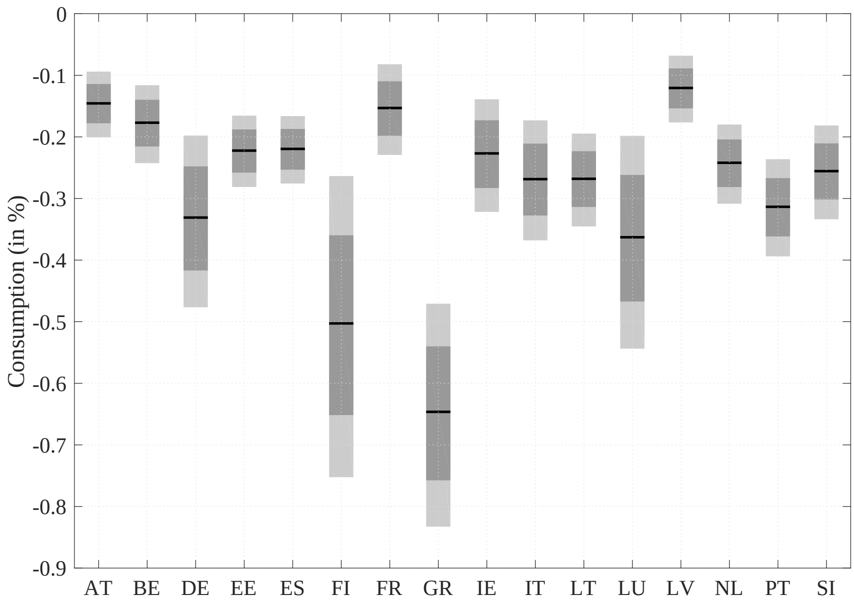

Figure 15: Impulse response functions of consumption 4 quarters after the shock. Alternative identification with business outlook to be unaffected on impact. Solid lines depict medians and the gray shaded areas denote 68 percent and 90 percent Bayesian credible intervals, respectively.

Prolonged restrictions on the mortgage rate

The analysis in section 4.3 attempts to quantify the effect of an interest rate that is not allowed to adjust in response to a house price bust shock. It inherits the disadvantage that the interest rate assumption applies to the impact quarter only. However, it might be interesting to see impulse response functions given an active restriction on the mortgage rate which lasts for the entire horizon. Identification of this shock would require as many zero restrictions on the mortgage rate as quarters considered for the analysis, which is generally infeasible. However, it is possible to define a scenario path which features a zero condition on the mortgage rate. In general, conditional projections may be a suitable tool when considering prolonged zero restrictions. Gustafsson et al. (2016) employ a scenario analysis in order to estimate the impact of a decline in house prices at the zero lower bound with respect to the Swedish economy. Figures 16 and 17 contrast the baseline impulse response functions with the median scenario paths. Simulations with respect to the Austrian and the Italian economy serve as representative examples. The scenario assumes that the baseline shock (S1) hits the economy in period 0 while the mortgage rate remains unchanged from quarter 1 onwards. Technically, implementation is accomplished via a conditional forecasting exercise according to Camba-Mendez (2012). The results show a relatively homogeneous pattern across countries. The impact of the house price bust shock on consumption is generally amplified if mortgage rates cannot adjust accordingly. With respect to GDP, the conclusion is less clear. In some countries such as Finland, Germany, Greece, Ireland, and Luxembourg, the restricted responses are close to the baseline case. Interestingly, the fall in the economic outlook of consumers seems to be dampened by the restricted mortgage rate. One may conjecture that consumers interpret a decrease in the mortgage rate to signal recessions. Consequently, higher mortgage rates

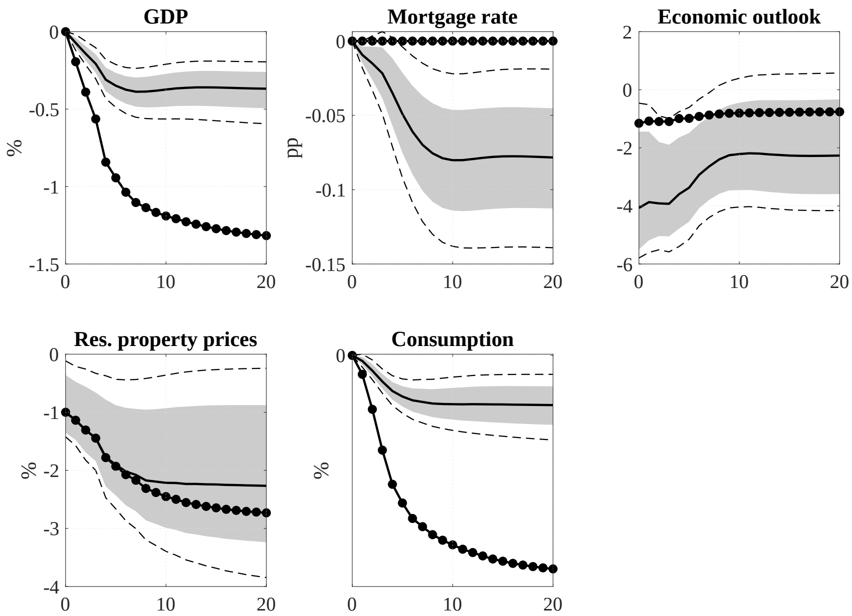

Figure 16: Quarterly cumulative impulse response functions for Austria contrasted with the zero mortgage scenario. Solid lines depict medians, circles indicate the median scenario path, and gray shaded areas denote 68 percent and dashed lines 90 percent Bayesian credible intervals, respectively.

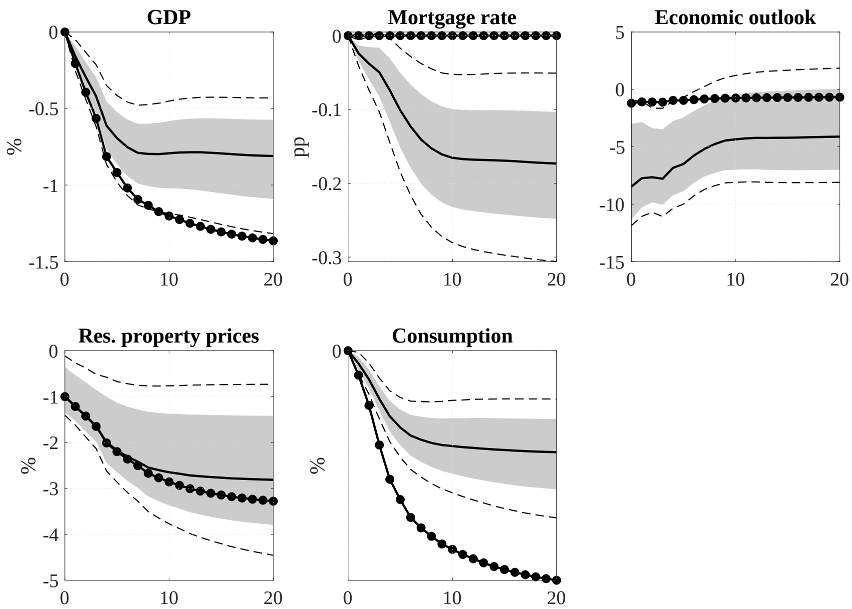

Figure 17: Quarterly cumulative impulse response functions for Italy contrasted with the zero mortgage scenario. Solid lines depict medians, circles indicate the median scenario path, and gray shaded areas denote 68 percent and dashed lines 90 percent Bayesian credible intervals, respectively.

enhance the economic outlook.

It has to be noted that, except for the impact period (quarter $h = 0$), all shocks are unrestricted, whereas the mortgage rate is restricted to zero thereafter (for quarters $h = 1, \ldots, 20$). The latter restriction implicitly defines paths for all shocks in the model in order to ensure that the required path of the mortgage rate is met. This includes the five non-identified shocks. In general, the composition of shocks that generate the scenario is non-unique, which implies that the paths of all other variables except for the mortgage rate may depend on the particular choice of shocks. However, the non-identified shocks may most likely constitute random disturbances that should not influence point estimates. If one were interested in the explicit sources of the scenario, it would be mandatory to identify the remaining 4 structural shocks in the VAR system, which is beyond the scope of this paper. Other techniques that try to account for the zero lower bound may circumvent this problem. Baumeister and Benati (2013) proposes a "zeroing out" approach where the coefficients in the structural interest rate rule are set to zero. Still the authors note that this procedure may be subject to the Lucas critique as parameters of the structural model are adjusted ex post. In general, the linear VAR model cannot capture zero lower bound episodes in a fully coherent way. As my results seem mostly plausible, the approximation error should not be substantial and I deem them reliable.

Agnostic identification

In order to put the identifying assumptions with regard to the structural shocks into perspective, I compute generalized impulse response functions (S4) with respect to a median one-percent change in house prices. Results for the GDP and the consumption response 4 quarters after the shock in figures 18 and 19 merely serve as an atheoretical benchmark for the structural analysis. Compared to the baseline experiment (S1), responses of GDP and consumption are more muted because variables are unrestricted. Across countries and compared to the baseline results, the effect on the GDP response is positive by about 0.1 to 0.5 percentage point and for consumption it is roughly between 0.03 and 0.1 percentage point higher. The difference with respect to impulse response functions seems to be more pronounced for countries that exhibit stronger effects of the shock, such as Finland, Germany, Greece, Ireland, or Luxembourg. The shape of the impulse response functions is also quite similar to the baseline case.[38] In fact, only the effect on the economic outlook seems to be more transitory and (the median) even turns positive during the second year after the shock. This may counteract the adverse nature of the shock a little bit whereas output and the mortgage rate, even without an explicit restriction, remain close to zero on impact. This finding supports the conclusion that adverse expectations considerably amplify the impact of house price declines.

5 Conclusion

This study employs a Bayesian structural vector autoregression model in order to quantify the impact of house price bust shocks on real economic activity and consumption for 16 euro area countries. Since one may conjecture that country specificities are important with

[38] Because of the limited value added, I have refrained from depicting the full paths of the responses. Results are available upon request.

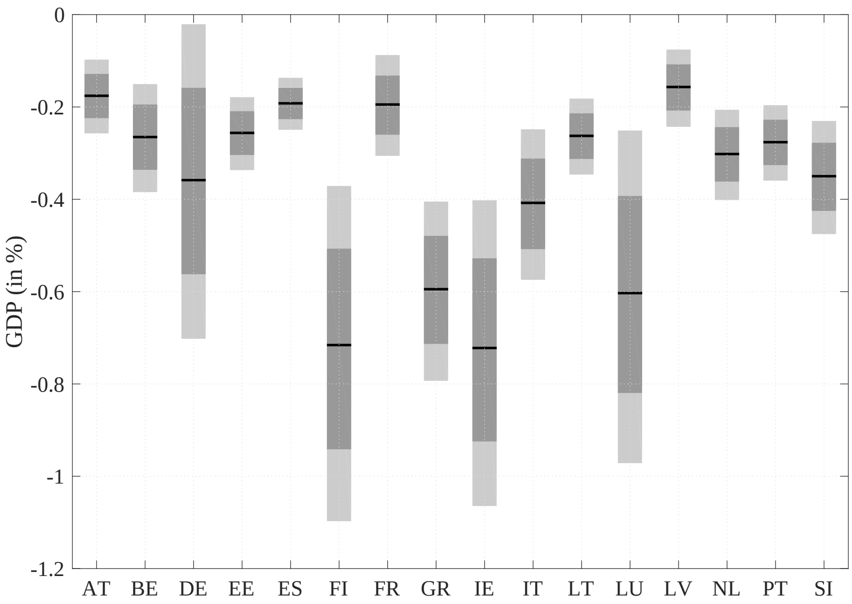

Figure 18: Impulse response functions of output 4 quarters after the shock. Alternative identification without restrictions (generalized impulse response functions). Solid lines depict medians and the gray shaded areas denote 68 percent and 90 percent Bayesian credible intervals, respectively.

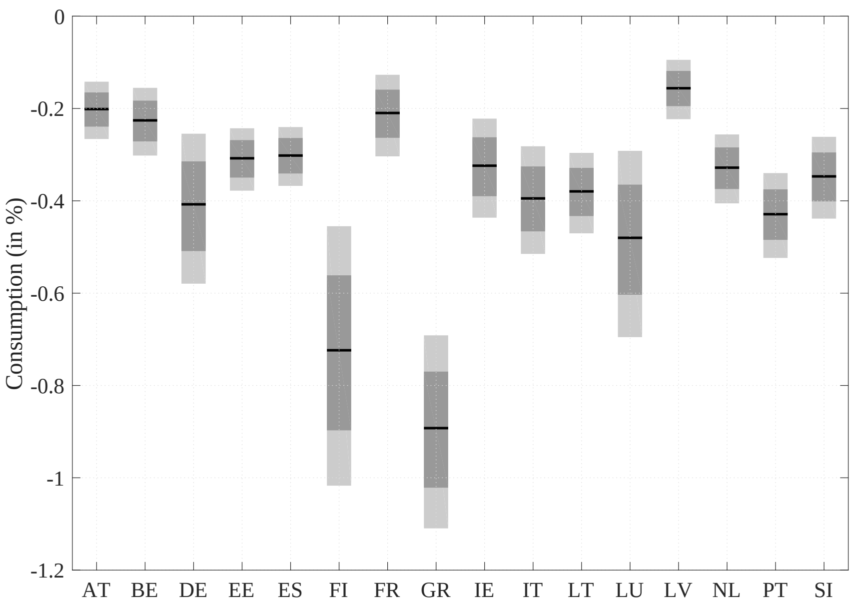

Figure 19: Impulse response functions of consumption 4 quarters after the shock. Alternative identification without restrictions (generalized impulse response functions). Solid lines depict medians and the gray shaded areas denote 68 percent and 90 percent Bayesian credible intervals, respectively.

respect to real estate market developments, the analysis is based on a novel econometric approach which overcomes the small-sample problem by considering an appropriate degree of homogeneity for each country-specific model.

Economically, the effect of a one-percent house price bust shock on output and consumption is substantial, as the GDP response after one year is in the range of -0.2 percent to -1.2 percent and that of consumption between -0.25 percent and about -1.2 percent. The shock likewise captures a deterioration of consumers' economic outlook which intensifies the severity of effects substantially. It incorporates properties of confidence shocks that may also reinforce negative effects in the emergence of crises. The respective dynamics show a quite persistent pattern. Both variables decline for about one year before stabilizing at the new lower level. Cross-country heterogeneity mainly impacts on the quantitative effects as opposed to the qualitative transmission channels. House price busts reduce household wealth and induce households to consume less and to demand less mortgage credit. House prices subsequently drop further, which amplifies the initial effect. Second-round effects seem to be important because the initial reactions are distinctly smaller. Flexible interest rates dampen the effect of house price shocks, whereas consumers' economic outlook implies a small amplification effect.

Econometrically, the partial pooling approach delivers robust estimates of impulse response functions, which confirm the reliability of the results. With respect to the prior distribution tightness parameter, the hierarchical Bayesian formulation delivers plausible results and enhances precision with respect to the posterior distribution of impulse response functions. Estimates increase with model complexity, which is approximated by the model's lag length and the number of variables included.

In terms of potential extensions to the approach, one generalization may be used to investigate cross-country heterogeneity in more detail. The econometric strategy assigns a single tightness parameter to the set of pooled cross-country observations. It appears promising to relax this assumption and to place a separate tightness parameter for each country represented in the prior distribution. The posterior distributions of those tightness parameters contain information on potential cross-country clusters in the sample.

References

Ahamada, I. and J. L. D. Sanchez (2013). A retrospective analysis of the house prices macro-relationship in the United States. *International Journal of Central Banking 9*(4), 153–174.

Arias, J. E., J. F. Rubio-Ramirez, and D. F. Waggoner (2014). Inference based on SVAR identified with sign and zero restrictions: Theory and applications. CEPR Discussion Papers 9796, C.E.P.R. Discussion Papers.

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

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

Berger, D., V. Guerrieri, G. Lorenzoni, and J. Vavra (2018). House prices and consumer spending. *Review of Economic Studies 85*(3), 1502–1542.

Bernanke, B. S. (2010, September). Statement. Board of Governors of the Federal Reserve System before the Financial Crisis Inquiry Commission.

Calza, A., T. Monacelli, and L. Stracca (2013). Housing finance and monetary policy. *Journal of the European Economic Association 11*, 101–122.

Camba-Mendez, G. (2012). Conditional forecasts on SVAR models using the Kalman filter. *Economics Letters 115*(3), 376–378.

Canova, F. (2005). The transmission of US shocks to Latin America. *Journal of Applied Econometrics 20*(2), 229–251.

Canova, F. (2007). *Methods for Applied Macroeconomic Research*. Princeton University Press.

Canova, F. and M. Ciccarelli (2013). Panel vector autoregressive models: a survey. Working Paper Series 1507, European Central Bank.

Case, K. E., J. M. Quigley, and R. J. Shiller (2013). Wealth effects revisited 1975-2012. *Critical Finance Review 2*(1), 101–128.

Del Negro, M. and F. Schorfheide (2003). Take your model bowling: forecasting with general equilibrium models. *FRB Atlanta Economic Review Economic Review 88*(4), 35–50.

Del Negro, M. and F. Schorfheide (2004). Priors from general equilibrium models for VARs. *International Economic Review 45*(2), 643–673.

Del Negro, M. and F. Schorfheide (2011). Bayesian macroeconometrics. In J. Geweke, G. Koop, and H. van Dijk (Eds.), *The Oxford Handbook of Bayesian Econometrics*, pp. 293–389. Oxford, UK: Oxford University Press.

Disney, R., J. Gathergood, and A. Henley (2010). House price shocks, negative equity, and household consumption in the United Kingdom. *Journal of the European Economic Association 8*(6), 1179–1207.

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

ESRB (2016). Vulnerabilities in the EU residential real estate sector. Technical report, European Systemic Risk Board Report.

Frey, C. and F. Mokinski (2016). Forecasting with Bayesian vector autoregressions estimated using professional forecasts. *Journal of Applied Econometrics 31*(6), 1083–1099.

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

Gustafsson, P., P. Stockhammar, and P. Österholm (2016). Macroeconomic effects of a decline in housing prices in Sweden. *Journal of Policy Modeling 38*(2), 242–255.

Hamilton, J. (1994). *Time Series Analysis*. Princeton, NJ: Princeton Univ. Press.

Jarociński, M. (2010). Responses to monetary policy shocks in the east and the west of Europe: a comparison. *Journal of Applied Econometrics 25*(5), 833–868.

Jarociński, M. and B. Maćkowiak (2013). Granger-causal-priority and choice of variables in vector autoregressions. Working Paper Series 1600, European Central Bank.

Jarociński, M. and B. Maćkowiak (2017). Granger causal priority and choice of variables in vector autoregressions. *The Review of Economics and Statistics 99*(2), 319–329.

Kadiyala, K. R. and S. Karlsson (1997). Numerical methods for estimation and inference in Bayesian VAR-models. *Journal of Applied Econometrics 12*(2), 99–132.

Koop, G. (2003). *Bayesian Econometrics* (1st ed.). John Wiley & Sons.

Koop, G. and D. Korobilis (2010). Bayesian multivariate time series methods for empirical macroeconomics. *Foundations and Trends(R) in Econometrics 3*(4), 267–358.

Koop, G. and D. Korobilis (2016). Model uncertainty in panel vector autoregressive models. *European Economic Review 81*(C), 115–131.

Lambertini, L., C. Mendicino, and M. T. Punzi (2013). Expectation-driven cycles in the housing market: Evidence from survey data. *Journal of Financial Stability 9*(4), 518–529.

Lambertini, L., C. Mendicino, and M. T. Punzi (2017). Expectations-driven cycles in the housing market. *Economic Modelling 60*(C), 297–312.

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

Nocera, A. and M. Roma (2017). House prices and monetary policy in the euro area: evidence from structural VARs. Working Paper Series 2073, European Central Bank.

Pesaran, H. H. and Y. Shin (1998). Generalized impulse response analysis in linear multivariate models. *Economics Letters 58*(1), 17–29.

Pesaran, M. H. and R. Smith (1995). Estimating long-run relationships from dynamic heterogeneous panels. *Journal of Econometrics 68*(1), 79–113.

Piazzesi, M. and M. Schneider (2009). Momentum traders in the housing market: Survey evidence and a search model. *American Economic Review 99*(2), 406–411.

Piazzesi, M. and M. Schneider (2016). *Housing and Macroeconomics*, Volume 2 of *Handbook of Macroeconomics*, Chapter 0, pp. 1547–1640. Elsevier.

Praet, P. (2019, February). On the importance of real estate statistics. Speech at the International Conference on Real Estate Statistics, Luxembourg, 21 February 2019.

Primiceri, G. E. (2005). Time varying structural vector autoregressions and monetary policy. *Review of Economic Studies 72*(3), 821–852.

Rickman, D. S. (1995). A Bayesian analysis of the use of pooled coefficients in a structural regional economic model. *International Journal of Forecasting 11*(3), 477–490.

Rodrigues, P. M. and R. F. Lourenço (2015). House prices: bubbles, exuberance or something else? Evidence from euro area countries. Working Papers w201517, Banco de Portugal, Economics and Research Department.

Sims, C. A. and T. Zha (1998). Bayesian methods for dynamic multivariate models. *International Economic Review 39*(4), 949–968.

Theil, H. and A. S. Goldberger (1961). On pure and mixed statistical estimation in economics. *International Economic Review 2*(1), 65–78.

Towbin, P. and S. Weber (2016). Price expectations and the US housing boom. Working Paper 6, Swiss National Bank.

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

Uhlig, H. (2017). Shocks, signs, restrictions, and identification. In B. Honoré, A. Pakes, M. Piazzesi, and L. Samuelson (Eds.), *Advances in Economics and Econometrics Vol 1 and 2: Eleventh World Congress*. Cambridge University Press.

Zellner, A. (1971). *Introduction to Bayesian Inference in Econometrics*. New York: John Wiley & Sons.

Zellner, A. (1986). On assessing prior distributions and Bayesian regression analysis with g-prior distributions. In P. Goel and A. Zellner (Eds.), *Bayesian Inference and Decision Techniques: Essays in Honor of Bruno de Finetti*, pp. 233–243. New York: Elsevier Science Publishers, Inc.

Zellner, A. and C. Hong (1989). Forecasting international growth rates using Bayesian shrinkage and other procedures. *Journal of Econometrics 40*(1), 183–202.

Series	Source	Mnemonic
GDP	ECB SDW	`MNA.Q.Y.*.W2.S1.S1.B.B1GQ._Z._Z._Z.EUR.LR.N`
Mortgage rate (outstanding)	ECB SDW	`MIR.M.*country*.B.A22.A.R.A.2250.EUR.O`
Business outlook indicator	OECD	`MEI.M.*.CSESFT02.STSA`
Residential property prices	BIS	`Q:*:N:628`
Consumption	ECB SDW	`MNA.Q.Y.*.W0.S1M.S1.D.P31._Z._Z._T.EUR.LR.N`

Table A.1: Listing of the data used for the analyses. "*" replaces the country abbreviation.

Appendix

A Additional data description

As the analysis employs a panel dataset with $5 \times 16 = 80$ time series, a direct graphical illustration seems lavish and is omitted. For the sake of transparency regarding the time series used, table A.1 references the data series and the respective sources.

B Dummy observations and the natural conjugate prior

For notational convenience the following derivation omits the country index. Consider the weighted likelihood of the dummy observations

$$f(\underline{\boldsymbol{Y}}|\boldsymbol{\beta},\boldsymbol{\Sigma}) = (2\pi)^{\frac{n\underline{T}}{2}} \det(\boldsymbol{\Sigma})^{-\frac{\underline{T}}{2}} \exp\left\{-\frac{1}{2}\mathrm{tr}\left[\boldsymbol{\Sigma}^{-1} \underbrace{(\underline{\boldsymbol{Y}} - \underline{\boldsymbol{X}}\boldsymbol{B})'(\underline{\boldsymbol{Y}} - \underline{\boldsymbol{X}}\boldsymbol{B})}_{\equiv\mathcal{R}}\right]\right\}. \tag{B.1}$$

The prior distribution's tightness parameter is implicitly contained in equation (B.1). Recall from the main text, that λ enters the model as follows

$$\underline{T} \equiv T^*\lambda, \quad \underline{\boldsymbol{Y}} \equiv \boldsymbol{Y}^*\sqrt{\lambda}, \quad \underline{\boldsymbol{X}} \equiv \boldsymbol{X}^*\sqrt{\lambda}.$$

Let us first rewrite the term in curly brackets

$$\begin{aligned}
\mathcal{R} &= (\underline{\boldsymbol{Y}} - \underline{\boldsymbol{X}}\boldsymbol{B} + \underline{\boldsymbol{X}}\,\underline{\boldsymbol{B}} - \underline{\boldsymbol{X}}\,\underline{\boldsymbol{B}})'(\underline{\boldsymbol{Y}} - \underline{\boldsymbol{X}}\boldsymbol{B} + \underline{\boldsymbol{X}}\,\underline{\boldsymbol{B}} - \underline{\boldsymbol{X}}\,\underline{\boldsymbol{B}}) \\
&= (\underline{\boldsymbol{Y}} - \underline{\boldsymbol{X}}\,\underline{\boldsymbol{B}})'(\underline{\boldsymbol{Y}} - \underline{\boldsymbol{X}}\,\underline{\boldsymbol{B}}) + (\boldsymbol{B} - \underline{\boldsymbol{B}})'\underline{\boldsymbol{X}}'\underline{\boldsymbol{X}}(\boldsymbol{B} - \underline{\boldsymbol{B}}) \\
&+ \underbrace{(\underline{\boldsymbol{Y}} - \underline{\boldsymbol{X}}\,\underline{\boldsymbol{B}})'\underline{\boldsymbol{X}}(\underline{\boldsymbol{B}} - \boldsymbol{B})}_{\equiv\mathcal{F}} + \underbrace{(\underline{\boldsymbol{B}} - \boldsymbol{B})'\underline{\boldsymbol{X}}'(\underline{\boldsymbol{Y}} - \underline{\boldsymbol{X}}\,\underline{\boldsymbol{B}})}_{\equiv\mathcal{F}'}
\end{aligned}$$

Set

$$\underline{\boldsymbol{B}} \equiv (\underline{\boldsymbol{X}}'\underline{\boldsymbol{X}})^{-1}(\underline{\boldsymbol{X}}'\underline{\boldsymbol{Y}})$$

in order to show

$$\begin{aligned}
\mathcal{F} &= (\underline{\boldsymbol{Y}} - \underline{\boldsymbol{X}}\,\underline{\boldsymbol{B}})'\underline{\boldsymbol{X}}(\underline{\boldsymbol{B}} - \boldsymbol{B}) \\
&= \left[\underline{\boldsymbol{Y}}' - \underline{\boldsymbol{Y}}'\underline{\boldsymbol{X}}(\underline{\boldsymbol{X}}'\underline{\boldsymbol{X}})^{-1}\underline{\boldsymbol{X}}'\right]\underline{\boldsymbol{X}}(\underline{\boldsymbol{B}} - \boldsymbol{B}) \\
&= (\underline{\boldsymbol{Y}}'\underline{\boldsymbol{X}} - \underline{\boldsymbol{Y}}'\underline{\boldsymbol{X}})(\underline{\boldsymbol{B}} - \boldsymbol{B}) \\
&= \boldsymbol{0}_{n,n}.
\end{aligned}$$

Define further

$$\underline{\boldsymbol{S}} \equiv \left(\underline{\boldsymbol{Y}} - \underline{\boldsymbol{X}}\underline{\boldsymbol{B}}\right)' \left(\underline{\boldsymbol{Y}} - \underline{\boldsymbol{X}}\underline{\boldsymbol{B}}\right), \quad \underline{\boldsymbol{V}}^{-1} \equiv \underline{\boldsymbol{X}}'\underline{\boldsymbol{X}}, \quad \text{and} \quad \underline{\nu} \equiv \underline{T} - n - 1$$

and substitute back $\mathcal{R}$. This yields

$$f\left(\underline{\boldsymbol{Y}}|\boldsymbol{\beta}, \boldsymbol{\Sigma}\right) = (2\pi)^{\frac{n\underline{T}}{2}} \exp\left\{-\frac{1}{2}\text{tr}\left[\boldsymbol{\Sigma}^{-1}\left(\boldsymbol{B} - \underline{\boldsymbol{B}}\right)' \underline{\boldsymbol{V}}^{-1}\left(\boldsymbol{B} - \underline{\boldsymbol{B}}\right)\right]\right\} \cdots$$
$$\cdots \times \det\left(\boldsymbol{\Sigma}\right)^{-\frac{\underline{\nu}+n+1}{2}} \exp\left[-\frac{1}{2}\text{tr}\left(\boldsymbol{\Sigma}^{-1}\underline{\boldsymbol{S}}\right)\right],$$

or

$$f\left(\underline{\boldsymbol{Y}}|\boldsymbol{\beta}, \boldsymbol{\Sigma}\right) = (2\pi)^{\frac{n\underline{T}}{2}} \underbrace{\exp\left\{-\frac{1}{2}\left(\boldsymbol{\beta} - \underline{\boldsymbol{\beta}}\right)'\left(\boldsymbol{\Sigma}^{-1} \otimes \underline{\boldsymbol{V}}^{-1}\right)\left(\boldsymbol{\beta} - \underline{\boldsymbol{\beta}}\right)\right\}}_{\text{kernel of } \mathcal{N}\left(\underline{\boldsymbol{\beta}}, \boldsymbol{\Sigma} \otimes \underline{\boldsymbol{V}}\right)} \cdots$$
$$\cdots \times \underbrace{\det\left(\boldsymbol{\Sigma}\right)^{-\frac{\underline{\nu}+n+1}{2}} \exp\left[-\frac{1}{2}\text{tr}\left(\boldsymbol{\Sigma}^{-1}\underline{\boldsymbol{S}}\right)\right]}_{\text{kernel of } \mathcal{IW}\left(\underline{\boldsymbol{S}}, \underline{\nu}\right)},$$

where $\underline{\boldsymbol{\beta}} \equiv \text{vec}(\underline{\boldsymbol{B}})$. This expression translates to the probability density function of a normal-inverse-Wishart distribution

$$\left(\boldsymbol{\beta}, \boldsymbol{\Sigma}\right)|\lambda \sim \mathcal{NIW}\left(\underline{\boldsymbol{\beta}}, \underline{\boldsymbol{V}}, \underline{\boldsymbol{S}}, \underline{\nu}\right). \tag{B.2}$$

with

$$\underline{\boldsymbol{V}}^{-1} = \underline{\boldsymbol{X}}'\underline{\boldsymbol{X}}, \quad \underline{\boldsymbol{B}} = \underline{\boldsymbol{V}}\underline{\boldsymbol{X}}'\underline{\boldsymbol{Y}}, \quad \underline{\boldsymbol{S}} = \left(\underline{\boldsymbol{Y}} - \underline{\boldsymbol{X}}\underline{\boldsymbol{B}}\right)'\left(\underline{\boldsymbol{Y}} - \underline{\boldsymbol{X}}\underline{\boldsymbol{B}}\right), \quad \underline{\nu} = \underline{T}$$

It can be shown that the normal-inverse-Wishart distribution serves as a natural conjugate prior distribution with respect to the VAR model likelihood function. The respective posterior distribution reads

$$\left(\boldsymbol{\beta}, \boldsymbol{\Sigma}\right)|\boldsymbol{Y}, \lambda \sim \mathcal{NIW}\left(\overline{\boldsymbol{\beta}}, \overline{\boldsymbol{V}}, \overline{\boldsymbol{S}}, \overline{\nu}\right). \tag{B.3}$$

where

$$\begin{aligned}
\overline{\boldsymbol{V}}^{-1} &= \underline{\boldsymbol{V}}^{-1} + \boldsymbol{X}'\boldsymbol{X} \\
\overline{\boldsymbol{B}} &= \overline{\boldsymbol{V}}\left(\underline{\boldsymbol{V}}^{-1} + \boldsymbol{X}'\boldsymbol{X}\widehat{\boldsymbol{B}}\right) \\
\overline{\boldsymbol{S}} &= \widehat{\boldsymbol{S}} + \underline{\boldsymbol{S}} + \widehat{\boldsymbol{B}}'\boldsymbol{X}'\boldsymbol{X}\widehat{\boldsymbol{B}} + \underline{\boldsymbol{B}}'\underline{\boldsymbol{V}}^{-1}\underline{\boldsymbol{B}} - \overline{\boldsymbol{B}}'\left(\underline{\boldsymbol{V}}' + \boldsymbol{X}'\boldsymbol{X}\right)\overline{\boldsymbol{B}} \\
\overline{\nu} &= T + \underline{\nu}
\end{aligned}$$

Equation (B.3) has been derived from the likelihood of a dummy variables prior distribution (B.1) and defines a natural conjugate prior distribution to the VAR model in equation (1).

The natural conjugate representation of the dummy observations prior in equation (B.2) gives us an additional perspective on the impact of the tightness parameter on the posterior distribution. It is well-understood and documented, e.g. by Koop and Korobilis (2010), that the prior distribution (B.2) yields the posterior distribution

$$(\boldsymbol{\beta}, \boldsymbol{\Sigma})\,|\boldsymbol{Y}, \lambda \sim \mathcal{NIW}\left(\overline{\boldsymbol{\beta}}, \overline{\boldsymbol{V}}, \overline{\boldsymbol{S}}, \overline{\nu}\right),$$

where

$$\begin{aligned}
\overline{\boldsymbol{V}}^{-1} &= \underline{\boldsymbol{V}}^{-1} + \boldsymbol{X}'\boldsymbol{X} \\
&= \lambda \boldsymbol{X}^{*\prime}\boldsymbol{X}^{*} + \boldsymbol{X}'\boldsymbol{X} \\
\overline{\boldsymbol{B}} &= \overline{\boldsymbol{V}}\left(\underline{\boldsymbol{V}}^{-1}\underline{\boldsymbol{B}} + \boldsymbol{X}'\boldsymbol{X}\widehat{\boldsymbol{B}}\right) \\
&= \overline{\boldsymbol{V}}\left(\lambda \boldsymbol{X}^{*\prime}\boldsymbol{X}^{*}\boldsymbol{B}^{*} + \boldsymbol{X}'\boldsymbol{X}\widehat{\boldsymbol{B}}\right) \\
\overline{\boldsymbol{S}} &= \widehat{\boldsymbol{S}} + \widehat{\boldsymbol{B}}'\boldsymbol{X}'\boldsymbol{X}\widehat{\boldsymbol{B}} + \underline{\boldsymbol{S}} + \underline{\boldsymbol{B}}'\underline{\boldsymbol{V}}^{-1}\underline{\boldsymbol{B}} - \overline{\boldsymbol{B}}'\left(\underline{\boldsymbol{V}}^{-1} + \boldsymbol{X}'\boldsymbol{X}\right)\overline{\boldsymbol{B}} \\
&= \widehat{\boldsymbol{S}} + \widehat{\boldsymbol{B}}'\boldsymbol{X}'\boldsymbol{X}\widehat{\boldsymbol{B}} + \lambda \boldsymbol{S}^{*} + \lambda \boldsymbol{B}^{*\prime}\boldsymbol{V}^{*-1}\boldsymbol{B}^{*} - \overline{\boldsymbol{B}}'\left(\lambda \boldsymbol{X}^{*\prime}\boldsymbol{X}^{*} + \boldsymbol{X}'\boldsymbol{X}\right)\overline{\boldsymbol{B}} \\
\overline{\nu} &= T + \underline{T} \\
&= T + \lambda T^{*}.
\end{aligned}$$

This representation confirms that the impact of the prior vanishes as the tightness parameter $\lambda \to 0$ and its impact increases as $\lambda \to \infty$. However, the analysis in the main text imposes $\lambda \in [0, 1]$ via the respective prior distribution, where $\lambda = 1$ resembles the pooled panel VAR model.

C Short-run effects

Section 4.2 depicts the effects of house price bust shocks on output and consumption four quarters after the shock hits the economy. A comparison of those variables' responses at shorter horizons may help to quantify dynamic amplification mechanisms. Figures C.1 and C.2 illustrate the respective reactions one quarter after the impact quarter. While the magnitude of the reaction is distinctly lower, the relative effect across countries is quite similar. With respect to output, the responses after one year are about four times the effect after one quarter. With respect to consumption, the responses are amplified by an even higher factor.

D Non-diagonal prior covariance

Previous approaches[39] proposed prior distributions for VAR models which feature a diagonal structure of the slope coefficient covariance matrix. One reason for this strategy may be that, as argued by Koop (2003, ch. 11), it is not obvious to formulate sensible a priori beliefs with respect to the covariance elements. Prior distributions implemented via dummy observations, like the one described in section 2.1, do not suffer from this disadvantage. In the following I show that given the DGP is a stationary VAR model (i) the covariance matrix of the slope coefficient matrix is a function of the autocovariance of the endogenous variables and (ii) the autocovariance is represented by a matrix $\boldsymbol{\Gamma}$ with

[39] See e.g. Jarociński (2010) or Koop and Korobilis (2016).

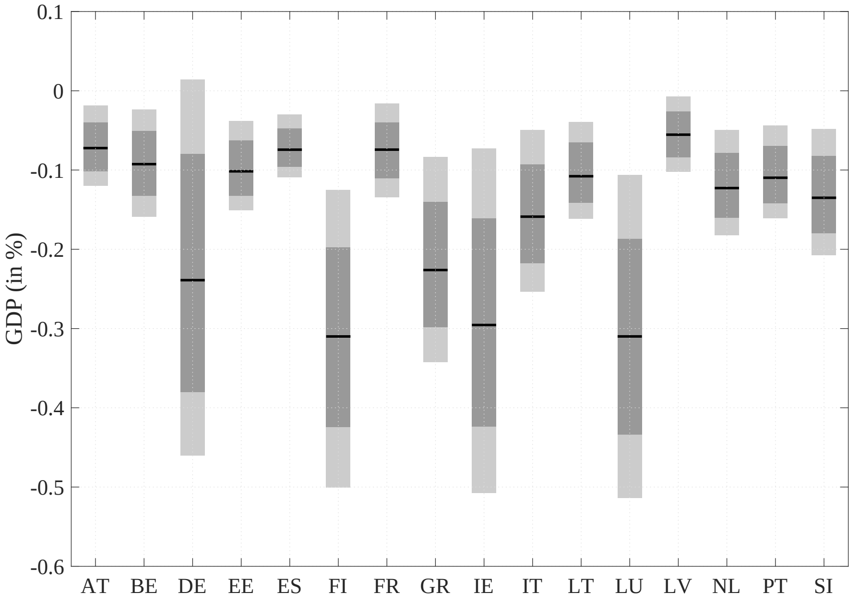

Figure C.1: Impulse response functions of GDP 1 quarter after the shock. Solid lines depict medians and the gray shaded areas denote 68 percent and 90 percent Bayesian credible intervals, respectively.

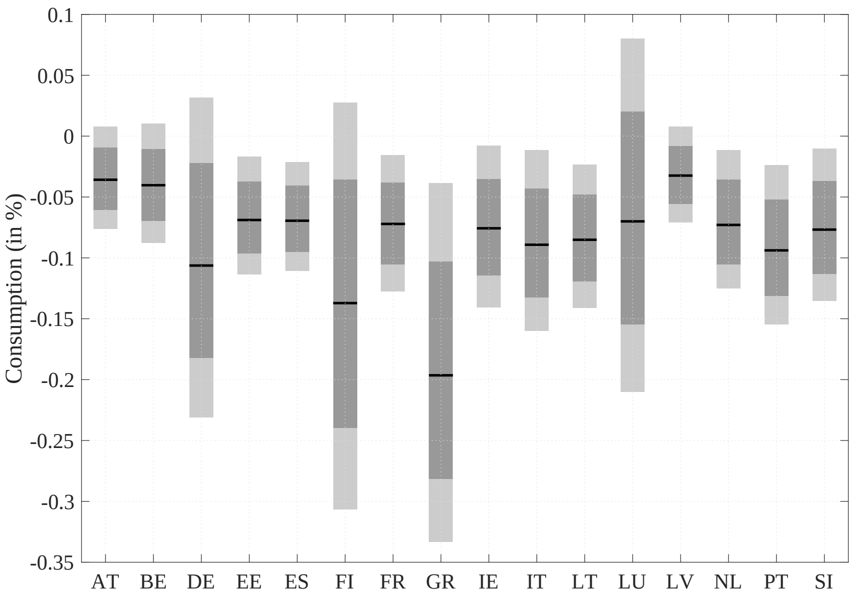

Figure C.2: Impulse response functions of consumption 1 quarter after the shock. Solid lines depict medians and the gray shaded areas denote 68 percent and 90 percent Bayesian credible intervals, respectively.

non-zero off-diagonal elements. A Monte Carlo analysis shows that the posterior slope coefficient covariance matrix is closer to the true matrix if the prior off-diagonal elements are derived from the cross section. The experiment shows that a slope covariance matrix with zero off-diagonal elements assigns more probability mass to unreasonable dynamics as mentioned by Del Negro and Schorfheide (2011).

D.1 Slope covariance matrix in autoregressive models

Assume the true DGP is known to be a stationary VAR with p lags as defined by equation (1). If data was simulated from that model, the corresponding likelihood function reads

$$f(\boldsymbol{Y}|\boldsymbol{B},\boldsymbol{\Sigma}) = (2\pi)^{\frac{nT}{2}} \det(\boldsymbol{\Sigma})^{-\frac{T}{2}} \exp\left\{-\frac{1}{2}\mathrm{tr}\left[\boldsymbol{\Sigma}^{-1}(\boldsymbol{Y}-\boldsymbol{X}\boldsymbol{B})'(\boldsymbol{Y}-\boldsymbol{X}\boldsymbol{B})\right]\right\}$$

The function may be rewritten in a similar way as has been done in section B

$$f(\boldsymbol{Y}|\boldsymbol{B},\boldsymbol{\Sigma}) = (2\pi)^{\frac{nT}{2}} \det(\boldsymbol{\Sigma})^{-\frac{T}{2}} \exp\left\{-\frac{1}{2}\left(\boldsymbol{\beta}-\widehat{\boldsymbol{\beta}}\right)'\left(\boldsymbol{\Sigma}^{-1}\otimes\boldsymbol{X}'\boldsymbol{X}\right)\left(\boldsymbol{\beta}-\widehat{\boldsymbol{\beta}}\right)\right\}. \quad \text{(D.1)}$$

Equation (D.1) defines a multivariate normal distribution with respect to the maximum likelihood estimate of the slope coefficients $\widehat{\boldsymbol{\beta}}$

$$\boldsymbol{\beta} \sim \mathcal{N}\left(\widehat{\boldsymbol{\beta}}, \boldsymbol{\Sigma}\otimes\widehat{\boldsymbol{V}}\right),$$

with $\widehat{\boldsymbol{\beta}} \equiv \mathrm{vec}(\widehat{\boldsymbol{B}}) = \mathrm{vec}[(\boldsymbol{X}'\boldsymbol{X})^{-1}\boldsymbol{X}'\boldsymbol{Y}]$ and $\widehat{\boldsymbol{V}} \equiv (\boldsymbol{X}'\boldsymbol{X})^{-1}$. Assume that the covariance matrix of the reduced form residuals $\boldsymbol{\Sigma}$ is known; then $\boldsymbol{X}'\boldsymbol{X}$ crucially determines the covariance matrix with respect to the coefficient estimate $\widehat{\boldsymbol{\beta}}$. In particular, since $\boldsymbol{X}$ exclusively contains the first p lags of $\boldsymbol{y}_t$, $\boldsymbol{X}'\boldsymbol{X}/T$ represents an estimate of the covariance matrix

$$\boldsymbol{\Gamma} \equiv \mathbb{E}\left(\boldsymbol{\mathcal{Y}}_{t-1}\boldsymbol{\mathcal{Y}}_{t-1}'\right) = \mathbb{E}\left(\boldsymbol{\mathcal{Y}}_t\boldsymbol{\mathcal{Y}}_t'\right),$$

where $\boldsymbol{\mathcal{Y}}_t \equiv [\boldsymbol{y}_t', \cdots, \boldsymbol{y}_{t-p+1}']'$. Consequently, $\boldsymbol{\Gamma}$ contains the autocovariance function of $\boldsymbol{y}_t$ up to lag $p-1$. In fact as $t\to\infty$ the estimator converges $\boldsymbol{X}'\boldsymbol{X}/T \to \boldsymbol{\Gamma}$.

In the following, I derive an expression for $\boldsymbol{\Gamma}$ that depends on the underlying DGP's coefficient matrices, $\boldsymbol{B}$ and $\boldsymbol{\Sigma}$. Recall, $\mathbb{E}(\boldsymbol{y}_t) = 0$ such that the companion form representation of the model reads

$$\boldsymbol{\mathcal{Y}}_t = \boldsymbol{F}\boldsymbol{\mathcal{Y}}_{t-1} + \boldsymbol{\mathcal{E}}_t,$$

where

$$\boldsymbol{\mathcal{E}}_t \equiv \begin{bmatrix} \boldsymbol{e}_t \\ \boldsymbol{e}_{t-1} \\ \vdots \\ \boldsymbol{e}_{t-p+1} \end{bmatrix}, \quad \boldsymbol{F} \equiv \begin{bmatrix} \boldsymbol{B}' & \\ \boldsymbol{\mathcal{I}}_{(p-1)n} & \boldsymbol{0}_{(p-1)n,n} \end{bmatrix}, \quad \text{and} \quad \boldsymbol{\Omega} \equiv \mathbb{E}(\boldsymbol{\mathcal{E}}_t\boldsymbol{\mathcal{E}}_t') = \begin{bmatrix} \boldsymbol{\Sigma} & \boldsymbol{0}_{n,(p-1)n} \\ \boldsymbol{0}_{(p-1)n,n} & \boldsymbol{0}_{n,n} \end{bmatrix}.$$

The covariance matrix of $\boldsymbol{\mathcal{Y}}_t$ is given by

$$\begin{aligned}\boldsymbol{\Gamma} &= \mathbb{E}\left[\left(\boldsymbol{F}\boldsymbol{\mathcal{Y}}_{t-1}+\boldsymbol{\mathcal{E}}_t\right)\left(\boldsymbol{F}\boldsymbol{\mathcal{Y}}_{t-1}+\boldsymbol{\mathcal{E}}_t\right)'\right]\\ &= \mathbb{E}\left(\boldsymbol{F}\boldsymbol{\mathcal{Y}}_{t-1}\boldsymbol{\mathcal{Y}}_{t-1}'\boldsymbol{F}'+\boldsymbol{\mathcal{E}}_t\boldsymbol{\mathcal{E}}_t'+\boldsymbol{F}\boldsymbol{\mathcal{Y}}_{t-1}\boldsymbol{\mathcal{E}}_t'+\boldsymbol{\mathcal{E}}_t+\boldsymbol{\mathcal{Y}}_{t-1}'\boldsymbol{F}'\right)\\ &= \boldsymbol{F}\mathbb{E}\left(\boldsymbol{\mathcal{Y}}_{t-1}\boldsymbol{\mathcal{Y}}_{t-1}'\right)\boldsymbol{F}'+\mathbb{E}\left(\boldsymbol{\mathcal{E}}_t\boldsymbol{\mathcal{E}}_t'\right)\\ &= \boldsymbol{F}\boldsymbol{\Gamma}\boldsymbol{F}'+\boldsymbol{\Omega}.\end{aligned}$$

The solution to this equation can be found in Hamilton (1994, ch. 10). It reads

$$\text{vec}\left(\boldsymbol{\Gamma}\right)=\left[\boldsymbol{\mathcal{I}}_{(np)^2}-\left(\boldsymbol{F}\otimes\boldsymbol{F}\right)\right]^{-1}\boldsymbol{\Omega}. \tag{D.2}$$

Equation (D.2) directly relates $\boldsymbol{\Gamma}$ to the VAR matrices $\boldsymbol{B}$ and $\boldsymbol{\Sigma}$ and implies that the off-diagonal elements of $\widehat{\boldsymbol{V}}$ are generally non-zero.

D.2 Monte Carlo exercise

A simulation experiment shows the impact of the prior covariance on its posterior counterpart. The model is chosen as simple as possible in order to keep the exposition clear. "Home country's" data is generated from an AR(2) model of the form

$$y_t=\rho_1 y_{t-1}+\rho_2 y_{t-2}+e_t,$$

with $e_t\sim\mathcal{N}(0,\sigma^2)$. The specification is chosen such that the parameters feature a stationary model, in particular $\rho_1=0.5$, $\rho_2=0.2$, and $\sigma^2=1$.

An artificial sample of T observations is generated from the above DGP. At the same time another "foreign country" model

$$y_t^*=\rho_1^* y_{t-1}^*+\rho_2^* y_{t-2}^*+e_t^*$$

with $e_t^*\sim\mathcal{N}(0,\sigma^2)$ is simulated. Its coefficients are similar to the "home country" model, i.e. $\rho_j*=w\cdot\rho_j$ for $j=\{1,2\}$,[40] where w is a parameter that controls the degree of heterogeneity. In the following, three estimates with respect to the coefficient covariance are computed with expanding sample size

1. "Maximum likelihood": $\boldsymbol{X}'\boldsymbol{X}/T$
2. "Full posterior": $\boldsymbol{X}'\boldsymbol{X}/T+\lambda\left(\boldsymbol{X}^*\right)'\boldsymbol{X}^*/T$
3. "'Diagonal posterior": $\boldsymbol{X}'\boldsymbol{X}/T+\lambda\text{diag}\left[\sum_{t=1}^{T}\left(\boldsymbol{y}_{t-1}^*\right)^2,\sum_{t=1}^{T}\left(\boldsymbol{y}_{t-2}^*\right)^2\right]$

Note that case 2 "full posterior" represents the strategy proposed in the main text. The respective derivation can be found in section B of the appendix. The tightness parameter is chosen $\lambda=0.5$. The reference point is the true covariance matrix which can be computed from equation (D.2)

$$\boldsymbol{\Gamma}=\begin{bmatrix}1.7094 & 1.0684\\ 1.0684 & 1.7094\end{bmatrix}.$$

[40] Other types of heterogeneity, e.g. $\rho_j^*=w+\rho_j$ and alternative calibrations (e.g. $\rho_1=0.3$, $\rho_2=0.4$) have been considered. Results are similar and available upon request.

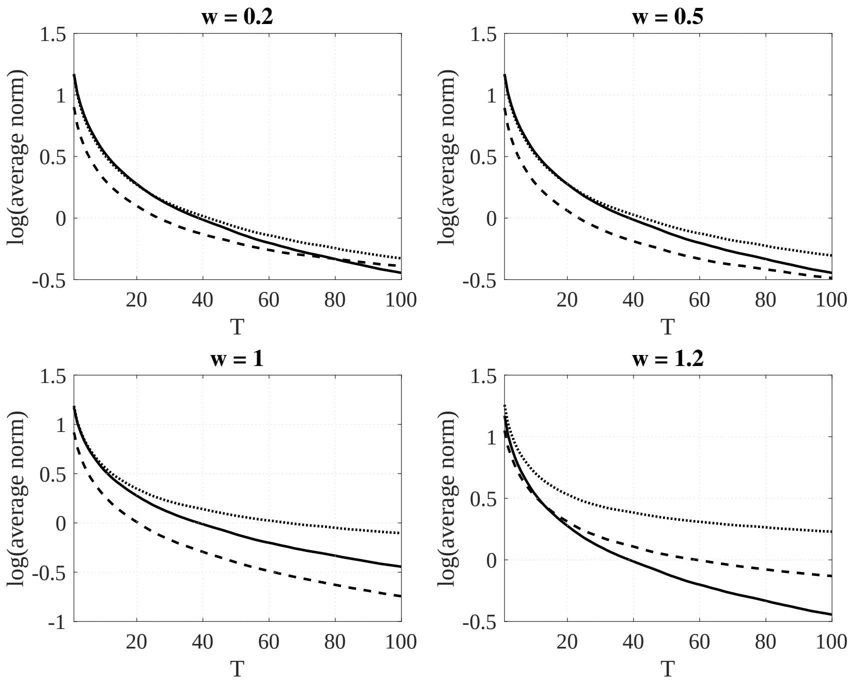

Figure D.1: Log of average norm of the difference between estimates of the coefficient covariance matrix and $\mathbf{\Gamma}$ as a function of the sample size for $\lambda = 0.5$. The solid line denotes the maximum likelihood estimator, the dashed line depicts the estimate according to strategy 2, and the dotted line represents the estimate with respect to strategy 3.

Within the exercise a total of 10,000 samples are generated and results are presented in terms of averages over those samples. In order to evaluate the estimates, the difference to the theoretical counterpart $\mathbf{\Gamma}$ is computed and the natural logarithm[41] of the Euclidean norm serves as the respective deviation measure. Figure D.1 shows that the maximum likelihood estimator approaches the true covariance matrix as the sample size T increases. The estimates with respect to strategies 2 and 3 generally do not approach $\mathbf{\Gamma}$ as $T \to \infty$ because the prior distorts the maximum likelihood estimate according to the heterogeneity of both country models. Consequently, the prior tightness parameter λ should, ceteris paribus, decrease in T. However, the "full posterior" estimate seems to be advantageous in small samples if the foreign country model is not too different from the home country model. In addition, the dashed line ("full posterior") provides better estimates compared to its "diagonal posterior" counterpart. The latter result is important because it illustrates the advantage of prior covariances which are not restricted to diagonal matrices.

A second exercise considers the role of the tightness parameter λ given a particular degree of heterogeneity $w = 0.4$. Figure D.2 depicts the respective results. Simulations show that setting $\lambda > 0$ improves the estimate of the covariance if the sample length is limited. However, in larger samples the researcher should probably choose a smaller values of λ. The intuition is that if there is "enough" data available, additional foreign country data tends to bias results. It follows that an "optimal" value with respect to the

[41] In order for the differences between lines to be better visible to the reader, I decided to depict the respective natural logarithm. The non-logarithmic counterpart of the maximum likelihood estimator converges towards zero as expected.

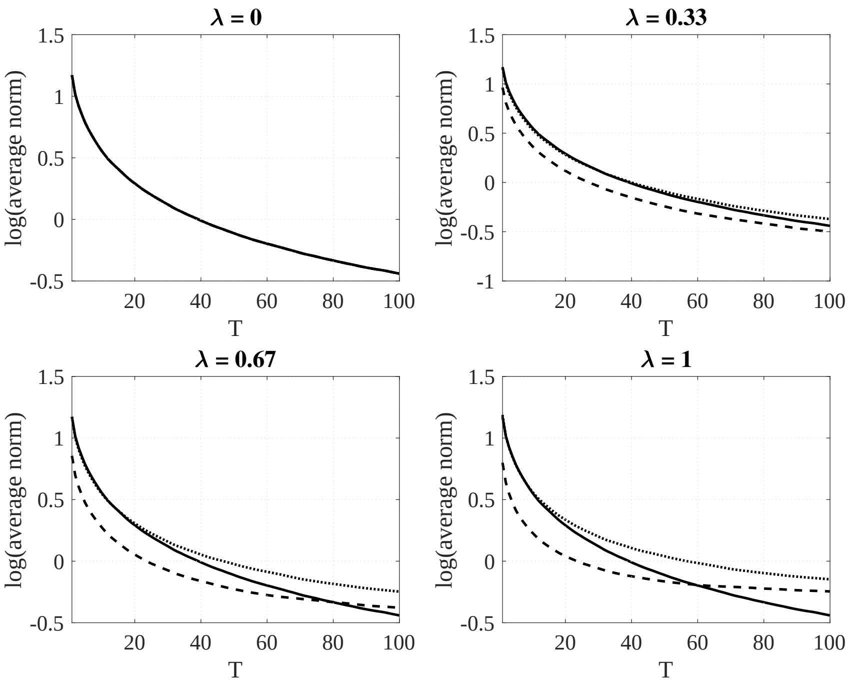

Figure D.2: Log of average norm of the difference between estimates of the coefficient covariance matrix and $\mathbf{\Gamma}$ as a function of the sample size for $w = 0.4$. The solid line denotes the maximum likelihood estimator, the dashed line depicts the estimate according to strategy 2 and the dotted line represents the estimate with respect to strategy 3.

tightness parameter should depend negatively on both the degree of heterogeneity and the sample size. This finding is confirmed by an additional exercises in section E of the appendix, where the baseline model is estimated with only $T = 40$ observations.

The calibration of the DGP may illustrate the point raised by Del Negro and Schorfheide (2011) who criticize the fact that diagonal covariance matrices of slope coefficients implicitly allocate more weight to unreasonable dynamics of the endogenous variables. Figure D.3 plots contours of the bivariate normal distribution $\mathcal{N}([\rho_1, \rho_2]', \mathbf{\Gamma})$ (left panels) and contours of the same distribution except that covariances are set to zero (right panels). The upper panels depict the boundaries with respect to the dynamic properties of the AR(2) model (see e.g. Hamilton, 1994, ch. 1). Within the solid triangle, dynamics are non-explosive. The region above the dashed line indicates non-oscillating behavior. Dynamics might be called "unreasonable" if parameter combinations (ρ_1, ρ_2) imply explosiveness. In order to evaluate which of the two distributions features more reasonable dynamics, a sample of 10^8 observations is simulated. Draws that fulfill the non-explosiveness conditions (illustrated by the dots in the lower panels) are counted. It turns out that the distribution that implies non-zero correlations between ρ_1 and ρ_2 generates about 17.5% more draws within the solid triangle implying reasonable dynamics. Similar results are obtained for alternative calibrations (e.g. $\rho_1 = 0.3$, $\rho_2 = 0.4$) of the AR(2) model as well as if non-complexity of eigenvalues is an additional requirement.

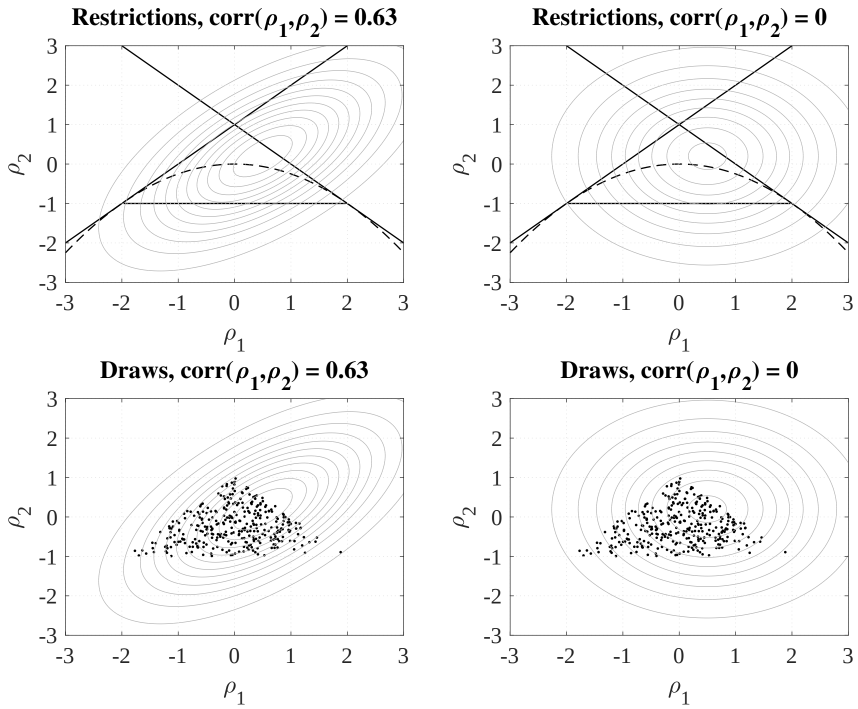

Figure D.3: Contours (gray lines) of bivariate normal distributions with respect to $[\rho_1, \rho_2]$. Panels to the left assume the theoretical covariance $\mathbf{\Gamma}$ and panels to the right set the off-diagonal elements of $\mathbf{\Gamma}$ to zero. The upper panels depict non-explosiveness regions and the lower panels illustrate non-explosive draws from the respective distributions.

E Short sample

In order to shed more light on the determinants of the tightness parameter posterior distribution $p(\lambda_c|\boldsymbol{Y}_c)$, the following exercise uses only the first 40 observations to estimate the baseline VAR model from the main text. The Monte Carlo study in section D.2 suggested that the estimate for λ_c should be a decreasing function in the sample size. Figure E.1 shows percentiles of the posterior distribution of λ_c. Comparing those estimates with figure E.1 in the main text shows that the short sample analysis requires tighter prior hyperparamters λ_c. This finding supports the above hypothesis and improves the plausibility and reliability of the econometric approach.

F Loose tightness parameters

As mentioned in the main text, if the value of λ_c is chosen too small, i.e. below the threshold $\widetilde{\lambda_c} = (p+1)n/T_t^*$, the prior distribution ends up being not well-defined. Analyses were conducted allocating zero prior weight with respect to $\lambda_c < \widetilde{\lambda_c}$. If the researcher wants to explicitly allow for smaller values of λ_c, the prior distribution has to be modified accordingly. Equation (4) becomes

$$p(\boldsymbol{\beta}_c, \boldsymbol{\Sigma}_c|\lambda_c) \propto [f(\boldsymbol{Y}_c^*|\boldsymbol{\beta}_c, \boldsymbol{\Sigma}_c)]^{\lambda_c} \times \det(\boldsymbol{\Sigma})^{-\frac{n+1}{2}} .$$

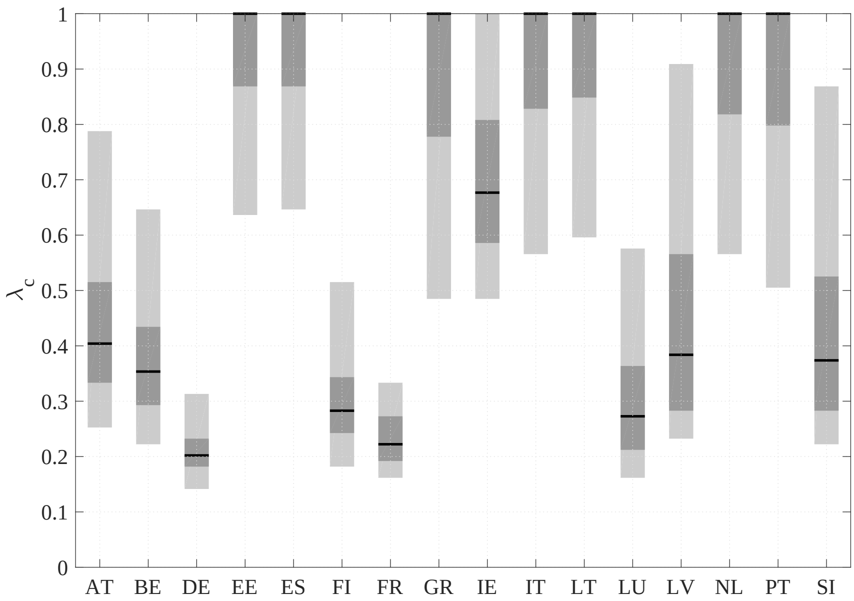

Figure E.1: Posterior distribution of the tightness parameters λ_c across countries. The gray shaded areas denote the 68 percent and the 90 percent Bayesian (highest probability density) credible intervals, respectively. Only the first 40 observations are used in the estimation procedure.

As $\lambda_c \to 0$, this specification leads to a diffuse prior distribution for the VAR model (see e.g. Kadiyala and Karlsson, 1997). Results with respect to the prior tightness are reported in figure F.2. The modified prior distribution generates similar results compared to the baseline analysis depicted in figure 1. With respect to most countries λ_c is only mildly lower. Estonia and and Lithuania represent examples where the model estimate is affected more strongly.

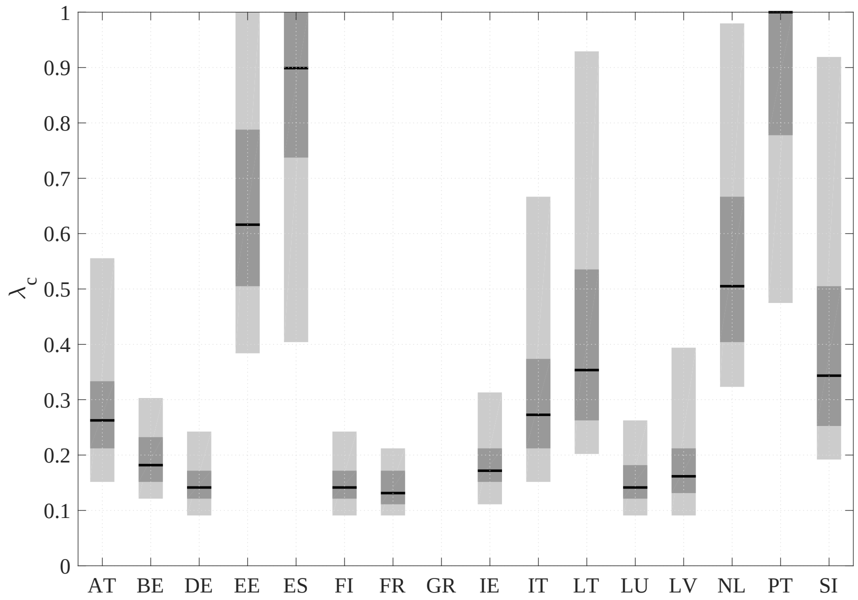

Figure F.2: Posterior distribution of the tightness parameters λ_c across countries. The gray shaded areas denote the 68 percent and the 90 percent Bayesian (highest probability density) credible intervals, respectively. Prior distribution accounts for low values of λ_c.