

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

Doireann Fitzgerald[†] Stefanie Haller[‡]

May 2012

Abstract

We use micro data on Irish producer prices to provide clean evidence on pricing-to-market across a broad range of manufacturing sectors. We have monthly observations on prices charged by the same plant for the same product to buyers in Ireland and the UK, two markets segmented by variable exchange rates. Assuming that relative marginal cost is constant across markets within a plant and a product, this allows us to observe the behavior of the markup in the UK market relative to the home market. To identify pricing-to-market that goes beyond what is mechanically due to price stickiness, we condition on episodes where prices change. When prices are invoiced in local currency, conditional on prices changing, the ratio of the markup in the foreign market to the markup in the home market increases one-for-one with depreciations of home against foreign currency and decreases one-for-one with appreciations of home against foreign currency, a very particular form of pricing-to-market.

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, CAED-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, Whitaker Square, Sir John Rogerson's Quay, Dublin 2, Ireland, stefanie.haller at esri.ie

1 Introduction

Pricing-to-market is a leading potential explanation for the high correlation of real and nominal exchange rates and the volatility and persistence of real exchange rates (see, e.g. Dornbusch (1987), Krugman (1987), and more recently, Bergin and Feenstra (2001), Atkeson and Burstein (2008) and Drozd and Nosal (2012a), (2012b)).¹ Producers price-to-market if they choose different markups for the same good in different markets, and in particular, if the ratio of the markups charged in two different markets moves around with the nominal exchange rate between the two markets. Though suggestive, much of the existing evidence on pricing-to-market is based on less-than-ideal data where it is not possible to disentangle pricing-to-market from price stickiness or local content, nor yet assess how widespread is this behavior. We use a novel data set on producer prices and a clean identification strategy to characterize the nature of pricing-to-market more precisely than the previous literature, and for a broad range of manufacturing sectors. We show that producers do indeed price-to-market, and moreover that that this goes beyond what is mechanically due to price stickiness, in the following sense. When prices are sticky in local currency, *conditional on changing prices*, producers choose markups such that the ratio of the markup in the foreign market to the markup in the home market increases one-for-one with home currency depreciations and decreases one-for-one with home currency appreciations. This implies that real exchange rates at the plant-product level track nominal exchange rates, even conditional on price changes. This behavior is pervasive across manufacturing plants and sectors with different characteristics, given local currency invoicing.

Our empirical analysis is guided by a simple partial equilibrium model of the producer’s pricing problem. We condition on the producer’s choice of invoice currency. If pricing is time-dependent, producers decide what new price to charge when they change prices. If there is a fixed cost of changing prices, producers additionally choose the timing of price changes. In both cases, the optimal price conditional on price adjustment (we refer to this throughout as the “optimal reset price”²) is a latent variable that is observed only in the event of a

¹The major competing explanations are sticky prices and local content of apparently “traded” goods - See Engel (2002) for a useful survey.

²The “reset price” terminology follows Bilal, Klenow and Malin (2011) who examine the role played by reset prices in explaining the performance of sticky-price models in a closed economy.

price change. The goal of our empirical analysis is to characterize how the relative markup between foreign and home markets implied by optimal reset prices responds to movements in nominal exchange rates between the two markets. By focusing on optimal reset prices, we isolate that part of pricing-to-market that is not mechanically due to price stickiness.

Our data is constructed by matching the annual plant census for Ireland with the monthly micro-data on producer prices used to generate the Irish producer price index. Clean identification of the behavior of markups relies on the fact that plants participating in the price survey are encouraged to provide matched prices charged to buyers from home and export markets for product categories sold to both markets, and that these prices are measured at the factory gate. Under the assumption that relative marginal cost across markets does not move around over time within a plant-product pair, this allows us to identify the behavior of relative markups. At the same time, the fact that we have monthly data allows us to observe the timing of price changes and hence condition on price changes. We focus on plants reporting parallel price series for the same product for buyers from the home market and from the UK, where the UK prices are invoiced in local currency. The reason for focusing on the UK and on the case of local currency invoicing is that for most of our sample, the data set does not identify the precise destination for export price quotes. We must use the currency of invoicing to identify the destination market, and Sterling-invoiced export prices account for the majority of foreign-currency invoiced export price quotes.³

We use two exercises to characterize pricing-to-market conditional on price changes. First we examine by how much producers adjust their price (for a given product) in the foreign market relative to the home market in response to exchange rate changes, conditional on changing prices in both markets simultaneously. Second, we examine how the probability that in a particular month, a producer who sells the same product in both markets increases (or decreases) the price in one market but not the other depends on the size and sign of exchange rate changes since the last price change. In both cases, by conditioning on the preceding price changes having been synchronized across both markets we can use fixed effects to control for changes in marginal cost. This strategy relies on the fact that in our

³Moreover, the UK is the modal export destination for Irish producers, and over 70% of industrial exports to the UK are invoiced in local currency.

data, synchronization of price changes across markets within plant-product pairs is the norm rather than the exception.

Using both approaches we find that if prices are sticky in local currency, there is pricing-to-market conditional on price adjustment, of a very particular type. The ratio of the optimal reset price in the foreign market to the optimal reset price in the home market increases one-for-one with depreciations of the home currency against the foreign currency, and decreases one-for-one with appreciations of the home currency against the foreign currency. Under the assumption of constant relative marginal cost, this implies that even conditional on price adjustment, the markup in the foreign market relative to the home market moves one-for-one with exchange rate movements. Our findings do not vary across industries or plants with different characteristics. They are also invariant to differences in pricing behavior (such as the frequency of adjustment) at the level of the plant-product pair. However we do find some evidence that the degree of pricing to market may be less extreme over horizons longer than 2 years.

For a small sub-sample of the data, we can perform the same exercises using prices that are sticky in home currency. In this case, we cannot reject the hypothesis that conditional on price changes, relative markups do not move around with exchange rates, exactly the opposite to what we find for local currency invoicing.

Our results have potentially important implications for the quantitative literature which aims to explain stylized facts about the behavior of real exchange rates. In documenting the comovement of reset prices with nominal exchange rates under local currency invoicing, our results point to a potential source of real exchange rate persistence and volatility that goes beyond the mechanical effect of price stickiness. For models with sticky prices, allowing for realistic behavior of reset prices along the lines we document could potentially improve performance in matching real exchange rate persistence. For models with flexible prices, our results place an important discipline on the nature of pricing-to-market, both in terms of the exact comovement of relative markups and exchange rates, and by the fact that this comovement does not vary systematically with industry or plant characteristics.

Our work relates to the existing literature on testing for pricing-to-market as follows. Our empirical strategy builds on Knetter (1989), who proposes identifying markup responses

to exchange rate movements by using fixed effects to control for marginal cost. Previous work that uses this approach⁴ has had to work with less-than-ideal data. The prices being compared do not necessarily apply to products produced by the same plant, prices are not always measured at the factory gate, and invoice currency and the timing of price changes are not usually observed. Though the results of this literature are consistent with pricing-to-market, it is not always obvious whether this is due to a failure of the identifying assumption of common marginal cost, to the presence of local content, to the mechanical effect of price stickiness, or to pricing-to-market as we define it. An exception is Burstein and Jaimovich (2009). They use weekly data on the prices paid to suppliers by a supermarket chain operating in the US and Canada for a narrow set of nonperishable nondurable goods, where goods are defined by UPC code, and where the authors confirm that they were sourced in the same country. Their findings are very similar to ours: real exchange rates at the level of the narrowly defined good move one-for-one with exchange rates.

There is also a structural literature on pricing-to-market, which makes use of rich data for very narrowly defined market segments. Within this literature, Goldberg and Hellerstein (2011), and Nakamura and Zerom (2010) find evidence of pricing-to-market conditional on price changes in the beer and coffee markets respectively. Relative to this literature, our results apply across a broad range of manufacturing sectors, and are based on relatively parsimonious identifying assumptions. Our work is also closely related to Gopinath, Itskhoki and Rigobon (2010), who find evidence of slow pass-through of exchange rate changes into US import prices invoiced in dollars, conditional on price adjustment. Because they observe prices in only one market, they cannot identify whether this is due to desired variation in relative markups or to the behavior of marginal cost. But as we show later, we can reproduce their findings using our data, implying that the pricing-to-market behavior we document is consistent with their findings on exchange rate pass-through.

The second section of the paper describes the data. The third section lays out a partial equilibrium model of the producer's pricing decision. The fourth section describes our empirical strategy. The fifth section presents our results. The final section discusses the results

⁴See Goldberg and Knetter's 1997 survey, and e.g. Goldberg and Verboven (2005) who use a related approach.

and concludes.

2 Our data

Our data comes from two sources, the Irish Census of Industrial Production (CIP) and the micro data collected for the purpose of constructing the Irish Producer Price Index (PPI) and Wholesale Price Index (WPI). Here, we briefly outline the salient features of both data sets. Additional details are provided in the Appendix. Survey and methodology documents are available at www.cso.ie.

The CIP is an annual census of producers in manufacturing, mining and utilities. All plants with 3 or more persons engaged are required to fill in a return. We use the data for the years 1995-2005. Of the variables collected in the CIP, those relevant for our purposes are the 4-digit industrial classification (NACE Revision 1.1), country of ownership, value of sales, share of sales exported (with some destination and currency invoicing information for export sales), employment, wage bill, energy and materials expenditures (with the share of imported materials, and some origin and invoicing information) and share of sales to related parties.

As is standard in other European countries and the US, the micro data used to construct the Irish PPI is collected through monthly surveys of plants. The sampling frame is the population of plants in the CIP. Plants selected to participate in the price survey 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. We have access to the monthly data from January 1995 through November 2006. The price data can be linked to the CIP data using a unique plant identifier. On average, 14% of CIP plants accounting for 38% of total CIP 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 (wholesalers who do not produce also participate in the price survey).

First-time participants in the PPI survey are asked to provide prices for their main products that are “suitable for pricing each month.” For each quote-line (to use the terminology of Klenow and Malin (2011)), they are asked to provide a product category, a detailed item

description, information on trading terms relevant to the price, the units for which the price is quoted, the destination market (for exports), any discounts and surcharges on the basic price, the invoice currency, and price in the invoice currency. Each subsequent month, participants receive a form where for all quote-lines from the previous month, the initial responses for these variables are already filled in, along with last month's price. They use this form to report the price invoiced for the product on the 15th of the current month.

The feature of the PPI survey that is crucial for our purposes, and that distinguishes it from similar surveys in other countries, is that prices for domestic sales and exports are collected using the same survey. Further, the initial survey form is set up in such a way as to explicitly solicit matched price quotes in home and export markets for each product category for which a plant reports prices. These product categories are classified using a system that is unique to the PPI survey. We have examined a sample of the product descriptions, and the classification is at the 6 or 8-digit level. In robustness checks we also make use of a more demanding match across home and export price quotes based on detailed item descriptions.

Our price data does have some shortcomings. First, several variables, including the trading terms, the units to which the price refers and, crucially, the destination market for exports, are reported only at the discretion of the respondent. Further, if provided, we observe these variables only for quote-lines present in the last cross-section (November 2006) as this part of the data file is overwritten every month. Given this, our baseline analysis uses the invoice currency to identify the precise destination market for exports, and hence must focus on the case of local currency invoicing. We restrict attention to export price quotes invoiced in Sterling (60% of foreign-currency-invoiced price quotes), where we are fairly certain that the currency identifies the destination market.⁵ We do make use of quote-lines present in the last cross-section where the destination is identified separately from the invoice currency to examine pricing-to-market under home currency invoicing, but the sample involved is small.

Other shortcomings are as follows. In reporting the current month's prices, participants

⁵We confirm using quote-lines present in the last cross-section that Sterling identifies the destination market as the UK with high probability. The destination description is sometimes imprecise e.g. "UK/Spain." We classify observations where such as these where the destination includes the UK, as prices that apply to the UK market. By this definition, only 4% (unweighted) of Sterling-invoiced export price quotes where the destination is identified are for destinations other than the UK.

are instructed to exclude delivery charges itemized separately on the invoice, but to report prices inclusive of delivery charges if they are included in the price and not separately itemized. There is no systematic indicator for whether or not delivery charges are included. Moreover, there is no requirement to flag prices for transactions between related parties (though we know from the CIP whether or not the plant reports any sales to related parties, and the majority of plants do not report such sales). Participants are instructed to discontinue a quote-line and replace it with another if the product or terms of sale are no longer available or representative. Relative to the micro data for the US PPI, we observe relatively little within-plant product turnover. Conversations with the CSO indicate that some product replacement may go unobserved by us, because the quote-line identifiers may not always be recoded in the event of a substitution, and the precise item description is observed only for quote-lines present in the last cross section, and only for that cross-section. Finally, unlike other micro price data sets such as the US PPI, we do not see missing observations in the middle of quote-lines. Conversations with the CSO indicate that when a price is missing, the previous price is brought forward. Unfortunately these filled-in observations are not flagged during our sample period.

The variables in the PPI data that we make use of are plant and product category identifiers, product descriptions, detailed item descriptions, prices expressed in domestic currency (adjusted for discounts or surcharges), the indicator for whether the price refers to a domestic sale or an export sale, the currency in which the price was originally quoted, and for prices not originally quoted in home currency, the exchange rate used to make the conversion from foreign currency.

Summary statistics on plants

On average, about 5000 plants appear in the CIP in each year, and between 550 and 900 of these participate in the PPI survey. Given our focus on prices in the home and UK markets, some facts about plant openness are of particular note. Table 1 reports key summary statistics on plant openness and size, for all plants in the CIP and for those in the matched sample, for 1995 and 2005, the first and last years of our sample. Further summary statistics on the plant data are available in the Appendix. Almost all plants sell something in the domestic

market. Meanwhile the UK market is a very important one, both in terms of fraction of plants participating, and fraction of total sales by the industrial sector. More than 85% of exporters export to the UK, and exports to the UK account for more than 10% of total industrial output. From the CIP, we also have information on the choice of invoice currency for sales to the UK market. At least 70% of industrial exports to the UK are invoiced in Sterling. These producers provide an ideal laboratory for examining the effects of exchange rate changes on pricing behavior, because pricing appropriately in home and export markets is important for the bottom line. At the same time, we have to be careful to take account of the fact that exchange rate movements can be a source of cost as well as demand shocks, as is clear from the importance of imported intermediates.

Summary statistics on prices

On average, in each month, each plant participating in the PPI survey reports prices for 1.6 distinct products (classified as described above). On average for each plant-product pair, 4.4 price quotes are reported. Within a plant-product pair, multiple price quotes are reported both because there may be quotes for multiple markets (home, export) and because within each market there may be multiple quotes. This adds up to between 4000 and 6000 distinct price quotes in any given month. In steady state (i.e. when we are sufficiently far from the beginning and the end of the sample that there is neither left-censoring nor right-censoring of quote-lines), the median quote-line is observed for between 80 and 90 months.

The behavior of producer prices in Ireland is broadly similar to that in six Euro-zone countries as reported in Vermeulen et al (2007). Table 2 reports the weighted mean frequency of price adjustment (calculated as the fraction of invoice currency prices that change, weighted by plant-market sales as described in the Appendix), for the sample as a whole, for home sales and exports separately, and, for exports by currency of denomination. Prices are sticky in invoice currency. For the sample as a whole, the weighted mean frequency of adjustment is 0.16. For domestic sales, the frequency is 0.19, while the frequency of adjustment of Sterling prices for Sterling-invoiced exports is 0.16. Vermeulen et al report weighted mean frequencies of adjustment in the range 0.15 to 0.25 for the six countries for which they have data.

In the Appendix, we report statistics on the frequency of price increases and decreases, and the size of price changes. These statistics further illustrate that the behavior of producer prices in Ireland is fairly typical of that in other European countries. Price increases are more frequent than price decreases, while the size of price increases and decreases is roughly symmetric, with the median increase being around 3%.

A feature of price-setting behavior that is crucial for our identification strategy is synchronization of price changes across quote-lines within a plant-product pair. In particular, we make use of cases where there is synchronization of price changes in the Irish and UK markets (identified by Sterling invoicing of exports). Summary statistics on synchronization of price changes are reported in Table 3. The first column reports the percentage of plant-product-months with more than one price quote where at least one price changes, i.e. the cases where there is potential for synchronization of price changes. The second column reports the percentage of plant-product-months with more than one price quote and at least one price change where there is exactly one price change. These are cases where price changes are not synchronized. In the full sample, these cases account for one fifth of plant-product-months with at least one price change. They account for just over a quarter of the sample where there is a price quote in both Irish and UK markets. The corollary is that for plant-product pairs with price quotes in both Irish and UK markets, just under three quarters of the time, episodes of price change tend to be synchronized. The third and fourth columns report on the degree to which this synchronization is imperfect or perfect (i.e. affects all quote-lines within a plant-product pair).

Exchange rates and other aggregate variables

In line with the way the price data is collected, to construct our key independent variable we use the spot exchange rate observed on the 15th of the month (or the date closest to the 15th of the month on which markets are open). The source is the Central Bank of Ireland. Our sample period covers an interval during which the home currency (first the Irish pound, and then the Euro) depreciates roughly 35% against Sterling, followed by a interval during which it appreciates by around 20%. Month-to-month fluctuations are substantially smaller. This is illustrated in Figure 1.

For robustness checks, we also make use of monthly data on the Irish and UK CPIs for the period 1997-2005, taken from the OECD. We also construct a monthly measure of real aggregate demand in the two markets over the period 1997-2005 as follows. We take data on seasonally adjusted quarterly nominal aggregate expenditure (GDP plus imports minus exports) for the Irish and UK economies from the OECD.⁶ We deflate by the quarterly CPI for the relevant country (also from the OECD). We then construct a 3-month moving average of the resulting real aggregate expenditure to use as a monthly variable.

3 Model

To provide context for our empirical work, we outline a partial equilibrium model of the producer’s pricing decision. Relative to Knetter (1989), we explicitly allow for the possibility of price stickiness, which makes the pricing problem dynamic. We condition on the choice of invoice currency,⁷ and for brevity, present only the case of local currency invoicing. We make two assumptions that allow us to consider separately the decision to change prices for each plant-product pair in each market: In each period, the plant faces marginal cost for each product that does not depend on the quantity of that product produced.⁸ Meanwhile, the fixed costs, if any, of changing prices are incurred separately for each product and market.

First, some notation: i indexes plant-product pairs, k indexes markets, and t indexes months. P_{t-1}^{ik*} is the local currency price charged by i in market k at date $t-1$ (as standard, a star indicates a price in foreign currency). It is an endogenous state variable of the producer’s problem. I_t^{ik} is an indicator variable that equals 1 if i changes its local currency price in market k at time t , and equals zero otherwise. \hat{P}_t^{ik*} is the local currency price i chooses if it does change its price in market k at date t (what we refer to as the optimal reset price). I_t^{ik} and \hat{P}_t^{ik*} are choice variables. Θ_t^{ik} is a vector of exogenous variables that

⁶Quarterly national accounts data for Ireland only begin in 1997.

⁷We document in the Appendix that there is substantial heterogeneity both within plants and across plants with similar characteristics in the choice of invoice currency for sales to the UK market. This suggests that information on the counterparty to a transaction may be necessary to fully understand invoice currency choice. Since we know nothing about counterparties, we do not pursue this question here.

⁸This assumption is stronger than the comparable assumption made by Knetter (1989), as in a menu cost environment, allowing marginal cost to vary with total quantity produced induces cross-market dependence in the timing of price changes.

affect i 's real gross profit from market k . F_t^i is a (real) cost that is incurred if i changes its price in market k at date t .

The problem of plant-product pair i in market k can be expressed as follows:

$$V(P_{t-1}^{ik*}, \Theta_t^{ik}) = \max_{\substack{I_t^{ik} \in \{0, 1\} \\ \hat{P}_t^{ik*} > 0}} [\Pi(P_t^{ik*}, \Theta_t^{ik}) - I_t^{ik} F_t^i + \beta \mathbb{E}V(P_t^{ik*}, \Theta_{t+1}^{ik})] \quad (1)$$

subject to

$$P_t^{ik*} = I_t^{ik} \hat{P}_t^{ik*} + (1 - I_t^{ik}) P_{t-1}^{ik*} \quad (2)$$

We can write the real gross profit flow from market k at date t as follows:

$$\Pi(P_t^{ik*}, \Theta_t^{ik}) = \frac{[E_t^k P_t^{ik*} - (1 + \tau^{ik}) C_t^i]}{P_t} q(P_t^{ik*}, \Theta_t^{ik}) \quad (3)$$

Here, $q(\cdot)$ is (residual) demand. E_t^k is the price of 1 unit of country k 's currency in terms of home currency. C_t^i is nominal marginal cost in plant-product pair i at date t , expressed in home currency. τ^{ik} is a per-unit multiplicative cost of serving market k that is assumed time-invariant. P_t is the home country price level, expressed in home currency. The vector Θ_t^{ik} includes E_t^k , C_t^i , τ^{ik} and P_t . It may also include variables that shift aggregate demand in market k , idiosyncratic shocks to plant-product pair i 's demand in market k , and depending on the market structure, competitors' prices in market k (under perfect competition or monopolistic competition, i takes competitors' prices as exogenous).

The way the problem is framed nests both state-dependent and time-dependent price setting as well as full price flexibility. To see this, note that if $F_t^i = 0$ with probability λ and equals infinity with probability $1 - \lambda$, we have Calvo time-dependent price setting, while if $F_t^i = 0$ always, we have flexible prices.

The solution to the problem (if it exists) yields two policy functions:

$$\hat{P}_t^{ik*} = P^*(\Theta_t^{ik}) \quad (4)$$

$$I_t^{ik} = I^*(P_{t-1}^{ik*}, \Theta_t^{ik}) \quad (5)$$

The optimal reset price \hat{P}_t^{ik*} is a latent variable, observed only if $I_t^{ik} = 1$. However I_t^{ik} is always observed.

Pricing-to-market

The model allows us to define what we pricing-to-market. Irrespective of whether or not the invoice currency price changes at date t , the observed price in market k , converted to home currency, can always be written as a gross markup over home currency marginal cost:

$$P_t^{ik} = E_t^k P_t^{ik*} \equiv \mu_t^{ik} (1 + \tau^{ik}) C_t^i \quad (6)$$

Let $k = H$ indicate the home market. There is pricing-to-market if μ_t^{ik}/μ_t^{iH} , the relative markup defined by (6), is not constant, and in particular, if μ_t^{ik}/μ_t^{iH} comoves with E_t^k , the nominal exchange rate between the home market and market k .

Flexible prices

If prices are flexible, the observed home currency price is always equal to the optimal reset price, expressed in home currency. Whether or not there is pricing-to-market, and the nature of comovement between μ_t^{ik}/μ_t^{iH} and E_t^k , then depends on the properties of $q(\cdot)$ and on the sensitivity of marginal cost to exchange rate shocks. In the Appendix, we review this dependence. We also work through the CES-monopolistic competition example (where markups are constant and there is no pricing-to-market) and an example based on linear demand with monopolistic competition (where there is pricing-to-market). Other mechanisms which have been used to generate pricing-to-market under flexible prices include the Translog and Kimball demand systems, additive distribution costs, non-monopolistic competition, and search and matching frictions which induce dynamic considerations even in the absence of nominal price stickiness.⁹

⁹See e.g. Bergin and Feenstra (2001), Gopinath, Itskhoki and Rigobon (2010), Gust, Leduc and Sheets (2009), Atkeson and Burstein (2008) and Drozd and Nosal (2012a).

Sticky prices

Under sticky prices, default pricing-to-market, i.e. what happens when prices do not change in either market, depends on the choice of invoice currency. Under local currency price stickiness, by default there is pricing-to-market, and μ_t^{ik}/μ_t^{iH} moves one-for-one with movements in E_t^k . Under home currency price stickiness, by default there is no pricing-to-market. But when prices are sticky, the nature and degree of pricing-to-market *conditional on price adjustment* depends not just on the properties of $q(\cdot)$ and cost sensitivity to exchange rate shocks, but also on the invoice currency, the nature and degree of price stickiness, and on the process for shocks hitting the producer. In the Appendix (section A.8), we illustrate this dependence by solving for the optimal reset price in two different special cases: CES demand, monopolistic competition and Calvo sticky prices, and linear demand, monopolistic competition and Calvo sticky prices.

Our empirical goal is to document the actual behavior of optimal reset prices, and in particular, how the relative markups implied by optimal reset prices comove with movements in nominal exchange rates. To do so, we estimate reduced form approximations to both (4) and (5). We refer to these as the “intensive margin” and the “extensive margin,” and now describe our precise strategy for each in turn.

4 Empirical strategy

4.1 Intensive margin

We start by assuming that the log optimal reset price, converted to home currency, can be approximated by the following expression (lower case letters refer to variables expressed in logs):

$$\hat{p}_t^{ik} = \alpha^{ik} + c_t^i + \beta^k \mathbf{z}_t^k + \varepsilon_t^{ik} \quad (7)$$

α^{ik} captures all time-invariant effects that are specific to plant-product pair i and market k , including time-invariant components of cost and the markup. c_t^i captures all effects that are common across markets within a plant-product pair i and a period t . This includes the time-varying component of marginal cost, which by assumption is common across all markets

within the plant-product pair. It also includes time-varying components of the markup that comove perfectly across markets. $\beta^k \mathbf{z}_t^k$ captures the sensitivity of the markup to time-varying aggregate variables which may differ across markets. It is precisely these sensitivities we are interested in. ε_t^{ik} captures shocks which are idiosyncratic to plant-product pair i and market k .

In principle, \mathbf{z}_t^k could include all observed aggregate variables that help the producer to forecast demand and costs over the horizon for which the price could be sticky. In our baseline analysis, we reduce \mathbf{z}_t^k to a single element, e_t^k . In this case β_e^k , the coefficient on e_t^k , combines the effect on the markup in market k of the nominal exchange rate and of other variables that are correlated with the nominal exchange rate. We choose this as the baseline specification because the exchange rate is precisely measured, is sufficient to identify time-varying relative markups, and in any case we will not be giving a structural interpretation to the estimated coefficients. We also examine robustness to using a richer set of variables in \mathbf{z}_t^k that includes the aggregate price level and real aggregate expenditure in the target market.

Because we observe price changes with greater precision than price levels, we estimate (7) in differences. We condition on invoice currency prices changing at time t and difference (7) by s_t^{ik} periods, where s_t^{ik} is the number of periods since the local currency price was last changed by plant-product pair i in market k . By differencing over the period between two price changes, we guarantee that the observed price change is equal to the change in the optimal reset price: $\Delta_{s_t^{ik}} p_t^{ik} = \Delta_{s_t^{ik}} \hat{p}_t^{ik}$. This yields the baseline estimating equation:

$$\Delta_{s_t^{ik}} p_t^{ik} = \theta_{t,s_t^{ik}}^i + \beta_e^k \Delta_{s_t^{ik}} e_t^k + \eta_{t,s_t^{ik}}^{ik} \quad (8)$$

where $\theta_{t,s_t^{ik}}^i$ is a plant-product-month-age-of-price fixed effect, and $\eta_{t,s_t^{ik}}^{ik} = \Delta_{s_t^{ik}} \varepsilon_t^{ik}$. The estimating equation with a richer set of controls is analogous. If prices change in both markets at date t , and $s_t^{i,IRL} = s_t^{i,UK}$, then $\Delta_{s_t^{ik}} c_t^i$ is absorbed by the fixed effect $\theta_{t,s_t^{ik}}^i$, and β_e^k , the parameter of interest, can be identified. This is illustrated in Figure 2. To be clear about how many observations are used to identify our estimates, we include only observations that fulfil the criterion of double synchronization of price changes in the regression sample.

Selection bias is potentially an issue in this setting, both because we use only price changes, and because those price changes are synchronized. We return to this issue after addressing the extensive margin of price adjustment.

In estimating (8), we exclude observations where the log change in the home currency price is greater than 2 in absolute value. We weight by plant-market-year sales shares, constructed using CIP data as described in the Appendix and cluster standard errors at the plant level.

4.2 Extensive margin

Suppose there is a fixed cost of changing prices. Producer i changes its local currency price in market k from P_{t-1}^{ik*} to \hat{P}_t^{ik*} if the gain in expected profit outweighs F_t^i , the cost of making the change. Remember that by definition of s_t^{ik} (the number of periods since the local currency price was last changed), the price inherited from $t - 1$ is exactly the optimal reset price at date $t - s_t^{ik}$: $P_{t-1}^{ik*} = \hat{P}_{t-s_t^{ik}}^{ik*}$. We assume that the policy function (5) can be approximated by a threshold rule of the following form:

$$I_t^{ik} = \begin{cases} 0 & \text{if } \underline{\rho}_t^i < \hat{p}_t^{ik*} - \hat{p}_{t-s_t^{ik}}^{ik*} < \bar{\rho}_t^i \\ 1 & \text{otherwise} \end{cases} \quad (9)$$

This requires that the gain in expected profit be increasing in the absolute value of the log difference between the optimal reset price and the price inherited from the previous period. If there are indeed fixed costs of changing prices, the assumption that the rule takes this form is not particularly restrictive, except in so far as the upper and lower thresholds are assumed to be the same across markets within a plant and a period.¹⁰

Given assumption (7) about the form of \hat{p}_t^{ik} , we have:

$$\hat{p}_t^{ik*} - \hat{p}_{t-s_t^{ik}}^{ik*} = \Delta_{s_t^{ik}} \hat{p}_t^{ik} - \Delta_{s_t^{ik}} e_t^k = \theta_{t,s_t^{ik}}^i + (\beta_e^k - 1) \Delta_{s_t^{ik}} e_t^k + \eta_{t,s_t^{ik}}^{ik}$$

This implies that we can substitute in and estimate (9) to obtain information on the value

¹⁰Random fixed costs and hence random thresholds would be necessary to explain the observed frequency of small invoice currency price changes in a purely fixed cost model of price adjustment.

of β_e^k : in particular, on whether $\beta_e^k \neq 1$.

Although we could explicitly estimate (9) allowing for the two thresholds, we choose instead to estimate separate equations for price increases and price decreases. There are two advantages to this approach. First, we can avoid placing any structure on the $\theta_{t,s_t^k}^i$. Second, as we discuss presently, the results from this approach can be interpreted under the alternative of time-dependent pricing as well as under state-dependent pricing. The drawback is that we cannot test whether the timing of price changes is endogenous.¹¹

The conditional logit is a convenient estimator. It allows us to condition out the $\theta_{t,s_t^k}^i$ and the cutoffs of the inaction region, i.e. treat them analogously to fixed effects in a linear setting, while still using a limited dependent variables approach. To implement this, we assume that $\eta_{t,s_t^k}^{ik}$ has a logistic distribution. We can then write:

$$\Pr[\text{increase}] = \Lambda \left(\psi_{t,s_t^k}^i + (\beta_e^k - 1) \Delta_{s_t^k} e_t^k \right) \quad (10)$$

$$\Pr[\text{decrease}] = \Lambda \left(\phi_{t,s_t^k}^i - (\beta_e^k - 1) \Delta_{s_t^k} e_t^k \right) \quad (11)$$

where $\Lambda(z) = \exp(z) / [1 + \exp(z)]$, $\psi_{t,s_t^k}^i = \theta_{t,s_t^k}^i - \bar{\rho}_t^i$ and $\phi_{t,s_t^k}^i = \underline{\rho}_t^i - \theta_{t,s_t^k}^i$. We construct the dependent variable in (10) by coding increases in destination currency prices as a one, while all other observations (decreases and no change) are coded zero. The dependent variable in (11) is constructed analogously.

The conditional logit uses plant-product pair i at date t to identify the coefficient on exchange rate changes only if two conditions are fulfilled. First, the preceding price change must have been synchronized across the Irish and UK markets ($s_t^{i,IRL} = s_t^{i,UK} = s_t^i$). Second, at date t there must be a price increase (or decrease) for at least one *but not all* quote-lines within the plant-product pair for which the last price change was synchronized. Two examples are illustrated in sets A and B in Figure 2. The conditional logit estimator does not make use of episodes where past price changes were synchronized, but either no prices change at t , or all prices change in the same direction. We test the robustness of our results by estimating the same equations using a linear probability model (i.e. $\Lambda(\cdot)$ linear). In all cases, we weight by plant-market-year sales, constructed as described in the Appendix, and

¹¹In any case, the results of such a test would be ambiguous in the case where the true $\beta_e^k = 1$.

cluster standard errors at the plant level.

Interpretation under time-dependent pricing

While the narrow motivation we provide for estimating (10) and (11) is based on assuming a fixed cost of price changes, the resulting estimates can be interpreted more generally. The intuition is as follows. For plant-product pairs where the last price change was synchronized across the two markets, the change in log marginal cost between t and the previous price change is the same in both markets. This implies that in the absence of any price change at t , relative markups must have drifted one-for-one with movements in exchange rates since the last price change (given local currency price stickiness). Conditional on changing prices in one or more markets at date t , as long as prices in the two markets do not change in the same direction, the *relative direction* of price changes, and how the relative direction is correlated with exchange rate movements since the last price change, can provide evidence as to whether producers want to offset ($\beta_e^k < 1$) or exacerbate (i.e. if $\beta_e^k > 1$) the baseline drift. This does not depend on why prices change at t .

The extensive margin exercises systematically examine whether in cases where at least one price changes, but prices do not all move in the same direction, the relative direction of price changes points to a desire to offset or exacerbate baseline markup drift. A zero coefficient on the exchange rate variable in each of the two extensive margin equations is consistent with producers not systematically choosing either to offset or exacerbate the drift. More details on this interpretation are provided in the Appendix (Section A.9). A corollary to this alternative interpretation is that the results from estimating (10) and (11) cannot be interpreted as evidence that the timing of price changes is endogenous, because an “event” is defined not as a price change, but a price change in a particular direction, with a price change in the opposite direction being defined as the event not taking place.

4.3 Selection issues in the intensive margin

We can only use our intensive margin strategy to identify the parameters of interest by conditioning on episodes of double synchronization of price changes. This has the potential

to induce selection bias if pricing is state-dependent. The potential selection problem has two dimensions.

First, if the fixed cost to change prices is paid on a plant-product-market basis, by conditioning on synchronization of price changes, it is more likely that we select episodes where price changes are induced by large accumulated changes in cost than by factors idiosyncratic to a particular market (such as $\Delta_{s_t^{ik}} e_t^k$ or $\eta_{t,s_t^{ik}}^{ik}$). It is not obvious how this should affect our intensive margin estimates. However the extensive margin is unaffected, so by comparing the two, we may have some idea of the degree of bias. It is also worth noting that as Table 3 documents, for whatever reason, synchronization of price changes is the default rather than the exception.

Second, within a plant-product pair and a market, conditioning on episodes where invoice currency prices actually change induces dependence of the conditional distribution of the error term $\eta_{t,s_t^{ik}}^{ik}$ on $\theta_{t,s_t^{ik}}^i + (\beta_e^k - 1) \Delta_{s_t^{ik}} e_t^k$. This is because the error term is truncated (from above in the case of invoice currency price decreases, and from below in the case of invoice currency price increases), and the extent of truncation in each case depends on $\theta_{t,s_t^{ik}}^i + (\beta_e^k - 1) \Delta_{s_t^{ik}} e_t^k$. How this affects our intensive margin estimates of β_e^k depends on a number of factors, including the true β_e^k , the conditional distribution of $\theta_{t,s_t^{ik}}^i$, asymmetry in the upper and lower thresholds for price adjustment, and the relative frequency of price increases and price decreases. We address this in greater detail in the Appendix (Section A.10), but to summarize the discussion there, if the true $\beta_e^k = 1$, conditional on $\theta_{t,s_t^{ik}}^i$, the distribution of $\eta_{t,s_t^{ik}}^{ik}$ does not depend on $\Delta_{s_t^{ik}} e_t^k$, and there is no bias due to within-plant-product-market selection. Again, this type of bias does not affect the extensive margin.

We do not attempt a formal selection correction, but instead compare results from the intensive margin with those from the extensive margin. If both sets of results point to similar values for β_e^k , this is suggestive that selection bias is not a problem in the intensive margin estimates.

5 Results

5.1 Intensive margin

Table 4 reports our baseline estimates of (8). The estimate of β_e^{UK} is almost exactly equal to one, significantly different from zero, and not significantly different from one, implying pricing-to-market, of a particular form. Conditional on adjusting prices, producers choose to have the ratio of the markup in the UK market to the markup in the home market increase one-for-one with depreciations of the Euro against Sterling, and increase one-for-one with appreciations of the Euro against Sterling.

Figure 3 illustrates one dimension of the variation that identifies β_e^{UK} . It plots the log change in the Sterling price in the UK against the log change in the home currency price in Ireland for the observations used in the baseline regression. An observation is a plant-product-month where prices change in both markets, and where the previous price change was also synchronized across markets.¹² Note that if $\beta_e^{UK} = 1$, the log change in the Sterling price in the UK should on average equal the log change in the home currency price in Ireland. The data does indeed line up around the 45° line. This figure also illustrates the fact that $\hat{\beta}_e^{UK} \simeq 1$ for both positive and negative price changes, and for price changes of different sizes.

A different dimension of the variation identifying the coefficient is illustrated in Figure 4, which plots histograms of the log changes in local currency prices and the log changes in exchange rates for the estimation sample. The standard deviation of price changes is large. Despite large cumulative swings in the exchange rate over the course of the sample period, on average the (absolute) size of exchange rate changes in the interval from one price change to the next is smaller than the size of the corresponding price changes.¹³

Producers' incentives to realign relative markups when they change prices may be related to how long they expect to wait before changing prices again, to the elasticity of the residual demand curve they face, and to the cost process they face.¹⁴ We explore these possibilities

¹²For observations where there is more than one price quote in a market, the mean log change across quotes within that market is plotted.

¹³This is similar to the findings of Burstein and Jaimovich (2009).

¹⁴We illustrate some of these in the calibration exercise in the Appendix.

by estimating (8) on many different subsamples of the data. The results are reported in the Appendix. We split the data by median frequency of price adjustment at the level of the plant-product pair, by actual age of prices, by type of good (following classifications by Vermeulen et al. (2007) and Rauch (1999)),¹⁵ by plant size as measured by number of employees, by quartiles of the share of sales exported to the UK, by quartiles of price-cost margins, by quartiles of the share of materials imported from the UK in variable cost, by whether or not the plant reports some sales (in any market) to related parties, and by ownership. In almost all cases, the estimate of β_e^{UK} is significantly different from 0 and not significantly different from 1.

Overall, the evidence from the intensive margin points to a particular form of pricing-to-market, where even conditional on adjusting prices in both markets, producers allow the markup in the foreign market relative to the markup in the home market to move one-for-one with movements in exchange rates. This implies that real exchange rates at the plant-product level track nominal exchange rates whether or not prices change. This behavior is pervasive across manufacturing plants and sectors with different characteristics, given local currency invoicing.

5.2 Extensive margin

Table 5 reports the results from estimating equations (10) and (11) on the baseline sample. As indicated in the table, if $\beta_e^{UK} < 1$ (i.e. producers want to offset the default effect of exchange rate movements on relative markups given local currency invoicing), we would expect a negative coefficient on the exchange rate change in the case of price increases, and a positive coefficient in the case of price decreases. Both in the case of price increases and price decreases, the baseline estimate of the coefficient on the exchange rate change is not significantly different from zero. Based on the extensive margin, we cannot reject the hypothesis that $\beta_e^{UK} = 1$. Note that logit estimation does not directly identify the scale of the parameters, but only whether they are positive, negative, or not significantly different from zero.

¹⁵The details of these classifications are provided in the Appendix.

We estimate on the same cuts of the data as in the intensive margin case. The results are reported in the Appendix. To summarize, the evidence from the different cuts of the data generally points to a value of β_e^{UK} that is not significantly different from 1. There are some exceptions, but in no case do the estimates from both the price increases equation and the price decreases equation point simultaneously to a value of β_e^{UK} significantly less than 1 or significantly greater than 1. Overall, the results from the extensive margin are consistent with those from the intensive margin: We cannot reject the null hypothesis that $\beta_e^{UK} = 1$

5.3 Selection bias

As discussed in section 4, the estimates of β_e^{UK} obtained using the intensive margin may be contaminated by selection bias. However since the extensive margin exercises yield similar implications for pricing-to-market to the intensive margin, and moreover, are consistent with $\beta_e^{UK} = 1$, the value for which the second type of selection is unlikely to be a problem, this reassures us that our intensive margin results are unlikely to be driven by selection. More details are provided in the Appendix (Section A.10).

5.4 Pricing-to-market under home currency invoicing

Our ability to examine pricing-to-market by producers invoicing in home currency is hampered by the fact that, since we do not generally know the destination of these exports, we do not know what is the correct exchange rate to use as the independent variable. But 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 survey responder. We use the sub-sample where the destination can be identified as the UK to implement both extensive and intensive margin exercises separately using exports to the UK invoiced in Sterling, and exports to the UK invoiced in Euros. We restrict the sample to the period 2003-2005, as the further back we go, the smaller and more selected is the set of quote-lines.¹⁶

Table 6 reports the results from estimating (8) separately for the local currency invoicing sub-sample and the home currency invoicing sub-sample. β_e^{UK} is imprecisely estimated in

¹⁶We do not use identified markets other than the UK, as the majority of observations in this category are for Euro-zone markets, against which there is no exchange rate variation in the years 2003-2005.

both cases. In the local currency case, the point estimate is bigger than 1, significantly different from zero at the 5% level, but not significantly different from one. In the home currency case, the point estimate is just less than 1 and it is not significantly different from zero or one at the 5% level.

The extensive margin exercise must be modified in the case of home currency invoicing so that the indicator variables are based on changes in the home currency price rather than changes in the destination currency price. This also modifies the interpretation of the coefficient on the exchange rate change. Under home currency invoicing, it is equal to β_e^{UK} in the case of price increases, and $-\beta_e^{UK}$ in the case of price decreases. The results from estimating (10) and (11) in the case of local currency invoicing, and their counterparts in the case of home currency invoicing are reported in Table 7. In the local currency invoicing case, the estimated coefficients on exchange rate changes are not significantly different from zero, so we cannot reject the null hypothesis that $\beta_e^{UK} = 1$. In the case of home currency invoicing, the estimated coefficients are not significantly different from zero, implying that in this case, we cannot reject the null hypothesis that $\beta_e^{UK} = 0$.

Taken together, the results from the intensive and extensive margin are consistent with pricing-to-market behavior differing by choice of invoice currency.

5.5 Robustness

We now describe the results of a large number of robustness checks. The relevant tables and the results of yet more robustness checks can be found in the Appendix.

Potential failures of the identifying assumption

Our identification of markup behavior relies on the assumption that within plant-product pairs, relative marginal cost is constant across the Irish and UK markets. There are several ways in which this assumption might fail, and we examine some of these.

First, the products sold in the two markets, though falling in the same product category and produced in the same plant, might differ in such a way that relative marginal cost could move around with (say) exchange rate movements. To reduce the likelihood that this could

happen, we perform the intensive and extensive margin exercises on a sample where the matching criterion - that the detailed description of the item be identical across the two markets - is more demanding than in the baseline analysis. The detailed description of the item may include brand name, product variety, quantities etc. This reduces the sample size, particularly for the extensive margin exercises. The intensive margin results (Table A.15) are unchanged under this more demanding matching strategy. On the extensive margin (Table A.16) we cannot reject the null of $\beta_e^{UK} = 1$ for increases, but we reject (in favor of $\beta_e^{UK} < 1$) for decreases.

Second, if delivery charges are included in the price, and they vary across the home and export market, this could lead to comovement of relative marginal cost with exchange rates, violating the identifying assumption. Plants participating in the pricing survey are instructed to exclude delivery charges that are itemized separately on the invoice. Implicitly, if delivery charges are not itemized separately, they are included in the price. For the set of quote-lines that are present in the last cross-section, we observe what the plant reports in the “trading terms” field on initial participation in the PPI survey. Sometimes plants report that delivery charges are, or are not, included in the reported price. We estimate our baseline specifications conditioning on delivery charges being included and delivery charges not being included. The results are reported in Tables A.17 and A.18. The samples involved are small, but the results are not systematically different from the baseline for either group.

A richer set of independent variables

We implement both intensive and extensive margin exercises using a richer set of independent variables, including the CPI and real demand in the target market. The coefficients on these variables are not very precisely estimated. Their inclusion does not change the baseline intensive margin result that the point estimate of β_e^{UK} is very close to 1, not significantly different from 1, and significantly different from 0. The extensive margin result that β_e^{UK} is not significantly different from 1 is also unchanged. The results are reported in Tables A.19 and A.20.

Longer horizons

Our numerical exploration of pricing-to-market conditional on price changes in a Calvo sticky price environment (reported in the Appendix Section A.8) suggests that β_e^{UK} may be time-horizon specific, and declining in the length of the time horizon. To investigate what happens at longer horizons, we make use of the first and last pair of synchronized price changes observed for each plant-product pair with quote-lines in both home and UK markets. We regress the log change in home currency price over this horizon on the contemporaneous log change in the relevant exchange rate. This allows us first, to examine price changes over longer time horizons than the baseline approach (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 sample of price changes are reported in Table A.21. 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. In the split by 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. This evidence suggests that in the case of local currency invoicing, the long run elasticity of relative markups with respect to nominal exchange rates may be lower than the short run elasticity, however it is not conclusive.

Dynamic adjustment

If there are strategic complementarities in pricing and price adjustment is not synchronized across competitors, it is possible that adjustment to exchange rate movements could be spread out over several price changes. Over and above the long horizon results just described, we investigate this possibility by including lagged exchange rate changes in the baseline estimating equation. The horizon over which each lag is taken is the same as the horizon over which the original difference is taken. We include one and two lags. Results are reported in Table A.22. The coefficients on the lagged exchange rate changes are never significantly

different from zero.

Other markets

An obvious question is whether the pricing-to-market behavior we identify is specific to the two markets we focus on. To test this, in addition to our baseline sample, we make use of all cases of parallel pairs of price quotes for home sales and for exports invoiced in non-Sterling foreign currencies. We use this broader sample to estimate (8), (10) and (11), where the independent variable is the change in the domestic exchange rate with the invoice currency.¹⁷ This need not be the same as the currency of the destination market, as, for example, if sales to Japan are invoiced in US dollars. The results (reported in Tables A.23 and A.24) are very similar to the baseline. On the intensive margin, the point estimate of β_e is very close to 1, not significantly different from 1, and significantly different from zero. On the extensive margin, we cannot reject the null hypothesis that $\beta_e = 1$.

What happens when we don't control for costs?

As noted in the Introduction, our work is related to Gopinath, Itskhoki and Rigobon (2010), who examine exchange rate passthrough conditional on price changes for US dollar-invoiced US imports. While they are unable to control for costs as we do, because they only observe prices in one market, we can perform an exercise very similar to theirs using our data. We restrict the sample to Sterling-invoiced exports, and estimate (8) with quote-line fixed effects instead of plant-product-month-age-of-price fixed effects. The results are reported in Table A.25. We find a point estimate of the coefficient on the exchange rate change equal to 0.76, with a standard error of 0.14. The point estimate is the same as that of Gopinath, Itskhoki and Rigobon when they pool across all import origins, though our estimate is considerably less precise than theirs.¹⁸

¹⁷We do not allow β_e to depend on the market in this case.

¹⁸Their estimates for different origin countries vary from 0.54 for Spain to 0.99 for Belgium.

6 Discussion and conclusions

The central contribution of this paper is to provide clean evidence that producers who sell in both home and export markets engage in pricing-to-market, of a particular type. When prices are sticky in local currency, even conditional on changing prices, producers allow the markup in the foreign market relative to the home market to increase one-for-one with depreciations of the home currency, and decrease one-for-one with appreciations of the home currency. Our findings apply to a wide range of industrial sectors, and across plants with different characteristics. The evidence on pricing-to-market for prices that are sticky in home currency is less clear-cut, but consistent with the opposite behavior: relative prices across markets do not co-move with nominal exchange rates. Relative to the past literature, our results point to an extreme form of pricing-to-market for producers invoicing in local currency. But though extreme, our results are similar to those of Burstein and Jaimovich (2009), a paper with a clean experiment comparable to ours. Moreover, the degree of exchange rate pass-through implied by our data is similar to that found by Gopinath, Itskhoki and Rigobon (2010). These similarities raise our confidence in the reliability of our results.

Our results raise two obvious questions. First, can they be rationalized in the context of a reasonable model? Second, what do they imply for the real exchange rate puzzles mentioned in the Introduction? As we mention in Section 3, in a flexible-price environment, the degree of pricing-to-market depends on properties of the residual demand faced by producers, along with the sensitivity of costs to exchange rates. We show in the Appendix (Section A.8) that in order to rationalize our findings in a flexible-price environment, one would have to assume that firms face a very strongly kinked residual demand. However our data is not generated by a flexible-price environment. In the Appendix we also show that when prices are sticky, the optimal degree of pricing-to-market conditional on price adjustment depends on other factors, including the choice of invoice currency, the nature of price stickiness, the expected frequency of price adjustment, and the joint process for demand and cost shocks hitting the firm. This suggests that it may be possible to rationalize our findings in the context of relatively standard demand and market structures.¹⁹

¹⁹Nakamura and Steinsson (2011) also provide a possible rationalization: Price adjustment may occur disproportionately around product replacements. Our data exhibits very little product churn relative to the

As regards the implications of our findings, there is a growing literature that explores quantitatively the impact of mechanisms such as non-monopolistic competition and customer markets in generating pricing-to-market in flexible price settings. Key papers in this literature include Atkeson and Burstein (2008) and Drozd and Nosal (2012a). Our results provide a way to discipline the behavior of markups in these models. In the light of our results, it may be worth trying to extend these models to allow for sticky prices, as our findings on pricing-to-market for prices invoiced in home currency compared to local currency point suggestively, if not conclusively, to a role for price stickiness in determining the behavior of prices even conditional on adjustment.

Given our findings on the behavior of reset prices, it may also be worthwhile to revisit the contribution of price stickiness to explaining real exchange rate puzzles, along the lines of Bilal, Klenow and Malin (2011). The literature on the topic of sticky prices and real exchange rate puzzles is substantial (See Bergin and Feenstra (2001), Chari, Kehoe and McGrattan (2002), Midrigan (2007), Steinsson (2008), Landry (2009), Carvalho and Nechio (2010), among others), but has had difficulty in matching the persistence of real exchange rates under reasonable frequencies of price adjustment. It is possible that under realistic behavior of reset prices along the lines we document, these models could do a better job of matching real exchange rate persistence.

References

- [1] Atkeson, A. and A. Burstein (2008), “Pricing-to-Market, Trade Costs and International Relative Prices,” *American Economic Review* 98 (5), 1998-2031.
- [2] Bergin, P. and R. Feenstra (2001), “Pricing-to-Market, Staggered Contracts and Real Exchange Rate Persistence,” *Journal of International Economics* 44, 333-359.
- [3] Bilal, M., P. Klenow and B. Malin (2011), “Reset Price Inflation and the Impact of Monetary Policy Shocks,” *American Economic Review*, forthcoming.

US import and export price micro-data they work with.

- [4] Burstein, A. and N. Jaimovich (2009), "Understanding Movements in Aggregate and Product-Level Real Exchange Rates," mimeo.
- [5] Carvalho, C. and F. Nechio (2010), "Aggregation and the PPP Puzzle in a Sticky-Price Model," *American Economic Review* 101 (6), 2391-2424.
- [6] Chari, V.V., P. Kehoe and E. McGrattan (2002), "Can Sticky Price Models Generate Volatile and Persistent Real Exchange Rates?" *Review of Economic Studies* 69, 533-563.
- [7] Dornbusch, R. (1987), "Exchange Rates and Prices," *American Economic Review* 77 (1), 93-106.
- [8] Drozd, L., and J. Nosal (2012a), "Understanding International Prices: Customers as Capital," *American Economic Review*, 102 (1), 364-395.
- [9] Drozd, L. and J. Nosal (2012b), "Pricing-to-Market in Business Cycle Models," mimeo.
- [10] Engel (2002), "Expenditure Switching and Exchange-Rate Policy," in M. Gertler and K. Rogoff, eds., *NBER Macroeconomics Annual 2002, Vol. 17*, Cambridge: NBER.
- [11] Goldberg, P. and R. Hellerstein (2011), "A Structural Approach to Identifying the Sources of Local-Currency Price Stability," *Review of Economic Studies*, forthcoming.
- [12] Goldberg, P. and M. Knetter (1997), "Goods Prices and Exchange Rates: What Have We Learned?" *Journal of Economic Literature* 35 (3), 1243-1272.
- [13] 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.
- [14] Gopinath, G., O. Itskhoki and R. Rigobon (2010), "Currency Choice and Exchange Rate Passthrough," *American Economic Review*, 100 (1), 304-336.
- [15] Gust, C., S. Leduc and N. Sheets (2009), "The Adjustment of Global External Balances: Does Partial Exchange Rate Pass-Through to Trade Prices Matter?" *Journal of International Economics* 79 (2), 173-185.

- [16] Klenow, P. and B. Malin (2011), “Microeconomic Evidence on Price Setting,” in B. Friedman and M. Woodford, eds., *Handbook of Monetary Economics*, Elsevier.
- [17] Knetter, M. (1989), “Price Discrimination by U.S. and German Exporters,” *American Economic Review* 79 (1), 198-210.
- [18] Krugman, P. (1987) “Pricing to Market When the Exchange Rate Changes,” in S. W. Arndt and J. Richardson, eds., *Real Financial Linkages Among Open Economies*, London: MIT Press.
- [19] Landry, A. (2009), “Expectations and Exchange Rate Dynamics: A State-Dependent Pricing Approach,” *Journal of International Economics* 78, 60-71.
- [20] Midrigan, V. (2007), “International Price Dispersion in State-Dependent Pricing Models,” *Journal of Monetary Economics* 54 (8), 2231-2250.
- [21] Nakamura, E. and D. Zerom (2010), “Accounting for Incomplete Pass-Through,” *Review of Economic Studies*, 77 (3), 1192-1230.
- [22] Nakamura, E. and J. Steinsson (2011), “Lost in Transit: Product Replacement Bias and Pricing to Market,” *American Economic Review*, forthcoming.
- [23] Rauch, J. (1999), “Networks Versus Markets in International Trade,” *Journal of International Economics* 48, 7-35.
- [24] Steinsson, J. (2008), “The Dynamic Behavior of the Real Exchange Rate in Sticky-Price Models,” *American Economic Review* 98 (1), 519-533.
- [25] 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	73	53	74	46
UK	67	48	63	39
% of sales exported 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
Employees, distribution across plants				
25th percentile	26	7	22	6
50th percentile	59	15	48	13
75th percentile	129	41	110	36

Note: PPI refers to the sample of CIP plants that participate in the PPI sample. CIP refers to the full sample of plants. Information on imports is based on the roughly 90% of the population of plants for which comparable information is available over the entire sample period. Information on the export currency is based on the roughly 95% of the population of plants where information is available over the entire sample period.

Table 2: Weighted Mean Frequency of Price Adjustment, Invoice Currency

	Unadjusted	Exit adjusted
total	0.16	0.18
Destination market		
home	0.19	0.20
export	0.14	0.17
Invoice currency for exports		
IEP/ EUR	0.11	0.13
STG	0.16	0.17

Note: The period covered is January 1995 - December 2005. The weighted mean frequency of price adjustment is calculated as $\sum_{t=1}^T \sum_i \omega_t^i I_t^i / \sum_{t=1}^T \sum_i \omega_t^i N_t^i$ where I_t^i is an indicator variable, equal to 1 if i 's price changed at t , and N_t^i equals 1 if i was present in the sample at t , whether or not its price was changed. Observations are weighted by the relevant plant's sales in the relevant market (home or export) expressed as a share of total within-sample sales in the year corresponding to date t . This implies that each month is given equal weight in calculating frequencies. Exit adjustment treats quote-line exit like a price change, i.e. I_t^i is set equal to 1 if the quote-line is no longer present in the sample at date $t + 1$.

Table 3: Synchronization of Price Changes Within Plant-Product Pairs

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

Note: The first column reports the fraction of plant-product-months where there is more than one price quote, and at least one of those price changes, where price changes refer to price changes in invoice currency. The second, third and fourth columns report the fraction of these cases that fall into each of three possible categories. The 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.

Table 4: Intensive Margin of Price Adjustment

$\Delta_{s^{ik}} e_t^k$	R ² -adj.	N	f.e.	clust
β				
1.01 (0.09) **	0.67	4212	1047	86

Note: Estimation method is OLS. Dependent variable is log change in home currency price since last price change. Regression includes a constant (coefficient and s.e. not reported) and the full set of plant-product-month-age-of-price fixed effects. Observations are weighted by sales at the level of the plant-market. Standard errors are clustered at the plant level. Standard errors are 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 5: Extensive Margin of Price Adjustment

Probability of a price increase					Probability of a price decrease						
$\Delta_{s_{it}^k} e_t^k$	ps-R ²	N	f.e.	clust	$\Delta_{s_{it}^k} e_t^k$	ps-R ²	N	f.e.	clust		
$\beta - 1$					$1 - \beta$						
0.59	(3.61)	0.00	4873	921	129	1.62	(4.99)	0.00	4564	875	103

Note: Dependent variable is indicator for increase or decrease in invoice currency price. This means in the case of price increases, the indicator equals one if the invoice currency price is increased, equals zero if the invoice currency price remains unchanged or is decreased. The case of price decreases is analogous. Estimator is conditional logit, conditioning on plant-product-month-age-of-price fixed effects. Observations are weighted by sales at the plant-market level. 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. A pseudo-R-squared is reported. The number of fixed effects indicates the number of plant-product-months used to identify the coefficient on exchange rates. The number of clusters indicates the number of plants used to identify the coefficient on exchange rates.

Table 6: Intensive margin by choice of invoice currency

Invoice curr.	$\Delta_{s_{it}^k} e_t^k$	R ² -adj.	N	f.e.	clust
destination	2.52 (1.22)*	0.68	741	235	40
home	0.97 (0.57)*	0.66	745	236	41

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. 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: Extensive margin by choice of invoice currency

		$\Delta_{s_{it}^k} e_t^k$	ps-R ²	N	f.e.	clust
Destination currency invoicing						
Pr(increase)	$\beta - 1$	10.97 (7.92)	0.00	586	134	30
Pr(decrease)	$1 - \beta$	-6.76 (11.57)	0.00	552	131	25
Home currency invoicing						
pr(increase)	β	22.96 (20.17)	0.02	640	151	42
pr(decrease)	$-\beta$	-13.01 (14.57)	0.00	503	129	29

Note: Sample is restricted to matched pairs of home sales and foreign sales where destination is identified as the UK in November 2006. Dependent variable is indicator for increase/decrease 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.




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




Figure 2: Type of price change episodes used to identify the intensive and extensive margins




Figure 3: Scatter plot of data used to identify the intensive margin

Note: Figure plots the log change in the Sterling price in the UK against the log change in the home currency price in Ireland. An observation is the change in prices for a plant-product-month where prices change in both markets and where the previous price change was also synchronized across markets. For observations where there is more than one price quote in a market, the mean log change across quotes within that market is plotted.




Figure 4: Size of price and exchange rate changes used to identify the intensive margin

Note: Figure plots histograms of the log changes in destination currency prices and the log changes in exchange rates for the sample used in the intensive margin estimation.