Essays on Models with Time-Varying Parameters for Forecasting and Policy Analysis

A thesis submitted in partial fulfillment of the requirements of the degree of Doctor of Philosophy (Ph.D.) in Economics

Fabrizio Venditti

School of Economics and Finance

Queen Mary, University of London

April 2016

a mia moglie Anna e ai miei bimbi Samuele e Davide

Declaration

I wish to declare

I, Fabrizio Venditti, confirm that the research included within this thesis is my own work or that where it has been carried out in collaboration with, or supported by others, that this is duly acknowledged below and my contribution indicated. Previously published material is also acknowledged below.

I attest that I have exercised reasonable care to ensure that the work is original, and does not to the best of my knowledge break any UK law, infringe any third party's copyright or other Intellectual Property Right, or contain any confidential material.

I accept that the College has the right to use plagiarism detection software to check the electronic version of the thesis.

I confirm that this thesis has not been previously submitted for the award of a degree by this or any other university.

Details of collaboration and publications:

Chapter 1 draws from joint work with Massimiliano Marcellino and Mario Porqueddu and published in the Journal of Business and Economic Statistics.[1]
Chapter 2 draws from joint work with George Kapetanios and Massimiliano Marcellino.
Chapter 3 draws from joint work with Marianna Riggi and published in the Journal of Economic Dynamics and Control.[2]

[1] Marcellino, M., Porqueddu, M. and F. Venditti, (2016) Short-Term GDP Forecasting With a Mixed-Frequency Dynamic Factor Model With Stochastic Volatility, Journal of Business & Economic Statistics, 34:1, 118-127

[2] Riggi, M. and F. Venditti (2015), The time varying effect of oil price shocks on euro-area exports, Journal of Economic Dynamics and Control, Elsevier, vol. 59(C), pages 75-94.

Chapter 4 draws from joint work with Marianna Riggi and published in International Finance.[3]

Signature: Fabrizio Venditti
Date: 4th of April, 2016

[3]Riggi, M and F. Venditti (2015), Failing to Forecast Low Inflation and Phillips Curve Instability: A Euro-Area Perspective. International Finance, 18: 47-68

Acknowledgements

I would like to thank my supervisors, George Kapetanios and Massimiliano Marcellino, for their guidance and support throughout my Ph.D. studies and the co-authors and friends with whom I have worked in the past four years, in particular Marianna Riggi, Mario Porqueddu, and above all Davide Delle Monache and Ivan Petrella, who have been a constant source of inspiration. I owe most of my professional growth to the interaction over the years with my colleagues in the Research Department of the Bank of Italy.

I would like to thank my parents, whose dedication to me and my brother has been an invaluable example that I try to follow every day; my brother Francesco, for patiently playing with me when we were kids, despite him being much better than me at most of the games we played; my children Samuele and Davide, whose morning smiles constantly reminded me where joy lies. Above all I would like to thank my wife Anna, for supporting me unconditionally in my choice of going back to studying, for taking care of our children while I was playing around with econometrics, for making me the happiest person on earth.

I acknowledge the School of Economics and Finance at Queen Mary, University of London, for generous funding and for giving me a second opportunity to pursue graduate studies.

Abstract

The aim of this thesis is the development and the application of econometric models with time-varying parameters in a policy environment.

The popularity of these methods has run in parallel with advances in computing power, which has made feasible estimation methods that until the late '90s would have been unfeasible. Bayesian methods, in particular, benefitted from these technological advances, as sampling from complicated posterior distributions of the model parameters became less and less time-consuming. Building on the seminal work by Carter and Kohn (1994) and Jacquier, Polson, and Rossi (1994), bayesian algorithms for estimating Vector Autoregressions (VARs) with drifting coefficients and volatility were independently derived by Cogley and Sargent (2005) and Primiceri (2005).

Despite their increased popularity, bayesian methods still suffer from some limitations, from both a theoretical and a practical viewpoint. First, they typically assume that parameters evolve as independent driftless random walks. It is therefore unclear whether the output that one obtains from these estimators is accurate when the model parameters are generated by a different stochastic process. Second, some computational limitations remain as only a limited number of time series can be jointly modeled in this environment. These shortcomings have prompted a new line of research that uses non-parametric methods to estimate random time-varying coefficients models. Giraitis, Kapetanios, and Yates (2014) develop kernel estimators for autoregressive models with random time-varying coefficients and derive the conditions under which such estimators consistently recover the true path of the model coefficients. The method has been suitably adapted by Giraitis, Kapetanios, and Yates (2012) to a multivariate context.

In this thesis I make use of both bayesian and non-parametric methods, adapting them (and in some cases extending them) to answer some of the research questions that, as a Central Bank economist, I have been

tackling in the past five years. The variety of empirical exercises proposed throughout the work testifies the wide range of applicability of these models, be it in the area of macroeconomic forecasting (both at short and long horizons) or in the investigation of structural change in the relationship among macroeconomic variables.

The first chapter develops a mixed frequency dynamic factor model in which the disturbances of both the latent common factor and of the idiosyncratic components have time varying stochastic volatility. The model is used to investigate business cycle dynamics in the euro area, and to perform point and density forecast. The main result is that introducing stochastic volatility in the model contributes to an improvement in both point and density forecast accuracy.

Chapter 2 introduces a nonparametric estimation method for a large Vector Autoregression (VAR) with time-varying parameters. The estimators and their asymptotic distributions are available in closed form. This makes the method computationally efficient and capable of handling information sets as large as those typically handled by factor models and Factor Augmented VARs (FAVAR). When applied to the problem of forecasting key macroeconomic variables, the method outperforms constant parameter benchmarks and large Bayesian VARs with time-varying parameters. The tool is also used for structural analysis to study the time-varying effects of oil price innovations on sectorial U.S. industrial output.

Chapter 3 uses a bayesian VAR to provide novel evidence on changes in the relationship between the real price of oil and real exports in the euro area. By combining robust predictions on the sign of the impulse responses obtained from a theoretical model with restrictions on the slope of the oil demand and oil supply curves, oil supply and foreign productivity shocks are identified. The main finding is that from the 1980s onwards the relationship between oil prices and euro area exports has become less negative conditional on oil supply shortfalls and more positive conditional on foreign productivity shocks. A general equilibrium model is used to shed some light on the plausible reasons for these changes.

Chapter 4 investigates the failure of conventional constant parameter models in anticipating the sharp fall in inflation in the euro area in 2013-2014. This forecasting failure can be partly attributed to a break in the

elasticity of inflation to the output gap. Using structural break tests and non-parametric time varying parameter models this study shows that this elasticity has indeed increased substantially after 2013. Two structural interpretations of this finding are offered. The first is that the increase in the cyclicality of inflation has stemmed from lower nominal rigidities or weaker strategic complementarity in price setting. A second possibility is that real time output gap estimates are understating the amount of spare capacity in the economy. I estimate that, in order to reconcile the observed fall in inflation with the historical correlation between consumer prices and the business cycle, the output gap should be wider by around one third.

Contents

List of Tables

List of Figures

Chapter 1

Short-term GDP forecasting with a mixed frequency dynamic factor model with stochastic volatility

1.1 Introduction

The conduct of monetary and fiscal policy relies on the timely assessment of current and future economic conditions. However, the task of providing a timely picture of the cyclical position is hindered by the publication lag of crucial economic indicators. GDP data, for example, are usually published with a significant delay both in the US and in the euro area. Important quantitative monthly indicators, like industrial production indexes, suffer from a similar publication delay. Survey data, on the other hand, provide very timely information as they are published roughly at the end of the reference month. Unfortunately, forecasts based on qualitative data only are known to be much less reliable than predictions based on quantitative information, see Banbura and Runstler (2011). The econometric literature has progressed significantly in the field of short-term forecasting in the past decade, and a number of tools have been developed, capable of dealing with the asynchronous timing of data releases, integrating data at different frequencies and dissecting the information content of monthly releases for tracking quarterly variables. Small and large scale factor models, in particular, have become the workhorse for short-term forecasting. Mariano and Murasawa (2003, henceforth MM03), building on the coincident indicator developed by Stock and Watson (1989), have proposed a unified framework for modeling quarterly GDP together with monthly indicators in a small scale factor model. The approach has recently been extended by Camacho and Perez-Quiros (2010) to accommodate real time issues and different GDP releases. On the large data side, research by Giannone, Reichlin, and Small (2008, henceforth GRS), Angelini, Camba-Mendez, Giannone, Reichlin, and Runstler (2011) and Banbura and Modugno (2010) has documented the predictive content of a large number of indicators for GDP growth. These models are nowadays used on a regular basis to inform decision makers both at Central Banks as well as in private institutions.[1]

Although the literature has moved very rapidly, there are still some gaps between the demands posed by policy makers and the answers that the models discussed above can provide. In particular, policy makers have become more and more interested in having not only point forecasts, but also a model based assessment of the uncertainty surrounding the outlook. This is testified by the number of Central Banks that have started publishing fan charts and confidence bands around their medium/long term forecasts (Bank of England, Bank of Canada, Norges Bank, South Africa Reserve, the Sveriges Riksbank, the Bank of Italy and the US Fed). Despite the growing preference for a probabilistic assessment of economic projections, however, the focus of short-term forecasting models is still on point forecasts.

Another open issue relates to parameter instability. As economic systems evolve and are hit by large shocks, the link between different indicators is likely to change over time. The issue of forecast failure in the presence of structural breaks, which has been explored extensively in the case of points forecasts, has been recently extended to density forecasting. In particular Jore, Mitchell, and Vahey (2010) show that changes in the underlying data generating process can severely worsen the accuracy of density forecasts produced with models with constant parameters. Building on an intuition that dates back to Sims (1993), Clark (2011) finds that allowing for stochastic shifts in the volatility of the shocks significantly increases the accuracy of density forecast produced by a BVAR.

In this paper we take stock of these issues and develop a mixed frequency small scale factor model that is suitable for producing density forecasts and that allows for time variation in some of the parameters. We start off with the basic setup of MM03 and twist it in two directions. First, we cast the model in a Bayesian estimation framework, which makes the model suitable for producing density forecasts. Second, following Clark (2011), we extend the model to allow for random shifts in the volatility of the underlying shocks.[2]

[1] An alternative approach to short-term forecasting with mixed frequency data is based on the MIDAS regressions introduced by Ghysels, Santa-Clara, and Valkanov (2004), see e.g. Clements and Galvão (2008), Foroni and Marcellino (2014) and Marcellino and Schumacher (2008) for macroeconomic applications. Mixed frequency VARs provide a third option, see e.g. Kuzin, Marcellino, and Schumacher (2011).

[2] See Baumeister, Liu, and Mumtaz (2013), Del Negro and Otrok (2008) and Korobilis (2013)

After showing how to estimate the model we turn to an empirical application in which we use a small number of monthly indicators to predict quarterly GDP growth in the euro area. We present three sets of results. First we show how macroeconomic releases not only improve point forecast accuracy but also increase the precision of density forecasts and reduce the width of forecast intervals. Second, we illustrate how, in a given quarter, our new tool can be applied not only to interpret the *news* content of monthly releases, like in Banbura and Modugno (2010), but also to assess how much *confidence* the model places on the revisions implied by the release of monthly indicators. Third, we design a (pseudo) real time out of sample forecasting exercise and evaluate both the point and density forecasts produced by the model. In line with Clark (2011) we find that the introduction of stochastic volatility leads to an improvement in both point and density forecast accuracy.

The paper is structured as follows. In section 2.2 we describe the model. In section 1.3 we discuss the main steps of the Gibbs sampler used for simulating the posterior distribution of the parameters. Section 1.4 presents the empirical application, with robustness checks discussed in 1.5. Section 1.6 concludes. An Appendix contains additional technical details and information.

1.2 The model

Let $Y_{q,t}$ be a quarterly series, which can be seen as a monthly variable with its value associated to the third month of the quarter and missing observations in the first two months, and $Y_{m,t}$ a vector of k monthly series $Y_{mj,t}$, for $j = 1, 2, \ldots, k$ (from this point onwards we use the convention that whenever we write mj we mean the j_{th} element in the vector of monthly variables, for $j = 1, 2, \ldots, k$). Now, as in MM03, let $Y_{q,t}$ be the geometric mean of a latent random variable $Y^{\star}_{qt}$ such that:

$$lnY_{q,t} = \frac{1}{3}(lnY^{\star}_{q,t} + lnY^{\star}_{q,t-1} + lnY^{\star}_{q,t-2}) \tag{1.1}$$

Filtering both sides with the filter $(1 - L^3)$, after some simple manipulation yields:

$$y_{q,t} = \frac{1}{3}y^{\star}_{q,t} + \frac{2}{3}y^{\star}_{q,t-1} + y^{\star}_{q,t-2} + \frac{2}{3}y^{\star}_{q,t-3} + \frac{1}{3}y^{\star}_{q,t-4} \tag{1.2}$$

where small case letters indicate growth rates over the previous three months: $y_{q,t} = \Delta_3 lnY_{q,t}$. We assume a dynamic (single) factor[3] model for the latent process

for examples of Bayesian dynamic factor models with stochastic volatility. In their models they do not handle mixed frequency data.

[3]The use of more than one common factor does not pose any additional technical difficulty as it would simply result in an enlargement of the state vector in the State Space representation of

$y^{\star}_{q,t}$ and the monthly observed variables $y_{m,t}$, that is:

$$\begin{pmatrix} y^{\star}_{q,t} \\ y_{m,t} \end{pmatrix} = \begin{pmatrix} \alpha^{\star}_1 \\ \alpha_2 \end{pmatrix} + \begin{pmatrix} \beta_q f_t \\ \beta_m f_t \end{pmatrix} + \begin{pmatrix} u_{q,t} \\ u_{m,t} \end{pmatrix} \tag{1.3}$$

Since the variable $y^{\star}_{q,t}$ is not observed, the model can be rewritten in terms of the observable variable $y_{q,t}$ using the identity (1.2), resulting in the following system of measurement equations:

$$\begin{pmatrix} y_{q,t} \\ y_{m,t} \end{pmatrix} = \begin{pmatrix} \alpha_1 \\ \alpha_2 \end{pmatrix} + \begin{pmatrix} \beta_1(\frac{1}{3}f_t + \frac{2}{3}f_{t-1} + f_{t-2} + \frac{2}{3}f_{t-3} + \frac{1}{3}f_{t-4}) \\ \beta_2 f_t \end{pmatrix} + \begin{pmatrix} \frac{1}{3}u_{q,t} + \frac{2}{3}u_{q,t-1} + u_{q,t-2} + \frac{2}{3}u_{q,t-3} + \frac{1}{3}u_{q,t-4} \\ u_{m,t} \end{pmatrix} \tag{1.4}$$

where $\alpha_1 = 3\alpha^{\star}_1$.

The law of motions of the factor and of the the idiosyncratic disturbances of the quarterly and monthly variables are described by the following:

$$\Phi_f(L)f_t = v_t e^{\lambda_{f,t}/2} \tag{1.5}$$

$$\Phi_q(L)u_{q,t} = \epsilon_{q,t}\sigma_q e^{\lambda_{q,t}/2} \tag{1.6}$$

$$\Phi_{mj}(L)u_{mj,t} = \epsilon_{mj,t}\sigma_{mj} e^{\lambda_{mj,t}/2} \qquad j = 1, \ldots, k \tag{1.7}$$

where v_t, $\epsilon_{q,t}$ and $\epsilon_{mj,t}$ are uncorrelated N(0,1) and the $\Phi_i(L)$ polynomials are lag polynomials of order p_i:

$$\Phi_i(L) = 1 - \phi^i_1 \mathrm{L} - \phi^i_2 \mathrm{L}^2 - \cdots - \phi^i_{p_i} \mathrm{L}^{p_i} \tag{1.8}$$

for $i = f, q, mj$. The log-volatilities $\lambda_{i,t}$ follow a driftless random walk:

$$\lambda_{i,t} = \lambda_{i,t-1} + \theta_{i,t}\sigma_{\lambda,i} \qquad \theta_{i,t} \sim N(0,1) \tag{1.9}$$

for $i = f, q, mj$, and are assumed to be independent across equations.[4] A more compact state space representation of the model is the following:

$$\boldsymbol{y}_t = F\boldsymbol{\mu}_t \tag{1.10}$$

$$\boldsymbol{\mu}_t = H\boldsymbol{\mu}_{t-1} + \eta_t \qquad \eta_t \sim N(0, Q_t) \tag{1.11}$$

$$\Lambda_t = \Lambda_{t-1} + \zeta_t \qquad \zeta_t \sim N(0, \Xi) \tag{1.12}$$

the model. The case of two factors is discussed theoretically and empirically in the Appendix.

[4]In the empirical application we consider as a robustness check an alternative specification in which the stochastic volatilities evolve as an AR(1).

where $\boldsymbol{y}_t$ collects both quarterly and monthly variables, the state vector $\boldsymbol{\mu}_t$ includes the unobserved factor f_t and the idiosyncratic components ($u_{q,t}$ and $u_{m,t}$), the matrix F collects the factor loadings, H collects the autoregressive parameters of the unobserved factors and of the idiosyncratic components, the time varying variance matrix Q_t is a diagonal matrix with elements $e^{\lambda_{f,t}}, \sigma_q^2 e^{\lambda_{q,t}}, \sigma_m^2 e^{\lambda_{mj,t}}$, Λ_t is the vector of drifting volatilities, Ξ is a diagonal matrix collecting the variances of the log-volatilities disturbances.

This model nests the one proposed by MM03, which can be recovered by shutting off the drifting volatilities, that is by setting $\Lambda_0 = 0$ and $\Xi = 0$. In this case the matrix Q_t is replaced by a time invariant matrix, Q.

To identify the model parameters some restrictions are placed. First, the scale of the factor loadings and of the factor cannot be separately identified, so we restrict the variance of the errors of the common factors to be 1 (see equation 1.5). Second, like Del Negro and Otrok (2008), we fix to zero the initial condition of the stochastic volatilities.

1.3 Model Estimation

The model is estimated with Bayesian methods using a Metropolis within Gibbs sampling procedure. The algorithm consists of six blocks, which we briefly describe. More details on the sampler can be found in the Appendix.

1.3.1 Steps 1 and 2: drawing F and the time constant elements of Q_t

Since the model disturbances are uncorrelated, elements of the F matrix can be drawn row by row (equation by equation). Take the i^{th} measurement equation:

$$y_{i,t} = F(i)\boldsymbol{\mu}_t = \beta(i,L)f_t + \Phi_i(L)^{-1}\epsilon_{i,t}\sigma_i e^{\lambda_{i,t}/2} \tag{1.13}$$

Conditioning on f_t, $\Phi_i(L)$ and $\lambda_{i,t}$, this is a standard regression with autocorrelated and heteroscedastic disturbances. Pre-multiplying by $\Phi_i(L)$ and dividing by $e^{\lambda_{i,t}/2}$ one obtains a standard regression model with homoscedastic, uncorrelated residuals. Positing a Normal-gamma conjugate prior, the conditional posterior for $\beta(i,L)$ and σ_i is also Normal-gamma.

1.3.2 Step 3: drawing H

The parameters in the transition matrix H can also be drawn row by row. Take the i^{th} transition equation:

$$\mu_{i,t} = \sum_{j=1}^{p_i} \phi_j \mu_{i,t-j} + \eta_{i,t} \tag{1.14}$$

Conditioning on $\mu_{i,t}$ and on the ith element of the Q_t matrix ($q_{i,t}$), this is a regression with heteroscedastic residuals. The residuals can be whitened by dividing by $q_{i,t}$. Positing a Normal prior for the regression coefficients the conditional posterior is also Normal.[5]

1.3.3 Step 4 and 5: drawing the stochastic volatilities

There are a number of methods for drawing the stochastic volatilities $\lambda_{i,t}$ and the related variances $\sigma_{\lambda,i}$. We employ the Jacquier, Polson, and Rossi (1994) algorithm, which involves drawing from a log-normal density and a Metropolis acceptance step. Details on the algorithm can be found in Cogley and Sargent (2005), Appendix B.2.5.

1.3.4 Step 6: drawing μ_t

Conditioning on all the other parameters and on the data, draws of the state vector are obtained via the disturbance smoother proposed by Koopman and Durbin (2003).

1.4 Empirical application: short-term forecasts of euro area GDP

We apply the model to the problem of forecasting euro area GDP growth at short horizons. Our information set consists of nine indicators, namely our target variable, which is the rate of growth of quarterly GDP, two Industrial Production indicators (the total index and the index for the Pulp and Paper sector), four surveys (the Germany IFO Business Climate Index, the Composite Purchasing Manager Index for the euro area, the Michigan Consumer Sentiment for the US, the euro area Economic Sentiment Indicator), the bilateral US dollar euro exchange rate and a the difference between the 3 months and the 10 years spread on US Government Bonds. Data start in January 1991 and end in May 2011. The indicators, listed in

[5] When drawing from the conditional posterior we discard explosive roots.

Table 1.1, were selected from a large pool of candidate series adapting the selection algorithm used by Camacho and Perez-Quiros (2010) to our Bayesian setting, see the Appendix for details. The empirical specification of the model also follows closely the one proposed by Camacho and Perez-Quiros (2010). In particular we use a single factor that summarizes the current state of the business cycle. In this setting the Industrial Production indexes and the interest rate spread load on the common factor contemporaneously. Survey data, on the other hand, are treated as if they were in phase with the year-on-year growth rate of the Industrial Production index, therefore loading a 11 terms moving average of the common factor.[6] We also let the bilateral exchange rate enter the model in year-on-year percentage growth, the rationale being that pricing to market is likely to buffer temporary exchange rate short-term movements with a variation in profit margins so that only more persistent changes impact on economic growth.

Our empirical analysis proceeds as follows. After a brief discussion on the priors, we present estimates obtained on the full sample, to gauge the relative contributions of the various indicators to the common factor and also to evaluate if the model actually captures any significant shifts in the variance of the common and idiosyncratic errors. We then turn to three empirical exercises. In the first one we discuss the typical situation of a forecaster that is required to update her forecasts at each new data release. In this context we replicate the analysis of *news* performed by Giannone, Reichlin, and Small (2008) and evaluate how point forecast accuracy is affected by data releases. We take advantage of the Bayesian nature of our model and extend GRS results to examine how new data affects density forecast accuracy and the width of forecast intervals. We then turn to a different concept of *news*, introduced in the literature by Banbura and Modugno (2010). We show how our set up adds a new dimension to their tool, as one can use draws from the posterior to derive a measure of uncertainty around the *news* content of each data (or block of data) release. Finally, we conduct an out-of-sample forecast exercise in which we assess the point and density forecasting performance of our model.

1.4.1 Priors

To set the prior hyperparameters we retain a three years training sample. Since

[6] Details on the state space matrices for this specification can be found in the Appendix. We also experimented with a different specification in which we relaxed this restriction and let the survey indicators load freely on 12 distributed lags of the common factor. This modification worsened slightly the results. Our intuition for this is that as the model is already heavily parametrized restricting the model space leads to more efficient estimates. Also, notice that our setup allows for serial correlation in the idiosyncratic components, so that any phase shift induced by our restriction will be picked up by the AR(2) structure of the idiosyncratic terms.

the model features an unobserved component that is common to the indicators, we obtain an initial estimate of the common factor f_t as the *cross-sectional* average of the monthly indicators, $\hat{f}_t$, over this training sample. Conditioning on this, an estimate of the factor loadings is obtained with an OLS regression of the indicators on $\hat{f}_t$. The prior distributions of the factor loadings are then centered around this $\hat{\beta}_{OLS}$ with a large variance, equal to $10^3 V(\beta_{OLS})$. By regressing the residuals of these regressions on their first two lags we also obtain an estimate of the autoregressive parameters of the idiosyncratic shocks to the observable indicators. The prior distributions of the ϕ are then centered on this estimate and their variance is set to $10^3 V(\phi_{OLS})$. Similarly, we use this training sample estimate of the factor $\hat{f}_t$ to set the prior mean and variance of $\phi_{f,1}, \phi_{f,2}$. Finally, we need to set the degrees of freedom and the scales of the prior inverse-Gamma distributions for the variances of the idiosyncratic shocks. For the constant terms, σ_q^2 and σ_m^2, we set the degrees of freedom to 1 and the scale parameters to the squared sum of the residuals of the OLS estimates obtained on the training sample data. For the time varying volatilities, $\sigma_{\lambda,i}^2$ we adopt a tighter belief, due to the fact that our sample is relatively short and that this is the only source of time variation in the model. Following Del Negro and Otrok (2008), we parametrize the degrees of freedom of these prior distributions as a fraction of the actual sample size and set it to $\frac{T}{10}$ (so that the weight of the prior relative to that of the data in determining the posterios is 1 to 10) and set the scale parameter to 0.1. We check the robustness of these assumptions in Section 1.5. The Gibbs sampler is initialized at the prior means.

1.4.2 Full sample results: loadings and volatilities

A first evaluation of the relative importance of the indicators that are included in the model is given by the full sample posterior estimates of the factor loadings (β), which are shown in Table 1.2. The highest posterior median weight (0.49) is given to the Industrial production index, followed by GDP (0.38) and by the Industrial production index in the Pulp and Paper sector. Survey data receive roughly the same weight (around 0.1), with a slight prevalence given to the PMI and the weakest contribution coming from the Michigan US Consumer Survey. The annual rate of change of the euro-dollar exchange rate and the US spread have a counter-cyclical effect on GDP. The sign of these two parameters is easily rationalized by considering that these indicators typically *lead* the business cycle, so that their correlation with current cyclical conditions (measured by the common factor) is negative.[7]

[7] An alternative way to look at the relative contribution of the indicators to the unobserved factor is through the Kalman filter weights derived in Koopman and Harvey (2003). The results

To see whether the model picks up any significant time variation in the variances of the common and idiosyncratic errors we plot the posterior median of selected members of Q_t together with their 68% confidence bands (Figure 1.1). Starting from the common factor (which can be seen as a measure of the underlying business cycle) the model identifies two shifts in volatility over the past twenty years. The former is a temporary increase at the beginning of the past decade, roughly around the brief recession experienced by the world economy in 2001. The latter, more persistent, starts between 2007/2008, and peaks in 2008, during the recent Great Recession. We next look at the hard indicators that receive the largest weights in the estimation of the common factor (GDP and IP). Visual inspection of the variances of the idiosyncratic shocks to these two indicators reveals that volatility has been rather stable over most of the sample, with the exception of the latest recession, when it surged significantly until 2008 to fall thereafter. Finally, the variance of the US spread shows a slight upward trend during the Nineties and a much more persistent increase during the 2007/2009 recession, consistently with the financial origins of the recent economic downturn.

1.4.3 News and forecasts 1

Given the mixed-frequency nature of our model, GDP forecasts are continuously updated as new monthly data become available. The impact of data releases on forecast revisions can be assessed using the methodology developed by GRS. To clarify the spirit of the exercise, the concept of vintage needs to be introduced. The Ω_{v_j} vintage is defined as:

$$\Omega_{v_j} = \{X_{it/v_j}; t = 1, \ldots, T_{iv_j}, i = 1, \ldots, n\} \tag{1.15}$$

that is the information set Ω_{v_j} is composed of n indicators available from month 1 to month T_{iv_j}, where the date for which the last observation is available varies across indicators. Within our model, a GDP forecast is obtained as an expectation of future GDP conditional on this information set.

Now consider a new vintage $\Omega_{v_{j+1}}$, which differs from the previous one for the release of a new observation of the i^{th} indicator:

$$\Omega_{v_{j+1}} - \Omega_{v_j} = X_{it/v_{j+1}} \tag{1.16}$$

confirm the findings in Banbura and Runstler (2011) indicating that in the months of the quarter when the dataset is “balanced", over half of the estimate of (unobserved) real activity growth depends on hard indicators, with soft indicators playing a role in the months when neither GDP nor Industrial Production are available.

The updated information entails a change in the conditioning set and, consequently, a forecast revision. Notice that we work with final data vintages in a pseudo real time context, that is, we do not consider data revisions but only new end of sample releases. This means that, starting from a given point in time, we let the information set gradually *expand*, one indicator at the time.

Data releases can occur at different intervals within the month but, for simplicity, we set up a stylized calendar in which the order of release of the various indicators is kept fixed within the month, see Table 1.3. From a given point in time we start enlarging our dataset by including new data on Industrial Production, typically published around the middle of each month. In the second month of each quarter, right after Industrial Production data are made available, GDP data are included in the information set. From the third week onwards survey data start being published by various sources. Surveys cannot be clearly ranked in terms of timeliness, since their release dates sometimes cross each other. We use the convention to place the IFO index release first, followed by the PMI, the Economic Sentiment Indicator and the US Michigan Consumer index. Finally we include exchange and interest rates, which enter the model as monthly averages of daily data.[8]

GRS evaluate how efficiently their large factor model incorporates data *news* in terms of Mean Squared Errors (MSE) reduction. Indeed, since successive vintages carry more information, one can reasonably expect to see a systematic fall in the forecast error variance as indicators are updated. Exploiting the Bayesian nature of our model we add two dimensions to this metrics. First, we look at the width of the forecast distribution at different horizons and investigate whether it shrinks as the information set expands. In a way this gives us some indication to whether the model forecast gains *confidence* as new information accrues and the forecast horizon decreases.

Second, we move beyond point forecast accuracy and evaluate the evolution of density forecast accuracy. To this end we use the log-score, that is the logarithm of the predictive density generated by the model evaluated at the outturn of the series.

We consider releases from January 2006 to May 2011 and forecast each quarter from the first month of the quarter to the first month of the subsequent one, that is we compute three nowcasts and one backast. For each month we update the vintages sequentially according to our stylized calendar, sample 1000 draws from the posterior, run the Kalman filter and smoother and, for each posterior draw, produce nine GDP estimates, corresponding to the release of each of the nine indicators, and consequently nine forecast errors. We compare our model forecasts with those of a

[8]This timing convention, which is the same used by GRS, somewhat penalizes financial variables as daily information on the dollar-euro and on the spread are disregarded.

naive constant growth model.[9]

In Figure 1.2 we show the evolution of the MSE within the month, relative to the MSE obtained with the naive model. In the first month the MSE falls monotonically within the month, albeit at a very slow rate. From the first to the second month there is a discrete jump corresponding to the publication of the Industrial production index. In the second month a large fall in the MSE occurs at the publication of the GDP for the previous quarter and the impact of survey and financial indicators weakens. From the third month onwards only Industrial production provides some further refinement of the GDP estimate. Also notice that throughout the forecast cycle the MSE ratio remains below one, reflecting the valuable content of conjunctural indicators. We next assess the evolution of forecast *confidence* over the forecasting cycle, as measured by the standardized interquartile range, that is the difference between the 75^{th} and the 25^{th} percentiles standardized by the median.[10] The evolution of the interquartile range over the forecast cycle, shown in Figure 1.3, reveals a clear downward tendency in the dispersion of GDP estimates, indicating that the *confidence* that the model places on its GDP forecasts increases as conjunctural information accrues. Moreover, soft data play an important role in driving the reduction in forecast dispersion, especially at the very beginning of the forecast cycle when a strong fall in forecast uncertainty occurs as the first surveys become available.

Finally, in Figure 1.4 we show the evolution of the log-score (crossed line) together with the log-score obtained with the constant growth model (dotted line). Density forecast accuracy monotonically increases at the release of each new indicator, indicating that as the forecast horizon shortens the model assigns (ex ante) a progressively higher probability to the actual GDP releases.

1.4.4 News and forecasts 2

In a recent paper Banbura and Modugno (2010) derive an alternative way to map directly *news* into forecast revisions. They motivate this alternative measure of *news* by noticing that in factor models the forecast of the unobserved factors is a weighted average of present and past observable indicators, with weights endogenously assigned by the Kalman smoother. When the information set is

[9] GRS provide evidence that the precision of the signal increases within the month as new data are released in both an in-sample and an out of sample exercise. Due to computational constraints we provide evidence only on the in-sample effect of *news*.

[10] We choose the interquartile range since it has some desirable statistical properties, in particular it is a robust statistics (i.e. it is not affected by outliers) and in a symmetric distribution it equals the median absolute deviation.

enriched by a new release, the Kalman smoother incorporates the new information by revising the weights assigned to *all the available indicators* making it impossible to discern whether an improvement in forecast accuracy is due to the new release or to a revision of the weights assigned to other indicators. They therefore devise a way to dissect more precisely the contribution of each release to forecast revisions. Their method is of particular interest in cases when, instead of considering the release of a single indicator, a whole block of data is released and the contribution of the *news* content of each single indicator needs to be assessed.

Our setup, by providing a quantification of the the uncertainty surrounding the news content of a new data (or block of data) release, provides a more complete picture of the forecast revision implied by the intra-monthly information flow. We illustrate this point using as a case study the GDP forecast of the second quarter of 2010. We start nowcasting this GDP release in the first half of April, when the February Industrial Production numbers become available. We update our forecasts twice a month until the first half of August, right before the first GDP estimate is published. The first by-monthly update coincides with the release of a string of hard data, the second with the publication of survey and of the monthly averages of financial indicators. The resulting forecast updates are shown in Figure 1.5. The bars below the dotted line depict the contribution of the release (the *news*) of each new indicator computed according to Banbura and Modugno methodology.[11]

At the beginning of the forecast cycle (mid-April) the prediction of the model stands quite far from the final outcome, as the model envisages barely positive growth against a GDP growth outturn of around 1%. Between the end of April and the middle of May positive signals coming from the survey first, and from Industrial production and the release of GDP data for Q1 afterwards, push the forecast progressively upwards. In May a false signal sent by the release of survey data depresses again GDP growth expectations. From June onwards, positive news from both soft and hard data set the model forecasts on the right track and GDP predictions start fluctuating more or less around 1%.

To complement the analysis with a measure of uncertainty on both (1) the overall revision implied by the release of an entire data block and (2) the contribution of each indicators to such revision, at each by-monthly update of our information set we draw 1000 forecasts from the predictive density and map each of these forecasts onto the *news*. In Figure 1.6 we report estimated kernel densities of the overall revision to the forecast due to the release of 'hard' (upper panel) and 'soft' (lower panel) data between April and July. To show what the individual contributions

[11] In order to evaluate the direct effect of news it is assumed that model parameters are unchanged between vintages.

look like we report in Figure 1.7 similar densities for two selected indicators, namely Industrial production and the Economic Sentiment Indicator, which appear to be responsible for most of the revisions over the forecast cycle.

From the comparison of these distributions with the information provided in Figure 1.5, the importance of having a tool to identify the credibility of forecast updates emerges quite clearly. In the second half of April and May, for example, the model picks up first a strong upward, then a strong downward revision due to the release of survey data, which can be largely attributed to *news* in the Economic Sentiment Indicator. However, results in Figures 1.6 and 1.7 show that in both months the overall revisions and the contribution of the Economic Sentiment Indicator to such revisions are measured with considerable uncertainty, calling for some caution in the interpretation of these forecast updates. In June and in July, on the other hand, as monthly information accumulates and the forecast horizon shortens, the dispersion of estimated revisions and contributions shrinks considerably.

1.4.5 Out of sample forecasting performance

Last, we conduct a (pseudo) out of sample forecast exercise. The design of the exercise is similar in spirit to the sequence of forecasts updates discussed in the previous section. For each quarterly GDP release we provide sixteen forecasts, starting from six months before the end the quarter of interest to one month afterwards (backcast). Taking as a target, for example, the third quarter of each year, we produce the first forecast in March and the last one in October. We update each of these projections twice a month, when, respectively, hard and soft data are released. The forecast exercise runs from the first quarter of 2006 to the last quarter of 2010. To appraise the contribution of the stochastic volatilities we contrast the forecasting performance of the proposed model with that of a restricted version in which the variance of the disturbances is constant.

Starting from point forecast evaluation, the evolution of the Root Mean Squared Forecast Errors (RMSFE) over the forecast cycle of the two specifications are shown in Figure 1.8. RMSFE are reported as a ratio to the RMSFE attained by a constant growth benchmark. Two observations are in order. First, time variation in the variances increases forecast accuracy, since the model with stochastic volatility has lower relative RMSFE over most of the forecast horizon, from the beginning to the first two nowcasts. From the end-month update of the second nowcast to the backcast, when more recent industrial production figures are released, the two models deliver instead broadly similar results. Second, the RMSFE of both models decline as the flow of information accumulates, yet the model with stochastic volatility exploits more efficiently early data releases and starts outperforming the naive benchmark (i.e. its RMSFE falls below 1) earlier in the forecast cycle.

Turning to density forecast evaluation, we look first at coverage rates, that is the frequency with which the actual outcome falls within a given confidence interval. If the model produces accurate density forecast actual GDP growth should fall 10% of the times within our 10% confidence interval, 20% of the times within our 20% confidence interval and so forth. Uncertainty can be measured through a t-test on the null hypothesis that the actual coverage equals the nominal one.[12] We look at backcast (projections one month after the end of the quarter), nowcast (projections during the quarter), and 1 step ahead forecast (projections for the next quarter). In Table 1.4 we report the coverage rate for the model without stochastic volatility. In some cases this specification produces far too wide confidence intervals. This is especially true in the case of the nowcast, when the test frequently rejects. Table 1.5 shows that adding stochastic volatility yields gains in density forecast accuracy as confidence intervals are usually well calibrated, especially when nowcasting.

Finally, we examine the normalized probability integral transforms (PITS) of the forecast errors. According to the testing framework developed by Berkowitz (2001), if the model forecast density matches the density that generated the data, the PITS should be independent standard normal. We follow Clark (2011) and test these conditions (zero mean, unit variance and no serial correlation) separately and jointly.

The p-values of the tests are presented in Table 1.6, in the top panel for the constant volatility model and in the bottom one for the model with stochastic volatility.[13] For each month we consider a mid-month and an end-month update, so that we have six different results for the nowcast and two results for the backast. Despite a few rejections of the individual tests, for both models the joint Normality/Independence hypothesis cannot generally be rejected for the nowcast. For the backast, again, the inclusion of stochastic volatility induces some gains, as the PITS do not display any serial correlation and the hypothesis that they are jointly Normally and Independently distributed cannot be refuted.

[12] As emphasized by Clark (2011) this test is slightly imprecise as it abstracts from parameter uncertainty

[13] The the testing framework developed by Berkowitz (2001) for the PITS applies to one step ahead forecast. We therefore consider only backast and nowcast, since forecasting into the next quarter actually implies two steps ahead predictions of GDP.

1.5 Robustness checks

We check the robustness of the results to some of our modeling choices. We start from assessing how sensitive the predictive accuracy of the model is to the choice of the parameters of the prior distributions. Since the priors on the factor loadings and on the AR coefficients are virtually flat, we concentrate on those of the log-volatilities. We experiment with two alternative prior scale parameters obtained by multiplying the baseline scale parameter by, respectively, 4 and $\frac{1}{4}$. This should give us a rough indication of whether further tightening or loosening this prior changes the results. We repeat the out of sample forecast exercise and report in the top panel of Figure 1.9 the resulting RMSFE (labeled *scale1* and *scale2*, respectively), together with the RMSFE of the baseline specification. The plot suggests that variations in the prior scales does not result in substantial changes in the model performance. Similar experiments conducted by altering the degrees of freedom (not reported for the sake of brevity) confirm these results.

As a second robustness check we modify the law of motion of the the stochastic volatilities by letting them follow an AR(1) process:

$$\lambda_{i,t} = \alpha_i \lambda_{i,t-1} + \theta_{i,t}\sigma_{\lambda,i} \qquad \theta_{i,t} \sim N(0,1) \tag{1.17}$$

For the parameter α_i we assume a Normal prior and try two values for the prior mean, 0.5 and 0.9, while keeping the prior variance at 1000. Again, the RMSFE attained by these two specifications do not differ significantly from those of the model in which the stochastic volatilities follow a random walk, as shown in the bottom panel of Figure 1.9.

We have also experimented with a model featuring two factors like in Frale, Marcellino, Mazzi, and Proietti (2011). In our application, however, the second factor turned out not to be well identified, as all the estimated loadings were not different from zero.

1.6 Conclusions

This paper introduces a mixed frequency factor model with stochastic volatility, and develops a Bayesian procedure for its estimation. The model deals with all the challenges faced by a forecaster that needs to produce updated quarterly GDP forecasts at each relevant data release, like data sampled at different frequencies and ragged-edge data. Differently from existing linear models, our setup allows for random shifts in the volatility of the errors.

DYNAMIC FACTOR MODEL, STOCHASTIC VOLATILITY

The method is applied to the problem of forecasting euro area GDP. The Bayesian setup allows an assessment of the uncertainty around the news content of monthly releases of hard, soft and financial indicators. Consistently with findings in the literature, we find that forecast accuracy improves significantly in connection with the release of monthly data as the forecast horizon decreases. Also, forecast uncertainty (measured by the width of the forecast distribution) progressively decreases as more information on the quarter of interest becomes available. Out-of-sample evidence indicates that the introduction of stochastic volatility contributes to higher point forecast accuracy and tends to yield more precise confidence intervals.

Table 1.1: Variable selection summary

Indicator	Country
GDP	Euro Area
Industrial Production	Euro Area
Industrial Production - Pulp/paper	Euro Area
Business Climate - IFO	Germany
Economic Sentiment Indicator	Euro Area
PMI composite	Euro Area
dollar-euro	US-Euro
10y-3m spread	US
Michigan Consumer Sentiment	US

Table 1.2: Factor Loadings - posterior estimates

Percentiles	25th	50th	75th
GDP	0.27	0.38	0.54
IP	0.40	0.49	0.60
IP-PULP	0.23	0.29	0.36
IFO	0.10	0.12	0.13
ESI	0.10	0.12	0.14
PMI	0.12	0.13	0.15
dollar-euro	-0.08	-0.05	-0.02
US-spread	-0.06	-0.04	-0.02
Michigan Consumer	0.04	0.06	0.08

Table 1.3: Stylized data release calendar

Indicator	Timing	Publication lag	Frequency
IP	$11^{th} - 15^{th}$ of month	2	Monthly
IP-PULP	$11^{th} - 15^{th}$ of month	2	Monthly
GDP	1 day after IP	2	Quarterly
IFO	$20^{th} - 30^{th}$ of month	0	Monthly
PMI	$20^{th} - 30^{th}$ of month	0	Monthly
ESI	$20^{th} - 30^{th}$ of month	0	Monthly
Michigan Consumer	Last Friday of the month	0	Monthly
dollar-euro	Last day of month(Monthly ave.)	0	Monthly
US-spread	Last day of month(Monthly ave.)	0	Monthly

Table 1.4: Coverage Rates - Model without Stochastic Volatility

Nom Cov	Backcast		Nowcast		1 step ahead	
	Coverage	P-value	Coverage	P-value	Coverage	P-value
0.1	0.25	**0.03**	0.15	0.10	0.14	0.15
0.2	0.32	0.10	0.30	**0.01**	0.23	0.46
0.3	0.36	0.39	0.39	**0.03**	0.36	0.18
0.4	0.48	0.32	0.48	**0.05**	0.48	**0.08**
0.5	0.57	0.37	0.58	**0.06**	0.55	0.22
0.6	0.68	0.26	0.67	0.11	0.62	0.62
0.7	0.75	0.45	0.77	**0.08**	0.70	0.94
0.8	0.84	0.47	0.84	0.20	0.73	**0.06**
0.9	0.89	0.78	0.88	0.46	0.76	**0.00**

Table 1.5: Coverage Rates - Model with Stochastic Volatility

Nom Cov	Backcast		Nowcast		1 step ahead	
	Coverage	P-value	Coverage	P-value	Coverage	P-value
0.1	0.05	**0.09**	0.12	0.55	0.09	0.57
0.2	0.25	0.45	0.23	0.48	0.24	0.27
0.3	0.39	0.25	0.34	0.30	0.36	0.17
0.4	0.45	0.48	0.41	0.89	0.45	0.31
0.5	0.48	0.77	0.46	0.38	0.54	0.38
0.6	0.61	0.86	0.57	0.50	0.56	0.40
0.7	0.77	0.26	0.67	0.50	0.69	0.76
0.8	0.86	0.23	0.79	0.76	0.73	**0.07**
0.9	0.93	0.41	0.88	0.40	0.76	**0.00**

Note to Tables 1.4 to 1.5. The table shows p-values for the test of the hypothesis that Nominal and estimated Coverage Probabilites are equal.

Table 1.6: Density forecast evaluation using the *pits*

Model without Stochastic Volatility

	Backcast		Nowcast					
	1	2	1	2	3	4	5	6
Mean	**0.01**	**0.01**	**0.03**	**0.02**	**0.03**	**0.02**	**0.05**	**0.00**
Variance	**0.02**	**0.02**	0.49	0.18	0.12	0.49	0.24	**0.01**
AR(1)	0.10	0.13	0.39	0.45	0.11	0.54	0.60	0.14
Joint	**0.03**	**0.05**	0.15	0.13	0.19	0.23	0.28	**0.04**

Model with Stochastic Volatility

	Backcast		Nowcast					
	1	2	1	2	3	4	5	6
Mean	**0.02**	**0.05**	**0.04**	0.11	0.22	**0.06**	**0.07**	**0.03**
Variance	**0.09**	**0.00**	0.84	0.70	0.76	0.71	0.99	**0.03**
AR(1)	0.43	0.48	0.13	0.15	0.81	0.70	0.48	0.52
Joint	0.25	0.26	**0.07**	0.21	0.67	0.22	0.31	0.17

Note to Table 1.6. The table displays p-values for the test of the hypotheses of zero mean, unit variance, no serial correlation and joint Normality/Indipendence of forecast errors at different horizons.

Figure 1.1: Stochastic volatility for the common factor and for selected variables

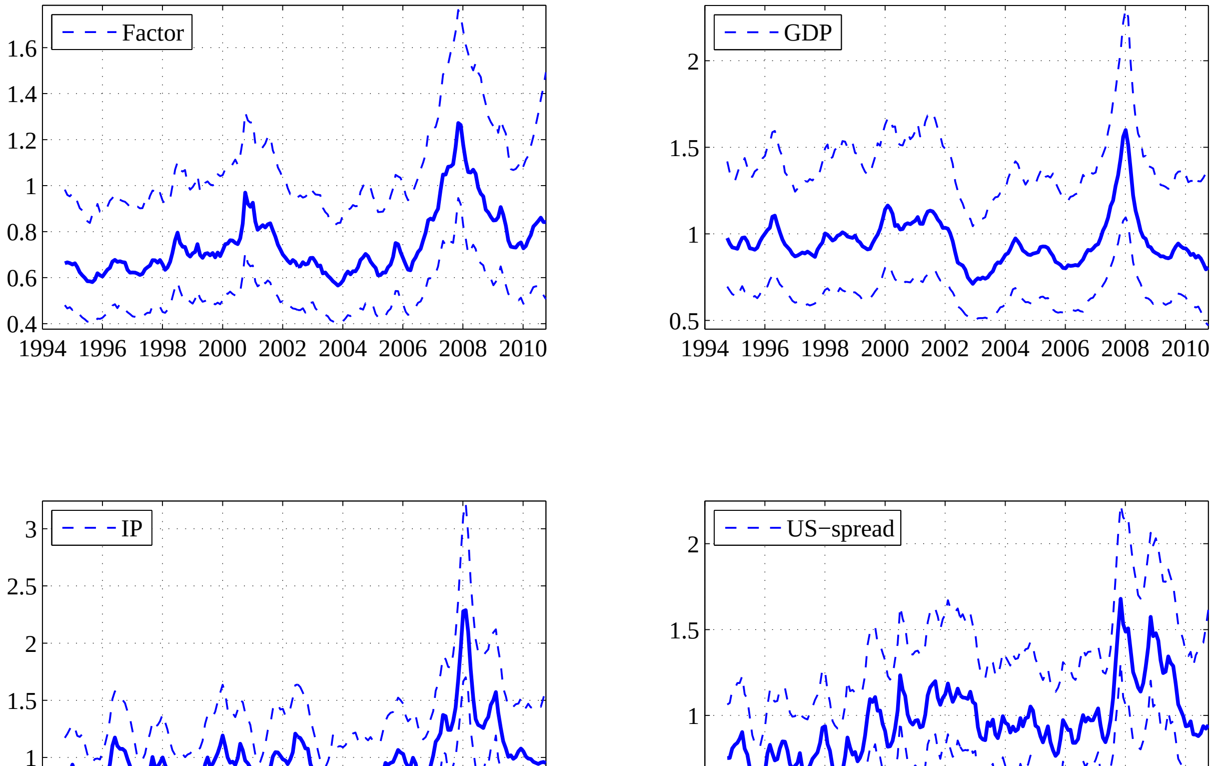

Figure 1.2: RMSE at different releases

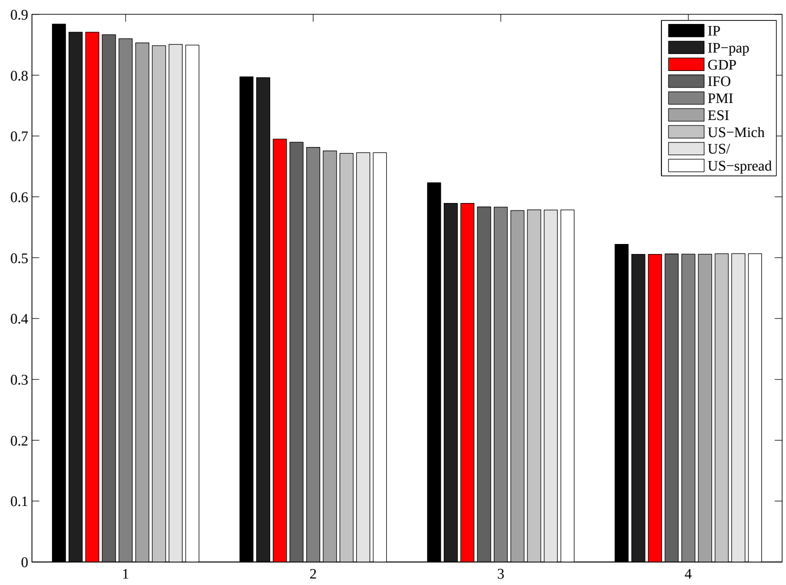

Note to Figure 1.2: the Figure shows the ratio of the RMSE of the factor model with stochastic volatility to that of a naive constant growth model for each of the indicated data release. Data releases follow the stylized calendar 1.3.

Figure 1.3: Forecast dispersion at different releases

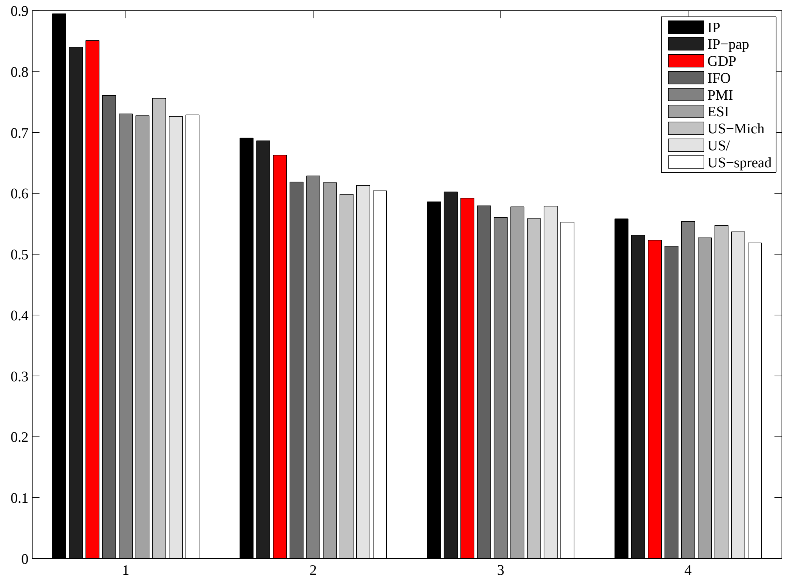

Note to Figure 1.3: the Figure shows the difference between the 75 and the 25 percentiles (both scaled by the median) of the forecast distribution obtained with the factor model with stochastic volatility updated at each data release. Data releases follow the stylized calendar 1.3.

Figure 1.4: Log-predictive score at different releases

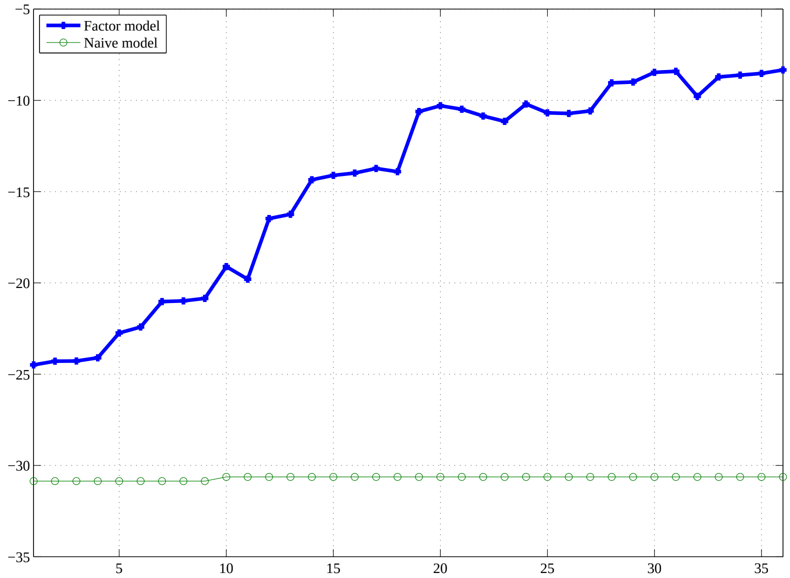

Note to Figure 1.4: the Figure shows the log-predictive score of the factor model with stochastic volatility updated at each data release and of the naive constant growth model. Data releases follow the stylized calendar 1.3.

Figure 1.5: Forecast revisions 2010Q2

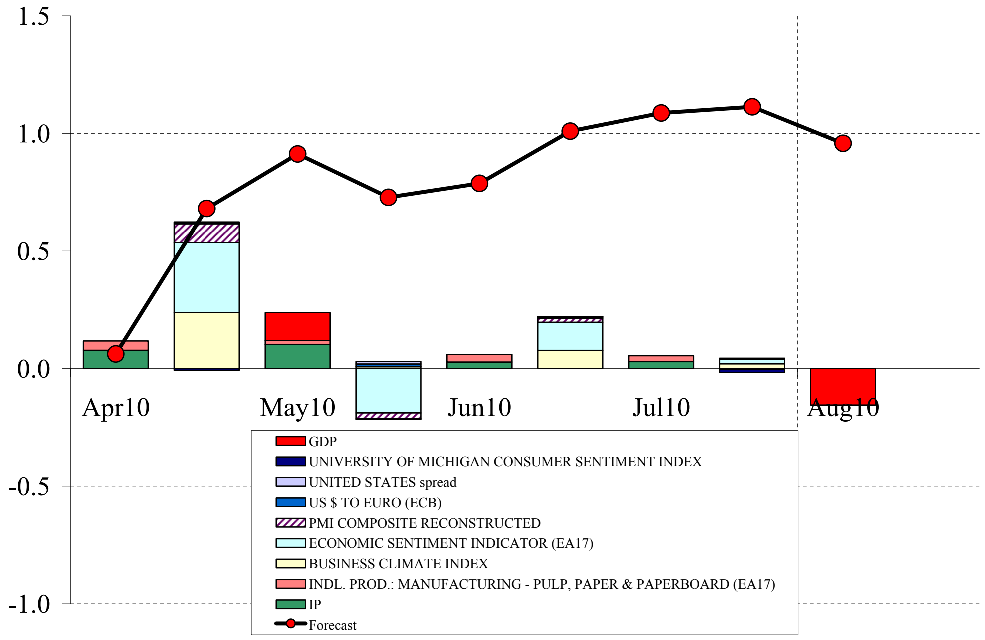

Note to Figure 1.5: the Figure shows the by-monthly GDP forecasts revisions relative to the second quarter of 2010 and the contributions of the new releases. The firs forecast update is at the end of April, the last update in the middle of August.

Figure 1.6: Revisions Density evolution 2010Q2

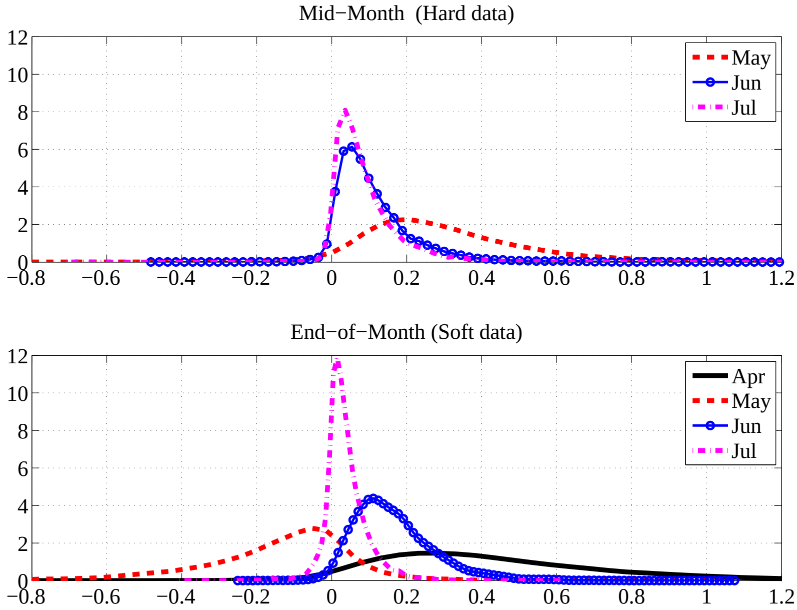

Figure 1.7: Revisions Density evolution 2010Q2, contribution of selected indicators

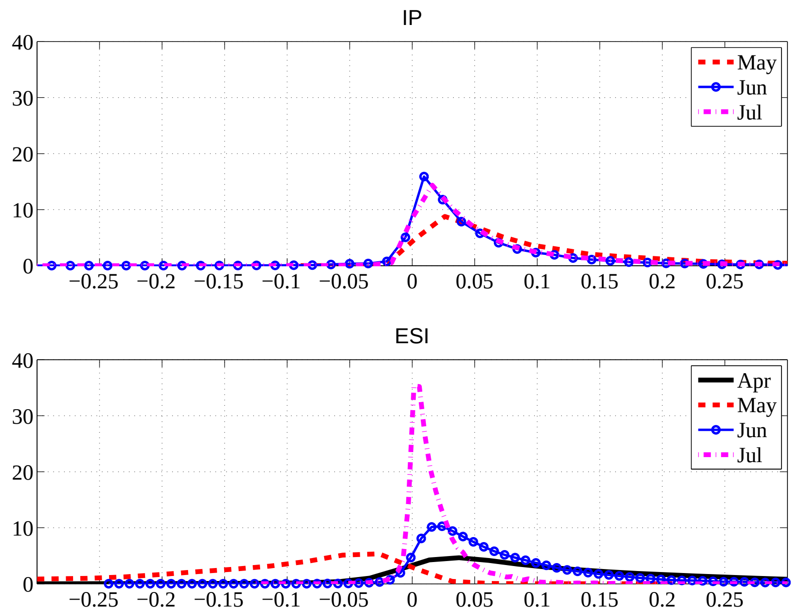

Note to Figures 1.6 and 1.7: The density estimation is based on a normal kernel function, using an optimal window parameter function of number of data points. The distribution is based on 1000 draws from the predictive density.

Figure 1.8: RMSE - baseline specification

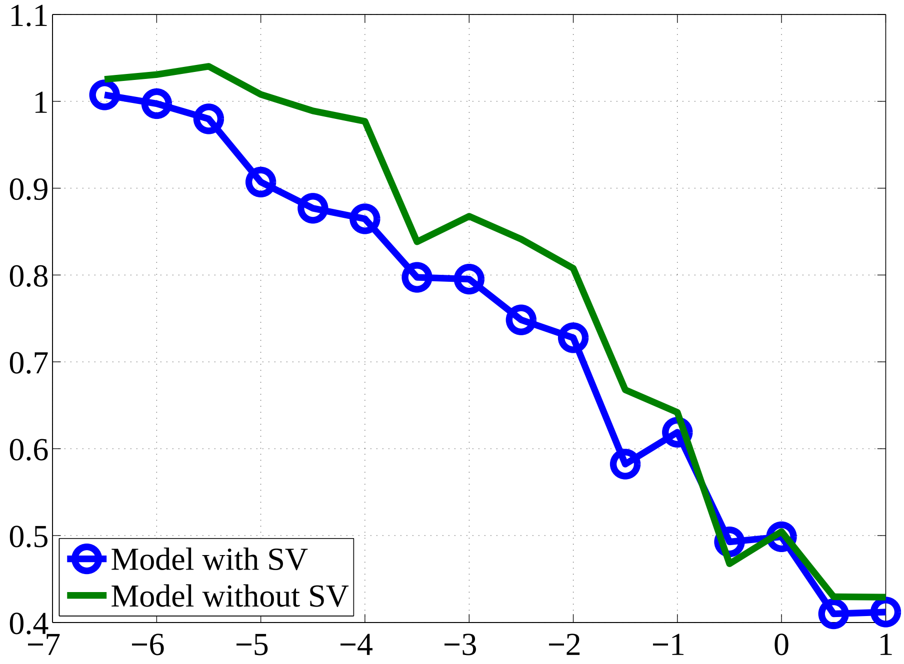

Figure 1.9: RMSE - alternative specifications

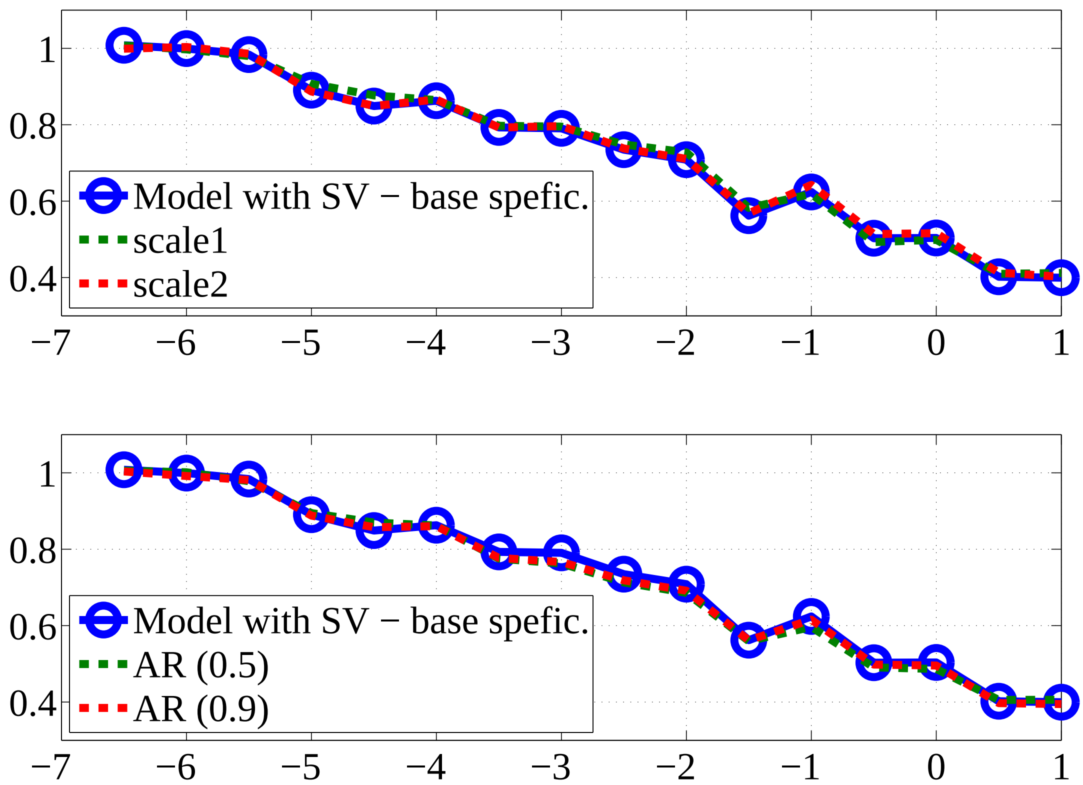

Note to Figures 1.8 and 1.9. The Figure shows the RMSFE obtained between the first quarter of 2006 and the last quarter of 2010. The forecast horizon goes from six months ahead to one month after the end of the quarter of interest (backast).

Appendix

1.A Details of the Gibbs sampler

We describe in more details the six blocks that compose our Gibbs sampler procedure. The sampler is based on the algorithm by Del Negro and Otrok (2008) modified to account for missing data and mixed frequencies.

1.A.1 Block 1: drawing the factor loadings $\beta_q, \beta_h, \beta_s$

In the first block of the Gibbs sampler we draw the factor loadings. Start from the measurement equation of the hard indicator:

$$y_{h,t} = \beta_h f_t + u_{h,t} \tag{1.18}$$

where the law of motion of the idiosyncratic shock is $u_{h,t} = \phi_{h,1} u_{h,t-1} + \phi_{h,2} u_{h,t-2} + \epsilon_{h,t} e^{\lambda_{h,t}/2}$ and $\epsilon_{h,t} \sim N(0, \sigma_h)$. Since we are conditioning on all the parameters, on the factor f_t and on the stochastic volatilities $\lambda_{h,t}$ we treat this as a regression with autocorrelated and heteroscedastic residuals. Now we quasi-difference the equation by filtering both sides with the filter $1 - \phi_{h,1} L - \phi_{h,1} L^2$ and divide each observation by $e^{\lambda_{h,t}/2}$:

$$y^\star_{h,t} = \beta_h x^\star_t + \epsilon_{h,t} \tag{1.19}$$

where $x^\star_t = (1 - \phi_{h,1} L - \phi_{h,1} L^2) f_t / e^{\lambda_{h,t}/2}$. We posit a Normal prior so thath the posterior is also Normal, see Kim and Nelson (1999) for a textbook treatment.

The case of survey variables can be treated in the same way after noticing that $x^\star_t = (1 - \phi_{h,1} L - \phi_{h,1} L^2) \sum_{j=0}^{11} f_{t-j} / e^{\lambda_{s,t}/2}$.

In the case of quarterly variables two adjustments are needed. First, since the variable is observed only every three months only these observations can be used for estimating the factor loading. Second, in the measurement equation an MA(4) regression error appears:

$$y_{q,t} = \beta_q w(L) f_t + w(L) u_{q,t} \tag{1.20}$$

where $w(L) = \frac{1}{3} + \frac{2}{3}L + L^2 + \frac{2}{3}L^3 + \frac{1}{3}L^4$. Furthermore the error term u_t is an AR(2) process $u_{q,t} = \phi_{q,1}u_{q,t-1} + \phi_{q,2}u_{q,t-2} + \epsilon_{q,t}e^{\lambda_{q,t}/2}$. We work out the variance covariance matrix of the error terms of equation (1.20), $\Phi(\phi_{q,1}, \phi_{q,2}, \sigma_q^2)$, which, at this step of the sampler, can be treated as known. Then we divide each observation by $e^{\lambda_{q,t}/2}$ and pre-multiply both sides of the equation by $\Phi^{-\frac{1}{2}}$ to obtain a standard regression with uncorrelated residuals. Assuming a normal prior, draws of β_q are obtained from a normal posterior.

1.A.2 Block 2: drawing $\phi_{f,1}, \phi_{f,2}, \phi_{q,1}, \phi_{q,2}, \phi_{h,1}, \phi_{h,2}, \phi_{s,1}, \phi_{s,2}$

To draw the AR parameters of the idiosyncratic shocks notice that, conditioning on the state vector μ_t, we can treat the common factor f_t and the residuals $u_{q,t}, u_{h,t}, u_{s,t}$ as known. The transition equations become standard regression problems which can be analyzed using the same steps used for drawing the factor loadings. We employ normal priors and rule out explosive roots by discarding draws if the roots of $\phi_j(L) = 0$ lie outside the unit circle.

1.A.3 Block 3: drawing the innovation variances $\sigma_f^2, \sigma_q^2, \sigma_h^2, \sigma_s^2$

We again proceed by treating the transition equations one at the time. Consider a generic element of the state vector $\mu_{i,t}$. Its law of motion is:

$$\mu_{i,t} = \phi_{i,1}\mu_{i,t-1} + \phi_{i,2}\mu_{i,t-2} + \eta_{i,t} \quad \eta_{i,t} \sim N(0, \sigma_i^2 e^{\lambda_{i,t}}) \tag{1.21}$$

For the innovation variance σ_i^2 we posit an inverse-Gamma prior $p(\sigma_i^2) = IG(n_i, s_i^2)$. Since the prior is conjugate it can be interpreted as adding n_i artificial observations to the state variable $\mu_{i,t}$. The prior embodies the belief that the sum of squared residuals of these artificial observations equals s_i^2:

$$s_i^2 = \frac{1}{n_i}\sum_{t=1}^{n_i}(\mu_{t,i}^\star - \phi_{i,1}\mu_{t-1,i}^\star - \phi_{i,2}\mu_{t-2,i}^\star)^2 \tag{1.22}$$

Given our assumption that the idiosyncratic disturbances are normal the posterior is also an inverse-Gamma, $IG(T + n_i, \frac{n_i s_i^2 + T d_i^2}{T + n_i})$ where:

$$d_i^2 = \frac{1}{n_i} \sum_{t=1}^{n_i} (\mu_{t,i} - \phi_{i,1}\mu_{t-1,i} - \phi_{i,2}\mu_{t-2,i})^2 \tag{1.23}$$

The weight of the prior is therefore proportional to the prior degree of freedom parameter n_i.

1.A.4 Block 4: drawing the state vector μ_t

Since the model can be cast in state space draws of the state vector can obtained via a state vector simulation smoother as in Carter and Kohn (1994) or with the disturbance smoother proposed by Koopman and Durbin (2003). We resort to the latter, which turns out to be slightly more efficient from a computational point of view.

1.A.5 Block 5: drawing $\lambda_{i,t}$

To sample the stochastic volatilities $\lambda_{i,t}$ notice that conditional on all parameters and on the states μ_t the orthogonal innovations $\eta_{i,t}/\sigma_{h,i}$ are observable. The $\lambda_{i,t}$ can then be sampled adopting the date-by-date blocking scheme developed by Jacquier, Polson, and Rossi (1994).[14].

1.A.6 Block 6: drawing $\sigma^2_{h,i}$

The final block of the sampler involves drawing the variances of the log-volatilities. Conditioning on the log-volatilities and postulating an inverse-Gamma prior distribution, the $\sigma^2_{h,i}$ can also be drawn from an inverse Gamma posterior.

1.B The selection of the monthly indicators

Small scale models have their own “curse of dimensionality": since they rely on a small set of indicators, they are prone to the criticism of potentially leaving out relevant information compared to factor models that use hundreds of time series. Part of the literature has, however, advocated the use of a models of small dimensions. Bai and Ng (2008) and Boivin and Ng (2006), for example, question the usefulness of ’too much information’ for forecasting purposes, showing that a number of variable selection techniques (already widely used in biomedical statistics

[14]Details on the algorithm, which involves a Metropolis Hastings step within the Gibbs sampler, can be found in Cogley and Sargent (2005) Appendix $B.2.5$

where the number of covariates is typically very large) give encouraging results when applied to economic time series. To make the choice of the indicators to be included in our model as objective as possible we proceed as follows. We start by considering a dataset of more than a hundred variables for the period 1987-2011

We then set a priori four *core* variables that we decide to include in the model, which are Industrial Production for the euro area (IP), the composite Purchasing Manager Index (PMI), the European Commission Economic Sentiment Indicator (ESI) and the Germany IFO Business Climate Index. To select the remaining variables, we calculate as a benchmark the percentage of GDP variance explained by the factor computed from the core variables only, as in Camacho and Perez-Quiros (2010), and design an algorithm for the selection of a set of additional indicators which maximize this statistic.

1. We evaluate datasets with all *core* variables and one other variable at a time in order to calculate the explained variance, and the probability that it is higher than in the dataset with *core* variables only. In this way we obtain a ranking of the other series.

2. We add a variable at a time, starting with the ones with an higher probability to increase the explained variance with respect to the benchmark; we keep the variable only if this probability increases. We end up with the small set of 8 variables described in the main text.

The specification we adopt follows Camacho and Perez-Quiros (2010) where surveys are modeled as a 12 terms moving average of the unobserved factor, while hard variables load the factor contemporaneously. This amounts to imposing that surveys are in phase with the year on year growth rate of Industrial Production (and of the other hard indicators).

Toget an idea of the state representation of the model while keeping notation to a minimum we present the case of a toy model with one quarterly variable, one hard indicator and one soft indicator in which all the idiosyncratic shocks follow an AR(2) process. The more general case can be easily derived from this example. The loading matrix F in the measurement equation (1.10) can be written as:

$$F = \begin{pmatrix} \beta_q \frac{1}{3} & \beta_q \frac{2}{3} & \beta_q & \beta_q \frac{2}{3} & \beta_q \frac{1}{3} & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \frac{1}{3} & \frac{2}{3} & 1 & \frac{2}{3} & \frac{1}{3} & 0 & 0 & 0 & 0 \\ \beta_h & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ \beta_s & \beta_s & \beta_s & \beta_s & \beta_s & \beta_s & \beta_s & \beta_s & \beta_s & \beta_s & \beta_s & \beta_s & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ \end{pmatrix}$$

where β_q, β_h and β_s are the loadings of, respectively, the quarterly variable, the hard and the soft indicators. The state vector is:

$$\mu_t = \begin{pmatrix} f_t & f_{t-1} & \dots & f_{t-11} & u_{q,t} & \dots & u_{q,t-4} & u_{h,t} & u_{h,t-1} & u_{s,t} & u_{s,t-1} \end{pmatrix}' \quad (1.24)$$

The transition matrix is:

$$H = \begin{pmatrix} \phi_{f,1} & \phi_{f,2} & 0 & 0 & 0 & \dots & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & \dots & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & \dots & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & 0 & 0 & \dots & \phi_{q,1} & \phi_{q,2} & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & \dots & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & \dots & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & 0 & 0 & \dots & 0 & 0 & 0 & 0 & 0 & \phi_{h,1} & \phi_{h,2} & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & \dots & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & \dots & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \phi_{s,1} & \phi_{s,2} \\ 0 & 0 & 0 & 0 & 0 & \dots & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \end{pmatrix} \quad (1.25)$$

Since the idiosyncratic shocks are collected in the state vector the matrix R_t is a (k+2) dimension zero matrix while the matrix Q_t is a diagonal matrix which collects all the variances:

$$Q_t = diag \begin{pmatrix} 1 & 0 & 0 & 0 & \dots & \sigma_q^2 e^{\lambda_{q,t}} & 0 & 0 & 0 & 0 & \sigma_h^2 e^{\lambda_{h,t}} & 0 & \sigma_s^2 e^{\lambda_{s,t}} & 0 \end{pmatrix} \quad (1.26)$$

1.B.1 The model with two factors

As a robustness check we have extend the baseline model to include a second factor, which we model as a (restricted) ARMA(2,2) process as in Frale, Marcellino, Mazzi, and Proietti (2011). The two monthly unobserved factors have the following reduced form representations:

$$\begin{aligned} (1 - \varphi_{11} L - \varphi_{12} L^2) f_{1,t} &= \varepsilon_{1t} \\ (1 - \varphi_{21} L - \varphi_{22} L^2) f_{2,t} &= (1 - \theta L)^2 \varepsilon_{2t} \end{aligned}$$

where $\varepsilon_{1t} \sim N(0, \sigma_{1\varepsilon})$ and $\varepsilon_{2t} \sim N(0, \sigma_{2\varepsilon})$. Frale et al. (2009) set $\theta = 0.5$, motivating such restriction as a way to enhance the fit at low frequencies, see also Morton and

Tunncliffe-Wilson (2004). We sketch the State Space representation of the modified model in a simple setup with three indicators (GDP, a hard monthly variable and a soft monthly variable). Since the presence of these extra MA terms produces a smoother factor we drop the 12 terms moving average representation of the loadings of the surveys in favour of a more standard contemporaneous relationship, so that the measurement matrix F is composed of the three following blocks:

$$F_1 = \begin{pmatrix} \beta_{q,1}\frac{1}{3} & \beta_{q,1}\frac{2}{3} & \beta_{q,1} & \beta_{q,1}\frac{2}{3} & \beta_{q,1}\frac{1}{3} \\ \beta_{h,1} & 0 & 0 & 0 & 0 \\ \beta_{s,1} & 0 & 0 & 0 & 0 \end{pmatrix}$$

$$F_2 = \begin{pmatrix} \beta_{q,2}\frac{1}{3} & \beta_{q,2}\frac{2}{3} & \beta_{q,2} & \beta_{q,2}\frac{2}{3} & \beta_{q,2}\frac{1}{3} & 0 & 0 \\ \beta_{h,2} & 0 & 0 & 0 & 0 & 0 & 0 \\ \beta_{s,2} & 0 & 0 & 0 & 0 & 0 & 0 \end{pmatrix}$$

$$F_3 = \begin{pmatrix} \frac{1}{3} & \frac{2}{3} & 1 & \frac{2}{3} & \frac{1}{3} & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \end{pmatrix}$$

where F_1, and F_2 collect the loadings on the first and second factor and F_3 the loadings on the idiosyncratic disturbances, which are also modelled as $AR(2)$ processes. Notice that F_2 has two extra column vectors of zeros, necessary to accommodate the MA terms in the second factor.

$$F = [F_1, F_2, F_3]$$

We adopt the $max(p, q+1)$ representation, see Durbin and Koopman (2006), which requires the slightly more general specification of the transition equations:

$$\mu_t = T\mu_{t-1} + R\eta_t$$

where $\eta_t \sim (0, Q_t)$, Q_t is a 5 dimensional diagonal matrix (collecting the variances of the 5 idiosyncratic terms, 3 for the observable indicators, 2 for the unobserved factors) and R is a 21×5 selection matrix. The matrix T is block diagonal with the following blocks:

$$T_1 = \begin{bmatrix} \varphi_{11} & \varphi_{12} & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \end{bmatrix}, T_2 = \begin{bmatrix} \varphi_{21} & 0 & 0 & 0 & 0 & 1 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ \varphi_{22} & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix},$$

$$T_3 = \begin{bmatrix} \varphi_{q,1} & \varphi_{q,2} & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \end{bmatrix}, T_4 = \begin{bmatrix} \varphi_{h,1} & \varphi_{h,2} \\ 1 & 0 \end{bmatrix}, T_5 = \begin{bmatrix} \varphi_{s,1} & \varphi_{s,2} \\ 1 & 0 \end{bmatrix}$$

the state vector is 21 dimensional:

$$\mu_t = [f_{1t}, f_{1t-1}, ..., f_{1t-4}, f_{2t}, f_{2t-1}, ...f_{2t-4}, z_{1,t}, z_{2,t}, u_{qt}, u_{qt-1}, ..., u_{qt-4}, u_{ht}, u_{ht-1}, u_{st}, u_{st-1}]$$

and the selection matrix R is

$$R = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & -2\theta & 0 & 0 & 0 \\ 0 & \theta^2 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 \end{bmatrix}$$

It can be seen that the two additional state variables z_t^1 and z_t^2 are:

$$\begin{aligned} z_{1,t} &= \varphi_{22} f_{2,t-1} + z_{2,t-1} - 2\theta\varepsilon_t \\ z_{2,t} &= \theta^2 \varepsilon_t \end{aligned}$$

Finally, $Q_t = (1, \sigma_{f_2}^2 e^{\lambda_{f_2,t}}, \sigma_q^2 e^{\lambda_{q,t}}, \sigma_h^2 e^{\lambda_{h,t}}, \sigma_s^2 e^{\lambda_{s,t}})$.

In our empirical application, however, this second factor is not well identified. In fact, as shown in Table 1.C.1, the loadings of the indicators on the second factor collapse to zero, with the first factor accounting for the co-movement in the data.

1.C News and forecast revisions

In their paper Banbura and Modugno (2010) derive a way to decompose a forecast revision as a linear function of news.

They denote as Ω_v a vintage of data corresponding to a statistical data release v, which as an example can be mid-month for industrial production and end of month for surveys, in order to define *news* as:

$$I_{v+1,j} = y_{i_j,t_j} - E[y_{i_j,t_j}|\Omega_v] \tag{1.27}$$

the surprise incorporated in a new data with respect to what was expected given information Ω_v. A forecast revision is defined as:

$$E[y_{k,t_k}|I_{v+1}] = E[y_{k,t_k}|\Omega_{v+1}] - E[y_{k,t_k}|\Omega_v] \tag{1.28}$$

and can be expressed as weighted average of news:

$$E[y_{k,t_k}|I_{v+1}] = B_{v+1}I_{v+1} = E[y_{k,t_k}I'_{v+1}]E[I_{v+1}I'_{v+1}]^{-1}I_{v+1} \tag{1.29}$$

where:

$$E[y_{k,t_k}I_{v+1,j}] = H_k E[(\mu_{t_k} - E(\mu_{t_k}|\Omega_v))(\mu_{t_j} - E(\mu_{t_j}|\Omega_v)')]H'_{i_j} \tag{1.30}$$

$$E[I_{v+1,j}I_{v+1,l}] = H_{i_j} E[(\mu_{t_j} - E(\mu_{t_j}|\Omega_v))(\mu_{t_l} - E(\mu_{t_l}|\Omega_v)')]H'_{i_l} \tag{1.31}$$

where $E[(\mu_{t_j} - E(\mu_{t_j}|\Omega_v))(\mu_{t_l} - E(\mu_{t_l}|\Omega_v)')]$ is the state vector covariance matrix obtained as a by-product of the Kalman Smoother.

Table 1.C.1: Factor loadings, two factor model

Percentiles	First factor			Second factor		
	75th	50th	25th	75th	50th	25th
GDP	0.44	0.27	0.08	0.00	0.00	0.00
Industrial Production	0.67	0.55	0.41	0.00	0.00	0.00
Industrial Production - Pulp/paper	0.34	0.28	0.21	0.00	0.00	0.00
Business Climate - IFO	0.04	0.02	0.00	0.00	0.00	0.00
Economic Sentiment Indicator	0.03	0.01	-0.01	0.00	0.00	0.00
PMI composite	0.02	0.01	-0.01	0.00	0.00	0.00
dollar-euro	0.01	-0.01	-0.04	0.00	0.00	0.00
10y-3m spread	0.03	0.01	0.00	0.00	0.00	0.00
Michigan Consumer Sentiment	0.09	0.05	0.02	0.00	0.00	0.00

Chapter 2

Large Time-Varying Parameter VARs: A Non-Parametric Approach

2.1 Introduction

In recent years macro-econometric research has been particularly active on two fronts. First, increasing availability of economic time series has prompted the development of methods capable of handling large dimensional datasets. Second a number of changes in the economic landscape (a renewed stream of oil price shocks, the Great Recession, unconventional monetary policy in most advanced countries) further stimulated work on models with time-varying parameters.

On the large models front typical solutions include data reduction and parameter shrinkage. Data reduction reduces the data space through linear combinations (factors) of the observed variables. This parsimonious representation of the data typically yields benefits in terms of estimation precision and forecasting. Shrinkage, on the other hand, constraints the parameter space within values that are (a priori) plausible. It therefore reduces estimation uncertainty, providing an alternative solution to the over-fitting problem. In the context of large Vector Autoregressions (VARs), for example, Banbura, Giannone, and Reichlin (2010) show that progressively tightening shrinkage as the cross-sectional dimension of the VAR increases, results in more accurate forecasts than those obtained on the basis of unrestricted VARs. Despite different premises, data reduction and shrinkage go in the same direction since, as shown by De Mol, Giannone, and Reichlin (2008), both methods stabilise OLS estimation by *regularising* the covariance matrix of the regressors.

Turning to time-varying parameters (TVP) models, a prolific line of research has grown in the Bayesian context, starting from the seminal work on VARs with time-varying coefficients and variances by Cogley and Sargent (2005) and Primiceri (2005). The estimation procedure of these models rests on the assumption that the VAR coefficients follow a random walk (or autoregressive) process. The assumed law of motion for the model parameters, coupled with the VAR equations, form a State Space system. Given the presence of time-varying second moments, a combination of Kalman filtering and Metropolis Hasting sampling is then used to deal with such models. The need to use the Kalman filter, however, limits the scale of the models, so that the numerous empirical applications that have followed this approach usually model a relatively small number of time series. Furthermore, in settings where the nature of the structural change is uncertain, methods based on simple data discounting could be more robust than Kalman filter based models.

These motivations are behind a stream of papers that in recent years have explored the performance of non parametric estimation methods for TVP-models. The viewpoint of this line of research is that the nature of time variation in the co-movement across time series is itself evolving, i.e. large infrequent breaks could coexist with periods of slow gradual time variation. Given this complexity, adaptive methods can deliver good forecasts and an accurate description of the structural relationships among macroeconomic variables at a relatively low computational cost.[1] In this framework Giraitis, Kapetanios, and Yates (2014) and Giraitis, Kapetanios, and Yates (2012) have developed non-parametric estimators for univariate and multivariate dynamic models. They show that, for a wide class of models in which the coefficients evolve stochastically over time, the path of the parameters can be consistently estimated by suitably discounting distant data and provide details on how to choose the degree of such discounting. Furthermore, being available in closed form, the estimator proposed by Giraitis, Kapetanios, and Yates (2012) in the context of VARs partly addresses the curse of dimensionality, as systems of seven variables are easily handled, see Giraitis, Kapetanios, Theodoridis, and Yates (2014).

The paths traced by the large model literature and by the TVP model literature have seldom crossed. Connections have been established for Factor Augmented VAR (FAVAR) models, see for example Eickmeier, Lemke, and Marcellino (2015) and Mumtaz and Surico (2012), while they are still scant in the VAR literature. A notable exception is represented by the paper by Koop and Korobilis (2013) where the restricted Kalman filter by Raftery, Karny, and Ettler (2010) is used to make a TVP-VAR suitable for large information sets. This approach, while solving

[1]A crucial issue in this framework is how to select the degree of data discounting. The problem is addressed by Giraitis, Kapetanios, and Price (2013), who show how to make this choice data dependent using cross-validation methods.

some issues, presents some shortcomings. First, the curse of dimensionality is only partially addressed since the parametric nature of the model implicitly limits its size. In practice, this framework cannot handle the large information sets employed in factor models (or FAVARs) or the large number of lags that are used when fitting medium-size VARs to monthly data like in Banbura, Giannone, and Reichlin (2010). Second, if the true data generating process is different from the postulated random walk type variation, the robustness of the Kalman filter to model misspecification is an obvious concern.

In this paper we propose an estimator that addresses, in a nonparametric context, both of these problems. Our idea is to start from the nonparametric estimator proposed by Giraitis, Kapetanios and Yates (2012), and adapt it to handle large information sets. To solve the issue of over fitting that arises when the size of the VAR increases, we recur to the mixed estimator by Theil and Goldberger (1960), which imposes stochastic constraints on the model coefficients, therefore mimicking in a classical context the role of the prior in Bayesian models. The resulting estimator, for which we derive asymptotic properties, mixes sample and non sample information to shrink the model parameters. It can be seen both as a generalisation to a time-varying parameter structure of the model by Banbura, Giannone, and Reichlin (2010) and as a penalised regression version of the estimator by Giraitis, Kapetanios and Yates (2012). The proposed method is, given its nonparametric nature, robust to changes in the underlying data generating process and for popular shrinkage methods delivers equation by equation estimation. This implies that the estimator can cope with systems as large as those analysed in the FAVAR and factor model literature.

Our estimator depends crucially on two parameters, the tuning constant that regulates the width of the kernel window used to discount past data, and the penalty parameter that determines the severity of the constraints imposed to control over fitting, akin to the prior tightness in Bayesian estimators. In the paper we explore a variety of cross-validation techniques to set these two parameters based on past model performance.[2] We also consider model averaging as an alternative strategy to deal with model uncertainty.

We next assess in Monte Carlo experiments the finite sample performance of our estimator, which turns out to be good, and compare it with the parametric estimator for TVP-VARs proposed by Koop and Korobilis (2013). We find that when the data

[2]The use of data discounting in regression models as a way to handle structural breaks is the focus of a large literature, see in particular Pesaran and Timmermann (2007), Pesaran and Pick (2011), and Rossi, Inoue, and Jin (2014). All these papers, however, are concerned with single equation regressions rather than with large models.

generating process matches exactly the one assumed in the parametric setup, the two estimators give broadly similar results. Yet, as we move away from this assumption, the performance of the parametric estimator deteriorates, while our non-parametric estimator proves quite robust to changes in the underlying data generating process.

After discussing the theoretical and finite sample properties of the non-parametric estimators, we examine their use through a number of applications. First, we explore whether time variation is indeed a necessary feature of the model to successfully forecast key macroeconomic variables using a large panel (up to 78 variables) of U.S. monthly time series. We organise the forecast exercise around three questions that have been central to the forecasting literature in recent years. The first one is whether time variation actually improves forecast accuracy. The second one is whether the performance of *medium-sized* VARs with time-varying parameters can be approximated by that of *large* VARs with *constant coefficients*. This question is motivated by the contrasting findings in Stock and Watson (2012), who find little evidence of parameter changes during the financial crisis in the context of a factor model, and those reported by Aastveit, Carriero, Clark, and Marcellino (2014), who provide substantial evidence of parameter changes in smaller dimensional VARs. This conflicting evidence suggests that parameter time variation can be due, at least partly, to omitted variables, so that enlarging the information set makes parameters' time variation unnecessary. Once we have established that time variation is indeed beneficial to forecast accuracy the third question is whether it pays off to go beyond a medium size system, i.e. if going from a 20 to a 78 TVP-VAR improves forecast accuracy for the small set of key variables that we are interested in.

The analysis indicates that the introduction of time variation in the model parameters yields an improvement in prediction accuracy over models with constant coefficients, in particular when forecast combination is used to pool forecasts obtained with models with different degrees of time variation and shrinkage. Our findings also indicate that, especially at longer horizons, medium-sized TVP-VARs perform better than a VAR with constant parameters that uses a large information set. Finally, we find that, in the context of TVP-VARs, going beyond 20 variables is not beneficial to forecast accuracy, in line with the results for the constant parameter case in studies such as Banbura, Giannone, and Reichlin (2010) and Koop (2013).

Our non-parametric large TVP-VAR is also useful for structural analysis. As an illustration we revisit, in the context of a large information set, the issue of the diminished effects of oil price shocks on economic activity, a question that has spurred a large number of studies in the applied macro literature in recent years, see for example Hooker (1999), Edelstein and Kilian (2009), Blanchard and Gali (2007), Blanchard and Riggi (2013) and Baumeister and Peersman (2013). The use of a large information set allows us to take a more granular view, allowing us to uncover some interesting findings on the evolving impact of oil price innovations on the output of different sectors of the U.S. industry. Specifically, we find that the declining role of oil prices in shaping U.S. business cycle fluctuations stems from lower effects on the production of durable materials, rather than on the automotive sector on which part of the literature has traditionally focused.

The paper is structured as follows. In Section 2.2 we describe the estimation method and derive its theoretical properties. In Section 2.3, we discuss cross-validation and model averaging. In Section 2.4 we assess the finite sample properties of our nonparametric method in Monte Carlo experiments and compare it with available parametric methods. In Section 2.5 we present the main forecasting exercise. In Section 2.6 we present an analysis of the time-varying impact of unexpected increases in the price of oil on U.S. industrial production. In Section 2.7 we summarise our main findings and conclude. Additional details are provided in Appendixes.

2.2 Setup of the problem

Let us consider a p-order VAR with n variables and time-varying (stochastic) coefficients:

$$\begin{aligned} \underset{1\times n}{y_t'} &= x_t'\Theta_t + \underset{1\times n}{u_t'}\,,\; t = 1, ..., T \\ \underset{1\times k}{x_t'} &= [y_{t-1}', y_{t-2,...,}'y_{t-p}', 1] \\ \underset{k\times n}{\Theta_t} &= [\Theta_{t,1}', \Theta_{t,2}', ..., \Theta_{t,p}', A_t']' \end{aligned} \tag{2.1}$$

where $k = (np + 1)$ is the number of random coefficients to be estimated in each equation so that at each time t there are nk parameters to be estimated, collected in the matrix Θ_t. For the time being, we assume that u_t is a martingale difference process with finite variance Σ_n.[3] A further crucial assumption is that Θ_t changes rather slowly, i.e., that:

$$\sup_{j\leq h} \|\Theta_t - \Theta_{t+j}\| = O_p\left(\frac{h}{t}\right). \tag{2.2}$$

A number of classes of models satisfy (2.2). For example, one such model is obtained by setting $\Theta_t = [\theta_{ij,t}]$, $\tilde{\Theta}_t = [\tilde{\theta}_{ij,t}]$ and letting $\tilde{\theta}_{ij,t} = \tilde{\theta}_{ij,t-1} + \epsilon_{\tilde{\theta},ij,t}$ and

[3] The issue of heteroschedasticity is discussed in Section 2.2.5.

$\theta_{ij,t} = \theta_{ij} \frac{\tilde{\theta}_{ij,t}}{\max_{1 \le i \le t} \tilde{\theta}_{ij,t}}$ for some constants θ_{ij} bounded between 0 and 1 and some set of stochastic processes $\epsilon_{\tilde{\theta},ij,t}$. This is an example of a bounded random walk model. We can allow for a wide variety of processes, $\epsilon_{\tilde{\theta},ij,t}$, making this class suitably wide.

Applying the *vec* operator to both sides of (2.1) we obtain:

$$\underset{n\times 1}{y_t} = \underset{n\times nk}{(I_n \otimes x_t')} \underset{nk\times 1}{\beta_t} + \underset{n\times 1}{u_t}, \tag{2.3}$$

where $\beta_t = vec(\Theta_t)$. Assuming persistence and boundedness[4] of the coefficients in Θ_t, Giraitis, Kapetanios, and Yates (2012, henceforth GKY) show that the path of the random coefficients is consistently estimated by the following kernel estimator:

$$\beta_t^{GKY} = \left[I_n \otimes \sum_{j=1}^{T} w_{j,t}(H) x_j x_j' \right]^{-1} \left[\sum_{j=1}^{T} w_{j,t}(H) vec\left(x_j y_j'\right) \right], \tag{2.4}$$

where the generic j^{th} element $w_{j,t}(H)$ is a kernel function with bandwidth H, used to discount distant data. Throughout the paper we use a Gaussian kernel:[5]

$$w_{j,t}(H) = \frac{K_{j,t}(H)}{\sum_{j=1}^{T} K_{j,t}(H)}, \tag{2.6}$$

$$K_{j,t}(H) = (1/\sqrt{2\pi}) exp\left[-\frac{1}{2} \left(\frac{j-t}{H} \right)^2 \right]. \tag{2.7}$$

One appealing feature of the estimator in (2.4) is that, given the Kronecker structure of the first term, it only requires the inversion of the $k \times k$ matrices $\sum_{j=1}^{T} w_{j,t} x_j x_j'$. In other words, estimation can be performed equation by equation that, as emphasised by Carriero, Clark, and Marcellino (2016) in a Bayesian context, substantially reduces the computing time.

A more compact notation is obtained by introducing the following notation: $X_{w,t} = W_{H,t} X$, where $W_{H,t} = diag(w_{1t}^{1/2}(H), ..., w_{Tt}^{1/2}(H))$ and the $T \times k$ matrix X is formed by stacking over t the vectors x_t'. Also, let us define $X_{ww,t} = W_{H,t} X_{w,t}$ and denote with Y the $T \times n$ matrix formed by stacking over t the vectors y_t'. The GKY

[4]More specifically, writing the VAR in companion form as a VAR(1) model, $Y_t = \Psi_t Y_{t-1}$, GKY assume that the spectral norm (that is the maximum absolute eigenvalue) of Ψ_t is strictly lower than 1.

[5]When forecasting, in order to preserve the pseudo real time nature of the exercise, we introduce an indicator function that assigns zero weight to the out of sample observations, so that only in sample information is used to estimate the parameters:

$$K_{j,t}(H) = (1/\sqrt{2\pi}) exp\left[-\frac{1}{2} \left(\frac{j-t}{H} \right)^2 \right] I(j \le t) \tag{2.5}$$

estimator can now be cast in the following matrix form:

$$\Theta_t^{GKY} = \left[X'_{w,t}X_{w,t}\right]^{-1}\left[X'_{ww,t}Y\right]. \tag{2.8}$$

2.2.1 Shrinkage through stochastic constraints

When the dimension of the system grows, it is desirable to impose some shrinkage on the model parameters to avoid an increase in the estimation variance (Hastie, Tibshirani, and Friedman, 2003). While in a Bayesian framework this can be achieved through the prior distribution, in a classical framework shrinkage can be performed by using the mixed estimator of Theil and Goldberger (1960). This is obtained by adding a set of stochastic constraints (i.e., constraints that hold with some degree of uncertainty) to model (2.3). The constraints are written as linear combinations of the parameter vector β_t plus a vector of noises, where the latter ensures that the constraints do not hold exactly. The complete model can be written as:

$$\underset{n\times 1}{y_t} = \underset{n\times nk}{(I_n \otimes x'_t)}\underset{nk\times 1}{\beta_t} + \underset{n\times 1}{u_t} \tag{2.9}$$

$$\sqrt{\lambda}\underset{nk\times 1}{r} = \sqrt{\lambda}\underset{nk\times nk}{R}\underset{nk\times 1}{\beta_t} + \underset{nk\times 1}{u_t^r}. \tag{2.10}$$

We assume that the errors u_t^r are a martingale difference process with finite variance and that their variance is proportional to that of the data, that is $var(u_t^r) = I_k \otimes \Sigma_n$. In other words, when the noise in the dynamic relationship between y_t and x_t has high variance, uncertainty about the constraints on the coefficients β_t also increases. As for the expected value of u_t^r, for the moment we leave it unspecified since it plays a crucial role in determining the bias of the estimator, as we show further below. Notice that both sides of equation (2.10) are pre-multiplied by a constant $\sqrt{\lambda}$. It is easy to see that this constant acts as a scaling factor of the variance of the stochastic constraints u_t^r.[6] Hence, low values of λ imply that the coefficient vector β_t is left relatively unrestricted; vice versa, high values of λ imply that the constraints in (2.10) hold relatively more tightly. Regarding the structure of the matrix R, we consider two cases. In the first case we assume that R has a Kronecker structure: $\underset{nk\times nk}{R} = (I_n \otimes \underset{k\times k}{\overline{R}})$. This case is of particular interest for two reasons. First, it holds for a number of popular shrinkage methods, like the Ridge regression and the Litterman prior. Second, it results in an estimator that can be cast in matrix form, hence being very efficient from a computational point

[6] Notice that, by multiplying both sides of (2.10) by $\frac{1}{\sqrt{\lambda}}$ the variance of the noise in (2.10) becomes $\frac{1}{\lambda}(I_k \otimes \Sigma_n)$.

of view and directly comparable to its unconstrained counterpart, i.e. the GKY estimator. Next, we consider the more general case where R does not have this particular structure.

2.2.2 Case 1: R has a Kronecker structure

If R has a Kronecker structure, the analysis of the estimator can proceed equation by equation. First, let us state the following definitions: $\underset{kn\times 1}{u_t^r} = vec(\underset{k\times n}{\overline{u}_t^r})$ and $\underset{kn\times 1}{r} = vec(\underset{k\times n}{\overline{r}})$. Also, since $\underset{nk\times nk}{R} = (I_n \otimes \underset{k\times k}{\overline{R}})$ and $\beta_t = vec(\Theta_t)$, it follows that $R\beta_t = (I_n \otimes \overline{R})vec(\Theta_t) = vec(\overline{R}\Theta_t)$. Hence the joint model in (2.9) and (2.10) can be expressed in matrix form as:

$$\underset{1\times n}{y_t'} = \underset{1\times k}{x_t'}\underset{k\times n}{\Theta_t} + \underset{1\times n}{u_t'}, \tag{2.11}$$

$$\sqrt{\lambda}\underset{k\times n}{\overline{r}} = \sqrt{\lambda}\underset{k\times k}{\overline{R}}\underset{k\times n}{\Theta_t} + \underset{k\times n}{\overline{u}_t^r}, \tag{2.12}$$

or more compactly:

$$\underset{(k+1)\times n}{y_t^*} = x_j'^*\Theta_t + u_t^*, \tag{2.13}$$

where $y_t^* = [y_t, \sqrt{\lambda}\overline{r}']'$, $x_j'^* = [x_t, \sqrt{\lambda}\overline{R}']'$, $u_t^* = [u_t, \overline{u}_t'^r]'$, and $var(vec(u_t^*)) = I_{nk+1} \otimes \Sigma_n$. The extended regression model in (2.13) can be analysed using the GKY estimator, with related properties. The estimator has the form:

$$\underset{k\times n}{\widehat{\Theta}_t} = \left(\sum_{j=1}^{T} w_{j,t}x_j^*x_j^{*\prime}\right)^{-1}\left(\sum_{j=1}^{T} w_{j,t}x_j^*y_j^*\right), \tag{2.14}$$

where, for simplicity, we have omitted the dependence of $w_{j,t}$ from the bandwidth H. Separating the contribution of the actual data from that of the constraints, the estimator can equivalently be written as:[7]

$$\underset{k\times n}{\widehat{\Theta}_t} = \left(\sum_{j=1}^{T} w_{j,t}x_jx_j' + \lambda\overline{R}'\overline{R}\right)^{-1}\left(\sum_{j=1}^{T} w_{j,t}x_jy_j' + \lambda\overline{R}'\overline{r}\right) \tag{2.15}$$

$$= \left(X_{w,t}'X_{w,t} + \lambda\overline{R}'\overline{R}\right)^{-1}\left(X_{ww,t}'Y + \lambda\overline{R}'\overline{r}\right). \tag{2.16}$$

It is worth making the following observations. First, when $\lambda = 0$ the constrained

[7]Notice that $\sum_{j=1}^{T} w_{j,t}x_j^*x_j^{*\prime} = \sum_{j=1}^{T} w_{j,t}\begin{bmatrix} x_t & \sqrt{\lambda}\overline{R}' \end{bmatrix}\begin{bmatrix} x_t' \\ \sqrt{\lambda}\overline{R} \end{bmatrix}$ and $\sum_{j=1}^{T} w_{j,t}x_j^*y_j^* = \sum_{j=1}^{T} w_{j,t}\begin{bmatrix} x_t & \sqrt{\lambda}\overline{R}' \end{bmatrix}\begin{bmatrix} y_t\prime \\ \sqrt{\lambda}\overline{r} \end{bmatrix}$.

estimator equals the unconstrained one: $\widehat{\Theta}_{t,GKY} = \left(X'_{w,t}X_{w,t}\right)^{-1}\left(X'_{ww,t}Y\right)$. Second, and vice versa, as $\lambda \to \infty$, $\widehat{\Theta}_t$ converges to the value implied by the constraints, that is $\Theta_t \to \Theta_C = (\overline{R}'\overline{R})^{-1}(\overline{R}'\overline{r})$. Hence, the constant term $\sqrt{\lambda}$ can also be interpreted as the weight of the sample size of the artificial observations (r and R) relative to T, the sample size of the observed data y_t and x_t. It is worth remarking that the value implied by the constraints (Θ_C) is time invariant. This means that the stochastic constraints *anchor* the evolution of Θ_t around a fixed value that is specified ex ante. Third, $\widehat{\Theta}_t$ can be expressed as the weighted sum of its unrestricted and restricted versions, $\widehat{\Theta}_{t,GKY}$ and Θ_C. To see this point, re-write (2.16) as follows:

$$\underset{k\times n}{\widehat{\Theta}_t} = \left(X'_{w,t}X_{w,t} + \lambda\overline{R}'\overline{R}\right)^{-1}\left[(X'_{w,t}X_{w,t})\widehat{\Theta}_{t,GKY} + (\lambda\overline{R}'\overline{R})\Theta_C\right] \quad (2.17)$$

$$= S_w^{-1}(X'_{w,t}X_{w,t})\widehat{\Theta}_{t,GKY} + S_w^{-1}(\lambda\overline{R}'\overline{R})\Theta_C, \quad (2.18)$$

where $S_w = \left(X'_{w,t}X_{w,t} + \lambda\overline{R}'\overline{R}\right)$.

The properties of $\widehat{\Theta}_t$ are derived in the following theorem.

Theorem 1 *Let the model be given by (2.13) where u_t^* is a martingale difference sequence with finite fourth moments. Let (2.2) hold and $H = o(T^{1/2})$. Let $X_{w,t}^* = W_{H,t}X^*$ where X^* is obtained by stacking over t the vectors $x_t'^*$, $X_{ww,t}^* = W_{H,t}X_{w,t}^*$, $\Gamma_{w,t}^* = p\lim \frac{1}{H}X_{w,t}^{*\prime}X_{w,t}^*$, $\Gamma_{ww,t}^* = p\lim \frac{1}{H}X_{ww,t}^{*\prime}X_{ww,t}^*$. Then,*

$$\left(\Gamma_{w,t}^{*-1}\Gamma_{ww,t}^{**}\Gamma_{w,t}^{*-1} \otimes \Sigma_n\right)^{-\frac{1}{2}}\sqrt{H}vec\left(\widehat{\Theta}_t' - \Theta_t' - \Theta_t'^B\right) \to^d N(0, I), \quad (2.19)$$

where $\Theta_t^B = p\lim S_w^{-1}\sqrt{\lambda}\overline{R}'\overline{u}_t^r = p\lim S_w^{-1}\lambda\overline{R}'\left(\overline{r} - \overline{R}\Theta_t\right)$ and $\Gamma_{ww,t}^{**}$ is defined in (2.43).

Proof. See Appendix 2.A. ■

Notice that the bias term $\Theta_t^B = p\lim S_w^{-1}\sqrt{\lambda}\overline{R}'\overline{u}_t^r = p\lim S_w^{-1}\lambda\overline{R}'\left(\overline{r} - \overline{R}\Theta_t\right)$ depends on the distance between the true parameters Θ_t from the constraints. Since the direction towards which we are shrinking is time-invariant it is possible that at some point in time the constraints $\overline{R}\Theta_t = \overline{r}$ actually hold, but in general they cannot hold at *each* point in time, or Θ_t would be time-invariant, which contradicts one of the assumptions.[8]

[8]The estimator could be made consistent by designing a penalty parameter λ that vanishes as T goes to infinity. This point is discussed in a related context in De Mol, Giannone, and Reichlin (2008), Appendix A.

2.2.3 Case 2: R does not have a Kronecker structure

Let us now turn to the more general case when R does not have a Kronecker structure. In this case the estimator can be written as:

$$\widehat{\beta}_t = \left[\left(\underbrace{I_n \otimes \sum_{j=1}^{T} w_{j,t} x_j x_j'}_{nk \times nk} \right) + \lambda R'R \right]^{-1} \left[\sum_{j=1}^{T} w_{j,t} \left(I_n \otimes x_j \right) y_j + \lambda R'r \right]$$
$$= \left[\left(I_n \otimes \sum_{j=1}^{T} w_{j,t} x_j x_j' \right) + \lambda R'R \right]^{-1} \left[\sum_{j=1}^{T} w_{j,t} vec(x_j y_j') + \lambda R'r \right]. \qquad (2.20)$$

A crucial difference between the estimator in (2.20) and the one in (2.16) is that the latter only requires the inversion of k dimensional matrices, which makes it computationally much faster. On the other hand, (2.20) can handle more general constraints. The properties of $\widehat{\beta}_t$ are derived in the following theorem.

Theorem 2 *Let the model be given by (2.9) and (2.10) where u_t is a martingale difference sequence with finite fourth moments. Let (2.2) hold, $H = o(T^{1/2})$. Let us define $\Phi = \frac{\lambda}{H} R'R$, $\Gamma_{w,t} = p\lim \frac{1}{H} \sum_{j=1}^{T} w_{j,t}(x_j x_j')$, $\Gamma_{ww,t} = p\lim \frac{1}{H} \sum_{j=1}^{T} w_{j,t}^2 x_j x_j'$ and $\beta_t^B = p\lim \left[\left(I_n \otimes \sum_{j=1}^{T} w_{j,t} x_j x_j' \right) + \lambda R'R \right]^{-1} \lambda R'r$. Then,*

$$\sqrt{H} \left[\left(I_n \otimes \Gamma_{w,t} + \Phi \right)^{-1} \left(\Sigma_n \otimes \Gamma_{ww,t} + \Phi \right) \left(I_n \otimes \Gamma_{wt} + \Phi \right)^{-1} \right]^{-1/2} \left(\widehat{\beta}_t - \beta_t - \beta_t^B \right) \to^d N\left(0, I\right) \qquad (2.21)$$

Proof. See Appendix 2.B. ■

2.2.4 Conditional Variance and MSE of the constrained estimator

In this subsection we analyse the effects of stochastic constraints on the variance of the penalised estimator relative to that of the unconstrained one. In particular, we show that the stochastic constraints have the unambiguous effect of lowering the variance of the non-parametric estimator. This is a standard finding in penalised estimators where a bias/variance trade off emerges, but it is a novel result in models with stochastic time varying coefficients. Notice that, since the model coefficients are stochastic, we need to condition the derivation of the variance of the estimators on a given realisation of the whole time series of the *true* coefficients $(\beta_1, \beta_2, \ldots, \beta_t, \ldots, \beta_T)$. Still, the ranking of the conditional variances derived in the theorem is unambiguous for any given realisation of the coefficients and therefore absolutely general.

Theorem 3 *Let the model be given by (2.9) and (2.10). Let* $E = I_n \otimes \sum_{j=1}^{T} w_{j,t} x_j x_j'$, $F = \lambda R'R$, *and* $G = \sum_{j=1}^{T} w_{j,t} vec(x_j y_j') = \sum_{j=1}^{T} w_{j,t}(I_n \otimes x_j) y_j'$. *Define a given realisation of the stochastic coefficients* $\beta^T = (\beta_1, \beta_2, \ldots, \beta_t, \ldots, \beta_T)$. *Then the constrained estimator can be written as*

$$\widehat{\beta}_t = [E + F]^{-1} \left[E \widehat{\beta}_{t,GKY} + \lambda R' r \right].$$

Further, we have that $avar(\widehat{\beta}_{t,GKY}|\beta^T) - avar(\widehat{\beta}_t|\beta^T)$ *is a positive semi-definite matrix, where the* $avar(.|.)$ *operator indicates the conditional asymptotic variance.*

Proof. See Appendix 2.C. ■

In Appendix 2.C we also discuss the ranking of the conditional Mean Square Error (MSE, given by the sum of the variance and of the squared bias) of the two estimators. We show that the stochastic constraints also lower the conditional MSE as long as there is sufficient collinearity in the dataset.

2.2.5 Time-varying volatilities

When both the error variances and the VAR coefficients change over time, variations in the parameters and in the variances can be confounded, see Cogley and Sargent (2005). An important implication is that if changes in the variances of the errors are neglected then the importance of variation in the VAR coefficients could be overstated. Giraitis, Kapetanios, and Yates (2014) show that the properties of their estimator are unaffected by the presence of stochastic volatilities as long as standard errors are studentised by an appropriate time-varying covariance matrix for the error terms. When performing structural analysis in a VAR context, GKY suggest to model time variation in the variance of the disturbances with a two-step approach. The method consists of fitting first an homoschedastic VAR, then estimating the time-varying volatilities on the residuals obtained in the first stage via the following kernel estimator:

$$\widehat{\Psi}_t = \sum_{j=1}^{T} w_{j,t}(H_\Psi) u_t u_t', \tag{2.22}$$

where the bandwidth parameter H_Ψ is not necessarily the same as the one used to estimate the coefficients. Orthogonalisation of the residuals is then based on the time-varying covariance matrix $\widehat{\Psi}_t$.

Our penalised kernel estimator can be adapted to account for changing volatilities along these lines, using the GLS correction proposed in Theil and Goldberger (1960). In the first step, the VAR coefficients are estimated using (2.20) and the resulting residuals are used to compute $\widehat{\Psi}_t$ as in (2.22). In a second step a GLS correction is applied:

$$\beta_t = \left[\sum_{j=1}^{T} w_{j,t}\left(\widehat{\Psi}_t^{-1} \otimes x_j' x_j\right) + \lambda R'R\right]^{-1} \left[\sum_{j=1}^{T} w_{j,t} vec\left(x_j' y_j' \widehat{\Psi}_t^{-1}\right) + \lambda R'r\right] \quad (2.23)$$

Notice that this GLS correction requires the inversion of potentially large matrices ($nk \times nk$), which slows down computation and limits the size of the VAR. In the empirical applications and in the Monte Carlo analysis discussed in Sections 2.5 and 2.4, where we experiment with relatively large systems, we therefore do not apply this correction. However, in section 2.5.5 we appraise the merits of this GLS correction in terms of forecasting performance. We find that the estimator that *does not* account for time-varying volatility actually produces more accurate forecasts.

In the next two sub-sections we describe how two popular shrinkage methods can be adapted to our setup. Notice that this is just an illustration of how the method can be used, not an exhaustive list of the constraints that can be applied. As long as the matrix R has column rank nk (i.e. the number of independent constraints is at least as large as the number of parameters of the VAR, but it can also be larger), the matrix $R'R$ appearing in the estimator will regularise estimation.

2.2.6 Ridge type shrinkage

The Ridge regression penalty shrinks all the parameters uniformly towards zero at a given penalty rate λ. A TVP-VAR with Ridge shrinkage can be obtained by setting $R = I_{nk}$ and $r = 0$, which consists of imposing the following stochastic constraints at each t:

$$0 = \sqrt{\lambda}\beta_t + u_t^r. \quad (2.24)$$

where the properties of u_t^r are as defined in Theorem 1.

The resulting estimator takes the form:

$$\beta_t^{Ridge}(\lambda, H) = \left[I_n \otimes \left(\sum_{j=1}^{T} w_{j,t} x_j' x_j + \lambda I_k\right)\right]^{-1} \left[vec\left(\sum_{j=1}^{T} w_{j,t} x_j' y_j'\right)\right] \quad (2.25)$$

Notice that, given the Kronecker structure of the constraints (as R is an identity matrix), estimation can proceed equation by equation and the estimator can be written in matrix form as:

$$\Theta_t^{Ridge}(\lambda, H) = \left[X'_{w,t}X_{w,t} + \lambda I_k\right]^{-1}\left[X'_{ww,t}Y\right] \tag{2.26}$$

2.2.7 Litterman type shrinkage

Some of the features of the Ridge penalty can be unappealing in the context of VAR models. First, the fact that all the coefficients are shrunk towards zero imposes a structure of serially uncorrelated data, which is at odds with the strong persistence that characterises most macroeconomic time series. Second, the same penalty is imposed on all the coefficients (including the intercept). Yet, having some flexibility in the penalisation of the different parameters could be desirable. A more general set of stochastic constraints, which produce the same effects that the Litterman prior has in a Bayesian framework,[9] is given by setting $\overline{r}$ and $\overline{R}$ as follows:

$$\overline{r} = \begin{pmatrix} \text{diag}(\delta_1\sigma_1, \delta_2\sigma_2, \delta_3\sigma_3, ..., \delta_n\sigma_n) \\ 0_{n(p-1)+1\times n} \end{pmatrix} \qquad \overline{R} = \begin{pmatrix} \overline{\Sigma} & 0 \\ 0 & \sigma_c^2 \end{pmatrix}, \text{ where}$$
$$\overline{\Sigma} = diag(1, 2, 3, ..., p) \otimes diag(\sigma_1, \sigma_2, ..., \sigma_n) \tag{2.27}$$

Notice that the vector $r = vec(\overline{r})$ towards which the VAR coefficients are driven by the constraints is generally different from zero. In empirical applications, for data in levels, the n values δ_i are typically set to 1, so that the model is pushed to behave like a multivariate random walk plus noise. Moreover, unlike in Ridge regressions, the precision of the constraints is not uniform across parameters but it is higher for more distant lags, as implied by the decay terms $(1, 2, ..., p)$. The scaling factors $\sigma_1^2, ..., \sigma_n^2$ appearing in $\overline{\Sigma}$ can be obtained by univariate regressions and the precision on the intercept σ_c^2 can be set to an arbitrarily small or large value, depending on the application.[10]

Summarising, by appropriately penalising the GKY estimator, some discipline on the VAR coefficients can be imposed through stochastic constraints a la Theil and Goldberger (1960). This makes the GKY method, originally designed for small/medium scale VARs, suitable for handling large n dataset. We have seen that the resulting estimator has a well defined asymptotic distribution under rather mild

[9] Karlsson (2012) distinguishes the Litterman prior from the more general Minnesota prior based on the assumptions on the covariance matrix of the VAR residuals, which is assumed to be diagonal in the Litterman prior, full in the more general Minnesota prior, see Kadiyala and Karlsson (1993, 1997).

[10] For example, in a Bayesian context, Carriero, Kapetanios, and Marcellino (2009) adopt a very tight prior centred around zero on the intercept in a large VAR, favouring a driftless random walk behaviour, to capture the behaviour of a panel of exchange rates.

conditions, and is generally more efficient than the unconstrained GKY estimator. Moreover, for popular shrinkage methods the resulting estimator can be cast in matrix form, with notable computational advantages.

The double nature of the estimator (being both nonparametric and penalised) is captured by its dependence on the two constants: H, the bandwidth parameter that determines the weight that each observation has as a function of its distance from t, and λ, a constant that determines the severity of the penalty. In the next section we discuss alternative solutions to the problem of determining these two parameters in empirical applications.

2.3 Model specification

The problem of setting appropriate values of λ and H can be tackled in two ways. The first is model selection, which typically rests on the optimisation of a given criterion. We describe two such criteria. The former adapts to our problem the procedure devised by Banbura, Giannone, and Reichlin (2010), and has an "in sample" fit flavour. The latter favours models with better out of sample performance and is inspired by the method proposed by Kapetanios, Labhard, and Price (2008) for assigning weights to different models in the context of forecast averaging. In the remainder of the paper we will refer to these two criteria as L_{fit} and L_{mse}. The second route consists of pooling the results obtained on the basis of a large range of different specifications. We describe each strategy in turn.

2.3.1 Model selection criteria

The L_{fit} criterion

The first criterion that we consider adapts to our problem the method by Banbura, Giannone, and Reichlin (2010). The intuition of the method is that, when forecasting with large datasets, some variables are more relevant than others. Over fitting should then be penalised up to the point where a large VAR achieves the same fit as that of a smaller VAR that only includes the key variables of interest. We adapt their criterion to the problem of choosing simultaneously λ and H. Formally, the criterion involves the following steps:

1. Pick a subset of n_1 variables of interest out of the n variables in the VAR.
2. Compute the in sample fit of a benchmark VAR with constant coefficients that only includes these n_1 variables.

3. Select λ and H to minimise the distance between the in sample fit of the large n variate VAR (featuring both time-varying parameters and shrinkage) and the benchmark VAR.

Formally, the loss function to be minimised is the following:

$$L_{fit}(\lambda, H) = \left| \sum_{i=1}^{n_1} \frac{rss_n^i(\lambda, H)}{var(y_{t,i})} - \sum_{i}^{n_1} \frac{rss_{n1}^i}{var(y_{t,i})} \right|$$

where the scaling by $var(y_{t,i})$ is needed to account for the different variance of the variables.

The L_{mse} criterion

As an alternative, λ and H can be selected at each point in time based on the predictive performance of the model in the recent past. The method, which has a cross-validation flavour, is similar in spirit to the one used by Kapetanios, Labhard, and Price (2008) to compute model weights in the context of forecast averaging. The necessary steps are the following:

1. Pick a subset of n_1 variables of interest out of the n variables in the VAR.[11]
2. At each step t in the forecast exercise and for each forecast horizon h consider a relatively short window of recent data $t-L-h, t-L-h+1, \ldots, t-1-h$ and compute the h steps ahead Mean Square Error (MSE) $mse_h^i(\lambda, H)$, for each $i \in n_1$.
3. Pick the values of λ and H that minimise the sum of these n_1 MSEs.

Formally, the loss function to be minimised is:

$$L_{mse}(\lambda, H) = \sum_{i=1}^{n_1} \frac{mse_h^i(\lambda, H)}{var(y_{t,i})}$$

where, again, the mse_i is scaled by the variance of y_i.

Practical considerations

In principle, standard optimisation algorithms could be used to minimise bot the L_{fit} and the L_{mse} criterion. However, we have often found that the minimum occurs at a kink. A problem of this type could arise because the stochastic constraints shrink

[11] Notice that, when using this criterion, we could set $n_1 = n$, that is we could focus on the whole set of variables in the VAR rather than only on a subset.

the VAR coefficients towards a constant parameter structure but time variation is also affected by the width of the kernel.

Since our estimator is easy to compute a feasible solution is represented by a grid search approach, along the lines of Carriero, Kapetanios, and Marcellino (2009) and Koop and Korobilis (2013). More specifically, in the empirical analysis that follows, we experiment with a wide (38 elements) grid for (the reciprocal) of λ, $\varphi = 1/\lambda$.

$$\varphi_{grid} = 10^{-10}, 10^{-5}, 10^{-4}, 10^{-3}, 10^{-2}, 10^{-2} + .3, 10^{-2} + 2 \times .3, 10^{-2} + 3 \times 03, \ldots, 1 \tag{2.28}$$

We suggest the use of a wider grid than the one used, for instance, by Koop and Korobilis (2013) because the stochastic constraints that we apply are binding at each point in time (rather than just at the initial condition like in Kalman filter based estimation methods), so that higher values of φ (i.e. lower values of λ) are needed to allow for meaningful time variation in the VAR coefficients.

Regarding the width of the kernel function $w_{j,t}$, we work with a six points grid for the tuning parameter H:

$$H_{grid} = 0.5, 0.6, 0.7, 0.8, 0.9, 1, \tag{2.29}$$

consistently with the parameterisation used in Monte Carlo experiments by Giraitis, Kapetanios, and Price (2013).

2.3.2 Pooling

An alternative strategy to model selection consists of pooling model estimates obtained with different values of λ and H. This could be particularly valuable in the context of forecasting. From a theoretical standpoint, the rationale for forecast pooling in the presence of structural breaks is offered for example by Pesaran and Pick (2011), who show that averaging forecasts over different estimation windows reduces the forecast bias and mean squared forecast errors, provided that breaks are not too small. In empirical applications pooling is typically found to be effective in improving forecast accuracy both in Bayesian (Koop and Korobilis, 2013) and in frequentist (Kapetanios, Labhard, and Price, 2008) settings.[12] In our context, model pooling could be based on relatively sophisticated weighting schemes, based on the selection criteria described in the previous subsections, or on simpler strategies like equal weights averaging.

[12] Kuzin, Marcellino, and Schumacher (2013) show that forecast pooling also works well in nowcasting GDP.

2.4 Finite sample properties

To assess the finite sample properties of our estimator we design a Monte Carlo exercise in which we contrast the forecasting performance of the non-parametric estimator with that of a popular parametric alternative.

We consider three alternative DGPs. In the first DGP (DGP-1) we assume that the coefficients follow a random walk plus noise process:

$$Y_t = \Lambda_t Y_{t-1} + \varepsilon_t$$
$$\Lambda_t = \Lambda_{t-1} + \eta_t$$

We make the stochastic process of the coefficients broadly consistent with a Litterman prior by bounding the first autoregressive parameter to lie between 0.85 and 1.[13] In the second one (DGP-2) we let the coefficients break only occasionally rather than at each time t:

$$\Lambda_t = (1 - I(\tau))\Lambda_{t-1} + I(\tau)\Lambda_{t-1} + \eta_t$$

The probability of the coefficients breaking equals a constant τ that we set to .025, implying that, with quarterly data, we would observe on average a discrete break once every ten years. We also relax the bounds on Λ_t and let them fluctuate randomly between 0 and 1.[14] In the third set of simulations (DGP-3) coefficients evolve as a sine functions and are bounded between -1 and 1:

$$\Lambda_t = sin(10\pi t/T) + \eta_t$$

In all DGPs we assume $\eta_t \backsim N(0,1)$ and random walk stochastic volatilities for the measurement equations:

$$\begin{aligned} \varepsilon_{it} &= u_{it} \exp(\lambda_{it}) \\ \lambda_{it} &= \lambda_{it-1} + \nu_{it} \end{aligned}$$

where $u_{it} \backsim N(0,1)$ and $\nu_{it} \backsim N(0, \sigma_\eta)$. We calibrate $\sigma_\eta = 0.01$.[15] For the remaining

[13] Details on how this is achieved are presented in Appendix 2.D.

[14] With tight boundaries (like the 0.85-1 interval imposed in DGP-1) the difference between coefficients that break only occasionally and coefficients that drift slowly, is negligible.

[15] Notice that this value for σ_η is quite large. Cogley and Sargent (2005) for example, assume that a priori σ_η is distributed as an inverse gamma with a single degree of freedom and scale parameter 0.01^2. Since the scale parameter can be interpreted as the (prior) sum of square residuals, this means that a priori they set the variance of the innovations to the log-volatility to $0.01^2/T$. Assuming $T = 100$, the prior variance is 10^{-6} , as opposed to our choice of 10^{-2}.

technical details on the design of the Monte Carlo exercise see Appendix 2.D

We assess the performance of our method based on the accuracy of one step ahead forecast errors. As a benchmark, we use the parametric estimator developed by Koop and Korobilis (2013), which we briefly describe in the next sub-section. While the controlled environment provided by the Monte Carlo exercise allows us to evaluate how robust the two methods are to different assumptions on the law of motion of the model parameters, a comparison of the two approaches based on actual data is presented later in section 2.5.

2.4.1 A parametric estimator

The model specification adopted by Koop and Korobilis (2013) follows closely the literature on (small) TVP-VARs in that it assumes a random walk evolution of the VAR coefficients. The model can then be cast in State Space, where the VAR equations

$$y_t = Z_t \beta_t + \varepsilon_t \tag{2.30}$$

serve as measurement equations and the unobserved states, the parameters β_t, evolve as driftless random walks plus noise:

$$\beta_{t+1} = \beta_t + u_{t+1}, \tag{2.31}$$

with $\varepsilon_t \sim N(0, \Sigma_t)$ and $u_t \sim N(0, Q_t)$. Also ε_t and u_t are independent of one another and serially uncorrelated. Even for medium-sized VARs the estimation algorithms developed by Cogley and Sargent (2005) and Primiceri (2005) become unfeasible due to computational complexity. To overcome these difficulties, following the literature on Adaptive Algorithms, see for example Ljung (1992) and Sargent (1999), Koop and Korobilis (2013) make two simplifying assumptions. The first one involves the matrix Q_t, which is specified as follows:

$$Q_t = \left(\frac{1-\theta}{\theta} \right) P_{t-1/t-1} \tag{2.32}$$

where $P_{t-1/t-1}$ is the estimated covariance matrix of the unobserved states β_{t-1} conditional on data up to $t-1$ and θ is a forgetting factor ($0 < \theta < 1$).[16] A similar simplifying assumption on Σ_t ensures that this matrix can be estimated by suitably

[16]Equation (2.32) basically states that the amount of time variation of the model parameters at time t is a small fraction of the uncertainty on the unobserved state β_t, so that large uncertainty on the value of the state at time t translates into stronger parameter time variation

discounting past squared one step ahead prediction errors:

$$\widehat{\Sigma}_t = \kappa \widehat{\Sigma}_{t-1} + (1-\kappa) v_t v_t' \tag{2.33}$$

where $v_t = y_t - Z_t \beta_{t/t-1}$. These assumptions make the system matrices Q_t and Σ_t (which are an *input* to the Kalman filter at time t) a function of the t-1 *output* of the Kalman filter itself. This recursive structure implies that, given an initial condition and the two constants θ and κ, an estimate of the coefficients β_t can be obtained through a single Kalman filter pass. Although it is laid out in a Bayesian spirit, the restrictions imposed on the Kalman filter recursions reduce the estimation procedure to a discounted least squares algorithm.

Before moving to the results of the Monte Carlo exercise let us make some remarks on the relative merits of the parametric approach compared to our non-parametric estimator. First, the use of a parametric model, and the simplifications imposed on the model structure to make the estimation feasible, do not come without costs. One potential pitfall is that the model assumes a very specific evolution for the model parameters. The driftless random walk assumption, widely used in econometrics and macroeconomics, does not have any other grounding than parsimony and computational convenience. If the true data generating process (DGP) is, however, very different from the one posited, the model is misspecified and this could result in poor performance. The second issue is that the curse of dimensionality is only partially solved. For 20 variables and 4 lags (a standard application in the large VAR literature with quarterly data) the stacked vector β_t contains 1620 elements. Larger model sizes (arising from a higher number of series in the system or by a higher number of lags, like the 13 lags conventionally used with monthly data in levels) are intractable in this setup. Finally, since the only source of time variation in the model is the prediction error, it can be shown that this forgetting factor algorithm boils down to an exponential smoothing estimator.[17] This means that the effect of the prior on the initial condition β_1 will die out relatively quickly. Also, the longer the sample size, the lower the effect of the prior on the parameter estimates. In contrast, the stochastic constraints that we use to penalise our estimator are effective at each point in time.

2.4.2 Monte Carlo results

For all the DGPs we fit our non-parametric estimator with Litterman-type stochastic constraints and average forecasts using equal weights across all the possible values

[17]See Delle Monache and Petrella (2014), Section 2.

of H and λ specified in the grids described in the previous Section. The parametric method also needs a prior on the initial value of the parameters, β_1 to discipline the estimation towards values that are a priori plausible. To keep the comparison with our method as fair as possible we also impose on the initial condition of the parametric model a Litterman type prior. The remaining details of the model specification of the parametric model are quite lengthy and are documented in the Appendix 2.E.

The results of the Monte Carlo exercise are shown in Table 2.1. The methods are compared in terms of 1 step ahead RMSE (relative to that of the parametric estimator and averaged across the n variables using either equal or inverse RMSE weights) for VARs of different sizes ($n = 7$ and $n = 15$) and for different sample sizes (100, 150 and 200). Throughout the exercise forecasts from our proposed estimator are obtained by equal weights averaging across different values for H and λ. Forecasts are computed on the second half of the sample, i.e. when T=100, forecasts are computed recursively for t=51 to t=100, when T=150 forecasts are computed recursively for t=76 to t=150 and so forth.

In the case of DGP-1 the performance of the two estimation methods is broadly comparable, with the parametric estimator improving slightly (by at most 2%) on the nonparametric one only for VARs of larger sizes. Notice that in this context one would expect the parametric estimator to have an edge, given the tight correspondence between the assumptions made by the model and the actual DGP. The gain attained by this method proves, however, negligible. When we move to DGP-2 and DGP-3 the relative performance of the nonparametric estimator improves steadily, with gains of the order of 15% in the case of DGP-3 and n=15. Although these DGPs are probably less representative of the typical relationship across macroeconomic time series, they do unveil some fragility of the Kalman filter based method, whose performance rapidly deteriorates when the behaviour of the coefficients moves further and further away from the random walk setting. The nonparametric estimator, on the other hand, not only proves robust to heteroschedastic errors but also to a wide range of different specifications of the coefficients.

Summing up, the results of the Monte Carlo analysis are quite supportive of the nonparametric estimator coupled with stochastic constraints. While this does not constitute conclusive evidence in favour of our non-parametric approach, we believe that its good theoretical and finite sample properties, combined with its computational efficiency, make it a very competitive benchmark for modelling and forecasting with large VARs, taking into consideration the possibility of time variation.

2.5 Forecast Evaluation

After evaluating the finite sample properties of our estimator by means of simulation experiments, we now explore its performance in the context of an extensive forecast exercise based on U.S. data. We first discuss the set-up of the exercise, next we present the results, evaluate the role of forecast pooling and of model size in the TVP context, and finally consider a comparison with the Koop and Korobilis (2013) approach.

2.5.1 Set-up of the exercise

Throughout the exercise we use Litterman type constraints, like Banbura, Giannone, and Reichlin (2010). The information set is composed of 78 time series spanning around five decades, from January 1959 to July 2013. Table 2.2 reports the list of the series used in the exercise together with the value of $\bar{r}$ used for each variable. Following the convention in the Bayesian literature we set to 1 the elements of $\bar{r}$ corresponding to variables that display a trend and to 0 those corresponding to variables that have a stationary behaviour (typically surveys). We examine the performance of VARs of two sizes. A medium sized VAR that includes only the 20 indicators that are highlighted in red in Table 2.2 and a large VAR that makes use of all the available information.[18]

We experiment with different model specifications obtained by intersecting various options for setting λ and H as summarised in Table 2.3. The table is organised in two panels. The top panel refers to model specifications that make use of the L_{fit} criterion, the bottom panel, on the other hand, to specifications based on the L_{mse} criterion. Starting from the top panel, the first set of models (M1 in Table 2.3) is obtained by fixing H at a given point in the grid and, conditional on this value of H, setting λ optimally at each t at the value that minimises the L_{fit} function. The second set of models (M2) are obtained as variants of M1 by choosing the λ that minimises L_{fit} in the pre-sample and then keeping it fixed for the rest of the exercise. In the third set of models (M3) the function L_{fit} is optimised at each t both with respect to λ and H. The fourth case (M4) is obtained as variant of M3 by choosing λ optimally in the pre-sample and then keeping it fixed for the rest of

[18]Koop and Korobilis (2013) also look at the performance of trivariate VARs with TVP. We do not pursue this route as over fitting is not an issue in small systems and in those cases the use of the unconstrained GKY estimator is appropriate.

the exercise. The remaining models (M5 to M8) are obtained by replacing the L_{fit} with the L_{mse} criterion. These different model specifications allow us to assess the importance of the various elements that characterise the proposed estimator.

The subset n_1 of variables of interest on which we focus the forecast evaluation is set to $n_1 = 3$, and we monitor the performance of three indicators of particular interest for monetary policy, i.e. the Fed Fund Rates (FEDFUNDS), the number of non farm payroll employees (PAYEMS) and CPI inflation (CPIAUCSL). We fix the lag length to 13 and retain 10 years of data (120 observations) as the first estimation sample. We then produce 1 to 24 months ahead pseudo real time forecasts with the first estimation sample ending in January 1970 and the last one ending in July 2011, for a total number of 499 forecasts. Finally, in the case of the L_{mse} criterion we need to choose L, that is the width of the short window of data on which to measure the predictive performance of the model. We set L =36 (corresponding to three years of data). As a benchmark we adopt the large Bayesian VAR (BVAR) with a Litterman prior and constant coefficients, which can be obtained as a restricted version of our estimator by shutting down the time variation in the VAR coefficients.

2.5.2 Results

As a first piece of evidence, in Figure 2.1 (for the 20 variables VAR) and Figure 2.2 (for the 78 variables VAR) we show the behaviour of the penalty parameter λ in the specifications where both λ and H are optimised over time with the L_{fit} criterion (specification M7 in Table 2.3).[19] As mentioned, high values of λ imply that the constraints hold more tightly, so that the VAR coefficients are less informed by the data. Starting from Figure 2.1, three distinct phases can be identified. In the first one λ starts from relatively low values and increases smoothly over time. In the 80s and throughout the Great Moderation it stays relatively constant around this value, to start falling again in the mid 1990s, with a steeper slope at the beginning of the 2000s. These results are broadly in line with those stressed in the literature on the predictability of macro times series before and after the Great Moderation. For example D'Agostino, Giannone, and Surico (2006) find that the predictive content for inflation and economic activity of common factors extracted from large panels weakened significantly during the Great Moderation, while in periods of higher volatility cross-sectional information proved more relevant for forecasting. Given the direct relationship between the relevance of cross sectional information and

[19] Results obtained using the L_{mse} are qualitatively similar, but for some data points the penalty parameter λ goes to infinity (i.e. the model is driven towards a multivariate random walk) making the visual result less clear.

λ, the results in Figure 2.1 send a similar message, as the contribution of cross-sectional information is progressively penalised by higher values of λ in the 1980s. When the dimension of the VAR increases (Figure 2.2) the optimal value of λ is higher, confirming the theoretical results in De Mol, Giannone, and Reichlin (2008) on the inverse relationship between the optimal level of shrinkage and cross-sectional dimensions in large panels. An inverse U shaped evolution of λ can be detected also in this case.

To verify that time variation in the coefficients is indeed useful for forecasting, we compare the performance of the 20 variables TVP-VAR with that of its constant coefficient counterpart. The results of this exercise are shown for the various model specifications in Table 2.4 where we report relative Root Mean Square Forecast Errors (RMSE). Values below 1, which imply that the introduction of time variation through the kernel estimator induces an improvement in prediction accuracy, are highlighted in grey. We assess the statistical significance in forecast accuracy through a Diebold-Mariano test (Diebold and Mariano, 1995) and *underline* the cases in which the null hypothesis can be rejected at the 10% significance. A bird-eye view of the table reveals that in many instances time variation increases forecast accuracy, as the majority of the cells (around 70% of the cases) report values below 1. However, the average improvement appears to be small as in most of the cases the gain is of the order of 5%. As a consequence, most of the differences in forecast accuracy are not significant, according to the Diebold Mariano test. Looking more in detail, three results emerge. First, time variation matters at long horizons for inflation and interest rates, while for employment the improvement is more consistent across different horizons. Second, the specifications that work best are those in which H is fixed at around 0.7 and λ is optimised in real time according to the L_{mse} criterion (M6 in the Table). In this case the TVP-VAR improves on the constant coefficients benchmark by more than 10% at long-horizons. Third, specifications in which both λ and H are optimised in real time (M3 and M8) do not perform well and, in fact, are often outperformed by the benchmark.

2.5.3 The role of forecast pooling

The substantial heterogeneity observed in the forecasting results across model specifications suggests that the performance of TVP-VAR could be further improved through forecast combination. Since combination schemes based on equal weights are usually found to perform remarkably well, we proceed by pooling forecasts through simple averaging.[20]

The results obtained by forecast pooling are summarised in Figure 2.3. The plots, which show the RMSEs of the combined TVP-VARs relative to the fixed coefficients benchmark, are organised in three panels corresponding to the three different target variables, CPI, Fed Fund Rates and employment. The six bars in each panel correspond to different forecast horizons, from 1 to 24 months ahead. Bars in grey identify the forecast horizons for which a Diebold-Mariano test does not reject the null hypothesis of equal forecast accuracy, while those in red denote the cases for which forecast accuracy is significantly different at the 10% confidence level.[21]

The forecasts obtained by pooling predictions from the different time-varying model specifications prove to be more accurate than those obtained from the benchmark at basically all horizons. Furthermore, according to the Diebold Mariano test, the improvement is statistically significant at the 10% confidence level, as evident from the large prevalence of red bars. There is also a tendency of the relative RMSEs to fall as the forecast horizon increases, as it was already apparent in the results displayed in Table 2.4, suggesting that time variation in the VAR coefficients is relatively more important for forecasting at longer than at shorter horizons. In Figure 2.4 we report the cumulative sum of squared forecast error differentials, computed as

$$CSSED_t = \sum_{j=1}^{t} (e^2_{j,BVAR} - e^2_{j,TVP-VAR}). \tag{2.34}$$

This statistics is very useful in revealing the parts of the forecast sample where the TVP-VAR accrues its gains. Positive and increasing values indicate that the TVP model outperforms the benchmark, while negative and decreasing values suggest the opposite. At relatively shorter horizons (top panels) the model with time-varying coefficients performs better than the one with constant parameters around economic downturns, as indicated by the jumps of the CSSED in periods classified by the NBER as recessions (grey shaded areas). At longer horizons (bottom panels), the gain is relatively uniform across the sample for interest rates and employment, while it is relatively concentrated in the 70s-80s for inflation.[22]

[20] More sophisticated weighting schemes, based on the selection criteria described in Table 2.3, deliver very similar results. The analysis is available upon request.

[21] The test is two sided so that bars in red and higher than 1 indicate that the forecast of the benchmark model is significantly more accurate.

[22] A fluctuation test over windows of 120 months shows that for interest rates and employment the gain of using the TVP-VAR is statistically significant over most of the sample. For inflation, the test rejects the null of equal predictive accuracy more sporadically.

2.5.4 The role of model size

To answer the question of whether enlarging the information set eliminates the need for time variation in the coefficients, we compare the performance of the 20 variables TVP-VAR with that of a fixed coefficients BVAR with 78 variables. The relative RMSEs reported in Figure 2.5 show that at shorter horizons (1 to 6 months ahead) the performance of the two models is overall comparable, although the time-varying model is more accurate in tracking interest rates. However, when we move to longer horizons, the performance of the TVP-VAR improves considerably. Looking at Figure 2.6 we find again that the CSSED tends to jump around recession periods. Hence, the importance of TVP is not (mainly) due to omitted variables.

The next issue that we want to explore is whether, in the context of a TVP-VAR, it pays off to go larger than around 20 variables, provided that the set of variables of interest is small. We tackle this question by comparing the performance of the 20 variables TVP-VAR with that of a 78 variables TVP-VAR. We find that, on the whole forecast sample, a medium-sized information set is sufficient to capture the relevant dynamics. The predictive accuracy of the 20 variables VAR is, in fact, typically higher than that of the larger model, especially for interest rates (see Figure 2.7). The evolution of the CSSED, shown in Figure 2.8, reveals that the accuracy gains of the 20 variables VAR are actually concentrated in the first part of the sample, and that from the 90s onwards, the performance of the two model sizes is very similar. This is an interesting finding that extends to a time-varying coefficients context the results obtained by Banbura, Giannone, and Reichlin (2010) in the case of constant coefficient VARs and those by Boivin and Ng (2006) in constant coefficient factor models.

2.5.5 Allowing for time varying volatilities

As a final exercise we contrast the predictive performance of the baseline 20 variables TVP-VAR with that of a TVP-VAR that also includes the GLS correction described in Section 2.5 to account for the presence of stochastic volatility in the data. For both models point forecasts are obtained by pooling predictions obtained on the basis of different values of ϕ and H. The results reported in Figure 2.9 show that the model that *does not* consider a GLS correction performs definitely better both at short and at longer horizons. Since the differences in terms of RMSEs are quite large, it is all the more important to understand in which part of the sample these gains are attained. The CSSED shown in Figure 2.10 reveal that the GLS correction is particularly detrimental for forecast accuracy in the first part of the sample, and in particular in the late Seventies. Still, in the remaining part of the sample, an upward trend in the CSSED is still well visible, implying that the relative performance of the two models is still in favour of the simpler method where no GLS correction is performed.[23] One reason for the poor performance of the GLS correction might be related to the fact that the kernel estimation of the covariance matrix can be imprecise for models of this size. To check for this possibility we run the exercise on smaller scale VAR including only seven variables. The results, shown in figures 2.11 and 2.12 show that, while in this setting the relative performance of the model with time varying volatilities improves, this GLS correction does not yield any significant gain in forecast accuracy.

2.5.6 Comparison with Koop and Korobilis (2013)

A comparison of the empirical performance of the nonparametric and parametric TVP-VAR needs to take into account the computational limitations to which the latter is subject. This means that a forecast competition based on monthly VARs with 13 lags, like those employed in the previous subsections, is unfeasible. We therefore proceed by taking quarterly transformations of the variables and specify a 20 variables VAR with 4 lags. The forecast exercise is similar to the one performed on monthly data, that is we produce 1 to 8 quarters ahead forecasts of the three key variables in our dataset, CPI, the Fed Fund Rates and payroll employment, with an out sample period ranging from 1970:q1 to 2013:q2 (167 data points).

Figure 2.13 presents the RMSEs of the kernel based estimator relative to those of the parametric one. Again, we use red bars to highlight the cases where a Diebold-Mariano test rejects the null hypothesis of equal forecast accuracy. Visual inspection of the graph reveals that the nonparametric estimator generates significantly better predictions for inflation and employment, while the parametric estimator is more accurate in forecasting short term interest rates. As for the remaining variables, the only case in which the Diebold-Mariano test rejects in favour of the parametric estimator is for the 10 year rate and for M1 at very short horizons, while for the remaining 10 indicators the evidence is either in favour of the nonparametric approach (red bars lower than one) or inconclusive (grey bars).

Summarising, the outcome of this extensive forecasting exercise provides further broad support for our method. Moreover, the fact that the estimator can accommodate a large information set has allowed us to address issues that could not be investigated with existing methods, such as the relationship between the size

[23]Results obtained by evaluating forecasts produced only from the mid Eighties onwards (not reported for brevity but available upon request) confirm this intuition as the simpler model still produces RMSEs that are 25 percent lower than those obtained with the model that allows for a GLS correction.

of the information set and parameters' time variation and the relevance of the model size in the context of models with time-varying parameters.

2.6 Structural analysis

Our non-parametric estimator can be useful also in the context of structural analysis when time variation in the parameters is considered to be an issue. As an illustration, we use the proposed method to estimate the time-varying responses of industrial production indexes to an unexpected increase in the price of oil. The changing response of key macroeconomic variables to unexpected oil price increases has been greatly debated in the past decade. In particular, using structural VARs and different identification assumptions a number of studies have found that oil price increases are associated with smaller losses in U.S. output in more recent years. While some of these studies have used sample-split approaches, like Edelstein and Kilian (2009), Blanchard and Gali (2007) and Blanchard and Riggi (2013), others have relied on Bayesian VARs with drifting coefficients and volatilities, see Baumeister and Peersman (2013) and Hahn and Mestre (2011). The latter approach, however, severely constraints the size of the system to be estimated so that only a small number of variables can be jointly modelled. Partly as a consequence of this constraint, available evidence on the break in the oil/output nexus mainly refers to aggregate GDP. Sectoral aspects, however, are equally relevant as the recessionary effect of oil price shocks is partly due to a costly reallocation of labor and capital away from energy intensive sectors (Davis and Haltiwanger, 2001). Where a more granular perspective is taken, like in Edelstein and Kilian (2009), special attention is paid to the role of the automotive sector, which is considered the main transmission channel of energy price shocks. Indeed as energy price increases reduce purchases of cars, and given that the dollar value of these purchases is large relatively to the energy they use, even small energy price shocks can cause large effects, an intuition formalised by Hamilton (1988). Given the importance of this sector, one would expect it to be the main responsible for the changing relationship between oil and GDP.

In this section we revisit this issue by extending the analysis conducted in Edelstein and Kilian (2009), based on a bivariate VAR and on a sample-split approach, to a large TVP-VAR setting in which energy prices are modelled jointly with industrial output in different sectors. In particular, we augment our baseline 20 variables VAR with 8 industrial production series split by market destination.[24]

[24]Since we are not concerned with forecasting, in the structural analysis presented in this

The additional series are Business Equipment, Consumer Goods and its two sub-components Durable (half of which is accounted for by Automotive products) and Nondurable (Food, Clothing, Chemical and Paper products), Final Product goods (Construction and Business Supplies and Defence and Space equipment), Material goods and its two sub-components Durable (Consumer and Equipment parts) and Nondurable (Textile, Paper and Chemical). A list of the series, together with their weight on the overall index, is reported in Table 2.5.

While the impact of energy prices on macroeconomic variables has typically been studied in small scale structural VARs, recent papers have investigated the issue in models of larger scale, see Stock and Watson (2012) and Aastveit (2014) for applications that use, respectively, a Structural Dynamic Factor Model and a FAVAR. Neither of these two papers, however, allows for time variation in the coefficients.

The identification of energy price shocks follows Edelstein and Kilian (2009), i.e. we assume that energy price shocks are exogenous relative to contemporaneous movements in the other variables in the system, which implies ordering the price of oil first in a recursive structural VAR.[25] Kilian and Vega (2011) provide a test of this assumption by regressing daily changes in the price of oil to daily news on a wide range of macroeconomic data and find no evidence of feedback from macroeconomic news to energy prices, concluding that energy prices are indeed predetermined with respect to the U.S. macroeconomy. A shortcoming of this approach is that it does not allow us to separate the source of variation behind oil price shocks, i.e. whether they are driven by supply rather than by demand.[26] In other words, our identified energy price shocks will be a linear combination of demand and supply shocks. However, given that we are interested in identifying the sectors that are central to the propagation of energy price shocks, rather than in determining the determinants of energy price fluctuations, the recursive identification is a valid working assumption as long as the effects of supply and demand shocks are not disproportionately different across sectors. Furthermore, this final section does not aim at pushing the frontier of the literature on the oil/macroeconomy nexus but rather at illustrating the potential usefulness for structural analysis of a large TVP-VAR.

Figure 2.14 shows that an innovation to the real price of oil generates a protracted

Section we follow Giraitis, Kapetanios, and Yates (2012) and use a two sided Gaussian kernel with smoothing parameter $H = 0.5$. Furthermore, the penalty parameter λ is chosen over the full sample through the L_{fit} criterion and changing volatilities are accounted for through the GLS correction in equation (2.23).

[25] A similar identification assumption is maintained by Blanchard and Gali (2007) and Blanchard and Riggi (2013).

[26] The debate on the relative role of supply and demand factors in determining oil prices dates back to Kilian (2009).

fall in overall industrial output. Furthermore, in line with the literature, the recessionary impact of an exogenous oil price disturbance is generally more severe in the Seventies than in later decades. Notice, however that the difference across the two sub-samples is entirely accounted for by the very early Seventies, a finding that cannot be uncovered with the simple sample-split strategy considered in Edelstein and Kilian (2009), Blanchard and Gali (2007) and Blanchard and Riggi (2013) and that validates the use of time-varying coefficients models.

The results for the individual sectors are reported in Figure 2.15. A number of interesting results emerge. First, in most sectors the effect of an unexpected increase in the real price of oil is generally negative in the first part of the sample, and the fall in production is much more pronounced in energy intensive segments, like Business Equipment, Durable Consumption and Material Goods. Second, most sectors display an attenuation of the recessionary impact of energy price shocks. In some of them unexpected increases in the real price of oil end up being associated with an expansion in production, consistently with the findings in Kilian (2009) that attribute energy price surprises in the 2000s to increased demand for commodities rather than to supply disruptions.[27] Again, most of the changes over time occur in more energy intensive sectors.

To assess the relative importance of each sector in explaining the changing pass-through of energy price shocks to overall industrial activity we proceed by weighing the IRFs in different sectors by their shares in overall industrial output (reported in Table 2.5). The resulting weighted IRFs are reported in Figure 2.16 where, for the sake of clarity, we only focus on the responses twelve months after the initial shock. When the relative weight of the various sectors is taken into account, the relevance for overall business cycle fluctuations of developments in the motor vehicles sector, which accounts for half of Durable Consumption, appears less relevant than that of other sectors. Instead, the response of overall industrial output to oil price shocks and its evolution over time are largely determined by that of the Durable Material sector, which includes intermediate goods for a wide range of final products. This outcome suggests that the increased efficiency in the energy use of automobiles has played a minor role in shaping the oil/output relationship in the U.S. over the past forty years. In turn, greater energy efficiency at the higher stages of the supply chain, as well as a larger role for demand shocks, are likely to be the driving forces behind changes in the relationship between oil prices and U.S. aggregate output.

[27] Blanchard and Gali (2007) also find that oil price innovations are associated with an increase in output after the 80s in France and in Germany, see Figure 7.6 therein.

2.7 Conclusions

In this paper we propose an estimator for large dimensional VAR models with flexible parameter structure, capable of accommodating breaks in the relationships among economic time series. Our procedure is based on the mixed estimator by Theil and Goldberger (1960), which imposes stochastic constraints on the model coefficients, and on the nonparametric VAR estimator proposed by Giraitis, Kapetanios, and Yates (2014). The use of stochastic constraints mimics in a classical context the role of the prior in Bayesian models, allowing to bypass the over-fitting problem that arises in large dimensional models.

We derive the asymptotic distribution of the estimator and evaluate the determinants of its efficiency. We also discuss various aspects of the practical implementation of the estimator, based on two alternative (fit and forecasting) criteria, and assess its finite sample performance in Monte Carlo experiments.

We then use the non-parametric estimator in a forecasting exercise where we model up to 78 U.S. macroeconomic time series. We find that the introduction of time variation in the VAR model parameters yields an improvement in prediction accuracy over models with a constant parameter structure, in particular when forecast combination is used to pool forecasts obtained with models with different degrees of time variation and penalty parameters. We also shed light on an issue that is central to the forecasting literature, namely how the size of the information set interacts with time variation in the model parameters. Specifically, we find that the relevance of time variation is not related to omitted variable problem and that, as in the constant parameter case, a medium-sized TVP-VAR is at least as good as a large TVP-VAR.

In a forecasting context, our non-parametric estimator compares well with the alternative parametric approach by Koop and Korobilis (2013), when using either actual or simulated data, and can handle a larger number of variables.

Finally, to illustrate the use of our method in structural analysis, we analyse the changing effects of oil price shocks on economic activity, a question that has spurred a large number of studies in the applied macro literature in recent years. We find that the declining role of oil prices in shaping U.S. business cycle fluctuations stems from changes related to Business Equipment and Materials sector, rather than from the automobiles sector as argued by part of the literature.

Overall, we believe that our findings illustrate how the econometric tool that we have proposed opens the door to a number of interesting analyses on forecasting and on the nonlinear transmission of shocks, which have been so far constrained by computational issues.

Table 2.1: 1 step ahead, relative RMSEs

T	Parametric	Non parametric	
		Inv. RMSE	Equal weights
DGP-1 (Random walk coefficients)			
n=7			
100	1	1.004	1.005
150	1	0.999	1.000
200	1	0.997	0.997
n=15			
100	1	1.021	1.024
150	1	1.012	1.013
200	1	1.006	1.007
DGP-2 (Occasionally breaking coefficients)			
n=7			
100	1	0.96	0.96
150	1	0.96	0.96
200	1	0.96	0.96
n=15			
100	1	0.96	0.96
150	1	0.95	0.95
200	1	0.94	0.94
DGP-3 (Sine function coefficients)			
n=7			
100	1	0.95	0.96
150	1	0.96	0.98
200	1	0.97	0.99
n=15			
100	1	0.87	0.88
150	1	0.86	0.87
200	1	0.85	0.86

Note to Table 2.1. The table shows the ratio between the one step ahead RMSE attained by, respectively, the nonparametric and the parametric model, averaged across the n variables. Forecasts are computed on the second half of the sample, i.e. when T=100, forecasts are computed recursively for t=51 to t=100, when T=150 forecasts are computed recursively for t=76 to t=150 and when T=200 forecasts are computed recursively for t=101 to t=200.

No.	Acronym (FRED database)	Description	SA	Logs	Prior mean
1	AAA	Interest rates on AAA bonds	Not Seasonally Adjusted	0	1
2	AHEMAN	Average Hourly Earnings Of Production And Nonsupervisory Employees: Manufacturing	Not Seasonally Adjusted	1	1
3	AWHMAN	Average Weekly Hours of Production and Nonsupervisory Employees:	Seasonally Adjusted	1	0
4	AWOTMAN	Average Weekly Overtime Hours of Production and Nonsupervisory	Seasonally Adjusted	1	0
5	BAA	Interest rates on BAA bonds	Not Seasonally Adjusted	1	1
6	CE16OV	Civilian Employment	Seasonally Adjusted	1	1
7	CPIAPPSL	Consumer Price Index for All Urban Consumers: Apparel	Seasonally Adjusted	1	1
8	CPIAUCSL	Consumer Price Index for All Urban Consumers: All Items	Seasonally Adjusted	1	1
9	CPILFESL	Consumer Price Index for All Urban Consumers: All Items Less Food &	Seasonally Adjusted	1	1
10	CPIMEDSL	Consumer Price Index for All Urban Consumers: Medical Care	Seasonally Adjusted	1	1
11	CPITRNSL	Consumer Price Index for All Urban Consumers: Transportation	Seasonally Adjusted	1	1
12	CPIULFSL	Consumer Price Index for All Urban Consumers: All Items Less Food	Seasonally Adjusted	1	1
13	DMANEMP	All Employees: Durable goods	Seasonally Adjusted	1	0
14	DSPIC96	Real Disposable Personal Income	Seasonally Adjusted	1	1
15	DPCERA3M086SBEA	Real personal consumption expenditures (chain-type quantity index)	Seasonally Adjusted	1	1
16	FEDFUNDS	Effective Federal Funds Rate	Not Seasonally Adjusted	0	1
17	GS1	1-Year Treasury Constant Maturity Rate	Not Seasonally Adjusted	0	1
18	GS10	10-Year Treasury Constant Maturity Rate	Not Seasonally Adjusted	0	1
19	GS5	5-Year Treasury Constant Maturity Rate	Not Seasonally Adjusted	0	1
20	HOUST	Housing Starts: Total: New Privately Owned Housing Units Started	Seasonally Adjusted Annual Rate	1	0
21	HOUSTMW	Housing Starts in Midwest Census Region	Seasonally Adjusted Annual Rate	1	0
22	HOUSTNE	Housing Starts in Northeast Census Region	Seasonally Adjusted Annual Rate	1	0
23	HOUSTS	Housing Starts in South Census Region	Seasonally Adjusted Annual Rate	1	0
24	HOUSTW	Housing Starts in West Census Region	Seasonally Adjusted Annual Rate	1	0
25	INDPRO	Industrial Production Index	Seasonally Adjusted	1	1
26	IPBUSEQ	Industrial Production: Business Equipment	Seasonally Adjusted	1	1
27	IPCONGD	Industrial Production: Consumer Goods	Seasonally Adjusted	1	1
28	IPDCONGD	Industrial Production: Durable Consumer Goods	Seasonally Adjusted	1	1
29	IPDMAT	Industrial Production: Durable Materials	Seasonally Adjusted	1	1
30	IPFINAL	Industrial Production: Final Products (Market Group)	Seasonally Adjusted	1	1
31	IPMAT	Industrial Production: Materials	Seasonally Adjusted	1	1
32	IPNCONGD	Industrial Production: Nondurable Consumer Goods	Seasonally Adjusted	1	1
33	IPNMAT	Industrial Production: nondurable Materials	Seasonally Adjusted	1	1
34	LOANS	Loans and Leases in Bank Credit, All Commercial Banks	Seasonally Adjusted	1	1
35	M1SL	M1 Money Stock	Seasonally Adjusted	1	1
36	M2SL	M2 Money Stock	Seasonally Adjusted	1	1
37	MANEMP	All Employees: Manufacturing	Seasonally Adjusted	1	0
38	NAPM	ISM Manufacturing: PMI Composite Index	Seasonally Adjusted	0	0
39	NAPMEI		Seasonally Adjusted	0	0
40	NAPMII	[]	Not Seasonally Adjusted	0	0
41	NAPMNOI	ISM Manufacturing: New Orders Index	Seasonally Adjusted	0	0
42	NAPMPI	[]	Seasonally Adjusted	0	0
43	NAPMSDI	[]	Seasonally Adjusted	0	0
44	NDMANEMP	All Employees: Nondurable goods	Seasonally Adjusted	1	0
45	OILPRICE	[]	Not Seasonally Adjusted	1	1
46	PAYEMS	All Employees: Total nonfarm	Seasonally Adjusted	1	1
47	PCEPI	Personal Consumption Expenditures: Chain-type Price Index	Seasonally Adjusted	1	1
48	PERMIT	New Private Housing Units Authorized by Building Permits	Seasonally Adjusted Annual Rate	1	0
49	PERMITMW	New Private Housing Units Authorized by Building Permits in the	Seasonally Adjusted Annual Rate	1	0
50	PERMITNE	New Private Housing Units Authorized by Building Permits in the	Seasonally Adjusted Annual Rate	1	0
51	PERMITS	New Private Housing Units Authorized by Building Permits in the South	Seasonally Adjusted Annual Rate	1	0
52	PERMITW	New Private Housing Units Authorized by Building Permits in the West	Seasonally Adjusted Annual Rate	1	0
53	PI	Personal Income	Seasonally Adjusted Annual Rate	1	1
54	PPIACO	Producer Price Index: All Commodities	Not Seasonally Adjusted	1	1
55	PPICRM	Producer Price Index: Crude Materials for Further Processing	Seasonally Adjusted	1	1
56	PPIFCG	Producer Price Index: Finished Consumer Goods	Seasonally Adjusted	1	1
57	PPIFGS	Producer Price Index: Finished Goods	Seasonally Adjusted	1	1
58	PPIITM	Producer Price Index: Intermediate Materials: Supplies & Components	Seasonally Adjusted	1	1
59	SandP	S&P 500 Stock Price Index		1	1
60	SRVPRD	All Employees: Service-Providing Industries	Seasonally Adjusted	1	1
61	TB3MS	3-Month Treasury Bill: Secondary Market Rate	Not Seasonally Adjusted	0	1
62	TB6MS	6-Month Treasury Bill: Secondary Market Rate	Not Seasonally Adjusted	0	1
63	UEMP15OV	Number of Civilians Unemployed for 15 Weeks & Over	Seasonally Adjusted	1	0
64	UEMP15T26	Number of Civilians Unemployed for 15 to 26 Weeks	Seasonally Adjusted	1	0
65	UEMP27OV	Number of Civilians Unemployed for 27 Weeks and Over	Seasonally Adjusted	1	0
66	UEMP5TO14	Number of Civilians Unemployed for 5 to 14 Weeks	Seasonally Adjusted	1	0
67	UEMPLT5	Number of Civilians Unemployed - Less Than 5 Weeks	Seasonally Adjusted	1	0
68	UEMPMEAN	Average (Mean) Duration of Unemployment	Seasonally Adjusted	1	0
69	UNRATE	Civilian Unemployment Rate	Seasonally Adjusted	0	0
70	USCONS	All Employees: Construction	Seasonally Adjusted	1	1
71	USFIRE	All Employees: Financial Activities	Seasonally Adjusted	1	1
72	USGOOD	All Employees: Goods-Producing Industries	Seasonally Adjusted	1	0
73	USGOVT	All Employees: Government	Seasonally Adjusted	1	1
74	USMINE	All Employees: Mining and logging	Seasonally Adjusted	1	0
75	USPRIV	All Employees: Total Private Industries	Seasonally Adjusted	1	1
76	USTPU	All Employees: Trade, Transportation & Utilities	Seasonally Adjusted	1	1
77	USTRADE	All Employees: Retail Trade	Seasonally Adjusted	1	1
78	USWTRADE	All Employees: Wholesale Trade	Seasonally Adjusted	1	1

Table 2.2: Data description

		λ	H
L_{fit}	M1	Optimized at each t	0.5 0.6 0.7 0.8 0.9 1
	M2	Fixed at pre-sample optimal level	0.5 0.6 0.7 0.8 0.9 1
	M3	Optimized at each t	Optimized at each t
	M4	Fixed at pre-sample optimal level	Optimized at each t
L_{mse}	M6	Optimized at each t	0.5 0.6 0.7 0.8 0.9 1
	M7	Fixed at pre-sample optimal level	0.5 0.6 0.7 0.8 0.9 1
	M8	Optimized at each t	Optimized at each t
	M9	Fixed at pre-sample optimal level	Optimized at each t

Table 2.3: Specifications for the TVP-VARs

Specifications for the TVP-VARs. The L_{fit} criterion is computed as $L_{fit}(\lambda, H) = \left|\sum_{i=1}^{n_1} \frac{rss_n^i(\lambda,H)}{var(y_{t,i})} - \sum_i^{n_1} \frac{rss_{n1}^i}{var(y_{t,i})}\right|$ where n_1 is a number of reference variables, rss_{n1} is the residual sum of squares obtained with an $n1$ variate VAR, and $rss_{n1}(\lambda, H)$ is the residual sum of squares obtained with the TVP-VAR. The L_{mse} criterion is computed as $L_{mse}(\lambda, H) = \sum_{i=1}^{n_1} \frac{mse_h^i(\lambda,H)}{var(y_{t,i})}$, where mse_h is the mean square prediction error h steps ahead obtained with the TVP-VAR.

LARGE TVP-VAR

Selection method		λ	H	CPI				Fed Funds Rates				Employment			
				Forecast horizon				Forecast horizon				Forecast horizon			
				1	6	12	24	1	6	12	24	1	6	12	24
L_{fit}	M1	Optimized	0.5	1.17	1.10	1.03	0.92	1.07	0.85	0.84	0.94	1.40	1.33	1.13	0.90
			0.6	1.12	1.01	1.02	0.99	1.02	0.90	0.93	0.96	1.28	1.25	1.12	0.90
			0.7	1.03	0.97	0.99	1.03	0.97	0.97	0.99	0.98	0.96	0.93	0.95	0.96
			0.8	1.01	0.98	0.99	0.97	0.99	0.97	0.96	0.95	0.97	0.94	0.95	0.95
			0.9	1.02	0.98	0.98	0.96	1.00	0.96	0.95	0.94	0.98	0.96	0.97	0.96
			1	1.02	0.98	0.98	0.95	1.00	0.96	0.95	0.93	0.98	0.97	0.98	0.96
	M2	Fixed	0.5	1.03	1.04	1.08	1.33	0.95	0.97	1.03	1.27	0.97	0.94	0.98	0.98
			0.6	1.00	0.97	1.03	1.13	0.93	1.02	1.07	1.12	0.97	0.94	0.98	1.04
			0.7	1.00	0.98	1.01	1.05	0.95	0.99	1.01	0.99	0.97	0.95	0.98	1.01
			0.8	1.01	0.99	0.99	0.99	0.97	0.98	0.97	0.96	0.97	0.96	0.97	0.98
			0.9	1.02	0.99	0.98	0.96	0.98	0.97	0.95	0.94	0.98	0.97	0.98	0.98
			1	1.02	0.99	0.98	0.95	0.98	0.96	0.95	0.94	0.98	0.98	0.99	0.98
	M3	Optimized	OPT	1.02	0.99	1.01	1.02	0.99	1.01	1.00	0.95	0.96	0.93	0.94	0.92
	M4	Fixed	OPT	1.01	0.98	1.01	1.03	0.96	1.01	1.01	0.96	0.98	0.97	1.00	1.02
L_{mse}	M6	Optimized	0.5	1.06	1.09	1.10	1.62	1.01	0.98	0.98	1.49	1.16	1.09	1.07	1.06
			0.6	1.08	1.00	0.89	0.84	0.99	0.96	0.90	0.98	1.12	1.05	0.98	0.93
			0.7	1.12	1.03	0.95	0.88	1.01	0.91	0.87	0.89	1.16	1.10	0.99	0.88
			0.8	1.15	1.06	0.98	0.88	1.01	0.88	0.84	0.85	1.18	1.13	1.00	0.84
			0.9	1.14	1.07	0.98	0.88	1.02	0.87	0.84	0.86	1.18	1.15	1.03	0.88
			1	1.14	1.05	0.98	0.88	1.02	0.87	0.84	0.86	1.19	1.16	1.04	0.90
	M7	Fixed	0.5	1.01	1.08	1.26	2.13	0.93	1.02	1.16	2.03	0.99	0.97	1.02	1.15
			0.6	1.01	0.97	1.02	1.11	0.93	1.00	1.05	1.11	0.97	0.93	0.96	1.01
			0.7	1.01	0.97	0.97	0.99	0.95	0.97	0.97	0.96	0.96	0.93	0.94	0.92
			0.8	1.03	0.98	0.96	0.93	0.98	0.94	0.93	0.92	0.98	0.95	0.94	0.89
			0.9	1.03	0.98	0.96	0.92	1.00	0.94	0.92	0.91	1.00	0.98	0.97	0.92
			1	1.04	0.98	0.95	0.92	1.00	0.93	0.92	0.91	1.00	0.99	0.98	0.92
	M8	Optimized	OPT	1.10	1.10	1.11	1.61	1.01	0.97	0.93	1.42	1.15	1.06	1.00	1.02
	M9	Fixed	OPT	0.99	1.04	1.16	2.08	0.95	1.00	1.02	1.99	0.99	0.92	0.95	1.10

Table 2.4: RMSE, TVP-VAR versus BVAR with 20 variables

Root Mean Square Forecast Errors: TVP-VARs versus constant coefficients BVAR (20 variables). The tables show the RMSE obtained by models with time varying parameters described in Table 2.3 relative to those obtained with the benchmark large BVAR with constant parameters. Values below 1 (shaded in grey in the table) imply that the model outperforms the benchmark. Values underlined indicate the cases in which the Diebold Mariano test rejects the null hypothesis of equal forecast accuracy at the 10% confidence level. The RMSE are computed on 499 out-of-sample forecast errors, from January 1970 to July 2013.

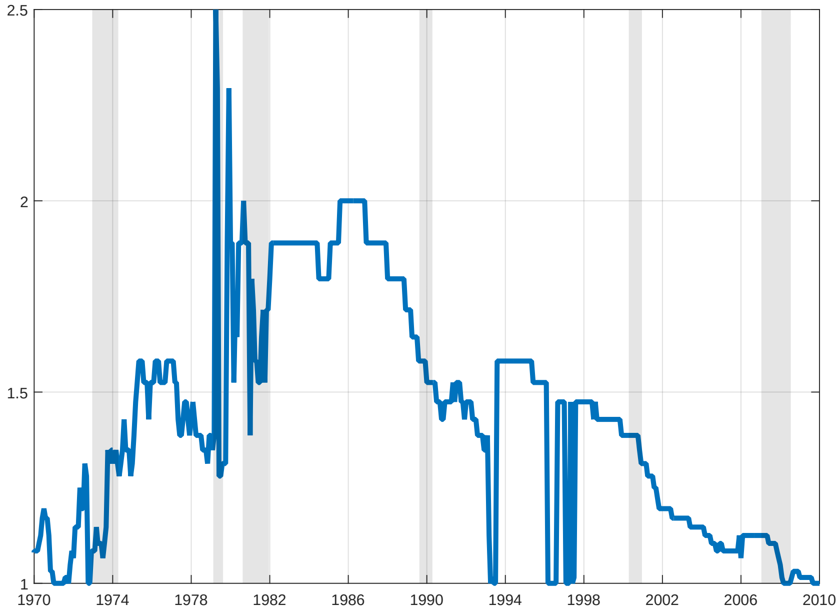

Figure 2.1: Optimal λ - 20 variables TVP-VAR

Optimal λ - 20 variables TVP-VAR. The figure shows the evolution of the value of λ optimized using the L_{fit} criterion in the TVP-VAR with 20 variables. Shaded areas indicate NBER-dated recessions.

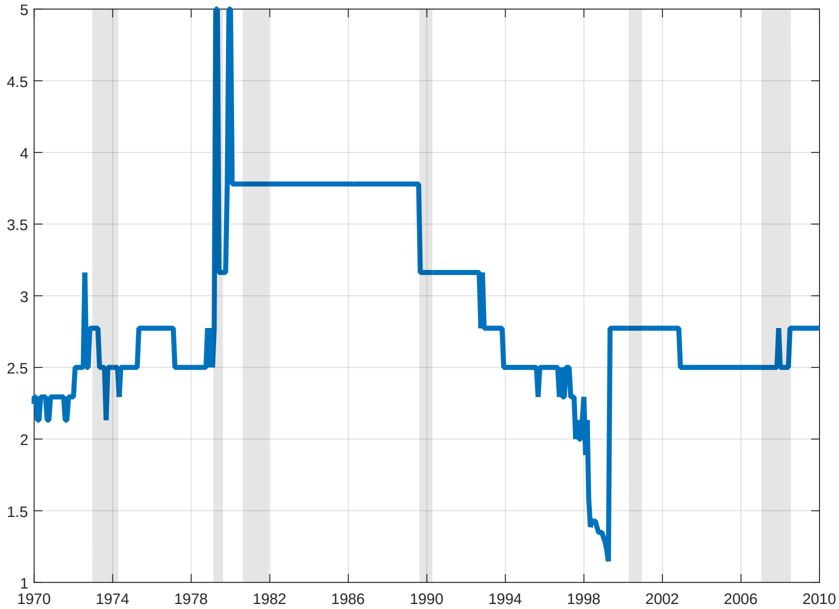

Figure 2.2: Optimal λ - 78 variables TVP-VAR

Optimal λ - 78 variables TVP-VAR. The figure shows the evolution of the value of λ optimized using the L_{fit} criterion in the TVP-VAR with 78 variables. Shaded areas indicate NBER-dated recessions.

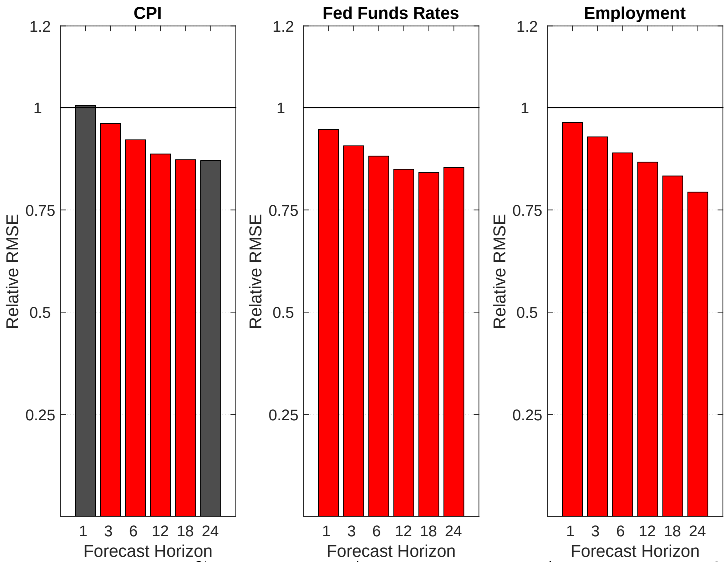

Figure 2.3: RMSE, TVP-VAR versus BVAR, 20 variables

Root Mean Square Forecast Errors: combined TVP-VARs versus constant coefficients BVAR (20 variables VARs). The bar plots show the RMSE obtained by equal weights forecast combination of models with time varying parameters relative to that obtained with the BVAR with constant coefficients. Values below 1 imply that the TVP model outperforms the benchmark. Bars in grey indicate the forecast horizons for which a Diebold-Mariano test does not reject the null hypothesis of equal forecast accuracy, those in red denote the cases for which forecast accuracy is significantly different at the 10% confidence level.

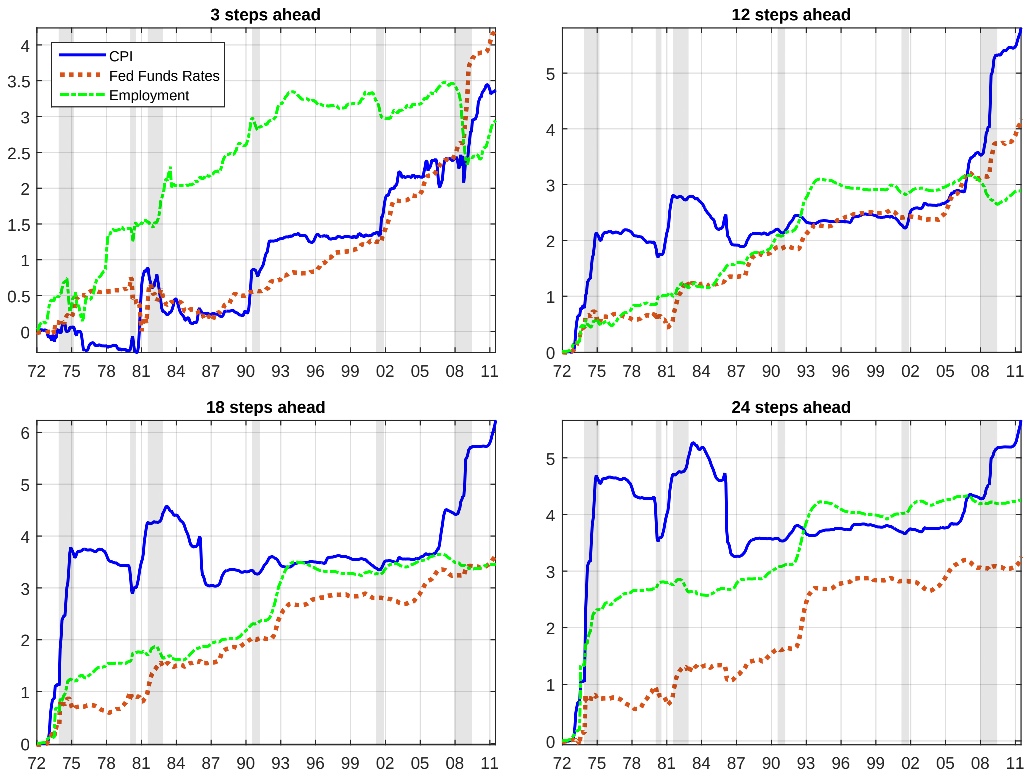

Figure 2.4: CSSED, TVP-VAR versus BVAR, 20 variables

Cumulative sum of squared forecast error differentials: combined TVP-VARs versus constant coefficients BVAR (20 variables VARs). The figure shows the Cumulative Sum of Squared Forecast Errors Differentials between the equal weights forecast combination of models with time varying parameters and the BVAR with constant coefficients. Positive and increasing values indicate that the TVP model outperforms the benchmark, while negative and decreasing values suggest the opposite. Shaded areas indicated NBER-dated recessions.

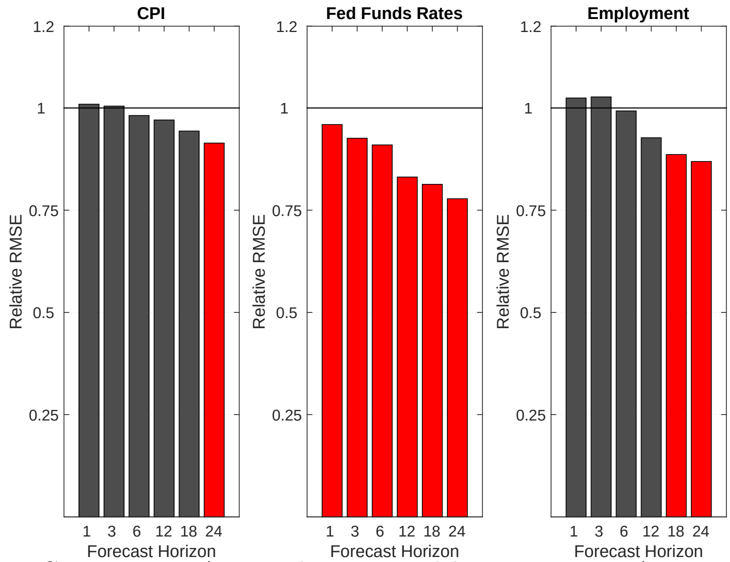

Figure 2.5: RMSE, TVP-VAR with 20 variables versus BVAR with 78 variables
Root Mean Square Forecast Errors: 20 variables combined TVP-VARs versus 78 variables constant coefficients BVAR. The bar plots show the RMSE obtained by equal weights forecast combination of 20 variables VARs with time varying parameters relative to that obtained with a 78 variables BVAR with constant coefficients. Values below 1 imply that the TVP model outperforms the benchmark. Bars in grey indicate the forecast horizons for which a Diebold-Mariano test does not reject the null hypothesis of equal forecast accuracy, those in red denote the cases for which forecast accuracy is significantly different at the 10% confidence level.

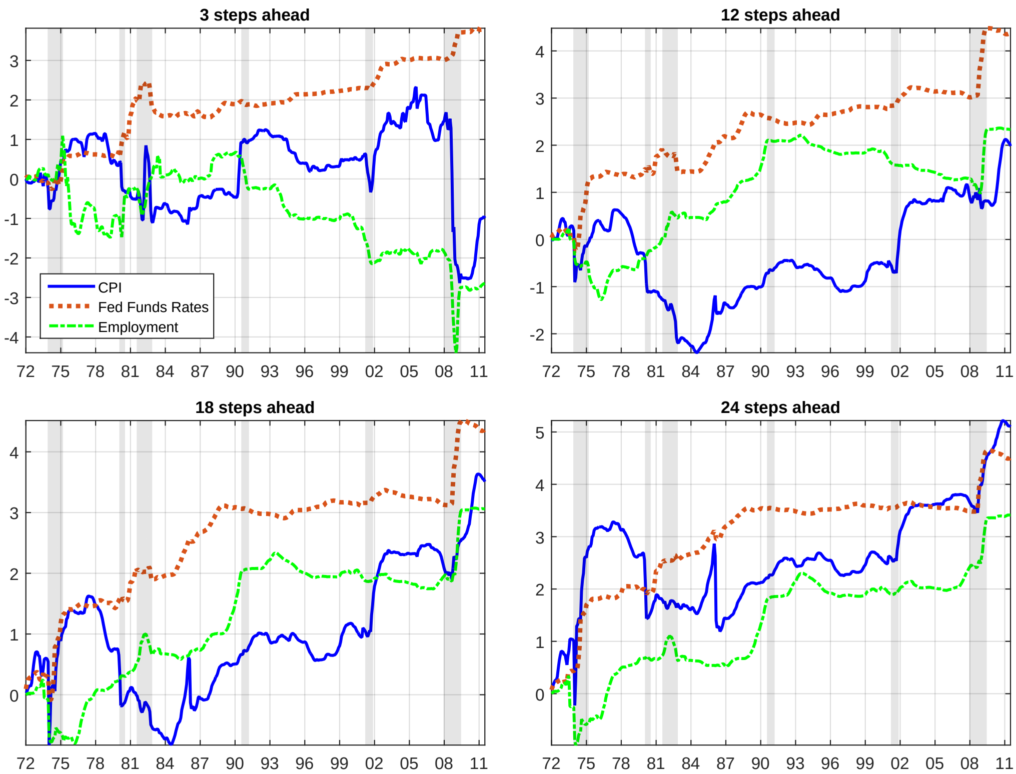

Figure 2.6: CSSED, TVP-VAR with 20 variables versus BVAR with 78 variables
Cumulative sum of squared forecast error differentials: 20 variables combined TVP-VARs versus 78 variables constant coefficients BVAR. The figure shows the Cumulative Sum of Squared Forecast Errors Differentials between the equal weights forecast combination of 20 variables VARs with time varying parameters and a 78 variables BVAR with constant coefficients. Positive and increasing values indicate that the TVP model outperforms the benchmark, while negative and decreasing values suggest the opposite. Shaded areas indicated NBER-dated recessions.

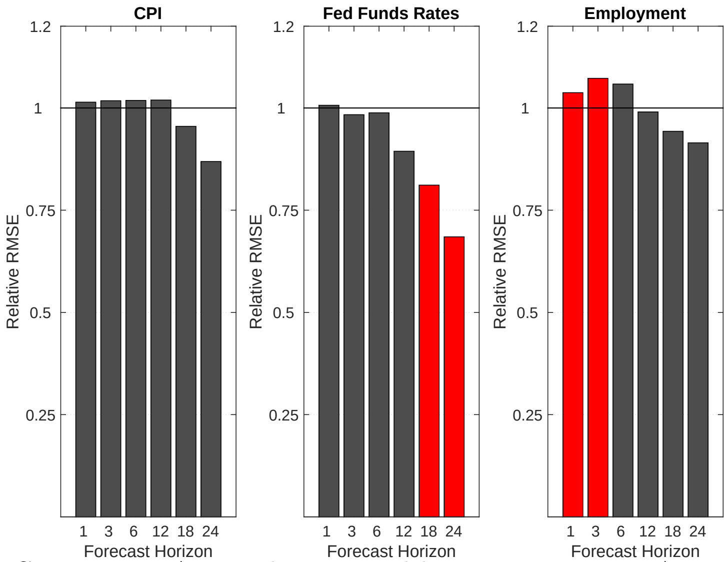

Figure 2.7: RMSE, TVP-VAR with 20 variables versus TVP-VAR with 78 variables
Root Mean Square Forecast Errors: 20 variables combined TVP-VARs versus 78 variables combined TVP-VARs. The bar plots show the RMSE obtained by equal weights forecast combination of 20 variables VARs with time varying parameters relative to that obtained by equal weights forecast combination of 78 variables VARs with time varying parameters. Values below 1 imply that the TVP model with 20 variables outperforms the benchmark. Bars in grey indicate the forecast horizons for which a Diebold-Mariano test does not reject the null hypothesis of equal forecast accuracy, those in red denote the cases for which forecast accuracy is significantly different at the 10% confidence level.

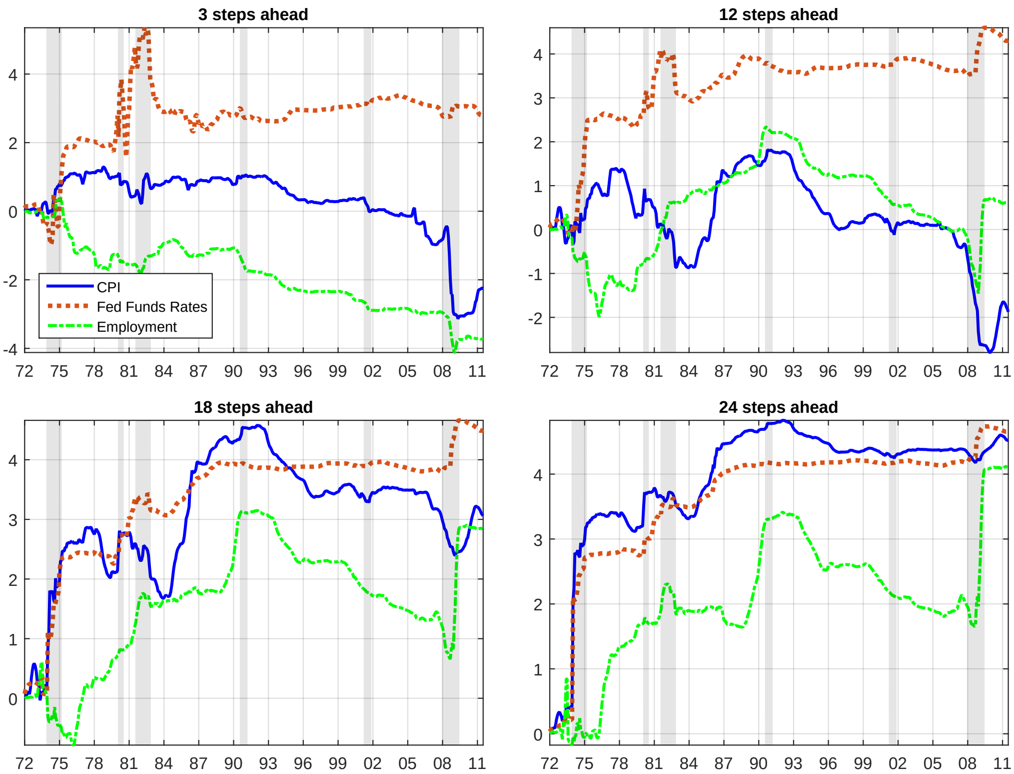

Figure 2.8: CSSED, TVP-VAR with 20 variables versus TVP-VAR with 78 variables
Cumulative sum of squared forecast error differentials: 20 variables combined TVP-VARs versus 78 variables combined TVP-VARs. The figure shows the Cumulative Sum of Squared Forecast Errors Differentials between the equal weights forecast combination of 20 variables VARs with time varying parameters and equal weights forecast combination of 78 variables VARs with time varying parameters. Positive and increasing values indicate that the TVP model with 20 variables outperforms the benchmark, while negative and decreasing values suggest the opposite. Shaded areas indicated NBER-dated recessions.

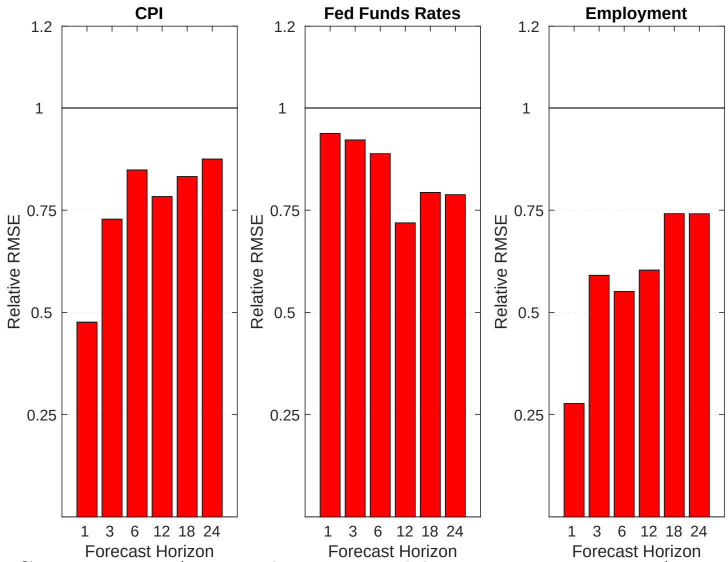

Figure 2.9: RMSE, TVP-VAR with 20 variables versus TVP-VAR with 20 variables and GLS correction

Root Mean Square Forecast Errors: 20 variables combined TVP-VARs versus 20 variables combined TVP-VARs with a GLS correction. The bar plots show the RMSE obtained by equal weights forecast combination of 20 variables VARs with time varying parameters relative to that obtained by equal weights forecast combination of 20 variables VARs with time varying parameters and the GLS correction described in Section 2.5. Values below 1 imply that the TVP model with 20 variables without GLS correction outperforms the benchmark. Bars in grey indicate the forecast horizons for which a Diebold-Mariano test does not reject the null hypothesis of equal forecast accuracy, those in red denote the cases for which forecast accuracy is significantly different at the 10% confidence level.

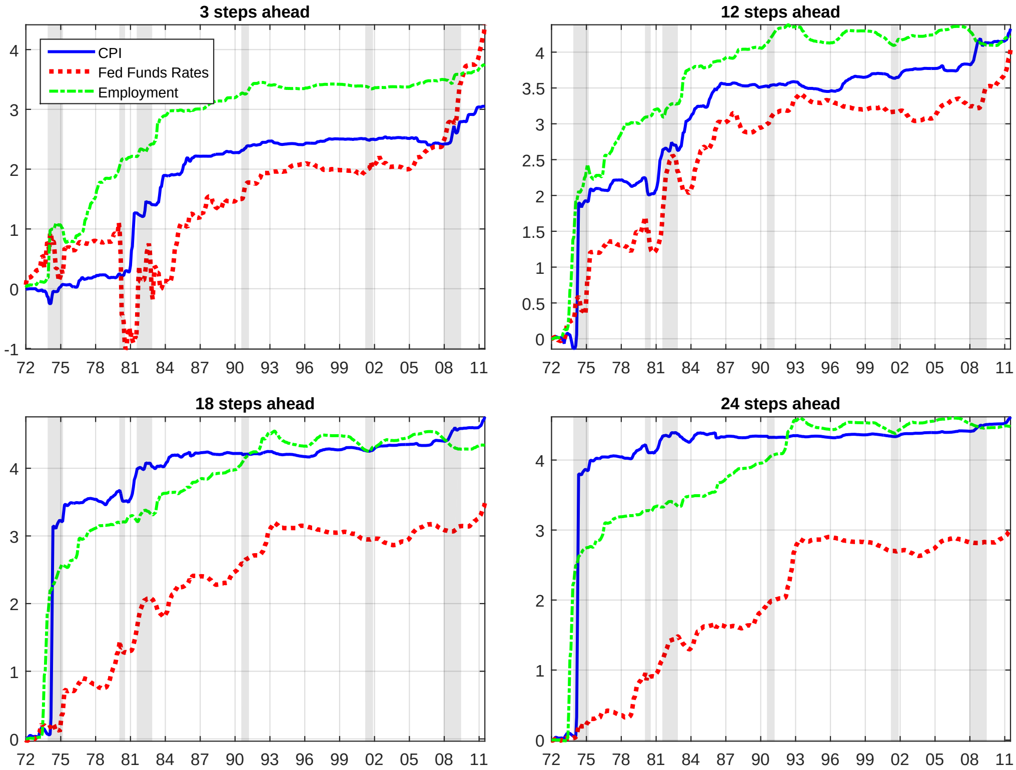

Figure 2.10: CSSED, TVP-VAR with 20 variables versus TVP-VAR with 20 variables and GLS correction

Cumulative sum of squared forecast error differentials: 20 variables combined TVP-VARs versus 20 variables combined TVP-VARs with a GLS correction. The figure shows the Cumulative Sum of Squared Forecast Errors Differentials between the equal weights forecast combination of 20 variables VARs with time varying parameters and equal weights forecast combination of 78 variables VARs with time varying parameters and the GLS correction described in Section 2.5. Positive

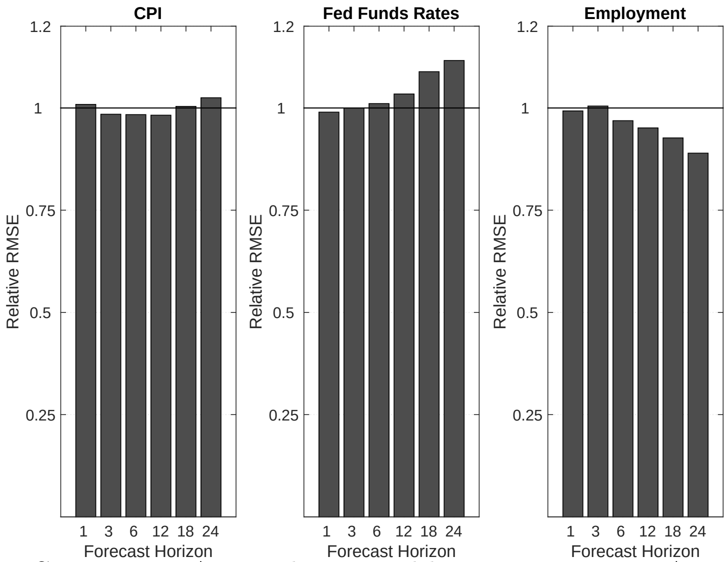

Figure 2.11: RMSE, TVP-VAR with 7 variables versus TVP-VAR with 7 variables and GLS correction

Root Mean Square Forecast Errors: 7 variables combined TVP-VARs versus 7 variables combined TVP-VARs with a GLS correction. The bar plots show the RMSE obtained by equal weights forecast combination of 7 variables VARs with time varying parameters relative to that obtained by equal weights forecast combination of 7 variables VARs with time varying parameters and the GLS correction described in Section 2.5. Values below 1 imply that the TVP model with 7 variables without GLS correction outperforms the benchmark. Bars in grey indicate the forecast horizons for which a Diebold-Mariano test does not reject the null hypothesis of equal forecast accuracy, those in red denote the cases for which forecast accuracy is significantly different at the 10% confidence level.

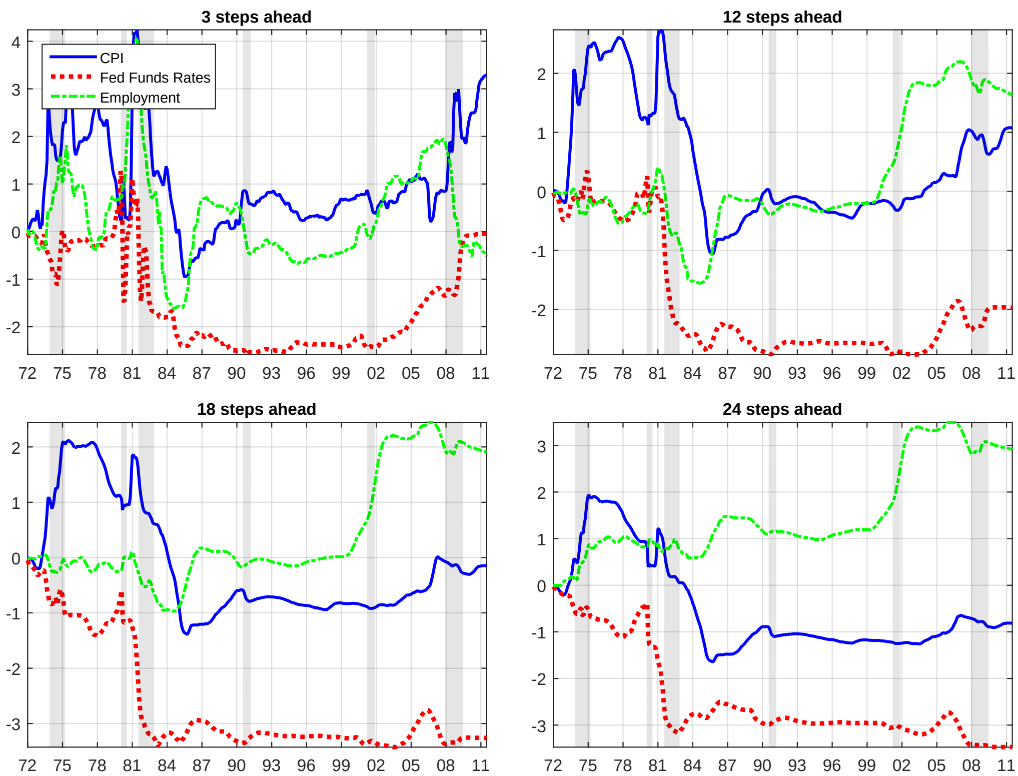

Figure 2.12: CSSED, TVP-VAR with 7 variables versus TVP-VAR with 7 variables and GLS correction

Cumulative sum of squared forecast error differentials: 7 variables combined TVP-VARs versus 7 variables combined TVP-VARs with a GLS correction. The figure shows the Cumulative Sum of Squared Forecast Errors Differentials between the equal weights forecast combination of 7 variables VARs with time varying parameters and equal weights forecast combination of 7 variables VARs with time varying parameters and the GLS correction described in Section 2.5. Positive and

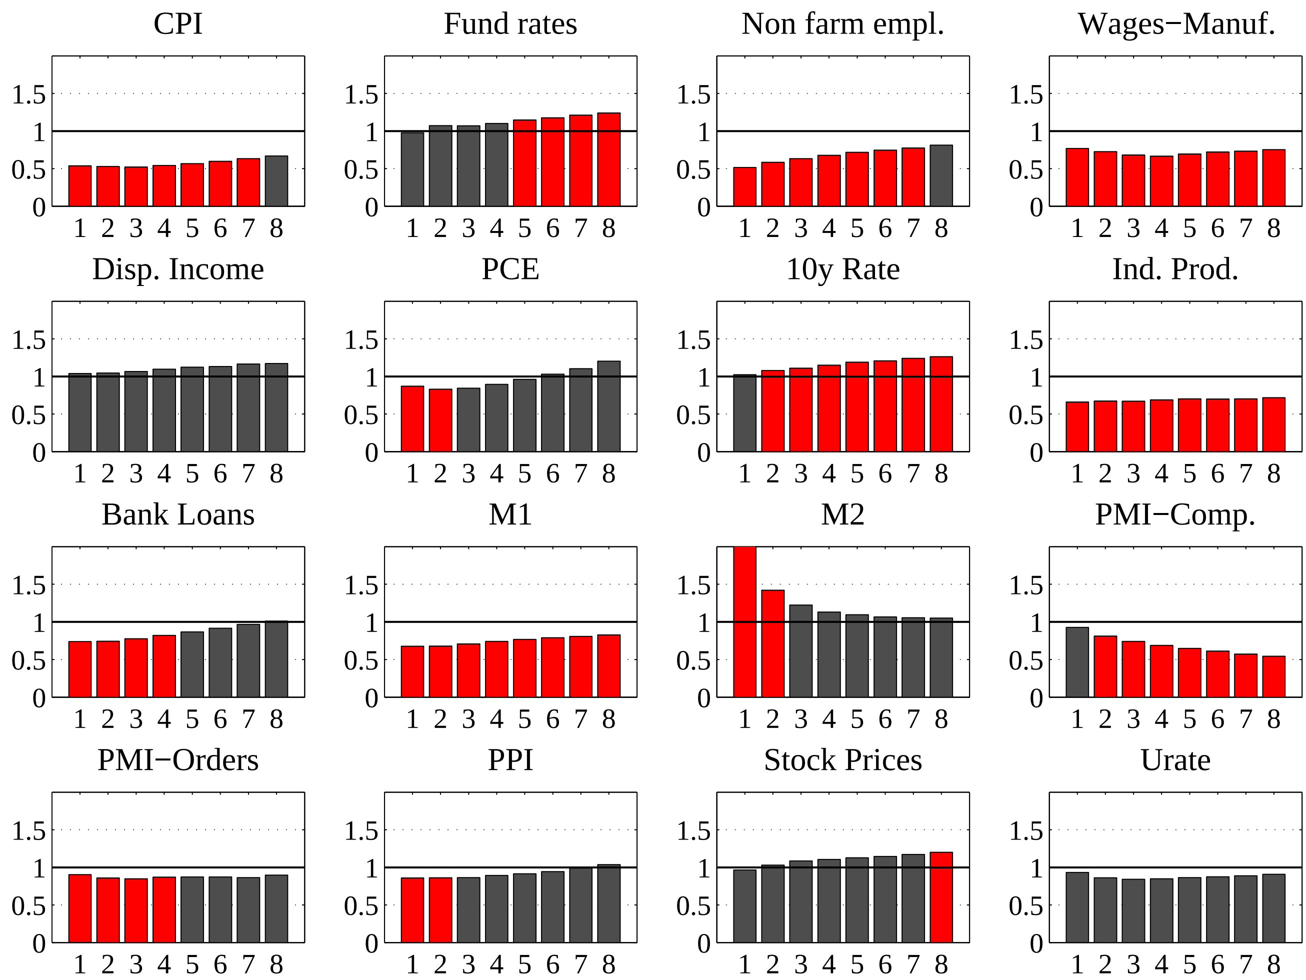

Figure 2.13: Forecast accuracy, nonparametric and parametric estimators

Note to Figure 2.13. The bar plots show the ratio between the RMSE attained by, respectively, the nonparametric and the parametric model. Values below 1 imply that the nonparametric model outperforms the parametric one. Bars in grey indicate that the Diebold-Mariano test does not reject the null hypothesis of equal forecast accuracy, while those in red denote the cases for which forecast accuracy is significantly different at the 10% confidence level.

Market group	Acronym	Weight
Industrial Production Index	INDPRO	100
Industrial Production: Business Equipment	IPBUSEQ	9.18
Industrial Production: Consumer Goods	IPCONGD	27.2
Industrial Production: Durable Consumer Goods	*IPDCONGD*	*5.59*
Industrial Production: Nondurable Consumer Goods	*IPNCONGD*	*21.62*
Industrial Production: Final Products (Market Group)	IPFINAL	16.58
Industrial Production: Materials	IPMAT	47.03
Industrial Production: Durable Materials	*IPDMAT*	*17.34*
Industrial Production: Nondurable Materials	*IPNMAT*	*11.44*

Table 2.5: Industrial production indexes by market group

Note to Table 2.5. The shares of market groups refer to 2011 Value added in nominal terms. Nondurable consumer goods includes Consumer Energy products, which account for 5.7% of total IP. We have excluded from the analysis Industrial Production of Energy Materials, which is part of the Materials (IPMAT) group and accounts for 18.3% of overall output. Source http://www.federalreserve.gov/releases/g17/g17tab1.txt

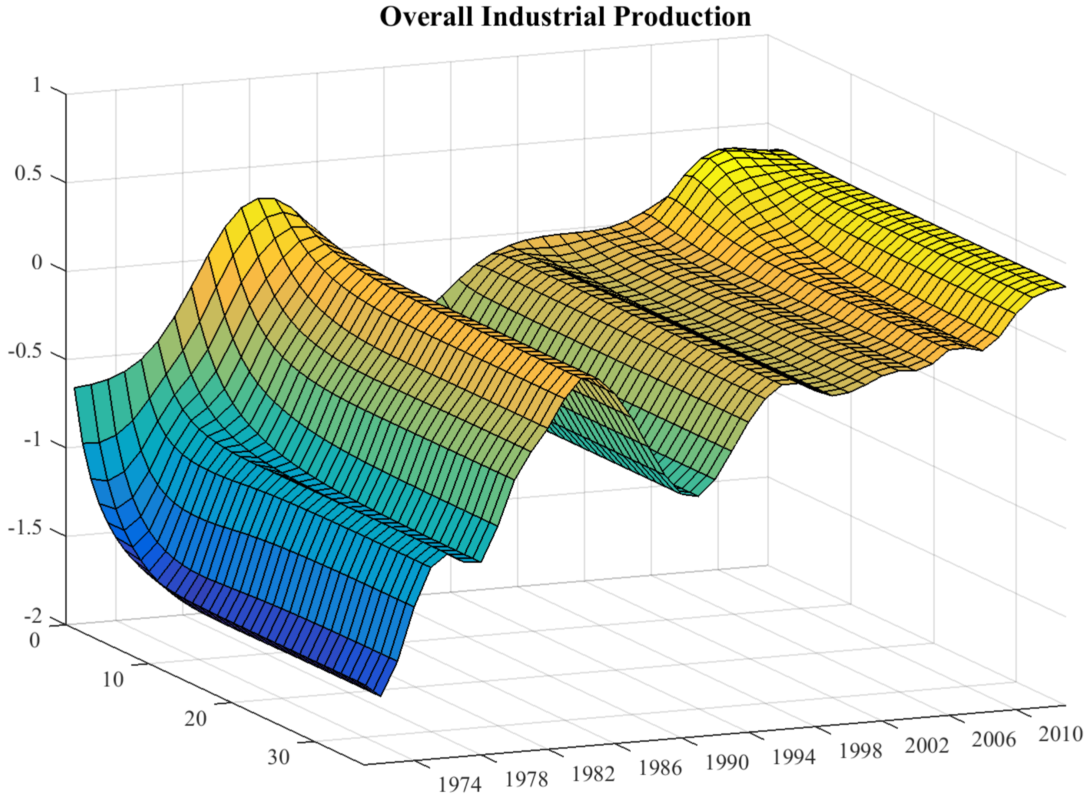

Figure 2.14: Response of Industrial production (overall index) to a 1% shock to the real price of oil

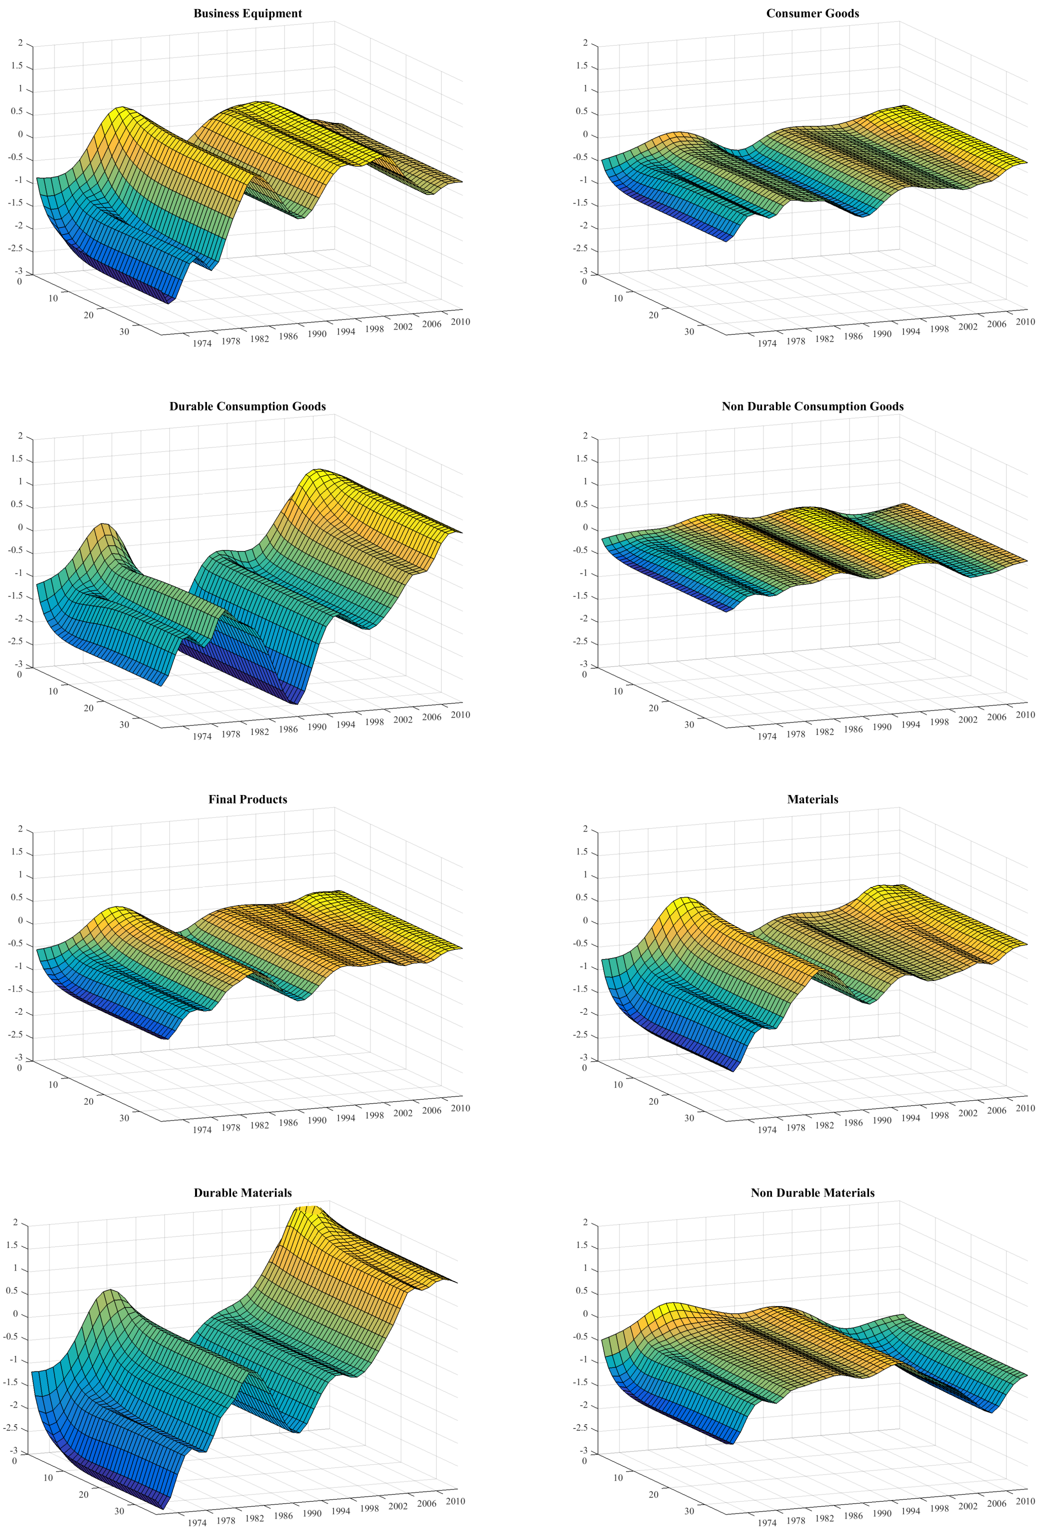

Figure 2.15: Response of Industrial production (sectors) to a 1% shock to the real price of oil

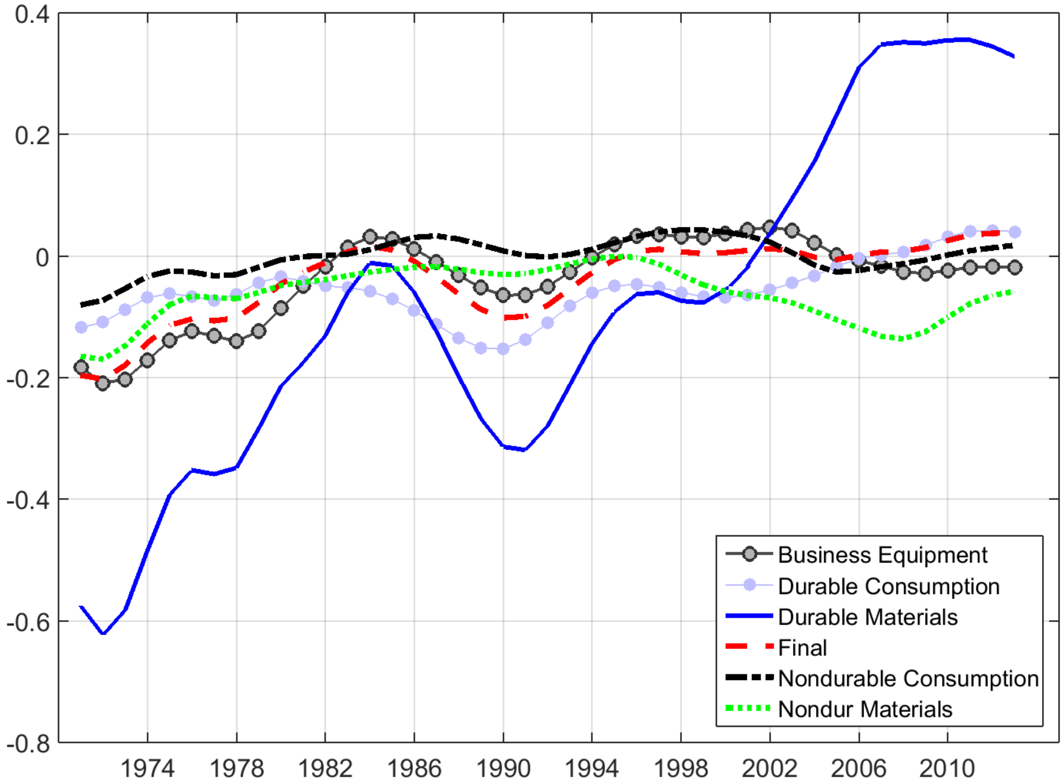

Figure 2.16: Contribution of selected sectors to the response of overall Industrial production (12 months out) to a 1% shock to the real price of oil

Appendix

2.A Proof of Theorem 1

Starting from (2.18), we have:

$$\begin{aligned}\widehat{\Theta}_t &= (S_w^{-1}X'_{w,t}X_{w,t})\widehat{\Theta}_{t,GKY} + S_w^{-1}(\lambda\overline{R}'\overline{R})\Theta_C & (2.35)\\ &= (S_w^{-1}X'_{w,t}X_{w,t})\widehat{\Theta}_{t,GKY} + S_w^{-1}(\lambda\overline{R}'\overline{R})\Theta_t + S_w^{-1}\sqrt{\lambda}\overline{R}'\overline{u}_t^r & (2.36)\end{aligned}$$

where we have used the fact that

$$\begin{aligned}\Theta_C &= (\overline{R}'\overline{R})^{-1}(\overline{R}'\overline{r}) & (2.37)\\ &= (\overline{R}'\overline{R})^{-1}\left(\overline{R}'\left(\overline{R}\Theta_t + \frac{1}{\sqrt{\lambda}}\overline{u}_t^r\right)\right) & (2.38)\\ &= \Theta_t + (\overline{R}'\overline{R})^{-1}\overline{R}'\frac{1}{\sqrt{\lambda}}\overline{u}_t^r & (2.39)\end{aligned}$$

Taking probability limits, recalling that $p\lim\widehat{\Theta}_{t,GKY} = \Theta_t$, and that $S_w = (X'_{w,t}X_{w,t} + \lambda\overline{R}'\overline{R})$ we have that:

$$p\lim\,(S_w^{-1}X'_{w,t}X_{w,t})\widehat{\Theta}_{t,GKY} + S_w^{-1}(\lambda\overline{R}'\overline{R})\Theta_t + S_w^{-1}\overline{R}'\lambda\overline{u}_t^r \;=\; \Theta_t + \Theta_t^B$$

To determine the normalising factor in (2.19), let us go back to the representation in (2.14) and let us take differences from the true parameter matrix Θ_t and from the bias term Θ_t^B. We obtain:

$$\begin{aligned}\widehat{\Theta}_t - \Theta_t - \Theta_t^B &= \left(\sum_{j=1}^T w_{j,t}x_j^*x_j^{*\prime}\right)^{-1}\left(\sum_{j=1}^T w_{j,t}x_j^*y_j^*\right) - \Theta_t - \Theta_t^B\\ &= \left(\sum_{j=1}^T w_{j,t}x_j^*x_j^{*\prime}\right)^{-1}\left(\sum_{j=1}^T w_{j,t}x_j^*(x_j^{*\prime}\Theta_j + u_j^*)\right) - \Theta_t - \Theta_t^B\\ &= \left(\sum_{j=1}^T w_{j,t}x_j^*x_j^{*\prime}\right)^{-1}\left(\sum_{j=1}^T w_{j,t}x_j^*x_j^{*\prime}\Theta_j\right) +\end{aligned}$$

$$+\left(\sum_{j=1}^{T} w_{j,t}x_j^*x_j^{*\prime}\right)^{-1}\left(\sum_{j=1}^{T} w_{j,t}x_j^*u_j^*\right)-\Theta_t-\Theta_t^B$$
$$=\left(\sum_{j=1}^{T} w_{j,t}x_j^*x_j^{*\prime}\right)^{-1}\sum_{j=1}^{T} w_{j,t}x_j^*x_j^{*\prime}\left(\Theta_j-\Theta_t\right)+$$
$$+\left(\sum_{j=1}^{T} w_{j,t}x_j^*x_j^{*\prime}\right)^{-1}\left(\sum_{j=1}^{T} w_{j,t}x_j^*u_j^*-\sum_{j=1}^{T} w_{j,t}x_j^*x_j^{*\prime}\Theta_t^B\right). \quad (2.40)$$

Now, if the bandwidth is $o(T^{1/2})$, then the term $\left(\sum_{j=1}^{T} w_{j,t}x_j^*x_j^{*\prime}\right)^{-1}\sum_{j=1}^{T} w_{j,t}x_j^*x_j^{*\prime}(\Theta_j-\Theta_t)$ is asymptotically negligible and we can focus on the element in (2.40).

First let us simplify the notation and let us write:

1. $\underbrace{X_{w,t}^{*\prime}X_{w,t}^{*}}_{k\times k}\equiv\underbrace{\sum_{j=1}^{T} w_{jt}\underbrace{x_j^*}_{k\times(k+1)}\underbrace{x_j^{*\prime}}_{(k+1)\times k}}_{k\times k}$

2. $\underbrace{\underbrace{X_{ww,t}^{*\prime}}_{k\times T(k+1)}\underbrace{U^*}_{T(k+1)\times n}}_{k\times n}\equiv\sum_{j=1}^{T} w_{j,t}\underbrace{\underbrace{x_j^*}_{k\times(k+1)}\underbrace{u_j^*}_{(k+1)\times n}}_{k\times n}$, where the $T(k+1)\times n$ matrix U^* is obtained by stacking over t the matrices $\underbrace{u_t^*}_{(k+1)\times n}$

3. $\underbrace{\Lambda^{\prime}}_{k\times k}\equiv\underbrace{(X_{w,t}^{*\prime}X_{w,t}^{*})}_{k\times k}\underbrace{S_w^{-1}}_{k\times k}\lambda\underbrace{\overline{R}^{\prime}}_{k\times k}$, which implies that $(X_{w,t}^{*\prime}X_{w,t}^{*})\Theta_t^B=(X_{w,t}^{*\prime}X_{w,t}^{*})S_w^{-1}\lambda\overline{R}^{\prime}\overline{u}_t^r=\Lambda^{\prime}\overline{u}_t^r$

Now, let us multiply the transpose of $\left(\widehat{\Theta}_t^{\prime}-\Theta_t^{\prime}-\Theta_t^{\prime B}\right)$ by $\sqrt{H}$, and let us focus on (2.40):

$$\sqrt{H}\left(\widehat{\Theta}_t^{\prime}-\Theta_t^{\prime}-\Theta_t^{\prime B}\right)=\left(\frac{1}{\sqrt{H}}\left(U^{\prime *}X_{ww,t}^{*}-\overline{u}_t^{\prime r}\Lambda\right)\right)\left(\frac{1}{H}X_{w,t}^{*\prime}X_{w,t}^{*}\right)^{-1}.$$

Taking vec of both sides yields:

$$\sqrt{H}vec\left(\widehat{\Theta}_t^{\prime}-\Theta_t^{\prime}-\Theta_t^{\prime B}\right)=\left(H\left(X_{w,t}^{*\prime}X_{w,t}^{*}\right)^{-1}\otimes I_n\right)vec\left(\frac{1}{\sqrt{H}}\left(U^{\prime *}X_{ww,t}^{*}-\overline{u}_t^{\prime r}\Lambda\right)\right). \quad (2.41)$$

The normalizing term that appears in Theorem 1 is the second moment of this expression. While it is easy to see how $\Gamma_{ww,t}^{*-1}$ is obtained, to derive $\Gamma_{ww,t}^{**}$, we need to analyze the asymptotic variance of the second term of (2.41), that is $\frac{1}{\sqrt{H}} vec\left(U'^{*} X^{*}_{ww,t} - \overline{u}'^{r}_{t} \Lambda\right)$. First, notice that:

$$\frac{1}{\sqrt{H}} vec\left(U'^{*} X^{*}_{ww,t} - \overline{u}'^{r}_{t} \Lambda\right) = \frac{1}{\sqrt{H}} \underbrace{\left(X'^{*}_{ww,t} \otimes I_n\right)}_{kn \times nT(k+1)} \underbrace{vec\left(U'^{*}\right)}_{nT(k+1)\times 1} - \frac{1}{\sqrt{H}} \underbrace{\left(\Lambda' \otimes I_n\right)}_{kn\times kn} \underbrace{u^{r}_{t}}_{kn\times 1} \tag{2.42}$$

where we have used the fact that $u^{r}_{t} = vec(\overline{u}'^{r}_{t})$.

There are four elements to be considered, namely two variances and two covariances. Defining again *avar* the asymptotic variance and *acov* the asymptotic covariance, we have the following results.

Term 1:

$$\begin{aligned} \frac{1}{H}\left(X'^{*}_{ww,t} \otimes I_n\right) avar\left(vec\left(U'^{*}\right)\right)\left(X^{*}_{ww,t} \otimes I_n\right) &= \frac{1}{H}\left(X'^{*}_{ww,t} \otimes I_n\right)\left(I_{T(k+1)} \otimes \Sigma_n\right)\left(X^{*}_{ww,t} \otimes I_n\right) \\ &= \frac{1}{H}\left(X'^{*}_{ww,t} X^{*}_{ww,t} \otimes \Sigma_n\right) \\ &= \Gamma^{*}_{ww,t} \otimes \Sigma_n \end{aligned}$$

Term 2:

$$\frac{1}{H}\left(\Lambda' \otimes I_n\right) avar(u^{r}_{t})\left(\Lambda \otimes I_n\right) = \frac{1}{H}\left(\Lambda' \otimes I_n\right)\left(I_k \otimes \Sigma_n\right)\left(\Lambda \otimes I_n\right) = \frac{1}{H}\left(\Lambda'\Lambda \otimes \Sigma_n\right)$$

Term 3:

$$\frac{1}{H}\left(X'^{*}_{ww,t} \otimes I_n\right) acov\left(\underbrace{vec(U'^{*})}_{nT(k+1)\times 1}, \underbrace{(u^{r}_{t})'}_{1\times kn}\right)\left(\Lambda \otimes I_n\right)$$

To analyse this, notice that:

$$vec(U'^{*}) = \begin{bmatrix} u_1 \\ u^{r}_{1} \\ u_2 \\ u^{r}_{2} \\ \cdots \\ u_t \\ u^{r}_{t} \\ \cdots \end{bmatrix}$$

This means that the relevant matrix will contain zeros everywhere but in correspondence of the vector u^{r}_{t} appearing in $vec(U'^{*})$, where it will equal $I_k \otimes \Sigma_n$.

Compactly, this can be written as:

$$\underbrace{acov\left(vec(U'^{*}),(vec(\overline{u}_t^r))'\right)}_{nT(k+1)\times kn} = \begin{bmatrix} 0_{[(t-1)(k+1)+1]\times k} \\ I_k \\ 0_{(T-t)(k+1)\times k} \end{bmatrix} \otimes \Sigma_n \equiv \underbrace{\Xi}_{T(k+1)\times k} \otimes \Sigma_n.$$

Plugging in this term, we have

$$\begin{aligned} \frac{1}{H}\left(X'^{*}_{ww,t}\otimes I_n\right) cov\left(\underbrace{vec(U'^{*})}_{nT(k+1)\times 1},\underbrace{(u_t^r)'}_{1\times kn}\right)(\Lambda\otimes I_n) &= \frac{1}{H}\left(X'^{*}_{ww,t}\otimes I_n\right)(\Xi\otimes\Sigma_n)(\Lambda\otimes I_n) \\ &= \frac{1}{H}\left(\underbrace{X'^{*}_{ww,t}\Xi\Lambda}_{k\times k}\otimes\Sigma_n\right) \end{aligned}$$

Term 4: is simply the transpose of Term 3.

Collecting terms we have that the main normalising term is:

$$\begin{aligned} \Gamma^{**}_{ww,t}\otimes\Sigma_n &= \left(\Gamma^{*}_{ww,t}\otimes\Sigma_n\right) + \frac{1}{H}\left[(\Lambda'\Lambda\otimes\Sigma_n) - \left(X'^{*}_{ww,t}\Xi\Lambda\otimes\Sigma_n\right) - \left(\Lambda'\Xi X^{*}_{ww,t}\otimes\Sigma_n\right)\right] \\ &= \frac{1}{H}\left[\left(H\Gamma^{*}_{ww,t} + \Lambda'\Lambda - X'^{*}_{ww,t}\Xi\Lambda - \Lambda'\Xi X^{*}_{ww,t}\right)\otimes\Sigma_n\right] \end{aligned}$$

From this the proof follows.

2.B Proof of Theorem 2

Replacing in (2.20) y_j and r with the processes implied by the model (2.9)-(2.10) we have:

$$\begin{aligned} \widehat{\beta}_t - \beta_t = &\left[\sum_{j=1}^{T} w_{j,t}(I_n\otimes x_j x_j') + \lambda R'R\right]^{-1}\left[\sum_{j=1}^{T} w_{j,t}(I_n\otimes x_j x_j' + \lambda R'R)(\beta_j - \beta_t - \beta_t^B)\right] + \\ &\left[I_n\otimes\sum_{j=1}^{T} w_{j,t}x_j x_j' + \lambda R'R\right]^{-1}\left[\sum_{j=1}^{T} w_{j,t}(I_n\otimes x_j)u_j + \sqrt{\lambda}R'u_t^r\right] \end{aligned}$$

where, again, the term that multiplies $(\beta_j - \beta_t)$ is negligible assuming that the bandwidth is $o_p(T^{1/2})$. The analysis of the estimation bias and the convergence to normality follows trivially as in Theorem 1, so we do not repeat it here.

2.C Proof of Theorem 3

Let us rewrite the constrained estimator as a linear combination of the unconstrained one and of the constraints. If we define $E = I_n \otimes \sum_{j=1}^{T} w_{j,t} x_j x_j'$, $F = \lambda R'R$,and $G = \sum_{j=1}^{T} w_{j,t} vec(x_j y_j') = \sum_{j=1}^{T} w_{j,t} (I_n \otimes x_j) y_j'$, then the unconstrained estimator is $\widehat{\beta}_{t,GKY} = E^{-1}G$. We can therefore write (2.20) as:

$$\widehat{\beta}_t = [E + F]^{-1} \left[E\widehat{\beta}_{t,GKY} + \lambda R'r \right]$$

Defining $C = avar(\widehat{\beta}_{t,GKY}|\beta^T)$ we have that for any vector q and $w = E\,[E+F]^{-1}\,q$

$$\begin{aligned} q'\left(avar(\widehat{\beta}_{t,GKY}|\beta^T) - avar(\widehat{\beta}_t|\beta^T\right) q &= q'(C - [E+F]^{-1}\,ECE\,[E+F]^{-1})q \\ &= w'(E^{-1}\,[E+F]\,C\,[E+F]\,E^{-1} - C)w \\ &= w'(\left[I + E^{-1}F\right] C \left[I + FE^{-1}\right] - C)w \\ &= w'(\left[C + E^{-1}FC\right] \left[I + FE^{-1}\right] - C)w \\ &= w'(C + E^{-1}FC + CFE^{-1} + E^{-1}FCFE^{-1} - C)w \\ &= w'(E^{-1}FC + CFE^{-1} + E^{-1}FCFE^{-1})w \geq 0, \end{aligned}$$

which proves the result.[28]

Next, we turn to a comparison of the (conditional) Mean Squared Error of the constrained and unconstrained non-parametric estimators. We see that the ranking is not clear-cut, unless a certain condition is satisfied.

Following Alkhamisi and Shukur (2008), let us analyse the *canonical*[29] version of model (2.13). Let Λ and Ψ be the eigenvalues/eigenvectors of $X_{w,t}^{*\prime} X_{w,t}^{*}$, i.e. $X_{w,t}^{*\prime} X_{w,t}^{*} = \Psi\Lambda\Psi'$ and $\Psi\Psi' = I_k$. Defining $\widetilde{y}_t = \sqrt{w_{t,t}} y_t$, $\widetilde{x}_t = \sqrt{w_{t,t}} x_t$, $\widetilde{u}_t = \sqrt{w_{t,t}} u_t$, the *weighted* regression model (in matrix form) is:

$$\underset{1\times n}{\widetilde{y}_t{}'} = \underset{1\times k}{\widetilde{x}_t{}'} \underset{k\times n}{\Theta_t} + \underset{1\times n}{\widetilde{u}_t{}'},$$
$$\underset{k\times n}{\sqrt{\lambda}\overline{r}} = \underset{k\times k}{\sqrt{\lambda}R} \underset{k\times n}{\Theta_t} + \underset{k\times n}{\overline{u}^r}.$$

Using $\Psi\Psi' = I_k$ we can write $\widetilde{z}_t' = \widetilde{x}_t{}'\Psi$ and $\Xi_t = \Psi'\Theta_t$, and re-state the model in canonical

[28]This can also be seen intuitively by noticing that $avar\left(\widehat{\beta}_t|\beta^T\right)$ has a lower bound at 0, attained when $\lambda \to \infty$, and an upper bound at $avar(\widehat{\beta}_{t,GKY}|\beta^T)$, corresponding to $\lambda = 0$. Furthermore $[E+F]^{-1}E$ falls monotonically as λ increases. It follows that $avar\left(\widehat{\beta}_t|\beta^T\right) \leq avar(\widehat{\beta}_{t,GKY}|\beta^T)$ for every positive value of λ.

[29]In the canonical form the regressors are orthogonalised, as clarified below.

form as:

$$\underset{1\times n}{\widetilde{y}_t{}'} = \underset{1\times k}{\widetilde{z}_t{}'}\,\underset{k\times n}{\Xi_t} + \underset{1\times n}{\widetilde{u}_t{}'},$$
$$\underset{k\times n}{\sqrt{\lambda}\overline{r}} = \underset{k\times k}{\sqrt{\lambda}\overline{R}}\,\underset{k\times n}{\Theta_t} + \underset{k\times n}{\overline{u}^r}.$$

The unconstrained estimator of Ξ_t is then:

$$\begin{aligned}\Xi_t^u &= (Z_{w,t}^{*\prime} Z_{w,t}^{*})^{-1}(Z_{w,t}^{*\prime} Y_{w,t}^{*}) \\ &= \left(\Psi' X_{w,t}^{*\prime} X_{w,t}^{*}\Psi\right)^{-1}(Z_{w,t}^{*\prime} Y_{w,t}^{*}) \\ &= \Lambda^{-1}(Z_{w,t}^{*\prime} Y_{w,t}^{*}),\end{aligned}$$

We can derive the conditional Mean Square Error of the unconstrained estimator as:

$$MSE(\Xi_t^u|\Xi_t) = V^u = \left(\frac{1}{H}\Lambda\right)^{-1}\left(\frac{1}{H}Z_{w,t}^{*\prime} W_{H,t} Z_{ww,t}^{*\prime}\right)\left(\frac{1}{H}\Lambda\right)^{-1}\otimes\Sigma_n = \overline{V}^u\otimes\Sigma_n. \quad (2.43)$$

The constrained estimator is

$$\Xi_t^c = (\Lambda + \lambda\overline{R}'\overline{R})^{-1}(Z_{w,t}^{*\prime} Y_{w,t}^{*} + \lambda\overline{R}'\overline{r}),$$

The conditional Mean Square Error of the constrained estimator is:

$$MSE(\Xi_t^c|\Xi_t) = V^c = \left(\frac{1}{H}\Lambda+\Phi\right)^{-1}\left(\frac{1}{H}Z_{w,t}^{*\prime} W_{H,t} Z_{ww,t}^{*\prime}+\Phi\right)\left(\frac{1}{H}\Lambda+\Phi\right)^{-1}\otimes\Sigma_n = \overline{V}^c\otimes\Sigma_n, \quad (2.44)$$

where $\Phi = \frac{\lambda}{H}\overline{R}'\overline{R}$. To study the conditions under which $\overline{V}^u - \overline{V}^c$ is a positive semidefinite matrix, let us consider equivalently the quadratic form:[30]

$$\begin{aligned}&\beta'(\Lambda+\Phi)\left[\overline{V}^u - \overline{V}^c\right](\Lambda+\Phi)\beta = \\ &= \beta'(\Lambda+\Phi)\Lambda^{-1}\underbrace{Z_{w,t}^{*\prime} W_{H,t} Z_{ww,t}^{*\prime}}_{\equiv A}\Lambda^{-1}(\Lambda+\Phi)\beta - \beta'\underbrace{Z_{w,t}^{*\prime} W_{H,t} Z_{ww,t}^{*\prime}}_{\equiv A} + \Phi\beta \\ &= \beta'\left[\left(I+\Phi\Lambda^{-1}\right)A\left(I+\Lambda^{-1}\Phi\right) - (A+\Phi)\right]\beta \\ &= \beta'\left[A + A\Lambda^{-1}\Phi + \Phi\Lambda^{-1}A + \Phi\Lambda^{-1}A\Lambda^{-1}\Phi - A - \Phi\right]\beta \\ &= \beta'\left[A\Lambda^{-1}\Phi + \Phi\Lambda^{-1}A + \Phi\Lambda^{-1}A\Lambda^{-1}\Phi - \Phi\right]\beta.\end{aligned}$$

Hence, if the quantity $\left[A\Lambda^{-1}\Phi + \Phi\Lambda^{-1}A + \Phi\Lambda^{-1}A\Lambda^{-1}\Phi - \Phi\right]$ is positive semidefinite it follows that $\overline{V}^u - \overline{V}^c$ is indeed a semidefinite positive matrix suggesting such a relationship for the mean square errors of the respective estimators. Notice that the required condition is related to the amount of collinearity among the regressors. Specifically, a high degree of collinearity in the time series collected in the matrix $X_{w,t}^{*}$ will push some eigenvalues of

[30] To simplify the notation, in what follows we have set $H = 1$. This assumption is immaterial, since this term can be factored out.

$X_{w,t}^{*\prime} X_{w,t}^{*}$ close to 0, therefore making Λ^{-1} tend to ∞.

2.D Monte Carlo exercise

In the Monte Carlo exercise the coefficients Λ_t are obtained through the following algorithm.

1. First, simulate $n^2 + n$ coefficients according to the chosen DGP, where n is the size of the VAR.
2. The coefficients are then bounded by the largest one, so as to be lower or equal to 1 in absolute value.
3. At each t the first n^2 coefficients are orthogonalised through the Grahm-Shmidt procedure and used to form an $n \times n$ orthonormal matrix. Call this matrix P_t.
4. The last n coefficients are used to form a diagonal matrix of eigenvalues L_t. At each point in time the elements of L_t (call them l_t) are transformed using the following function

$$\widetilde{l_t} = 0.5(1 + \theta_L + eps) + 0.5(1 - \theta_L - eps)\arctan(l_t)/\arctan(1) - eps$$

   where l_t are the input eigenvalues (that by construction lie between -1 and 1), θ_L is the desired lower bound and eps is a small constant that ensures that the upper bound is $1 - eps$, that is strictly below 1. The resulting function is relatively smooth, as it can be seen in Figure 2.D.1 where $\widetilde{l_t}$ is plotted against the possible values of l_t (on the x axis there are 201 equally spaced points between -1 and 1) and $\theta_L = 0.85$.
5. Construct $\Lambda_t = P_t \widetilde{L_t} P_t'$

Figure 2.D.1: Constrained simulated coefficients

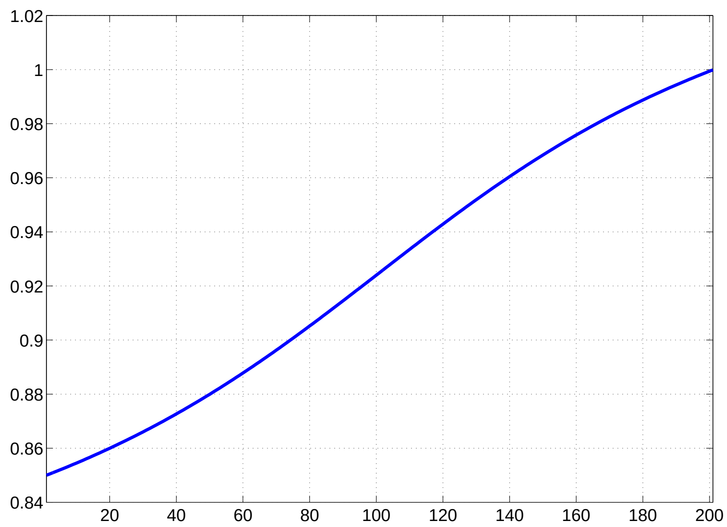

2.E Dynamic model selection

The estimation method of the parametric model suggested by Koop and Korobilis depends on a number of constants, the so called forgetting factor, θ, the prior tightness for the initial conditions λ and the constant κ. To select θ and λ we follow their dynamic model selection (DMS) algorithm. First the forgetting factor θ is made time-varying as follows:

$$\theta_t = \theta_{\min} + (1 - \theta_{\min})L^{f_t} \tag{2.45}$$

where $f_t = -NINT(\widehat{\varepsilon}'_{t-1}\widehat{\varepsilon}_{t-1})$, $NINT$ is the rounding to the nearest integer function, $\widehat{\varepsilon}_{t-1}$ are the one step ahead forecast errors, $\theta = 0.96$, $L = 1.1$ (values calibrated to obtain a forgetting factor between 0.96 and 1). As for the prior tightness we use a grid of J values. Each point in this grid defines a new model. Weights for each model j (defined $\pi_{t/t-1,j}$) are obtained by Koop and Korobilis as a function of the predictive density at time $t-1$ through the following recursions:

$$\pi_{t/t-1,j} = \frac{\pi^{\alpha}_{t-1/t-1,j}}{\sum_{l=1}^{J} \pi^{\alpha}_{t-1/t-1,l}} \tag{2.46}$$

$$\pi_{t/t,j} = \frac{\pi_{t/t-1,j}\, p_j(y_t|y^{t-1})}{\sum_{l=1}^{J} \pi_{t/t-1,l}\, p_l(y_t|y^{t-1})} \tag{2.47}$$

where $p_j(y_t|y^{t-1})$ is the predictive likelihood. Since this is a function of the prediction errors and of the prediction errors variance, which are part of the output of the Kalman filter, the model weights can be computed at no cost along with the model parameters estimation. Note that here a new forgetting factor appears, α, which discounts past predictive likelihoods and is set to 0.99. The constant κ is set to 0.96 throughout the exercise. At each point in time, forecast are obtained on the basis of the model with the highest weight $\pi_{t/t-1,j}$.

Chapter 3

The time varying effect of oil price shocks on euro-area exports

3.1 Introduction

The evolving relationship between the price of oil and the macroeconomy has been at the center of a lively debate in the empirical literature. A number of studies have indeed documented a decline in the importance of energy price shocks for economic activity, arguing that this relationship has changed around the mid-Eighties, see Hooker (1999), Blanchard and Gali (2007) and Edelstein and Kilian (2009).

According to some of these papers, changes in the structure of the economies and in the conduct of monetary policy have progressively insulated advanced economies from the negative effects of exogenous energy price increases. Blanchard and Gali (2007), for instance, point to more effective monetary policy, to a fall in the share of oil in both production and consumption and to lower real wage rigidities as plausible causes for this structural break. Blanchard and Riggi (2013) formally explore the quantitative relevance of each of these explanations through an estimated DSGE model, and find support for the role of vanishing wage indexation and improved monetary policy credibility. Other studies argue that the change in the relationship between oil prices and the macroeconomy reflects an evolution in the composition of the shocks underlying oil price fluctuations. According to this literature the role of exogenous flow supply shocks to crude oil, whose effect on economic activity is unambiguously depressive, has only a marginal role in determining oil price fluctuations. In contrast, shocks to the demand for oil associated with global activity booms have been responsible for most increases in the real price of oil from the mid-Seventies onwards (Kilian, 2009). This alone could explain why oil price increases are not necessarily associated with recessions. For example, part of the literature argues that the surge in the real price of oil observed in the 2000s can be mainly related to the rapid growth in emerging economies.

A natural corollary of this latter view is that, for oil importing countries, the *external* channel could be as relevant as the *domestic* one in understanding the effects of oil price shocks on macroeconomic activity. As booming economies lift the real price of oil they also stimulate *trade* and *exports*, which, in turn, could more than offset the adverse impact of higher energy prices on domestic demand components. Moreover, this mechanism is likely to have been reinforced by the remarkable increase in trade integration favoured by globalization. Considering, for instance, the euro area economy taken as a whole, which is the focus of the present study, between 1970 and 2010 the rate of growth of real exports has significantly outpaced that of GDP, see Figure 1. As a consequence the share of exports on GDP as well as the relevance of foreign demand fluctuations for domestic growth have constantly increased over time and foreign demand has become more and more important for domestic growth.

Interestingly, however, while the reaction of domestic demand to energy price shocks has been extensively studied, the relationship between trade and oil price fluctuations has been largely overlooked in the literature.[1] Our paper contributes to filling this gap. We start our investigation by looking at simple correlations between euro-area exports and the real price of oil. This preliminary analysis reveals a change in the co-movement between the two variables. In particular, while in the Seventies the correlation between euro-area exports and the real price of oil was basically nil, since the mid-Eighties it has become positive and significantly different from zero. This change might simply reflect an increase in the relative importance of expansionary global shocks (which stimulate both oil demand and global trade) in accounting for oil price variability in more recent periods. This hypothesis is supported, for instance, by findings in Kilian (2009) and Baumeister and Peersman (2013). A complementary explanation, which constitutes the focus of our paper, is that, conditional on each shock, the relationship between the price of oil and euro-area exports has varied over time owing to some structural changes that have influenced the joint dynamics of the two macroeconomic variables. We investigate the plausible structural sources of such a change using both a theoretical model and a Bayesian time-varying parameter structural vector autoregression with stochastic volatility (TVP-VAR). The theoretical framework is used to model the interplay between the oil market and exports and pick out the robust features of the

[1] One exception is Kilian, Rebucci, and Spatafora (2009), whose focus, however, is on the the impact on the external balance of oil price shocks.

impact responses of a number of endogenous variables to two structural shocks. The former is a flow oil supply shock, capturing the effect of an unexpected disruption in the production of oil. The latter is a foreign (from the euro area point of view) productivity shock that drives up the demand for oil. Simulations of the theoretical model provide us with useful restrictions on the signs of the reaction of the variables of interest to these shocks. In the spirit of Kilian and Murphy (2014) we combine these restrictions with plausible bounds on the price elasticity of oil supply and on the price elasticity of oil demand to identify these two shocks in the TVP-VAR, and analyze empirically how their effects on euro area exports has evolved over time. The structural analysis conducted on the basis of the TVP-VAR reveals that, conditional on each shock, the co-movement between the real price of oil price and euro area exports has indeed varied over time. In particular, conditional on negative oil supply shocks the association between the real price of oil and exports has become less negative, while, following a foreign productivity shock, a stronger positive co-movement has emerged.

We finally try to rationalize these changes using our theoretical model. We focus on a number of channels. First, a stronger trade relationship with emerging countries, whose growth has recently driven oil price increases. Second, a fall in the quantitative importance of oil in the world economy. Third, an increase in competitive pressures in the product market. Model simulations suggest that, in combination, these factors could potentially account for the changes documented by our empirical analysis.

The rest of the paper is organized as follows. Section 2 presents the basic stylized facts that motivate our analysis. Section 3 lays out the theoretical model and presents its predictions on the response of some endogenous variables to oil supply and demand shocks. Section 4 describes the empirical evidence obtained from the VAR with time varying coefficients and stochastic volatility. Section 5 uses the theoretical model in order to assess the potential for the three hypotheses listed above to explain our empirical findings. Section 6 concludes.

3.2 Motivation

We start our investigation by looking at raw correlations between real exports and the real price of oil. Real exports are chain linked export volumes, as measured in the euro area quarterly national accounts. The real price of oil is obtained by first converting in euros the U.S. dollar price of crude Brent, then deflating the resulting nominal price (denominated in euros) by the euro-area consumption deflator.[2] We use data between the second quarter of 1970 and the fourth quarter of 2013 and take log changes of both variables. Table 1 shows the correlation coefficients between these variables in two selected sub-samples. The former runs between 1970 to 1984, the latter between 1985 and 2013. The choice of the cutoff date (1984) is motivated by the findings in Hooker (1999), Blanchard and Gali (2007) and Edelstein and Kilian (2009) who identify a break in the relationship between the real price of oil and economic activity around the mid-Eighties, both for the U.S. and for the largest euro-area economies.[3] In the first sub-sample no clear co-movement emerges between real exports and the real price of oil, as the correlation coefficient stands at 0.12 and it is not significantly different from zero. In the second part of the sample, instead, a strong positive correlation (0.44) emerges. To further investigate changes in the relationship between these two variables, we run a regression of the rate of growth of real exports on the rate of growth of the real price of oil, allowing for the coefficient of the latter to change at a given point in time. First, we use a dummy variable that equals 0 between 1970 and 1984 and 1 thereafter, and interact it with the real price of oil. The results of this exercise are reported in the first two columns of Table 2. Between 1970 and 1984 the regression coefficient of the real price of oil is estimated at 0.07, and is not significantly different from zero. Its interaction with the shift dummy, on the other hand, displays a coefficient of 0.16, and is significantly different from zero at the 10 percent confidence level, indicating a stronger positive association between export and oil price growth after 1984. Columns 3 and 4 report the results of a similar exercise conducted using a dummy indicator that takes a value of 1 after 1989. The outcome is broadly similar, with the regression coefficient turning from 0.10 before 1989 to 0.24 after that year. Again, the interaction of the shift dummy and the real price of oil results in a coefficient that is significantly different from zero at conventional confidence levels.

In summary this preliminary exploration of the data indicates that a significant change in the reduced form relationship between the real price of oil and euro-area exports has occurred around the second half of the Eighties. In the next section we turn to a structural theoretical model that will provide us with some guidance for a deeper structural analysis of this issue.

[2]For the years prior to the Monetary Union an estimate of the exchange rate between the euro and the U.S. dollar, as well as of the consumption deflator and of real exports, is provided by the ECB Area Wide Model.

[3]We do not rely on formal break tests to date the change in the regression coefficients since these tests have been shown to have very low power in detecting time variation when the parameters of the true data generating process behave as a random walk Benati (2007).

3.3 The theoretical model

Our model is a variant of Clarida, Gali, and Gertler (2002), extended to consider the role of oil price dynamics in the spirit of Campolmi (2008) and Lipinska and Millard (2012). In our economy there are two oil importing countries, home (H) and foreign (F). They differ in size and share identical preferences, technology and market structure, though shocks may be imperfectly correlated. H has a mass of households n, whereas F has a mass $(1-n)$. In each country, production takes place in two stages. There is a continuum of intermediate goods firms each producing a differentiated input. These firms are monopolistic competitors and set nominal prices in a staggered fashion. Final goods producers are perfectly competitive. They combine intermediate inputs into final output, which they sell to households. The number of final goods firms within each country equals the number of households, whereas the number of intermediate goods firms is normalized at unity in each country[4]. Oil is used by the intermediate firms in the two countries, H and F, as an input in production together with employment. Within each economy households consume a domestically produced good and a good imported from the other country. In both countries households have access to a complete set of Arrow-Debreu securities which can be traded both domestically and internationally. An oil exporting country sells its endowment of oil and spends the associate revenues on consumption of goods from both H and F. Oil price is determined in equilibrium.

We make two assumptions about the oil market. First, oil is non-storable. This is clearly an ad-hoc hypothesis - as oil is in fact a storable commodity - made on the ground of simplicity. A recent relevant contribution by Unalmis, Unalmis, and Unsal (2012) incorporates speculative oil storage into a general equilibrium framework, giving rise to a dynamic link among oil inventories, storers' expectations of oil price and the spot price. This set up allows to study the impact of a storage demand shock, which decreases the availability of oil in place, therefore increasing the oil price. Through the lens of our model, this disturbance would be interpreted as an oil supply shock.

The second simplifying assumption is the exogeneity of oil supply. This assumption has some grounding in the empirical literature, where there is wide agreement that the price elasticity of oil supply is indeed close to zero (Hamilton, 2009). This is because changing oil production is highly costly and, given the uncertainty about the state of the crude oil market, oil producers do not revise the

[4] As shown by Clarida, Gali, and Gertler (2002) this assumption ensures that final goods producers face the same technology.

production level in response to high-frequency changes in demand Kilian (2009).[5] This assumption is maintained by a large part of the theoretical literature, see Campolmi (2008) Lipinska and Millard (2012) and Unalmis, Unalmis, and Unsal (2012). Some recent work, however, has shown that a mildly positively sloped short-run supply curve is indeed consistent with historical evidence Kilian and Murphy (2012). To take into account these recent developments, in the empirical section we relax this assumption and allow for a mildly upward sloping supply curve, using the elasticity bounds estimated by Kilian and Murphy (2012).

Given our objectives, we focus on two sources of cyclical fluctuations driving up oil prices: an oil supply shrinkage and an increase in foreign productivity. The latter is meant to capture the dynamic effects of an oil demand increase fostered by faster foreign growth.

3.3.1 Preferences

The representative household i in country H maximizes:

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left\{ \frac{(C_t(i) - hC_{t-1})^{1-\sigma}}{1-\sigma} - \frac{N_t(i)^{1+\phi}}{1+\phi} \right\} \tag{3.1}$$

where β is the discount factor, N_t denotes the household's i hours of labor and $C_t \equiv \Theta C_{F,t}^{\gamma} C_{H,t}^{1-\gamma}$ is a composite index of consumption of home and foreign goods, with $\Theta \equiv \gamma^{-\gamma} (1-\gamma)^{-(1-\gamma)}$. $C_{F,t}$ is an index of consumption of imported goods produced by F, $C_{H,t}$ is an index of consumption of domestic goods and $\gamma \equiv (1-n)\chi$ denotes the weight of imported goods in the consumption basket of households located in country H. The latter depends on the relative size of F and on χ, which is the degree of trade openness of H.

Households are concerned with "catching up with the Joneses": there is a certain degree of external habit persistence $h \in [0,1)$; C_{t-1} is the aggregate consumption level in period $t-1$. The period budget constraint is given by

$$P_{H,t}C_{H,t} + P_{F,t}C_{F,t} + Q_t^B B_t = B_{t-1} + W_{H,t}N_t + \Pi_t$$

where $P_{H,t}$ is the domestic price index, $P_{F,t}$ is a price index for foreign goods (in domestic currency), $W_{H,t}$ is the nominal wage, Π_t are profits, Q_t^B is the price of a one-period nominally riskless bond, paying one unit of domestic currency and B_t denotes the quantity of that asset purchased in period t. The optimal

[5] Note, however, that our assumption is stronger than in Kilian (2009), who only imposes a one month vertical oil supply curve.

allocation of expenditures between imported and domestically produced goods implies $P_{H,t}C_{H,t} = (1-\gamma)\, P_tC_t$ and $P_{F,t}C_{F,t} = \gamma P_tC_t$, where $P_t = P_{H,t}^{1-\gamma}P_{F,t}^{\gamma}$ is the consumer price index. In an analogous manner, the composite index of consumption in the foreign economy is $C_t^* \equiv \Theta^* \left(C_{F,t}^*\right)^{1-\gamma^*} \left(C_{H,t}^*\right)^{\gamma^*}$, where $\gamma^* \equiv n\chi^*$ and χ^* is the degree of trade openness in F.[6] The optimal allocation of expenditures implies $P_{H,t}^*C_{H,t}^* = \gamma^* P_t^* C_t^*$ and $P_{F,t}^*C_{F,t}^* = (1-\gamma^*)\, P_t^*C_t^*$, where $P_{F,t}^*$ denotes the price of foreign goods denominated in the producer's currency, $P_{H,t}^*$ denotes the price of domestic goods denominated in the foreign currency and P_t^* is the consumer price index in F denominated in foreign currency $P_t^* = P_{H,t}^{*\gamma^*} P_{F,t}^{*(1-\gamma^*)}$.

The law of one price implies that $P_{F,t} = \mathcal{E}_t P_{F,t}^*$, where $\mathcal{E}_t$ is the nominal exchange rate. The real exchange rate R_t is defined by $R_t \equiv \frac{\mathcal{E}_t P_t^*}{P_t}$.

3.3.2 Risk sharing

Under complete markets, the efficiency conditions for bonds' holdings by residents in F reads:

$$Q_t^B = \beta \mathbb{E}_t \frac{P_t^* \mathcal{E}_t}{P_{t+1}^* \mathcal{E}_{t+1}} \frac{\left(C_t^* - hC_{t-1}^*\right)^{\sigma}}{\left(C_{t+1}^* - hC_t^*\right)^{\sigma}} \tag{3.2}$$

Equating (3.2) with the Euler equations for both the home and foreign economies and log-linearizing around the steady state yields the familiar expression for the wedge between domestic and foreign interest rates:

$$i_t - i_t^* = \mathbb{E}_t \left[\pi_{t+1} - \pi_{t+1}^*\right] - r_t + r_{t+1} \tag{3.3}$$

where we have denoted with r the proportional deviation from steady state of the real exchange rate and with π and π^* the CPI inflation rate in H and F, respectively.

3.3.3 Firms

Final goods producers are perfectly competitive. Each of them produces output by using a continuum of intermediate goods as input, according to the CES technology $Y_t = \left(\int_0^1 Y_t(f)^{\frac{\epsilon-1}{\epsilon}} df\right)^{\frac{\epsilon}{\epsilon-1}}$ where Y_t denotes aggregate output and $Y_t(f)$ is the input produced by intermediate goods firm f. Both variables are expressed in per capita terms. Profit maximization, taking the price of the final good $P_{H,t}$ as given, implies

[6] The absence of home bias would require that the weight of domestically produced goods in home consumption basket $(1-\gamma)$ is equal to the weight of imported goods in foreign consumption basket (γ^*); home bias in consumption would require that $(1-\gamma) = 1-(1-n)\,\chi > \gamma^* = n\chi^*$.

the set of demand equations $Y_t(f) = \left(\frac{P_{H,t}(f)}{P_{H,t}}\right)^{-\epsilon} Y_t$ and the domestic price index $P_{H,t} = \left(\int_0^1 P_{H,t}(f)^{1-\epsilon} df\right)^{\frac{1}{1-\epsilon}}$.

Intermediate goods firms are monopolistic competitors, produce a differentiated intermediate good and set nominal prices in a staggered fashion. Each of them produces with the following technology:

$$Y_t(f) = N_t(f)^{\alpha_n} M_{H,t}(f)^{\alpha_m} \tag{3.4}$$

where $M_{H,t}(f)$ is oil used by firm f and $N_t(f)$ is labor input used by firm f (both normalized by population size). Firms take the price of both inputs as given. Accordingly, cost minimization implies $\frac{M_t^H(f)}{N_t(f)} = \frac{\alpha_m}{\alpha_n}\frac{W_{H,t}}{P_{m,t}}$, where $P_{m,t}$ is the nominal price of oil. Firm's f's nominal marginal cost is given by $\Psi_t(f) = \frac{W_{H,t}}{\alpha_n(Y_t(f)/N_t(f))} = \frac{P_{m,t}}{\alpha_m(Y_t(f)/M_{Ht}(f))}$. We assume that each period only a fraction of intermediate goods firms $(1-\theta)$, selected randomly, reset prices. The remaining firms keep their prices unchanged.

Intermediate goods firms in country F produce with the following technology:

$$Y_t^*(f) = A_t^* N_t^*(f)^{\alpha_n} M_{F,t}(f)^{\alpha_m} \tag{3.5}$$

where A_t^* is a productivity factor common across firms. We assume a first order autoregressive process $A_t^* = \left(A_{t-1}^*\right)^{\rho_A} e^{u_t^a}$, where u_t^a is an i.i.d. shock to foreign technology level.

3.3.4 Oil market

As in Lipinska and Millard (2012), oil is costless to transport and is non storable. Oil producer does not have access to world capital markets and simply uses the revenues from its production of oil to purchase final goods produced by H and F. The representative consumer in the oil producing country maximizes the following utility:

$$\max \mathbb{E}_t \sum_{r=0}^{\infty} \beta^r C_{t+r}^O$$

Aggregate consumption $C_t^O \equiv \Gamma\left(C_{F,t}^O\right)^{\varpi_O}\left(C_{H,t}^O\right)^{1-\varpi_O}$ is a composite consumption index of goods produced by H and F, where $\Gamma \equiv \varpi_O^{-\varpi_O}(1-\varpi_O)^{-(1-\varpi_O)}$, ϖ_O is the share of F-produced goods in the consumer's basket, $C_{H,t}^O$ is the consumption of the H-produced goods and $C_{F,t}^O$ is the consumption of F-produced goods. The consumer's budget constraint is given by $P_{m,t}M_t^s = C_{H,t}^O P_{H,t} + C_{F,t}^O P_{F,t}$, where M_t^s is oil endowment as defined below. Oil demand of the world economy is:

$$M_t^d = n\int_0^1 M_{H,t}(i)\,di + (1-n)\int_0^1 M_{F,t}(i)\,di = \frac{\alpha_m}{\alpha_n}\left(n\frac{W_{Ht}}{P_{m,t}}N_t + (1-n)\frac{W_{Ft}^*}{P_{m,t}^*}N_t^*\right) \tag{3.6}$$

The oil endowment M_t^s is assumed to follow a first order autoregressive process $M_t^s = \left(M_{t-1}^s\right)^{\rho_m} e^{u_t^m}$, where u_t^m is an i.i.d. shock to oil supply. Equilibrium in the oil market requires:

$$P_{m,t} = \frac{\alpha_m}{\alpha_n}\left(\frac{nW_{Ht}N_t + (1-n)W_{Ft}N_t^*}{M_t^s}\right) \tag{3.7}$$

3.3.5 Aggregate resource constraint

We can write the aggregate resource constraint in home and foreign country as follows:

$$nY_t = nC_{H,t} + (1-n)C_{H,t}^* + C_{H,t}^O \tag{3.8}$$

$$(1-n)Y_t^* = (1-n)C_{F,t}^* + nC_{F,t} + C_{F,t}^O \tag{3.9}$$

3.3.6 Monetary policy

We assume that the central bank of the home economy follows an interest rate rule:

$$I_t = \left[(\pi_t)^{\phi_\pi}(Y_t)^{\phi_x}\right] \tag{3.10}$$

where $I_t \equiv 1 + i_t$ is the nominal interest rate. A symmetric rule is assumed for F.

Steady state relations and optimality conditions are shown in Appendix A.

3.3.7 Model consistent impact sign restrictions

We use our theoretical frame to pick out the robust features of the responses of some endogenous variables to the two random disturbances in the model. We follow Canova and Paustian (2011), Dedola and Neri (2007) and Lippi and Nobili (2012) and carry out a Monte Carlo simulation on the relevant parameters of our theoretical model, assuming that they are uniformly and independently distributed over wide ranges. In particular, we draw 10000 vectors of the structural parameters from the uniform densities reported in Table 2, for each draw we save the responses to an oil supply shock and to a foreign productivity shock and compute the median, the 5th and 95th percentiles of the resulting distribution of impulse responses.

Table 2 reports the ranges of the uniform distributions for the parameters of the model that are simulated and the values imposed on the remaining calibrated parameters. Concerning the former group, the ranges of the uniform densities are sufficiently wide so as to cover all the reasonable values that these parameters can take. The degree of price stickiness is drawn over the interval that includes both a flexible prices' scenario (when $\theta = 0.1$ firms adjust their prices each quarter) and an almost completely rigid prices' scenario (when $\theta = 0.95$ firms adjust their prices once every twenty quarters). The parameter capturing the degree of habits is drawn from the range over which it is theoretically defined $h \in [0, 1)$. The ranges for the degree of trade openness in H and F include, at one extreme, the possibility that the country is completely closed to foreign trade (when χ and $\chi^* = 0$) and, at the other extreme, a high degree of trade openness (when χ and $\chi^* = 1$) and the possibility of no home bias in consumption. The share of foreign goods in the oil exporter country's consumers' basket is drawn from a uniform density over the support between $\varpi_O = 0$ (implying that oil exporters consume only $H-$produced goods) and $\varpi_O = 1$ (implying that oil exporters consume only $F-$produced goods). The ranges from which the coefficients of the Taylor rule and the inverse of Frisch elasticity are drawn encompass most calibrated and also estimated values used in the literature: $\phi_\pi \in [1.1, 5]$, $\phi_x \in [0, 1]$ and $\phi \in [0.1, 2]$. The elasticity of substitution among differentiated goods ϵ is drawn over the interval between 3 (implying a desired markup of 50%) and 11 (implying a desired markup of 10%). The range for the share of oil in production α_m is $[0.01, 0.04]$. The latter encompass the plausible values according to our computations based on OECD input output tables.[7] As for the serial correlation of the shock processes, we consider the same support used in Lippi and Nobili (2012), namely any value between and including 0.5 and 0.999, thus allowing also for near-random walks.

The remaining four parameters are calibrated. We set the discount factor in line with standard calibrations of DSGE models ($\beta = 0.998$). By calibrating $n = 0.5$ we assume that F and H have the same size. We set the risk aversion $\sigma = 0.1$, consistently with long lasting effects of shocks. In line with standard practice, we calibrate the elasticity of output with respect to labor input α_n to a value of 2/3.

The results of the model simulations are shown in Figures 2 and 3 , where we report the impulse responses of the four variables that we will use in the empirical

[7]For further details see Section 5.2.

analysis: the real price of oil, real exports of the home country, foreign output and global oil production to the two structural shocks. A feature common to these responses is that they can revert to steady state more or less slowly, given rather uninformative densities over the persistence of shocks ρ_m and ρ_a.

Figure 2 shows the effects of an oil supply shock, normalized to yield a 10 percent reduction in oil production on impact. The results fit well the conventional wisdom about the implications of an oil supply shortfall: the real price of oil increases and foreign GDP declines persistently. The response of exports of the home country H is mostly negative, although, for some combination of the model parameters, a positive response can not be ruled out. Figure 3 illustrates the impulse responses to a positive productivity innovation in the foreign country, F. The real price of oil goes up, as the demand for oil increases, and exports rise benefiting from the foreign expansion. Our theoretical model is silent on the implications of this shock on oil production, because, as discussed above, it makes the assumption of a vertical oil supply curve.

In the next section, these qualitative indications will provide a starting point for disentangling empirically two structural sources of fluctuations in oil prices: oil supply shocks and oil demand increases driven by foreign productivity shocks.

3.4 Empirical evidence

We jointly model four variables: the real price of oil, real exports, foreign GPD, and the global production of crude oil. The definition of real exports and of the real price of oil has been provided in Section 2. To construct a measure of foreign GDP we aggregate GDP volumes (at the price levels and PPP of 2000) from all available countries, excluding euro area economies. Further details are provided in Appendix B. Finally, global oil production is measured by the world production of crude oil (quarterly average of barrels per day). All variables are in log-changes. When information is available at the monthly frequency (like for the nominal price of oil and for oil production), it is aggregated at the quarterly frequency by taking quarterly averages. Foreign GDP is constructed from annual data and interpolated at the quarterly frequency using the Chow-Lin methodology. The sample period runs from the second quarter of 1970 to the fourth quarter of 2013. In Appendix B we provide further details on data sources and transformations.

3.4.1 The model

We specify the following VAR(p) model:

$$y_t = B_{0,t} + B_{1,t}y_{t-1} + B_{2,t}y_{t-2} + ... + B_{p,t}y_{t-p} + u_t, \qquad u_t \sim N(0, \Sigma_t)$$

We stack the VAR coefficients in the vector $\theta_t = vec([B_{0,t}, B_{1,t}, B_{2,t}, ..., B_{p,t}]')$ and assume that they evolve according to the law of motion $p(\theta_t/\theta_{t-1}, Q) = \mathbb{I}(\theta_t)f(\theta_t/\theta_{t-1,}Q)$, where the indicator function $\mathbb{I}(\theta_t)$ rejects unstable draws and the function $f(\cdot)$ is a multivariate Gaussian distribution such that, conditional on past information, θ_t is normally distributed with mean θ_{t-1} and variance Q. Based on these assumptions the VAR has the following state space representation:

$$\begin{aligned} \underset{n\times 1}{y_t} &= (I_n \otimes x_t)\,\theta_t + u_t, \qquad & u_t &\sim N(0, \Sigma_t) \\ \theta_t &= \theta_{t-1} + \epsilon_t, & \epsilon_t &\sim N(0, Q) \end{aligned} \tag{3.11}$$

where the row vector $x_t = [1, y'_{t-1}, y'_{t-2}, \ldots, y'_{t-p}]$ collects the intercept and the lags of the endogenous variables. In line with the literature we set $p = 2$. The VAR's reduced form innovations in (3.11) follow a multivariate Gaussian with zero mean and time varying covariance matrix Σ_t. The matrix Σ_t is further partitioned as $\Sigma_t = A_t^{-1}H_tA_t^{-1\prime}$, where A_t is lower triangular matrix with ones on the mail diagonal, and H_t is a diagonal matrix, that is:

$$A_t = \begin{bmatrix} 1 & 0 & 0 & 0 \\ \alpha_{21,t} & 1 & 0 & 0 \\ \alpha_{31,t} & \alpha_{32,t} & 1 & 0 \\ \alpha_{41,t} & \alpha_{42,t} & \alpha_{43,t} & 1 \end{bmatrix}, \quad H_t = \begin{bmatrix} h_{1,t} & 0 & 0 & 0 \\ 0 & h_{2,t} & 0 & 0 \\ 0 & 0 & h_{3,t} & 0 \\ 0 & 0 & 0 & h_{4,t} \end{bmatrix}$$

Collecting the $n(n-1)/2$ time varying elements of A_t in the vector a_t and the n time varying elements of H_t in the vector h_t, we further assume that:

$$\begin{aligned} a_t &= a_{t-1} + \eta_t, \qquad & \eta_t &\sim N(0, \Omega_a) \\ \log(h_t) &= \log(h_{t-1}) + e_t, & e_t &\sim N(0, \Omega_h) \end{aligned}$$

The structure of the covariance matrices Ω_a and Ω_h is the following:

$$\Omega_a = \begin{bmatrix} S_1 & 0_{1\times 2} & 0_{1\times 3} \\ 0_{2\times 1} & S_2 & 0_{2\times 3} \\ 0_{3\times 1} & 0_{3\times 2} & S_3 \end{bmatrix}, \quad \Omega_h = \begin{bmatrix} \sigma_1 & 0 & 0 & 0 \\ 0 & \sigma_2 & 0 & 0 \\ 0 & 0 & \sigma_3 & 0 \\ 0 & 0 & 0 & \sigma_4 \end{bmatrix}$$

where $S_1 = cov(\eta_t^1), S_2 = cov(\eta_t^2, \eta_t^3)$ and $S_3 = cov(\eta_t^4, \eta_t^5, \eta_t^6)$. This implies that the non-zero non-unit elements of A_t are independently distributed across rows, while being correlated within rows. Block independence is assumed for the random errors, so that their joint distribution is:

$$\begin{bmatrix} \varepsilon_t \\ \epsilon_t \\ \eta_t \\ e_t \end{bmatrix} \sim N(0, V), \quad V = \begin{bmatrix} I_4 & 0 & 0 & 0 \\ 0 & Q & 0 & 0 \\ 0 & 0 & \Omega_a & 0 \\ 0 & 0 & 0 & \Omega_h \end{bmatrix}$$

where $\varepsilon_t = A_t^{-1} H_t^{1/2} u_t$. The model is estimated with Bayesian methods, through a Gibbs sampling algorithm. The exact steps of the algorithm are described in details in a number of papers, see for example Benati and Mumtaz (2007), pages 9 to 12, and therefore will not be repeated here. Further details on the application of the method to our specific case are provided in Appendix B.

3.4.2 Identification of the structural shocks

Orthogonal structural shocks ε_t are recovered from the reduced form residuals u_t through the following relationship:

$$u_t = A_{0,t}\varepsilon_t, \qquad \varepsilon_t \sim N(0, I_n) \tag{3.12}$$

so that $\Sigma_t = A_{0,t}A'_{0,t}$. Consistently with the theoretical model laid out in Section 3, we identify two structural shocks, namely a disturbance to the supply of oil and a foreign (from the euro area point of view) productivity shock. To disentangle these two shocks we follow Baumeister and Peersman (2013) but complement their approach with insights from Kilian and Murphy (2012). In more details, we employ a mix of restrictions on the signs of the first two columns of the structural impact matrix $A_{0,t}$ as well as on the relative magnitude of some of its elements, so as to ensure that the slopes of the oil demand and of the oil supply curve have the correct sign and fall within plausible values. Furthermore, we add to these two sets of restrictions a constraint on the dynamic response of the real price of oil conditional on a supply shock. The procedure to identify the structural shocks is based on the algorithm proposed by Rubio-Ramirez, Waggoner, and Zha (2010), modified to take into account the additional constraints.

Sign restrictions on impact

The theoretical model provides us with a set of robust restrictions on the sign of the response of the variables included in our VAR to both a negative shock to the supply of oil and to a positive foreign productivity shock.

Following an unexpected disruption in oil supply the real price of oil increases, while oil production and foreign output decrease on impact. The response of euro area exports is very likely to be negative, yet given some uncertainty in the sign of the response produced by the DSGE model we prefer no to impose any sign on the impact response of this variable.

In response to a foreign productivity shock, which stimulates both the demand for oil and for euro area goods, the real price of oil, euro area exports and foreign output unambiguously increase. Regarding oil production, our theoretical model would imply a zero response on impact, consistently with assumptions in Kilian (2009). Kilian and Murphy (2012) challenge the hypothesis of a vertical short-term supply curve and allow for a positively sloped supply schedule, but impose a tight upper bound on the impact price elasticity of oil supply. In our empirical investigation we follow their indications and constrain the response of oil supply to be non-negative on impact, an assumption consistent both with a vertical (like in the theoretical model) and with a positively sloped supply curve in the short term. A summary of the identifying impact sign restrictions is reported in Table 3.

Even when they are derived from a theoretical model, impact sign restrictions have been shown to be too weak to provide reasonable estimates of the effects of oil demand and oil supply shocks. Kilian and Murphy (2012) and Kilian and Murphy (2014) suggest to further narrow down the set of admissible structural models by imposing both dynamic sign restrictions and elasticity bounds. In the remaining sub-sections we explain the additional identifying restrictions that we impose to reduce the set of admissible structural impact matrices $A_{0,t}$.

Dynamic sign restrictions

Following Kilian and Murphy (2014) we restrict the response of the real price of oil to an oil supply shock to be positive for at least five quarters (starting in the impact period) after the initial shock. This restriction rules out the possibility that an unanticipated oil supply disruption has a very short lived effect on the price of oil and that, following an increase on impact, the price of oil actually falls within a year from the shock.

Elasticity bounds

We further reduce the set of admissible structural impact matrices $A_{0,t}$ through restrictions on the price elasticity of oil supply and of oil demand.

Starting from the oil supply curve, an estimate of the impact price elasticity can be obtained from the ratio of the response of oil production relative to the response of the real price of oil to an oil demand shock, which in our theoretical model corresponds to an unexpected change in foreign TFP.Kilian and Murphy (2014) argue that a plausible upper bound to such elasticity (which is required to be non-negative) stands at 0.025. This estimate corresponds to the ratio of the percentage increase in oil supply and the percentage increase in the real price of oil observed in August 1990, when the First Gulf War burst out. Since the spike in the price of oil recorded in this particular occasion can be seen as truly exogenous, the corresponding change in the production of oil traces the price elasticity of supply. In their study Kilian and Murphy find that their findings are robust to higher values of this bound, up to 0.1. In our analysis we pick the higher end of the [0.025-0.1] range considered by these authors for two reasons. First, we use quarterly, rather than monthly data and it is therefore reasonable to assume that, at this lower frequency, producers have more opportunity to change supply in response to demand disturbances. Second, and more importantly, we measure oil prices in euros. Since an increase in the nominal price of oil in U.S. dollars is historically accompanied by a depreciation of the U.S. dollar[8], the sensitiveness to structural shocks of the nominal price of oil is likely to be lower when the latter is measured in euros, therefore requiring an upward adjustment of the oil supply elasticity. The magnitude of such adjustment (by a factor of 4, from 0.025 to 0.1) is consistent with the fact that the standard deviation of the percentage change in the U.S. dollars nominal price of oil is around four times as high as the one computed on the nominal price of oil in euros (22.5 as opposed to 4.5). Given the ordering of the variables (real price of oil, exports, foreign GDP and oil production) and of the shocks (oil supply shock first, oil demand shock second), the oil supply elasticity bound, together with the sign restriction described above, implies the following constraint on the admissible structural matrix $A_{0,t}$: $0 \leq A_{0,t}^{(4,2)}/A_{0,t}^{(1,2)} \leq 0.1$.

Next, we turn to the price elasticity of demand (which can be inferred by the ratio of the impact response of oil production to an oil supply shock relative to the

[8]In our dataset, between the second quarter of 1970 and the fourth quarter of 2013, we observe an unconditional correlation of -0.4 between the price of oil in U.S. dollars and the dollar/euro exchange rate.

impact response of the real price of oil). When imposing a price elasticity of supply of 0.1, Kilian and Murphy (2014) find an impact oil demand elasticity between -0.24 and -0.76 with 64 percent confidence.[9] A similar range of values is reported by Baumeister and Peersman (2013). Using a VAR with time varying parameters, they estimate the oil demand elasticity to have declined in absolute value from -0.6 to -0.1 in the past four decades. Taking into account this evidence we rule out structural impact matrices $A_{0,t}$ that yield an impact elasticity of demand more negative than -0.8. Taking also into account the impact sign restrictions discussed above, we require the elements of $A_{0,t}$ to satisfy these set of inequalities $0 \geq A_{0,t}^{(4,1)}/A_{0,t}^{(1,1)} \geq -0.8$ for $A_{0,t}$ to be in the set of admissible structural models.

The algorithm

To summarize, the identification procedure consists of the following steps.

1. Obtain a draw of the reduced form parameters $\widetilde{\theta}_t$ and $\widetilde{\Sigma}_t$ from the posterior distribution.

2. Let $\widetilde{\Sigma}_t = P_t D_t P_t'$ be the eigenvalue-eigenvector decomposition of the covariance matrix at time.

3. Draw an $(n \times n)$ independent standard normal matrix $\overline{X}$ and compute the QR decomposition $\overline{X} = QR$ where $QQ' = I$ and R is a diagonal matrix with positive elements on the main diagonal.

4. Generate a candidate structural impact matrix $A_{0,t} = P_t D_t^{\frac{1}{2}} Q'$ and compute the impulse response functions based on $A_{0,t}$. If the impulse responses satisfy the restrictions described above, record $A_{0,t}$ and the corresponding impulse response functions, otherwise discard them.

We apply the above algorithm until we retain 500 admissible $A_{0,t}$ at each time t.

Summarizing the dynamic effects of structural shocks

When there is more than one structural admissible model, like in our case, summarizing the dynamic effects of structural shocks poses a conceptual challenge. Fry and Pagan (2011), in fact, point out that the conventional choice of reporting

[9] See Table II in Kilian and Murphy (2014). The model analyzed in Kilian and Murphy (2014) includes oil inventories. This allows them to distinguish the price elasticity of oil in production from the price elasticity of oil in use, where the latter is by construction lower (in absolute value) due to adjustments in inventories following a supply shock. We do not make such a distinction, implicitely equating production and consumption.

the vector of median posterior responses is fundamentally flawed. It is indeed very likely that, at different horizons, the median responses will be generated by different structural models, making it impossible to give the results a structural interpretation. Their suggested solution is to focus on the structural impact matrix $A_{0,t}$ that generates Impulse Response Functions (IRFs) that show minimum distance with respect to the median posterior response. This ensures that the vector of responses is generated by a single structural model that can be seen, in some sense, as the most representative one. However, Inoue and Kilian (2013) criticize this proposal by noticing that the median vector is not a well defined measure of central tendency, so that, even if the strategy proposed by Fry and Pagan (2011) were to result in a perfect match between the IRFs of a single structural model and the median posterior response, there would be no compelling reason to focus on such a model. A practical way to address this shortcoming is to search for a model that minimizes that distance from the mean - rather than from the median - response, see Inoue and Kilian (2013) for a discussion. This choice has some theoretical appeal given that a quadratic loss function would be minimized at the posterior mean and that the mean is a well defined statistical concept in the case of vectors. We therefore follow this route and represent results from our structural VAR based on the admissible model that minimizes (at each point in time, given the time-varying nature of the model parameters) the distance from the mean response.[10]

3.4.3 Estimation Results

We start by investigating whether and how the response of euro-area exports to unexpected falls in the production of crude oil and to foreign TFP shocks has changed over time.

In Figure 4 we display the IRFs of exports to the two identified structural shocks. The top panel of this plot shows that, despite the fact that no sign restriction was placed on the impact effect on exports, the response of exports to a shock to the supply of oil is consistently negative throughout the sample. However, an upward tendency can be detected in the profile of the IRFs, so that the effect of the oil supply shock results more negative in the 70s than in the following decades, indicating that euro-area economic activity has been progressively more insulated from the recessionary impact of oil supply disruptions also thanks to a less negative response of foreign sales. In the top panel of Figure 5 we report the same information

[10]A shortcoming of relying on the posterior mean is that there is no way to construct credible sets around the central tendency. The neatest solution would be to use the method by Inoue and Kilian (2013) that minimizes the distance from the posterior mode. This, however, poses important technical challenges, especially for a partially identified VAR model like the one we employ here.

but average the IRFs over four separate sub-samples, roughly corresponding to the four decades under investigation.[11] The smoother profile of the time varying IRFs magnifies the difference between the 70s and the following decades.

We next turn to the dynamic effect of a structural shock to foreign TFP, shown in the bottom panel of Figures 4. Again, we detect notable time variation in the response of export (constrained to be positive on impact) to an unexpected increase in external demand. The sensitiveness of euro area exports to such a shock has increased smoothly over time, stabilizing around historically high levels after 2000. This gradual increase is all the more evident from the bottom panel of Figure 5.

After having documented an attenuation in the negative effects of oil supply shocks on euro-area exports and an amplification of the positive impact of foreign demand shocks, we now turn to examine changes in the co-movement between exports and the real price of oil conditional on the two identified shocks. Following den den Haan and Sumner (2004) we investigate the issue by looking at the conditional covariance between the variables of interest, which can be computed as the product of the response functions at each horizon k, cumulated over the previous horizons.[12]

The estimated conditional covariances between euro-area exports and the real price of oil are shown in Figures 6 for the single quarters, and in 7 for the four decades in our sample. The top panels of these Figures show that, conditional on an oil supply shock, the co-movement between exports and the real price of oil has become markedly less negative over time, especially when the 70s are compared to the following decades. Conditional on a foreign TFP shock, which causes an unexpected increase on oil demand, the covariance between export volumes and the real price of oil has become more positive over time, with a peak at the end the Nineties, early 2000s. Although this tendency has partially reversed in recent years, the covariance is substantially higher at the end of the sample than in the 70s, early 80s.

Summing up, based on a VAR with time varying parameters and stochastic volatilities we have provided evidence of structural change in the transmission of identified oil supply and oil demand shocks to the joint dynamics of euro-area exports and of the real price of oil. In recent decades oil price increases due to an oil supply disruption have been coupled with a more muted fall of export volumes, while a stronger association between oil price and export increases has emerged conditional

[11] Since our sample extends to 2013 the fourth sub-sample has 12 observations more than the first three. Excluding these observations does not change the results.

[12] See Kilian (2008) for an application of this methodology to the co-movement of consumer prices and output conditional on exogenous oil supply shocks.

on an unexpected surge in global output. These variations in conditional second moments point to the existence of at least some structural changes that have affected the joint dynamics of euro area exports and the real price of oil over last decades.

3.5 Interpreting structural changes

In this section we use the theoretical model developed in Section 3 to shed light on the plausible mechanisms behind the structural changes uncovered by our empirical analysis. The three plausible explanations, not mutually exclusive, on which we focus are the consolidation of the trade relationship with emerging economies, the decrease in the share of oil in production and lower markups.[13]

3.5.1 Trade relationship with emerging countries

The stronger is the trade relationship between euro area and emerging countries, the larger will be the positive response of euro area exports to faster growth in emerging economies. A first gauge of the importance of this channel can be gained by looking at the evolution of the shares of exports towards Asian economies in the four largest euro area countries.

Figure 8 shows that export shares towards Asia have indeed notably increased since the end of the 1980s. This increase is glaring in Germany, and, to a lesser extent, in France, where the share of exports directed to Asian emerging countries has more than tripled over the last two decades, reaching around 10 and 8 percent, respectively, compared to about 6 in Italy and 4 in Spain. The privileged trade relationship with emerging economies might have played an important role in the macroeconomic performance of Germany during the 2000s, when Asia (and more in particular China) progressively became a major sales market for German goods.

Figures 9 and 10, panels A, illustrate the changes in the responses of exports in H to the two different sources of fluctuations in oil prices, when χ^* becomes

[13]We also explored a fourth channel: a new advantageous flood of petrodollars towards the euro-area. A surge in oil prices leads to a redistribution of income from oil consuming to oil producing countries and the use that the latter make of their revenues can considerably affect global imbalances and, thus, the overall impact of rising oil prices on oil importing economies. Quantifying how the oil-revenues are spent is somewhat problematic. Still, Higgins, Klitgaard, and Lerman (2006) report suggestive evidence that the geography of petrodollar recycling has changed: oil exporters are importing more goods from the euro-area today than they were 25 years ago and fewer goods from the US. We have studied the changes in the responses of exports when $(1 - \varpi_O)$, i.e. the parameter capturing the preference of the oil producing economy for goods produced by H, becomes higher. The implications conditional on foreign productivity shocks are negligible. Considering the responses to an oil supply shock, to have a significant change in the response of exports the fraction of petrodollars recycled back home to purchase H- produced goods should have increased from zero to 100%.

higher. In the model χ^* is the deep parameter capturing the preference in F for goods produced by H.[14]

When higher oil prices are driven by faster growth in F, the positive conditional correlation between the real price of oil and exports in H can be amplified by a tighter trade relationship with F (Figure 10, panel A). However, this structural change is conducive to larger negative response of exports to oil supply shocks (Figures 9 panel A), an implication of the model that contrasts starkly with the evidence shown in Figure 5.

The rationale is as follows. Oil supply shocks are recessionary for both H and F and thus lead in both countries to a contraction in demand, that involves both consumption of domestic goods and well as consumption of imported goods. Oil supply shocks are instead expansionary for the oil producing economy, that accordingly increases its consumption of both $H-$ and $F-$ produced goods. From the point of view of H, its exports towards F fall while its exports towards the oil producing economy rise. When the share of H-produced goods in the consumption basket of F goes up (i.e. when χ^* rises), all other things held constant, the ratio of exports in H towards F over total exports increases and, at the same time, the ratio of exports in H towards the oil producing country over total exports falls. This explains why, when χ^* rises, the contractionary effects of the oil supply shock on exports in H become larger.

3.5.2 Lower oil shares

Edelstein and Kilian (2009) document that the overall share of energy expenditures in US consumer spending fell steadily throughout the 1980s and 1990s, reaching a low at the beginning of 2000s; Kilian and Vigfusson (2013) focus on the share of crude oil in U.S. GDP showing that the latter peaked in 1980/81, reached a trough in the late 1980s and began to rebound only after 2003, although it never reached the levels of 1980s. Consistently, Blanchard and Galì (2010) provide evidence that the share of oil in both US production and consumption is smaller today than it was in the 1970s.

The decline in the quantitative importance of oil appears to be present in a large number of countries. By using input-output tables from the OECD database, we find that the shares of oil in production have markedly decreased in the euro area

[14] Parameters' calibration behind the responses shown in Figures 9 and 10 is consistent with the one used to get sign restrictions in Section 3.7: those parameters that are kept fixed in Section 3.7 are calibrated at the same values as indicated in Table 3; for those parameters that are Monte-Carlo simulated in Section 3.7 we choose a value inside the ranges given in Table 3.

as well (Figure 11).[15]

While Blanchard and Gali (2007) conjecture that this decline can account for a part of the decrease in the US macroeconomic volatility conditional on oil shocks, Edelstein and Kilian (2009) show that fluctuations in the share of energy in consumption cannot explain the declining importance of energy price shocks. In order to highlight the implications of lower oil shares, we simulate our model under different values of α_m. Figure 9, panel B illustrates that such a structural change reduces the negative impact of an oil supply disruption on exports. The reason is that it mitigates the recessionary effects of oil supply shocks on oil importing economies trading with each other.

By contrast, the reduction of oil shares in the world economy has a negligible impact on the response of exports to foreign productivity shocks (Figure 10; panel B). Indeed, in this case the bulk of the exports' movements in H depends on the cyclical expansion in F, which is almost unaffected by the change in the shares of oil.

3.5.3 Lower markups

There is widespread agreement that the global integration of the real and financial markets, the new ICT technologies and the process of European integration led to a strong and sharp increase in competitive pressures for the euro-area economies. The latter can be captured in our theoretical model by an increase in the elasticity of substitution among differentiated goods ϵ, that implies lower desired markups $(\frac{\epsilon}{\epsilon-1})$ in oil importing economies.

Figure 9 panel C shows that an increas in competitive pressure indeed reduces the impact of a negative oil supply shock on the exports of the oil importing countries. This happens because lower desired markups make the Phillips curves flatter, thus dampening the inflationary effect of oil price increases and consequently their recessionary effect. Milder recessions imply that the contraction of exports in these countries, that trade with each other, turns out to be smaller.

As displayed by Figure 10 panel C, the positive response of exports to faster foreign growth is amplified with lower desired markups, although the quantitative impact of this structural change is almost negligible.

[15] Figure 11 does not display the strong co-movement between the oil shares and the oil price detected by Kilian and Vigfusson (2013) because of the low frequency nature of the time series we have constructed. Notice, however, that the pattern is very similar to that shown for the U.S. by Kilian and Vigfusson (2013): in France and Italy - for which we have data prior to 1980- the share of oil peaked in 1980/81, in all the major euro area countries it reached a trough at the beginning of 1990s, and began to rebound systematically thereafter, although remaining well below the level of 1980s.

3.6 Conclusions

This paper provides novel evidence on how the relationship between fluctuations of the real price of oil and exports in the euro area has changed over the last 40 years. The conjecture that motivates our analysis is that the milder recessionary impact of oil price shocks, amply documented in the literature, might partly stem from changes in the way euro-area foreign sales respond to the structural shocks that drive the real price of oil. Our analysis is based on the interaction between a theoretical model and a structural VAR with time-varying coefficients and stochastic volatility. The theoretical model allows us to derive a set of restrictions on the signs of the impact response of the real price of oil, real exports, foreign (from the euro-area point of view) output and global oil production to two structural shocks: an unexpected fall in the supply of oil and a shock to foreign total factor productivity that also raises the demand for oil. By complementing these restrictions with plausible bounds on the price elasticity of oil supply and of oil demand, we identify these two shocks through our structural VAR, and study how their effect on euro area exports and on the co-movement between exports and the real price of oil has changed over time.

Our estimates indicate that the co-movement between euro-area exports and the real price of oil has become less negative conditional on oil supply shocks, and more positive conditional on oil demand shocks, pointing to the existence of some structural change influencing the joint dynamics of these two variables. Through the lens of our theoretical model, we qualitatively explore the role of a number of factors in accounting for such changes: larger export shares towards emerging countries, stronger competitive pressure in the product market and a reduction in the quantitative importance of oil in production. We show that a stronger trade relationship with emerging countries can potentially explain the increase in the positive correlation between euro area exports and the real price of oil conditional on foreign productivity shocks. However, such a structural change, taken in isolation, would lead to an even larger drop in exports following an oil supply shock, in contrast with our evidence. Instead, the more muted effect of an adverse oil supply shock can be accounted for by higher competitive pressure in the product market and by a decrease in the share of oil in production. We conclude that, in combination, these three factors could make a good job in explaining our evidence on the variations in the joint response of exports and energy prices to identified shocks. Our analysis bears important policy implications, as it adds an international dimension to the assessment of the impact of oil price fluctuations on the macroeconomy, often confined to the domestic demand channel. For example, our results indicate that the stimulus for the euro-area economy stemming from the fall in the price of oil observed since the summer of 2014 is likely to be mild. On the one hand, insofar as this decrease reflects weakening global trade, it will be associated with lower foreign demand. Moreover, for a given shock to global output, the loss of foreign sales is likely to be more marked than in previous decades. On the other hand, to the extent that falling energy prices reflect also an increase in the supply of oil, the positive effect exerted on GDP through higher exports is likely to be negligible. Finally, our findings raise an interesting question on whether the fluctuations in the price of oil that we have observed since 2003 have played an important role in widening cross-country imbalances within the Monetary Union, an issue that we leave for future research.

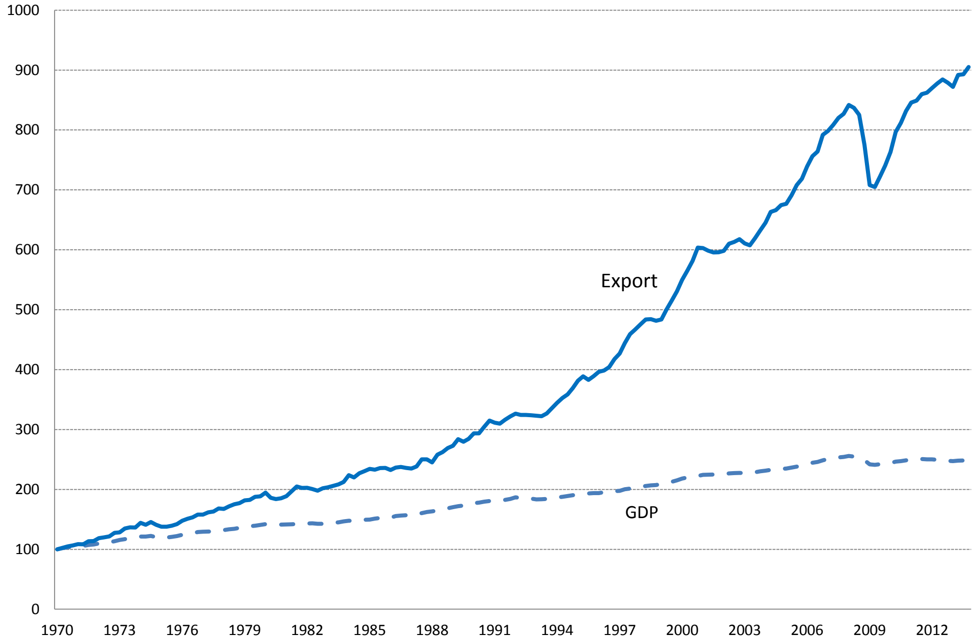

Figure 3.1: Real GDP and real exports in the euro area (1970q1=100)

	1970-1984	1985-2013
ρ	0.12	0.44
p-val	0.33	0.00

Table 3.1: Table 1: Correlation between the growth rate of euro-area real exports and the growth rate of the real price of oil in the indicated sub-samples

	1 Coefficient	2 p-val	3 Coefficient	4 p-val
intercept	5.49	0.00	5.23	0.00
Δ oil price	0.07	0.37	0.10	0.07
D1984	-0.78	0.52		
Δ oil price $\times$ D1984	0.16	0.07		
D1989			-0.67	0.56
Δ oil price $\times$ D1989			0.15	0.05

Table 3.2: Regression analysis. The dependent variable is the rate of growth of real exports, as defined in Section 2. The price of oil is the real price of crude oil, as defined in Section 2. D1984 is a dummy variable that equals 0 between the second quarter of 1970 and the fourth quarter of 1984, 1 otherwise. D1989 is a dummy variable that equals 0 between the second quarter of 1970 and the fourth quarter of 1989, 1 otherwise. The sample goes from the second quarter of 1970 to the fourth quarter of 2013 (175 observations).

	SIMULATED PARAMETERS	RANGE OF VALUES
θ	Price stickiness	$[0.1, 0.95]$
χ	Degree of trade openness in H	$[0.0, 1.0]$
χ^*	Degree of trade openness in F	$[0.0, 1.0]$
ϖ_O	Share of F-goods in the oil exporter country' consumers basket	$[0.0, 1.0]$
h	Habit	$[0.0, 0.999]$
ϕ_π	Taylor coefficient on inflation	$[1.1, 5.0]$
ϕ_x	Taylor coefficient on the output gap	$[0.0, 1.0]$
ϵ	Elasticity of substitution among differentiated goods	$[3, 11]$
ϕ	Inverse of the Frisch elasticity	$[0.1, 2.0]$
α_m	Oil's share in production	$[0.01, 0.04]$
ρ_m	Persistence of oil supply shock	$[0.5, 0.999]$
ρ_a	Persistence of foreign productivity shock	$[0.5, 0.999]$
	CALIBRATED PARAMETERS	
β	Intertemporal discount factor	0.998
σ	Risk aversion	0.1
α_n	Labor's share in production	2/3
n	Mass of households in H	0.5

Table 3.3: Ranges over which the indicated parameters are drawn in the Monte Carlo simulation of the theoretical model.

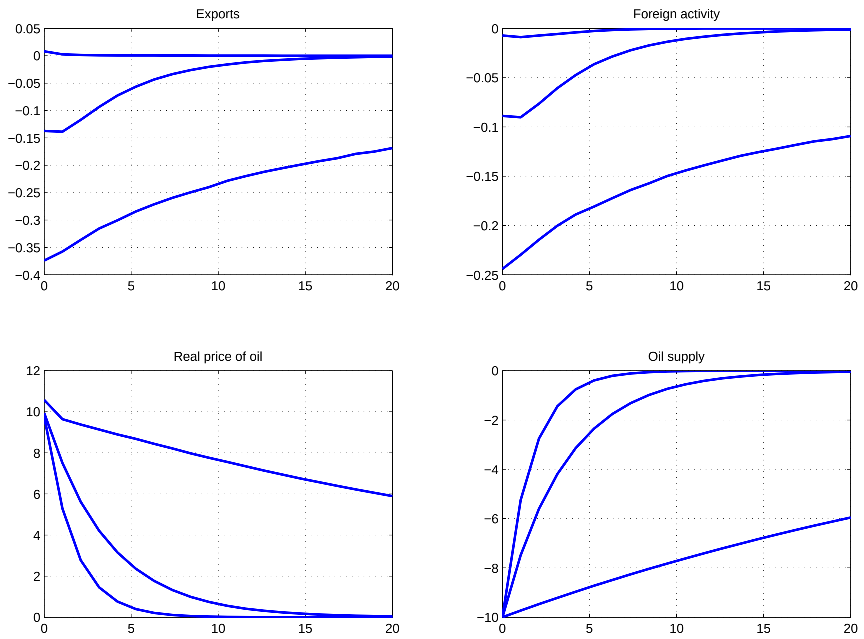

Figure 3.2: Theoretical IRFs to an oil supply shock. The Figure reports the median, the 5th, and 95th percentiles of the IRFs distribution.

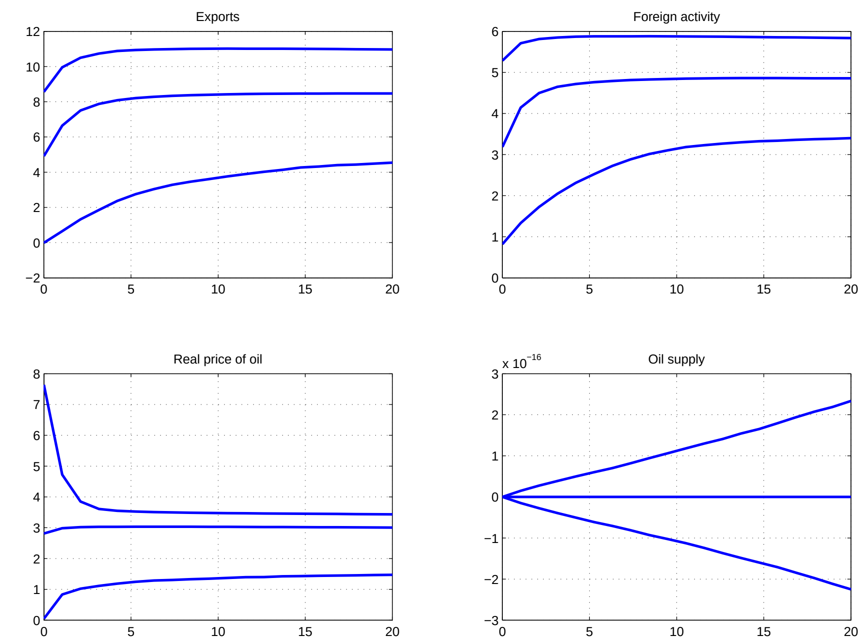

Figure 3.3: Theoretical IRFs to a foreign TFP shock. The Figure reports the median, the 5th, and 95th percentiles of the IRFs distribution.

	Structural shocks	
VAR variables	oil supply	foreign TFP
real price of oil	+	+
real export		+
foreign output	-	+
oil supply	-	+

Table 3.4: Impact sign restrictions on the IRFs of the endogenous variables. A + (or−) indicates that the impulse response of the variable of interest is restricted to be positive (negative) on impact. A blank entry indicates that no restriction is imposed on the response.

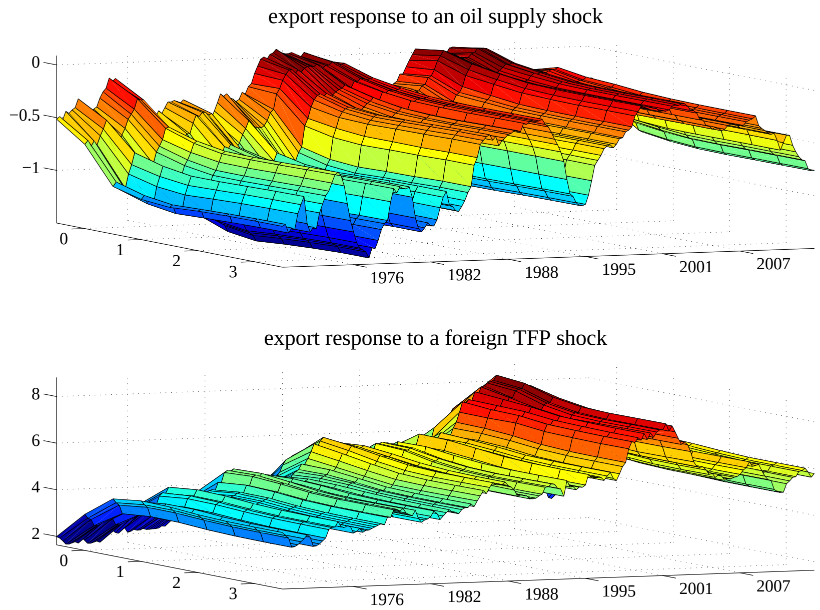

Figure 3.4: Cumulative Impulse Response Functions (IRFs) of export volumes. In each quarter the IRF is selected as the one that is closest to the mean IRF among those derived from the admissible structural models given the identifying restrictions. IRF to an oil supply shock are normalised to yield a 1% impact increase in oil prices. IRF to a foreign output shock are normalised to yield a 1% impact increase in foreign output.

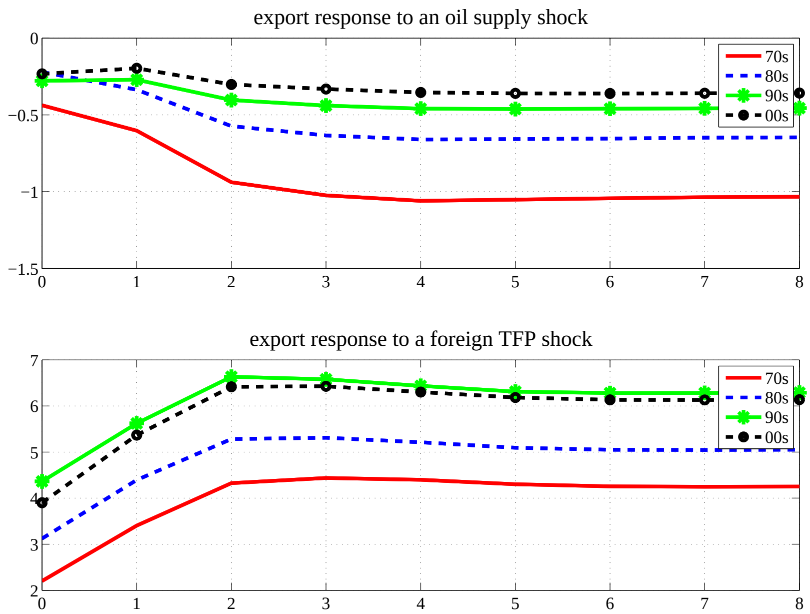

Figure 3.5: Cumulative Impulse Response Functions (IRFs) of export volumes in different decades. Each line shows the the time varying IRFs averaged over the quarters in the indicated decade. In each quarter the IRF is selected as the one that is closest to the mean IRF among those derived from the admissible structural models given the identifying restrictions. IRF to an oil supply shock are normalised to yield a 1% impact increase in oil prices. IRF to a foreign output shock are normalised to yield a 1% impact increase in foreign output.

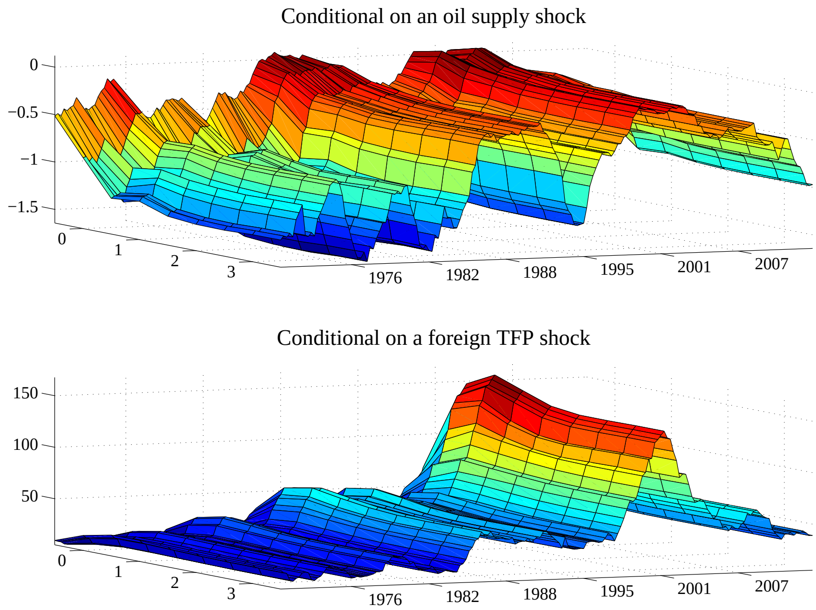

Figure 3.6: Covariances of the real price of oil and of export volumes conditional on structural shocks. Conditional covariances are obtained as the product of the cumulative IRFs of the price of oil and of exports to each structural shock, following den Haan and Sumners (2004).

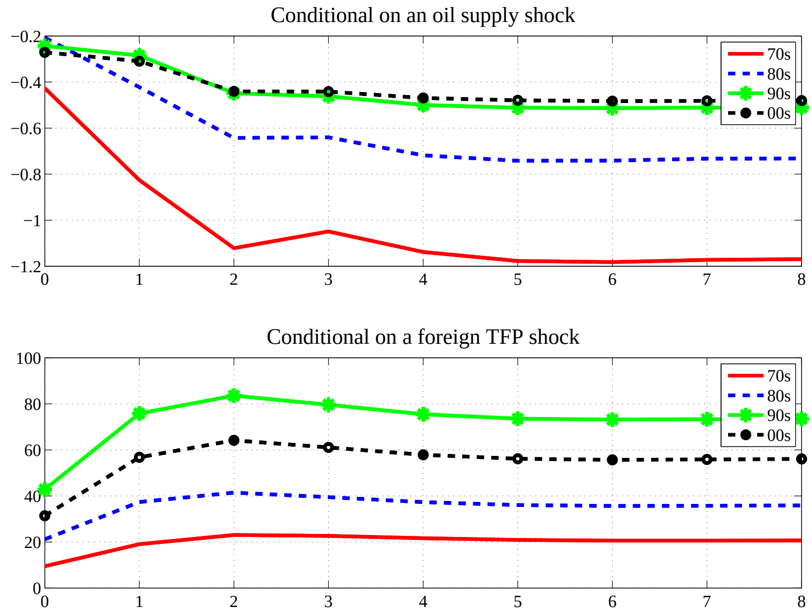

Figure 3.7: Covariances of the real price of oil and of export volumes conditional on structural shocks in different decades. Each line shows the the time varying conditional covariance averaged over the quarters in the indicated decade. Conditional covariances are obtained as the product of the cumulative IRFs of the price of oil and of exports to each structural shock, following den Haan and Sumners (2004).

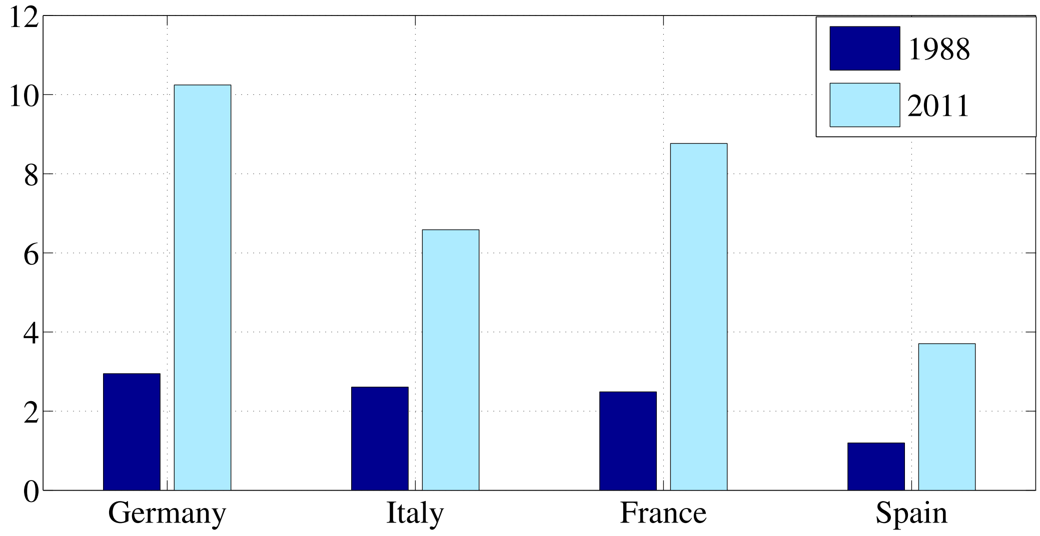

Figure 3.8: Export shares towards Asian Emerging Countries. The share of exports is computed as the ratio of exports (in current values) by destination country over total exports. Asian Emerging Countries are China, South Korea,Hong Kong, Malaysia, Singapore, Taiwan, Thailand.

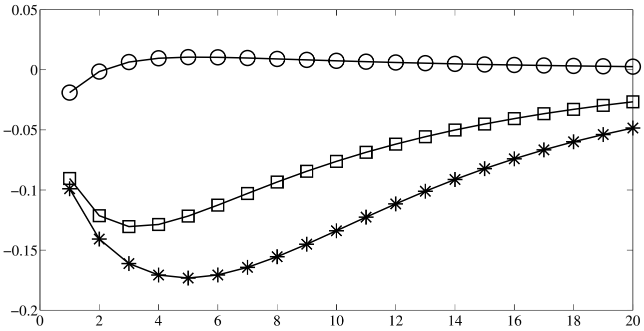

a. Stronger trade relationship with EME
− ◯ − χ^*= 0.05
− ⊡ − χ^*= 0.30
− ∗ − χ^*= 1.00

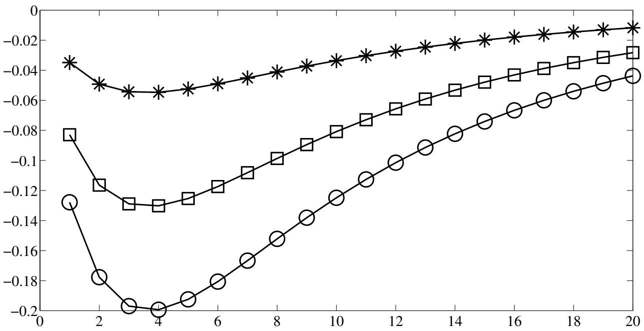

b. Lower oil shares
− ◯ − α_m= 0.040
− ⊡ − α_m= 0.025
− ∗ − α_m= 0.010

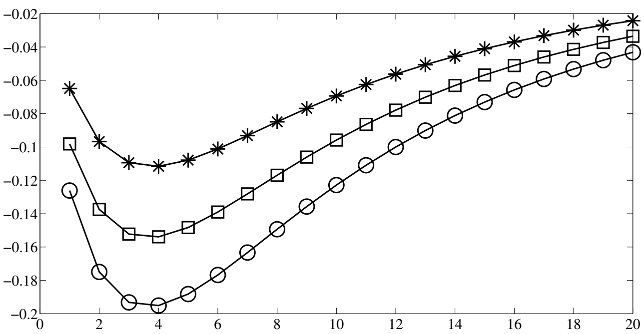

c. lower markups
− ◯ − ϵ= 3.0 (desired markup 50%)
− ⊡ − ϵ= 6.0 (desired markup 20%)
− ∗ − ϵ= 11.0 (desired markup %)

Figure 3.9: Theoretical IRFs to an oil supply shock. IRFs are obtained by simulating the theoretical model under a 10 per cent reduction in oil supply for different calibrated values of χ^*(panel a), α_m (panel b) and ϵ (panel c). The remaining parameters are calibrated consistently with the Monte Carlo simulation performed in Section 3.7.

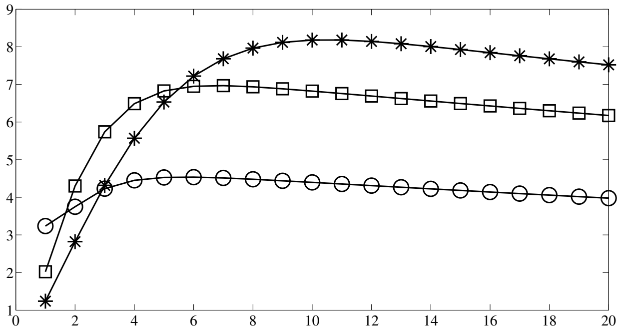

a. Stronger trade relationship with EME
$-\bigcirc-$ $\chi^{*}= 0.05$
$-\boxdot-$ $\chi^{*}= 0.30$
$-*-$ $\chi^{*}= 1.00$

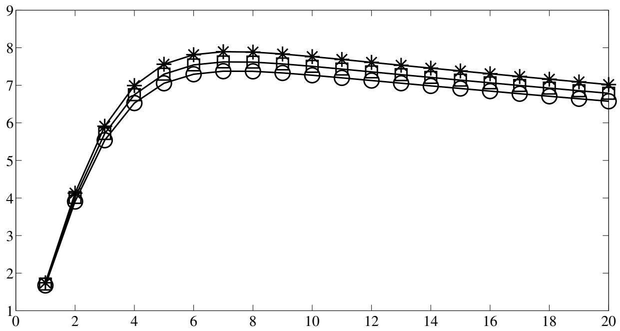

b. Lower oil shares
$-\bigcirc-$ $\alpha_m= 0.040$
$-\boxdot-$ $\alpha_m= 0.025$
$-*-$ $\alpha_m= 0.010$

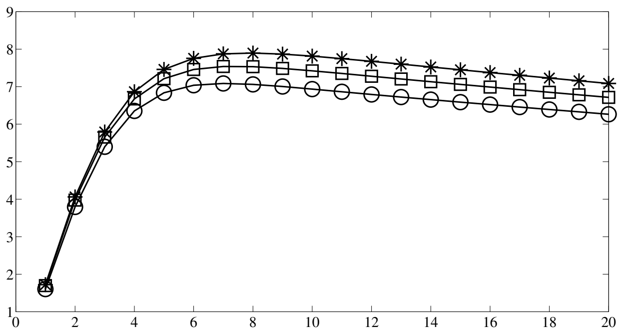

c. lower markups
$-\bigcirc-$ $\epsilon= 3.0$ (desired markup 50%)
$-\boxdot-$ $\epsilon= 6.0$ (desired markup 20%)
$-*-$ $\epsilon= 11.0$ (desired markup %)

Figure 3.10: Theoretical IRFs of exports to an oil demand shock. IRFs are obtained by simulating the theoretical model under foreign productivity shock, normalized to increase foreign GDP by 1 percent, for different calibrated values of χ^{*}(panel a), α_m (panel b) and ϵ (panel c). The remaining parameters are calibrated consistently with the Monte Carlo simulation performed in Section 3.7.

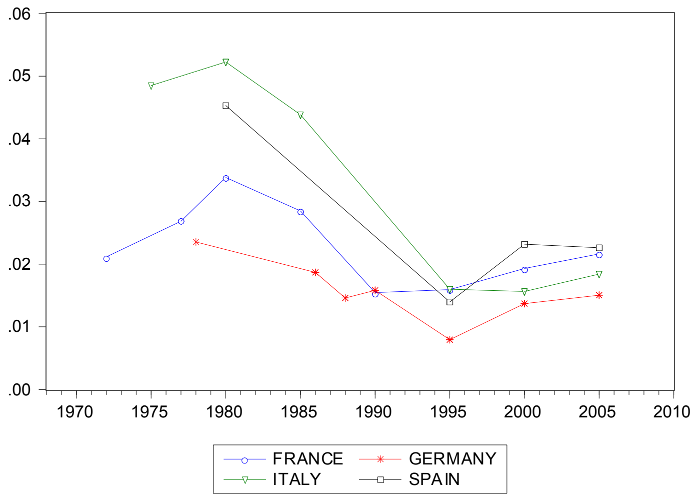

Figure 3.11: Oil shares dynamics in the euro-area. We compute oil shares by using input-output tables from the OECD database, following the methodology discussed in Blanchard and Galì (2010). The latter are available for the following years: France 1972, 1977, 1980, 1985, 1990 1995, 2000, 2005; Germany 1978, 1986, 1988, 1990, 1995, 2000, 2005; Italy 1975, 1980, 1985, 1995, 2000, 2005; Spain 1980, 1995, 2000, 2005. For each country and for each year, we proceed as follows. We split the industries into two large categories: the oil producing and the non-oil producing category. Depending on the country, the former is made up of one or more sectors. As an example, say that the oil producing category is made up of two sectors, A and B, and call C the rest of the economy, which is made up of the rest of the industries. The sum of output and imports of sector A (B) can be split between a certain amount x_A (x_B) for domestic final uses, and a certain amount y_A (y_B) for intermediates, of which z_A (z_B) goes to A and/or B and $y_A - z_A$ ($y_B - z_B$) goes to the non-oil category C. Say that the country's value added is v and the value added by A (B) is v_A (v_B). Then the shares of oil in production α_m can be computed as follows: $\alpha_m = \frac{(y_A - z_A)+(y_B - z_B)}{v-(v_A+v_B)+(y_A-z_A)+(y_B-z_B)}$

Appendix

3.A First order conditions

The optimality conditions implied by the maximization of (3.1) subject to the budget constraint are the stochastic Euler equation:

$$1 = \frac{\beta}{Q_t^B}\mathbb{E}_t\left\{\frac{P_t}{P_{t+1}}\frac{(C_t - hC_{t-1})^\sigma}{(C_{t+1} - hC_t)^\sigma}\right\} \tag{3.13}$$

and the intratemporal optimality condition:

$$\frac{W_{H,t}}{P_t} = N_t^\phi (C_t - hC_{t-1}) \tag{3.14}$$

The optimal price-setting strategy for the typical firm resetting its price in period t is:

$$E_t\left\{\sum_{k=0}^{\infty}\theta^k Q_{t,t+k} Y_{t+k/t}\left(\widetilde{P}_{H,t} - \frac{\epsilon}{\epsilon - 1}\Psi_{t+k/t}\right)\right\} = 0 \tag{3.15}$$

where $Q_{t,t+k}$ is the stochastic discount factor for nominal payoffs, $\widetilde{P}_{H,t}$ denotes the price newly set at time t, $Y_{t+k/t}$ and $\Psi_{t+k/t}$ are the level of output and the (nominal) marginal cost in period in period $t + k$ for a firm that last set its price in period t and $\frac{\epsilon}{\epsilon-1}$ measures the desired markup.

In the oil producing country, the optimal allocation of expenditures between F and H-produced goods implies:

$$C_{H,t}^O = (1 - \varpi_O)\, C_t^O \frac{P_t^O}{P_{H,t}} \tag{3.16}$$

$$C_{F,t}^O = \varpi_O C_t^O \frac{P_t^O}{P_{F,t}} \tag{3.17}$$

where P_t^O denotes the aggregate price level in the oil-producing economy $P_t^O = P_{H,t}^{1-\varpi_O} P_{F,t}^{\varpi_O}$.

Steady state relations

From cost minimization we can write:

$$\frac{M^H}{M^F} = \frac{NW_H}{N^*W_F} \tag{3.18}$$

Equation (3.14) implies that $W_H = N^{\phi}\left[C - hC\right]P$ and $W_F = N^{*\phi}\left[C^* - hC^*\right]P^*\mathcal{E}$. Note that

$$P^*\mathcal{E} = (P_H^*\mathcal{E})^{\gamma^*}(P_F^*\mathcal{E})^{(1-\gamma^*)} = (P_H)^{\gamma^*}(P_F)^{(1-\gamma^*)} \tag{3.19}$$

Accordingly we can write:

$$\frac{W_H}{W_F} = \frac{N^{\phi}}{N^{*\phi}}\frac{C}{C^*}\frac{P_H^{1-\gamma-\gamma^*}}{P_F^{1-\gamma^*-\gamma}} = \frac{N^{\phi}}{N^{*\phi}}\frac{C}{C^*}S^{-(1-\gamma^*-\gamma)} \tag{3.20}$$

where $S \equiv \frac{P_F}{P_H}$ is the *terms of trade* between the domestic economy and F, i.e. the relative price of foreign goods. Combining (3.18) and (3.20) we get:

$$\frac{M^H}{M^F} = \frac{N^{1+\phi}}{N^{*1+\phi}}\frac{C}{C^*}S^{-(1-\gamma^*-\gamma)} \tag{3.21}$$

We define $\Gamma \equiv \frac{C}{Y}$ and $\Gamma^* \equiv \frac{C^*}{Y^*}$. Accordingly

$$\frac{C}{C^*} = \frac{\Gamma Y}{\Gamma^* Y^*} = \frac{\Gamma N^{\alpha_n} M_H^{\alpha_m}}{\Gamma^* N^{*\alpha_n} M_F^{\alpha_m}} \tag{3.22}$$

Taken into account that $M_{Ht} = \frac{\alpha_m}{\mathcal{M}_t^p}\frac{Y_t}{P_{m,t}}P_{H,t}$ and $M^F = \frac{\alpha_m}{\mathcal{M}_t^{*p}}\frac{Y_t^*}{P_{m,t}}P_{F,t}$, we get:

$$\frac{Y}{Y^*S} = \frac{M^H}{M^F} \tag{3.23}$$

Combining (3.23) with (3.21) and (3.22) yields:

$$\frac{N^{1+\phi}}{N^{*1+\phi}}\frac{\Gamma}{\Gamma^*}S^{(\gamma^*+\gamma)} = 1 \tag{3.24}$$

We conclude that the steady state terms of trade can be written as:

$$S = \left[\frac{\Gamma^*}{\Gamma}\left(\frac{N^*}{N}\right)^{1+\phi}\right]^{\frac{1}{\gamma^*+\gamma}} \tag{3.25}$$

Combining (3.21) with (3.22) we can write:

$$\left(\frac{M^H}{M^F}\right)^{1-\alpha_m} = \left(\frac{N}{N^*}\right)^{1+\phi+\alpha_n} S^{-(1-\gamma^*-\gamma)}\frac{\Gamma}{\Gamma^*}$$

which implies, using (3.25):

$$\frac{M^H}{M^F} = \left(\frac{N}{N^*}\right)^{\frac{1+\phi+\alpha_n}{(1-\alpha_m)}+(1+\phi)\frac{(1-\gamma^*-\gamma)}{(\gamma^*+\gamma)(1-\alpha_m)}} \left(\frac{\Gamma}{\Gamma^*}\right)^{\frac{1}{(\gamma^*+\gamma)(1-\alpha_m)}} \tag{3.26}$$

To recover $\frac{C_H}{Y}$ and $\frac{C_F^*}{Y^*}$, we consider $P_{H,t}C_{H,t} = (1-\gamma)\, P_t C_t$ and $P^*_{F,t}C^*_{F,t} = (1-\gamma^*)\, P^*_t C^*_t$ in steady state and get:

$$\frac{C_H}{Y} = \Gamma\,(1-\gamma)\, S^{\gamma} \tag{3.27}$$

$$\frac{C^*_{F,t}}{Y^*} = \Gamma^*\,(1-\gamma^*)\, S^{-\gamma^*} \tag{3.28}$$

Note also that considering $P^*_{H,t}C^*_{H,t} = \gamma^* P^*_t C^*_t$ in steady state together with (3.23), one obtains:

$$\frac{C_H^*}{Y} = \frac{M^F}{M^H}\gamma^* S^{(-\gamma^*)}\Gamma^*$$

and similarly

$$\frac{C_F}{Y^*} = \frac{M^H}{M^F}\gamma\Gamma S^{\gamma} \tag{3.29}$$

Finally, in order to recover $\frac{C_H^O}{Y}$, let consider (3.16) together with the consumer's budget constraint:

$$\frac{C_H^O}{Y} = (1-\varpi_O)\,\frac{P_m M^s}{Y P_H}$$

Taking into account (3.7) we can write:

$$\frac{C_H^O}{Y} = (1-\varpi_O)\left(\frac{nM^H P_m}{YP_H} + \frac{(1-n)\, M^F P_m}{YP_H}\right) \tag{3.30}$$

By combining (3.30) with $M^H = \alpha_m \frac{P_H}{P_m}\frac{\epsilon-1}{\epsilon} Y$ we get:

$$\frac{C_H^O}{Y} = (1-\varpi_O)\left[n\frac{\epsilon-1}{\epsilon}\alpha_m + (1-n)\, S\frac{Y^*}{Y}\frac{\epsilon-1}{\epsilon}\alpha_m\right] \tag{3.31}$$

that can be combined with (3.23) to get:

$$\frac{C_H^O}{Y} = (1-\varpi_O)\,\frac{\epsilon-1}{\epsilon}\alpha_m\left[n + (1-n)\,\frac{M^F}{M^H}\right] \tag{3.32}$$

Similarly, in order to recover $\frac{C_F^O}{Y^*}$, let consider (3.17) together with the consumer's budget constraint:

$$\frac{C_F^O}{Y^*} = \varpi_O \frac{P_m M^s}{Y^* P_F} \tag{3.33}$$

Taking into account (3.7) and $M^H = \alpha_m \frac{P_H}{P_m} \frac{\epsilon-1}{\epsilon} Y$ we can write:

$$\frac{C_F^O}{Y^*} = \varpi_O \left(n \frac{Y}{Y^* S} \frac{\epsilon - 1}{\epsilon} \alpha_m + (1-n) \frac{\epsilon - 1}{\epsilon} \alpha_m \right) \tag{3.34}$$

Finally, considering (3.23) one obtains:

$$\frac{C_F^O}{Y^*} = \varpi_O \left(n \frac{M^H}{M^F} \frac{\epsilon - 1}{\epsilon} \alpha_m + (1-n) \frac{\epsilon - 1}{\epsilon} \alpha_m \right) \tag{3.35}$$

3.B Data and prior distributions

- Oil supply for the years 1973-2013 is world crude oil production (thousand barrels per day) reported in the April 2015 Monthly Energy Review by the U.S. Energy Information Administration. For the years 1970-1972 we use annual data from the December 2010 International Petroleum Monthly (also published by the Energy Information Administration) and convert it to the quarterly frequency using a quadratic spline.

- Euro area exports are chain linked volumes from the ECB Area Wide Model (AWM) database.

- The real price of oil is obtained by first converting in euros the price of crude Brent, then deflating the resulting nominal price (denominated in euros) by the euro area GDP deflator. The nominal price of oil in U.S. dollar (Brent quality), the euro/U.S. dollar exchange rate and the GDP deflator are all obtained from the ECB-AWM database.

- Foreign GDP is constructed following the methodology described in the Appendix in Hahn and Mestre (2011). Annual and quarterly GDP data are taken from the Economic Outlook, OECD database (GDP volumes at the price levels and PPP of 2000). Starting from the sum of annual GDP, the quarterly series is obtained by interpolating with the Chow Lin methodology. Data are available for the following countries: Australia, Austria, Brazil, Canada, China, Czech Republic, Chile, Denmark, Estonia, Finland, UK, Hungary, India, Iceland, Japan, Korea, Mexico, Norway, New Zealand, Poland, the Slovak Republic, Sweden, Turkey, and the United States.

The prior distributions for the initial values of the states (θ_0, a_0 and λ_0) of the VAR are postulated to be all normal, and independent both from one another, and from the distribution of the hyperparameters. The calibration of the mean and of the variance of the prior distribution of θ_0 follows the standard approach of using the output of the estimate of a time-invariant VAR over a training period. Since our data sample is relatively short, rather than discarding data we use the period 1970-1985, a strategy suggested for example, by Canova and Ciccarelli (2009). The rest of the procedure to set up the priors follows step by step Benati and Mumtaz (2007), pages 9 to 11. The simulation algorithm used to obtain the posterior distribution of the parameters is also taken from Benati and Mumtaz (2007), pages 11 and 12. In the Monte Carlo Markov Chain procedure we use 10.000 replications and discard the first 5.000.

Chapter 4

Failing to forecast low inflation and Phillips curve instability: a euro-area perspective

> *The ECB never expects inflation to deviate from the target of just under 2 per cent. Yet each month inflation undershoots, and the ECB is apparently taken by surprise.*
>
> Münchau W., 2014[1]

4.1 Introduction

Debate over the Phillips curve has gained momentum since the 2008 financial crisis. In the course of the recession that followed that crisis, a puzzle had emerged, in that inflation in advanced countries had not fallen as much as a traditional Phillips curve, as discussed by Williams (2010) and Ball and Mazumder (2011). The decline of euro area inflation between 2013 and 2014 is pointing in the opposite direction. Following the sovereign debt crisis, the euro area fell into a severe recession, which generated sizeable output losses in the countries more directly involved, in particular Greece, Spain, Portugal, Italy and Ireland. The recession was followed by a sharp fall in consumer price inflation, with core (net of food and energy) inflation dropping in the euro area to historically low levels in mid-2014. Two features stand out in this rapid inflation decline. First, it is broad based across countries, although relatively more intense those that have been hit the hardest by the sovereign debt crisis. Second, it was not anticipated by professional forecasters. This is particularly surprising if we

[1]Münchau, W (2014), Draghi is running out of legal ways to fix the euro, Financial Times, 17 August.

consider that the fall in economic activity that most of the euro area countries have experienced after 2011 has generated significant gaps between actual and potential output in these economies.

Two plausible explanations, not mutually exclusive, can be put forward. One is that forecasters underestimated the output gap over this horizon. This hypothesis relates to the usual difficulty of separating trends from cycles in real time, a task made even more difficult by the severity of the shock to GDP. The issue of quantifying structural and cyclical factors behind economic activity is crucial for the conduct of monetary policy and it is at the center of the policy debate, as testified by the 2014 Jackson Hole speech by ECB President Draghi.[2] A second possibility is that forecasters conditioned on an *accurate* measure of the output gap (where by *accurate* we mean the output gap that would have been available to them ex-post) but the response of inflation was stronger than estimated with data up to 2012. This second hypothesis, which has so far found less echo in the debate, is the focus of our paper. Drawing from the econometric literature, which has long identified structural breaks as the main cause of forecast failure, we investigate through structural break tests and time varying parameter models whether the recent deep and long lasting fall in economic activity has been accompanied by an increased sensitivity of euro area inflation to cyclical conditions, measured by the coefficient of the output gap in a backward looking Phillips curve.

We find that the sensitivity of inflation to business cycle conditions has indeed increased from 2013 onwards. This is consistent both with the muted response of consumer price to the global recession in 2008-2009[3] and with the sudden decrease in inflation that followed (albeit with some delay) the sovereign debt crisis. An analysis of the sub-aggregates of the consumer price index shows that this feature holds for both goods and services, i.e. tradable and non tradable products.[4]

Our findings are in line with the evidence put forward in a number of papers that investigate the inflation-unemployment relationship in the U.S. .Stock and Watson (1989), for instance, find that unemployment is more useful for predicting inflation in recessions than in booms, a feature also highlighted in Olivei and Barnes (2004). Stella and Stock (2012), using a multivariate unobserved component model that implies a time varying Phillips curve, find that since 2008 the slope of the curve has become steeper.

We provide two alternative explanations for our findings. The first is that the

[2]See `http://www.ecb.europa.eu/press/key/date/2014/html/sp140822.en.html`.

[3]For Italy, for example, estimates based on a DSGE model find that the Phillips curve was relatively flat up to 2012, see Riggi and Santoro (2015).

[4]For simplicity of exposition in the paper we will simply call *goods* the *non-energy industrial goods* subcomponent of the consumer price index.

crisis could have induced some changes in the structure of the economy that could have favoured a stronger responsiveness of prices to the output gap. We show that in a new Keynesian Phillips curve a rise in inflation cyclicality stems either from lower nominal rigidities, i.e., a higher frequency of price adjustment, or from weaker strategic complementarities in price setting, which could result from a significant fall in the number of firms in the economy. This latter channel arises in the model because an exogenous decrease in the number of firms implies lower elasticity of demand and higher desired markups. A second explanation is that even the ex-post output gap measures are *underestimating* the amount of slack in the economy. This, in turn, would be picked up as a change in the model parameters due to an omitted variable bias. We derive an estimate of the output gap that is consistent both with the observed fall in inflation and with the lower correlation between inflation and the output gap estimated before 2013. This counterfactual output gap is significantly wider, by around one third, than the one currently estimated by international Institutions. A third factor potentially at work is a downward adjustment of inflation expectations, which could be feeding back to actual inflation. The importance of this mechanism cannot be assessed within the theoretical model (given the hypothesis of rational expectations) and, in the absence of a reliable measure of expectations, it is also hard to gauge empirically, although a robustness check (in which we control for inflation *forecasts* elicited from professional forecasters) leaves unaltered our baseline results.

The paper is structured as follows. Section 4.2 motivates the paper by discussing how forecasters overestimated inflation in 2013 and 2014. Section 4.3 presents the empirical analysis. Section 4.4 discusses alternative interpretations of the evidence. Section 4.5 concludes. Appendix A provides additional material.

4.2 The inflation surprise

The pronounced slowdown in euro area inflation in 2013 and 2014 was not correctly predicted by forecasters. Figure 4.1 shows actual inflation between 2001 and 2013, together with 4 steps ahead inflation forecast errors computed (as the difference between actual and expected inflation) on the basis of the Consensus Economics survey.[5] In the figure we also present the price of oil (in euros). Three features stand out:

- Between 2001 and 2008, when consumer price inflation overshot the ECB

[5] We use the quarterly survey of professional forecasters conducted by Consensus Economics in March, June, September and December, which provides forecasts for the next seven quarters for a number of macro variables.

target and fluctuated slightly above 2.0%, professional forecasters were systematically surprised on the upside. There are two plausible explanations for this outcome. First, at the end of the Nineties, many euro area countries had pursued disinflationary policies (mainly by restraining wage growth) in order to comply with the Maastricht criteria. However, after joining the Monetary Union, these policies were relaxed, thus fostering inflation rates (Busetti, Forni, Harvey, and Venditti, 2007). Second, between 2003 and 2008, oil (and other commodity) prices were subject to a sequence of positive shocks, with brent prices more than doubling from 30 to 70 euros per barrel, providing continuous upward pressure on euro area inflation.

- After the unexpected collapse of oil prices that followed the financial crisis, inflation fell sharply and forecast errors turned negative for the whole of 2009. This was the first spell of negative errors observed since 2001. As oil prices returned to pre-crisis levels starting in 2010, forecast errors once again turned positive.
- In 2013 and 2014, following the sovereign debt crisis, inflation slowed down gradually and a second spell of negative forecast errors was recorded. Comparing the two episodes of negative inflation surprises (the one in 2009 and the one in 2013-2014) two differences can be observed. First, the most recent one is more persistent, as no sign of reversion in forecast errors has yet emerged. Second, it has occurred in the context of stable oil prices. These features suggest that professional forecasters failed to predict low inflation in the euro area because they were mostly surprised by the slackening of core (net of food and energy) inflation, i.e. the inflation component that is more related to cyclical conditions. This intuition is further reinforced by looking at oil price futures collected in February 2012, also presented in Figure 4.1, which show that the relative stability of oil prices in the next two years was largely expected by the markets, so that no negative surprise stemmed from oil commodity prices.[6]

As forecast failure in the econometric literature is frequently associated with structural breaks, we investigate whether the recent negative, persistent, inflation surprise is associated with a change in the elasticity of core inflation to the output gap in a backward looking Phillips curve, of the type commonly used to for forecasting (Stock and Watson, 2008). The next section turns to an empirical investigation of this hypothesis.

[6] Oil price futures in euros are obtained under the assumption of constant euro/US dollar exchange rate from the first quarter of 2012 onwards.

4.3 Empirical evidence

Our empirical analysis is based on the following backward looking Phillips curve:

$$\pi_t = \mu + \sum_{j=1}^{k} \beta_j \pi_{t-1} + \gamma y_{t-1} + \Gamma' z_t + \eta_t \tag{4.1}$$

where π_t indicates (quarter on quarter, seasonally adjusted and annualized) consumer price inflation, y_t is a measure of economic slack and z_t is a vector of other explanatory variables. In our application we set $k = 2$ as two lags are more than enough to capture the persistence of the inflation process.

We consider six different measures of inflation. The first three are the core Harmonized Index Consumer prices (HICP) net of food and energy (Core) and its two sub-components, goods and services. The other three are the corresponding indicators net of the impact of indirect taxation (defined CoreX, GoodX and ServicesX in the rest of the paper), which are computed by Eurostat under the assumption that indirect tax increases are passed through fully and immediately to final consumer prices. The importance of such indicator has risen in recent years, owing to the sequence of indirect taxation hikes with which a number of countries have tried to reduce fiscal deficits and restore market confidence.[7] They are therefore relevant for our study since the actual inflation rate could have been kept temporarily high by indirect tax increases.

We interact these inflation measures with output gaps computed by the European Commission (EC), the Organisation for Economic Co-operation and Development (OECD) and the International Monetary Fund (IMF). These output gaps are shown in Figure 4.2. In our analysis we consider data from 1999 to the third quarter of 2014. We choose to discard data prior to the inception of the euro motivated by the findings in Benati (2008), according to which the inflation targeting pursued by the ECB has significantly changed the statistical properties of the inflation process, so that any findings obtained using data before 1999 are unlikely to shed any light on current inflation developments.

4.3.1 End of sample instability tests

The first analysis we conduct is based on structural break tests. Since we are interested in parameters instability at the end of the sample, conventional break tests like those of Andrews, Lee, and Ploberger (1996) are not well suited to our

[7] Notice that if VAT increases are not passed through to final prices these indicators provide a lower bound of the actual inflation rate net of tax increases.

purpose, due to the fact that the number of observations in the period of potential change is low compared to the sample size. Also the extension to the end-of sample case by Andrews (2003) only has power when the change-point is known. Busetti (2012) addresses these issues and introduces a number of new tests designed to have high power at the end of the sample when the location of the break is not known a priori. The improvement in power is obtained by either limiting the possibility of a change-point to the last part of the sample or by giving increasing weight to the likelihood that a break will occur as the end of the sample is approached. In our application we will focus on two versions of the Locally Most Powerful (LMP) test proposed by Busetti (2012). These tests are designed to have power against the alternative of random walk type variation in the model parameters, a widely used assumption in models with time varying coefficients (Cogley and Sargent, 2005).

Given a linear regression like the one in equation (4.1), involving T observations and k regressors collected in a vector x_t, the LMP statistics has the following form:

$$L_\pi = \hat{\sigma}^2 (T - \pi T)^{-2} \sum_{t=\pi T+1}^{T} S_t' V^{-1} S_t$$

where $\hat{\sigma}^2 = \hat{u}_t' \hat{u}_t/(T-k)$, $\hat{u}_t$ are the regression residuals, $S_t = \sum_{j=t}^{T} \hat{u}_j x_j$, $V = T^{-1} \sum_{t=1}^{T} x_t x_t'$, and π is the last fraction of the sample where the break is supposed to have occurred. The two tests that we use are functions of this statistics and are computed as:

$$Sup - L = \underset{\pi \epsilon \Pi}{Sup} (L_\pi)$$

$$Exp - L = \log \int_{\pi \epsilon \Pi} \exp(L(\pi)) \pi d\pi$$

We apply these two tests for $\pi = 0.10$ and 0.25 (i.e. the last 10 and 25% of the sample),the fractions for which critical values have been tabulated by Busetti (2012). Overall, we consider the 18 different specifications that can be obtained by interacting the six measures of inflation with the three output gaps that we have selected. The results of the analysis are shown in Table 4.1. The table is organized in two vertical panels corresponding to the baseline specification (in which we do not add any control variable z_t) and to an alternative specification in which we add as a control variable the percentage change of non-energy import prices[8], to control for the effect of the exchange rate on consumer prices. In each cell we report 1 if the null hypothesis of coefficients stability is rejected at the 10% confidence level, 0

[8]This is estimated as the residual of a regression of the percentage change of the import deflator to the percentage change of oil prices in euros.

otherwise. The results can be summarized as follows:

1. When using the aggregate core index (Core) no evidence of instability emerges. On the contrary, when the underlying core inflation components are considered separately (Goods and Services), both the exp-L and the sup-L tests detect a break in the model parameters in the last portion of the sample, a result that holds regardless of the measure of output gap considered and whether or not import prices are included in the regression.
2. When we clean the price indices of the upward pressure of recent indirect tax increases, evidence of instability emerges also for the aggregate core inflation index (CoreX) and it is confirmed for the prices of services (ServicesX). Again this result is spread across different measures of output gap and it is not affected by the inclusion of import prices. In this case, however, evidence of a break is not picked up by the tests for the prices of goods (GoodsX).
3. Overall, a significant fraction of the stability tests (66%) suggests that some instability in the inflation-output nexus has indeed emerged in recent years. The figure is quite high especially if one considers the difficulty that break tests have in detecting parameter shifts that are slow and gradual, as evidenced by ?.

4.3.2 Time varying parameter models

To further investigate the hypothesis of parameter instability we now relax the assumption of constant parameters and specify a time varying coefficient model:

$$\pi_t = \mu_t + \sum_{j=1}^{k} \beta_{j,t}\pi_{t-1} + \gamma_t y_{t-1} + \Gamma_t' z_t + \eta_t \tag{4.2}$$

Parameter estimates will produce a path for the coefficients, therefore allowing us to gauge the direction of the change signalled by the break tests.

Given the large number of models under analysis we use a non-parametric estimator, which is computationally much less cumbersome than the Bayesian methods customarily used in the context of models with time varying parameters. The nonparametric approach has long been used in econometrics in the case of deterministic structural change. It has been recently extended to the case of stochastic time variation by Giraitis, Kapetanios, and Yates (2014). The idea of this estimator is that, in the presence of structural change, older data should be discounted in favour of more recent information. This is achieved by weighting observations with decaying weights when computing sample correlations. Collecting the right hand variables of equation (4.2) in the column vector X_t, the dependent variable in Y_t and the time varying parameters in the vector ρ_t, the estimator has the form:

$$\rho_t = \left[\sum_{j=1}^{T} \omega_{j,t} X_j X_j'\right]^{-1} \left[\sum_{j=1}^{T} \omega_{j,t} X_j Y_j\right]$$

The sample moments are therefore discounted by the function $\omega_{j,t}$:

$$\omega_{j,t} = cK\left(\frac{t-j}{H}\right) \tag{4.3}$$

where c is an integration constant and $K\left(\frac{t-j}{H}\right)$ is the kernel function determining the weight of each observation j in the estimation at time t. This weight depends on the distance to t normalized by the bandwidth H. Giraitis, Kapetanios, and Price (2013) show that the estimator has desirable frequentist properties and suggest the optimal bandwidth value $H = \sqrt{T}$. We follow their suggestion and estimate the parameters ρ_t using a Gaussian kernel and set $H = \sqrt{T}$. Although, asymptotically, the estimator is Normally distributed, we derive confidence bands via bootstrap simulations, given the low number of observations in our sample.[9]

The estimated evolution of the output gap coefficient (γ_t) is shown in Figure 4.3, which is organized in four panels. The left hand panels show estimates obtained using, respectively, core inflation (top) and core inflation net of indirect taxation (bottom) and a baseline specification with no additional control variables. The right hand panels display analogous estimates obtained controlling for import prices. In each plot we report the 15th, 50th and 85th percentiles of the empirical distribution of the estimated γ_t, together with the estimate obtained with a constant coefficient model and data up to 2012q4.[10]

In all cases, the median estimate of γ_t shows an increasing tendency from the end of 2012, to a value of around 0.25/0.30. This is almost three times as large as the estimate obtained from a fixed coefficient model. Notice that this latter estimate is also well below the 15th percentile of the empirical distribution of γ_t from 2013 onwards. To assess which component of inflation is driving these results, we inspect the estimated gap coefficients for goods and services separately, as shown in Figures 4.4 and 4.5. Results on the subcomponents are overall in line with those of the

[9] When computing confidence bands we also allow for changes in the variance of the errors η_t. At each point in time of the bootstrap simulation we therefore draw the errors from a Normal distribution with mean zero and variance σ_t^2. We estimate also σ_t^2 with nonparametric methods as suggested by Giraitis, Kapetanios, and Yates (2012): $\sigma_t^2 = 1/T \sum_{j=1}^{T} \omega_{j,t} u_j^2$.

[10] Since we have three different measures of the output gap, we account for output gap uncertainty by pooling bootstrap estimates from specifications with different output gaps and compute the percentiles on the empirical distribution of the estimated coefficients.

aggregate as the responsiveness of both goods and services prices to the business cycle has increased markedly in recent years. When controlling for import prices, in particular, a significant discontinuity appears in 2013-2014.

To explore a possible role for inflation expectations we augment the baseline specification with a forward looking inflation measure, i.e., expected inflation 6 quarters ahead, as surveyed by Consensus Economics. The results are presented in Figure 4.6 for the core index, and in Figures 4.7 and Figure 4.8 for goods and services. The inclusion of inflation forecasts results in an increase in estimation uncertainty (relatively more pronounced for the prices of goods) but does not remove the upward trend of the median estimates at the end of the sample.

Finally, a break in the inflation/output gap relationship could also involve other parameters of equation 4.2, like the intercept and the dynamics, also with detrimental effects on forecast accuracy (Hendry and Mizon, 2014). We therefore explore whether the persistence of the inflation process, measured by the sum of the autoregressive coefficients, $\beta_1 + \beta_2$, or the long run mean, $\mu_t/(1 - \beta_1 - \beta_2)$, have changed in recent years.[11] The analysis reveals that the long-run mean of core inflation has remained steady around its historical average (1.5%, Figure 4.9). Also the sum of the autoregressive coefficients, has stayed rather stable around zero since 2006, confirming the the results obtained by Benati (2008) who finds that the serial correlation of inflation is typically zero in monetary areas with a well defined nominal anchor, like the medium-term ECB inflation target.

4.4 Interpretation of the evidence

Having documented an increase in the sensitivity of inflation to the output gap we discuss possible interpretations of such evidence along two lines. First, we go through the theoretical pricing model by Sbordone (2007) and explore which changes in the structure of the economy would lead to an increase in the slope of the Phillips curve. An alternative explanation is that the nonlinearity in the parameters of the empirical model is simply indicating an underestimation of the actual output gap. On this respect we provide an estimate of the gap that would result in a stable Phillips curve.

We start from a discussion of the Phillips curve implied by the model by Sbordone (2007):[12]

$$\pi_t = \beta \mathbb{E}_t \pi_{t+1} + \zeta \widehat{s}_t \tag{4.4}$$

[11] Notice that our long-run mean estimate, obtained in the baseline specification with no control variables, implies a zero long-run forecast for the output gap, y_t.

[12] Model's details are provided in the Appendix.

where π_t denotes inflation, β is the discount factor, $\widehat{s}_t$ denotes real unit labor cost (where a hat indicates the log- deviation from the steady state) and ζ is a convolution of deep parameters capturing the sensitivity of price changes to variations in real unit labor costs, which, in this class of models, are related to the output gap by an approximate log-linear relationship. Notice that the above equation is purely forward looking, while the model used in the empirical analysis has a backward looking nature. This is not a major issue, since our aim is not taking equation (4.4) to the data, but using it to organize a discussion on the possible sources of increased inflation cyclicality.

As shown in the Appendix, the slope coefficient can be defined as:

$$\zeta \equiv \frac{(1-\alpha\beta)(1-\alpha)}{\alpha} \frac{1}{1+\overline{\theta}(N)\left[\overline{\epsilon}_{\mu}(N)+\overline{s}_{y}(N)\right]} \tag{4.5}$$

where β is the discount factor, α is the degree of price stickiness ($\frac{1}{1-\alpha}$ is the average price duration), N is the number of firms, $\overline{s}_y(N)$ denotes the elasticity of marginal costs to the firm's own output, $\overline{\theta}(N)$ is the steady state elasticity of the firm's own output demand to its relative price and $\overline{\epsilon}_{\mu}(N)$ is the elasticity of the markup function to the firm's market share evaluated at steady state.

On the basis of (4.5) we can thus disentangle the different channels through which a steepening of the Phillips curve could have occurred:

1. *Lower nominal rigidities.* More frequent prices changes (i.e., lower α) induce a steeper Phillips curve.

2. *Lower elasticity of marginal cost to the firm's own output.* To understand this mechanism suppose there is a positive shock to real marginal costs $\widehat{s}_t$. This induces an increase in prices and a loss in demand. The latter, in turn, produces a fall in marginal costs (due to decreasing returns to scale) that will partially offset the initial shock and, therefore, reduce the need to adjust prices. It follows that a lower elasticity of marginal costs to output ($\overline{s}_y(N)$) requires a relatively larger price adjustment.

3. *Lower steady state elasticity of the firm's own output demand to its relative price.* The mechanism is akin to the one described in the previous point. For a lower steady state elasticity of demand ($\overline{\theta}(N)$), the loss in demand resulting from the initial adjustment to a shock to $\widehat{s}_t$ is milder, hence inducing a relatively larger price adjustment.

4. *Lower elasticity of the markup function evaluated at steady state.* When the elasticity of substitution between differentiated goods is decreasing in

the relative quantity consumed of the variety, firms face a price elasticity of demand that is increasing in their good's relative price. This makes the desired markup increasing in the firm's relative market share (decreasing in firms' relative price). If the elasticity of the markup function evaluated at steady state ($\overline{\epsilon}_{\mu}(N)$) decreases, the Phillips curve steepens. Indeed, when the elasticity of demand is increasing in the relative price, firms are reluctant to change their price as they would face a more elastic demand curve than firms whose relative price declines as a result of price fixity.

This model therefore suggests two possible explanations for an increase in ζ. One explanation is lower nominal rigidities, i.e., a higher frequency of price adjustment (smaller α), which could have been favoured, for instance, by structural reforms in stressed countries. Empirical evidence on recent changes in the frequency of price adjustment in the euro area is, however, scarce and characterized by mixed results. Moreover, it only covers data prior to 2013. For example, for Italy, Fabiani and Porqueddu (2013) show that in the period between 2006 and 2012 the average duration of consumer prices in Italy has indeed declined to five months, from eight months between 1996 and 2001, indicating that increased sensitivity of prices to cyclical conditions might be partly accounted for by lower nominal rigidities. On the other hand, Berardi, Gautier, and Le Bihan (2013) find that during the Great Recession, the patterns of price adjustment in France were only slightly modified: the frequency, average size and dispersion of price decreases increased only marginally.[13] Ongoing research at the Eurosystem level through a new wave of the Wage Dynamics Network[14] will provide better data and more evidence on this issue.

The second explanation rests on the three remaining channels, known in the literature as *strategic complementarities*. As shown in the Appendix they vary with the number of firms; hence so does the slope of the Phillips curve. When the number of firms decreases, the steady state elasticity of demand $\overline{\theta}$ goes down (in line with the general intuition that the larger the number of goods that are traded in the market, the more likely it is that demand declines in response to a small increase in prices); this tends to increase inflation cyclicality. By contrast, the elasticity of the mark-up function $\overline{\epsilon}_{\mu}$ and the elasticity of the marginal cost to firm's own output $\overline{s}_y$ go up and this tends to result in lower inflation cyclicality. If the first effect dominates the

[13] By using the CPI research database collected by the Bureau of Labor Statistics, Vavra (2013) explores the business cycle properties of the distribution of price changes in the US and find that while price change dispersion (i.e. the second moment of the price change distribution) is strongly counter-cyclical, the rise in the frequency of adjustment during recessions is modest. Dixon, Luintel, and Tian (2014) find similar results for the UK.

[14] See `https://www.ecb.europa.eu/home/html/researcher_wdn.en.html`.

other two, inflation cyclicality will increase as N falls. To sum up:

$$\zeta \equiv \frac{(1-\alpha\beta)(1-\alpha)}{\alpha} \frac{1}{1+\overline{\theta}(\underset{+}{N})\left[\overline{\epsilon}_{\mu}(\underset{-}{N})+\overline{s}_{y}(\underset{-}{N})\right]} \tag{4.6}$$

The combination of these effects shapes the relationship between the slope of the Phillips curve and the number of firms, as shown in Figures 4.10 and 4.11 under two different calibrations for the parameters controlling the elasticity and the curvature of the demand function taken from the literature (see the Appendix for details). It turns out that under these calibrations the relationship between inflation cyclicality and the number of firms is almost everywhere negative.

A formal test of the hypothesis linking consumer prices and the number of firms in the economy is difficult because of poor data quality regarding business demography in the euro area. Keeping these caveats in mind, some preliminary analysis on available data indicates that, in the case of Italy and Spain, the sovereign debt crisis induced a significant reduction in the number of firms. This suggests that the fact that strategic complementarities played a role in the steepening of the Phillips curve cannot be ruled out.

An alternative interpretation of the increase in γ_t hinges on the fact that the output gap is a latent variable, whose measurement is rather problematic especially during a deep recession such as the one that has hit the euro area since 2011. Measurement errors in the output gap estimates might have contributed to the finding of a Phillips curve steepening in the more recent quarters. A question that arises is how wide the gap should be in order to explain the observed fall in inflation in the context of a stable Phillips curve. To provide an answer we construct alternative output gaps assuming that, starting from 2011Q3, cyclical developments have been more adverse than assessed by current estimated, and re-estimate our baseline specification until we obtain a stable estimate of γ_t . Results are shown in Figure 4.12. Red lines are the output gap estimated by the EC (upper panel) and the corresponding estimated profile of γ_t (lower panel).[15] Blue dashed lines illustrate the counterfactual output gaps and the corresponding estimates of γ_t . What emerges is that, if the finding of the increased inflation cyclicality was entirely attributable to an underestimation of the amount of spare capacity in the economy, the actual euro area output gap would be around -4%, 1.5 percentage points wider than currently measured.

[15] Results obtained on the basis of the OECD and IMF gaps are very similar.

4.5 Conclusions

The bout of disinflation between 2013 and 2014 has been broad based across the euro area and more intense in those countries that have been hit the hardest by the sovereign debt crisis. Despite the persistent economic weakness, professional forecasters largely failed to predict the decline in inflation: those surveyed by Consensus Forecast systematically over predicted average inflation for 2013 and 2014. In this paper we explore, from an empirical point of view, whether this over-prediction can be partly attributed to a structural break of inflation cyclicality. Time varying estimates of the elasticity of inflation to the output gap reveal that in 2013 and 2014 there has been a significant increase in the sensitivity of inflation to the business cycle.

A steepening of the Phillips curve might have resulted from changes in the structure of the economy. In this respect either lower nominal rigidities, due perhaps to structural reforms in stressed countries, or a decrease in strategic complementarities in price setting, related to the fall in the number of firms in the economy as a consequence of the two recent recessions, could have led to a higher elasticity of consumer prices to the output gap. An alternative explanation is that the structure of the economy has not really changed but the gap between actual and potential output is wider than currently measured. We show that a downward adjustment of the output gap by about one third could in fact rationalize the observed fall in inflation. Only more data, especially at the firm level, on wage and price setting after the Sovereign debt crisis will help to sort the issues.

Figure 4.1: Inflation, forecast errors and oil prices

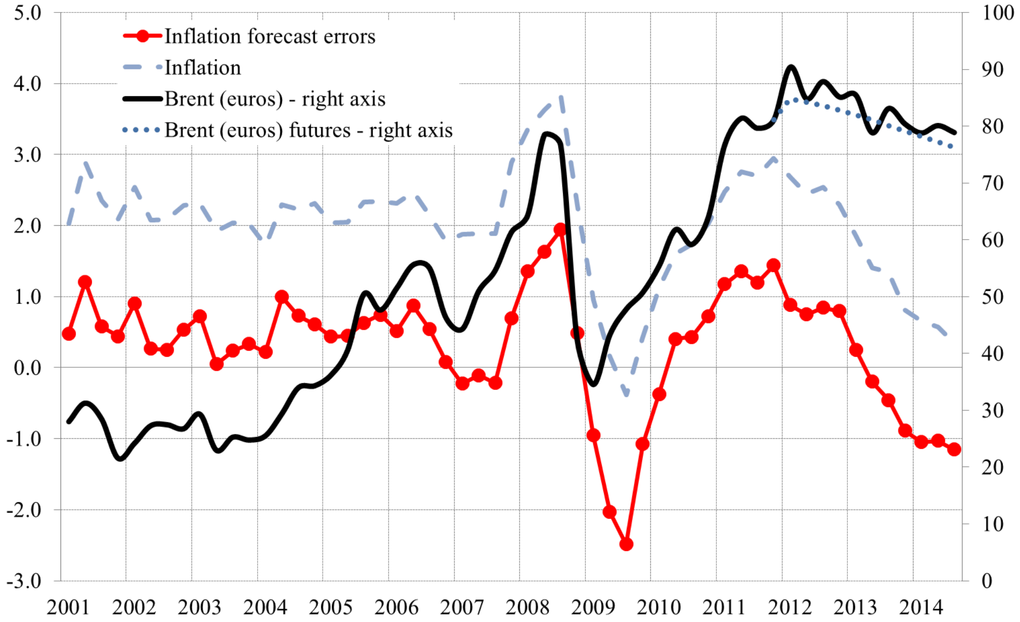

Note to Figure 4.1. The forecast errors are computed on the basis of the quarterly survey of professional forecasters conducted by Consensus Economics in March, June, September and December, which provides forecasts over the next seven quarters.

Figure 4.2: Output gaps

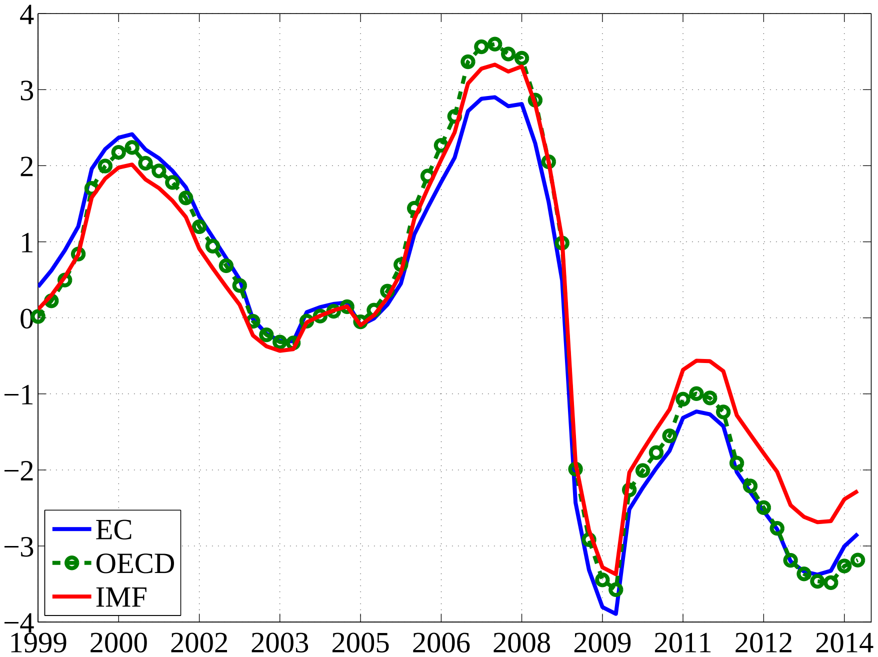

Note to Figure 4.2. EC data are from the 2014 Spring forecasts. OECD data are from the 2014 Interim Autumn Economic Assessment. IMF data are from the 2014 October World Economic Outlook. Annual data are interpolated at the quarterly frequency through a quadratic polynomial.

Table 4.1: End of sample instability tests, rejections at the 10% confidence level

		Baseline			Controlling for import prices		
Test	Π	Gap-EC	Gap-OECD	GAP-IMF	Gap-EC	Gap-OECD	GAP-IMF
			Core			Core	
exp-L	75	0	0	0	0	0	0
exp-L	90	0	0	0	0	0	0
sup-L	75	0	0	0	0	0	0
sup-L	90	0	0	0	0	0	0
			Goods			Goods	
exp-L	75	1	1	1	1	1	1
exp-L	90	1	1	1	1	1	1
sup-L	75	1	1	1	1	0	1
sup-L	90	1	1	1	1	1	1
			Services			Services	
exp-L	75	1	1	1	1	1	1
exp-L	90	1	1	1	1	1	1
sup-L	75	1	1	1	1	1	1
sup-L	90	1	1	1	1	1	1
			CoreX			CoreX	
exp-L	75	1	1	1	1	1	1
exp-L	90	1	1	1	1	1	1
sup-L	75	1	1	1	1	1	1
sup-L	90	1	1	1	1	1	1
			GoodsX			GoodsX	
exp-L	75	0	0	0	0	0	0
exp-L	90	0	0	0	0	0	0
sup-L	75	0	0	0	0	0	0
sup-L	90	0	0	0	0	0	0
			ServicesX			ServicesX	
exp-L	75	1	1	1	1	1	1
exp-L	90	1	1	1	1	1	1
sup-L	75	1	1	1	1	1	1
sup-L	90	1	1	1	1	1	1

Note to table 4.1. In each cell we report 1 if the test statistics is higher than the 10% critical values tabulated in Busetti (2012), 0 otherwise.

Figure 4.3: Slope of the Phillips curve: core inflation

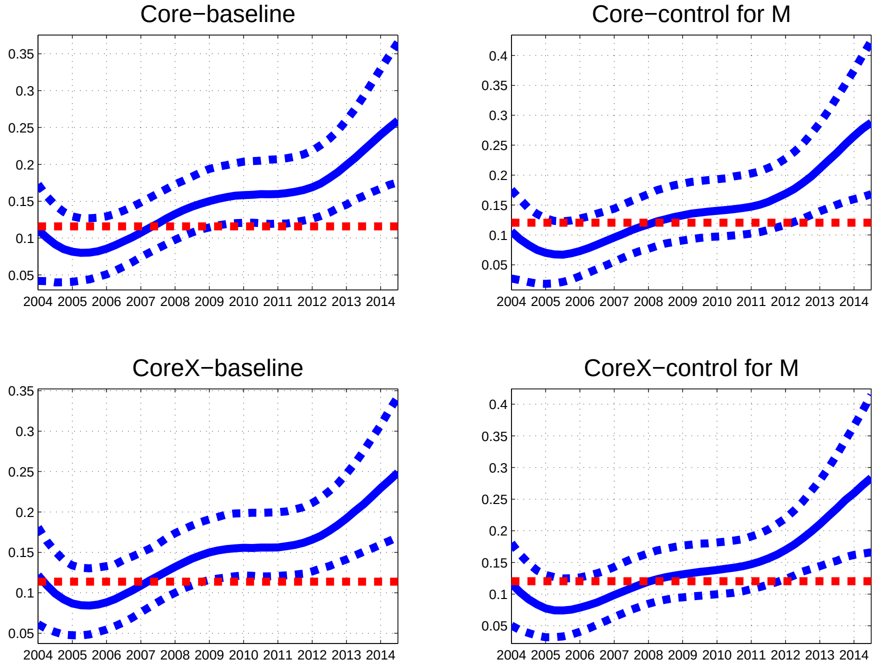

Figure 4.4: Slope of the Phillips curve: goods prices

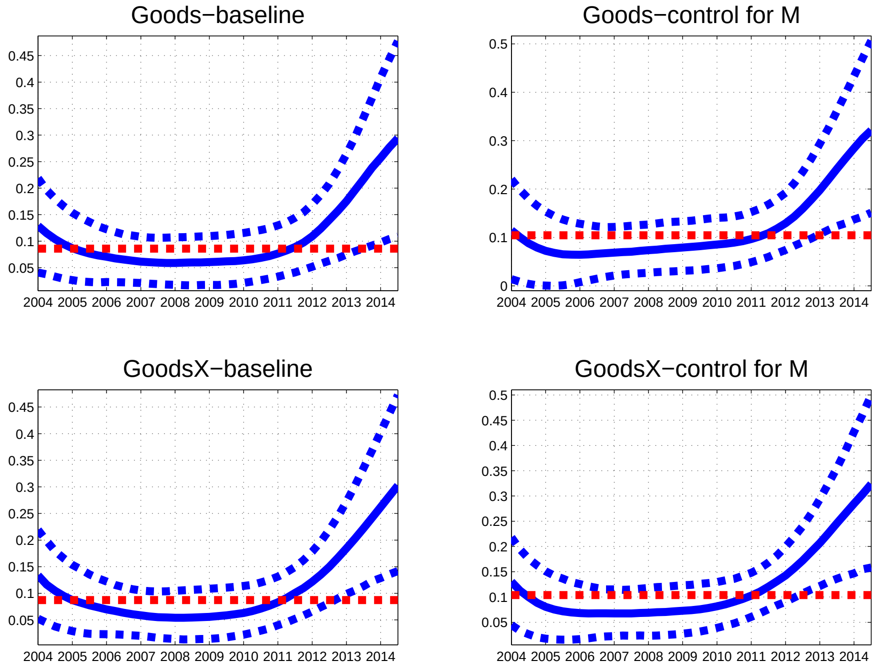

Figure 4.5: Slope of the Phillips curve: services prices

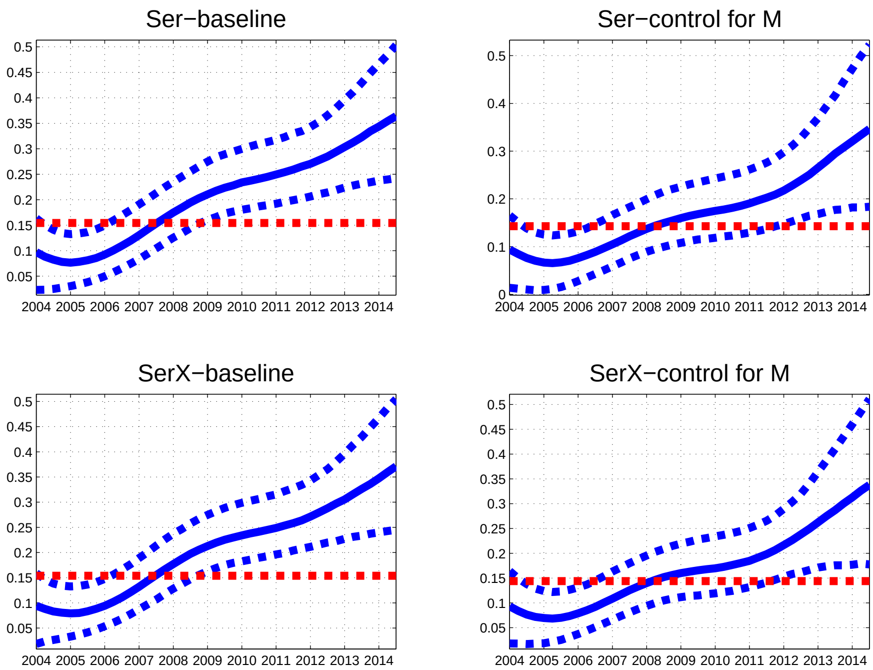

Figure 4.6: Slope of the Phillips curve: core inflation, controlling for expectations

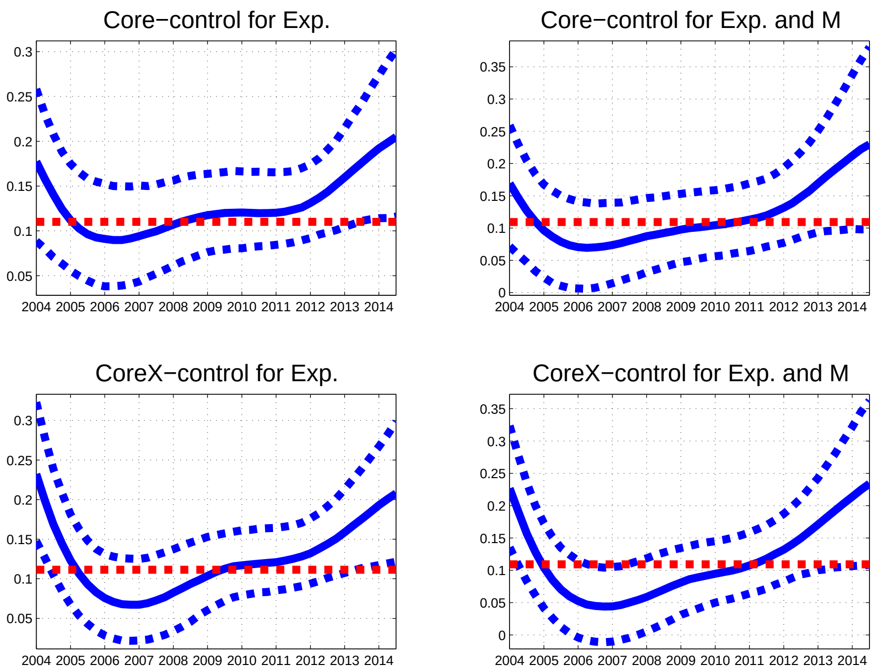

Figure 4.7: Slope of the Phillips curve: services prices, controlling for expectations

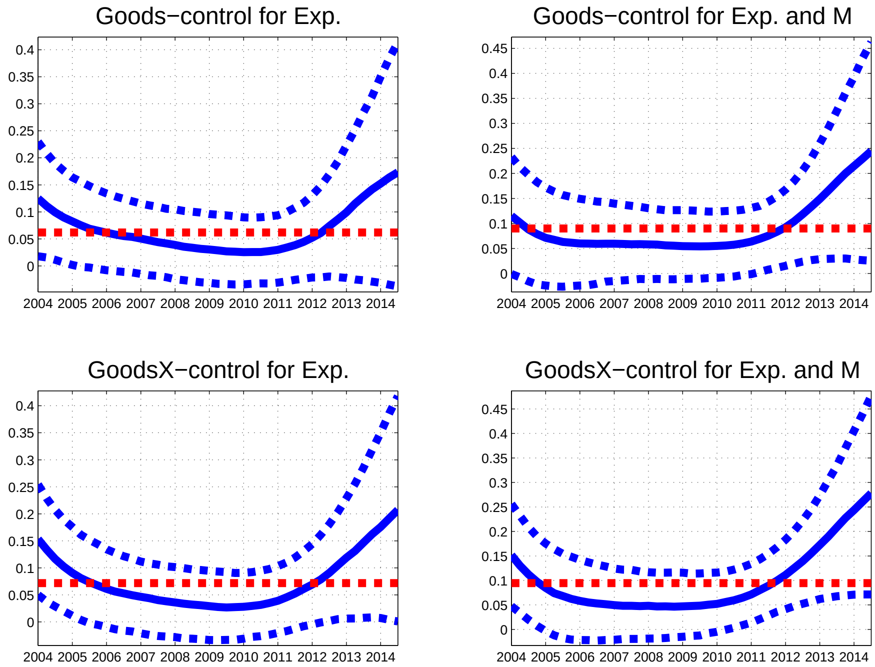

Figure 4.8: Slope of the Phillips curve: goods prices, controlling for expectations

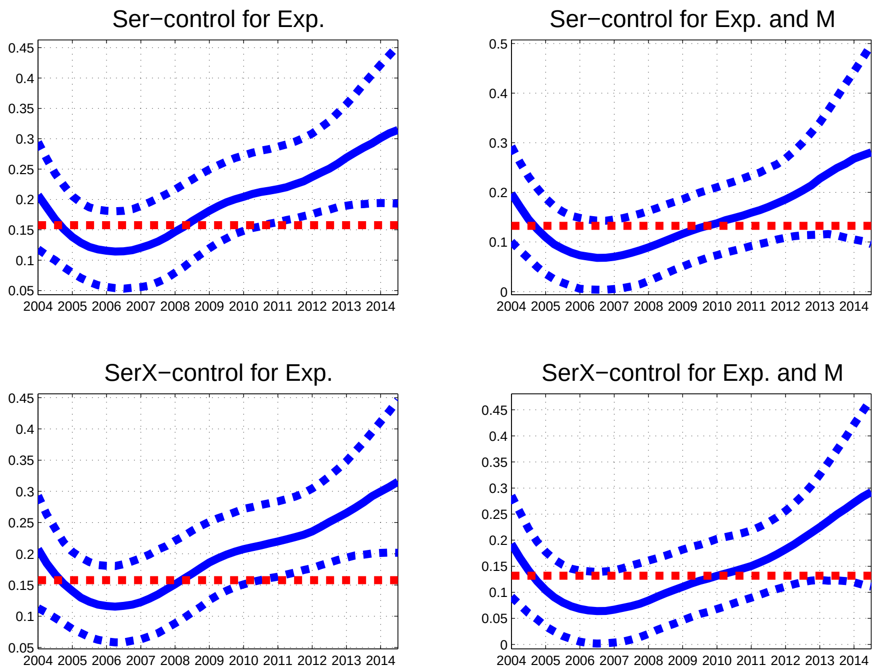

Figure 4.9: Long run mean and persistence

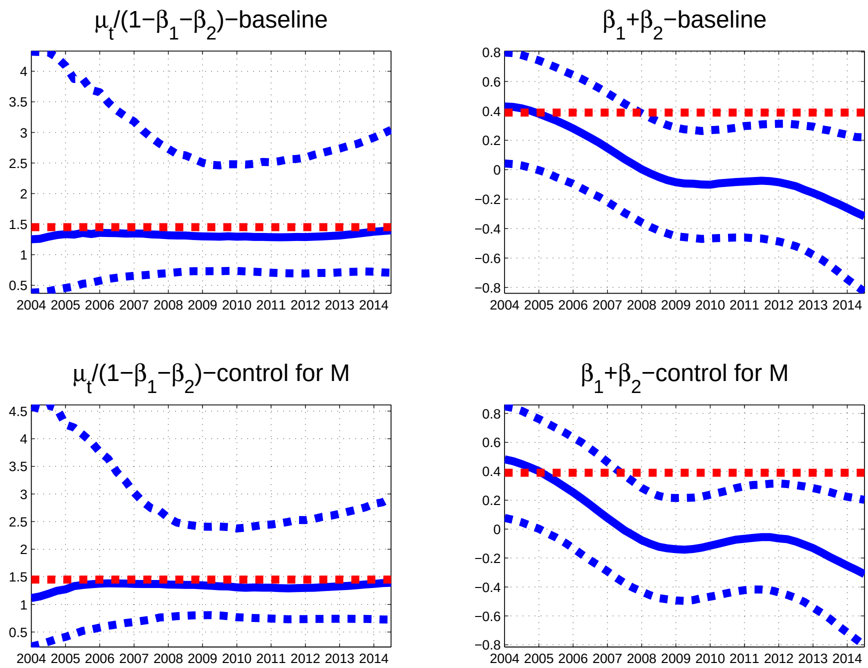

Figure 4.10: Slope of the Phillips curve and number of firms ($\eta = -2$, $\gamma = 1.14$)

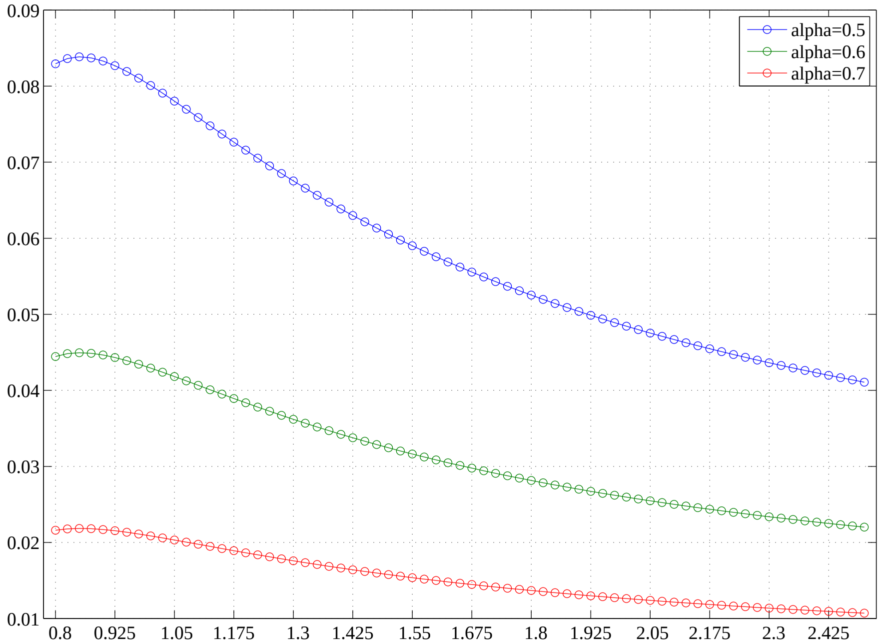

Figure 4.11: Slope of the Phillips curve and number of firms ($\eta = -3$, $\gamma = 1.07$)

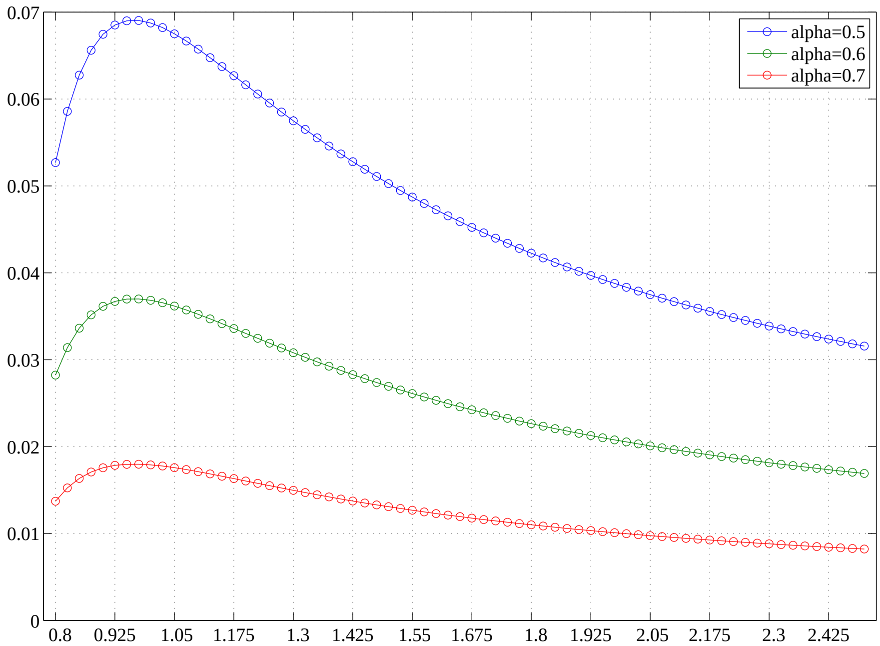

Figure 4.12: Counterfactual output gaps

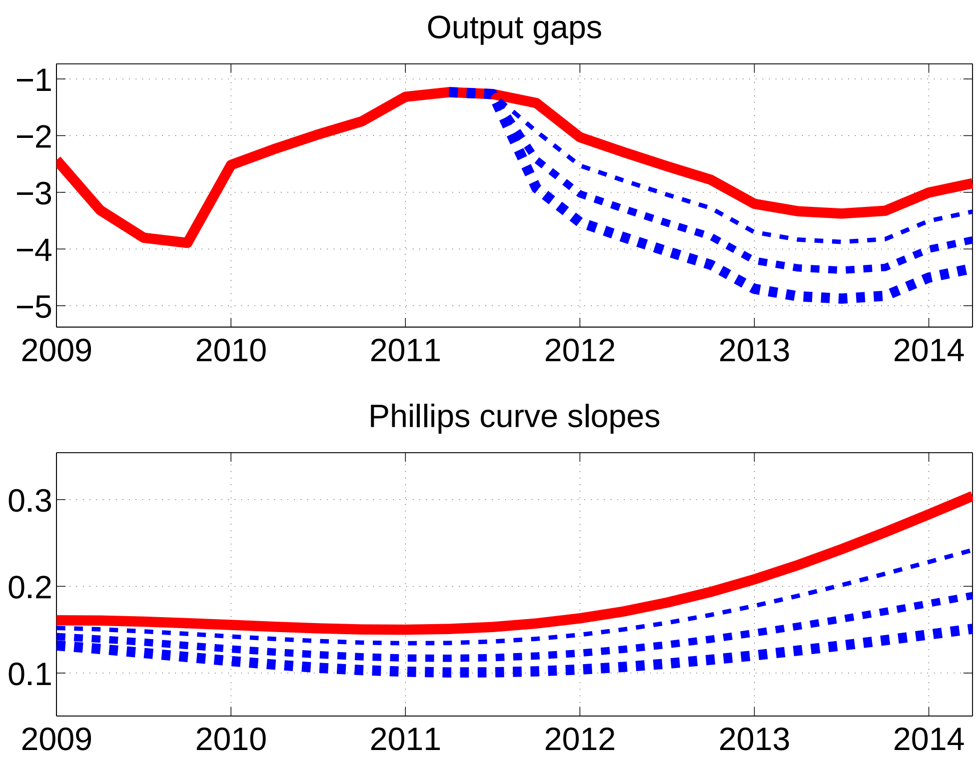

Appendix

4.A The theoretical model

We consider the theoretical framework developed by Sbordone (2007), that extends the Kimball's model in an environment where the number of firms is variable. Households' utility is defined over an aggregate C_t of differentiated goods $c_t(i)$, implicitly defined as

$$\int_{\Omega} \psi\left(\frac{c_t(i)}{C_t}\right) di = 1 \tag{4.7}$$

where $\psi(\cdot)$ is an increasing strictly concave function and Ω is the set of all potential goods produced. Note that the standard CES preferences are nested within this specification and the Kimball aggregator reduces to the Dixit-Stiglitz when $\psi\left(\frac{c_t(i)}{C_t}\right) = \left(\frac{c_t(i)}{C_t}\right)^{\frac{\theta-1}{\theta}}$ for some $\theta > 1$.

Each firm produces a differentiated good. We assume that the set of firms is $[0, N]$ and, thus, $c_t(i) = 0 \ \forall i > N$. The household must decide how to allocate its consumption expenditures among the different goods: $\min_{\{c_t(i)\}} \int_0^N p_t(i) c_t(i)\, di$ s.t. $\int_0^N \psi\left(\frac{c_t(i)}{C_t}\right) di = 1$. From the FOC to this problem one obtains the demand for each good i:

$$c_t(i) = C_t \psi'^{-1}\left(p_t(i) \Lambda_t C_t\right) \ \forall i \in [0, N]$$

where Λ_t is the Lagrangian multiplier for constraint (4.7), that is implicitly defined by $\int_0^N \psi\left(\psi'^{-1}\left(p_t(i)\Lambda_t C_t\right)\right) di = 1$

The aggregate price index is the cost of a unit of the composite good: $P_t \equiv \frac{1}{C_t} \int_0^N p_t(i) c_t(i) di$.

We assume that firm i produces with the following technology:

$$y_t(i) = h_t(i)^{1-a} - \Phi \tag{4.8}$$

where Φ is a fixed cost. Accordingly firm's i real marginal cost $s_t(i)$ is:

$$s_t\left(y_t(i);\Gamma_t\right) = \frac{1}{1-a}\frac{W_t}{P_t}\left(y_t(i)+\Phi\right)^{\frac{a}{1-a}} \tag{4.9}$$

where Γ_t indicates aggregate variables that enter into the determination of firms' marginal costs, W_t, is nominal wage and P_t is the aggregate price.

Following the formalism proposed in Calvo (1983), each firm may reset its price only with probability $(1-\alpha)$ in any given period, independently of the time elapsed since the last adjustment ($\frac{1}{1-\alpha}$ is the expected average duration of prices). A firm re-optimizing in period t will choose the price $p_t(i)$ by maximizing the expected string of profits over the life of the set price.

$$\mathbb{E}_t\left\{\sum_{j=0}^{\infty}\alpha^j Q_{t,t+j}\left[p_t(i)Y_{t+j}\psi'^{-1}\left(\frac{p_t(i)}{\widetilde{P}_{t+j}}\right) - C\left(Y_{t+j}\psi'^{-1}\left(\frac{p_t(i)}{\widetilde{P}_{t+j}}\right);\Gamma_{t+j}\right)\right]\right\} \tag{4.10}$$

where $C\left(\cdot\right)$ is the firm's cost function and $\widetilde{P}_t \equiv \frac{1}{\Lambda_t P_t}$. Combining the first order condition associated with the problem above with the aggregate price dynamics yields the following Phillips curve:

$$\pi_t = \beta\mathbb{E}_t\pi_{t+1} + \zeta\widehat{s}_t \tag{4.11}$$

where π_t denotes inflation, β is the discount factor, $\widehat{s}_t$ denotes real unit labor cost (where a hat indicates the log- deviation from the steady state) and ζ is a convolution of deep parameters. Our goal is to evaluate how the number of producing firms N affects the slope coefficient ζ. Let define $x = \psi'^{-1}\left(\frac{1}{N}\right)$ the relative share in the symmetric steady state, i.e. a steady state with symmetric prices $(p_t(i) = p_t \forall i)$. Then, we can define the steady state elasticity of demand:

$$\overline{\theta} = -\frac{\psi'(x)}{x\psi''(x)} \tag{4.12}$$

and the elasticity of the mark-up function evaluated at steady state:

$$\bar{\epsilon}_\mu = \frac{x\mu'(x)}{\mu(x)}$$

where $\overline{\mu} \equiv \frac{\overline{\theta}}{\overline{\theta}-1}$ is the steady state desired markup. The slope coefficient can be defined as:

$$\zeta \equiv \frac{(1-\alpha\beta)(1-\alpha)}{\alpha}\frac{1}{1+\overline{\theta}(N)\left[\bar{\epsilon}_\mu(N)+\overline{s}_y(N)\right]} \tag{4.13}$$

where $\overline{s}_y(N) = \frac{a}{1-a}\left[\frac{xY}{xY+\Phi}\right]$ denotes the elasticity of the marginal cost to firm's own output (Y is steady state aggregate output). We now turn to examine how the number of firms N affects these channels. To this aim we need to choose a functional form for $\psi(x)$. As in Sbordone (2007) we assume the one proposed by Dotsey and King (2005).

$$\psi(x) = \frac{1}{(1+\eta)\,\gamma}\left[(1+\eta)\,x-\eta\right]^{\gamma} - \frac{1}{(1+\eta)\,\gamma}\,(-\eta)^{\gamma} \tag{4.14}$$

In this case the steady state relative share x is:

$$x \equiv \psi^{-1}(\frac{1}{N}) = \frac{1}{1+\eta}\left\{\left(\frac{(1+\eta)\,\gamma}{N} + (-\eta)^{\gamma}\right)^{\frac{1}{\gamma}} + \eta\right\} \tag{4.15}$$

The latter is clearly decreasing in N. The steady state mark-up $\overline{\mu}$ is:

$$\overline{\mu} = \frac{\eta-(1+\eta)\,\psi^{-1}(\frac{1}{N})}{\eta-\gamma\,(1+\eta)\,\psi^{-1}(\frac{1}{N})}$$

In order to see the dependence of the slope on N, we need to study how $\overline{\theta}(N)$, $\overline{\epsilon}_{\mu}(N)$ and $\overline{s}_y(N)$ vary with N. The steady state elasticity $\overline{\theta}$ is:

$$\overline{\theta} = \frac{\eta-(1+\eta)\,\psi^{-1}(\frac{1}{N})}{(\gamma-1)\,(1+\eta)\,\psi^{-1}(\frac{1}{N})} \tag{4.16}$$

which is decreasing in the steady state relative share x and, thus, increasing in N. This is in line with the general intuition that more goods are traded in a market more likely it is for the demand to decrease more in response to a small increase in prices. The elasticity of mark-up $\overline{\epsilon}_{\mu}$, that determines how much the steady state mark-up varies for small variation in N, is

$$\overline{\epsilon}_{\mu} = \frac{\eta\,(\gamma-1)\,(1+\eta)\,\psi^{-1}(\frac{1}{N})}{\left[\eta-(1+\eta)\,\psi^{-1}(\frac{1}{N})\right]\left[\eta-\gamma\,(1+\eta)\,\psi^{-1}(\frac{1}{N})\right]} \tag{4.17}$$

It can be demonstrated that this elasticity is a decreasing function of N.[16] Finally the elasticity of the marginal cost to firm's own output is:

$$\overline{s}_y(N) = \frac{a}{1-a}\left[\frac{xY}{xY+\Phi}\right] \tag{4.18}$$

It can be demonstrated that, assuming a fairly standard log-utility $u(C,h) = \log C - \frac{1}{1+\nu}h^{1+\nu}$, the steady state aggregate output is the solution to $xY + \Phi =$

[16] Indeed $\frac{\partial log\mu}{\partial \log N} = -\overline{\epsilon}_{\mu}\frac{\partial \log x}{\partial \log N}$ and $log\mu$ is a convex function of $logx$. Indeed because $\mu\,(x)$ is an increasing function of x, it is not possible for $log\mu$ to be a concave function of $logx$ as this would require $log\mu$ to be negative for positive and small enough x. If $log\mu$ must be convex at least for small values of x, it is convenient to assume that it is globally convex function of $logx$.

$\left[\frac{1-a}{\overline{\mu}xYN^{1+v}}\right]^{\frac{1-a}{\nu+a}}$ and $\overline{s}_y(N)$ is decreasing in N. To sum up:

$$\zeta \equiv \frac{(1-\alpha\beta)(1-\alpha)}{\alpha} \frac{1}{1+\overline{\theta}(\underset{+}{N})\left[\overline{\epsilon}_\mu(\underset{-}{N})+\overline{s}_y(\underset{-}{N})\right]} \tag{4.19}$$

To evaluate the quantitative impact of changes in the number of firms on inflation cyclicality, the function $\psi(x)$ needs to be parameterized, i.e. one has to choose particular values for the parameters γ and η. While the literature does not provide much guidance for what are the most plausible values for these two parameters, we follow Sbordone (2007) and Levin, Lopez-Salido, and Yun (2007) by choosing a combination of them that guarantees a plausible value for the mark-up in steady state $\overline{\mu}$, where the relative share x is equal to 1. Figures 10 and 11 in the main text show the relationship between the slope of the Phillips curve and the number of firms under two different combinations of γ and η, both consistent with a steady state markup of 16%: $\eta = -2$ and $\gamma = 1.14$ and $\eta = -3$ and $\gamma = 1.07$.

Bibliography

AASTVEIT, K. A. (2014): "Oil price shocks in a data-rich environment," *Energy Economics*, 45(C), 268–279. 75

AASTVEIT, K. A., A. CARRIERO, T. E. CLARK, AND M. MARCELLINO (2014): "Have Standard VARs Remained Stable since the Crisis?," Working Paper 1411, Federal Reserve Bank of Cleveland. 51

ALKHAMISI, M., AND G. SHUKUR (2008): "Developing Ridge Parameters for SUR Model," *Communications In Statistics-Theory And Methods*, 37(4), 544–564. 97

ANDREWS, D. W. K. (2003): "End-of-Sample Instability Tests," *Econometrica*, 71(6), 1661–1694. 145

ANDREWS, D. W. K., I. LEE, AND W. PLOBERGER (1996): "Optimal changepoint tests for normal linear regression," *Journal of Econometrics*, 70(1), 9–38. 144

ANGELINI, E., G. CAMBA-MENDEZ, D. GIANNONE, L. REICHLIN, AND G. RUNSTLER (2011): "Short-term forecasts of euro area GDP growth," *Econometrics Journal*, 14(1), C25–C44. 15

BAI, J., AND S. NG (2008): "Forecasting economic time series using targeted predictors," *Journal of Econometrics*, 146(2), 304–317. 40

BALL, L., AND S. MAZUMDER (2011): "Inflation Dynamics and the Great Recession," *Brookings Papers on Economic Activity*, 42(1 (Spring), 337–405. 140

BANBURA, M., D. GIANNONE, AND L. REICHLIN (2010): "Large Bayesian vector auto regressions," *Journal of Applied Econometrics*, 25(1), 71–92. 48, 50, 51, 61, 68, 72

BANBURA, M., AND M. MODUGNO (2010): "Maximum likelihood estimation of factor models on data sets with arbitrary pattern of missing data," Working Paper Series 1189, European Central Bank. 15, 16, 20, 24, 45

Banbura, M., and G. Runstler (2011): "A look into the factor model black box: Publication lags and the role of hard and soft data in forecasting GDP," *International Journal of Forecasting*, 27(2), 333–346. 14, 22

Baumeister, C., P. Liu, and H. Mumtaz (2013): "Changes in the effects of monetary policy on disaggregate price dynamics," *Journal of Economic Dynamics and Control, Elsevier*, 37(3), 543–560. 15

Baumeister, C., and G. Peersman (2013): "Time-Varying Effects of Oil Supply Shocks on the US Economy," *American Economic Journal: Macroeconomics*, 5(4), 1–28. 51, 74, 102, 113, 116

Benati, L. (2007): "Drift and breaks in labor productivity," *Journal of Economic Dynamics and Control*, 31(8), 2847–2877. 104

——— (2008): "Investigating Inflation Persistence Across Monetary Regimes," *The Quarterly Journal of Economics*, 123(3), 1005–1060. 144, 148

Benati, L., and H. Mumtaz (2007): "U.S. evolving macroeconomic dynamics: a structural investigation," Working Paper Series 0746, European Central Bank. 113

Berardi, N., E. Gautier, and H. Le Bihan (2013): "More Facts about Prices: France Before and During the Great Recession," Working papers 425, Banque de France. 150

Berkowitz, J. (2001): "Testing Density Forecasts, with Applications to Risk Management," *Journal of Business and Economic Statistics*, 1(19), 465–474. 27

Blanchard, O. J., and J. Gali (2007): "The Macroeconomic Effects of Oil Price Shocks: Why are the 2000s so different from the 1970s?," in *International Dimensions of Monetary Policy*, NBER Chapters, pp. 373–421. National Bureau of Economic Research, Inc. 51, 74, 75, 76, 101, 104, 121

Blanchard, O. J., and M. Riggi (2013): "Why are the 2000s so different from the 1970s? A structural interpretation of changes in the macroeconomic effects of oil prices," *Journal of the European Economic Association*, 11(5), 1032–1052. 51, 74, 75, 76, 101

Boivin, J., and S. Ng (2006): "Are more data always better for factor analysis?," *Journal of Econometrics*, 132(1), 169–194. 40, 72

Busetti, F. (2012): "On detecting end-of-sample instabilities," Temi di discussione (Economic working papers) 881, Bank of Italy, Economic Research and International Relations Area. 145, 154

Busetti, F., L. Forni, A. Harvey, and F. Venditti (2007): "Inflation Convergence and Divergence within the European Monetary Union," *International Journal of Central Banking*, 3(2), 95–121. 143

Calvo, G. A. (1983): "Staggered prices in a utility-maximizing framework," *Journal of Monetary Economics*, 12(3), 383–398. 162

Camacho, M., and G. Perez-Quiros (2010): "Introducing the euro-sting: Short-term indicator of euro area growth," *Journal of Applied Econometrics*, 25(4), 663–694. 14, 20, 41

Campolmi, A. (2008): "Oil price shocks: Demand vs Supply in a two-country model," MNB Working Papers 2008/5, Magyar Nemzeti Bank (Central Bank of Hungary). 105, 106

Canova, F., and M. Paustian (2011): "Business cycle measurement with some theory," *Journal of Monetary Economics*, 58(4), 345–361. 109

Carriero, A., T. Clark, and M. Marcellino (2016): "Large Vector Autoregressions with asymmetric priors and time-varying volatilities," manuscript, Queen Mary University of London. 53

Carriero, A., G. Kapetanios, and M. Marcellino (2009): "Forecasting exchange rates with a large Bayesian VAR," *International Journal of Forecasting*, 25(2), 400–417. 60, 63

Carter, K., and R. Kohn (1994): "On Gibbs Sampling for State Space Models," *Biometrika*, 81(2), 541–553. 5, 40

Clarida, R., J. Gali, and M. Gertler (2002): "A simple framework for international monetary policy analysis," *Journal of Monetary Economics*, 49(5), 879–904. 105

Clark, T. E. (2011): "Real-Time Density Forecasts From Bayesian Vector Autoregressions With Stochastic Volatility," *Journal of Business & Economic Statistics*, 29(3), 327–341. 15, 16, 27

Clements, M. P., and A. B. Galvão (2008): "Macroeconomic Forecasting With Mixed-Frequency Data," *Journal of Business & Economic Statistics*, 26, 546–554. 15

Cogley, T., and T. J. Sargent (2005): "Drift and Volatilities: Monetary Policies and Outcomes in the Post WWII U.S," *Review of Economic Dynamics*, 8(2), 262–302. 5, 49, 58, 65, 145

D'Agostino, A., D. Giannone, and P. Surico (2006): "(Un)Predictability and macroeconomic stability," Working Paper Series 0605, European Central Bank. 69

Davis, S. J., and J. Haltiwanger (2001): "Sectoral job creation and destruction responses to oil price changes," *Journal of Monetary Economics*, 48(3), 465–512. 74

De Mol, C., D. Giannone, and L. Reichlin (2008): "Forecasting using a large number of predictors: Is Bayesian shrinkage a valid alternative to principal components?," *Journal of Econometrics*, 146(2), 318–328. 48, 56, 70

Dedola, L., and S. Neri (2007): "What does a technology shock do? A VAR analysis with model-based sign restrictions," *Journal of Monetary Economics*, 54(2), 512–549. 109

Del Negro, M., and C. Otrok (2008): "Dynamic factor models with time-varying parameters: measuring changes in international business cycles," Staff Reports 326, Federal Reserve Bank of New York. 15, 18, 21, 38

Delle Monache, D., and I. Petrella (2014): "Adaptive Models and Heavy Tails," Birkbeck Working Papers in Economics and Finance 1409, Birkbeck, Department of Economics, Mathematics & Statistics. 66

den Haan, W. J., and S. W. Sumner (2004): "The comovement between real activity and prices in the G7," *European Economic Review*, 48(6), 1333–1347. 118

Diebold, F. X., and R. S. Mariano (1995): "Comparing Predictive Accuracy," *Journal of Business & Economic Statistics*, 13(3), 253–63. 70

Dixon, H. D., K. B. Luintel, and K. Tian (2014): "The impact of the 2008 crisis on UK prices: what we can learn from the CPI microdata," Cardiff Economics Working Papers E2014/7, Cardiff University, Cardiff Business School, Economics Section. 150

Dotsey, M., and R. G. King (2005): "Implications of state-dependent pricing for dynamic macroeconomic models," *Journal of Monetary Economics*, 52(1), 213–242. 163

Edelstein, P., and L. Kilian (2009): "How sensitive are consumer expenditures to retail energy prices?," *Journal of Monetary Economics*, 56(6), 766–779. 51, 74, 75, 76, 101, 104, 120, 121

Eickmeier, S., W. Lemke, and M. Marcellino (2015): *Journal of the Royal Statistical Society, Series A, forthcoming.* 49

Fabiani, S., and M. Porqueddu (2013): "La flessibilitá dei prezzi in Italia: evidenze per il periodo 2006-2012," mimeo, Bank of Italy. 150

Foroni, C., and M. Marcellino (2014): "A comparison of mixed frequency approaches for nowcasting Euro area macroeconomic aggregates," *International Journal of Forecasting, Elsevier*, 30(3), 554–568. 15

Frale, C., M. Marcellino, G. L. Mazzi, and T. Proietti (2011): "EUROMIND: a monthly indicator of the euro area economic conditions," *Journal of the Royal Statistical Society Series A*, 174(2), 439–470. 28, 42

Fry, R., and A. Pagan (2011): "Sign Restrictions in Structural Vector Autoregressions: A Critical Review," *Journal of Economic Literature*, 49(4), 938–60. 116, 117

Ghysels, E., P. Santa-Clara, and R. Valkanov (2004): "The MIDAS Touch: Mixed Data Sampling Regression Models," CIRANO Working Papers 2004s-20, CIRANO. 15

Giannone, D., L. Reichlin, and D. Small (2008): "Nowcasting: The real-time informational content of macroeconomic data," *Journal of Monetary Economics*, 55(4), 665–676. 15, 20

Giraitis, L., G. Kapetanios, and S. Price (2013): "Adaptive forecasting in the presence of recent and ongoing structural change," *Journal of Econometrics*, 177(2), 153–170. 49, 63, 147

Giraitis, L., G. Kapetanios, K. Theodoridis, and T. Yates (2014): "Estimating time-varying DSGE models using minimum distance methods," Bank of England working papers 507, Bank of England. 49

Giraitis, L., G. Kapetanios, and T. Yates (2012): "Inference on multivariate stochastic time varying coefficient and variance models," mimeo. 5, 49, 53, 75, 147

——— (2014): "Inference on stochastic time-varying coefficient models," *Journal of Econometrics*, 179(1), 46–65. 5, 49, 58, 77, 146

Hahn, E., and R. Mestre (2011): "The role of oil prices in the euro area economy since the 1970s," Working Paper Series 1356, European Central Bank. 74

Hamilton, J. D. (1988): "A Neoclassical Model of Unemployment and the Business Cycle," *Journal of Political Economy*, 96(2), 593–617. 74

——— (2009): "Causes and Consequences of the Oil Shock of 2007-08," *Brookings Papers on Economic Activity*, 40(1 (Spring), 215–283. 105

Hastie, T., R. Tibshirani, and J. Friedman (2003): *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*, no. 9780387952840 in Statistics. Springer. 54

Hendry, D. F., and G. E. Mizon (2014): "Unpredictability in economic analysis, econometric modeling and forecasting," *Journal of Econometrics*, 182(1), 186–195. 148

Higgins, M., T. Klitgaard, and R. Lerman (2006): "Recycling petrodollars," *Current Issues in Economics and Finance*, 12(Dec). 119

Hooker, M. A. (1999): "Oil and the macroeconomy revisited," Finance and Economics Discussion Series 1999-43, Board of Governors of the Federal Reserve System (U.S.). 51, 101, 104

Inoue, A., and L. Kilian (2013): "Inference on impulse response functions in structural VAR models," *Journal of Econometrics*, 177(1), 1–13. 117

Jacquier, E., N. G. Polson, and P. E. Rossi (1994): "Bayesian Analysis of Stochastic Volatility Models," *Journal of Business & Economic Statistics*, 12(4), 371–89. 5, 19, 40

Jore, A. S., J. Mitchell, and S. P. Vahey (2010): "Combining forecast densities from VARs with uncertain instabilities," *Journal of Applied Econometrics*, 25(4), 621–634. 15

Kapetanios, G., V. Labhard, and S. Price (2008): "Forecasting Using Bayesian and Information-Theoretic Model Averaging: An Application to U.K. Inflation," *Journal of Business & Economic Statistics*, 26, 33–41. 61, 62, 63

KILIAN, L. (2008): "A Comparison of the Effects of Exogenous Oil Supply Shocks on Output and Inflation in the G7 Countries," *Journal of the European Economic Association*, 6(1), 78–121. 118

——— (2009): "Not All Oil Price Shocks Are Alike: Disentangling Demand and Supply Shocks in the Crude Oil Market," *American Economic Review*, 99(3), 1053–69. 75, 76, 101, 102, 106, 114

KILIAN, L., AND D. P. MURPHY (2012): "Why Agnostic Sign Restrictions Are Not Enough: Understanding The Dynamics Of Oil Market Var Models," *Journal of the European Economic Association*, 10(5), 1166–1188. 106, 113, 114

——— (2014): "The Role Of Inventories And Speculative Trading In The Global Market For Crude Oil," *Journal of Applied Econometrics*, 29(3), 454–478. 114, 115, 116

KILIAN, L., A. REBUCCI, AND N. SPATAFORA (2009): "Oil shocks and external balances," *Journal of International Economics*, 77(2), 181–194. 102

KILIAN, L., AND C. VEGA (2011): "Do Energy Prices Respond to U.S. Macroeconomic News? A Test of the Hypothesis of Predetermined Energy Prices," *The Review of Economics and Statistics*, 93(2), 660–671. 75

KILIAN, L., AND R. J. VIGFUSSON (2013): "Do Oil Prices Help Forecast U.S. Real GDP? The Role of Nonlinearities and Asymmetries," *Journal of Business & Economic Statistics*, 31(1), 78–93. 120, 121

KIM, C., AND R. NELSON (1999): *State Space Models with Regime Switching*. MIT Press, Cambridge, Massachussets. 38

KOOP, G., AND D. KOROBILIS (2013): "Large time-varying parameter VARs," *Journal of Econometrics*, 177(2), 185–198. 49, 50, 63, 65, 68, 73, 77

KOOP, G. M. (2013): "Forecasting with Medium and Large Bayesian VARS," *Journal of Applied Econometrics*, 28(2), 177–203. 51

KOOPMAN, S. J., AND J. DURBIN (2003): "Filtering and smoothing of state vector for diffuse state-space models," *Journal of Time Series Analysis*, 24(1), 85–98. 19, 40

KOOPMAN, S. J., AND A. HARVEY (2003): "Computing observation weights for signal extraction and filtering," *Journal of Economic Dynamics and Control*, 27(7), 1317–1333. 21

KOROBILIS, D. (2013): "Assessing the Transmission of Monetary Policy Using Time-varying Parameter Dynamic Factor Models," *Oxford Bulletin of Economics and Statistics*, 75(2), 157–179. 15

KUZIN, V., M. MARCELLINO, AND C. SCHUMACHER (2011): "MIDAS vs. mixed-frequency VAR: Nowcasting GDP in the euro area," *International Journal of Forecasting*, 27(2), 529–542. 15

——— (2013): "Pooling Versus Model Selection For Nowcasting Gdp With Many Predictors: Empirical Evidence For Six Industrialized Countries," *Journal of Applied Econometrics*, 28(3), 392–411. 63

LEVIN, A., J. D. LOPEZ-SALIDO, AND T. YUN (2007): "Strategic Complementarities and Optimal Monetary Policy," CEPR Discussion Papers 6423, C.E.P.R. Discussion Papers. 164

LIPINSKA, A., AND S. MILLARD (2012): "Tailwinds and Headwinds: How Does Growth in the BRICs Affect Inflation in the G-7?," *International Journal of Central Banking*, 8(1), 227–266. 105, 106, 108

LIPPI, F., AND A. NOBILI (2012): "Oil And The Macroeconomy: A Quantitative Structural Analysis," *Journal of the European Economic Association*, 10(5), 1059–1083. 109, 110

LJUNG, L. (1992): "Applications to Adaptive Algorithms,," in *Stochastic Approximations and Op timization of Random Systems, edited by L. Ljung, Georg Pflug, and Harro Walk*, pp. 95–113. Birkhauser. 65

MARCELLINO, M., AND C. SCHUMACHER (2008): "Factor-MIDAS for now- and forecasting with ragged-edge data: A model comparison for German GDP," CEPR Discussion Papers 6708, C.E.P.R. Discussion Papers. 15

MARIANO, R. S., AND Y. MURASAWA (2003): "A new coincident index of business cycles based on monthly and quarterly series," *Journal of Applied Econometrics*, 18(4), 427–443. 14

MUMTAZ, H., AND P. SURICO (2012): "Evolving International Inflation Dynamics: World And Country-Specific Factors," *Journal of the European Economic Association*, 10(4), 716–734. 49

OLIVEI, G. P., AND M. L. BARNES (2004): "Inside and Outside Bounds: Threshold Estimates of the Phillips Curve," Econometric Society 2004 Australasian Meetings 295, Econometric Society. 141

Pesaran, M. H., and A. Pick (2011): "Forecast Combination Across Estimation Windows," *Journal of Business & Economic Statistics*, 29(2), 307–318. 50, 63

Pesaran, M. H., and A. Timmermann (2007): "Selection of estimation window in the presence of breaks," *Journal of Econometrics*, 137(1), 134–161. 50

Primiceri, G. (2005): "Time Varying Vector Autoregressions and Monetary Policy," *The Review of Economic Studies*, 1(19), 465–474. 5, 49, 65

Raftery, P., M. Karny, and P. Ettler (2010): "Online prediction under model uncertainty via dynamic model averaging: Application to a cold rolling mill," *Technometrics*, 52, 52–66. 49

Riggi, M., and S. Santoro (2015): "On the slope and the persistence of the Italian Phillips curve," *International Journal of Central Banking, forthcoming.* 141

Rossi, B., A. Inoue, and L. Jin (2014): "Optimal Window Selection in the Presence of Possible Instabilities," Discussion Paper 10168, CEPR. 50

Rubio-Ramirez, J. F., D. F. Waggoner, and T. Zha (2010): "Structural Vector Autoregressions: Theory of Identification and Algorithms for Inference," *Review of Economic Studies*, 77(2), 665–696. 114

Sargent, T. J. (1999): *The Conquest of American Inflation*, Cambridge Books. Princeton University Press. 65

Sbordone, A. M. (2007): "Globalization and Inflation Dynamics: the Impact of Increased Competition," NBER Working Papers 13556, National Bureau of Economic Research, Inc. 148, 161, 163, 164

Sims, C. A. (1993): "A Nine-Variable Probabilistic Macroeconomic Forecasting Model," in *Business Cycles, Indicators and Forecasting*, NBER Chapters, pp. 179–212. National Bureau of Economic Research, Inc. 15

Stella, A., and J. H. Stock (2012): "A state-dependent model for inflation forecasting," International Finance Discussion Papers 1062, Board of Governors of the Federal Reserve System (U.S.). 141

Stock, J., and M. Watson (1989): "New Indexes of coincident and leading economic indicators," in *NBER Macroeconomics Annual*, pp. 351–393. Blanchard, O. and S. Fischer (eds). MIT Press, Cambridge, MA. 14, 141

Stock, J. H., and M. W. Watson (2008): "Phillips Curve Inflation Forecasts," NBER Working Papers 14322, National Bureau of Economic Research, Inc. 143

——— (2012): "Disentangling the Channels of the 2007-09 Recession," *Brookings Papers on Economic Activity*, 44(1 (Spring), 81–156. 51, 75

Theil, H., and A. Goldberger (1960): "On pure and mixed statistical estimation in econometrics," *International Economic Review*, 2, 65–78. 54, 59, 60

Unalmis, D., I. Unalmis, and F. Unsal (2012): "On the Sources and Consequences of Oil Price Shocks; The Role of Storage," Imf working papers, International Monetary Fund. 105, 106

Vavra, J. (2013): "Inflation Dynamics and Time-Varying Volatility: New Evidence and an Ss Interpretation," *The Quarterly Journal of Economics*, 129(1), 215–258. 150

Williams, J. C. (2010): "Sailing into headwinds: the uncertain outlook for the U.S. economy," Speech 85, Federal Reserve Bank of San Francisco. 140