Aggregate real exchange rate persistence through the lens of sectoral data[*]

By Laura Mayoral[a] and Maria Dolores Gadea[b]

[a]Institute for Economic Analysis, (CSIC); [b]Dept. of Applied Economics, U. of Zaragoza

September 2010

Abstract

In this paper, a novel approach to analyze real exchange rate (RER) persistence and its sources is presented. The tight relation between aggregate and sectoral persistence recently investigated by Mayoral (2009) allows us to break down aggregate RER persistence into the persistence of its different subcomponents. Using highly disaggregated data for a group of 15 EU countries, it is shown that the distribution of sectoral persistence is highly heterogeneous and very skewed to the right so that a limited number of sectors are responsible for the high levels of persistence observed at the aggregate level. Accordingly, quantile regression has been employed to investigate whether traditional theories, such as the lack of arbitrage due to nontradability or imperfect competition combined with price stickiness, are able to account for the slow reversion to parity of RERs.

KEYWORDS: Real exchange rates, persistence, PPP puzzle, heterogeneous dynamics, aggregation bias, nontradability, imperfect competition, pricing-to-market.

JEL CLASIFFICATION: E31; F31; D40

1. INTRODUCTION

Most of the empirical literature on purchasing power parity (PPP) and real exchange rate (RER) persistence focuses on the analysis of aggregate RER data, where the latter are constructed with

1

aggregate price indices. After years of close scrutiny, the general consensus is that aggregate RERs may converge to parity in the long run, but that the rate at which this happens is very slow, with half-lives (HL) in the range of 3 to 5 years (Frankel and Rose, 1996, Lothian and Taylor, 1996 and Murray and Papell, 2005). Thus, while the high volatility of real exchange rates could potentially be explained by monetary or financial shocks, the rate of reversion to parity seems to be too slow to be compatible with plausible nominal rigidities, giving rise to the so-called *PPP puzzle* (Rogoff, 1996).

Several avenues have been pursued to shed more light on this issue. Recent literature has focused on the analysis of disaggregate real exchange rates, cf. Crucini and Shintani (2008), Imbs et al. (2005), Crucini et al. (2005), Cheung et al. (2001), Yang (1997), Knetter (1993), etc. One of the common findings of these papers is that there is a considerable degree of heterogeneity across sectors. With respect to sectoral persistence and its relation with that observed at the aggregate level, the empirical findings appear to be disparate. Some authors have found large divergences between sectoral and aggregate reversion rates. Using Eurostat data, Imbs et al. (2005) report standard HL estimates for aggregate RERs in the range of 3-5 years and considerably lower HL estimates, around 1 year, when sectoral data is employed. They claim that the PPP puzzle arises as a consequence of an *aggregation bias* that affects aggregate estimates due to the high degree of heterogeneity in sectoral RERs which neither standard time series nor panel data techniques are able to control. On the other hand, it has also been pointed out that the *aggregation bias* appears not to be a robust feature in the data. Crucini and Shintani (2008) analyze a micro-panel of local currency prices of individual retail goods and services in major cities and find that the median level of persistence (across goods) is similar to that obtained for the aggregate RER (HL around 12-19 months).

The theoretical results recently presented in Mayoral (2009) help to clarify the contrasting empirical findings outlined above. She has studied the relations between measures of persistence computed at different aggregation levels and has shown that there is a close connection between them. In particular, in a linear setting similar to that considered in Imbs et al. (2005) and Crucini and Shintani (2008), it has been proven that the impulse response function (IRF) computed with aggregate data equals the average of the sectoral impulse responses.[1] A similar relation holds for

---

[1]It has also been shown that sectoral heterogeneity does not necessarily induce a bias in aggregate estimates

the scalar measures associated with the IRF, such as the cumulative impulse response (CIR). Gadea and Mayoral (2009) have used these results to show that, in fact, this relation between IRFs holds very closely for Imbs et al. (2005)'s dataset, implying that aggregate and average sectoral speeds of reversion to parity are very similar.

These theoretical results are the starting point of this paper. They imply that, since aggregate persistence -as measured by the IRF or the associated scalar tools- is completely determined by the behavior of the sectors, the aggregate HL or CIR can be estimated by using either aggregate or sectoral data. However, by using sectoral data, it will be possible to break down aggregate persistence into the persistence of its different subcomponents, thereby obtaining a lot of valuable information about the sources of aggregate persistence. Another interesting implication of these results is that they unveil the nature of the relation between sectoral and aggregate persistence: the aggregate response to a shock is the average of the individual responses and, since averages are very non-robust measures, a situation where most sectors present relatively quick reversion to parity but where a few of them are highly persistent, is compatible with a highly persistent aggregate RER.

The goal of this paper is to shed further light on the causes of the slow reversion to parity of aggregate exchange rates through the analysis of sectoral data. To do this, a novel methodology has been employed on a group of EU-15 real exchange rates (defined against the U.K. pound) for which highly disaggregate price data is available (aggregate prices are broken down into 94 sectors). Our strategy is twofold. Firstly, aggregate persistence is broken down into the persistence of its different components and the distribution of sectoral shock responses is analyzed. This will allow us to identify interesting features about the sources that determine aggregate persistence and to show that aggregate persistence is determined, to a great extent, by the behavior of a limited number of sectors in the upper quantiles of the distribution of persistence. Secondly, we investigate more thoroughly the factors that account for the slow reversion to parity. Several theories have been proposed to explain the slow speed of convergence of RERs. Among these theories, the lack of arbitrage in nontradable goods and the existence of imperfect competition due to pricing-to-market have prominent roles. Using data on market structure and international trade, a standard and a quantile panel regression analysis to test these theories have been performed. Since it is

computed using standard time series techniques.

3

well known that the resulting estimates of various effects on the conditional mean of sectoral persistence are not necessarily indicative of the size and nature of these effects on the upper tail of the distribution, special emphasis has been placed on explaining the behavior of the upper quantiles of the distribution of sectoral persistence because, as mentioned before, they determine, to a large degree, the persistence observed at the aggregate level. In this way, a more complete picture of covariate effects can be provided by estimating a family of conditional quantile functions.

Our main results can be summarized as follows. Firstly, a high degree of heterogeneity as well as high and positive skewness in the speed of reversion of EU sectoral RERs are documented. By breaking down aggregate persistence into the persistence of its different components and arranging them according to their contribution to the aggregate HL, it is shown that 70% of the least persistent sectors only account, on average, for 20% of the HL. By contrast, the most persistent 8 % are able to account for 40% of the HL. This means that understanding the nature of the PPP puzzle is equivalent to understanding what drives the persistence of the upper quantiles of the distribution of sectoral persistence. Sectors belonging to the durable category are the ones that show the lowest speed of reversion to parity: on average, they account for more than 40% of the long-run cumulative effect of shocks to aggregate RERs. Accordingly, it is found that tradables are in general more persistent than nontradables in this dataset. The conditional distribution of sectors in the upper quantiles have also been explored. The results of this analysis, in agreement with the aforementioned, show that the durable category is heavily overrepresented among the top 30% most persistent sectors. On average, its weight within this group exceeds its unconditional weight by 35%. Moreover, the correlation between these "excess weights" that durable goods have in the upper quantiles and the aggregate country HLs is greater than 70%, suggesting a strong relation between aggregate RER persistence and the existence of highly persistent durable goods in the corresponding country.

Secondly, our quantile panel regression analysis shows that variables related to the price stickiness of the final good and the market structure and the degree of competition of the intermediate inputs have a significant effect on sectoral persistence. Furthermore, the impact of these variables tends to be more important the higher the quantile considered. Interestingly, once the market structure of the intermediate inputs has been taken into account, that of the final goods does not appear

4

to be important in explaining sectoral persistence. Finally, variables related with the tradability of the goods are not significant either, implying that traditional theories that attribute the slow speed of reversion of RERs to the existence of nontraded goods in the consumption basket do not explain EU current trade patterns very well. These conclusions are in agreement with modern trade theories (cf. Chari et al., 2002, Carvalho and Nechio, 2008).

The outline of the paper is as follows. Section 2 presents a brief overview of the literature that has dealt with RER persistence at different levels of aggregation. Section 3 introduces the different datasets used in this paper, provides estimates of the IRFs computed with both aggregate and sectoral data and tests whether these estimates are actually equal, as theory predicts. Section 4 explores the distribution of sectoral persistence while Section 5 analyzes whether the traditional theories (lack of arbitrage due to nontradability or imperfect competition and price stickiness) are able to explain the distribution of sectoral persistence. Section 6 concludes. An Appendix presents additional explanations about the data and the econometric techniques used throughout the article.[2]

2. HETEROGENEITY, AGGREGATION AND PERSISTENCE.

The literature that studies the behavior of RERs has extensively documented two well known facts. Firstly, aggregate RERs may converge to parity in the long run, although the rate at which it happens is very slow (Rogoff, 1996). Secondly, impediments to arbitrage, nominal rigidities and market structure vary considerably across sectors in the economy and, since these impediments are usually believed to be behind cross-country price differences, they give rise to important heterogeneity in the speed of reversion to parity across sectors.

Starting with the contribution by Imbs et al. (2005a) (IMRR henceforth), several papers have looked at the causal relation between heterogeneity in sectoral exchanges rates and the slow speed of reversion observed at the aggregate level. IMRR argue that the slow speed of reversion to parity of aggregate RERs is consistent with a much faster average adjustment of disaggregate relative prices: estimates of the persistence of sectoral RERs are, on average, considerably smaller than those of the aggregate rate itself. According to their view, estimates of aggregate persistence are

based upon the implicit assumption that relative prices converge to parity at the same speed. The failure to allow for heterogeneity in adjustment dynamics at the sectoral level induces a positive bias in persistence estimates when aggregate data is employed.

To illustrate their main arguments, a very simple model for the sectoral exchange rate in sector i, $q_{i,t}$, which allows for heterogeneous speed of price adjustment, has been considered,

$$q_{i,t} = \gamma_i + \alpha_i q_{i,t-1} + \nu_{i,t} \tag{1}$$

$$\nu_{i,t} = \rho_i u_t + \varepsilon_{i,t} \tag{2}$$

where $q_{i,t} = p_{i,t} - p_{i,t}^* - s_t$, $p_{i,t}$ and $p_{i,t}^*$ are the logs of the price of good i in the domestic and foreign countries respectively, and s_t is the log of the domestic currency price of foreign currency. The parameter α_i varies across sectors and captures the heterogeneity in price adjustment dynamics, while ρ_i measures the impact of the common shock u_t on sector i. These coefficients are assumed to be random draws from some random variables, α and ρ. The support of α cannot exceed 1 and, without loss of generality, the expected value of ρ is normalized to 1. The processes u_t and $\varepsilon_{.,t}$ represent an aggregate and an idiosyncratic shock, respectively, and are assumed to be mutually-independent, zero-mean $i.i.d.$ processes. Finally, $\varepsilon_{.,t}$ is assumed to be independent from α.

The aggregate RER can be obtained as the expected value of the sectoral RERs (Stoker 1984, Lewbel, 1994). Under the previous assumptions, the moving average representation of $q_{i,t} = \sum_{j=0}^{\infty} \alpha_i^j L^j (\rho u_t + \varepsilon_{i,t})$. Thus, it is easy to obtain the dynamics of the aggregate process as the expected value over the distribution of sectors of the former expression[3]

$$Q_t = E_s(q_t) = E_s \left(\rho + \alpha \rho L + \alpha^2 \rho L^2 ... \right) u_t$$

$$+ E_s \left(1 + \alpha L + \alpha^2 L^2 ... \right) E_s(\varepsilon_{.,t}) = \sum_{j=0}^{\infty} E_s(\alpha^j \rho) u_{t-j}, \tag{3}$$

where $E_s(.)$ denotes expectation across the distribution of sectors. To quantify the size of the *aggregation bias*, IMRR compare estimates of the HL computed at different degrees of aggregation. For a simple model such as in (1), IMRR compute the *sectoral IRF* as the IRF of the process $\bar{q}_t = \bar{\alpha} \bar{q}_{t-1} + a_t$, where $\bar{\alpha} = E_s(\alpha)$, which is given by

$$IRF_{\bar{q}}(h) = \bar{\alpha}^h \tag{4}$$

---

[3]It is also assumed that $\sum_{j=0}^{\infty} E\left(\alpha^j\right)^2 < \infty$. The stringent assumptions in this section are only imposed to simplify the explantion but the main arguments are valid in a more general setting (see Mayoral, 2009 for details).

The IRF of Q_t to a unitary shock in u_t can be easily obtained from (3) as follows

$$IRF_Q(h) = E_s(\alpha^h \rho). \tag{5}$$

IMRR claim that HLs obtained from (5) are considerably larger than those obtained from (4). Using Eurostat data, they report estimates of the aggregate HL in the range of 3-5 years, while those based on sectoral data are considerably smaller, around 1 year. Thus, they claim that the PPP puzzle is due to a positive *bias* in aggregate estimates arising because individual heterogeneity is not explicitly taken into account by standard estimates.

Chen and Engel (2005) provide a number of criticisms of the methods of IMRR. From a theoretical point of view, they argue that the analytical results provided by IMRR to justify the aggregation bias argument are different from the claims made in their empirical work. While in the latter they focus on the behavior of the HL across aggregation levels, analytically they demonstrate that the first autocorrelation of Q_t is larger than the average of the first autocorrelations of $q_{i,t}$. They notice that, firstly, there is not a direct relationship between the first autocorrelation and the HL in the general case and, secondly, that the first correlation of Q_t can be larger or smaller than the average of the first correlations of $q_{i,t}$, as opposed to the claims made by IMRR. From an empirical perspective, they argue that Eurostat data is plagued with measurement error that could make the series appear much less persistent than what they actually are. Using similar data as IMRR, but with corrections for data entry errors and for small-sample bias, Chen and Engel show that the aggregation bias, as measured by the difference between the HLs in the aggregate model and that obtained from an IRF similar to that in (4) is, in fact, small.[4]

These results have been corroborated by Crucini and Shintani (2008). Using a different data set, an extensive annual micro panel on local currency prices of individual retail goods and services in major cities from 1990 to 2005, and a different estimation strategy, they find that aggregation bias is not a robust feature in the data. While they are able to find some evidence for the U.S., they fail to do so for the rest of the international locations in their data set.

Gadea and Mayoral (2009, GM henceforth) have put forward a further objection to IMRR's methods. Building on the results in Mayoral (2009), GM propose to measure average sectoral

---

[4]However, Imbs et al (2005b) replied to this paper and showed that their results survive each of the criticisms raised by CE.

persistence by averaging the individual IRFs, rather than by averaging the AR coefficients and then, computing the IRF of the resulting process as IMRR do. For the simple models considered here, the impulse response of the sector i to a unitary change in the aggregate shock u_t, h periods ahead, is given by

$$IRF_{q_i}^u(h) = \alpha_i^h \rho_i, \tag{6}$$

and averaging across sectors it is obtained that

$$\overline{IRF_q^u}(h) = E_s(\alpha^h \rho). \tag{7}$$

Notice that if α and ρ are independent then

$$\overline{IRF_q^u}(h) = E_s(\alpha^h), \tag{8}$$

since $E_s(\rho) = 1$.

GM argue that the different persistence behavior between aggregate and sectoral exchange rates reported by IMRR is not due to an upward bias in aggregate data estimates but rather, to a negative bias affecting IMRR's sectoral persistence estimates. Since the IRF is a highly non linear function, averaging the AR coefficients and then computing the IRF, as IMRR do, or averaging the individual responses, may yield very different results. In fact, Jensen's inequality ensures that, for most empirically relevant cases, the former measure is larger than the latter. The intuition of this result is clear: the IRF is a convex function, hence, when averaging the individual responses, highly persistent sectors increase the mean considerably. However, in the computations of IMRR, the latter sectors are eliminated in the first stage when the model coefficient estimates are averaged. This translates, not surprisingly, into lower estimates of persistence. Defining the aggregation bias (AB) as the difference between the aggregate and the average of the sectoral IRFs to a unitary shock in u_t, it is easy to see that[5]

$$AB = IRF_Q(h) - \overline{IRF_q^u}(h) = 0, \text{ for all } h. \tag{9a}$$

---

[5]See Mayoral (2009) for a complete description of this result developed in a more general model. She has also shown that other persistence measures routinely employed in applications, such as the sum of the autoregressive coefficients or the largest autoregressive root, are not invariant to aggregation. Furthermore, as long as a sufficiently long autoregression is fit to the data (as IMRR actually do) standard time series methods can be used to obtain consistent aggregate and sectoral estimates under heterogeneity.

Using similar data and estimation strategy as those employed in IMRR's paper, Gadea and Mayoral (2009) have shown that the aggregation bias is not statistically different from zero.[6] Furthermore, standard estimates of the HL, in the range of 3-5 years are obtained using either sectoral or aggregate data.

Carvalho and Nechio (2010, henceforth CN) have introduced a theoretical model with sticky prices that departs from the existing literature by allowing for heterogeneity in the frequency of price changes across sectors. Combined with local currency pricing, these differences make sectoral RERs have heterogeneous dynamics. To gauge the impact of allowing for heterogeneity, they compare the persistence implied by two models sharing the same average frequency of price changes, one that allows for heterogeneity and another that does not. They denote as *Total Heterogeneity Effect (THE)* the difference in persistence implied by these two models. This quantity is further broken down into two components: a *counterfactuality effect* (CE) and an *aggregation effect (AE)*. Using the IRF(h) as persistence measure, CN define *THE, CE* and *AE* as follows

$$THE = IRF_Q\left(h\right) - IRF_{\bar{q}}\left(h\right) = CE + AE,$$

where

$$CE = \overline{IRF_q^v}\left(h\right) - IRF_{\bar{q}}\left(h\right),$$

$$AE = IRF_Q\left(h\right) - \overline{IRF_q^v}\left(h\right), \tag{10}$$

and $\overline{IRF_q^v}\left(h\right)$ is the average over the sectors of the response of $q_{i,t}$ to a unitary change in the reduced-form shock $\nu_{i,t}$, given by

$$\overline{IRF_q^v}\left(h\right) = E_s\left(\alpha^h\right). \tag{11}$$

CN use this decomposition to account for the different results found in the literature. They argue that different authors have estimated different objects. While Imbs et al. (2005) provide estimates of *THE*, GM and Crucini and Shintani (2008) provide estimates that are closer to the AE. CN argue that both AE and CE are strictly greater than zero but that CE accounts for the greater part of *THE* in IMRR's data while the AE is fairly small.

---

[6]Mayoral (2009) has shown that, as long as a sufficiently long autoregression is fit to the data (as IMRR actually do) standard time series methods can be used to obtain consistent aggregate and sectoral estimates under heterogeneity.

Finally, notice that the AE and the AB, defined in (10) and in $(9a)$, respectively, are very close definitions. The difference between them stems from the shocks which are assumed to change when computing the IRFs. CN consider a unitary change in the aggregate shock, u_t, and a unitary change in the reduced form shock, $\nu_{i,t}$, to compute the aggregate and the individual IRFs, respectively. On the other hand, in the computation of AB the aggregate and the individual IRFs are both based on unitary changes in the aggregate shock u_t.

Notice further that, under independence of α and ρ, $\overline{IRF_q^v}(h) = \overline{IRF_q^u}(h)$, as (8) and (11) show. This implies that $AE = AB = 0$, that is, there is no aggregation effect even if the individual IRFs are computed with respect to their reduced form shock $\nu_{i,t}$. If α and ρ are not independent, the direction of the correlation between α_1^h and ρ will determine whether AE is positive or negative. In CN's model, this correlation is actually positive (meaning that higher $\alpha's$ obtain greater weights, on average) and this is why they obtain that $AE > 0$. If the correlation were negative, however, the opposite would be obtained. CN's observation that AE is small in Eurostat data can be interpreted as saying that the correlation between α^h and ρ is close to zero.

In our opinion, considering aggregate and sectoral persistence referring to the same common shock has two advantages: firstly, it enables a better understanding of the role of aggregation on the transmission of shocks, since the same shock is hitting both the sectors and the aggregate process; secondly, it allows us to pin down a close relation between sectoral and aggregate behavior, since AB=0. As a consequence, aggregate persistence can be estimated using either sectoral or aggregate data. This relation will be exploited in the rest of the paper to break down aggregate RER persistence into the persistence of the sectors. In addition, using sectoral data to obtain aggregate estimates yields estimators that are considerably more efficient, as the next section will show.

The following section presents the data and the econometric models employed in our empirical exercise. It also provides estimates of AB and AE and tests whether these quantities are equal to zero.

3. AGGREGATE AND SECTORAL RER PERSISTENCE

This section describes the data and the econometric methods employed in the empirical analysis. It is shown that measures of persistence based on the IRF and estimated with aggregate or sectoral data are, in fact, very close, as the theory presented in the previous section predicts. Similar persistence values to previous studies have been obtained for most countries, with half lives of deviations from PPP falling well within the "*consensus view*" of 3 to 5 years.

3.1. The data

Eurostat Harmonized Index of Consumer Prices (HICP) for 11 European countries ranging from 1996:1 to 2007:12 have been employed. These countries and their corresponding abbreviations are Austria (AU), Belgium (BE), Denmark (DK), Finland (FI), France (FR), Germany (GE), Italy (IT), Netherlands (NL), Spain (SP), Sweden (SW) and United Kingdom (UK).[7] Eurostat provides data corresponding to different levels of aggregation. We focus on the most disaggregate level, which contains prices relative to 94 sectors.[8] Nominal exchange rates are obtained from the Main Economic Indicators of the OECD and RERs are defined against the U.K. pound. Aggregate RERs have been constructed by aggregating sectoral RERs using Eurostat price weights.[9]

In order to account for the persistence of European exchange rates, three additional datasets have been employed: the Comtrade (United Nations Commodity Trade Statistic Database), the OECD Structural Analysis Statistics (STAN, Edition 2008) and the Input-Output Tables (IOT) from the OECD. The first of these provides information about trade flows for individual countries at six digits of disaggregation and allows us to calculate different indicators that capture the trade features of each sector. The second contains data about value added, gross output, labor costs and other industry indicators that will be employed to analyze the market structure of the sectors. Finally,

---

[7]We originally considered all the EU-15 countries. However, four of these countries (Portugal, Luxembourg, Greece and Ireland) present important data availability limitations for the other datasets employed in this paper. For this reason, we decided to drop them from the analysis.

[8]Price data on some sectors is missing for some countries. See Eurostat for more details.

[9]The weights employed to aggregate sectoral RERs are the average over the period 1996-2007 of those used by Eurostat to aggregate sectoral prices in country *c*. The original aggregate price indices provided by Eurostat were also employed to compute aggregate RERs. Results obtained with either type of aggregate data were almost identical.

the IOT allows us to determine the type and proportion of inputs employed in the production of the final goods in our price dataset. More details on these databases are provided in the Appendix.

3.2. Econometric models and methods

It is assumed that sectoral RERs follow a linear specification, similar to that in (1) and (2). More specifically, the RER of sector i in country c at time t, $q_{i,t}$, is given by[10]

$$q_{i,t} = \alpha_{i,0} + \sum_{k=1}^{K} \alpha_{i,k} q_{i,t-k} + \nu_{i,t}, \qquad \text{for } t = 1, ..., T; \ i = 1, ..., N. \tag{12}$$

$$\nu_{i,t} = \rho_i u_t + \varepsilon_{it} \tag{13}$$

As shown by Lewbel, the aggregate of (12) follows an AR(∞) process. Thus,

$$Q_t = \sum_{j=0}^{\infty} A_i Q_{t-i} + u_t. \tag{14}$$

Sectoral IRFs, defined as the responses of $q_{i,t}$ to unitary changes in u_t and ν_t, have been considered. The impulse response of sector i to a unitary change in u_t or ν_t is given, respectively, by $IRF_{q_i}^{u}(0) = \rho_i$ and $IRF_{q_i}^{\nu}(0) = 1$ and

$$IRF_{q_i}^{g}(h) = \sum_{k=1}^{h} \alpha_{i,k} IRF_{q_i}^{g}(h-k) \text{ for } g = \{u, \ \nu\} \text{ and } h \geq 1. \tag{15}$$

Average sectoral persistence is computed as the mean of the individual IRFs, that is, $\overline{IRF_{q}^{g}}(h) = \sum_{i=1}^{N} \omega_i IRF_i^{g}(h)$, for $g = \{u, \nu\}$, where ω_i are Eurostat sectoral weights.

The IRF of Q_t to a unitary change in u_t is given by $IRF_Q(0) = 1$ and

$$IRF_Q(h) = \sum_{j=1}^{h} A_j IRF_Q(h-j), \text{ for } h \geq 1. \tag{16}$$

---

[10] As a preliminary analysis, panel unit root tests have been applied to EU RERs (Levin et al., 2002 and Im et al., 2003). The results, not reported for reasons of space, show that the unit root hypothesis could be rejected in most cases, in line with the previous literature (Frankel and Rose, 1996, Wei and Parsley, 1996 and Taylor and Sarno, 1998). In accordance with these results, the estimation of the parameters has been restricted so that the sum of the autoregressive coefficients is smaller than 1.

Since the IRF is a vector of numbers, it is customary to use scalar measures instead. Two of these scalar tools are employed: the half life (HL) and the cumulative impulse response (CIR). The HL is defined as the value of the IRF that satisfies

$$IRF_Q(t, h = HL) = 0.5, \tag{17}$$

or, alternatively, $\overline{IRF_q^u}(h = HL) = 0.5$ if sectoral data is employed (see Mayoral, 2009). As in Kilian and Zha (2002), the HL is defined as the largest value of HL such that $\mathrm{IRF}_Q(t, HL - 1) \geq 0.5$ and $\mathrm{IRF}_Q(t, HL + 1) < 0.5$.

The cumulative impulse response (CIR) at horizon h is defined as

$$CIR(h) = \sum_{l=0}^{h} IRF_Q(t, l), \tag{18}$$

or, equivalently, $CIR(h) = \sum_{l=0}^{h} \overline{IRF_q^u}(l)$.

3.2.1. Estimation.—

To obtain estimates of the IRFs and the scalar measures of persistence, the following methods have been employed. Under sectoral heterogeneity, Q_t might display complicated dynamics, as equation (14) shows. Following Berk (1974), AR(K) with large K have been fitted to aggregate RERs.[11] Since RERs are, in general, highly persistent, bias-corrected estimates have been computed (see the Appendix). Then, estimates of the IRF, denoted as $\widehat{IRF_Q}$, have been obtained based on expression (16).

To obtain $\widehat{\overline{IRF_q^v}}(h)$, bias-corrected $AR(k)$ processes have been fitted to sectoral data and the output has been plugged to (15) to obtain estimates of $\overline{IRF_q^v}(h)$.

Computation of $\widehat{\overline{IRF_q^u}}(h)$ requires estimates of all the model parameters, $\alpha_i s$ and ρ_i. The following approach has been adopted: firstly, bias-corrected estimates of the AR parameters and of the common shock u_t, $\hat{u}_t$, have been obtained. Next, $\hat{u}_t$ has been introduced in the AR equation as

---

[11] Berk (1974) has demonstrated that consistent estimates of an AR(∞) process can be obtained provided an autoreggresion of order K is fitted to the data, where K tends to infinity with the sample size. Furthermore, if the general-to specific (GTS) approach (Ng and Perron, 1995) is employed to select K, consistent and asymptotically normal estimates of the model parameters can be obtained (Kursteiner, 2005). Thus, the GTS method (with a maximum number of AR terms equal to 36) has been employed throughout.

an additional regressor and ρ_i has been estimated by OLS. To estimate u_t, three approaches have been considered: firstly, sectoral $RERs$ have been fitted as described above, the residuals of these estimations have been averaged and this average has been taken as an estimate of u_t.[12] Secondly, $\hat{u}_t$ has been obtained from the residuals of the aggregate RERs equations. A third estimate of u_t has been obtained by factor analysis. A Monte Carlo experiment has been performed to assess the performance of these methods. It turns out that the three of them perform well and quite similarly. The second approach has been employed in the empirical analysis since it is the less demanding from a computational point of view. (See the Appendix for details on the estimation and the Monte Carlo experiment).

Confidence intervals for the IRFs have been obtained using Boostrap methods. Details of the calculations can be found in the Appendix.

3.2.2. Testing.—

Formal tests to check whether persistence estimates obtained from sectoral or aggregate data are the same have been implemented. In addition, whether the two sectoral IRFs, $\overline{IRF_q^v}(h)$ and $\overline{IRF_q^u}(h)$, are actually the same function, has also been tested. The null hypotheses of these tests have been formulated as follows:

$$\text{Test A. } H_0^A : \sum_{h=0}^{G} (\overline{IRF_q^v}(h) - \overline{IRF_q^u}(h))^2 = 0$$

$$\text{Test B. } H_0^B : \sum_{h=0}^{G} (IRF_Q(h) - \overline{IRF_q^u}(h))^2 = 0$$

$$\text{Test C. } H_0^C : \sum_{h=0}^{G} (IRF_Q(h) - \overline{IRF_q^v}(h))^2 = 0$$

H_0^A entails that the two sectoral IRFs described above are identical up to lag G. H_0^B and H_0^C correspond to the hypothesis that AB=0 and AE=0, respectively. The horizon G has been set to 60 (5 years). Boostrap techniques have been employed to test the former hypotheses (see the Supplementary Material in the Appendix for detailed explanations).

---

[12] Notice that these residuals are an estimate of the average of the common and the idiosyncratic component and that, by the law of large numbers, the latter term tends to zero in probability.

3.3. Empirical results

Figures 1a, 1b present the plots of $\widehat{IRF_Q}$, $\widehat{\overline{IRF_q^u}}(h)$ and $\widehat{\overline{IRF_q^v}}(h)$ as well as their confidence bands at the 5% significance level for each of the countries in our dataset.[13]

These plots show that, although European Union markets are very integrated, EU RERs are highly persistent. Estimates of the IRFs based on aggregate or sectoral data are, in general, very similar, as predicted by the theoretical results in Section 2. In addition, estimates of $\widehat{\overline{IRF_q^u}}(h)$ and $\widehat{\overline{IRF_q^v}}(h)$ are strinkigly similar. Interestingly, estimates based on sectoral data are considerably more efficient than those based on aggregate persistence, as shown by the fact that aggregate CIs are much wider than sectoral ones.

Table I displays some summary statistics corresponding to the estimated IRFs, namely, the HL and four values of the $CIR_c(h)$ (for h corresponding to 1, 2, 3 and 5 years). Agg, Sec^u and Sec^v correspond to measures obtained from $\widehat{IRF_Q}$, $\widehat{\overline{IRF_q^u}}(h)$ and $\widehat{\overline{IRF_q^v}}(h)$, respectively. The values of the HLs are in line with those found in previous studies, with most HLs located in the so-called Rogoff's interval, that is, displaying values between 3 to 5 years. In addition, the different measures of aggregate persistence computed with aggregate or disaggregate data are very close, corroborating the graphs in Figures 1a, 1b.

(Figure 1a about here)

(Figure 1b about here)

---

[13] Confidence bands for $\widehat{\overline{IRF_q^v}}(h)$ are omitted for the sake of clarity, since they are very similar to those of $\widehat{\overline{IRF_q^u}}(h)$.

TABLE I

RER PERSISTENCE WITH AGGREGATE AND SECTORAL DATA

		AU	BE	DK	FI	FR	GE	IT	NL	SP	SW
	Agg	37.52	40.18	44.08	36.04	37.21	31.32	60.51	44.19	118.53	22.38
HL	Sec^v	36.57	37.26	37.33	12.02	37.61	18.33	36.99	37.32	85.47	20.06
	Sec^u	36.75	37.32	36.99	12.06	37.66	18.39	37.23	37.42	91.11	20.30
	Agg	9.55	9.05	10.41	8.64	10.00	9.53	10.35	9.91	10.92	8.87
CIR(12)	Sec^v	8.60	8.59	9.10	8.03	9.03	8.29	8.61	8.93	9.25	8.15
	Sec^u	8.55	8.59	8.96	7.92	8.98	9.31	8.59	8.78	9.22	8.38
	Agg	16.02	15.37	18.55	13.90	17.11	15.54	19.60	17.64	21.48	15.94
CIR(24)	Sec^v	13.90	14.30	15.28	12.54	15.52	13.24	14.68	14.80	16.90	14.28
	Sec^u	13.87	14.30	14.99	12.36	15.49	13.32	14.74	14.64	16.95	14.81
	Agg	22.59	23.04	27.42	18.50	23.96	21.28	31.19	26.56	35.94	17.88
CIR(36)	Sec^v	18.99	20.22	21.24	16.56	22.04	18.04	20.96	21.02	25.97	17.48
	Sec^u	19.08	20.21	20.76	16.29	22.03	18.18	21.21	20.83	26.07	18.33
	Agg	27.91	31.82	37.80	18.48	28.11	24.33	48.88	37.52	67.57	15.65
CIR(60)	Sec^v	26.76	29.14	30.61	19.75	30.23	24.66	28.38	28.88	42.55	22.39
	Sec^u	27.08	29.14	29.86	19.51	30.31	24.89	29.20	28.91	43.09	23.35

Agg, Sec^v and Sec^u correspond to measures obtained from $\widehat{IRF_Q}$, $\widehat{\overline{IRF_q^v}}(h)$ and $\widehat{\overline{IRF_q^u}}(h)$, respectively.

In view of Table I, it is not surprising that the null hypotheses H_0^A, H_0^B and H_0^C could not be rejected for any of the countries in our data set.[14] This has two important practical implications: firstly, the fact that H_0^A cannot be rejected implies that considering sectoral IRFs that refer to changes in the aggregate or the reduced-form shock is irrelevant in our sample. Since $\widehat{\overline{IRF_q^v}}(h)$ are easier to compute, the rest of the paper will use these measures as estimates of sectoral persistence. To simplify the notation the superindex 'v' will be dropped hereinafter. Secondly, the fact that neither H_0^B nor H_0^C could be rejected implies that estimates of aggregate persistence can be obtained using either sectoral or aggregate data.

[14]Figures are reported in the Supplementary Material appendix.

4. THE DISTRIBUTION OF SECTORAL PERSISTENCE

The results obtained in the previous section motivate the remainder of our empirical analysis. Since the persistence of the aggregate real exchange rate is a simple function of that of the sectors, a lot of information about the sources of aggregate persistence can be obtained by analyzing sectoral behavior.

This section describes the distribution of sectoral persistence and breaks down aggregate RER persistence into the persistence of its components. Sectoral persistence is highly heterogeneous and skewed to the right, that is, a reduced number of sectors are highly persistent while the rest are considerably less so. Since aggregate response is the average of the individual responses, the former characteristics imply that persistence is likely to be driven by a few sectors at the right tail of the distribution of persistence. A careful examination of the characteristics of these sectors and what causes them to be so persistent may provide valuable insight into the forces that shape aggregate behavior. This section provides a description of these sectors while Section 5 uses quantile regression to investigate more formally the factors that can account for the persistence in the upper quantiles.

4.1. Heterogeneity and skewness in RER sectoral data

Many studies have documented the existence of a high degree of heterogeneity in sectoral RER data (see Imbs et al., 2005, Crucini and Shintani, 2008, for recent references).

The dataset considered in this paper is not an exception. Figure 2 displays the scatter plots of HLs and CIR(h)s at several horizons, namely, $h = \{12, 36, 60\}$ months, corresponding to all sectors and all countries in our dataset. The ordering of the sectors corresponds to the average country. Vertical lines depict the values of the former persistence measures for a particular sector and the different countries in the sample.

(Figure 2 about here)

As Figure 2 shows, cross-sectional heterogeneity is very important and it increases considerably with the horizon of the IRF. For the HL, country heterogeneity is also considerable, although it is

17

smaller than the sectoral one.

Another aspect of the distribution of persistence that is worth emphasizing is the fact that it is highly skewed to the right, implying that most sectors present moderate inertia while a few of them are very persistent. This implies that the latter sectors drive average persistence upwards in spite of the fact that most sectors are considerably less so.

To illustrate this fact, the following exercise has been carried out. Firstly, for each country in our dataset, sectoral RERs have been ranked according to their contribution to the HL. The contribution of sector i to the HL has been computed as the value of its corresponding IRF evaluated in $h = HL$, times its relative weight, that is, $\omega_i IRF_i (h = HL)$. Secondly, sectors have been grouped into 5 categories, such that the contribution of each group amounts to 20% of the total HL. Thus, the first bin gathers the first N_1 sectors such that

$$\frac{\sum_{i=1}^{N_1} \omega_i IRF_{c,i} (t,h)}{\sum_{i=1}^{N} \omega_i IRF_{c,i} (t,h)} = 20\%.$$

The remaining groups have been computed in a similar way. Figures 3a and 3b display the graphs containing the percentage of sectors in each of these bins. The first one, that only accounts for 20% of the aggregate HL, contains 69% of the sectors, on average. By contrast, the percentage of sectors in the fifth and fourth bins is only 3 and 5%, respectively, on average. It follows that around 8% of sectors account for a large part of the HL (40%). Thus, a careful examination of the composition of this right tail of the persistence distribution will provide valuable information as to the sources of aggregate RER persistence.[15]

(**Figure 3a about here**)

(**Figure 3b about here**)

4.3. Decomposing aggregate persistence by groups of sectors

The relation between the aggregate and the sectoral IRFs established in (**??**) allows us to quantify the contribution of each sector or group of sectors to aggregate persistence. To analyze the relation between persistence and sector characteristics, sectors have been grouped into two long-established

---

[15]See Appendix A for a description of the composition of bins four and five.

categories. The first one classifies them as food (F), durable (D), nondurable (ND), services (S) and energy (E). The second one, following the traditional dichotomy, as traded (T) and nontraded (NT).[16] Next, the contribution of these groups to total persistence is computed. The relative composition of the top 30% most persistent sectors has also been analyzed since, as mentioned in the previous section, they are responsible to a large extent for the high persistence observed in aggregate RERs.

To evaluate the relative contributions of the groups, we have proceeded as follows: the aggregate IRF can be broken down as

$$IRF_Q(h) = \sum_{j=1}^{J} \sum_{i=1}^{N_j} \omega_i^j IRF_{q_i}^u(h),$$

where J denotes the number of groups considered, ω_i^j is the weight associated with sector i in group j, and N_j is the number of sectors in group j, with $\sum_{j=1}^{J} N_j = N$. The percentage contribution of group j to the aggregate IRF at horizon h, denoted as $PC_j(h)$, can be computed as

$$PC_j(h) = \frac{\sum_{i=1}^{N_j} \omega_i^j IRF_{q_i}(h)}{\sum_{i=1}^{N} \omega_i IRF_{q_i}(h)}.$$

Thus, the percentage contribution of group j to the aggregate cumulative response, $\text{PC-CIR}_j(h)$, is defined as

$$\text{PC-CIR}_j(h) = \sum_{r=1}^{h} PC_{c,j}(r). \tag{19}$$

The relative contribution of group j in country c to the HL, denoted as PC-HL_j, has been computed in the following way

$$\text{PC-HL}_j = \frac{\sum_{i=1}^{N_j} \omega_i^j IRF_{c,i}(h = HL)}{\sum_{i=1}^{N} \omega_i IRF_{c,i}(h = HL)}. \tag{20}$$

Table II presents the estimated values of the PC-HL and the PC-CIR(h) corresponding to the two categories considered above. Columns 2 to 6 display the average across countries of the PC-CIR$_j(h)$, for different values of h, and of PC-HL$_{c,j}$ defined in (19) and (20) respectively, for each of the groups. In addition, the first column of this table presents the average across countries of the weights that Eurostat assigns to each of these categories in order to build the price index. For

---

[16]Housing and services are considered as nontraded (with the exception of air travel and financial services) while all other sectors are considered as traded.

example, the average Eurostat weight (across countries) of all the products labeled as food is 23%. This column has been included in order to be able to evaluate whether the percentage contribution to total persistence is larger or smaller than the percentage weight in the aggregate RER.

It is possible to assess the evolution of the persistence of shocks over time across groups of sectors by looking at the different horizons of the CIR in columns 3 to 6. In the short run (one year), the impact of shocks in all groups is very similar, as shown by the fact that the contribution to CIR(12) of each group is almost equal to its corresponding initial weight. However, as longer horizons are analyzed, the picture changes substantially. Durable goods become the group with the highest contribution to long-run persistence. They account for 43% of the total cumulative effect of shocks in the long-term (CIR(84)). Moreover, their contribution to the cumulative response relative to the initial weight of the group increases substantially as longer horizons are considered. For instance, their contribution to CIR(60) (38%) and CIR(84) (43%) exceeds their corresponding initial weight (27%) by 41% and 59%, respectively. Within this group, the *electronic products* and the *clothing and personal effects* subcategories are the most persistent components. Their long-run contribution to the CIR exceeds their initial weights by 100% and 80%, respectively. On the other hand, the contribution of the services and energy sectors to aggregate persistence decreases when distant horizons are considered. Their contribution to $CIR(60)$ and $CIR(84)$ is only 60% and 50% that of their initial weight, respectively, for energy, and 76% and 61% in the case of services. This result is quite surprising because services are usually believed to be behind the high persistence of aggregate RER. It is also worth highlighting that the contribution of food and nondurables remains fairly constant over time and very similar to its initial weight.

The traded/nontraded goods categories also display a clear pattern: the percentage contribution of the traded goods category to total persistence is bigger than its initial weight and increases with the horizon considered. For instance, the contribution to long-run persistence, as measured by CIR(60) and CIR(84), exceeds its initial weight by 14% and 20%, respectively. So, the nontraded category seems to be less persistent than the traded one, as shown by the fact that its contribution to CIR(60) and CIR(84) is only 0.76% and 0.68% of its initial weight. Nevertheless, the discrepancies between the traded and nontraded groups appear to be considerably smaller than those obtained with the first classification described above, suggesting that the characteristics of these groups of

goods might not be that different, as has been pointed out by other authors (see Engel, 1999, Chari et al., 2002, Crucini and Shintani, 2008).

TABLE II

CONTRIBUTION TO AGGREGATE PERSISTENCE (IN %)

	Weights	PC-HL	PC-CIR(12)	PC-CIR(36)	PC-CIR(60)	PC-CIR(84)
1. FOOD	23	23	26	24	23	21
- Food	18	17	20	19	18	17
- Alcohol and Tobacco	5	0	6	5	5	5
2. ENERGY	10	6	9	8	6	5
3. NONDURABLES	6	8	6	7	8	8
4. DURABLES	27	41	25	31	38	43
- Clothing and personal effects	10	19	10	13	16	18
- Durables for the dwelling	6	8	5	6	8	9
- Motor vehicles	7	9	7	8	9	10
- Electronic products	2	3	2	2	3	4
- Recreational and cultural	2	2	1	2	2	2
5. SERVICES	34	23	34	29	26	21
- Services relating to the dwelling	10	7	9	9	7	6
- Transport	5	3	4	4	3	3
- Financial services	1	3	2	2	3	3
- Recreational and cultural services	15	8	16	12	11	9
- Other services	2	2	2	2	2	0
TOTAL	1	1	1	1	1	1
TRADED	62	74	61	66	71	74
NONTRADED	38	26	39	34	29	26
TOTAL	1	1	1	1	1	1

Summarizing, it seems that persistence is dominated by sectors in the durable category. If this is true, the durable group should be overrepresented in the right tail of the distribution of

21

sectoral persistence with respect to the whole sample. To investigate this issue further, the relative composition of the 30% most persistent sectors has been examined. The following simple exercise has been carried out. For each country, we have computed the weight of each of the groups considered above within the 30% most persistent sectors. In order to do this, we have simply added the relative weights assigned by Eurostat to the sectors in a particular category and have divided these figures by their sum so that the resulting weights add up to 1. Then, we have subtracted the weight that the corresponding category has in the whole sample.

The results can be summarized as follows: On average, sectors in the food, energy and nondurable categories are equally represented in the upper quantiles and in the overall distribution, as shown by the short bars associated with these groups. Durable goods and services, however, are over and underrepresented, respectively, in the right tail of the distribution. On average, the weight of the durable goods category in the upper quantiles exceeds its unconditional weight by 35%. Moreover, the correlation between these "excess weights" of durables and the country HLs (first column in Table II) is 72%, suggesting a strong link between the persistence of the durable component and aggregate RER persistence. It follows that the more overrepresented durables are among the most persistent sectors in country c, the more persistent the corresponding aggregate RER. This relation is very significant.

The opposite is found for the service category: its weight in the upper quantiles is 25% smaller than its corresponding unconditional weight. Accordingly, it is not surprising to see that it is the traded category, and not the nontraded one, which is overrepresented among the most persistent sectors, as opposed to what traditional theories would suggest. The weight of the traded category exceeds its unconditional weight by almost 20%, on average.

The results in this section imply that traditional theories based on the traded/nontraded dichotomy will fail to account for the persistence observed in European exchange rates. Aggregate persistence seems to be driven by the persistence of highly tradable sectors, as goods in the durable group are. In the following section, we carry out a more formal investigation using quantile regression analysis to determine the factors that account for the observed levels of persistence.

22

5. ACCOUNTING FOR RER PERSISTENCE

Explanations of the slow convergence to PPP have traditionally been related to one (or several) of the following theories: barriers to trade, such as tariffs or transportation costs, that can be high enough to prevent some goods and services from being traded and, therefore, arbitraged; imperfect competition practices, such as pricing to market (PTM), combined with price stickiness that are able to create a wedge between the prices of the same good sold in different markets, violating the Law of One Price (LOP) which is a building block of PPP; and different consumption preferences across countries that mean that inflation measurements are computed on different consumption baskets, so there is no reason for exchange rate changes to offset official measures of inflation differences.

The goal of this section is to investigate empirically whether these theories can account for the persistence observed in European RERs. In doing so, we will pay particular attention to explaining the behavior of the most persistent sectors since, as shown in the previous section, they are crucial to understanding the behavior of the aggregate RER.

By concentrating on harmonized sectoral price data, we can discard different consumption preferences as a source of deviations from PPP since we consider disaggregate prices referring to a homogeneous basket of goods. Thus, in the following, we concentrate on the first two potential explanations: the fact that some goods are nontraded and so cannot be arbitraged, and the lack of perfect competition combined with price stickiness in goods markets.

The conventional approach emphasizes the fact that many goods in the consumption basket are not traded (Salter, 1959 and Swan, 1960). Since the forces of arbitrage are, at best, weak on these goods, volatile and persistent aggregate RERs are to be expected. However, the empirical support for this theory is mixed. Engel (1999) suggests decomposing the RER into two components: one due to changes in two countries' relative prices of nontradable to tradable goods and the other due to changes in the countries' relative price of tradables. Under the classical theory, the second component obeys the LOP so all the movements in RER are due to movements in the first component. However, Engel (1999) finds that the opposite is true: nearly all the variability of the RER can be attributed to the second component. Similar findings have also been reported by Chari et al. (2002). Using a micro panel of local currency prices of individual retail goods and services

in major cities, Crucini and Shintani (2008) have found that the median HL of nontraded goods is higher than that corresponding to traded goods, although the difference is very small (2 to 6 months, depending on the geographical area considered).

In response to this mixed evidence, a new theory that combines pricing-to-market (PTM) behavior and nominal rigidities has emerged to account for the scarce effect of exchange rate movements on traded goods prices. Betts and Devereux (1996) made the initial contribution along these lines by augmenting the framework introduced by Obstfeld and Rogoff (1995) to allow for PTM. Bergin and Feenstra (2001) incorporate preferences that exhibit the property that the elasticity of demand is not constant. They show that this property is important to allow PTM and staggered contracts to generate endogenous persistence. In their model, a fraction ϕ of the goods are nontraded internationally and they emphasize that a lower degree of openness increases persistence. The intuition of this result is that, under a home currency depreciation, home producers will tend to lower the price they charge abroad. If home goods play a large role in the foreign consumption basket, then the foreign price index will go down, offsetting the impact of the depreciation of home currency on the real exchange rate. Nevertheless, the model is not able to reproduce the high persistence levels observed in the data.

Also in the spirit of Obstfeld and Rogoff (1995), Chari et al. (2002) consider a model where all goods can be traded and where PTM and staggered contracts occur at the intermediate goods level. This model is successful in replicating the volatility observed in RER data but it cannot match the degree of persistence observed in RER data unless implausible parameter values are assumed at the outset. By introducing heterogeneity in price stickiness across sectors, Carvalho and Nechio (2008) are able to match the observed RER persistence in a model which, otherwise, presents similar features to Chari et al. (2002)'s.

Hairault and Sopraseuth (2003) provide an integrated model with PTM and nontradables and show that both variables are deeply intertwined but their model suggests that PTM is more important than non tradables in accounting for RER volatility. See also Faruqee (1995), Chang and Devereux (1999) and Devereux and Engel (1998) for related references.

In that follows, we empirically identify the main determinants of the persistence observed in European RERs. More specifically, we analyze whether the widespread presence of nontraded

goods in the consumption basket (or in the inputs used in the elaboration of the final good), the prevalence of PTM and nominal rigidities, either at the final and the intermediate goods level, or a combination of both factors can account for the slow reversion to parity of European relative prices. Our analysis primarily focuses on analyzing the behavior of the upper quantiles of the distribution since, as shown in the previous section, they shape, to a large extent, the persistence observed at the aggregate level. We also explore how the explanatory power of the different variables considered to account for RER persistence change at different horizons of the evolution of the shock.

We now introduce the set of dependent and independent variables used to test the theories above, the econometric techniques employed in our empirical exercise and the results obtained.

5.1. Independent and dependent variables

The explanatory variables considered belong to three categories: those that are a proxy for market structure and imperfect competition, those that try to measure the degree of openness of the sectors and some control variables. As mentioned in Section 3, three additional databases have been employed to elaborate these variables: the United Nations Commodity Trade Statistic Database (UN Comtrade), the OECD Structural Analysis Statistics (STAN, 2008 edition) and the Input-Output Tables (IOT) from the OECD.[17] Additional details on the construction of the regressors are provided in the Appendix.

Variables related to market structure.

Intra-Industry Trade (IIT). As emphasized by Faruqee (1995), under PTM, the persistence of the real exchange rate increases as the elasticity of substitution between home and foreign varieties increases. The intuition is clear: as the elasticity of substitution rises, the domestic monopolistically competitive exporting firms have to maintain their prices more in line with those of the foreign firms, thus increasing price rigidities in local currency terms. Since domestic and foreign products are more substitutable under intra-industry trade, ceteris paribus, a greater degree of intra-industry trade leads to more persistent exchange rates.

---

[17] Data employed to construct the explanatory variables refer to 2003 for the STAN and the UN Comtrade databases and to 2002 for the IOT one.

We follow Grubel and Lloyd (1975) and define the IIT index for sector i in country c as

$$IIT_{c,i} = 1 - \frac{\sum_{i=1}^{N} |X_{c,i} - M_{c,i}|}{\sum_{i=1}^{N} (X_{c,i} + M_{c,i})}, \tag{21}$$

where $X_{c,i}$ ($M_{c,i}$) represents total exports (imports) of sector i in country c.[18]

As shown by Chari et al. (2002), the existence of PTM at the intermediate goods level can also have a big impact on the persistence of the relative price of the final good i even when it is sold in a perfectly competitive market. To capture this effect, we have also computed an index that measures the degree of IIT associated with the intermediate items needed to produce good i. To do this, we have calculated the weighted average of IIT indices (where each of them is defined as in (21)) for each of the inputs employed in the elaboration of good i, where the weights are the relative contribution of the corresponding input g to the production of good i, as stated by the Input-Output tables of country c. Thus, the intermediate goods IIT index (denoted as Input-IIT), is computed as

$$Input\text{-}IIT_{c,i} = \sum_{g=1}^{G} \omega_g IIT_{c,g}$$

where ω_g and G denote the share of good g and the total number of inputs involved in the production of good i, respectively, and $IIT_{c,g}$ is computed as in (21).

Price-cost margin (PCM). Imperfect competition will typically involve market segmentation and price discrimination across the destination markets (Goldberg and Knetter, 1997). A classical measure of imperfect competition is the price-cost margin variable, which approximates the degree of profitability of an industry. Thus, the lower the value of the PCM, the fiercer the competition in this sector. The PCM relative to sector i in country c is given by

$$PCM_{c,i} = \frac{VA_{c,i} - W_{c,i}}{VA_{c,i} + CM_{c,i}}, \tag{22}$$

where VA_i^c is the total value added of sector i (the value of total production minus the cost of materials) in country c, $W_{t,i}^c$ is labor compensation and $CM_{c,i}$ denotes the cost of materials. We have also computed the PCM associated with the inputs employed in the production of good i in a similar fashion to *Input-IIT$_{c,i}$*. It is defined as

---

[18] Data on trade of commodities and of services has been obtained from the UN Comtrade and from the OECD, respectively.

$$Input\text{-}PCM_{c,i} = \sum_{g=1}^{G} \omega_g PCM_{c,g},$$

where $PCM_{c,g}$ is the price-cost margin of input g defined as in (22).

Price Stickiness. It has also been emphasized that, without price stickiness, a model of PTM cannot generate adequate persistence (Chang and Devereux, 1998). To proxy price stickiness, we have considered the volatility of sectoral inflation (VOL). If the prices of sector i are very sticky, low levels of sectoral inflation volatility will be expected and viceversa. We compute the volatility of inflation for sector i in country c as

$$VOL_i = std(INFL_{t,i}^c),$$

where std denotes standard deviation and $INFL_{c,i,t}$ is the inflation rate of sector i in country c, defined as

$$INFL_{c,i,t}^c = 1200\,(p_{c,i,t} - p_{c,i,t-1}) \tag{23}$$

Variables related to the tradability of goods.

Openness. Conventional wisdom suggests that the more traded goods are, the more important the forces of arbitrage are and, therefore, the degree of openness should have a negative impact on RER persistence. In a different set-up, Bergin and Feenstra (2001) and Faruqee (1995) emphasize that, under PTM and nominal rigidities, an increase in openness fosters price adjustment when changes in the exchange rate take place, offsetting the impact of exchange rate movements and thus reducing RER persistence.

The degree of openness of sector i is measured as

$$OP_{c,i} = \frac{X_{c,i} + M_{c,i}}{GDP_{c,i}}. \tag{24}$$

where $GDP_{c,i}$ is the total GDP of sector i in country c for year 2003.

As for the group of variables related to market structure, we have also computed the degree of openness of the intermediate inputs, defined as

$$Input\text{-}OP_i^c = \sum_{g=1}^{G} \omega_g OP_g^c,$$

27

where $OP_{t,g}^c$ denotes the degree of openness of the intermediate good g, computed as in (24).[19]

Control variables

Finally, we have also controlled for another set of variables that may have an impact on RER persistence.

Inflation. It has been argued that a higher inflation rate can lead to a more rapid price adjustment (Ball and Mankiw, 1994) and, thus, to a lower degree of nominal rigidities. Some studies have shown that PPP tends to hold well for high inflation countries (McNown and Wallace, 1989) and that a higher level of inflation is associated with a lower level of real exchange rate persistence (Cheung and Lai, 2000). To control for the level of inflation, we have considered the variable $\text{INFL}_{t,i}^c$, defined in (23), that measures the inflation in sector i of country c.

Following previous studies, we also considered other control variables, such as government spending and the volatility of the exchange rate (see Cheung et al, 2001). However, these variables were not significant and did not seem to have any important impact on the coefficients of the remaining variables so, for the sake of brevity, we do not report the estimation output corresponding to models containing these variables.

Dependent variables

Our main dependent variable is the sectoral CIRs defined as

$$CIR_{c,i}(h) = \sum_{l=0}^{h} IRF_{c,i}(t,l), \text{ for } i = 1, ..., N.$$

To capture the explanatory power of the independent variables at different moments of the lifetime of the shocks, several values of h have been considered, namely, $h = \{12, 36, 60, 84\}$, in order to measure the short (h=12), medium ($h = 36$) and long-run ($h = \{60, 84\}$) effect of shocks. For completeness, sectoral HLs have also been considered as the dependent variable.

---

[19]We have also tried other approaches to measure the degree of openness of the final and the intermediate goods. We have considered transportation costs, proxied by the distance between national capitals; trade barriers have also been measured as in Anderson and Wincoop (2003) and Novy (2008). None of these variables turned out to be significant or had any significant impact on the coefficients of the other variables so we decided to drop them from the analysis.

5.2. Econometric methods

In order to examine the empirical relations between RER persistence and the various theories outlined above, we have carried out both a standard and a quantile panel regression analysis. Standard regression methods only provide a single summary measure for the conditional distribution of the dependent variable (the conditional mean) given the predictors. However, it has been recognized that the corresponding estimates are not necessarily indicative of the response of the dependent variable to the regressors on other parts of the conditional distribution. Since we are particularly interested in explaining the behavior of the most persistent sectors, the use of quantile regression techniques will provide us with a more complete picture of the covariate effects at the right tail of the distribution of RER persistence.

Following Koenker (2004), we consider the following model for the conditional quantile τ associated with the response of the corresponding persistence measure in sector i of country c :

$$Q_{y_{c,i}}\left(\tau|x_{c,i}\right) = \alpha_i + x'_{c,i}\beta\left(\tau\right); \; i = 1,...,N; c = 1,...,C, \tag{25}$$

where $(y_{c,i},\ x_{c,i})$ denote the values of the dependent and independent variables, respectively. The role of the $\alpha's$ is to control for individual unobserved heterogeneity that has not been adequately captured by other regressors in the model. However, α is assumed to be the same for all the quantiles and thus, is not allowed to depend on τ. This is because the number of observations in each sector is small and, therefore, it would be unrealistic to try to estimate a $\tau - dependent$ individual effect. Thus, α captures an individual specific location-shift effect. On the other hand, the effect of the regressors on $y_{c,i}$, $\beta\left(\tau\right)$ is allowed to depend on τ but it is assumed to be the same across sectors and countries. Since, in general, $x_{c,i}$ will contain an intercept, the estimated constant term will be the sum of two components, one depending on τ but not on i and the other depending only on i. See the Appendix for further details on the estimation estrategy.

With respect to the standard panel regression analysis, the following model has been considered

$$y_{c,i} = \theta_c + x'_{c,i}\beta + u_{c,i},$$

where the parameters have been estimated using the fixed-effects estimator, as the Hausmann test rejected the hypothesis of consistency of the random effects estimator.

5.3. Results

Two models have been estimated: M1 includes all the regressors described in Section 5.1 while M2 only contains those variables that turned out to be significant in the regressions performed in M1. We have estimated conditional quantile regressions for all the deciles although, for the sake of brevity, full regression results are only presented for the quantiles $\tau_\kappa = \{0.5, 0.7, 0.9\}$.[20] Similar models have also been estimated using standard panel techniques. Tables IIIa and IIIb report the corresponding figures while Figure 4 displays the evolution of the estimated coefficients over the different quantiles considered. The dashed line in these graphs represents the value of the standard panel coefficient estimate of the corresponding variable in a regression where the dependent variable is CIR(h), $h = \{36.60, 84\}$ months. The solid line depicts the quantile regression estimates for the quantiles $\{0.1, 0.2,...0.9\}$. Finally, Table IV reports the goodness of fit statistics defined in the previous section corresponding to M1.

The main results can be summarized as follows. The most important group of variables to account for RER persistence appear to be those related to the market structure of the inputs. In particular, the IIT variable associated with the intermediate goods (Input-IIT) has the expected positive sign and is always significant at the 5% significance level in all the models considered. Input-PCM also shows a positive relation with sectoral RER persistence that is generally significant, especially in the medium and long run and when higher quantiles are considered. As Figure 5 shows, the coefficients associated with these variables tend to be larger, the farther the horizon of the CIR considered and the higher the quantile.

Interestingly, once we have controlled for the market structure of the intermediate inputs, the market structure of the final good turns out not to be important in explaining RER persistence. Both IIT and PCM usually have the expected positive signs but they are not significant.

In line with the theoretical predictions, we find that a high degree of price stickiness is associated with a higher degree of persistence, as captured by the negative sign of the variable VOL, that is highly significant in all models.[21] Interestingly, the coefficients of the quantile regression parameters

---

[20]Results corresponding to the other quantiles are available upon request.

[21]A variable capturing the volatility of inflation of the intermediate inputs was also introduced in the regressions but it did not turn out to be significant.

30

turn out to be considerably larger in absolute value as farther horizons and higher quantiles are considered.

With respect to the variables that capture the degree of tradability, *openness* appears, in general, with the expected negative sign for the quantiles to the right of the median although its sign is positive to the left of this value. The variable that captures the degree of tradability of the inputs, *Input-OP*, usually presents a positive sign, especially in the upper quantiles, suggesting a positive relation between the degree of openness of the inputs and RER persistence. This result is not as surprising as it might look at first glance: internationally traded inputs are those exposed to PTM practices and, as seen before, this seems to be a key determinant of RER persistence. Nevertheless, neither *OP* nor *Input-OP* are significant in any of the models considered. The lack of significance of these variables confirms recent findings that suggest that traded and nontraded goods have similar characteristics (Engel, 1999 and Chari et al., 2002) and, therefore, this distinction is not key in accounting for persistence. Another aspect that may have a role in explaining the lack of significance of these variables is that CPI data, even at the very disaggregate level considered in this paper, does not allow us to completely disentangle traded and nontraded goods because the price of traded goods often involves nontraded components, such as marketing and distribution services. In addition, the European Union is an area where trade barriers are very low since tariffs have been eliminated and trade costs are relatively small and, therefore, one could expect that the traded/non traded categories are less important than in other geographical areas. However, it is noteworthy that, in spite of Europe's market integration, deviations from PPP have similar persistence levels to most other areas, as Table I illustrates.

Finally, inflation has a positive impact on RER persistence. This result suggests that, once price variability has been taken into account, a higher level of inflation is related to higher persistence levels. This result is not surprising since, in general, high inflation countries such as Spain or Italy are usually those presenting the highest persistence levels.

The results of the standard panel regression are very much in line with the discussion above: the variables associated with the market structure of the inputs and with the price stickiness of the final good are always significant and present the expected sign. However, the market structure of the final good, as well as its degree of tradability, do not appear to be important to account for the

persistence of shocks to sectoral RERs.

Although the values of the coefficients from the quantile regressions and standard panel regressions are sometimes quite different, the confidence bands of the latter usually include the estimated quantile coefficients (with some exceptions for Input-IIT, INFL and VOL). This is probably due to the fact that panel coefficients have been estimated using only a small number of countries and, therefore, these estimates present, in general, large standard deviations.

Finally, in spite of their simplicity, the models considered in this section are able to explain an important proportion of the variability of the different dependent variables: around 65% and 45% for the panel and quantile regressions, respectively (see Table A.2 in the Appendix).

Summarizing, our results are in agreement with theoretical models such as Chari et al. (2002) and Carvalho and Nechio (2008). PTM at the intermediate goods level and price stickiness seem to be the key determinants of RER persistence. The classical dichotomy that classifies goods into traded and nontraded appears not to account for European RER persistence.

TABLE IIIa

QUANTILE REGRESSION RESULTS (I)

	τ_κ	CIR, Short run		CIR, Medium Run		CIR, Long run				HL	
		CIR(12)		CIR(36)		CIR(60)		CIR(84)			
		M1	M2	M1	M2	M1	M2	M1	M2	M1	M2
	0.5	-0.02	-	-0.10	-	0.12	-	-0.29	-	0.49	-
OP	0.7	-0.00	-	-1.05	-	-1.71	-	-0.70	-	-0.78	-
	0.9	-0.11	-	-1.13	-	-1.96	-	-2.59	-	-1.74	-
	panel	0.00	-	-0.08	-	0.15	-	0.06	-	0.45	-
	0.5	0.23	-	0.92	-	-0.08	-	-0.48	-	2.11	-
Inputs-OP	0.7	0.20	-	1.64	-	1.14	-	0.12	-	3.51	-
	0.9	0.94	-	15.73	-	22.61	-	6.62	-	-5.38	-
	panel	1.66	-	5.57	-	7.13	-	9.15	-	-17.85	-

Notes: *, ** denote significance at the 10 and 5% level, respectively.

TABLE IIIb

QUANTILE REGRESSION RESULTS (II)

| | τ | CIR, Short run | | CIR, Medium Run | | CIR, Long run | | | | HL | |
| | | CIR(12) | | CIR(36) | | CIR(60) | | CIR(84) | | | |
		M1	M2	M1	M2	M1	M2	M1	M2	M1	M2
IIT	0.5	0.52	-	1.37	-	1.14	-	2.18	-	-2.83	-
	0.7	0.45	-	0.86	-	0.76	-	0.02	-	-4.99	-
	0.9	0.71	-	1.69	-	3.94	-	-0.71	-	-28.43	-
	panel	0.33	-	-0.02	-	0.52	-	1.26	-	-6.29	-
Input -IIT	0.5	1.43**	1.40**	8.89**	8.83**	15.04**	15.41**	16.39**	17.56**	25.41**	22.41**
	0.7	0.97**	1.33**	8.60**	9.90**	20.10**	24.35**	26.56**	26.07**	31.77**	31.88**
	0.9	1.73*	1.53**	10.27**	9.06**	26.22**	27.13**	53.52**	50.90**	72.01**	55.60**
	panel	1.83**	1.66**	10.54**	9.90**	18.85**	18.96**	24.00**	24.04**	39.13**	35.53**
PCM	0.5	1.49	-	5.81	-	2.09	-	-7.97	-	-30.66**	-
	0.7	0.88	-	2.92	-	-2.82	-	-5.09	-	-37.32*	-
	0.9	0.94	-	15.73	-	22.61	-	6.62	-	-5.38	-
	panel	1.66	-	5.57	-	7.13	-	9.15	-	-17.85	-
Input -PCM	0.5	3.81*	4.34	18.01*	19.83*	25.15	28.59	35.00	29.41	87.16**	52.08**
	0.7	4.52*	5.13**	24.76**	23.08**	51.54**	60.86**	33.64*	33.33*	127.23**	107.13**
	0.9	11.04**	10.02**	256.57	34.12*	49.90	61.24**	89.26*	77.24*	201.00	308.75**
	panel	3.96*	4.53**	18.65*	21.82**	45.21*	49.02**	66.38*	74.22**	108.72**	105.67**
VOL	0.5	-0.39**	-0.38**	-1.15**	-1.12**	-1.24**	-1.39**	-1.51**	-1.52**	-1.04**	-1.02**
	0.7	-0.37**	-0.34**	-1.16**	-1.14**	-1.70**	-1.63**	-2.02**	-2.15**	-1.41**	-1.39**
	0.9	-0.22**	-0.26**	-0.92*	-0.98**	-2.08**	-2.14**	-2.59**	-2.91**	-0.78	0.48
	panel	-0.45**	-0.44**	-1.29**	-1.27**	-1.94**	-1.94**	-2.80**	-2.80**	-2.60*	-2.72**
INFL	0.5	0.08	0.07	1.04**	0.94**	2.25**	2.18**	2.52**	2.55**	3.98**	3.95**
	0.7	0.11	0.09	0.86**	0.94**	1.96**	2.16**	2.36**	2.49**	3.25**	3.32**
	0.9	0.04	0.03	0.67**	0.80**	1.52**	1.80**	1.85**	1.69**	2.76**	2.89**
	panel	0.05	0.05	0.92**	0.99**	2.35**	2.47**	3.50**	3.67**	2.91**	3.33**

Note: *, ** denote significance at the 10 and 5% level, respectively.

(Figure 4 about here)

6. CONCLUSIONS

This paper analyzes sectoral RER persistence for a group of EU-15 countries with the aim of shedding further light on the factors that drive aggregate persistence. We make use of recent theoretical results that establish a link between the aggregate and sectoral impulse response functions which allows us to decompose aggregate persistence into the persistence of its different subcomponents (Mayoral, 2009). We explore the forces that shape aggregate persistence by analyzing the characteristics of the sectors in the upper tail of the distribution of persistence. It is shown that the distribution of sectoral persistence has a large variance and is highly skewed to the right. We show that, as a consequence of the high skewness, the slow reversion to parity of aggregate RERs is driven by a few highly persistent sectors. Interestingly, sectors in the durable category are the most persistent ones and they alone account for more than 40% of the cumulative effect of shocks in the long run. Furthermore, there is a strong link between the overrepresentation of durable goods in the upper quantiles of the distribution of sectoral persistence and the persistence of aggregate RERs as measured by the HL. The correlation between these quantities is greater than 70%, suggesting that understanding why the durable goods category is so persistent is key to explaining aggregate HLs. Further research is needed to explain this phenomenon.

Using trade and industry data, we test whether the traditional theories (nontradability, market structure and price stickiness) are able to account for the pattern of persistence observed in sectoral data. Our results suggest that persistence in the upper quantiles is explained by factors that have to do with the market structure in the intermediate goods market. Since the behavior of the upper quantiles determine, to a large extent, the persistence observed at the aggregate level, we conclude that pricing to market and price stickiness are two key factors in explaining the slow reversion to PPP in the EU. Obviously, these conclusions can be biased by the fact that the EU market is highly integrated and, thus, transaction costs and trade barriers are, in general, low. Further research is needed to clarify whether these conclusions can be extended to other economic areas.

REFERENCES

Anderson, J. and E. van Wincoop, 2003. "Gravity with gravitas: a solution to the border puzzle", *American Economic Review* 93, 170-192.

Ball, L. and N.G. Mankiw, 1994. "Asymmetric price adjustment and economic fluctuations", *Economic Journal* 104, 247-261.

Bergin, P.R. and R.C. Feenstra, 2001. "Pricing-to-market, staggered contracts, and real exchange rate persistence", *Journal of International Economics* 54, 333–359.

Berk, K.N., 1974. "Consistent autoregressive spectral estimates", *The Annals of Statistics* 2, 489-502.

Betts, C. and M.B. Devereux, 1996. "The exchange rate in a model of pricing-to-market", *European Economic Review* 40, 1007-1021.

Carvalho C. and F. Nechio, 2010. "Aggregation and the PPP puzzle in a sticky-price model", forthcoming in *TheAmerican Economic Review*.

Chang, R. and M. Devereux, 1998. "Dynamic collusion, pricing to market, and real exchange rates", Mimeo, Federal Reserve Bank of Atlanta and University of British Columbia.

Chari, V.V., P. J. Kehoe and E.R. McGrattan, 2002. "Can sticky price models generate volatile and persistent real exchange rates?", *Review of Economic Studies* 69, 533-563.

Cheung, Y.-W. and K.S. Lai, 2000. "On cross-country differences in the persistence of real exchange rates", *Journal of International Economics* 50, 375-397.

Cheung, Y.-W., M. Chinn and E. Fujii, 2001. "Market structure and the persistence of sectoral real exchange rates" *International Journal of Finance and Economics* 6, 95-114.

Crucini, M., Ch.I. Telmer and M. Zachariadis, 2005. "Understanding European real exchange rates", *American Economic Review* 95, 724-738.

Crucini, M.J. and M. Shintani, 2008. "Persistence in law of one price deviations: evidence from micro-data", *Journal of Monetary Economics* 55, 629-644.

Devereux, M.B. and Ch. Engel, 1998. "Fix vs floating exchange rates: how price setting affects the optimal choice of exchange rates regimes?", *NBER Working Paper* 6867.

Engel, Ch., 1999. "Accounting for U.S. real exchange rate changes", *Journal of Political Economy* 107, 507-538.

Faruqee, H., 1995. "Pricing to market and the real exchange rate", *IMF Working Papers* 95/12.

Frankel, J.A.and A.K. Rose, 1996. "A panel project on purchasing power parity", *Journal of International Economics* 40, 229-240.

Gadea, M.D. and L. Mayoral, 2009. "Aggregation is not the solution: the PPP puzzle strikes back", *Journal of Applied Econometrics* 24, 875-894.

Goldberg, P.K. and M. Knetter, 1997. "Goods prices and exchange rates: what have we learned?", *Journal of Economic Literature* 35, 1243-72.

Grubel, H. G., and P. Lloyd,. 1975. *Intra-industry trade: the theory and measurement of international trade in differentiated products.* London: MacMillan.

Hairault, J.O. and Th. Sopraseuth, 2003. "Exchange rate determination in a model of pricing–to–market and nontradeables", manuscript.

Im, K. S., M. H. Pesaran, and Y. Shin, 2003. "Testing for Unit Roots in Heterogeneous Panels", *Journal of Econometrics* 115, 53–74.

Imbs, J, A. Muntaz, M.O. Ravn and H. Rey, 2005a. "PPP strikes back: aggregation and the real exchange rate", *Quarterly Journal of Economics* 120, 1-43.

Imbs, J, A. Muntaz, M.O. Ravn and H. Rey, 2005b. "Aggregation bias does explain the PPP puzzle", Mimeo.

Knetter, M., 1993. "International comparisons of price-to-market behavior", *American Economic Review* 83, 473-86.

Kilian, L., 1998. "Small-sample confidence intervals for impulse response functions", *Review of Economics and Statistics* 80, 218-30.

Kilian, L. and T. Zha, 2002. "Quantifying the uncertainty about the half-life of deviations from PPP", *Journal of Applied Econometrics* 17, 107-125.

Koenker, R., 2004. "Quantile regression for longitudinal data", *Journal of Multivariate Analysis* 91, 74–89.

Koenker, R. and J. Machado, 1999. "Goodness of fit and related inference processes for quantile regression", *Journal of American Statistical Association* 94, 1296-1310.

Kursteiner, G.M., 2005. "Automatic inference for infinite order vector autoregression", *Econometric Theory* 21, 85-115.

Lamarche, C., 2006. "Robust penalized quantile regression estimation for panel data", Ph.D. Dissertation, University of Illinois at Urbana-Champaign.

Levin, A., C. F. Lin, and C. Chu, 2002. "Unit Root Tests in Panel Data: Asymptotic and Finite-Sample Properties", *Journal of Econometrics* 108, 1–24.

Lewbel A. 1994. "Aggregation and Simple Dynamics", *American Economic Review* 84, 905-918.

Lothian, J.R. and M.P. Taylor, 1996. "Real exchange rate behavior: the recent float from the perspective of the past two centuries", *Journal of Political Economy* 104, 488-509.

Mayoral, L., 2008. "Heterogeneous dynamics, aggregation and the persistence of economics shocks", *UFAE and IAE Working Papers* 08.

McNown, R. and M.S. Wallace, 1989. "National price levels, purchasing power parity, and cointegration: a tests of four high inflation economies", *Journal of International Money and Finance* 8, 533-545.

Murray, C. and D.H. Papell, 2005. "The purchasing power parity puzzle is worse than you think", *Empirical Economics* 30, 783-790.

Ng, S. and P. Perron, 1995. "Unit root tests in ARMA models with data dependent methods for selection of the truncation lag", *Journal of the American Statistical Association* 90, 268-281.

Novy, D., 2008. "Gravity redux: measuring international trade costs with panel data", University of Warwick.

Obstfeld, M. and K. Rogoff, 1995. "Exchange rates dynamics redux", *Journal of Political Economics* 105, 624-660.

Rogoff, K., 1996. "The purchasing power parity puzzle", *Journal of Economic Literature* 34, 647-68.

Salter, W., 1959. "Internal and external balance: the role of price and expenditure effects", *Economic Record* 35, 226-38.

Swan, T., 1960. "Economic control in a dependent economy", *Economic Record* 36, 51-66.

Stoker, T.M., 1984. "Completeness, distribution restrictions, and the form of aggregate functions", *Econometrica* 52, 887-907.

Taylor, M. and L. Sarno, 1998. "The behavior of real exchange rates during the post Bretton Woods period", *Journal of International Economics* 46, 281–312.

Yang, J., 1997. "Exchange rate pass-through into U.S. manufacturing industries", *Review of Economics and Statistics* 79, 95-104.

Wei, S.-H. and D. Parsley, 1996. "Convergence to the Law of the One Price without Trade Barriers or Currency Fluctuations", *Quarterly Journal of Economics* 111, 1211-1236.
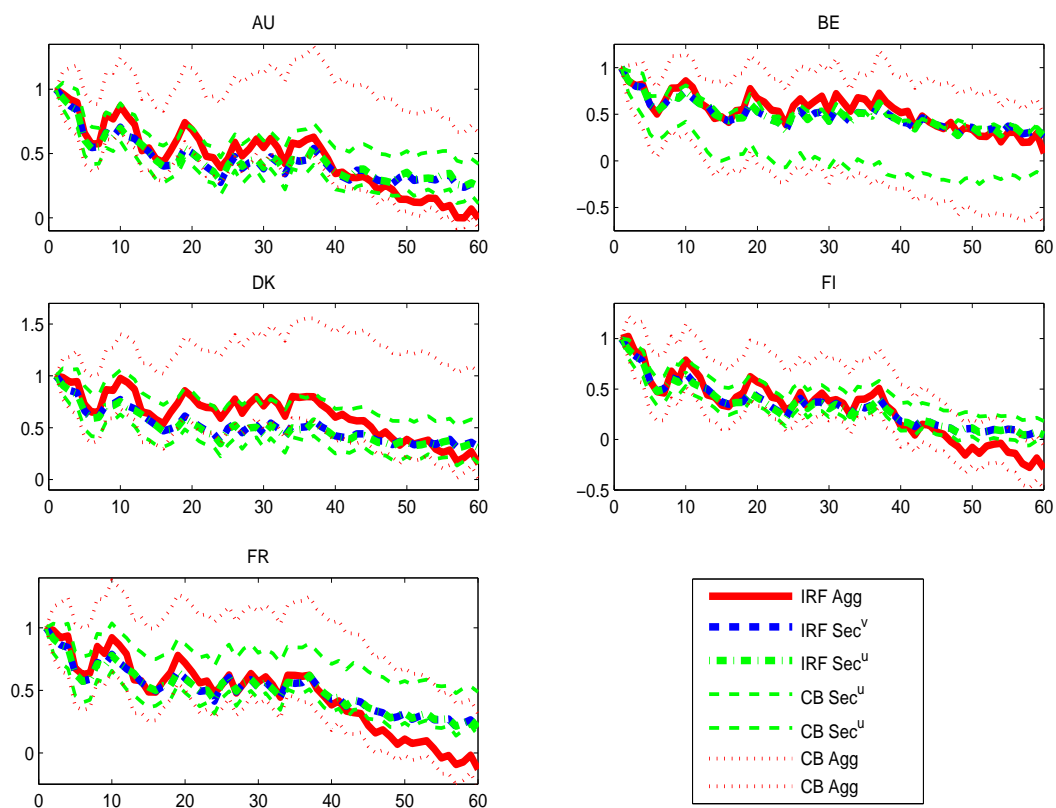
FIGURES

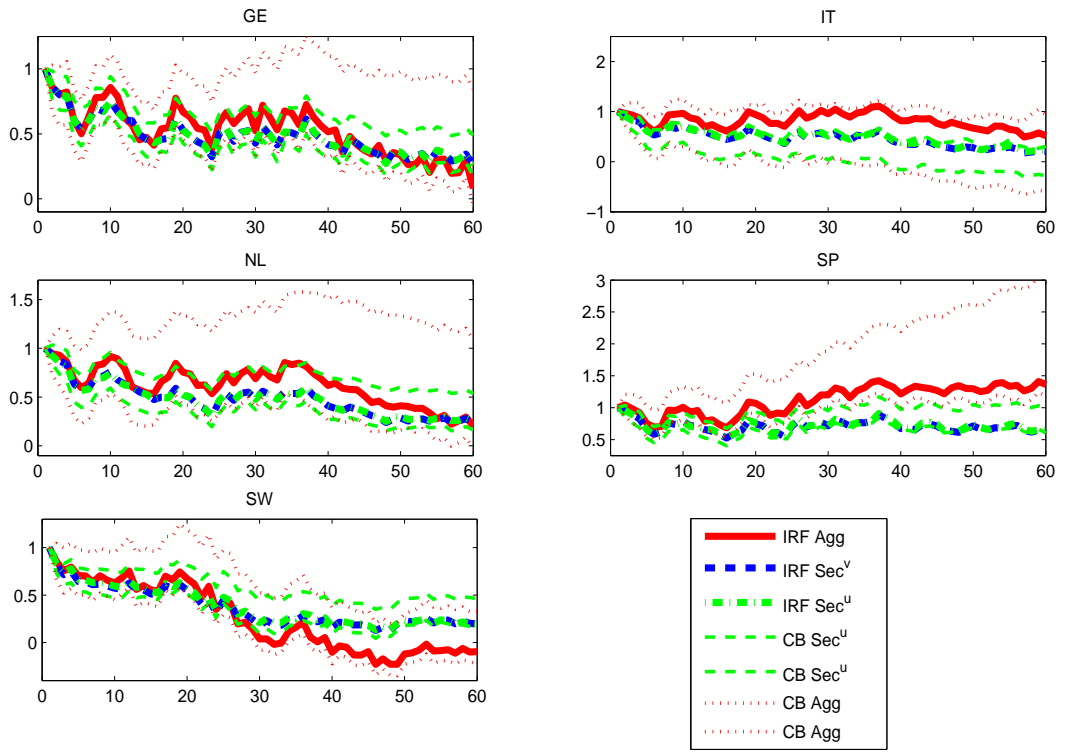FIG. 1A. Estimated IRFs using aggregate and sectoral data.

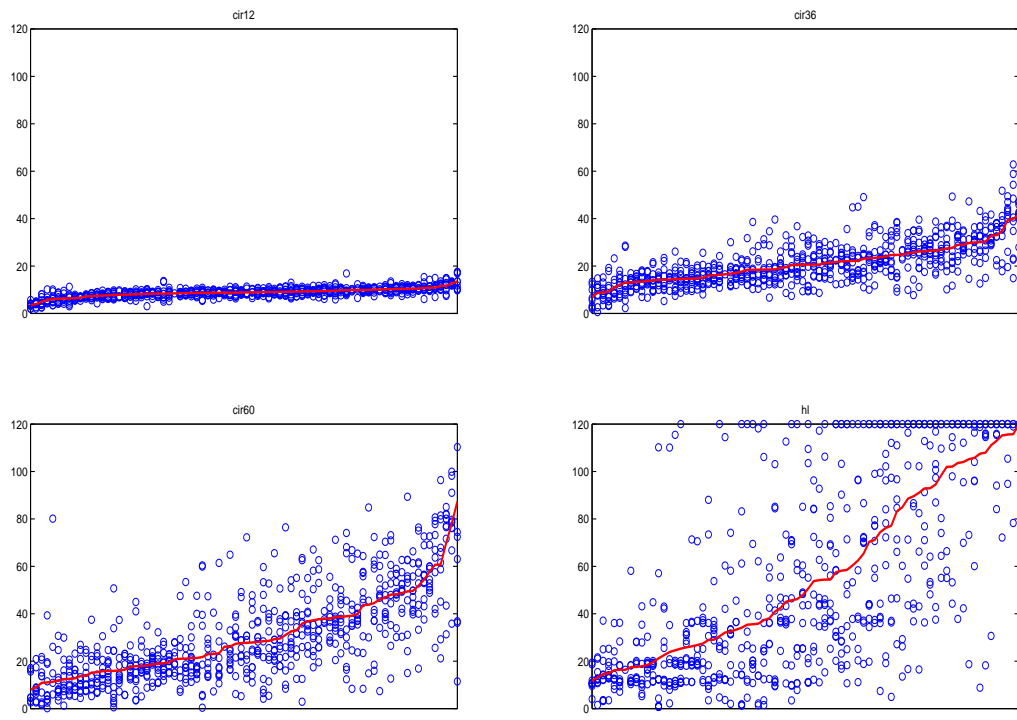FIG. 1B. Estimated IRFs using aggregate and sectoral data.
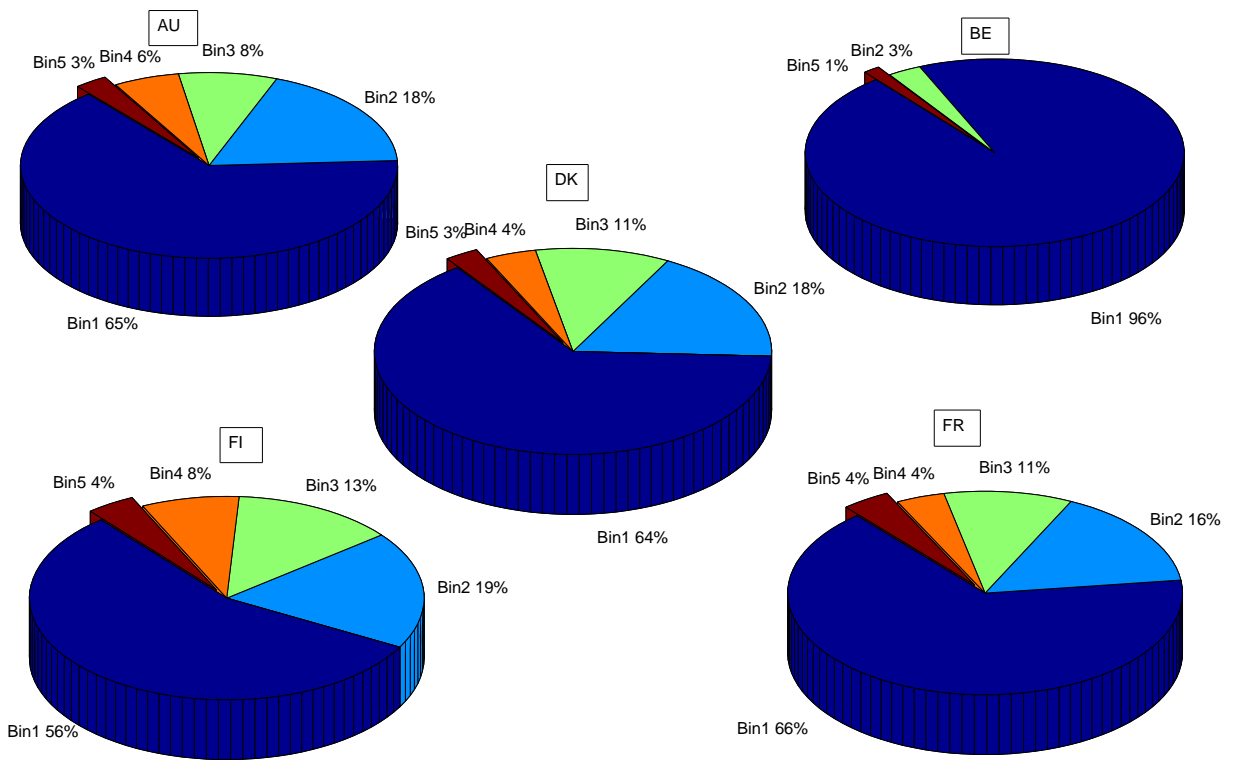
FIG 2. Heterogeneity in sectoral RER persistence.

FIGURE 3A. % of sectors in bins 1-5, ranked by contribution to the HL.
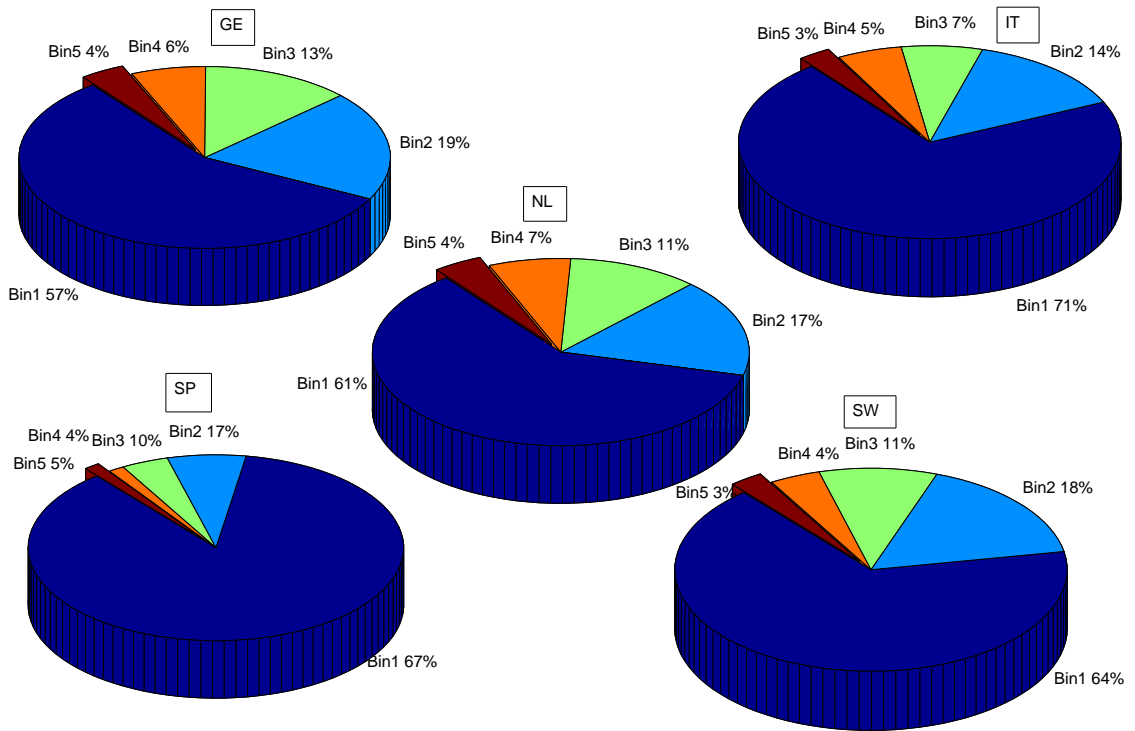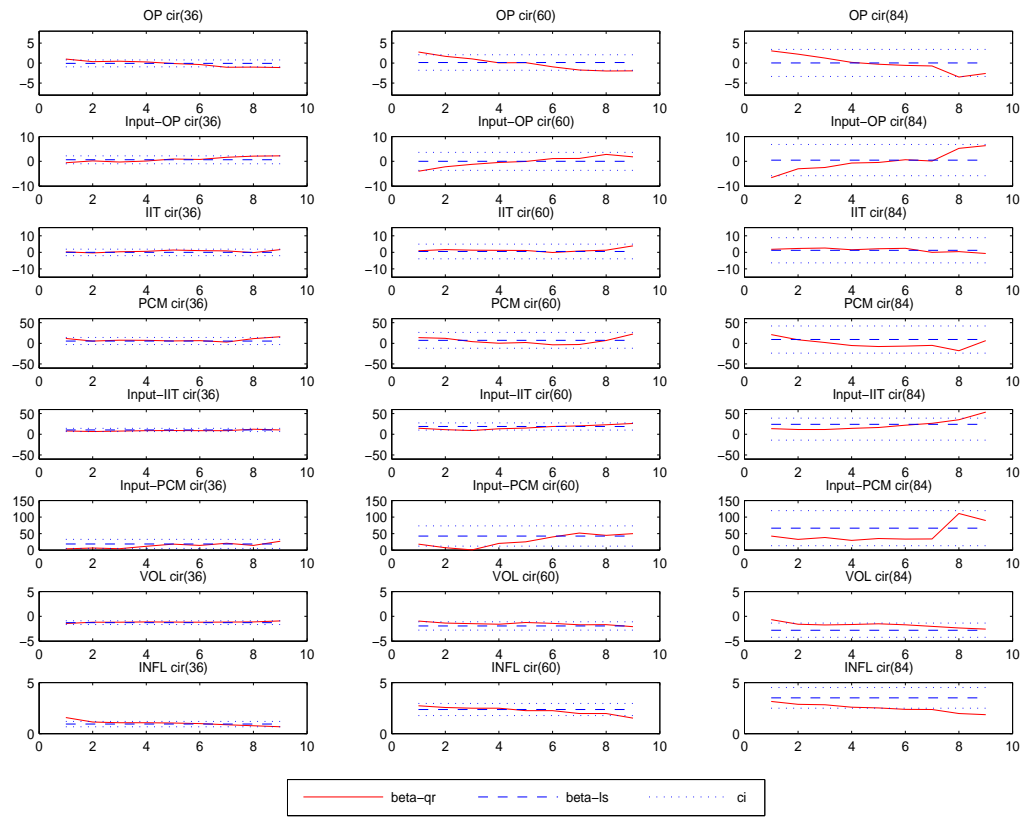
FIGURE 3B. % of sectors in bins 1-5, ranked by contribution to the HL.

43

FIGURE 4. Quantile coefficient estimates.