Pricing-to-Market: Evidence From Plant-Level Prices*

Doireann Fitzgerald† Stefanie Haller‡

September 2009

Abstract

We document pricing-to-market by plants which sell simultaneously in two markets that are segmented by variable exchange rates. Because we can match price quotes for the same product sold by the same plant across the two markets, we can cleanly identify relative markup responses to exchange rate movements by using fixed effects to control for unobserved marginal cost changes. For prices invoiced in destination currency, we find that, conditional on prices changing in both markets, relative markups move one-for-one with exchange rate changes. Consistent with this, the probability of prices changing in one market but not the other does not depend on the behavior of the exchange rate. This suggests that plants engage in an extreme form of pricing-to-market, where relative markups inherit the random walk behavior of exchange rates.

JEL Classification: F31, F41, L11, L16

*This work makes use of data from the Central Statistics Office, Ireland, which is CSO copyright. The possibility for controlled access to confidential micro data sets on the premises of the CSO is provided for in the Statistics Act 1993. The use of CSO data in this work does not imply the endorsement of the CSO in relation to the interpretation or analysis of the data. This work uses research datasets which may not exactly reproduce statistical aggregates published by the CSO. We thank the staff of the CSO for making this project possible. Stefanie Haller is grateful for financial support for this research from the Irish Research Council for the Humanities and Social Sciences. Doireann Fitzgerald is grateful for financial support from the NSF under grant number 0647850. We are grateful to Manuel Amador, Ariel Burstein, Nir Jaimovich, Pete Klenow and to seminar participants at the ESRI, CEAD-Budapest, Stanford and the NBER Summer Institute for helpful comments and suggestions. All remaining errors and omissions are our responsibility.

†Department of Economics, Stanford University, Stanford, CA 94305, dfitzger at stanford.edu

‡Economic and Social Research Institute, Sir John Rogerson's Quay, Dublin 2, Ireland, stefanie.haller at esri.ie

1 Introduction

Pricing-to-market refers to the situation where producers choose different and time-varying markups across markets which are segmented by variable exchange rates. Such behavior potentially plays an important role in explaining the relationship between nominal exchange rate fluctuations and the real economy. To date, evidence on pricing-to-market is strongly suggestive, but suffers from some important limitations, relying as it does on a specific sector, on data aggregated across producers or over time, or on strong modeling assumptions. The goal of this paper is to exploit a unique set of price data to provide clean evidence on this behavior across a broad range of sectors, using only a bare minimum of assumptions about the nature of the environment in which the producers operate.

We match data from the annual plant census for Ireland with the monthly micro-data on producer prices used to produce the Irish producer price index. We focus on plants reporting prices for sales of the same product to both Ireland and the UK, two markets that are segmented by variable exchange rates. For these plants, we restrict attention to prices invoiced in destination currency, the predominant case, as the currency of invoicing is what allows us to identify the destination market for exports at the level of individual price quotes. We address two questions: First, how do producers who sell the same product in both markets adjust their markup in one market relative to the other in response to exchange rate changes, conditional on actually changing prices in both markets? Second, given that price adjustment is infrequent in our monthly price data, does the probability that they adjust prices in one market but not the other (i.e. take action to realign relative markups following drift due to exchange rate movements) depend on the size and sign of exchange rate changes?

Our empirical strategy exploits the structure of our data. We observe parallel monthly price series for Irish and UK sales for the same product produced and sold by the same Irish plant. Price changes are infrequent, but changes are often synchronized across markets within plant-product pairs. Our identifying assumption is that over a given time interval, the percentage change in marginal cost is the same across the two markets for a given plant-product pair. Under this assumption, we can cleanly identify relative markup responses to exchange rate movements by focusing on consecutive pairs of price changes that are synchronized across markets, and controlling for unobserved marginal cost changes (and changes in markups that are common across markets) using fixed effects.

Our data has several advantages over different types of micro-data used in the previous literature to test for pricing-to-market. First, we have prices for the same product in different markets collected on the producer side rather than on the consumer side. Moreover, we know that the prices we observe are for products that are actually traded, not just potentially traded. Together, these two features allow us to eliminate the role of local non-traded content in driving relative price movements. Second, we know that the price quotes we observe are for products that, though sold in different markets, are produced in the same plant. This justifies the assumption that marginal cost changes are the same across markets over a given time interval, which is crucial to identifying relative markup variation. Third, our data is on prices rather than unit values, and at a sufficiently high frequency to allow us to look separately at the intensive margin (the size of price changes conditional on adjustment) and the extensive margin (the timing of price changes). This allows us to disentangle pricing-to-market and price stickiness. Fourth, our data covers a broad range of industrial sectors. This gives our results general applicability. Finally, we can link pricing behavior not just to product characteristics, but also to plant characteristics. This allows us to test a number of hypotheses about pricing-to-market that have not been tested before.

We find that, conditioning on prices actually changing, relative markups move one-for-one with exchange rate changes. Given that we focus on the case of destination currency invoicing, this implies that the co-movement of relative markups with exchange rates when conditioning on prices changing is the same as the default co-movement when prices do not change. If relative prices do not respond to exchange rate movements conditional on price adjustment, one would not expect the probability of price adjustment to depend on exchange rates. Consistent with this, we find that the probability that a producer changes the price in one market but not the other does not depend on the behavior of the exchange rate since the last (synchronized) price change. These two sets of results are based on overlapping but distinct subsets of the data. Together, they imply that when prices are invoiced in destination currency, relative markups across markets inherit the random walk behavior of exchange rates.

Our results are invariant to differences in product characteristics. It does not matter whether the product in question is a consumer good, a capital good or an intermediate. They are invariant to differences in plant characteristics. It does not matter whether the UK market is important or unimportant to the plant, whether imported intermediates account for a high or a low fraction of variable costs, whether price-cost margins at the plant level are large or small, whether the plant itself is large or small, or whether it is domestic or foreign-owned. Our findings are also invariant to differences in pricing behavior, such as the length of the time interval between price changes, and the median frequency of price changes at the plant-product pair level. In addition, we find no conclusive evidence in favor of delayed responses to exchange rate changes at the level of individual price quotes once marginal cost changes are controlled for.

Our findings are consistent with anecdotal evidence that, in the face of exchange rate fluctuations, exporters may favor maintaining market share over protecting margins. We present evidence that price-cost margins are substantial for most of the plants in our data, consistent with pricing behavior of the type we document being a sustainable strategy for exchange rate fluctuations on the order of those we observe in our sample period. We also note that many plants are to some extent hedged against exchange rate fluctuations, in the sense that they source a non-trivial fraction of intermediates in the UK as well as selling to the UK.

Methodologically, our approach to identifying pricing-to-market is closely related to that proposed by Knetter (1989). Knetter and others apply this using sector-level unit value data and find evidence of what they call local currency price stabilization, particularly in cases where prices are likely to be invoiced in destination currency. These results are summarized in Goldberg and Knetter's (1997) survey.[1] As these authors are aware, aggregation across sectors, across invoice currencies and over time makes it tricky to interpret their results as evidence of pricing-to-market at the plant level - composition effects and short-run price stickiness may be responsible. The results of this literature generally point to a degree of

[1]Gil-Pareja (2003) applies the same strategy to the European car industry with similar results. Using a related approach, Goldberg and Verboven (2005) also find evidence of pricing-to-market in the European car industry.

pricing-to-market less extreme than what we find. This could be due to differences in the level of aggregation, in the way in which the data is collected, or in the type of sectors covered.

Our strategy for identifying pricing-to-market contrasts with the structural approach of Goldberg and Hellerstein (2008) and Nakamura and Zerom (2008), who, like us, find evidence of pricing-to-market conditional on prices changing.[2] Relative to what they do, our methodology has the advantage of not requiring strong assumptions about preferences, market structure, cost functions etc., and we have the data to apply it across a broad range of sectors. The tradeoff is that the parameters that we estimate do not have a precise structural interpretation. Nor can we decompose incomplete pass-through into the shares attributable to local cost components, pricing-to-market and price stickiness, as they do, both because we do not have data on retail prices, and because of our reduced-form approach. We see our contribution as a complement to the structural literature, in that it can provide guidance as to the type of structure that is most likely to fit the data.

The most similar papers to ours in terms of the type of data used are Nakamura and Steinsson (2008) and Vermeulen et al. (2007), who make use of the micro data underlying the US and various European producer price indices respectively, and Gopinath and Rigobon (2008), Gopinath, Itskhoki and Rigobon (2008), Gopinath and Itskhoki (2009) and Nakamura and Steinsson (2009) who make use of the micro data underlying the US import and export price indices. Among these, the paper most closely related to ours is Gopinath, Itskhoki and Rigobon (2008). They find evidence of slow pass-through of exchange rate changes into US import prices invoiced in dollars, conditional on price adjustment. However because their data does not allow them to link prices in different destinations to the same individual plants, they cannot identify this as being due to pricing-to-market. A paper that makes use of different data to ours (wholesale prices at the UPC level faced by a supermarket chain operating in the US and Canada), but suggests similar behavior of producer prices is Burstein and Jaimovich (2009).

Our results provide some support for models of real exchange rate behavior that as-

[2] Goldberg and Verboven (2001) also find evidence of pricing-to-market in the context of a structural model of the European car market, but their data is not at sufficiently high frequency to control for price stickiness.

sume producers face residual demand curves with variable elasticity, as variable elasticity is necessary to generate pricing-to-market. Variable elasticities can be modeled in several different ways. In the literature on exchange rates, Bergin and Feenstra (2001), Burstein, Eichenbaum and Rebelo (2007) and Gust, Leduc and Sheets (2008), among others, model them as arising from the structure of preferences. Atkeson and Burstein (2008) and Burstein and Jaimovich (2009) generate variable elasticities from assumptions about the nature of market structure. A different approach is taken by Alessandria (2004) and Drozd and Nosal (2008), who examine real exchange rate behavior in models where search frictions make it costly for producers to build market share. To assess whether our findings are consistent with a particular model of variable elasticities or market structure, it would be necessary to calibrate a model nesting that structure with sticky prices, something that is beyond the scope of this paper. Of the literature on the subject, only the first set of papers - those that use preferences as a shortcut to model the nature of residual demand - nest sticky prices. These models suggest that relative markups should respond differently to current and lagged exchange rate changes, in a way that is not strongly supported in our data.

The next section of the paper describes our data set. The third section describes our empirical strategy. The fourth section presents our results. The final section discusses our results, and concludes.

2 Our data

Our data comes from two sources. The first source is the Irish Census of Industrial Production (CIP). This census of manufacturing, mining and utilities takes place annually at both the firm and plant level (over 90% of plants belong to firms with only one plant in Ireland). All plants with 3 or more employees are required to fill in a return. The data available to us covers the period 1991 to 2005.

Of the variables collected, those relevant for our purposes are the 4-digit industrial classification (NACE Rev. 1.1), country of ownership, value of sales, share of sales exported (with some destination and currency invoicing information), employment, wage bill, materials and energy expenditures (with share of materials imported, and some origin and invoicing information) and share of sales to related parties. Further details are provided in a data appendix.

The second source is the micro data collected for the purpose of constructing the Producer Price Index (PPI). The sampling frame for this data is the population of plants in the CIP. Plants selected to participate in the PPI do so on a long-term basis, though there is periodic resampling from the CIP to maintain coverage following attrition in the original sample and entry of new plants into the CIP. Participants report prices monthly, and we have access to the data from January 1995 through November 2006. The price data can be linked to the CIP plant data using a unique plant identifier. On average, 14% of CIP plants accounting for 38% of sales participate in the PPI sub-sample in any given year, while 89% of price observations can be matched to a plant in the CIP (this is probably due to the fact that wholesalers who do not produce also participate in the price survey).

Plants participating in the PPI survey are asked to provide transactions prices for their main products in both home and export markets. The precise request is for prices drawn from invoices dated on the 15th of the month in question. Participants are instructed to exclude delivery charges that are itemized separately on the invoice. In addition to a detailed description of the product, the first time they report a particular quote-line (to use the terminology of Klenow and Kryvtsov (2008)), participants are encouraged to provide information on terms of sale such as order size and type of customer (e.g. related party or non-related party), units for which the price is quoted, and destination market (for exports). They are asked to discontinue a quote-line and replace it with another if the product or terms of sale are no longer available or representative.

The variables in the PPI data that are relevant for our purposes are the price expressed in domestic currency, the currency in which the price was originally quoted, the exchange rate used to make the conversion from foreign currency, and whether the price refers to a domestic sale or an export sale. Unfortunately, the price-determining variables, including, crucially, units and the destination market for exports, are reported at the discretion of the respondent. Further, if provided, they are available only for price quotes present in the last cross-section (November 2006). As a result, we focus on price changes rather than price levels, and on domestic currency-invoiced prices in the home market and Sterling-invoiced prices in the export market. This latter condition identifies the destination as the UK with high probability.[3]

In the PPI, products are classified using system that is unique to this data. The level of aggregation at which the same product can be compared within and across plants is somewhere between a 4-digit and an 8-digit PRODCOM code (we have examined the product descriptions, and the level of aggregation varies across products). Our empirical strategy relies on the assumption that changes in marginal cost are the same across the home and export market within a plant-product pair as defined by this classification.

Summary statistics on plants

We now briefly present some relevant summary statistics. More information is available in the data appendix. Table 1 reports statistics on a wide range of plant characteristics, for all plants and for plants in the matched sample, for 1995 and 2005. These statistics illustrate the fact that Irish plants are very open on both the output and the input side. As such, they provide an ideal laboratory for examining the effects of exchange rate changes on pricing behavior.

Some facts are of particular note. In examining pricing-to-market behavior, we focus on plants simultaneously quoting prices for the same product in both Irish and UK markets, invoicing in destination currency. The UK market is an important one for Irish producers, both in terms of fraction of plants participating, and fraction of total sales by the industrial sector. For sales to the UK market, the CIP collects information on the choice of invoice currency. At least 70% of industrial exports to the UK are invoiced in Sterling. Hence, although our focus on the Irish and UK markets and on destination currency invoicing is one of necessity, it is not unduly restrictive.

While we do not observe marginal cost at the level of individual price quotes, and hence cannot calculate the level of markups, we can calculate price-cost margins at the plant level. The measure we use is total sales less total variable cost (wage bill plus materials and fuel

[3] For quote-lines that are present in the last cross section of the price data (November 2006) and where the destination of exports is identified, we can check the validity of our assumption that Sterling-invoiced price quotes are exported to the UK. For only 4% (unweighted) of Sterling-invoiced export price quotes in this sample is the reported destination clearly not the UK.

expenditures) divided by sales. These margins are quite substantial.[4] This fact will be of relevance in interpreting our results.

Summary statistics on prices

Table 2 provides summary statistics on the hierarchical structure of the matched price data. The average plant reports prices for 1.6 products, and for each product, reports an average of 4.4 price quotes. We do not go into detail here on the frequency and size of price adjustment (more details are available in the data appendix). Briefly, the behavior of producer prices in Ireland is broadly similar to that in six Euro-zone countries along the dimensions of weighted mean frequency and size of price adjustment, and is not too dissimilar from that in the US. Table 3 reports the mean frequency of price adjustment, for the sample as a whole, for home sales and exports separately, and, for exports by currency of denomination. As is clear from this table, prices are sticky in invoice currency, and the degree of stickiness is roughly comparable in the home market and for export sales invoiced in Sterling.

One feature of price-setting behavior that is important for our identification strategy is synchronization of price changes across markets within a plant-product pair. Summary statistics on this type of synchronization are reported in Table 4. The first column reports the fraction of the relevant population (i.e. plant-product-months with more than one price quote) where at least one price changes. The second column reports the fraction of these cases where there is exactly one price change. The third column reports the fraction where at least two, but not all prices change (this applies only to cases with three or more price quotes). The fourth column reports the fraction for which all prices change. This table illustrates that there is substantial, if not perfect synchronization of price changes across quotes within plant-product pairs. This is also true when we restrict the sample only to plant-product-pair-months with quotes both in home currency in the home market and in Sterling in the export market.

[4]They appear somewhat larger than those reported for US industry, e.g.in Domowitz, Hubbard and Petersen (1986).

Exchange rate variation

Figure 1 illustrates that our sample period covers a long period of depreciation of roughly 35%, followed by an appreciation of 20% of the domestic currency against Sterling.

3 Empirical strategy

From the perspective of a producer serving two markets that are segmented by variable exchange rates, changes in the nominal exchange rate act as a shifter of relative demand across markets. In addition, exchange rate movements may affect marginal costs. Our goal is to document how producers act to adjust the markup in one market relative to the other in response to these shocks, without imposing a lot of structure on the underlying environment.

To isolate the effect on relative markups, we rely on the assumption that the percentage change in marginal cost over a given time interval is the same for all markets served by a given plant-product pair.[5] This requires that the definition of a product be sufficiently narrow, and time-invariant. It also relies on prices being measured at the factory gate rather than inclusive of delivery charges, or on delivery charges being constant over the relevant time horizon.

We make several additional simplifying assumptions. We condition on continuing participation in both home and foreign markets, and on destination currency invoicing. When estimating markup responses conditional on price changes, we are agnostic about the nature of price stickiness, though we are required to take a stand when estimating the response of the probability of a price change to exchange rate changes.

In what follows, we refer to the exercise where we examine the response of relative markups to exchange rate changes, conditional on prices changing as the *intensive margin.* We refer to the exercise where we examine the relationship between the timing of price changes and exchange rate changes as the *extensive margin.* Before describing these, we introduce some terminology.

[5]Without additional assumptions, we cannot isolate the markup response to the relative demand shift as distinct from changes in marginal cost, as markups may respond differently to the same marginal cost shock in the two different markets.

Consider a plant-product pair indexed by i that sells in destination markets indexed by k, where $k \in \{IRL, UK\}$ The desired *reset* price[6] in market k, expressed in *home* currency, is equal to the marginal cost times the desired gross reset markup, which we refer to as the "desired markup:"

$$\hat{p}^i_{kt} = mc^i_t \hat{\mu}^i_{kt}$$

By reset price, we mean the price the producer would choose if forced to change its price at time t, assuming that the behavior of its competitors is held fixed, and given rational expectations about shocks and about its own and its competitors' future behavior. The desired reset price depends on the nature of price stickiness and the process for shocks, as well as the residual demand and marginal costs faced by the producer. Let $s^i_t(t)$ be the number of periods since the *destination* currency price for plant-product pair i in market k at time t was last changed. In general, $\Delta_{s^i_k(t)} \ln \hat{p}^i_{kt}$, the change in the desired home currency price between time t and the last time the destination currency price was changed, is a latent variable, and this is indicated by a carat.

Intensive margin

We can observe desired changes in *home* currency prices only when *destination* currency prices are reset at t. To examine the response of the desired markup to exchange rate changes, we condition on destination currency prices changing at t and estimate:

$$\Delta_{s^i_k(t)} \ln p^i_{kt} = \alpha + \theta^i_{t,s^i_k(t)} + \beta \Delta_{s^i_k(t)} \ln e_{kt} + \varepsilon^i_{kt,s^i_k(t)} \tag{1}$$

The plant-product-pair-month-age-of-price fixed effect $\theta^i_{t,s^i_k(t)}$ picks up changes in marginal cost between t and $t - s^i_k(t)$ as well as any changes in the desired markup that are the same across markets served by plant-product pair i. β is the elasticity of the desired relative markup across markets with respect to changes in nominal exchange rates. If producers desire constant relative markups, $\beta = 0$. If there is pricing-to-market $\beta \neq 0$. Since we condition on destination currency pricing, the default (i.e. when producers do not change

[6] This terminology is taken from Bils, Klenow and Malin (2009).

prices) is $\beta = 1$.

Given the structure of the fixed effects, information from plant-product pair i is used to identify β if price changes are synchronized across Irish and UK markets at t, conditional on the previous price changes in those two markets also having been synchronized. This is illustrated in Figure 2. As we document in Table 4, synchronization of price changes within plant-product pairs is relatively common. However, the requirement that two consecutive sets of price changes be synchronized does restrict the size of the sample we can use to estimate (1) and raises the possibility that our results may be affected by sample selection. We return to this issue later.

In estimating (1), we exclude observations where the log change in home currency price is greater than 2 in absolute value. We weight by turnover shares, constructed as described in the online appendix, and cluster standard errors at the plant level.

Extensive margin

To estimate the response of the probability of a price change to exchange rate changes, we must take a stand on the nature of price stickiness. We assume that producers follow an (S,s) rule, without going into more detail about what might motivate such behavior.

The destination currency price is what producers choose to reset. If the latent desired change in the *home* currency price between t and $t - s_k^i(t)$ takes the form (1), the latent desired change in *destination* currency price over the same time interval is given by

$$\Delta_{s_k^i(t)} \ln \hat{p}_{kt}^{i*} = \alpha + \theta_{t,s_k^i(t)}^i + (\beta - 1)\, \Delta_{s_k^i(t)} \ln e_{kt} + \varepsilon_{kt,s_k^i(t)}^i \tag{2}$$

where a star indicates that this is a destination currency price rather than a home currency price, and a carat indicates that the price is a latent variable, since in this case, we do not condition on observing a price change at t. If producers follow an (S,s) rule, large positive and negative values of $\Delta_{s_k^i(t)} \ln \hat{p}_{kt}^{i*}$ induce them to reset prices. In between, there is a band of inaction. Allowing for the possibility that the (S,s) band is not symmetric, we posit the rule (noting the distinction between desired and realized changes):

$$\begin{array}{lll} \ln p_{kt}^{i*} \uparrow & \text{if} & \Delta_{s_k^i(t)} \ln \hat{p}_{kt}^{i*} > \bar{\rho} > 0 \\ \Delta_{s_k^i(t)} \ln p_{kt}^{i*} = 0 & \text{if} & \underline{\rho} \leq \Delta_{s_k^i(t)} \ln \hat{p}_{kt}^{i*} \leq \bar{\rho} \\ \ln p_{kt}^{i*} \downarrow & \text{if} & \Delta_{s_k^i(t)} \ln \hat{p}_{kt}^{i*} < \underline{\rho} < 0 \end{array} \tag{3}$$

To operationalize this, we assume that $\varepsilon_{kt,s_k^i(t)}^i$ has a logistic distribution. Making use of (3) and (2), we can then write:

$$\Pr\left[\ln p_{kt}^{i*} \uparrow\right] = \Lambda\left(\alpha - \bar{\rho} + \theta_{t,s_k^i(t)}^i + (\beta - 1)\,\Delta_{s_k^i(t)} \ln e_{kt}\right) \tag{4}$$

$$\Pr\left[\ln p_{kt}^{i*} \downarrow\right] = \Lambda\left(-\alpha + \underline{\rho} - \theta_{t,s_k^i(t)}^i - (\beta - 1)\,\Delta_{s_k^i(t)} \ln e_{kt}\right) \tag{5}$$

where $\Lambda(z) = \exp(z)/1 + \exp(z)$.

We estimate (4) and (5) as two separate conditional logit models.[7] The conditioning procedure that eliminates the fixed effects ($\theta_{t,s_k^i(t)}^i$) implies that information from plant-product pair i is used to identify the coefficient on exchange rate changes only if the last price change was synchronized across the Irish and UK markets, while at date t there is a price increase (or decrease) in one market but not the other. This is illustrated in Figure 3. We note that the sub-sample used to identify the extensive margin and the sub-sample used to identify the intensive margin do not overlap perfectly.

We are principally interested in whether the probability of increasing or decreasing prices in one market but not the other depends on exchange rate changes (i.e., is $\beta - 1 \neq 0$), and if this is the case, what is the sign of the effect. Note that logit estimation does not identify the scale of $(\beta - 1)$ and $-(\beta - 1)$, but only whether the effect of changes in exchange rates on the probability that prices change is positive, negative, or zero.

As in the case of the intensive margin, in estimating (4) and (5), we weight by turnover shares and cluster standard errors at the plant level.

[7]For (4), we code increases in destination currency prices as a one, while all other observations (decreases and no change) are coded zero. The dependent variable in (5) is constructed analogously.

4 Results

Intensive margin

Tables 5 and 6 report the results from estimating the response of relative markups to exchange rate movements, conditioning on consecutive synchronized price changes (i.e. from estimating equation (1)). The first row of Table 5 gives the baseline estimates. The number of plants used to identify the coefficient is indicated by the number of clusters, while the number of plant-product-months is indicated by the number of fixed effects. The estimated β is positive and significantly different from zero, but not significantly different from one. Conditional on changing prices, the producers in our sample engage in pricing-to-market. Further, the estimated coefficient is almost exactly equal to 1, the elasticity of relative markups with respect to exchange rate changes when destination currency prices are *not* changed. This implies an extreme form of pricing-to-market.[8] We now illustrate the variation that identifies β, describe what happens when we condition on aspects of pricing behavior and product and plant characteristics, and address some measurement issues, dynamic adjustment and selection.

Figure 4 illustrates one dimension of the variation that underlies the baseline estimate of β. It plots a histogram of the changes in destination currency prices and exchange rates for the estimation sample. There is a lot of variation in the size of price changes. But although there are large cumulative swings in the exchange rate over the course of the sample period, the absolute size of exchange rate changes in the interval from one price change to another is on average a good deal smaller than that of price changes.[9] Figure 5 illustrates a different dimension of the variation in the estimation sample. It plots the log change in the Sterling price in the UK against the log change in the home currency price in Ireland,[10] along with the 45^o line. This illustrates that the estimated β of 1 is consistent with changes in destination currency prices being approximately equal in both markets. This figure also illustrates that

[8]Even if delivery charges are included in prices, and they co-move with exchange rates, this alone cannot explain our finding that β is not significantly different from 1 unless delivery charges are enormous relative to the f.o.b. price.

[9]This is consistent with the findings of Burstein and Jaimovich (2009).

[10]For cases where there is more than one price quote in a particular market, the mean log change across quotes is used.

the β coefficient is similar for positive and negative exchange rate changes, and for large and medium-sized price changes.[11]

Pricing, product and plant characteristics

Theory suggests that producers' incentives to realign relative markups when they change prices may be related to how long they expect to wait before changing prices again.[12] We do not find conclusive evidence of differences in pricing-to-market behavior along this dimension of price characteristics. The second panel of Table 5 reports the results from estimating equation (1), splitting the data by median frequency of price adjustment at the level of the plant-product pair. For all frequency groups, the estimated β is significantly different from zero and not significantly different from 1. This panel of the table also reports the results splitting the data by the gap between the first and second synchronized price changes. It seems plausible that realignment of markups to bring them back into line would be more likely when prices have not changed for a long time, but in all cases, the estimated β is significantly different from zero, and not significantly different from 1.

Theory also suggests that, conditional on market segmentation, the desired degree of pricing-to-market should vary systematically with the elasticity of the residual demand curve. The elasticity of residual demand may vary by type of product. Since the identifying sample is small to start with, we follow Vermeulen et al. (2007) in using a relatively aggregated classification of 4-digit sectors by end-use, and Rauch (1999) in using a tripartite classification of 4-digit sectors by degree of differentiation.[13] The third panel of Table 5 reports the results from estimating equation (1), splitting the data by type of product for these two classifications. For the Vermeulen classification, the point estimates of β are all in the neighborhood of 1 (with the exception of consumer non-food non-durables, where the coefficient is identified from relatively few observations), and in each case, significantly different from zero, but not significantly different from 1. There is no strong evidence that behavior varies across products at this level of aggregation. For the Rauch classification, surprisingly, the

[11] These results are reported in the data appendix.

[12] Gopinath and Itskhoki (2009) discuss in detail how exchange rate pass-through is related to the frequency of price adjustment.

[13] The details of these classifications are provided in the data appendix.

estimated β is always significantly different from zero and not significantly different from 1, even for the organized exchange (homogeneous) group. Having looked in detail at the products assigned to each category, we attribute this to the fact that the Rauch classification does not capture the intended product attributes in this data.

Theory also suggests that the incentive to engage in pricing-to-market may differ by measurable plant characteristics. For example, the elasticity of residual demand may be different for large and small plants. In the first panel of Table 6, we split the data by plant size (we exclude the largest plants for confidentiality reasons). The estimated β does not differ significantly across size classes, and is significantly different from 0, and not significantly different from 1 in all cases. It is also possible that pricing strategies may differ between plants for which the export market is a core market, and those for which it is not. In the second panel of Table 6, we split the data by quartiles of the share of sales exported to the UK. For the upper 3 quartiles we find no significant differences in the estimated β - they are significantly different from 0, not significantly different from 1. For the lowest quartile, the coefficient is less than 1, and significantly different from 1 and 0, but we note that it is identified from data from only 4 plants, and relatively few price observations. In panel 3 of the table, we split the data by quartiles of price-cost margins. The sustainability of pricing-to-market of the type we document should be related to this variable. We find no significant differences in the estimated β across these groups.

A producer's ability to hedge along different dimensions might also affect its pricing decisions. We have no information on hedging through financial instruments, but plants can hedge exposure to exchange rate risk through their sourcing of materials.[14] In the second panel of Table 6, we split the data by quartiles of the share of materials imported from the UK in variable cost (materials plus energy plus wage bill). Again, we find no significant differences in the estimated β across these groups. In the third panel, we also split the data by ownership, since the scope for hedging and incentives to price-to-market may differ across home and foreign-owned plants. However the pricing-to-market behavior of domestic, foreign and UK-owned plants is statistically indistinguishable.

[14] Sterling invoicing for imports appears roughly as prevalent as Sterling invoicing for exports, reinforcing the role of UK sourcing as a hedge.

One issue of concern is the fact that we do not know whether the prices we observe are for transactions between related parties. Luckily, the CIP collects information on the fraction of sales that is to related parties. We estimate (1) using only domestic-owned plants, splitting the data by whether some or no sales to related parties are reported. This is reported in the third panel of Table 6. The estimate of β is significantly different from 0 and not significantly different from 1 in both cases.

Dynamic adjustment

Given that prices are sticky, and that there may be strategic complementarities in pricing behavior, producers might prefer to spread the adjustment of relative markups to exchange rates movements over several price changes. This could contribute to our finding of an extreme form of pricing-to-market, since our baseline intensive margin exercise relies on consecutive synchronized price changes, with no intervening price changes allowed. To some extent, we get at this issue by splitting the data along the dimension of the length of the interval between the first price change and the second. The longer this interval, the more likely it is that the producer will adjust to shocks. As we note above, the estimates of β for short and long time horizons in Table 5 are statistically indistinguishable, and if anything, the point estimates go in the opposite direction to what the dynamic adjustment hypothesis would suggest.

To further investigate the possibility of dynamic adjustment of relative markups, we perform two additional exercises. First, we add a second lag of the exchange rate change to (1), where the lag length is the same as for the original independent variable:

$$\Delta_{s_k^i(t)} \ln p_{kt}^i = \alpha + \theta_{t,s_k^i(t)}^i + \beta_1 \Delta_{s_k^i(t)} \ln e_{kt} + \beta_2 \Delta_{s_k^i(t)} \ln e_{kt-s_k^i(t)} + \varepsilon_{kt,s_k^i(t)}^i \tag{6}$$

The results from estimating this equation are presented in Panel 1 of Table 7. We find that β_2 is close to and not significantly different from 0.

The other way we investigate this issue is as follows. Rather than focusing on consecutive pairs of synchronized price changes, we make use of only the first and last synchronized price changes for each plant-product pair with quote-lines in both home and UK markets. This allows us first, to use price differences over longer time horizons (more than two years) and second, to compare cases with greater and fewer numbers of intermediate price adjustments. The tradeoff is a reduced sample size. The results from performing the intensive margin exercise on this set of price changes are reported in Panel 2 of Table 7. Here, the evidence is inconclusive. The coefficient on the exchange rate change is imprecisely estimated. In the split by time horizon, it is never significantly different from 0 or from 1, though the point estimates are above 1 for horizons shorter than 2 years and below 1 for horizons greater than or equal to 2 years, which is somewhat suggestive. In the split by fewer or greater number of intermediate adjustments, the point estimates of β are in both cases below 1, but not significantly different from 1 or from 0 at the 5% level.

Potential selection issues

Because our estimates of the pricing-to-market parameter are based on a selected subset of price changes, we address the possibility that these observations are somehow special, and that the coefficient estimate is driven by this. There are two types of selection at work. First, we use only observations for plant-product pairs with simultaneous price quotes in the home market in home currency, and in the UK market in Sterling. Quotes of this type account for 19% of the sample of price changes (which are themselves 11%, unweighted, of the full price sample). We are not worried about this type of selection, because it is precisely this case that is relevant where producers in a small country export to a large country, invoicing in the large country's currency (e.g. the case for most of US imports).

The second type of selection is potentially more important. In order to identify β, we require that two consecutive sets of price changes be synchronized across the two markets. As we have documented, synchronization of price changes is frequent within plant-product pairs. Moreover, conditioning on double synchronization leaves us with a respectable 39% of price changes fulfilling the simultaneous home and UK quotes requirement. However, it could still be that non-synchronized price changes are different from synchronized price changes in a way that distorts the picture our results paint of pricing-to-market behavior.

To address this, we control for the time interval between price changes (the main observable along which we find selection)[15] by splitting the data along this variable. The results from estimating (1) on the resulting subsamples are reported in Table 5. As we have already noted, our estimates of β do not vary substantially across these groups - in all cases they are significantly different from zero, and not significantly different from one. We further note that our extensive margin results, which make use of a slightly different subset of the data (together, the intensive and extensive margin account for 59% of the sample of price changes in plant-product pairs with simultaneous home and Sterling-invoiced export quotes), are consistent with the baseline intensive margin estimates, suggesting that our results are unlikely to be driven by our focus on synchronized price changes.

Extensive margin

Table 8 reports the results from estimating how the probability of a price increase or price decrease responds to changes in the exchange rate, i.e. equations (4) and (5). The coefficient on the exchange rate change is not significantly different from zero, both in the case of increases and the case of decreases. The probability that producers change price in one market but not the other does not appear to respond to the change in the exchange rate since the last synchronized price change. This is perfectly consistent with our finding that, conditional on changing prices, producers do not readjust relative markups to respond to the changes in exchange rates since the last time prices were changed. Our confidence in our results on the intensive margin is increased by this, especially since the extensive margin is identified using a somewhat different set of observations from that used to identify the intensive margin.

In results available in the data appendix to this paper, we estimate (4) and (5), splitting the data by pricing, product and plant characteristics exactly as for the intensive margin. There are a few cases where the coefficient on the exchange rate change is significantly different from zero at the 5% level. But there is no systematic pattern to the signs of these coefficients. Consistent with the intensive margin results, there is no systematic evidence of heterogeneity across product and plant characteristics.

[15] See the data appendix for details of selection on observables.

Robustness

We perform a number of robustness checks of our results. We briefly describe some of them here. Full details are available in the data appendix to the paper.

One question that we have not so far been able to answer is how pricing-to-market differs by invoice currency. Given that prices are sticky, that the default behavior of relative markups is different under home and destination currency invoicing, and that invoice currency is a choice variable, one might expect it to differ. For quote-lines that are present in the last cross section of the price data (November 2006) the destination of exports is identified at the discretion of the respondent to the PPI survey. We use the sub-sample where the destination is reported as the UK to estimate (1) separately using exports to the UK invoiced in Sterling, and exports to the UK invoiced in Euros. We restrict this exercise to the period 2003-2005.

The results are reported in Table 9. β is relatively imprecisely estimated in both cases. In the Sterling case, the point estimate is bigger than 1, significantly different from zero at the 5% level, but not significantly different from one. In the Euro case, the point estimate is just less than 1 and it is not significantly different from zero or one at the 5% level. On the extensive margin, in neither case do we find evidence that the probability of an invoice currency price change in one market but not the other depends on the change in exchange rates since the last price change (see the data appendix for details). This is consistent with $\beta = 1$ in the Sterling case and $\beta = 0$ in the Euro case. We conclude that there is weak evidence that pricing-to-market behavior differs by choice of invoice currency. But better data would be required to answer this question definitively.

Another obvious question is whether the pricing-to-market behavior we identify is specific to producers selling in the Irish and UK markets. To test this, we additionally make use of cases of parallel pairs of price quotes for home sales and for exports invoiced in foreign currencies besides Sterling. We use this broader sample to estimate equation (1), where the independent variable is the change in the domestic exchange rate with the invoice currency. This need not be the same as the currency of the destination market, as, for example, in the case of sales to Japan that are invoiced in US dollars. We find that β is significantly different from zero, and not significantly different from one, just as we find in the baseline sample. On the extensive margin too, the results do not change.

In addition to these exercises, we check the robustness of our results to controlling for home and foreign inflation at various different levels of aggregation, to using real rather than nominal exchange rates, and to controlling for the change in aggregate imports of the relevant market (Ireland or UK) as a proxy for the level of demand. We also use forward exchange rates (where the horizon is chosen to match the median frequency of price adjustment within the plant-product pair) rather than spot exchange rates. Our results are robust to all of these modifications, the details of which are available in the data appendix.

5 Discussion and conclusions

The goal of this paper is to provide clean evidence on pricing-to-market at the plant level, while placing only the bare minimum of structure on the environment in which producers operate. We exploit data on the prices quoted by producers in Ireland for sales of a given product to two different markets segmented by exchange rate changes (Ireland and the UK). Under the assumption that percentage changes in marginal cost are the same across different markets within the same plant-product pair, this allows us to use fixed effects to isolate the response of relative markups across the two markets to movements in exchange rates. Since prices are sticky, we explore the nature of price adjustment both at the intensive and extensive margin.

We focus on prices invoiced in destination currency, the dominant case for exports from Ireland to the UK. In the case of destination currency invoicing, when prices do not change, relative markups drift one-for-one with the exchange rate. We find that even conditional on prices being changed in both markets, relative markups still move one-for-one with the exchange rate. This implies that producers do not take advantage of their decision to change prices to bring relative markups back into line. This behavior is remarkably consistent, both across plants with very different characteristics, and across sectors. Moreover, passive responses to exchange rate changes are not just a feature of the case where price adjustment is synchronized across markets. On the extensive margin, we find that the probability that producers change prices in one market but not the other does not depend on the drift in markups due to exchange rate movements since the last price change.

Qualitatively, our results are consistent with previous research, which finds evidence of pricing-to-market that goes in the same direction. Quantitatively, the pricing-to-market behavior we document is more extreme than that found in the previous literature. This is likely to be related to the differences between our data and the data used in previous studies. Principal among these is the fact that we use prices at the level of the plant-product pair rather than unit values at the level of the industry. If there are shifts in either the within-plant or cross-plant composition of sales, this will be reflected in unit values, but not in our prices. Additionally, we focus only on the case of destination currency invoicing, and by selection, this makes it more likely that we will find pricing-to-market than not.

The contrast between our results and those of the previous literature highlights the fact that using price data at a highly disaggregated level entails a tradeoff in the absence of corresponding quantity data. While on the one hand, we would like to eliminate composition effects in order to understand how transactions prices behave, on the other hand, the economic significance of that behavior depends precisely on the size of the resulting composition effects. Unfortunately, though our price data are highly disaggregated, they are not well-suited to exploring whether there is adjustment along non-price margins such as quantities, contract terms and product characteristics. Besides the fact that we do not have the requisite information on these other dimensions, respondents to the pricing survey are not required to provide prices for all transactions, or even given a metric for judging the representativeness of the sample they do provide. This makes it difficult to assess the full significance of the behavior we document.

Nevertheless, our results are consistent with anecdotal evidence (Blinder et al. (1998) and Fabiani et al. (2005)) that pricing behavior is affected by the tradeoff firms perceive between increasing current profits through higher markups, and increasing future profits through maintaining or building market share. This tradeoff may lead firms to engage in pricing strategies that might appear sub-optimal in the short run. The size of price-cost margins in the plants in our sample suggests that for exchange rate movements of the relevant magnitude, the pricing behavior we document is medium-run sustainable even if it is not necessarily statically optimal: Price-cost margins are substantial (see Table 1) while exchange rate movements between one price change and another are usually quite small (see Figure 4).

As regards a quantitative assessment of the ability of such a tradeoff to rationalize findings such as ours, recent work that takes steps in this direction includes Kleshchelski and Vincent (2009) for a closed economy, and Drozd and Nosal (2008) for an open economy. Neither of these papers directly addresses the question of whether the extreme form of pricing-to-market we document can be rationalized by reasonable parameter values, and doing so is beyond the scope of this paper. However we are hopeful that future research will throw light on this question.

References

[1] Alessandria, G. (2004), “International Deviations From the Law of One Price: The Role of Search Frictions and Market Share,” *International Economic Review*, 45 (4), 1263-1291.

[2] Atkeson, A. and A. Burstein (2008), “Pricing-to-Market, Trade Costs and International Relative Prices,” *American Economic Review* 98 (5), 1998-2031.

[3] Blinder, A., E. Canetti, D. Lebow and J. Rudd (1998), *Asking About Prices: A New Approach to Understanding Price Stickiness*. New York, Russell Sage Foundation.

[4] Bergin, P. and R. Feenstra (2001), “Pricing-to-Market, Staggered Contracts and Real Exchange Rate Persistence,” *Journal of International Economics* 44, 333-359.

[5] Bils, M., P. Klenow and B. Malin (2009), “Reset Price Inflation and the Impact of Monetary Policy Shocks,” mimeo.

[6] Burstein, A., L. Eichenbaum and S. Rebelo (2007), “Modeling Exchange Rate Passthrough After Large Devaluations,” *Journal of Monetary Economics* 54 (2), 346-368.

[7] Burstein, A. and N. Jaimovich (2009), "Understanding Movements in Aggregate and Product Level Real Exchange Rates," mimeo.

[8] Domowitz, I., R. G. Hubbard and B. Petersen (1986), "Business Cycles and the Relationship between Concentration and Price-Cost Margins," *Rand Journal of Economics* 17 (1), 1-17.

[9] Drozd, L., and J. Nosal (2008), "Understanding International Prices: Customers as Capital," mimeo.

[10] Fabiani, S., M. Druant, I. Hernando, C. Kwapil, B. Landau, C. Loupias, F. Martins, T. Matha, R. Sabbatini, H. Stahl and A. Stokman (2005), "The Pricing Behaviour of Firms in the Euro Area: New Survey Evidence," ECB Working Paper 535.

[11] Gil-Pareja, S., (2003), "Pricing to Market Behavior in European Car Markets," *European Economic Review* 47, 945-962.

[12] Goldberg, P. and R. Hellerstein (2008), "A Framework for Identifying the Sources of Local-Currency Price Stability with an Empirical Application," mimeo.

[13] Goldberg, P. and M. Knetter (1997), "Goods Prices and Exchange Rates: What Have We Learned?" *Journal of Economic Literature* 35 (3), 1243-1272.

[14] Goldberg, P. and F. Verboven (2001), "The Evolution of Price Dispersion in the European Car Market," *Review of Economic Studies* 68 (4), 811-848.

[15] Goldberg, P. and F. Verboven (2005), "Market Integration and Convergence to the Law of One Price: Evidence from the European Car Market," *Journal of International Economics* 65 (1), 49-73.

[16] Gopinath, G., O. Itskhoki and R. Rigobon (2008), "Currency Choice and Exchange Rate Passthrough," *American Economic Review*, forthcoming.

[17] Gopinath, G. and O. Itskhoki (2009), "Frequency of Price Adjustment and Passthrough," *Quarterly Journal of Economics*, forthcoming.

[18] Gopinath, G. and R. Rigobon (2008), "Sticky Borders," *Quarterly Journal of Economics*, 123 (2).

[19] Gust, C., S. Leduc and N. Sheets (2008), "The Adjustment of Global External Balances: Does Partial Exchange Rate Pass-Through to Prices Matter?" mimeo.

[20] Klenow, P. and O. Kryvtsov (2008), "State-Dependent or Time-Dependent Pricing: Does it Matter for Recent US Inflation?" *Quarterly Journal of Economics* 123 (3), 863-904.

[21] Kleshchelski, I. and N. Vincent (2009), "Market Share and Price Rigidity," *Journal of Monetary Economics* 56 (3), 344-352.

[22] Knetter, M. (1989), "Price Discrimination by U.S. and German Exporters," *American Economic Review* 79 (1), 198-210.

[23] Nakamura, E. and D. Zerom (2008), "Accounting for Incomplete Pass-Through," mimeo.

[24] Nakamura, E. and J. Steinsson (2008), "Five Facts About Prices: A Reevaluation of Menu Cost Models," Quarterly Journal of Economics 123 (4), 1415-1464.

[25] Nakamura, E. and J. Steinsson (2009), "Lost in Transit: Product Replacement Bias and Pricing to Market," mimeo.

[26] Rauch, J. (1999), "Networks Versus Markets in International Trade," *Journal of International Economics* 48, 7-35.

[27] Vermeulen, P., D. Dias, M. Dossche, E. Gautier, I. Hernando, R. Sabbatini and H. Stahl (2007), "Price Setting in the Euro Area: Some Stylized Facts From Producer Price Data and Producer Surveys," ECB Working Paper 727.

Table 1: Summary Statistics on Plants

	1995		2005	
	PPI	CIP	PPI	CIP
% of plants exporting to				
Anywhere	55	61	75	72
UK	18	16	12	11
% plants invoicing UK sales in				
IEP/ EUR only	15	22	24	18
STG only	65	56	44	33
Other only	4	4	2	1
mix	15	18	30	48
% UK sales invoiced in				
IEP/ EUR	11	9	16	20
STG	84	77	70	69
Other	5	13	14	11
% plants importing materials from				
Anywhere	75	58	83	54
UK	69	49	73	48
% of imported intermediates from				
Anywhere	40	48	64	59
UK	16	15	18	16
% plants by ownership				
Foreign	34	16	29	13
UK	7	3	3	2
Price-cost margins, distribution across plants				
25th percentile	0.20	0.20	0.22	0.23
50th percentile	0.30	0.31	0.33	0.35
75th percentile	0.42	0.42	0.45	0.47
Employees, distribution across plants				
25th percentile	26	7	22	6
50th percentile	59	15	48	13
75th percentile	129	41	110	36

Note: Information for imports is based on the roughly 90% of the population for which comparable information is available over the entire time period. Information for the export currency is based on the roughly 95% of the population where information is available over the entire time period.

Table 2: Hierarchical structure of price data

	Plants	Plant-product pairs	Quote-lines	Obs
1995	669	1102	4887	54009
1996	647	1068	4795	52163
1997	627	1040	4658	51022
1998	595	1010	4807	49198
1999	555	947	4174	46327
2000	580	977	4496	46910
2001	653	1074	4929	50017
2002	808	1235	5456	53224
2003	876	1326	5819	59740
2004	852	1295	5368	58658
2005	836	1243	4995	56428
total	1213	1946	12232	577696

Note: Number of distinct values of each category observed in the relevant time-frame.

Table 3: Weighted mean frequency of price adjustment in invoice currency

	Unadjusted	Exit adjusted
total	16	18
Destination market		
home	19	20
export	14	17
Invoice currency for exports		
IEP/ EUR	11	13
STG	16	17

Note: Obs. weighted by plant sales in home/export market as appropriate. Equal weighting within plant-market-years. Period covered is Jan 1995-Dec 2005. Exit adjustment treats quote-line exit like a price change.

Table 4: Synchronization of price changes within plant-product pairs

	% of plant-prod-mths with >1 quote and >=1 price change	Of which		
		One price changes	>1 but < all change	All prices change
Full sample	16	21	28	51
Irl & UK sample	20	28	45	28

Note: Irish and UK sample includes plant-product-months with at least one home currency price quote in the home market and at least one Sterling-invoiced price quote in the export market. Price changes refer to price changes in invoice currency.

Table 5: Intensive margin I

	$\Delta_{s_k^i(t)} ln(e_{tk})$		R^2-adj.	N	# f.e.	# clust
Baseline						
	1.014	0.086**	0.669	4212	1047	86
Panel 1						
Median frequency of price adjustment of plant-product pair						
freq<0.3	1.069	0.123**	0.638	1226	293	68
0.3 < freq < 0.5	0.927	0.150**	0.605	1089	313	8
0.5 < freq < 0.7	1.027	0.329**	0.665	1269	242	7
freq > 0.7	0.996	0.385**	0.689	628	199	5
Time interval between synchronized price changes						
1 month	0.773	0.380**	0.657	748	381	42
2-5 months	0.932	0.115**	0.799	1149	516	46
6-11 months	0.961	0.035**	0.881	1377	578	71
12+ months	1.033	0.282**	0.636	938	441	64
Panel 2						
Type of product (Vermeulen et al. 2007)						
consumer food products	0.913	0.165**	0.555	1300	393	20
cons non-food non durab	1.859	0.603**	0.432	29	13	6
cons durables	1.062	0.187**	0.682	147	28	8
intermediates	1.113	0.234**	0.813	1481	404	31
capital goods	1.048	0.095**	0.602	1251	207	22
Type of product (Rauch 1999)						
organized exchange	1.030	0.416**	0.528	985	292	10
reference priced	0.873	0.166**	0.849	378	146	7
differentiated	1.141	0.151**	0.559	1666	345	46

Note: Estimation method is OLS. Dependent variable is log change in home currency price since last price change. All regressions include a constant and the full set of plant-product-month-age-of-price fixed effects. Observations are weighted by sales. Standard errors are clustered at the plant level. Standard errors in brackets. Two stars indicates significantly different from zero at the 5% level, one star indicates significantly different from zero at the 10% level.

Table 6: Intensive margin II

	$\Delta_{s_k^i(t)} ln(e_{tk})$		R^2-adj.	N	# f.e.	# clust
Baseline						
	1.014	0.086**	0.669	4212	1047	86
Panel 1						
Number of employees						
<20	1.252	0.124**	0.856	283	98	10
20-29	1.067	0.232**	0.665	894	188	24
50-249	1.062	0.116**	0.554	2129	477	44
250-499	0.952	0.172**	0.613	752	227	7
Panel 2						
Quartiles of share of sales exported to the UK						
Q1	0.729	0.028**	0.433	68	25	4
Q2	0.937	0.054**	0.776	117	25	11
Q3	1.065	0.215**	0.576	1865	378	34
Q4	1.013	0.112**	0.708	2162	619	57
Quartiles of share of variable cost imported from the UK						
Q1 & Q2	1.029	0.256**	0.631	1334	413	32
Q3	1.033	0.131**	0.704	617	170	38
Q4	1.002	0.066**	0.670	2213	445	40
Panel 3						
Ownership						
domestic	1.042	0.126**	0.573	2842	791	64
foreign	0.963	0.084**	0.817	1370	256	24
UK	0.990	0.011**	0.963	175	27	8
Quartiles of price-cost margins						
Q1	0.977	0.166**	0.605	1757	466	39
Q2	0.924	0.125**	0.699	1178	294	38
Q3	1.612	0.546**	0.813	951	208	34
Q4	0.990	0.105**	0.400	326	79	22
Share of sales to related parties						
None	1.114	0.257**	0.578	1731	466	38
Some	1.111	0.110**	0.670	438	111	14

Note: Estimation method is OLS. Dependent variable is log change in home currency price since last price change. All regressions include a constant and the full set of plant-product-month-age-of-price fixed effects. Observations are weighted by sales. Standard errors are clustered at the plant level. Standard errors in brackets. Two stars indicates significantly different from zero at the 5% level, one star indicates significantly different from zero at the 10% level.

Table 7: Intensive margin of price adjustment: Dynamics

	$\Delta_{s_k^i(t)} ln(e_{tk})$		$\Delta_{s_k^i(t)} ln\left(e_{t-s_k^i(t),k}\right)$		R^2-adj.	N	# f.e.	# clust
Panel 1								
Two lags of exchange rate changes								
	0.934	0.084**	-0.054	0.110	0.666	3794	960	67
Panel 2								
First and last synchronized price changes, all observations								
	0.781	0.511			0.391	677	116	96
First and last synchronized price changes, by interval between first and last								
<6 months	1.594	5.027			0.155	43	23	20
6-11 months	1.376	0.878			0.438	65	27	22
12-23 months	1.808	1.343			0.468	141	40	35
24+ months	0.736	0.612			0.361	428	75	66
First and last synchronized price changes, by number of intervening price changes								
<2	0.858	0.463 *			0.329	400	58	49
≥2	0.760	0.660			0.302	277	58	50

Note: Estimation method is OLS. Dependent variable in first panel is log change in home currency price since last price change. Dependent variable in subsequent panels is log change in home currency price between first and last synchronized price change for matched pair of home and UK quotes. All regressions include a constant and the full set of plant-product-month-difference interval fixed effects. Observations are weighted by sales. Standard errors are clustered at the plant level. Standard errors in brackets. Two stars indicates significantly different from zero at the 5% level, one star indicates significantly different from zero at the 10% level.

Table 8: Extensive margin

	Increases		Decreases	
	coeff.	s.e.	coeff.	s.e.
$\Delta_{s_k^i(t)} \ln e_{uk,t}$	0.59	(3.61)	1.62	(4.99)
N	4,873		4,564	
# f.e.	921		875	
# clusters	129		103	
Pseudo-R^2	0.00		0.00	
χ^2 (p-value)	0.03	(0.87)	0.11	(0.75)

Note: Dependent variable is indicator for increase in invoice currency price. Estimator is conditional logit, conditioning on plant-product-month-age of price fixed effects. Observations are weighted by sales. Standard errors in brackets. Standard errors are clustered at the plant level. Two stars indicates significantly different from zero at the 5% level, one star indicates significantly different from zero at the 10% level.

Table 9: Robustness

Invoice curr.	$\Delta_{s_k^i(t)} ln(e_{tk})$		R^2-adj.	N	# f.e.	# clust
home	0.969	0.569*	0.662	745	236	41
destination	2.521	1.222*	0.677	741	235	40

Note: Sample is restricted to matched pairs of home sales and foreign sales where destination is identified as the UK in November 2006. Estimation method is OLS. Dependent variable is log change in home currency price since last price change. All regressions include a constant and the full set of plant-product-month-age-of-price fixed effects. Observations are weighted by sales. Standard errors are clustered at the plant level. Standard errors in brackets. Two stars indicates significantly different from zero at the 5% level, one star indicates significantly different from zero at the 10% level.

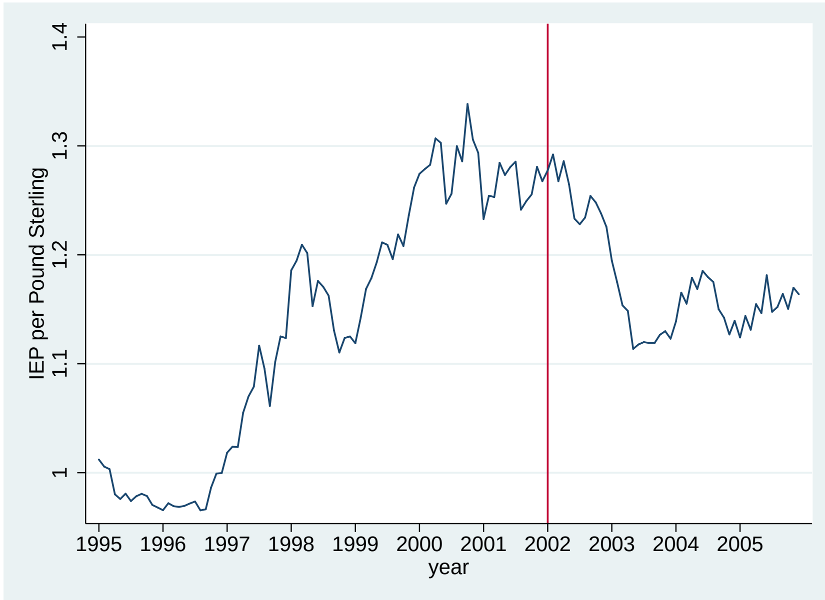

Figure 1: IEP/EUR per Pound Sterling over the sample period

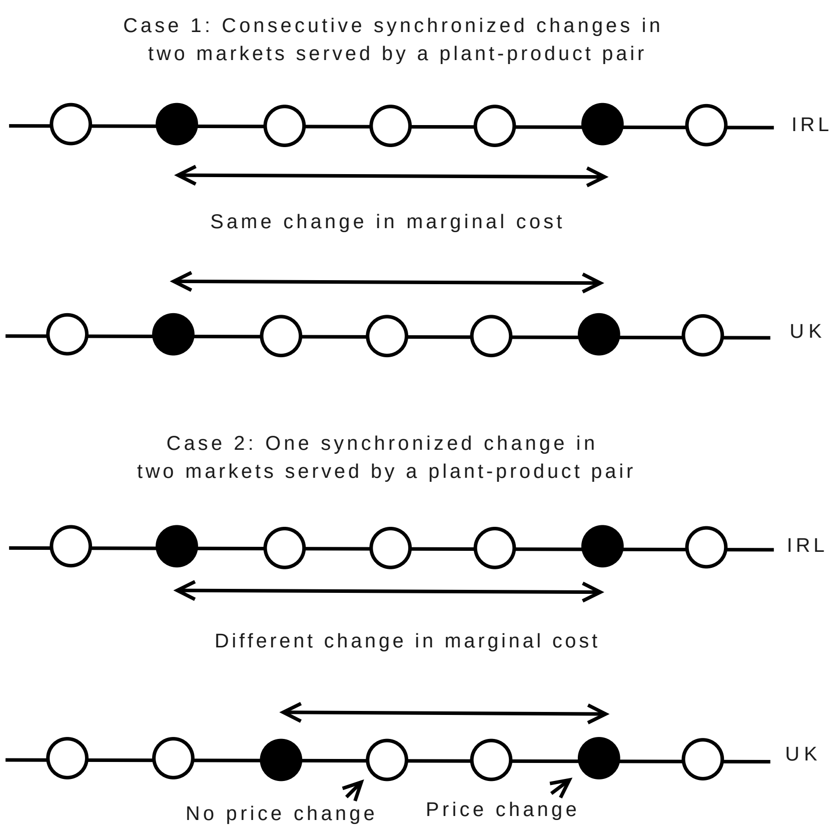

Figure 2: When is β identified?

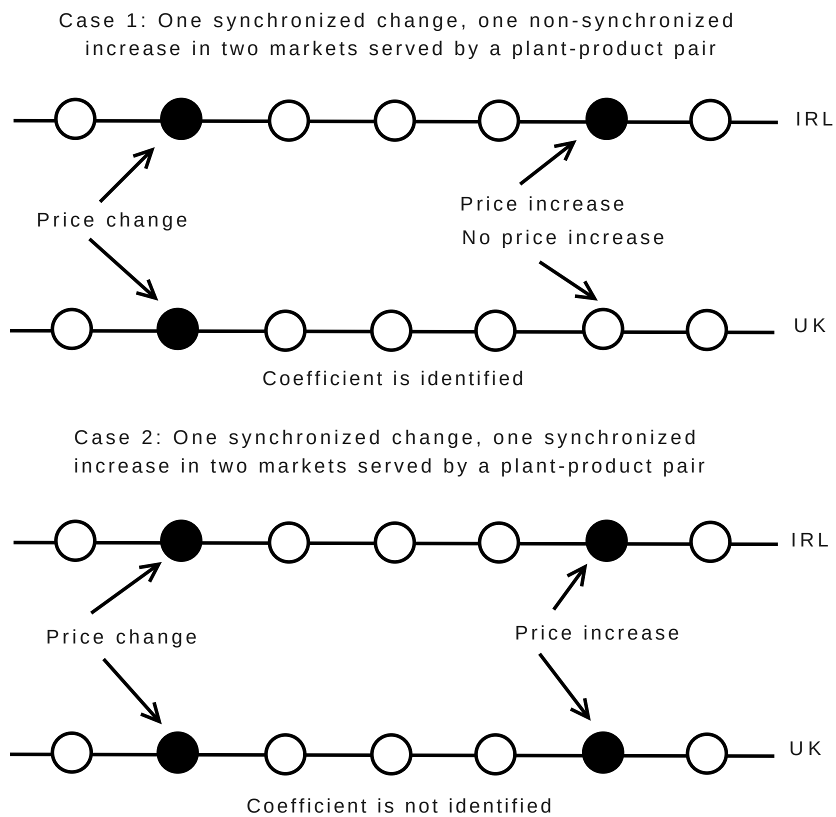

Figure 3: When is $(\beta - 1)$ identified?

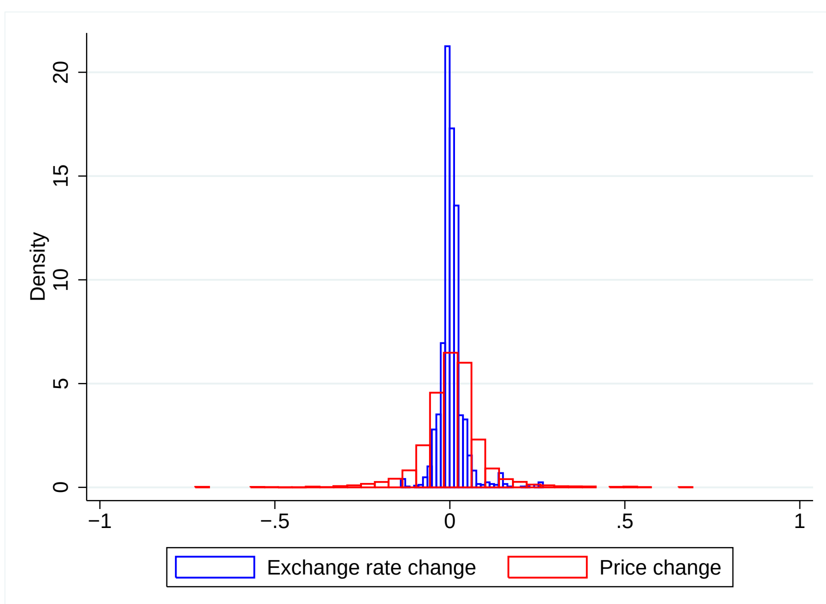

Figure 4: Identifying variation in invoice currency prices and exchange rates

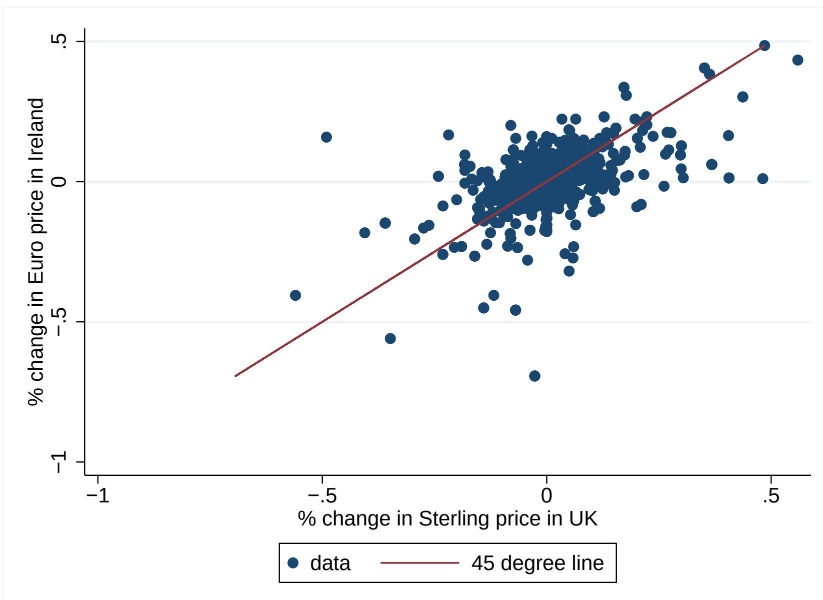

Figure 5: Change in Sterling price in UK against change in Euro price in Ireland