Price Dispersion in the Euro Area: The Case of a Symmetric Oil Price Shock

Kai Carstensen* Oliver Hülsewig† Timo Wollmershäuser‡

February 21, 2009

Abstract

This paper studies the causes of price dispersion in the euro area emerging in response to a shock that hits all member countries symmetrically. We use a pooled VAR model which is estimated over the period 1996–2007 to generate impulse responses of a range of price and wage variables to an oil price shock. We split our sample of countries into two disjoint groups according to the impact of the oil price shock on the overall price level. While cross–country heterogeneity in the short–run pass–through can be attributed to a different oil intensity of production and a different weight of energy in the consumption basket, heterogeneity in the medium–run response of consumer prices is mainly due to a different response of wages and salaries in the industry sector, which can be attributed to different degrees of price and wage rigidities in the member countries.

JEL classifications: C32, C33, E31

Key words: Pass–through, oil price shock, euro area inflation, heterogeneity, pooled VAR model.

*LMU and Ifo Institute for Economic Research, Poschingerstr. 5, 81679 München, Germany.

†Ifo Institute for Economic Research, Poschingerstr. 5, 81679 München, Germany.

‡Corresponding author. CESifo and Ifo Institute for Economic Research, Poschingerstr. 5, 81679 München, Germany. *Email:* <Wollmershaeuser@ifo.de>

1 Introduction

The European Central Bank (ECB) is mandated to maintain price stability in the euro area as a whole. However, unless the European Monetary Union (EMU) is a homogenous economic entity, stabilization at the aggregate level does not preclude that the member countries are affected in different ways by a common exogenous disturbance, such as an oil price shock (European Central Bank, 2005). As a particular concern for monetary policy, differences in national economic structures and rigidities may cause price adjustments to deviate from each other across countries. While temporary price dispersion is not necessarily harmful to euro area economies, persistent heterogeneity in the adjustment of prices can trigger undesirable consequences such as long–lasting distortions in the development of relative prices across the member states that can lead to welfare losses for individual countries.

Potential causes for inflation dispersion in the euro area are discussed, among others, by European Central Bank (2003) and Hofmann and Remsperger (2005). First, inflation differentials in a currency union may arise from price level convergence, either because tradable goods prices converge as a result of increased trade integration or because non–tradable goods prices converge in the wake of real income convergence. Second, countries may be hit by asymmetric shocks as discussed extensively in literature on optimum currency areas (see, e.g., Mongelli, 2005, and the references therein). Finally, symmetric shocks may evoke different adjustment processes across countries.

In this paper, by analyzing the effects of oil price shocks that originate on world oil markets and hit all countries to the same extent, we concentrate on the latter question, namely, whether a symmetric shock causes asymmetric price adjustment in the euro area. There are at least two possible reasons for this scenario. Most obviously, different weights of crude oil in the production process and in the consumption basket will provoke different reactions to an oil price shock. However, if the economies are flexible, the adjustment of relative prices will take only a few quarters to restore an equilibrium.

In contrast, price dispersion may also arise from different degrees of national structural inefficiencies such as imperfect factor mobility or price and wage setting rigidities that give rise to long–lasting distortions in relative prices after exogenous shocks. The empirical evidence on the heterogeneity of price and wage adjustments across euro area countries is large: While for example Altissimo, Ehrmann, and Smets (2006) and Angeloni and Ehrmann (2007) (just to mention two out of a wide range of papers) show that inflation persistence varies markedly across countries, Arpaia and Pichelmann (2007) and Andersson et al. (2008) provide evidence on heterogenous wage formation processes in the euro area.

The existence of different degrees of national structural inefficiencies has important consequences for the single monetary policy in a currency union. Even if the ECB succeeds to restore price stability in the medium term, this does not guarantee that all economies operate at the efficient frontier (Dellas and Tavlas, 2005). For an inflation targeting central bank in a currency union Benigno (2004) showed that if the degrees of rigidities are different across countries, monetary policy should attach a higher weight to the country with a higher degree of rigidity.

This paper empirically investigates the effects of an oil price shock on consumer prices. Based on a pooled vector autoregressive (VAR) model that includes 11 euro area countries, we generate impulse responses of consumer prices to an unanticipated increase in the oil price taking account of additional macroeconomic variables – real GDP, the nominal short–term interest rate and the real effective exchange rate – that might be relevant.[1] We also extent our baseline specification by a number of auxiliary variables, which include producer prices, several sub–components of the consumer price index and real labor costs of the industry and service sector to gain an insight into the transmission of the oil price shock through the economy.

Since the time period of the data that we consider is short, ranging from 1996Q1 to 2007Q4, we use cross–sectional information to obtain better estimates of the parameters of each unit. We select this period for three reasons. First, harmonized data, which allows a detailed cross–country analysis of the transmission

[1] VAR–based analyses of the pass–through from oil price shocks to consumer prices in the euro area are scarce. We are only aware of one paper by Hahn (2003) who studies the impact of oil price shocks (and other external shocks) to euro area inflation using aggregate euro area data for the period from 1970 to 2002. She finds that external shocks explain a large fraction of the variance of prices.

of an oil price shock to consumer prices, is provided by Eurostat only from 1996 on. Second, the behavior of labor unions in response to oil price shocks today is very different from the 1970s and 1980s, which must be taken appropriately into account when estimating VAR models over the last two or three decades (Blanchard and Galí, 2007). Finally, we also expect a common reaction of the national central banks to oil price shocks in the pre–EMU period, as most of the euro area countries successfully terminated their disinflation policy in the mid–1990s and followed a fixed exchange rate policy vis–à–vis the German mark from then on.

To detect heterogeneities in the transmission of a structural shock in the context of a pooled VAR model we suggest a data–driven approach that clusters countries into disjoint groups according to the impact of the oil price shock on the overall price level. We split our sample of countries into two groups – a *high pass–through group* and a *low pass–through group* – that are endogenously identified by using a distance measure, which is determined by the absolute value of the difference between cumulated impulse responses of the overall price level. We consider the responses of consumer prices over different simulation horizons, focusing on the oil price pass–through into national prices in the short– and the medium–run after the occurrence of the shock.

Our findings indicate that the response of consumer prices is very heterogenous across member countries. Differences in the short–run pass–through can be attributed to the weight of energy in the consumption basket and in the production process. However, heterogeneity in the medium–run response of consumer prices is mainly due to different responses of wages and salaries in the industry sector, which is an indication of different degrees of price and wage rigidities in the member countries.

The remainder of the paper is organized as follows. In Section 2, the pooled VAR model for our sample of countries is presented. We generate impulse responses of a range of price and wage variables to an oil price shock. Section 3 sets out our approach of identifying disjoint groups of countries. We describe our methodology and compare the impulse responses of the groups of countries. Section 4 provides concluding remarks.

2 Estimation Approach

We employ a pooled VAR model for the euro area of the form:

$$X_t = c + \sum_{j=1}^{p} A_j X_{t-j} + \varepsilon_t, \tag{1}$$

where X_t is a matrix of endogenous variables, c is a matrix of country specific constant terms, A is a matrix of autoregressive coefficients, p is the number of lags and ε_t is a matrix of error terms. The matrix X_t consists of five columns:

$$X_t = [oil_t, \quad sti_t, \quad gdp_t, \quad reer_t, \quad hicp_t], \tag{2}$$

where oil_t denotes the oil price measured in euro, sti_t denotes the nominal short–term interest rate, which serves as the policy instrument of the central bank, gdp_t denotes real GDP, $reer_t$ denotes the real effective exchange rate, and $hicp_t$ denotes the overall price level as measured by the Harmonized Consumer Price Index (HICP). Each column is a stacked vector of country variables, consisting of $M \cdot T$ rows, where M denotes the number of countries and T is the number of observations corrected for the number of lags p.

Our sample comprises $M = 11$ euro area countries: Austria (AT), Belgium (BE), Finland (FI), France (FR), Germany (DE), Greece (GR), Ireland (IE), Italy (IT), the Netherlands (NL), Portugal (PT) and Spain (ES). The data is taken from the Eurostat database, except for the real effective exchange rate, which is from the database of the Bank for International Settlements, and the oil price, which is from the Reuters EcoWin database.[2] The *hicp* series has been seasonally adjusted using Census X12. All variables are expressed in logs except the nominal short–term interest rate that is expressed in percent. The variables are linearly de–trended. The data runs from 1996Q1 to 2007Q4 which is the maximum sample for which harmonized euro area data are available. We use a lag order of $p = 4$, which ensures that the residuals are free of first–order autocorrelation as indicated by the LM test of Baltagi (2008).

The pooled VAR model is estimated via OLS. This is appropriate because the time series dimension (48 quarters) is large relative to the cross section dimen-

[2] The oil price refers BFO crude oil, which is composed of the three North Sea grades Brent, Forties and Oseberg.

sion (11 countries). Hence, we do not have to resort to panel GMM techniques advocated for panels with large M and small T.

On the basis of the VAR model (1) we generate impulse responses of the variables to an oil price shock, which is identified by imposing a triangular orthogonalization with the oil price ordered first. This implies that the other variables are contemporaneously affected by an oil price shock but the oil price is not immediately affected by other shocks. This is justified because the oil price is determined on the world market and most likely not by the developments in a specific country, particularly not in the very short run.

2.1 Results of the Baseline Specification

Figure 1 displays the impulse responses over the first 20 quarters after an oil price shock, which corresponds to the estimated standard deviation of the residuals of the oil equation in the VAR model. The solid lines denote impulse responses, which are calculated as the median of a bootstrap procedure with 500 replications. The shaded areas are the related 68% confidence bands. All variables are expressed in percent terms, except for the nominal short–term interest rate which is expressed in basis points at an annual rate (100 basis points equal one percent).

When a ten percent oil price shock hits the economy, there is an immediate increase in the general price level of roughly 0.1 percent. This is not a surprise because the oil price has a direct and instantaneous impact on the prices for heating oil and gasoline. At the same time, the short–term interest rate is shifted upwards, presumably in an attempt of the central bank to counter the inflationary impulse, and the real exchange rate devaluates. In the subsequent quarters, the price effect increases further to a maximum of 0.15 percent after four quarters before it slowly dies out. This reflects the well-known sluggishness of a wide array of consumer prices. For example, the prices for alternative energy sources like natural gas and for services like transport typically react with a delay. In addition, it may take some time, until producers pass on their cost increases to consumers. Finally, second–round effects operating through wage negotiations are implemented with some delay. Mirroring the price reaction, the short–term interest rate hike continues for some quarters. Subsequently, it is quickly taken

Figure 1: Impulse Responses to an Oil Price Shock — Baseline Specification

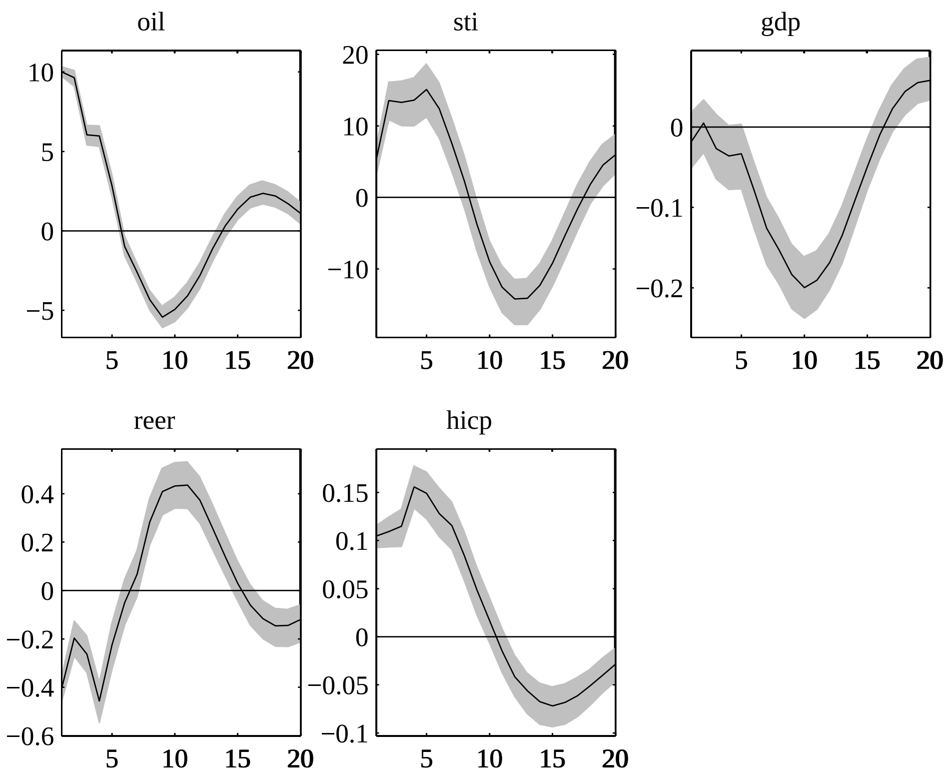

Notes: Orthogonalized impulse responses to an oil price shock. The solid lines denote impulse responses, which are calculated as median of a bootstrap procedure with 500 replications. The shaded areas are the related 68% confidence intervals. All variables are expressed in percent terms, except for the nominal short–term interest rate which is expressed in basis points at an annual rate. The horizontal axis is in quarters.

back and reaches a trough after 12 quarters. This probably reflects the attempt of the central bank to stabilize the economy as real output starts to decline after two, and reaches a trough after ten, quarters.

2.2 Results of Extended Specifications

To get more insights into the transmission process after an oil price shock, we extend the VAR model with additional variables. We do this by adding only one variable at a time. This prevents an overfitting and guarantees that we do not run out of degrees of freedom. Specifically, we estimate the extended specification (with the additional variable ordered last)

$$X_t = \left[oil_t, \;\; sti_t, \;\; gdp_t, \;\; reer_t, \;\; hicp_t, \; z_t\right], \tag{3}$$

where z_t denotes one of the following variables: the HICP component only including energy items ($hicp_{nrg}$), the HICP component only including unprocessed food ($hicp_{foodu}$), the HICP component only including processed food ($hicp_{foodp}$), the HICP component only including non–energy industrial goods ($hicp_{igood}$), the HICP component only including services ($hicp_{serv}$),[3] the domestic producer price index of total industry (excluding construction) ($dppi$), the wages and salaries component of the nominal labor cost index of sections C to K of the NACE Rev. 1 nomenclature (industry excluding construction and services excluding public administration) deflated by the HICP ($rlci$), the wages and salaries component of the nominal labor cost index of total industry (excluding construction, sections C to E of the NACE Rev. 1 nomenclature) deflated by the HICP ($rlci_{ind}$), and the wages and salaries component of the nominal labor cost index of services (excluding public administration, sections F to K of the NACE Rev. 1 nomenclature) deflated by the HICP ($rlci_{serv}$). All series are taken from the Eurostat database. The *hicp* and *dppi* series have been seasonally adjusted using Census X12.

The reasons for choosing these variables as additional endogenous variables in the VAR model are twofold. First, we would like to analyze the direct transmission of oil price shocks to consumer prices by focussing on the price response of energy goods, which are part of the consumption basket and which account

[3]These are the five principal components of the HICP that are published by Eurostat.

for on average 9 percent of households' final monetary consumption expenditure. Second, with the remaining variables we would like to shed some light on the indirect transmission of oil price shocks to other components of the consumer price index. On the one hand, higher costs for energy inputs in production are likely to have an impact producer prices, which predict upstream pressures on the prices for non–energy goods and services. On the other hand, the response of wage costs to oil price increases is an indicator for understanding the importance of second–round effects on prices.

The response of each of the additional variables to an oil price shock is shown in Figure 2.[4] Not surprisingly, the energy component of the HICP increases instantaneously by about 1.5 percent and takes a development very similar to that of the oil price. The price of processed food also rises on impact but reaches the maximum reaction only after five quarters and remains elevated for more than ten quarters. This indicates that it is difficult for producers to pass on increased costs to consumers. Interestingly, the reaction of unprocessed food prices is quite different. Being unaffected on impact, they slightly decrease for two quarters before they rise considerably, reaching a maximum of 0.5 percent after 7 quarters. This behavior might reflect that in the short–run the reduction in real income dominates, leading to lower consumption demand. In the medium term, however, energy costs, which are an important determinant of food prices, necessitate a price hike. In contrast, the prices for industrial goods and for services show almost no reaction in the first three quarters. Subsequently, they increase significantly and peak after 8 to 9 quarters at around 0.07 and 0.09 percent, respectively. Hence, the reaction pattern of the HICP components differ markedly. In particular, the strong initial response of the total HICP index is mainly driven by the prices for energy and for processed food.

Additional insights can be obtained from an analysis of producer prices and labor costs. Producer prices go up instantaneously after an oil price shock. However, they transmit with a considerable lag into the prices for industrial goods and services. This indicates that markups are temporally depressed before firms are able to raise retail prices. Total real wage costs initially drop by 0.1 percent

[4]As the responses of the five baseline variables are almost unaffected by the extension of the VAR model, they are not shown again.

before they start rising gradually. Hence, nominal wage costs seem to react with some delay. However, the patterns are different between industry and services. The negative effect on real wage costs in industry is pronounced and long–lasting, while real wage costs in services drop only on impact. They even rise above trend after about 10 quarters. Overall, real wages do not seem to be extremely sticky. In particular, they show a significant and, compared to the HICP response, quantitatively important instantaneous reaction to an oil price shock.

Figure 2: Impulse Responses to an Oil Price Shock — Extended Specification

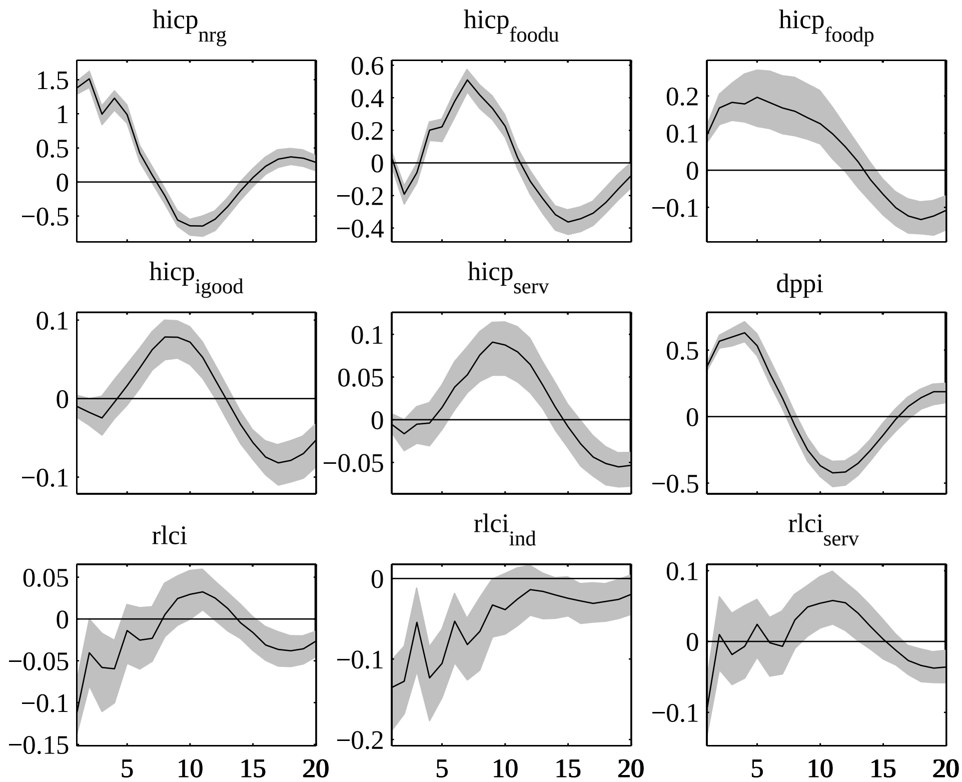

Notes: Orthogonalized impulse responses to an oil price shock. The solid lines denote impulse responses, which are calculated as median of a bootstrap procedure with 500 replications. The shaded areas are the related 68% confidence intervals. All variables are expressed in percent terms. The horizontal axis is in quarters.

3 Heterogeneity across Countries

So far, we have estimated the pooled VAR under the assumption that systematic country differences can be explained by different intercepts. However, there is considerable evidence that the euro area countries are more heterogenous. The two most obvious reasons for heterogenous country–specific reactions to an oil price shock are, first, the different weights of oil in both production and household consumption and, second, different levels of wage rigidity. In order to take this into account, we cluster the countries into two different groups according to the response of the overall price level to an expansionary oil price shock. Hence, the countries are divided into a *high pass–through group* and a *low pass–through group*. Since our approach is novel, we describe the methodology more explicitly.

3.1 Methodology

In principle, one can think of the reaction of overall price level to an oil price shock as a general function of the country–specific characteristics. This implies that the VAR parameters depend on these characteristics and, hence, the impulse responses differ from country to country. Therefore, countrywise estimation would be optimal. Unfortunately, the precise estimation of impulse response coefficients within the VAR framework requires a relatively large number of observations. Since for the reasons outlined above a our sample does not start before 1996, we need to construct country panels in order to increase the number of observations by using the cross–section dimension. To facilitate an easy distinction between such country panels, we consider only two of them, namely a *high pass–through group* and a *low pass–through group*. Hence, the question we have to answer in this section is how to allocate the countries in our sample to one of these two groups. This is achieved in three steps.

1. Step: Define and Estimate the Distance between Sup–panels To quantify the difference between any two sub–panels of countries, we need to define a distance measure. As we are interested in the different impulse responses of the overall price level after an oil price shock, we use

$$d = \left| \sum_{i=q_1}^{q_2} \alpha_{1i} - \sum_{i=q_1}^{q_2} \alpha_{2i} \right| , \tag{4}$$

where α_{1i} and α_{2i} are the median responses of the overall price level of the first and second sub–panel, respectively, i periods after the occurrence of the shock. We consider the response lags $q_1 = 1$ to $q_2 = 2$ (*short–run pass–through*) and $q_1 = 3$ to $q_2 = 12$ (*medium–run pass–through*). Hence, the distance measure in expression (4) reflects the absolute value of the difference between the cumulated impulse responses.

At first sight, it is now straightforward to allocate each country to either the *high pass–through group* or the *low pass–through group*. One can simply estimate all possible pairs of sub–panels and choose the pair with the largest distance. This approach resembles a cluster algorithm, where the number of clusters is fixed and the distance between the cluster centers (i.e., the impulse response coefficients) is maximized. However, we have to bear in mind that the impulse response coefficients are not observed but estimated. Hence, choosing the maximum distance pair only would contaminate the choice by a considerable portion of randomness. In fact, we find that there a many different pairs of sub–panels that exhibit similar distance measures.

Therefore, we proceed as follows. We estimate pooled VAR models for all possible pairs of sub–panels, which contain at least three countries to ensure enough degrees of freedom for each sub–panel.[5] Overall the number of pairs of sub–panels amounts to 957.[6] For all pairs of sub–panels we generate impulse responses to a oil price shock and calculate the distance measure.

2. Step: Select Pairs of Sub–panels with Significant Distance Measure Then, we identify all pairs of sub–panels that exhibit a significant distance

[5] As before the VAR models are estimated for the same set of variables as in the baseline specification (oil price, nominal short–term interest rate, real GDP, real effective exchange rate and overall price level) using a lag length of $p = 4$. The oil price shock is identified by imposing a triangular orthogonalization.

[6] Notice that in our panel the total number of disjoint pairs of sub–panels amounts to 1024 ($= 2^{11}/2$). Given that we consider only pairs of sub–panels containing at least three countries, this reduces the number of pairs to 957, since there are one combination without any country, 11 combinations with only one country and 78 combinations – $(10 \times 11)/2$ – with two countries.

measure, where significance is detected as follows. Assume that the estimated impulse response coefficients $\widehat{\alpha}_{1i}$ and $\widehat{\alpha}_{2i}$ asymptotically follow a normal distribution. Then the sums of the coefficients considered for the distance measure, denoted by $\widehat{s}_1 = \sum_{i=q_1}^{q_2} \widehat{\alpha}_{1i}$ and $\widehat{s}_2 = \sum_{i=q_1}^{q_2} \widehat{\alpha}_{2i}$, are also asymptotically normal. Under the null hypothesis that all pairs of sub–panels are identical and have the same sum of population coefficients $s = \sum_{i=q_1}^{q_2} \alpha_i$, the only systematic difference in the estimation results is the size of the panel from which they are estimated.

The sums of the estimated coefficients should be approximately distributed as:

$$\widehat{s}_1 - s \sim N\left(0, \sigma^2/(N_1 T)\right) \tag{5}$$

$$\widehat{s}_2 - s \sim N\left(0, \sigma^2/(N_2 T)\right), \tag{6}$$

where N_1 is the size of the first sub–panel, N_2 is the size of the second sub–panel, T is the number of observations corrected for the number of lags p in the VAR model and σ^2 is the population variance that is assumed to be constant across countries. Furthermore, assuming that the countries are independent, we can apply a classical two–sided difference test using the statistic: $d = \widehat{s}_1 - \widehat{s}_2$. Under the null hypothesis, the statistic is approximately normally distributed with mean zero and variance:

$$\mathrm{Var}(d) = \sigma^2/(N_1 T) + \sigma^2/(N_2 T) = (1/N_1 + 1/N_2)\, \sigma^2/T. \tag{7}$$

Since σ^2 is unknown, we estimate the population variance from expression (7) by noting that:

$$\sigma^2 = T\mathrm{Var}(d)/\left(1/N_1 + 1/N_2\right), \tag{8}$$

where the sample variance of the distance measure $\mathrm{Var}(d)$ is calculated from the numerous realizations of d. Given the estimate of σ^2, we construct a t–statistic and compare it with the corresponding 95% critical value of the t–distribution.

As a result, we have identified all those pairs of sub–panels that are significantly different from each other. If there was no significant difference at all, we would conclude that all countries show the same response of the overall price level to an oil price shock and terminate the analysis here. However, we find 353 (266) significant distance measures for the response lags $q_1 = 1$ to $q_2 = 2$ ($q_1 = 3$ to

$q_2 = 12$). In contrast to using only the maximum–distance pair, we thus consider all the different ways to split the panel of countries into significantly different sub–panels. Thereby, we alleviate the problem that the impulse response coefficients, and hence the distance measure, are subject to estimation uncertainty. However, this approach in turn raises the question how to allocate a single country to either the *high pass–through group* or the *low pass–through group*.

3. Step: Allocate each Country to either the High or the Low Pass–through Group The allocation problem is tackled in the final step. Using the pairs of sub–panels with a significant distance measure we calculate the frequency that a specific country belongs to the sub–panels with the stronger reaction of the overall price level to an oil price shock. If this frequency is above a threshold that is determined below, then the respective country is allocated to the *high pass–through group*, otherwise it is allocated to the *low pass–through group*.

The idea behind this rule is as follows. Assume there are three "true" high pass–through countries. Then we should expect that the distance measure is maximized when these three countries are put into one sub–sample and all the others in the other sub–sample. However, due to sampling error, a different pair of sub–samples may actually exhibit the largest distance. Using our approach, we may at least expect to find each of the three high pass–through countries to be more often in the high pass–through sub–sample than any of the other countries.

To accomplish this, we now derive the threshold for the frequency that a specific country belongs to the high pass–through sub–panels. From the previous step we know which pairs of sub–samples are significantly different from each other. Now we count how many times each country is in a high pass–through sub–sample. A priorily, each country has the same chance to be a high pass–through country. Hence, under this null hypothesis there is, for each pair of sub–panels, a 50 percent chance that a specific country is in the high pass–through sub–panel. Now assume that there are a total of N_c different pairs of sub–panels of which n exhibit a significant distance measure. Then, for each country, the number of times it is in the high pass–through sub–panel resembles a random experiment, where n draws without replacement are taken from a population of size N_c that is composed of 50 percent white (=high pass–through) and 50 percent black (=low pass–through) elements. Accordingly, the frequency x – that a particular country is found to be in the *high pass–through group* – follows a hypergeometric distribution: $f(x; N_c, N_c/2, n)$, where the number of pairs N_c depends on the total number of countries M and the minimum size of a sub–panel.[7]

Finally, from the hypergeometric distribution we derive a 95% critical value for the frequency that a particular country belongs to the *high pass–through group*. If any country is selected more often, it is unlikely that this is due to pure chance. Hence, we allocate these countries to the *high pass–through group*. All other countries are allocated to the *low pass–through group*.

3.2 Identified Country Groups

Table 1 summarizes the country clusters for the two time horizons. While some countries either belong to the group with a high pass–through (Austria, Finland, France and Germany) or a low pass–through (Greece and Portugal) in both the short run and the medium run, the other countries show either a high pass–through in the short run and a high pass–through in the medium run (Belgium, Ireland and Spain) or vice versa (Italy and the Netherlands). Figure 3 plots the relative frequencies of belonging to the high pass–through group for the two horizons together with the average critical value of the hypergeometrical distributions.

Table 1: Country Clusters

Horizon	High pass–through group	Low pass–through group
1-2 quarters	AT, BE, FI, FR, DE, IE, ES	GR, IT, NL, PT
3-12 quarters	AT, FI, FR, DE, IT, NL	BE, GR, IE, PT, ES

[7]Let us denote the the minimum size of a sub–panel by m. Then the number of possible pairs of sub–panels can be calculated as $N_c = \sum_{i=m}^{M-m} \begin{pmatrix} M \\ i \end{pmatrix}$. In our case, with $M = 11$ countries and a minimum sub–panel size of $m = 3$, we have $N_c = 2048$ pairs. Of these pairs, we have to estimate only $2048/2 = 1024$ because, e.g., the ordering of the pair $A = \{1, 2, 3, 4, 5, 6\}, B = \{7, 8, 9, 10, 11\}$ or $A = \{7, 8, 9, 10, 11\}, B = \{1, 2, 3, 4, 5, 6\}$ is irrelevant, while *ex ante* either A or B could be the strong reaction sub–panel.

Figure 3: Relative frequency of belonging to the high pass–through group for different horizons following the shock

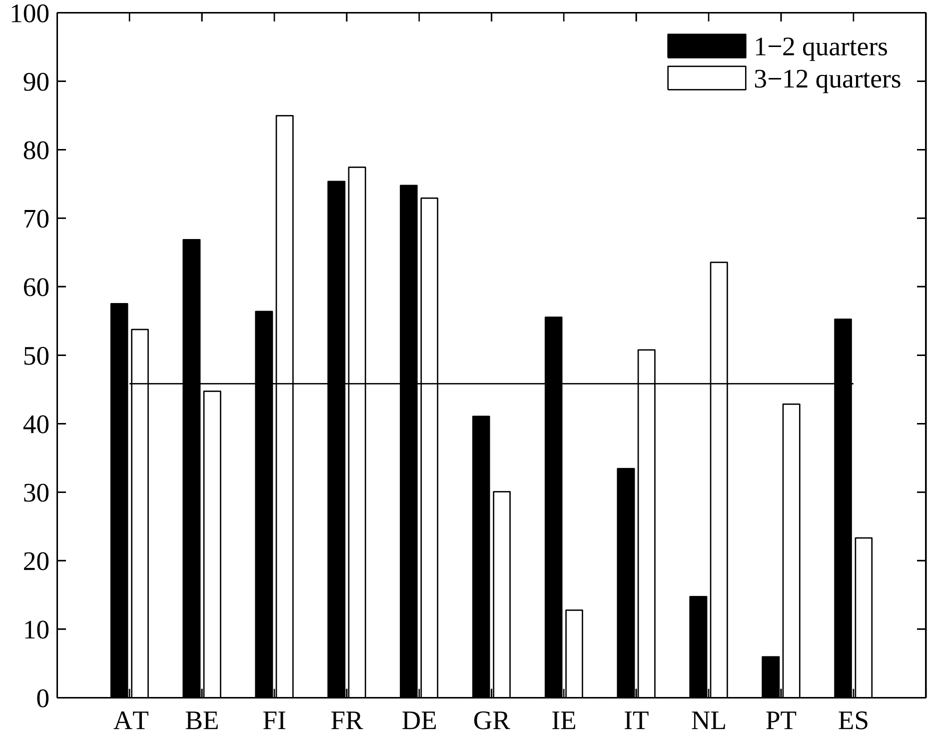

Notes: For each country the bars show the relative frequency in percent of belonging to the high pass–through group in the short–run and in the medium–run. The total number of combinations of countries n that show a significantly higher distance is 353 and 266, respectively. The horizontal line denotes the average critical value for x (over the two horizons) that a country is in the high pass–through group.

3.3 Impulse Responses when Countries are Clustered according to Short–run Pass–through

Now we re-estimate the pooled VAR model for the pair of country groups that is clustered according to the short–run pass–through of the oil price shock to the overall price level. The impulse responses are shown in Figure 4. To facilitate a comparison with the initial assumption that all countries are equal, the confidence regions for the impulse responses estimated from the baseline VAR model are reported as shaded areas.

First note that the oil price development is almost indistinguishable between the groups. This confirms that the groups are hit by an identical shock and differences in the responses of other variables are due to the structural heterogeneity between them. The most notable differences arise at the HICP responses. In the low pass–through group the HICP instantaneously rises by 0.05%, whereas in the high pass–through group the increase amounts to 0.15%. The impact reaction of both groups is outside the confidence region estimated from the baseline VAR model which confirms that there is heterogeneity. The response difference between the groups amounts to 0.10 percentage points and lasts for four quarters. In the following quarters, the reaction of the high pass–through group dies out more quickly than that of the low pass–through group. In contrast to the HICP reaction, the real exchange rate response is more pronounced in the low pass–through group. The responses of the short–term interest rate and of GDP are very similar and statistically indistinguishable. All this indicates that in the economies of the high pass–through group prices rather than real exchange rates carry the burden of adjustment while in the economies of the low pass–through group it is vice versa.

To analyze the reasons behind the differences, we report the impulse responses of the 9 variables estimated with the extended specification in Figure 5. To facilitate a comparison between the country groups, we also report the differences between the impulse responses in Figure 6. The initial difference in the response of the HICP is largely driven by the energy component of the HICP. While energy prices rise by 1.6 percent on impact in the high pass–through group, they only rise by 1.1 percent in the low pass–through group. This can be explained by the

Figure 4: Impulse Responses to an Oil Price Shock — Countries Clustered According to Short–run Pass–through

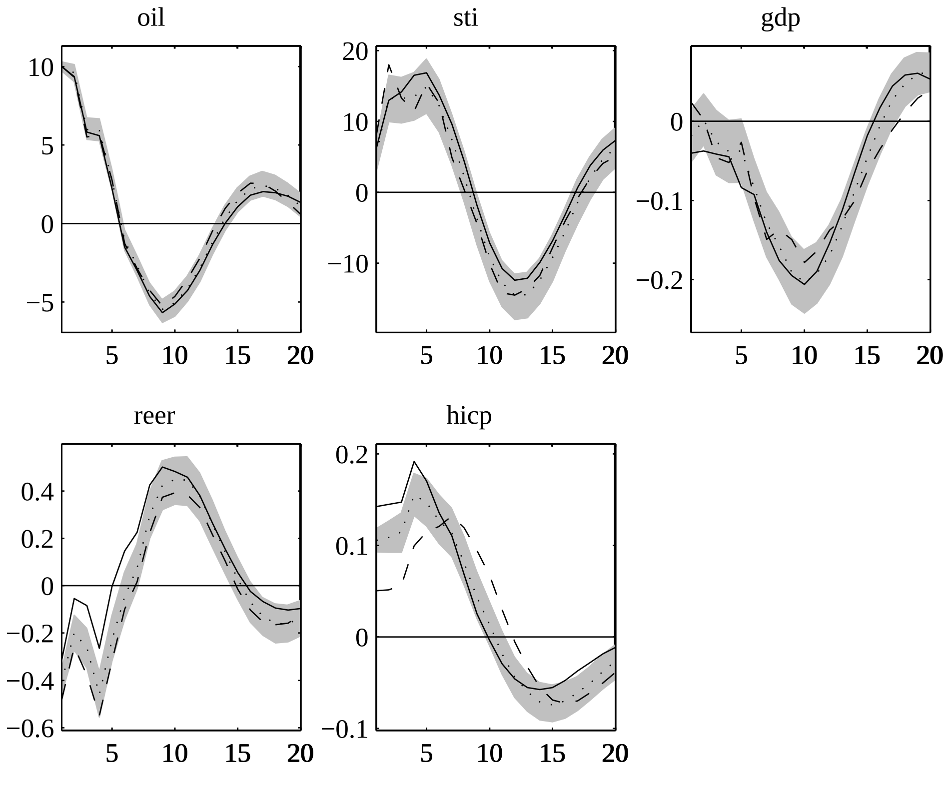

Notes: Orthogonalized impulse responses to an oil price shock. The lines denote impulse responses, which are calculated as median of a bootstrap procedure with 500 replications. The dotted lines refer to the estimation of the complete panel including 11 countries of the euro area, the solid lines to the high pass–through group (AT, BE, FI, FR, DE, IE, ES), and the dashed lines to the low pass–through group (GR, IT, NL, PT). The shaded areas are the 68% confidence intervals resulting from the estimation of the complete panel. All variables are expressed in percent terms, except for the nominal short–term interest rate which is expressed in basis points at an annual rate. The horizontal axis is in quarters.

Figure 5: Impulse Responses to an Oil Price Shock — Countries Clustered According to Short–run Pass–through — Extended Specifications

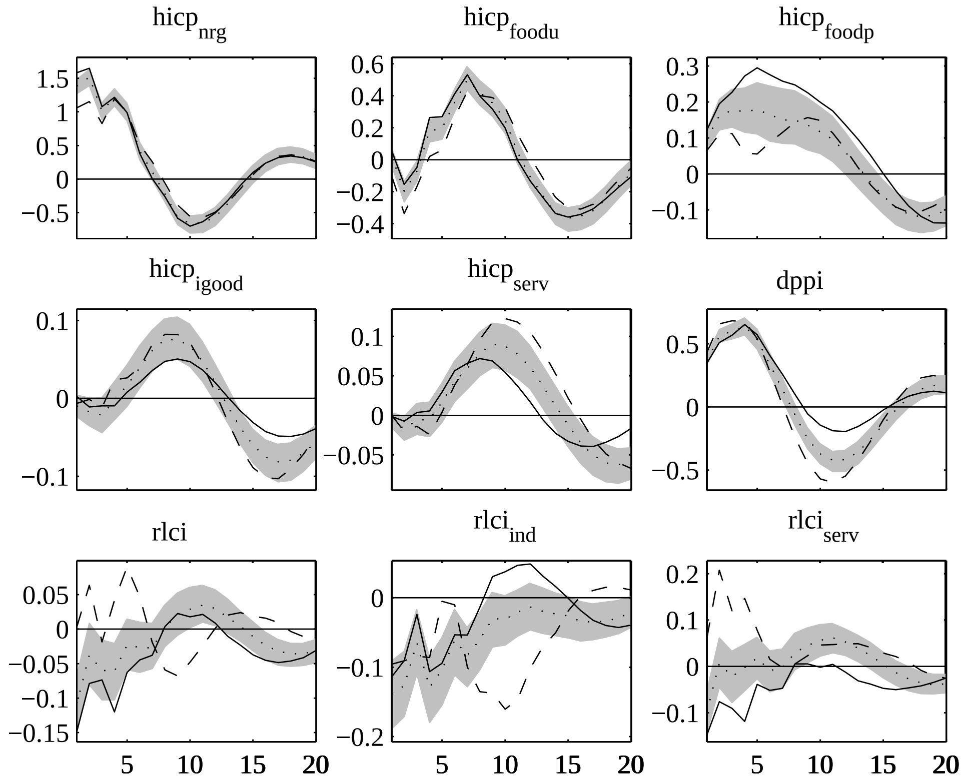

Notes: Orthogonalized impulse responses to an oil price shock. The lines denote impulse responses, which are calculated as median of a bootstrap procedure with 500 replications. The dotted lines refer to the estimation of the complete panel including 11 countries of the euro area, the solid lines to the high pass–through group (AT, BE, FI, FR, DE, IE, ES), and the dashed lines to the low pass–through group (GR, IT, NL, PT). The shaded areas are the 68% confidence intervals resulting from the estimation of the complete panel. All variables are expressed in percent terms. The horizontal axis is in quarters.

Figure 6: Impulse Response Differential between High Pass–through and Low Pass–through Group — Countries Clustered According to Short–run Pass–through

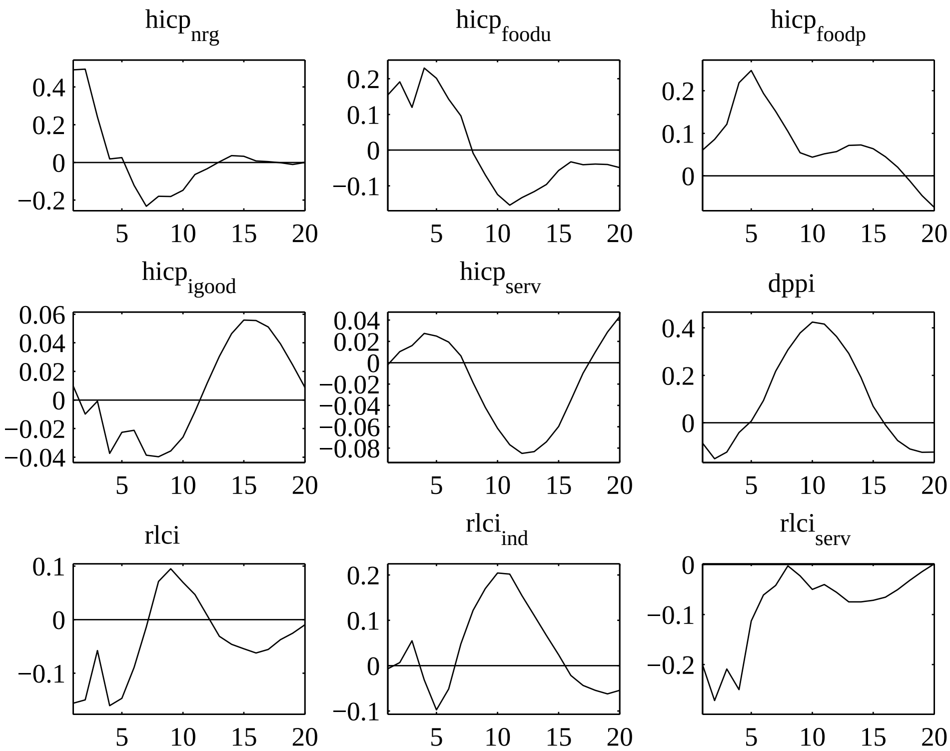

Notes: The difference between the impulse responses is expressed in percentage points. The horizontal axis is in quarters.

different weights of energy in the HICP basket. In the high pass–through group the energy component accounts for on average 9.2 percent as opposed to 8 percent in the low pass–through group (see Figure 10 in the Appendix). Different tax rates on mineral oil may also play a role. Among the other HICP components, the price reactions of processed and unprocessed food also contribute to a stronger overall price increase of the high pass–through group. In contrast, the prices for industrial goods and services, which do not react immediately, show a stronger medium–term response in the low pass–through group. This can partly be explained by the fact that oil consumption of the industry sector (as defined by the International Energy Agency) relative to its gross value added amounts to 7.9 percent in the low pass–through group but only 5.8 percent in the high pass–through group (averages over the years 2004-2006, see Figure 11 in the Appendix). As a final piece of evidence, real labor costs are strongly depressed in the high pass-through group while they remain largely unaffected on impact, and even rise thereafter, in the low pass–through group. This seems to indicate that the economies with a high degree of short–term real wage persistence show the more curbed overall price response to an oil price shock.

3.4 Impulse Responses when Countries are Clustered according to Medium–run Pass–through

We the estimate the VAR models for the two groups identified as the medium–run high pass–through countries (AT, FI, FR, DE, IT, NL) and low pass–through countries (BE, GR, IE, PT, ES). Figure 7 shows that the oil price development after the shock is again almost indistinguishable between groups. However, the HICP response is markedly different in size but not so much in the overall shape. The maximum effect in the high pass–through group is 0.18 percent as opposed to 0.13 percent in the low pass-through group. Moreover, the responses lie outside the confidence region estimated from the baseline VAR and thereby confirms that there is significant heterogeneity. The difference in price responses is accompanied by differences in short–term GDP responses and medium–term real exchange rate responses while the development of the interest rate is very similar across groups. In particular, the stronger slump in GDP in the low pass-through group indicates that the adjustment in this group works relatively more through the real side of the economy.

Again, a more detailed analysis is possible with the extended specification, see Figures 8 and 9. They show that the heterogenous medium–run response of the HICP is mainly due to the evolution of prices for services and—to a lesser extent—for non–energy industrial goods. While the reaction of service prices reaches its peak after nine to ten quarters in both groups, the maximum effect in the high pass–through group is 0.12 percent compared to only 0.04 percent in the low pass–through group. The prices for non–energy industrial goods reach their peak one or two quarters earlier. The differential between the two groups is on average 0.04 percentage points in the second and third year following the shock.

Looking at the determinants of the consumer prices, it also turns out that the producer prices react slightly stronger in the low pass–through group on impact. Again, this might be explained by the higher dependence on oil in relation to gross value added in industrial production (7.5 percent compared to 5.7 percent in the high pass–through group). However, from the fourth quarter on the increase in domestic producer prices turns out to be more pronounced in the high pass–through group. While in both groups producer prices peak after four quarters, producer prices fall much faster in the low pass–through group, resulting in a maximum differential of more than 0.15 percentage points after eight quarters.

This medium–run heterogeneity is mainly caused by large differentials in the evolution of real labor costs in the industry sector. While in the high pass–through group real labor costs remain more or less unchanged in the first year following the shock, implying that nominal wages and the price level increased proportionally, real industry labor costs in the low pass–through group dropped immediately by -0.2 percent and remained at this low level for around one year, before gradually returning to baseline. Thus, the countries with a higher degree of real industry wage rigidity exhibit a larger medium–run reaction of producer prices and non–energy industrial goods prices.

The response of wages and salaries in the service sector, i.e. the sector which is commonly characterized as being excluded from international competition, is different from their response in the industry sector. While in both country groups real labor costs drop by approximately 0.1 percent on impact, they subsequently

Figure 7: Impulse Responses to an Oil Price Shock — Countries Clustered According to Medium–run Pass–through

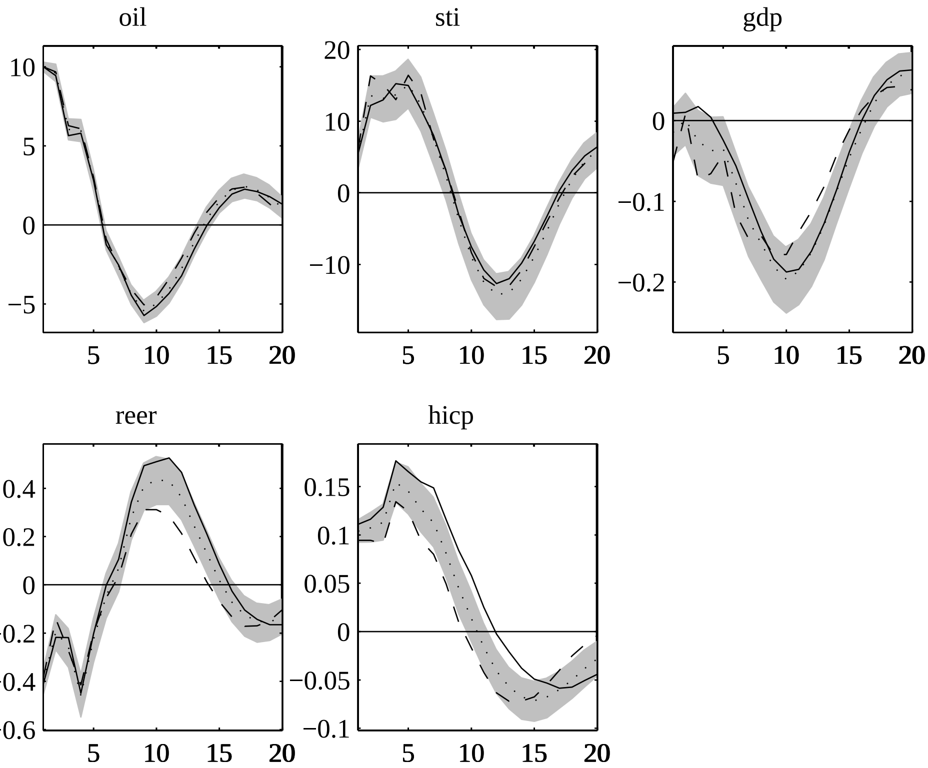

Notes: Orthogonalized impulse responses to an oil price shock. The lines denote impulse responses, which are calculated as median of a bootstrap procedure with 500 replications. The dotted lines refer to the estimation of the complete panel including 11 countries of the euro area, the solid lines to the high pass–through group (AT, FI, FR, DE, IT, NL), and the dashed lines to the low pass–through group (BE, GR, IE, PT, ES). The shaded areas are the 68% confidence intervals resulting from the estimation of the complete panel. All variables are expressed in percent terms, except for the nominal short–term interest rate which is expressed in basis points at an annual rate. The horizontal axis is in quarters.

Figure 8: Impulse Responses to an Oil Price Shock — Countries Clustered According to Medium–run Pass–through — Extended Specification

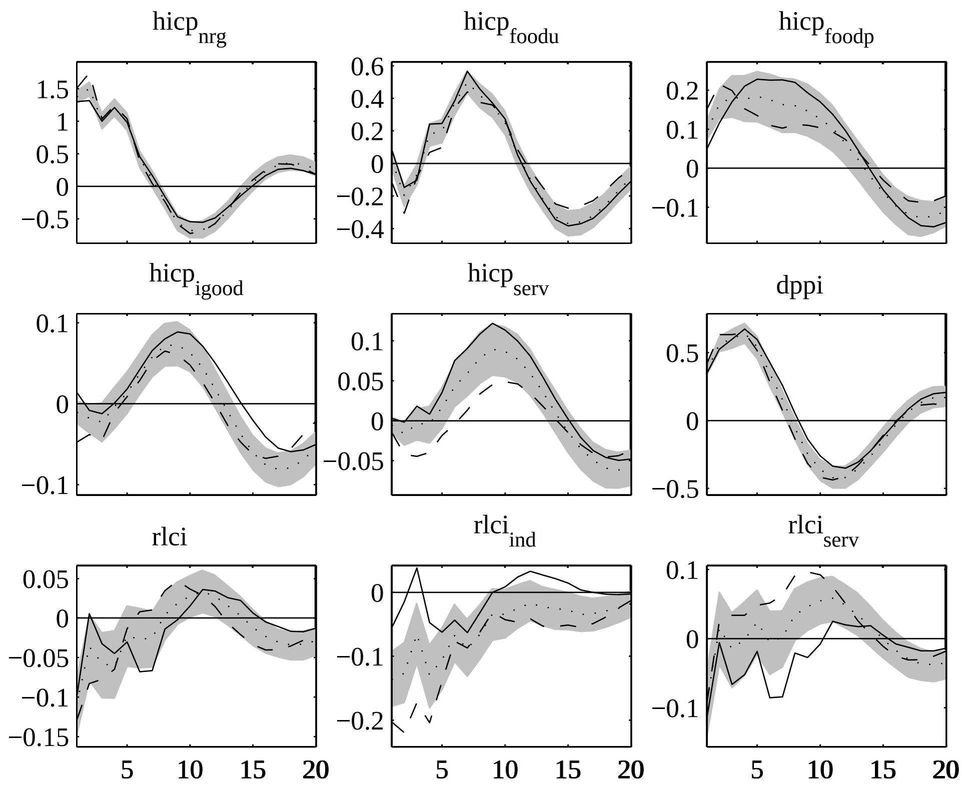

Notes: Orthogonalized impulse responses to an oil price shock. The lines denote impulse responses, which are calculated as median of a bootstrap procedure with 500 replications. The dotted lines refer to the estimation of the complete panel including 11 countries of the euro area, the solid lines to the high pass–through group (AT, FI, FR, DE, IT, NL), and the dashed lines to the low pass–through group (BE, GR, IE, PT, ES). The shaded areas are the 68% confidence intervals resulting from the estimation of the complete panel. All variables are expressed in percent terms. The horizontal axis is in quarters.

Figure 9: Impulse Response Differential between High Pass–through and Low Pass–through Group — Countries Clustered According to Medium–run Pass–through

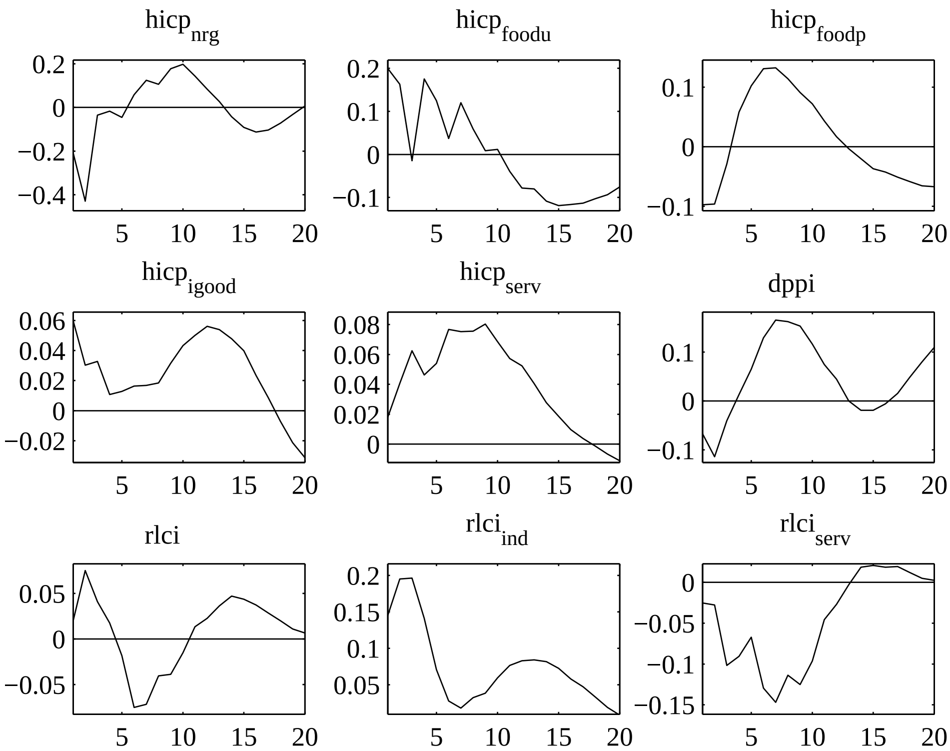

Notes: The difference between the impulse responses is expressed in percentage points. The horizontal axis is in quarters.

rise above baseline in the low pass–through group, whereas they continue to be negative for around two years in the high pass–through group. However, this evolution seems to have no impact on the prices of services in the consumption basket, where the pass–through of the oil price shock is significantly higher in the high pass–through group. An explanation for this result is that the final price of services is not only determined by wage costs in the service sector, but also by the costs of other intermediate inputs (energy, industrial goods), which are affected more by the increase in oil prices in the high pass–through group.

4 Concluding Remarks

This paper studies the causes of price dispersion in the euro area emerging in response to a shock that hits all member countries symmetrically. We use a pooled VAR model which is estimated over the period 1996–2007 to generate impulse responses of a range of price and wage variables to an oil price shock. To detect heterogeneities in the transmission of a structural shock in the context of a pooled VAR model we suggest a data–driven approach that clusters countries into disjoint groups according to the impact of the oil price shock on the overall price level. We split our sample of countries into two groups – a *high pass–through group* and a *low pass–through group* – that are endogenously identified by using a distance measure, which is determined by the absolute value of the difference between cumulated impulse responses of the overall price level. We consider the responses of consumer prices over different simulation horizons, focusing on the oil price pass–through into national prices in the short– and the medium–run after the occurrence of the shock.

Our findings indicate that the response of consumer prices is very heterogenous across member countries. Differences in the short–run pass–through can be attributed to the weight of energy in the consumption basket and in the production process. This type of price dispersion should, however, not be a cause for concern for the ECB per se. As long as the the economies are flexible enough, the adjustment of relative prices will only take a few quarters before restoring a new equilibrium.

The policy–relevant finding of this paper is that heterogeneity in the medium–

run response of consumer prices is mainly due to different responses of wages and salaries in the industry sector, which is an indication of different degrees of price and wage rigidities in the member countries and which could give rise to long–lasting distortions in relative prices after exogenous shocks. The existence of different degrees of national structural inefficiencies has important consequences for the single monetary policy in a currency union. Even if the ECB succeeds to restore price stability in the medium term, this does not guarantee that all economies operate at the efficient frontier (Dellas and Tavlas, 2005). For an inflation targeting central bank in a currency union Benigno (2004) showed that if the degrees of rigidities are different across countries, monetary policy should attach a higher weight to the country with a higher degree of rigidity.

References

ALTISSIMO, F., M. EHRMANN, AND F. SMETS (2006): "Inflation Persistence and Price–setting Behaviour in the Euro Area – A Summary of the IPN Evidence," Occasional Paper Series 46, European Central Bank.

ANDERSSON, M., A. GIESECK, B. PIERLUIGI, AND N. VIDALIS (2008): "Wage Growth Dispersion across the Euro Area Countries – Some Stylised Facts," Occasional Paper Series 90, European Central Bank.

ANGELONI, I., AND M. EHRMANN (2007): "Euro Area Inflation Differentials," *Topics in Macroeconomics*, 7(1), 1509–1509.

ARPAIA, A., AND K. PICHELMANN (2007): "Nominal and Real Wage Flexibility in EMU," *International Economics and Economic Policy*, 4(3), 299–328.

BENIGNO, P. (2004): "Optimal Monetary Policy in a Currency Area," *Journal of International Economics*, 63(2), 293–320.

BLANCHARD, O., AND J. GALÍ (2007): "The Macroeconomic Effects of Oil Shocks: Why are the 2000s so different from the 1970s?," NBER Working Paper 13368, National Bureau of Economic Research.

DELLAS, H., AND G. TAVLAS (2005): "Wage Rigidity and Monetary Union," *Economic Journal*, 115(506), 907–927.

EUROPEAN CENTRAL BANK (2003): "Inflation Differentials in the Euro Area: Potential Causes and Policy Implicatios," Occasional Paper September, European Central Bank.

——— (2005): "Monetary Policy and Inflation Differentials in a Heterogenous Currency Area," *Monthly Bulletin*, May, 61–77.

HAHN, E. (2003): "Pass–through of External Shocks to Euro Area Inflation," Working Paper Series 243, European Central Bank.

HOFMANN, B., AND H. REMSPERGER (2005): "Inflation Differentials among Euro Area Countries: Potential Causes and Consequences," *Journal of Asian Economics*, 16(3), 403–419.

Mongelli, F. (2005): "'New' Views on the Optimum Currency Area Theory: What is EMU Telling Us?," *Journal of Common Market Studies*, 43, 607–635.

5 Appendix

Figure 10: Weight of the Energy Component in the HICP

Source: Eurostat.

Figure 11: Industry Oil Consumption Relative to Gross Value Added

Notes: Oil consumption is measured in thousands of metric tons. Gross value added refers to gross value added in the industry sector including construction and is measured in current prices.
Source: International Energy Agency, Eurostat.