

The Impact of Monetary Union on Trade Prices*

By

Robert Anderton (European Central Bank), Richard E. Baldwin (Graduate Institute for International Studies) and Daria Taglioni (Graduate Institute for International Studies)

January 2003

Abstract

Two seemingly unconnected empirical results suggest an intriguing mechanism. First, economic integration helps harmonize prices internationally, with trade being the primary channel (Rogoff 1996, Goldberg and Knetter 1997). Second, monetary union may greatly increase the amount of trade among members (Rose 2001). Putting these together, we see that formation of a monetary union may induce changes that help harmonise inflation rates. The effect might be large if the elimination of exchange rate volatility simultaneously leads to a large increase in intra-union trade *and* a big increase in the speed at which price shocks are transmitted across members' goods markets. The problem is that standard estimates of price transmission speed suggest that trade's price-homogenising effect operates too slowly to matter much. Some new empirical evidence, however, suggests that a reduction in exchange rate variability reduces the variability of international price differences. Moreover, the effect seems to be highly nonlinear, and monetary union seems to have an effect even controlling for exchange rate volatility.

This paper is a first attempt to piece together part of this mechanism, namely the impact of monetary union (and exchange rate volatility more generally) on the international transmission of price shocks via the imported/exported inflation channel. In doing this we generate specific testable hypotheses and confront these with a number of data sets on European trade prices.

* In September 2002, this paper was presented at the ECB's Working Group on Econometric Modelling ("Second Pillar Seminar") and at the European Trade Study Group Annual Conference, Kiel, Germany. We therefore thank seminar participants for their very useful comments that have now been incorporated into the paper, particularly Dimitris Sideris (Bank of Greece) who was a discussant. The views expressed in this paper are the authors' and do not necessarily correspond to those of the European Central Bank.

Corresponding author: Robert Anderton; e-mail address robert.anderton@ecb.int

Word count: 8,400 words

1. Introduction

Two seemingly unconnected empirical results suggest an intriguing mechanism. First, economic integration helps harmonize prices internationally, with trade being the primary channel (Rogoff 1996, Goldberg and Knetter 1997). Second, monetary union may greatly increase the amount of trade among members (Rose 2001). Putting these together, we see that formation of a monetary union may induce changes that help harmonise inflation rates. The effect might be large if the elimination of exchange rate volatility simultaneously leads to a large increase in intra-union trade *and* a big increase in the speed at which price shocks are transmitted across members' goods markets.

The problem is that standard estimates of price transmission speed (Rogoff 1996, Goldberg and Knetter 1997) suggest that trade's price-homogenising effect operates too slowly to matter much – so slowly that many find the domestic-foreign price gap to be a random walk. Some new empirical evidence (Parsley and Wei 2001, and Asplund and Friberg 2001) suggests that a reduction in exchange rate variability reduces the variability of international price differences. Moreover, the effect seems to be highly nonlinear, and monetary union seems to have an effect even controlling for exchange rate volatility.

This paper is a first attempt to piece together part of this mechanism, namely the impact of monetary union (and exchange rate volatility more generally) on the international transmission of price shocks via the imported/exported inflation channel. In doing this we generate specific testable hypotheses and confront these with a number of data sets on European trade prices.

In our simple model, a change in exchange rate volatility – formation of a monetary union, for example – induces trading firms to change their behaviour in a way that tends to harmonize aggregate price movements between nations. Our model combines elements from new work by Friberg (2001) with elements drawn from the old sunk-cost hysteresis literature (Baldwin 1988, Baldwin and Krugman 1989, Dixit 1990), and especially its application to European monetary union (Baldwin 1991).

The paper begins by presenting the mechanism linking exchange rate volatility and trade price pass-through (Section 2). As part of this, Section 2 studies pricing behaviour under the two usual extreme assumptions of perfectly segmented markets and perfectly integrated markets. It then proceeds to endogenise the degree of segmentation and to study the implications of this for pass-through. Section 3 takes the model's main predictions to the data and empirically tests the hypotheses. The final section presents our concluding remarks.

2. Theoretical underpinnings

Firms engaged in international trade frequently set different prices for their goods in different markets. Actual and potential arbitrage, however, tends to limit the extent of such price gaps. If prices diverge sufficiently – and the goods sold in the different markets are sufficiently good substitutes – third parties will be tempted to arbitrage price gaps by re-exporting, re-importing goods, or by simply re-directing exports before they reach their intended destination. What all this goes to say is that unless firms explicitly engage in strategies to make price arbitrage expensive or illegal, their ability to price discrimination will naturally be limited.

The theoretical framework in this paper studies the optimal degree of market segmentation from the perspective of trading firms, by looking at the cost and benefits of 'buying' a given degree of market segmentation (defined as the maximum price gap that cannot be arbitrated). As it turns out, exchange rate volatility has an important impact on the degree of optimal market segmentation (Friberg 2001). The logic involves three steps.

First, discriminatory pricing is always profitable when demand elasticities vary across markets, but exchange rate variation is surely a far more important real-world motive for wanting to price discriminate. When the exchange rate gets a long way from its steady state value, firms will have a very large profit-incentive to ‘price to market’, i.e. to charge very different prices in different markets. Since this is true for both appreciations and depreciations, the value of being able to price discriminate is a U-shaped function of the level of the exchange rate, with the nadir of the U at the long-run exchange rate. The second step is to note that the expected value of this U-shaped function changes with the volatility of the exchange rate. In essence, a more volatile exchange rate means that one more frequently observes exchange rates where it would be very profitable to price discriminate. The final step is to consider the cost of market segmentation. Assuming that a higher degree of market segmentation involves corporate strategies that are progressively more expensive, profit maximising firms will choose a degree of segmentation (i.e. a maximum price gap) where the marginal benefit of further segmentation just equals its marginal cost. Because reduced exchange rate volatility lowers the marginal benefit of segmentation, firms will find it optimal to lower the degree of market segmentation in response to dampened exchange rate variability.

2.1. Basic Model

To illustrate the main economic logic of our analytic framework, we start with the simplest possible partial equilibrium model. We suppose there are only two nations and we consider a market where a monopolist (located in the ‘home’ country) sells in both markets. The market in each nation is a priori identical with demand in the two markets given by:

$$(1) \quad q = a - p; \quad q^* = a - p^*$$

where p and q represent the price and quantity sold respectively; variables without an “*” are home-nation variables, while those with an “*” are foreign nation variables. Without loss of generality, units are chosen so that the slopes of the identical demand curves are unity.

The firm’s production technology involves constant marginal costs such that:

$$(2) \quad c[w, q + q^*] = c(q + q^*); \quad c \equiv wa_q$$

where c is the constant marginal cost, which depends on w , the home nation wage, and a_q , the unit input labour requirement, both of which are taken as constant.

For simplicity’s sake, we ignore all forms of trade costs (natural and manmade), so the firm’s two-market objective function is:

$$(3) \quad \max_{p, p^*} pq + \frac{p^* q^*}{s} - c(q + q^*)$$

where ‘ s ’ is the exchange rate (s is a mnemonic for spot rate) defined as foreign currency units per unit of home currency. Note that a fall in ‘ s ’ raises the competitiveness of home-nation goods in the foreign market, i.e. is a depreciation of the home currency.

2.2. Polar segmentation assumptions

Trade theory typically views international markets as integrated – prices must be identical in all markets, i.e. third-degree price discrimination is not possible – or segmented, i.e. firms can set prices independently. To fix ideas, we first work out the firm’s problem assuming perfectly segmented markets, i.e. assuming no constraint on the choice of p and p^* .

2.2.1 Perfectly segmented markets

When markets are perfectly segmented, the firm is free to choose a price in each market. The solution to (3) thus involves two first order conditions, which solve to:

$$(4) \quad p = \frac{a+c}{2}; \quad p^* = \frac{a+cs}{2}$$

Since (3) is strictly concave in p and p^* , the first order conditions are necessary and sufficient, so (4) gives the profit maximisation prices with perfectly segmented markets. Notice that a sufficiently high ' s ' will yield a p^* that involves zero foreign sales. Specifically, the choke-off ' s ' equals a/c ; for spot rates higher than this, optimal foreign sales are zero. Using these optimal prices in the objective function, the profit function under the perfectly segmented markets hypothesis, which we denote as Π^{PSM} , is:

$$(5) \quad \Pi^{PSM} = \left(\frac{a-c}{2}\right)^2 + \frac{1}{s} \left(\frac{a-cs}{2}\right)^2; \quad 0 < s \leq \frac{a}{c}$$

for $s > a/c$, the profit function consists only of the first term in the expression since no foreign sales are made; 'PSM' is a mnemonic for perfectly segmented markets.

Expression (5) shows that the profit function is inherently asymmetric with respect to the exchange rate, s . As the spot rate rises, the home-based firm loses competitiveness in the foreign market and eventually stops exporting. (With linear demand this happens at a finite s , but with general demand, exports limit to zero as $s \rightarrow \infty$.) But as the spot rate declines, the profitability of exports sales increases without bound in terms of home country currency. This has nothing to do with functional forms because it is driven by the translation of foreign currency earnings into local currency units.

2.2.2 Perfectly integrated markets

The other extreme assumption is perfectly integrated markets, i.e. where firms are constrained to choose a single price for sales to the two markets – presumably for reasons having to do with potential resale among consumers. The firm's problem in this case is identical to that of the previous sections with the sole addition of the integrated market constraint, namely $p=p^*$. Solving (3) subject to $p=p^*$ yields optimal home and foreign prices (which differ only by s). Note that as ' s ' falls, the foreign market becomes increasingly attractive, so eventually the firm sets a price that implies zero home sales. The value of this choke-off s is $c/(2a-c)$. Plugging the optimal prices back into (3), the profit function under the perfectly integrated markets hypothesis, which we denote as Π^{PIM} , is:

$$(6) \quad \Pi^{PIM} = \frac{(2a-c(1+s))^2}{4(1+s)}; \quad \frac{c}{2a-c} \leq s$$

for s below the choke-off level, the firm sells only to foreigners and chooses the optimal price for foreign sales as given in (4), with PIM the mnemonic for perfectly integrated markets.




Figure 1: Difference between profits under PSM and PIM assumptions

2.2.3 Friberg's insight

We turn now to showing the insight, due to Friberg (2001), that the reward to being able to engage in segmented-market pricing versus integrated-market pricing depends in a large part on the exchange rate's probability distribution. This general result can be directly illustrated with our simple model. Simplifying, the difference between (5) and (6) becomes:

$$(7) \quad \Pi^{PSM} - \Pi^{PIM} = \frac{a^2(1-s)^2}{4s(1+s)}; \quad \frac{c}{2a-c} \leq s \leq \frac{a}{c}$$

While this is not everywhere convex – because of the asymmetric effects of appreciation and depreciation – it is convex in the neighbourhood of the long-run exchange rate $s=1$. The left panel of Figure 1 plots the two profit functions against s ; the higher curve is Π^{PSM} and the lower curve is Π^{PIM} . Clearly the wedge between the two, and thus the value of being able to price discriminate, increases as ' s ' deviates from its long-run value (which is here normalised to unity). The right panel plots the $\Pi^{PSM} - \Pi^{PIM}$ directly against s ; this clearly shows the intrinsic asymmetry involved in the PSM-vs-PIM wedge.




Figure 2: Volatility and the expected value of the PSM-vs-PIM wedge

The next step is to show that the expected value of the PSM-vs-PIM wedge is increasing in the volatility of the exchange rate. If the wedge were a strictly convex function, this result would be a straightforward application of Jensen's inequality. Given the asymmetry, however, the wedge will not in general be strictly convex; it is probably for this reason that Friberg (2001) is unable to prove a general

relationship between exchange rate volatility and the value of market segmentation. Nevertheless, he conjectures that such a relationship exists (see his Corollary 2 and discussion).

Using our simple functional forms and an additional simple assumption on the stochastic process driving the exchange rate, we can establish a link between volatility and the reward to segmentation by direct calculation. Assuming that the exchange is uniformly distributed on a support $[1-\sigma/2, 1+\sigma/2]$, the expected value of the PSM-vs-PIM wedge is:

$$(8) \quad \int_{1-\sigma/2}^{1+\sigma/2} (\Pi^{PSM} - \Pi^{PIM}) f(s) ds = \frac{a^2}{4\sigma} \left(\ln \left(\frac{(\sigma + 4)^4 (2 - \sigma)}{(4 - \sigma)^4 (2 + \sigma)} \right) - \sigma \right); \quad f(s) = 1/\sigma$$

where σ measures exchange rate volatility. Note that $\sigma \leq 2$, since $s \geq 0$. Inspection of (8) shows:

Result 1: The expected value of having perfectly segmented markets, as opposed to perfectly integrated markets, increases as exchange rate volatility increases.

A more complete characterisation of the link between volatility and expected value of the PSM-vs-PIM wedge is shown in Figure 2 (here $a=10$ and $c=1$). Notice that the relationship is highly nonlinear; the impact of volatility reduction on the value of having perfect segmentation gets weaker as the level of volatility falls. Next, consider an extension that allows us to endogenise market segmentation.

2.3. Variably segmented markets

The classification of markets into perfectly segmented or perfectly integrated is theoretically convenient, but not terribly realistic. More generally, one expects that firms can engage in a whole array of strategies that make cross-market re-sale increasingly difficult.

Market Segmentation Technology

The key aspect of segmentation is the inability of consumers to re-sell goods across markets since this prevents arbitrageurs from exploiting and eroding price differences. The European Union (EU) does not allow firms to forbid re-sale within the EU, so segmentation requires some sort of product differentiation, perhaps combined with government regulation. For example, a German food company that wishes to price discriminate between the German and French markets, but not between the German and Dutch markets, could produce one version of the good with a German-and-Dutch-language label and another with a French-language-only label. This labelling would make it expensive to arbitrage price differences, but not impossible unless this differentiation were combined with government regulations on the sale of food labelled in the ‘wrong’ language and a ban on third party re-labelling. Implementing this strategy may involve a fixed cost in adjusting the production line; an investment that may be worthwhile, the exchange rate is very volatile.

In particular, suppose the segmentation technology involves a one-time sunk cost – call it F – that does not depreciate, so the firm will sink the cost if the expected value of doing so exceeds the sunk cost. Analysis of this sort of investment decision under a variety of stochastic processes is well understood. To illustrate our main points as simply as possible, assume the exchange-rate process over time is independently and identically distributed (iid) over the support $1-\sigma/2$ to $1+\sigma/2$ (so σ is a measure of volatility). In this case, the firm undertakes the segmenting investment if and only if pF exceeds the one-period expected value of the PSM-vs-PIM wedge.

With this simple market segmentation technology, Figure 2 shows that reducing exchange rate volatility would make it less likely that firms find segmentation profitable. For example, if the sunk cost facing a particular firm is F , then our measure of volatility would have to fall below 0.5 to make segmentation unattractive to the firm. Volatility would have to fall below 0.25 to make it unprofitable when the fixed cost was F^* . Notice that the impact of reduced volatility on segmentation is nonlinear here, but the form of the nonlinearity indicates that there would be little difference between very low volatility

and a true monetary union. Since this seems to be counterfactual, given the empirical studies discussed above, we must go beyond the Friberg framework that relies on sunk costs.

Variably Segmented Markets Assumptions

To model firms' choice over variable segmentation, we suppose that firms can choose a strategy that allows them to charge cross-market price gap of up to "G" without fear of re-sale; to be concrete, we assume that G is fixed in foreign currency terms (G is a mnemonic for gap). More specifically, for a given G, the firm acts as if markets are perfectly segmented when 's' is such that they want to charge prices within the following range:

$$(9) \quad sp + G \leq p^*, \quad p^* \leq sp - G$$

The left-hand inequality is relevant when s is low, so the home price looks very attractive to foreign consumers when translated into foreign currency terms – attractive enough to make it worthwhile paying the arbitrage cost of G per unit bought. The right-hand inequality is relevant when s is high enough to make home residents consider re-importing the good from abroad despite the arbitrage cost.

We refer to this framework as "variably segmented markets" and note that it will allow us to endogenously determine the extent of market segmentation. The segmentation technology is such that the firm faces a rising cost of implementing strategies that raise G. To be concrete and to simplify calculations, we assume the segmentation cost function is:

$$(10) \quad C[G] = aG^2 / 2$$

This implies that the marginal cost of a higher G is positive and rising in G. The idea here is that it becomes increasingly costly for a firm to institute strategies that permit it to charge progressively wider gaps.

2.3.1 Modelling variable market segmentation

The assumption of variably segmented markets (VSM) yields a richer problem for the firm. To explore this we consider the expected value of having a particular level of G. The firm will face three types of exchange rates. If s is close enough to its long-run level (unity), then it can act as if the market were perfectly segmented, since its profit-maximising prices imply a cross-market price gap that is too small to arbitrage. In this case, its problem and solution is identical to the one in (5). If, however, s is sufficiently high, the solution to its profit maximisation problem will involve the constraint that $p^* = sp - G$, and if s is sufficiently low, the constraint will be $p^* = sp + G$. The critical values of s are easily determined, since when the inequalities in (9) hold with equality, the firm will be pricing according to (4). Using this fact, we find the critical values to be:

$$(11) \quad s_{hi} = \frac{1}{1 - 2G/a}, \quad s_{lo} = 1 - 2G/a$$

This implies that the width of the hysteresis band, i.e. the band where arbitrage provides zero pressure for price convergence, is:

$$(12) \quad \text{band width} = \frac{4G(a - G)}{(a - 2G)a}$$

notice that this hysteresis band widens with G at an increasing rate, i.e. the band is convex in G.

In the two regions of s where the arbitrage constraint binds, optimal prices are:

$$(13) \quad p_{hi} = \frac{2(a + sG) + c(1 + s)}{2(1 + s)}, p_{lo} = \frac{2(a - G) + c(1 + s)}{2(1 + s)}$$

The optimal prices in the band are given by (4). Plugging these optimal prices into the definition of profits yields three different profit functions that we do not report since they are not particularly revealing.




Figure 3: VSM vs. PIM profit functions, and marginal benefit and cost of G

We can illustrate the impact of variably segmented markets (VSM) with the left panel of Figure 3. For a given G , the high and low bounds on s will be fixed and for realisations of the exchange rate within this range, the difference between VSM and PIM is identical to the difference between PSM and PIM as in the right panel of Figure 1. When s is higher than s_{hi} , the firm is able to do somewhat better under VSM than it would with perfectly integrated markets, but not as well as if it had perfectly segmented markets. This is shown with the bottom dashed line on the right side of the diagram. Similarly, when s is below s_{lo} , the firm does better than PIM but worse than PSM, as shown by the lower right-side dashed line.

If the firm 'buys' a higher level of segmentation, as measured by G' , then its profits will unambiguously be higher or equal to those with G . The diagram shows this with the upper dashed lines on the right and left sides of the diagram. Notice that the benefit of higher G shows up only for exchange rates that are sufficiently far from the long-run level.

2.4. Endogenous, variable market segmentation

Since the exchange rate is a random variable, optimising firms will choose the optimal G by considering the expected value of having a given level of market segmentation. The expected value is the sum of three parts. For s between $1 - \sigma/2$ (the minimum realisation of s) and s_{lo} , the relevant profit function will be profits maximised subject to the left-hand price constraint in (9). For s between s_{lo} and s_{hi} , the relevant profit function is (5), and for s between s_{hi} and $1 + \sigma/2$ (maximum realisation of s) the profit function is profits maximised subject to the right-hand price constraint in (9). Integrating over the three ranges and subtracting the result from the expected value of perfectly integrated markets (which is assumed to apply when $G=0$ is chosen), we get:

$$\begin{aligned}
EV = & \frac{G}{s} \ln\left(\frac{s^a (1+s)^G}{s(1+s)^{2s}}\right) \Big|_{1-\frac{\sigma}{2}}^{s_{lo}} \\
& + \frac{1}{4s} \left(s - \ln\left(\frac{(1+s)^4}{s}\right)\right) \Big|_{s_{lo}}^{s_{hi}} + \frac{G}{s} \left(s(a-G) + \ln\left(\frac{(1+s)^G}{(1+s)^{2s}}\right)\right) \Big|_{s_{hi}}^{1+\frac{\sigma}{2}}
\end{aligned}
\tag{14}$$

where s_{lo} and s_{hi} are defined in (11).

It is important to note that this calculation assumes that G is such that the ordering of integration limits is correct, namely that $1-\sigma/2 \geq s_{lo}$ and $1+\sigma/2 \geq s_{hi}$. Since s_{lo} and s_{hi} depend upon G , this puts restrictions on the permissible G . Intuitively, the point is that no firm would ever find it profitable to buy a G that was so high that no price constraint was ever strictly binding. The specific restriction is that optimal G 's will be less than $a\sigma/2(2+\sigma)$.

Volatility and segmentation in the MB MC diagram

To illustrate the optimal choice of G , and thus the optimal hysteresis bandwidth, we plot the marginal benefit of G – i.e. the partial of (14) with respect to G – and the marginal costs in the right panel of Figure 3. As per (10), MC is rising. The MB curve is more involved. For a given level of exchange rate volatility, an increase in G widens the hysteresis band. Figure 3 helps us to evaluate the impact of this on EV. As the band widens, some realisations of s will correspond to a high profit level (the solid PSM vs PIM line instead of the dashed line), so a higher G will clearly raise EV. However, each successive increase in G affects a progressively narrower range of s . As a consequence, the marginal benefit of raising G falls with the level of G , as shown in the right panel of Figure 3.

The optimal G corresponds to the intersection of the MC and MB curves. For sufficiently high exchange rate volatility, for example σ' , the firm will choose a positive G' . As the figure shows, lowering σ to, say σ'' , will result in a progressively lower optimal G , viz. G'' . Importantly, this model predicts segmentation only reaches zero when the level of is zero. The analysis can be summarised as follows:

Result 2: The degree of endogenous market segmentation – as measured by the width of the hysteresis band – falls as the level of exchange rate volatility falls.

The testable implication of this result is clear. The estimated bandwidth should be narrower when using data for nations that have experienced relatively little bilateral exchange rate variation.

3. A Natural Experiment: The DM Bloc

The theoretical model suggests that pass-through should be linked to bilateral exchange rate volatility, but bringing this prediction to the data is problematic. The logic of our model suggests that pass-through is the result of conscious investments by firms. Since some of these investments are presumably long lasting – for example, the decision to produce differentiated products for export markets, or lobbying for regulatory segmentation – it is somewhat difficult to know how to measure volatility. Firms make the investment today in expectation of tomorrow's volatility, so what we need is a measure of how volatile firms expected the exchange rate to be in the next few years. While there are a variety of ways of dealing with this econometrically, we choose instead to rely on a natural experiment.

With the breakdown of the Bretton Woods exchange rate system in the 1970s, Europe adopted a string of managed exchange rate systems. European bilateral exchange rates nevertheless fluctuated widely with one important exception. Right from the beginning of post-Bretton Woods period, a group of countries – the deutschemark (DM) bloc – maintained quite tight bilateral rates. For example, the DM-guilder bilateral rate was not realigned between 1983 and the start of the monetary union and the bilateral exchange rate fluctuated within a very narrow band during this period, even narrower than the maximum

allowed under the European Monetary System's Exchange Rate Mechanism (ERM). Similarly, the Danish crown, Belgium franc, and Austrian schilling have had very stable relationships with the DM since the late 1980s.




Figure 4: Bilateral exchange rate volatility, DM bloc and others, 1960-1994

Importantly, this arrangement was stable and widely expected to continue. As Figure 4 shows, the DM bloc exchange rates stayed close throughout the long sequence of exchange rate crises in the 1970s, 1980s and 1990s. Firms engaged in intra-DM bloc trading should, therefore, have had sufficient time and incentives to adjust their segmenting strategies to the lower degree of exchange rate volatility. What this means is that historical data should reveal a difference in trade-pricing behaviour between Germany and the DM bloc nations and between Germany and other nations.

3.1. Null hypothesis: naïve linear model

Our theoretical model generates a number of testable predictions, but the most obvious are: (a) pass-through behaviour should depend upon bilateral exchange rate volatility, and (b) there exists a band or threshold which creates nonlinear adjustment to long-run equilibrium (resulting from a greater tendency for arbitrage outside the band compared to inside the band). Thus as a first cut, we investigate whether we can reject the null hypothesis that pass-through behaviour is unrelated to exchange rate volatility.

3.1.1 The naïve empirical model

To test the null hypothesis we need an empirical model and here we adopt the “linear model” that has been used countless times in the purchasing power parity (PPP) and law of one price (LOOP) literature (see Rogoff 1996 for an overview). This model focuses on the gap between prices charged in different locations measured in the same currency (i.e., in this case, the gap between a country's export price and its domestic price). Taking ‘z’ as the price gap, the model is usually written as:

$$(15) \quad \Delta z_t = \lambda z_{t-1} + e_t$$

where e_t is an iid error term. Thus, we assume that z_t follows an AR(1) process. For convenience, the model is not specified in terms of the underlying AR(1) parameter ρ but rather the speed-of-convergence parameter, $\lambda \equiv \rho - 1$. That is, λ obeys $-1 \leq \lambda \leq 0$ and it is the convergence speed in the sense that price differentials are reduced each period by a fraction λ plus an error term. A λ equal to zero implies perfect persistence (random walk); a λ equal to minus one indicates zero persistence (iid behaviour).

The data

We test this simple AR(1) specification using a panel of German import price data from 17 industrialised trade partners covering 10 manufacturing (2-digit) sectors in the period 1975-1995 (see data appendix for details). Specifically, we denote p_t^1 and p_t^2 as the log price levels of a good in two locations at time t . In our dataset, p_t^1 is the German import price of goods imported from country 2, and p_t^2 is the producer price in country 2 (the exporting country) both measured in the exporter's currency, and the difference between these two prices is the price gap.

A non-parametric test based on Cuzick (1985) fully rejects the presence of a trend in the data. Since the data are in index form, even in a world of perfect arbitrage, eventual mismatches could produce a constant gap. In order to overcome this potential problem, we demean the data. Specifically, we generate the price gap, z_t , as the OLS residual from the regression of $\ln(p_t^1 - p_t^2)$ on a constant (however, qualitatively identical results are obtained using non-demeaned data series¹).

Testing for unit roots

Equation (15) is roughly equivalent to the usual specifications used for testing for a unit root, and standard unit root tests would be the best way of testing whether λ is equivalent or different to zero. Therefore, the obvious first step would be to also use equation (15) to establish whether the price gap data are stationary. However, a problem arises because our main hypothesis is that the price gap moves in a nonlinear fashion due to threshold effects. Several papers demonstrate that standard unit tests can give misleading results in the presence of nonlinear or threshold effects, hence unit root tests based on equation (15) would be unreliable given the threshold effects predicted by our theoretical model.² Although there are several papers that show how to test for a unit root in the presence of nonlinearities or threshold effects,³ they are all derived for use in time series analysis and can not be easily translated for use in panel estimation.⁴ Accordingly, we do not test whether the price gap panel-data are stationary as the results could be quite misleading.

Results

Following the logic of the DM bloc natural experiment, we distinguished between import suppliers belonging to the so called DM-block and other EU countries traditionally displaying highly volatile bilateral exchange rates vis-à-vis Germany. We grouped data together for imports from Austria, Belgium (and Luxemburg), Denmark, and the Netherlands in one sub-sample (the DM bloc), and data on imports from France, Greece, Ireland, Italy, Portugal, Spain and United Kingdom in another.

As Table 1 shows, the difference in coefficients for the two groups using the naïve linear model (equation 15) is clear.⁵ The price gaps for intra-DM bloc trade are much less persistent (i.e., λ nearer to minus one) than the gaps for German imports from other EU nations. Specifically, it takes about a year and a half for half of a price shock to be reabsorbed (half life of 1.52 periods) for intra-DM bloc trade, while it takes more than 20 months for the ones coming from the rest of the EU.⁶

Of course the fact that the two different datasets produce different estimates does not allow us to conclude that the difference is due to exchange rate volatility since, for example, the intra-DM bloc might involve a substantially different set of traded goods. We can say, however, that the fact that adjustment is

¹ See Table 3 in appendix 3.

² For example, Peel and Sarno (2001) demonstrate the very low power of standard univariate unit-root tests to reject a false null hypothesis of unit root behaviour when the true model is non-linearly mean reverting.

³ See, for example, Caner and Hansen (2001), Enders and Granger (1998), Shin and Lee (2001), etc.

⁴ For example, testing for unit roots with threshold effects in the time series literature requires the computation of new critical values, etc, which is clearly beyond the scope of this paper.

⁵ As mentioned earlier, we do not test the data for unit roots because we do not have the correct panel-data critical values. Nevertheless, the large magnitude of the t-statistics for λ shown in Table 1, suggest that the critical values would have to be extremely high if we were to reject the hypothesis of stationarity.

⁶ More formally, a log likelihood ratio test shows that the difference between the parameters is statistically significant at the 5% level of significance.

quicker in the less volatile exchange rate environment is certainly consistent with our basic notion that firms choose to invest less in segmenting markets where the bilateral exchange rate is less volatile.⁷

Table 1: Testing the naïve linear model

<i>Linear AR1 specification</i>		
	<u>Lambda</u>	<u>AR1 half-life</u>
All	-0.328 (0.006) ^{***}	1.74
DM Block	-0.367 (0.007) ^{***}	1.52
Other EU	-0.337 (0.017) ^{***}	1.69

Notes: Standard errors in parentheses. *** indicates statistically significant at the 1% level. Panel estimates on annual data for period 1975-1995 for a total dataset covering potentially 17 trade partners and 10 manufacturing sectors. However, in this case, "All" represents the DM bloc countries plus the other EU countries.

Additional supporting evidence can be found by splitting the whole sample of import data into two sets of observations, one displaying low volatility and the second high volatility. In this exercise, we also include data on imports from Canada, Finland, Japan, Sweden, Switzerland and United States as well as the DM bloc and other EU nations. The resulting split is not country-specific since some suppliers may be classed as low-volatility in some years and high volatility in other years. Running the simple naïve linear model on these two separate samples also produces significantly different coefficients, with price gaps in the more volatile dataset displaying greater persistence (see results on right hand side of Table 3 in appendix 3). Again, this is in line with our theoretical priors.

3.2. Testing for pass-through hysteresis bands

Confirmation that pass-through was quicker for trade flows marked by lower volatility provides a very rough check on our model, but the model makes much more precise predictions. We turn now to testing those predictions.

The bedrock of our theoretical model is the existence of a band where price gaps are much harder to arbitrage. The simple model we worked with went further by assuming that arbitrage is perfect and instantaneous outside the band but impossible inside the band. Roughly speaking this would correspond to a price gap that followed a random walk with reflecting barriers. Plainly, instantaneous price changes are not the norm due to all sorts of adjustment costs. This leads us to introduce a new element, namely slow adjustment of the price gap outside of the band. In particular, our hypothesis is that when price gaps are small enough, there is either very slow or no convergence, while price differentials larger than a certain threshold are arbitrated away and hence decay according to a stable AR process.

In essence, the theoretical model allows for nonlinear adjustment to long-run equilibrium. It predicts a sort of regime switching, where the applicable regime depends upon the size of the gap. In the statistical literature, this sort of adjustment process is called a threshold autoregressive (TAR) process. Recent research shows that multiple-regimes threshold models are more powerful explanations for modelling purchasing power parity (PPP) in the presence of international trade and arbitrage costs.

⁷ Although this paper focuses on the impact of exchange rate volatility on the persistence of deviations away from the law of one price, there are many other reasons for the empirical failure of the law of one price, for example: transportation costs, trade barriers, imperfect information about prices in different locations, differences in national preferences (for a comprehensive survey, see Goldberg and Knetter, 1997). Accordingly, we choose our comparison group as 'other EU' as these countries should have similar transport costs and trade barriers as the DM bloc – hence differences in λ are more likely to be due to differences in exchange rate volatility rather than other factors.

Before proceeding to a structural estimation of the TAR process, we provide evidence for one of the most fundamental predictions of our band-hypothesis, namely that the price-gap adjustment process should be nonlinear, especially when trade costs are high.

3.2.1 The Tsay test for nonlinearity

A customary way of testing for nonlinearities in TAR literature is a non-parametric test proposed by Tsay (1989) whose detailed procedure is reported in the Appendix. Our results for the Tsay test report all significant maximal χ^2 (for the DM-block at the 10 percent level of significance). We conclude that we can reject the hypothesis of linearity in the data.⁸ Furthermore, the hypothesis that data are nonlinear in the presence of higher trading costs is confirmed.

3.3. TAR estimation

As next step, we estimate our nonlinear data using the TAR methodology (see Appendix for details). While in general the adjustment process may display many regimes, evidence of a TAR model in a PPP or LOOP model suggests to initially adopt the parsimonious specification that assumes just one, symmetric band with one speed of convergence λ inside the band and another outside the band:

$$(16) \quad \Delta x_t = \begin{cases} \mathbf{I}^{out} (x_{t-1} + b_{lo}) + \mathbf{e}_t^{out}, & x_{t-1} < b_{lo} \\ \mathbf{I}^{in} x_{t-1} + \mathbf{e}_t^{in}, & b_{lo} \leq x_{t-1} \leq b_{hi} \\ \mathbf{I}^{out} (x_{t-1} - b_{hi}) + \mathbf{e}_t^{out}, & x_{t-1} > b_{hi} \end{cases}$$

where \mathbf{I}^{in} and \mathbf{I}^{out} are the speeds of convergence in and out of the band and the b 's define the band edges. To keep things simple we assume that the band is symmetric around a zero price gap, so we have only three parameters to estimate: the two different speeds of convergence \mathbf{I} 's and one band edge b .

Conceptually, estimation of the TAR is trivial if one knows the ' b ' a priori. After all, the TAR essentially says that the adjustment is driven by two separate models – one for gaps inside the band and one for gaps outside the band. If we knew the value for b , we could partition the data into two populations and run a simple AR(1) on the separate sample to obtain consistent estimates of the λ 's. We must instead estimate it. Following a procedure employed, among others, in Taylor (1997), we construct a whole array of provisional partitions of the data that we use to find the best TAR approximation to the true model. Namely, we define $L_{AR}(\mathbf{I}, \mathbf{s})$ as the estimated log likelihood function for the null AR(1) linear model, and define $L_{TAR}(\mathbf{I}^{in}, \mathbf{I}^{out}, \mathbf{s}^{in}, \mathbf{s}^{out}, b)$ as the estimated log likelihood function of the alternative TAR model for a given b . Estimation proceeds via a grid search on the value of b which maximises the log likelihood ratio $LLR = 2(L_{TAR} - L_{AR})$. Since the TAR model is locally linear, computationally, for any given b , the maximum likelihood estimation of the TAR corresponds to an OLS estimation on partitioned samples of case data wholly inside or wholly outside the threshold.

In keeping with the natural experiment logic, we estimate a TAR for German import trade prices vis-à-vis the DM bloc trade partners and a separate TAR for Germany's trade with other EU nations.

The results of this procedure are reproduced in Table 2. These results are consistent with the main predictions of our empirical model. In particular, we find that the band-width (b) is greater for trade marked by more volatile exchange rate movements, and that the price gap again exhibits greater persistence for German imports from 'other EU' in comparison to German imports from the DM bloc. Moreover, in line with our priors, we find that the adjustment process is faster outside the band (compared to inside) for the 'Other EU' sample: inside the band the half-life is 1.85 years while outside it is 1.68

⁸ Maximal F and χ^2 statistics are: (i) entire dataset $\chi^2(1) = 8.78$ with $\text{prob} > \chi^2 = 0.003$ (ii) DM Block $\chi^2(1) = 2.71$ with $\text{prob} > \chi^2 = 0.099$ and (iii) other EU countries $\chi^2(1) = 76.88$ with $\text{prob} > \chi^2 = 0.000$

years and this difference is statistically significant. For the DM bloc trade, however, the two adjustment processes are basically identical with the λ 's being -0.392 and -0.391 .⁹

Table 2: Summary results for TAR estimation (fixed effect panel)

	<i>b</i>	<i>I</i> ^{out}	half-life out of band	<i>I</i> ⁱⁿ	half-life in band
DM Block	0.460	-0.392 (0.011)***	1.394	-0.391 (0.014)***	1.399
Other EU	0.734	-0.338 (0.036)***	1.683	-0.313 (0.016)***	1.845

Note: Estimation of TAR(2,2,1) using a maximum likelihood grid search over *b*.

4. Concluding Remarks

This paper presents a model where firms choose the extent to which they can bilaterally ‘price to market’ by investing in mechanisms that increase the cost of cross-border price arbitrage. Examples of such mechanisms include differentiating products by market, or lobbying for regulator controls that make third-party arbitrage expensive. As Friberg (2001) showed, the degree of bilateral exchange rate volatility is an important determinant of how profitable such segmentation investments may be. Our model goes beyond this by allowing firms a continuous choice of progressively more expensive segmentation strategies that result in their being able to maintain correspondingly larger price gaps. A straightforward prediction of our model – and one that is empirically testable – is that the no-arbitrage bands should be wider for bilateral trading relationships where firms anticipate a higher level of exchange rate volatility.

To test this prediction we rely on a ‘natural experiment’ from Europe. While most intra-European bilateral exchange rates were volatile in the 1980s and 1990s, one group of countries – the DM bloc – consistently maintained very narrow margins of fluctuation. We therefore estimate separate threshold autoregressive (TAR) processes for intra-DM bloc trade and Germany’s trade with other EU nations. Our results, namely that the intra-DM bloc band is narrower, and that the degree of persistence of the price gap for intra-DM bloc trade is lower relative to Germany’s other trade, are consistent with the theoretical predictions. The implication is that monetary union could produce changes in corporate strategies that result in faster cross-border transmission of price movements via the import and export price channel which, in turn, would tend to homogenise price movements across the member countries of a monetary union.

References

- Arellano, M. and Bond, S. (1991), “Some tests of specification for panel data: Monte Carlo evidence and an application to employment equations”, *Review of Economic Studies*, Vol. 58, pp. 277-297.
- Asplund, M. and R. Friberg, 2001, “The law of one price in Scandinavian duty-free stores,” *American Economic Review*, 91, 4, pp 1072-1083.
- Baldwin, R., 1991, “On the Microeconomics of the European Monetary Union,” *European Economy*, Special Issue No. 1, “One market one money”.

⁹ Another potential criticism of our econometric methodology might be that our results are inconsistent and biased given that our T dimension is short, while our N dimension is significantly large (as pointed out by Nickell, 1981). One possible way of overcoming these problems is by using GMM estimation as suggested by Arellano and Bond (1991). However, we did not follow the GMM approach as this becomes computationally complex within a TAR estimation framework, particularly with respect to choosing the optimal set of instruments.

- Baldwin, R., 1988, "Hysteresis in Import Prices: The Beachhead Effect," **American Economic Review**, 78, 4, pp 773-785.
- Baldwin, R. and P. Krugman, "Persistent Trade Effects of Large Exchange Rate Shocks," with Paul Krugman, **Quarterly Journal of Economics**, pp 635-654, No. 419.
- Caner, M. and Hansen, B.E. (2001), "Threshold autoregressions with a unit root", **Econometrica**, vol. 69, pp. 1555-1596.
- Dixit, A. (1990), "Trade policy with imperfect information
- Enders, W. and Granger, C.W.J. (1998), "Unit root tests and asymmetric adjustment with an example using the term structure of interest rates", **Journal of Business and Economic Statistics**, Vol. 16, pp. 304-311.
- Friberg, Richard (2001), "Two monies, two markets? Variability and the option to segment" **Journal of International Economics**, 55, pp 317-327.
- Goldberg, P. and M. Knetter, 1997, "Goods prices and exchange rates: What have we learned?" **Journal of Economic Literature**, XXXV, pp 1243-1272.
- Nickell, S. (1981), "Biases in dynamic models with fixed effects", **Econometrica**, No. 47, pp. 1249-66.
- Obstfeld M. and A.M. Taylor "Nonlinear Aspects of Goods-Market Arbitrage and Adjustment: Heckscher's Commodity Points Revisited", **Journal of the Japanese and International Economies** Vol. 11, No. 4, December 1997 pp. 441-479.
- Parsley D. C. and S.J. Wei, 2001. "Limiting Currency Volatility to Stimulate Goods Market Integration: A Price Based Approach," **NBER Working Papers** 8468.
- Rogoff, K., 1996, "The purchasing power parity puzzle," **Journal of Economic Literature**, XXXIV, pp 647-668.
- Rose, A. 2001 "One Money, One Market: Estimating the Effect of Common Currencies on Trade" **Economic Policy**.
- Shin, D.W. and Lee. O. (2001), "Tests for asymmetry in possibly non-stationary time series data", **Journal of Business and Economic Statistics**, Vol. 19, pp. 233-244.
- Taylor, A., 2001, "Potential Pitfalls for the PPP Puzzle? Sampling and Specification Biases in Mean-Reversion Tests of the LOOP" **Econometrica**, 69, pp. 473-498.
- Taylor, M., D. Peel and L.Sarno "Non-linear mean reversion in real exchange rates: Towards a solution to the purchasing power parity puzzles" , **International Economic Review**, 42, 4, pp 1015-1042.
- Tsay (1989) "Testing and modeling Threshold Autoregressive Processes", **Journal of American Statistical Association**, vol. 84, no.405.

Appendix 1: Data description

Our dataset consists of disaggregated annual producer and import price data covering the period 1975-1995. We calculated price gaps as the difference, in absolute values, between the natural logarithms of the German bilateral import price and the corresponding exporter's producer price indices (base year 1990). Import prices are for 52 (*4-digit ISIC rev.2*) German manufacturing sectors and producer prices from 17 industrialized trade partners (remaining EU countries, Canada, Japan and US) relating to 8 (*2-digit ISIC rev.2*) sectors. Import prices are taken from the OECD's International Trade by Commodities Statistics (ITCS) and have been transformed from SITC rev.3 to ISIC rev.2 in order to be compatible with the industrial classification of producer prices. Producer prices are from the OECD Indicators of Industry and Services database. All data are measured in national currency of the exporter. Exchange rate volatility is measured as the standard deviation of the first difference of the end of period monthly natural log of the nominal exchange rate (IFS) in the year preceding period t .

Appendix 2: Procedure for estimating and testing

The Tsay test for non-linearity:

Tsay (1989) proposes a very general specification test that compares a generic TAR alternative against a linear AR null. The general form of a TAR model, according to the terminology chosen by Tsay is: $X_t = a_0^i + a_1^i X_{t-1} + e_t^i$ if $r(i) > X_{t-1} > r(i-1)$. Where $r(i)$ are non trivial thresholds. This test works with *cases of*

data $(x_t - x_{t-1})$ with $t=1, \dots, T$. A case data is composed of a *pair of temporally consecutive changes in the price gaps*, e.g. $x_t - x_{t-1}$ is one case and $x_{t-1} - x_{t-2}$ is another. The case data are ordered according to the lagged value of the price gap, that decides between which thresholds $r(i)$ each case is located. This process is repeated twice, first in ascending and then in descending order. The reason why the test should be run with both increasing and decreasing ordering of the arranged case data is of practical nature and lies in the fact that, especially in small samples, the case data may not fall in all of the regimes delineated by every threshold value. In this paper, we adopt that practice and we report only the most significant p-values for each of the two (descending and ascending) F-tests. The test goes as follows: we run an AR(1) on the ordered data and generate a series of recursive residuals e_{ts} on all the case data. We then regress the predictive recursive residuals on the dependent variable (the contemporaneous price gap) by OLS, i.e. $e_{ts} = a_0 + a_1 x_t + u_{ts}$. If these residuals are orthogonal to the dependent variable, then we can reject non-linearity in the adjustment process. The orthogonality test amounts to calculating the conventional F statistic for this regression, which for large N approximates a χ^2 random distribution. The intuition of the test is that if there is regime switching, only the first 'n' cases (falling under the first regime) will show orthogonality properties, while the remaining (N-n) cases will follow a different behaviour revealing the point of switching which should be related to the size of the price gap.

TAR estimation:

This appendix describes the TAR estimation method used in the paper. The equation takes the form:

$$(17) \quad \Delta x_t = \begin{cases} \mathbf{I}^{out}(x_{t-1} + b_{lo}) + \mathbf{e}_t^{out}, & x_{t-1} < b_{lo} \\ \mathbf{I}^{in} x_{t-1} + \mathbf{e}_t^{in}, & b_{lo} \leq x_{t-1} \leq b_{hi} \\ \mathbf{I}^{out}(x_{t-1} - b_{hi}) + \mathbf{e}_t^{out}, & x_{t-1} > b_{hi} \end{cases}$$

where x_t the price gap at time t and x_{t-1} its lagged value, λ^{in} and λ^{out} are the speeds of convergence in and out of the band while b_{hi} and b_{lo} define respectively the upper and lower band edge. To keep things simple we assume that the band is symmetric around a zero price gap, so we have only three parameters to estimate: the two λ 's and the absolute value of b . The aim is to test the likelihood function of equation (16) against the AR(1) model $\Delta x_t = \lambda x_{t-1} + e_t$. This can be done by maximising the likelihood ratio $LLR = 2(L_{TAR} - L_{AR1})$ where L_{TAR} and L_{AR1} are maximum likelihood of the TAR and AR(1) models, respectively.

Procedure

Not having the values for the band edges 'b' in hand, we must estimate them as a first step by means of a search algorithm. The technique that we used for calculating the TAR is the following. **Step 1:** order the absolute values of price gaps by size; **step 2:** partition the ordered data into steps of the smallest possible width marked by candidate thresholds b_k . In our estimation we eliminate partitions with 12 or fewer observations in either part of the partition. The only limitation to our search is that, in a normal distribution, a commodity point is unlikely to be close to the 0th or 100th percentile. For these extreme points there are not sufficient observations to provide an efficient estimate, therefore we select an interval stretching from the 10th to the 90th percentile within which to look for the candidate thresholds. **Step 3:** for each candidate threshold b_k partition the sample in observation wholly inside and outside the band and calculate the maximum likelihood TAR. Since the TAR model is locally linear, computationally, for any given b , the maximum likelihood estimation of the TAR corresponds to an OLS estimation on partitioned samples of case data wholly inside or wholly outside the threshold. **Step 4:** After calculating the maximum likelihood estimation of the AR(1) model, compute the likelihood ratio $LLR = 2(L_{TAR} - L_{AR1})$. Specifically, we define $L_{AR}(\mathbf{I}, \mathbf{s})$ as the estimated log likelihood function for the null AR(1) linear model, and define $L_{TAR}(\mathbf{I}^{in}, \mathbf{I}^{out}, \mathbf{s}^{in}, \mathbf{s}^{out}, b)$ as the estimated log likelihood function of the alternative TAR model for a given b . **Step 5:** estimation ends with a grid search on the value of the band edge b that maximises the log likelihood ratio $LLR = 2(L_{TAR} - L_{AR})$.

Appendix 3: Results using non-demeaned data

Table 3: Summary results for TAR on country groups and data partitioned according to degree of volatility (panel, fixed effects, non-demeaned data)

<u>ALL</u>		C	l	Half-life	<u>ALL</u>		C	l	Half-life
	AR1 λ	0.46 (.009)	-0.33 (.006)	1.74		AR1 λ	0.46 (.009)	-0.33 (.006)	1.74
	λ (out)	0.79 (.017)	-0.40 (.008)	1.36		λ (out)	0.79 (.017)	-0.40 (.008)	1.36
	λ (in)	-0.19 (.014)	-0.39 (.019)	1.41		λ (in)	-0.19 (.014)	-0.39 (.019)	1.41
<u>DM</u>		C	l	Half-life	<u>Low-Vol.</u>		C	l	Half-life
	AR1 λ	0.46 (.014)	-0.37 (.011)	1.52		AR1 λ	0.53 (.013)	-0.39 (.009)	1.40
	λ (out)	0.82 (.040)	-0.36 (.017)	1.56		λ (out)	0.72 (.023)	-0.38 (.011)	1.46
	λ (in)	-0.11 (.010)	-0.36 (.019)	1.56		λ (in)	-0.20 (.024)	-0.33 (.027)	1.70
<u>Other EU</u>		C	l	Half-life	<u>Hi-Vol.</u>		C	l	Half-life
	AR1 λ	0.42 (.014)	-0.29 (.008)	2.02		AR1 λ	0.31 (.017)	-0.21 (.010)	2.90
	λ (out)	1.20 (.134)	-0.32 (.032)	1.81		λ (out)	1.01 (.038)	-0.47 (.017)	1.09
	λ (in)	0.30 (.012)	-0.31 (.010)	1.89		λ (in)	-.08 (.022)	-0.41 (.04)	1.31

Notes: Standard errors in parentheses. C represents the constant. “All” represents Canada, Finland, Japan, Sweden, Switzerland and United States as well as the DM bloc and other EU nations.

Table 3 above shows the results using fixed effects, but: (a) uses non-demeaned data; (b) shows the naïve estimate of λ (i.e., AR1 λ) as well as estimates of λ inside the band [i.e., λ (in)] and outside the band [i.e., λ (out)]; and (c) also experiments with different groupings of the data. For example, the left-hand-side of the table shows the results for the same two sub-groups as Table 1 in the main text (i.e., ‘DM bloc’ and ‘other EU’), while the right-hand-side splits the sample into two sets of observations, one displaying low-volatility and the other high-volatility (the resulting split is not country specific since the same import supplier may be classed as low-volatility in some years and high-volatility in other years).

Starting with the left-hand-side of the table and reading down vertically, we see that the non-demeaned data give qualitatively the same results as Table 1 in the main text. That is, for the naïve AR1 specifications (i.e., AR1 λ), we see that the DM bloc has the largest estimated λ - the fastest speed of response – in line with our theoretical predictions (i.e., an estimated λ of -0.37 for the DM bloc, compared with -0.33 for all groups and -0.29 for the other EU countries). However, unlike our earlier results using de-meaned data, using the TAR technique the estimated λ inside and outside of the bands are not statistically different for any of the groups. Meanwhile, turning to the right-hand-side of Table 3, we see that the naïve AR1 specifications (i.e., AR1 λ) show that the low-volatility observations exhibit a higher estimated λ compared to the high volatility observations (i.e., -0.39 compared to -0.21). Again, this is in line with our theoretical priors that the persistence of price differentials is greater for high volatility exchange rates relative to less-volatile bilateral exchange rates. Moreover, the results for both the high and low-volatility groupings show that λ is larger outside of the band in comparison to inside the band, which is again in line with our theoretical priors.