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

Doireann Fitzgerald[†] Stefanie Haller[‡]

December 2011

Abstract

We document pricing-to-market by producers who sell the same product to buyers in two markets that are segmented by variable exchange rates. We can cleanly identify desired relative markup responses to exchange rate movements because we observe prices charged by the same plant for the same product to buyers from the two different markets at a monthly frequency. The matched price quotes allow us to use fixed effects to control for unobserved marginal cost changes. The high frequency allows us to focus only on episodes where prices change, hence separating out desired markup variation from the default behavior of relative markups due to price stickiness. For prices invoiced in destination currency, we find that desired relative markups move one-for-one with exchange rate changes. This implies that producers invoicing in destination currency engage in an extreme form of pricing-to-market, with relative markups inheriting 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, 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 refers to the situation where for the same good, producers choose markups that vary across markets 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. The goal of this paper is to exploit a unique set of data on producer prices to provide clean evidence on this behavior across a broad range of sectors, making only a bare minimum of assumptions about the nature of the environment in which the producers operate. In our data, as is typical for producer prices, prices are sticky. We distinguish the default behavior of markups, due solely to price stickiness, from desired pricing-to-market, where conditional on price adjustment, new prices may be chosen such that relative markups co-move with exchange rates. We estimate the desired co-movement of relative markups across markets with nominal exchange rates in the case where prices are sticky in destination currency.

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. On the producer price survey form, respondents are encouraged to provide matched price quotes in home and export markets for product categories sold to both markets. We focus on plants reporting parallel price series for sales of the same product to Ireland and the UK, two markets segmented by variable exchange rates. We restrict attention to matched price series where export prices are invoiced in Sterling, the predominant case for sales to the UK market, as we must use the currency of invoicing to identify the destination market for exports at the level of individual price quotes. We assume that within a plant-product pair, relative marginal cost across markets is invariant to movements in exchange rates, even if the level of marginal cost does move with exchange rates. Under this assumption, having matched price quotes in the two markets allows us to cleanly identify producers' desired co-movement of relative markups with exchange rates.

We characterize the response of desired relative markups to exchange rate changes using two different approaches that rely on episodes where prices change. First, since price adjustment is infrequent, we estimate 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. Second, we examine by how much 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 changing prices in both markets simultaneously. In both cases, by conditioning on the price changes having been previously synchronized across markets, we can control for changes in marginal cost using fixed effects, hence isolating the co-movement of relative markups with exchange rates. This strategy makes use of the fact that synchronization of price changes across markets within plant-product pairs is common in our data.

These two approaches to characterizing the response of desired relative markups to exchange rate changes make use of overlapping but non-identical subsamples of the data. Using both approaches we find that producers desire a one-for-one co-movement of relative markup with nominal exchange rates. Given that we focus on the case of destination currency invoicing, this implies a desired co-movement of relative markups with exchange rates that is the same as the default co-movement when prices do not change. This is an extreme form of pricing-to-market, with relative markups inheriting the random walk behavior of exchange rates.

Within the sample, our findings are invariant to differences in pricing behavior at the level of the plant-product pair, such as the length of the time interval between consecutive price changes and the median frequency of price changes. They are also invariant to differences in product characteristics. There is some weak evidence that smaller plants may engage in a less extreme form of pricing-to-market than larger plants. But otherwise, pricing-to-market behavior is very similar across different categories of plants. It is the same 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, and whether the plant is domestic or foreign-owned.

Our finding of a positive desired elasticity of relative markups with respect to movements in exchange rates is not surprising. As already noted, when prices are sticky, the default co-movement of relative markups with exchange rates under destination currency invoicing is equal to 1. In contrast, the default co-movement under home currency invoicing is equal to 0. Given that the invoice currency is a choice variable, it is natural to expect the desired elasticity of relative markups with respect to exchange rates to be closer to 1 than to 0 in cases where the producer chooses to invoice in destination currency. For a small sub-sample of the data, we can identify the destination market for exports without relying on destination currency invoicing. In this sample, we find evidence that the desired elasticity of relative markups with respect to exchange rate changes for prices invoiced in home currency is not significantly different from 0, the default elasticity under home currency invoicing. This suggests that selection on invoice currency choice is important to understanding our results.

It is somewhat surprising that the degree of desired pricing-to-market we estimate is so extreme. In evaluating this result, it is important to bear several things in mind. First, the absolute size of the exchange rate changes we use to identify markup elasticities is small, considerably smaller than the absolute size of the corresponding price changes. We cannot test for nonlinearity in responses because we do not observe extreme movements in exchange rates over the course of the sample. Second, and related to this point, price-cost margins for the plants in our data are large. This suggests that competitive pressures may be insufficient to eliminate market segmentation. Moreover, it implies that the risks of producers being driven to negative profits by small movements in exchange rates are slim. 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 market. Finally, we present some evidence that over horizons of two years or more, when prices are allowed to change multiple times, producers may adjust relative markups to offset to some degree the drift due to movements in nominal exchange rates (i.e. the degree of pricing-to-market may be less extreme). Connected with this point, it is important to remember that the plants in our data operate in a sticky-price environment. We estimate desired markup responses to exchange rate movements conditional on the stickiness of prices. When prices are sticky, strategic complementarities in pricing behavior may lead producers to desire limited adjustment of relative markups to movements in exchange rates, even at horizons as long as a year or more.

To summarize, we present evidence that highlights the importance of invoice currency choice as a determinant of exchange rate pass-through. When changes in exchange rates are not very large, it appears that producers who invoice in destination currency do not want to readjust relative markups to eliminate the drift caused by exchange rate movements. Hence, even conditional on price adjustment, relative markups move one-for-one with changes in exchange rates.

Relative to the existing literature on pricing-to-market, we innovate in our ability to cleanly document desired pricing-to-market under price stickiness. We can do this because of a unique combination of features of our data set. First, we have prices collected at the producer side, not the consumer side, for products that are actually traded, not just potentially traded. This allows 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 relative marginal cost is constant across markets over a given time interval, which is crucial to identifying relative markup variation. Third, our data are on prices rather than unit values, and at a sufficiently high frequency to allow us to observe the timing of price changes. We also know the invoice currency. This allows us to disentangle desired pricing-to-market from default pricing-to-market that arises purely due to destination currency invoicing combined with price stickiness. Fourth, our data covers a broad range of industrial sectors. This gives our results general applicability. Finally because we match price and plant data, we can link pricing behavior not just to product characteristics, but also to plant characteristics.

The fixed effects approach to identifying relative markup variation in the context of pricing-to-market was first proposed and implemented by Knetter (1989). We extend Knetter's approach to take account of the fact that prices are sticky, and that we observe the timing of price changes. Knetter and others implement the basic strategy using annual sector-level unit value data and find evidence of what they refer to as 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] The results

[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.

of this literature generally point to a degree of pricing-to-market less extreme than what we find. But as these authors are aware, aggregation across plants, across invoice currencies and over time makes it tricky to interpret their results as evidence that producers desire relative markup variation in response to exchange rate changes. The evidence we provide goes more directly to this point.

Our strategy for identifying desired pricing-to-market contrasts with the structural approach of Goldberg and Hellerstein (2010), applied to beer, and Nakamura and Zerom (2010), applied to coffee. Like us, these authors find evidence of pricing-to-market conditional on prices changing.[2] Relative to their methodology, our approach has the advantage of not requiring strong assumptions about preferences, market structure, cost functions etc. In addition, our data is not restricted to a particular narrowly defined sector. The tradeoff is that the key parameter that we estimate does not have a precise structural interpretation. We see our contribution as a complementary to this 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 Vermeulen et al. (2007), Nakamura and Steinsson (2008) and Goldberg and Hellerstein (2009), who make use of the micro data underlying the producer price indices for various European countries and the US, 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. But they cannot identify this as being due to desired variation in relative markups. Two papers that make use of different data (wholesale prices at the UPC level faced by a supermarket chain operating in the US and Canada), but document behavior of wholesale prices across markets that is fully consistent with what we find are Burstein and Jaimovich (2009) and Gopinath, Gourinchas, Hsieh and Li (2010).

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

[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.

producers face residual demand curves with variable elasticity, as variable elasticities are necessary to generate desired markup variation. 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. An alternative 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. A related paper that focuses on a closed economy setting is Kleshchelski and Vincent (2009). 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. However it is not clear that these models can match quantitatively the very slow adjustment to lagged changes in nominal exchange rates that we document.

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. All plants with 3 or more employees are required to fill in a return. We make use of the CIP data for the years 1995 to 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 (we have some destination and currency invoicing information for export sales), employment, wage bill, materials and energy expenditures (we also have the share of imported materials, and some origin and invoicing information) and share of sales to related parties. Further details are provided in the data appendix.

The second source is the micro data collected for the purpose of constructing the Irish Producer Price Index (PPI). As is standard in other European countries and the US, this data is collected in monthly surveys of plants. The sampling frame for this survey 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. We have access to the monthly 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 (wholesalers who do not produce also participate in the price survey).

The PPI survey methodology is broadly similar to that used in other European countries and the US. First-time participants are asked to provide prices for their main products that are "suitable for pricing each month." As well as prices, they are invited to provide a detailed product description, 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). Each subsequent month, participants receive a form where 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. Discounts and surcharges on the "basic price" are reported separately, and incorporated into the prices we use. In reporting the current month's prices, participants are instructed to exclude delivery charges that are itemized separately on the invoice. There is no requirement to flag prices for transactions between related parties. Participants are instructed to discontinue what we shall refer to as a quote-line (using the terminology of Klenow and Malin (2010)) and replace it with another if the product or terms of sale are no longer available or representative. Further description of the data are available in a data appendix. Survey documents are available at www.cso.ie.

The feature of the PPI survey that is crucial for our purposes (and 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.

The variables in the PPI data that we make use of are plant and product identifiers and detailed descriptions, the price expressed in domestic currency, an 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. Product categories are classified using a system that is unique to the PPI survey. They are defined at a level that is more disaggregated than the 4-digit NACE categorization that is applied to plants, but the level of aggregation is not consistent across sectors.[3] Unlike other micro price data sets such as the US PPI, we do not see missing observations in the middle of quote-lines, nor is there a flag in the data-set for missing values that are filled in by the CSO.

Some variables, including the units to which the price refers and, crucially, the destination market for exports, are reported at the discretion of the respondent. Further, if provided, we observe these variables only for price quotes present in the last cross-section (November 2006) as this part of the data file is overwritten every month. In the absence of a destination market indicator, we use the invoice currency to identify the precise destination market for exports. This means that we cannot examine pricing-to-market for exports invoiced in home currency. It also restricts the destination markets for which we can examine pricing-to-market, since for vehicle currencies such as the US dollar, the currency does not identify the market. This leads us to focus our attention on export price quotes invoiced in Sterling. We confirm that Sterling identifies the destination market as the UK with high probability.[4] Roughly 60% of foreign-currency-invoiced export price quotes are invoiced in Sterling, while most of the other foreign currencies that precisely identify a destination market account for only a handful of price observations. Our identification strategy requires us to focus on plant-product pairs with price quotes in both the home and the destination market, and

[3] We have examined a sample of the product descriptions, and the product definition appears to be at the 6 or 8-digit level.

[4] We check this by examining quote-lines that are present in the last cross section of the price data (November 2006) and where the destination of exports is explicitly identified. Some of these reported destinations are less precise than we would like - for example, "UK/Spain." We classify observations 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 are for destinations other than the UK.

fortunately it is common for plants reporting price quotes in Sterling in the export market also to report prices quotes in the home market for the same product.

Summary statistics on plants

On average, about 5000 plants appear in the CIP in each year, while between 550 and 900 of these additionally participate in the PPI survey. Table 1 reports statistics on a range of plant characteristics, 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.

The first set of statistics illustrates the fact that Irish producers are very open on both the output and the input side. Roughly half of plants export, while over 50% of total industrial output is exported. More than half of plants use imported intermediates, with around 50% of total expenditure on intermediates being expenditure on imports. These data 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 to the bottom line for these plants. 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.

Some facts about plant openness are of particular note. 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.

A substantial fraction of the plants for which we have data are very small, as is natural since we have data on the universe of plants. The median plant in the CIP has fewer than 20 employees, while the median plant for which we have matched price data has between 50 and 60 employees. At the same time, our plants cover a wide variety of sectors, with the biggest concentrations (in value terms) being in Food, Beverages and Tobacco, Chemicals and Electrical Machinery, which together account for 75% of total output.

Finally, we cannot calculate the level of markups, but we can calculate price-cost margins. The measure we use is total sales less total variable cost (wage bill plus materials and fuel expenditures) divided by total sales. These margins are quite substantial - the median is around 30% - somewhat larger from those reported for US industry.[5]

Further summary statistics on the plant data are available in the data appendix.

Summary statistics on prices

Table 2 provides summary statistics on the hierarchical structure of the matched plant-price data set. On average, each participating plant 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 may be reported both because there are 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 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 3 reports the weighted mean frequency of price adjustment (calculated as the fraction of prices that change, weighted by sales at the plant level), for the sample as a whole, for home sales and exports separately, and, for exports by currency of denomination. The frequencies reported are for the adjustment of prices expressed in invoice currency. 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 data 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

[5] e.g. in Domowitz, Hubbard and Petersen (1986).

decreases, but not markedly so, 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 this behavior are reported in Table 4. The first column reports the percentage of plant-product-months with more than one price quote where at least one price changes. These are the cases where there may, or may not, be 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. These cases account for just over a quarter of the sample we are interested in. The corollary is that just under three quarters of the time, episodes of price change tend to be synchronized across Irish and UK quote-lines within plant-product pairs. 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 rate variation

Our sample period covers a long period during which the home currency (first the Irish pound, and then the Euro) depreciates roughly 35% against Sterling, followed by a period during which it appreciates by around 20%. Month-to-month fluctuations are substantially smaller. This is illustrated in Figure 1.

3 Empirical strategy

From the perspective of a producer who sells a product in a number of markets that are segmented by variable exchange rates, changes in nominal exchange rates shift relative demand across markets. Exchange rate movements may also affect costs. Our goal is to document producers' desired response of relative markups across markets in response to such shocks.

We exploit the structure of our data to do this without having to make strong assumptions about the nature of the underlying environment.

In the presence of price stickiness there are three margins that determine de facto markup responses to exchange rate movements. First, the choice of the currency in which the sticky price is set determines the default co-movement of relative markups with nominal exchange rates. Second, the producer may choose whether and in which direction to change prices in response to exchange rate movements. Third, conditional on deciding to change prices, the producer can decide how much to change prices as a function of the behavior of exchange rates. The limitations of our data force us to condition on destination currency invoicing since the invoice currency is what allows us to identify the precise destination market for exports, so we focus on the latter two margins. In what follows, we refer to the exercise where we examine the relationship between the timing and sign of price changes and exchange rate changes as the *extensive margin*. 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*. Both of these two exercises will allow us to estimate desired relative markup variation in response to exchange rate changes.

Our key assumption is that relative marginal cost is constant for all markets served by a given plant-product pair. This is what allows us to isolate markup responses to exchange rate shocks.[6] Interpreted strictly, this assumption requires that the definition of a "product" be both narrow and time-invariant. It also relies on prices being measured at the factory gate rather than inclusive of delivery charges, or alternatively, on delivery charges being constant over the relevant time horizon. As noted above, the PPI survey form explicitly solicits matched price quotes in home and export markets for each product category for which a plant reports prices. We rely on this natural matching of price quotes within plant-product pairs to argue that the assumption of product homogeneity is reasonable for our data.[7] As regards the treatment of delivery charges, we are on somewhat weaker ground. Respondents are asked to report prices exclusive of delivery charges if they are invoiced separately, but we

[6]We cannot isolate the relative markup response to the relative demand shift as distinct from changes in marginal cost, as markup responses to an identical marginal cost shock may differ across markets.

[7]Spot checks of the detailed descriptions provided for each quote-line within a plant-product pair support this assumption.

have no information on whether prices are usually charged inclusive or exclusive of delivery charges. However we will later argue that our results cannot be explained by the behavior of delivery charges.

Before describing our empirical strategy in detail, we introduce some notation. Since prices are sticky, it will be important to distinguish between *observed* prices and what we will refer to as *desired* prices. Let i index plant-product pairs, and let k index destination markets. In our case, $k \in \{IRL, UK\}$. We use p_t^{ik} to denote the observed price of plant-product pair i in market k at time t, expressed in home currency. The desired price is the price the producer would choose if forced to change prices at t.[8] This is *not* the same as the optimal flexible price in an environment where all prices are fully flexible. In choosing a new price conditional on adjustment in a sticky-price environment, the producer anticipates that this price and the prices of his competitors will be sticky in the future. We use a carat to distinguish a desired price from an observed price, so $\hat{p}_t^{ik}$ is the desired price expressed in home currency. By definition, the observed home currency price equals the desired home currency price if the invoice currency price (destination currency price in our case) of i in market k is changed at t.

It will also be important to distinguish between prices expressed in home currency and prices expressed in destination currency. We use a star to distinguish prices expressed in destination currency from those expressed in home currency. So p_t^{ik*} is the observed price of i in market k at time t, expressed in the currency of market k. Meanwhile $\hat{p}_t^{ik*}$ is the desired price of i in market k at time t, expressed in the currency of market k. In the case where $k = IRL$, we have $p_t^{ik*} = p_t^{ik}$ and $\hat{p}_t^{ik*} = \hat{p}_t^{ik}$. In addition, we need to keep track of when destination currency prices were last changed. Suppose we are in period t. Let s_t^{ik} be the number of months since the destination currency price for plant-product pair i in market was last changed. For example, if the price changes at date 0, but not at date 1, then irrespective of whether the price is changed at date 2 or not, the number of periods since the last price change is equal to 2.

By definition of the markup, the desired home currency price is equal to the marginal cost expressed in home currency times the desired (gross) markup over marginal cost. We assume

[8] Bils, Klenow and Malin (2009) refer to this as the *reset* price.

that within a plant-product pair, relative marginal cost is constant, so we write this identity as: $\hat{p}_t^{ik} = \gamma^{ik} mc_t^i \hat{\mu}_t^{ik}$. Assuming that desired relative markups may respond to exchange rate movements, we can approximate the log change in the desired home currency price between t and $t - s_t^{ik}$ as follows:

$$\Delta_{s_t^{ik}} \ln \hat{p}_t^{ik} = \alpha + \theta_{t,s_t^{ik}}^i + \beta \Delta_{s_t^{ik}} \ln e_t^k + \varepsilon_{t,s_t^{ik}}^{ik} \tag{1}$$

The term $\theta_{t,s_t^{ik}}^i$ (a plant-product-pair-month-age-of-price fixed effect) captures the log change in marginal cost between t and $t - s_t^{ik}$ (i.e. since the last time the invoice currency price was changed) as well as any changes in the desired markup that are the same across all markets served by plant-product pair i over this horizon. The term $\Delta_{s_t^{ik}} \ln e_t^k$ is the log change in the nominal exchange rate between the home market and market k between t and $t - s_t^{ik}$. This term is equal to zero when $k = IRL$. The coefficient β then captures the desired response of the markup in market k relative to the home market to movements in exchange rates. It is the key parameter we are interested in characterizing. If producers desire constant relative markups (or if price discrimination is not possible), $\beta = 0$. Otherwise, there is "pricing-to-market." β is not a primitive parameter. It depends on the market structure, the nature and degree of price stickiness, and on the process for demand and cost shocks from all sources. It is useful to note that the choice of destination currency invoicing implies that the default (i.e. conditional on prices not changing in either market) is for relative markups to move one-for-one with movements in exchange rates. We now outline our two approaches to characterizing the desired relative markup response, β.

Extensive margin

Prices are sticky in destination currency our data. We use (1) to derive a relationship between the probability that the destination currency price is increased (or decreased), and movements in exchange rates. The derivation is presented under the assumption of a menu cost of changing prices, but we argue that the resulting estimating equations have a useful interpretation under any form of price stickiness.

If the change in the log desired *home* currency price between t and $t - s_t^{ik}$ takes the form

(1), we subtract $\Delta_{s_t^{ik}} \ln e_t^k$ from both sides to obtain the change in the log desired *destination* currency price over the same time interval:

$$\Delta_{s_t^{it}} \ln \hat{p}_t^{ik*} = \Delta_{s_t^{it}} \ln p_t^{ik*} - \Delta_{s_t^{ik}} \ln e_t^k = \alpha + \theta_{t,s_t^{ik}}^i + (\beta - 1)\, \Delta_{s_t^{ik}} \ln e_t^k + \varepsilon_{t,s_t^{ik}}^{ik} \tag{2}$$

Suppose that producers increase destination currency prices if the change in the desired price is sufficiently positive, and reduce them if the change in the desired price is sufficiently negative. Substituting in from (2) this yields the rule:

$$\begin{aligned}
\Delta_{s_t^{ik}} \ln p_t^{ik*} > 0 \quad &\text{if} \quad \alpha + \theta_{t,s_t^{ik}}^i + (\beta - 1)\, \Delta_{s_t^{ik}} \ln e_t^k + \varepsilon_{t,s_t^{ik}}^{ik} > \bar{\rho}_t^i > 0 \\
\Delta_{s_t^{ik}} \ln p_t^{ik*} = 0 \quad &\text{if} \quad \underline{\rho}_t^i \leq \alpha + \theta_{t,s_t^{ik}}^i + (\beta - 1)\, \Delta_{s_t^{ik}} \ln e_t^k + \varepsilon_{t,s_t^{ik}}^{ik} \leq \bar{\rho}_t^i \\
\Delta_{s_t^{ik}} \ln p_t^{ik*} < 0 \quad &\text{if} \quad \alpha + \theta_{t,s_t^{ik}}^i + (\beta - 1)\, \Delta_{s_t^{ik}} \ln e_t^k + \varepsilon_{t,s_t^{ik}}^{ik} < \underline{\rho}_t^i < 0
\end{aligned} \tag{3}$$

A rule of this form is optimal if there are fixed costs of changing prices, and expected profits are concave in (destination currency) prices. In (3), we make the additional assumption that the cutoffs of the inaction region are the same across markets within a plant-product pair.

Rule (3) suggests that we estimate two separate equations, one for price increases, and one for price decreases. The conditional logit is a convenient estimator, as it allows us to treat $\theta_{t,s_t^{ik}}^i$ and the cutoffs of the inaction region as fixed effects. To implement this, we assume that $\varepsilon_{t,s_t^{ik}}^{ik}$ has a logistic distribution. We can then write:

$$\Pr\left[\Delta_{s_t^{ik}} \ln p_t^{ik*} > 0\right] = \Lambda\left(\psi_{t,s_t^{ik}}^i + (\beta - 1)\, \Delta_{s_t^{ik}} \ln e_t^k\right) \tag{4}$$

$$\Pr\left[\Delta_{s_t^{ik}} \ln p_t^{ik*} < 0\right] = \Lambda\left(\phi_{t,s_t^{ik}}^i - (\beta - 1)\, \Delta_{s_t^{ik}} \ln e_t^k\right) \tag{5}$$

where $\Lambda(z) = \exp(z) / [1 + \exp(z)]$, $\psi_{t,s_t^{ik}}^i = \alpha - \bar{\rho}_t^i + \theta_{t,s_t^{ik}}^i$ and $\phi_{t,s_t^{ik}}^i = -\alpha + \underline{\rho}_t^i - \theta_{t,s_t^{ik}}^i$. To estimate (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. In estimating (4) and (5), we weight by sales shares, constructed as described in the data appendix, and cluster standard errors at the plant level. Note that logit estimation does not identify the scale of $(\beta - 1)$ or $-(\beta - 1)$, but only whether the effect of changes in exchange rates on the probability that prices change is positive, negative, or not significantly different from zero.

Given the structure of the fixed effects, 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 last price change must have been synchronized across the Irish and UK markets ($s_t^{i,IRL} = s_t^{i,UK} = s_t^i$).[9] 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 $s_t^{i,IRL} = s_t^{i,UK} = s_t^i$. The method does not make use of cases where $s_t^{i,IRL} = s_t^{i,UK} = s_t^i$, but no prices change at t, or where $s_t^{i,IRL} = s_t^{i,UK} = s_t^i$ but all prices change in the same direction. This is illustrated by sets A and B in Figure 2. For clarity, in estimating (4) and (5), we include only observations that can be used to identify the coefficient on exchange rate changes.

The intuition for estimating (4) and (5) under more general forms of price stickiness is as follows. For plant-product pairs where the last price change was synchronized across the two markets, the change in marginal cost between t and the previous price change is the same in both markets. This implies that if destination currency prices do not change in either market at t, relative markups will move one-for-one with movements in exchange rates. If producers wish to offset these changes in relative markups (i.e. if $\beta < 1$), they must change prices in at least one of the two markets. Moreover, if $\beta < 1$, the *direction* of price changes conditional on at least one price changing must be such as to (weakly) undo the drift in relative markups due to exchange rate movements since the last time prices were changed. The prediction about the direction of relative price movements does not depend on the existence of menu costs of price changes.

To illustrate with a particular example, suppose $s_t^{i,IRL} = s_t^{i,UK} = s_t^i$, and $\Delta_{s_t^i} \ln e_t^{UK} > 0$ (i.e. the home currency depreciates against Sterling between t and $t - s_t^i$). Absent a change in destination currency prices, this increases the markup in the UK market relative to the home market. If $\beta < 1$, the producer may choose to change at least one price in such a way as to realign relative markups in the desired direction. This requires either a reduction of the Sterling price in the UK, an increase in the home currency price at home, or both. First consider the case of a reduction in the Sterling price in the UK market, but no reduction in the home currency price in the Irish market. This implies a positive correlation between the

[9] As Table 4 documents, synchronization of price changes is common.

probability of a reduction of the destination currency price and the change in the exchange rate, and hence a positive coefficient on the exchange rate change in (5). Second, consider the case of an increase in the home currency price at home, but no increase in the Sterling price in the UK market. Since $\Delta_{s_t^i} \ln e_t^{IRL} = 0$, an increase in the home currency price in the home market shows up as a positive fixed effect $\psi^i_{t,s_t^{ik}}$, in (4). The failure to increase the Sterling price in the UK market then implies a negative coefficient on the exchange rate change in (4).

In contrast, if $\beta = 1$, producers find it desirable for relative markups to move one-for-one with exchange rates. In this case changes in exchange rates alone will not induce producers to change prices, and the probability of price increases or decreases will be unrelated to exchange rate changes. The case where $\beta > 1$ is the converse of that where $\beta < 1$.

Intensive margin

The exercise just described makes use of episodes where prices need not change in both markets, but it has the disadvantage that it does not allow us to fully characterize β, only to test whether $\beta \gtrless 1$. If we condition on destination currency prices changing in both the home and UK markets at date t, while maintaining the assumption that $s_t^{i,IRL} = s_t^{i,UK} = s_t^i$, we can can directly estimate:

$$\Delta_{s_t^{ik}} \ln p_{kt}^{ik} = \alpha + \theta^i_{t,s_t^{ik}} + \beta \Delta_{s_t^{ik}} \ln e_t^k + \varepsilon^{ik}_{t,s_t^{ik}} \tag{6}$$

We can do this because when destination currency prices change, observed home currency prices are equal to desired home currency prices. We use fixed effects to control for $\theta^i_{t,s_t^{ik}}$. In estimating (6), we exclude observations where the log change in home currency price is greater than 2 in absolute value. We weight by sales shares and cluster standard errors at the plant level.

Given the structure of the fixed effects, information from plant-product pair i is used to identify β only if prices change in both Irish and UK markets at date t, and the previous price changes were also synchronized. This is illustrated by set C in Figure 2. As we document in Table 4, synchronization of price changes within plant-product pairs is the norm rather than the exception. However, the requirement that two consecutive sets of price changes be synchronized does restrict the size of the sample we can use to estimate (6). As Figure 2 also illustrates, there is some overlap between the samples used to characterize β in the extensive margin exercise and those used to identify β in the intensive margin exercise, but this overlap is not perfect. For clarity, in the estimation sample we include only observations that can be used to identify the coefficient on exchange rate changes.

4 Results

Extensive margin

Tables 5 and 6 report the results from estimating equations (4) and (5) on the baseline sample and various different subsamples. The first row of each table gives the baseline estimates, pooling all the identifying observations. The number of plants used to identify the coefficient on exchange rate changes is indicated by the number of clusters, while the number of plant-product-months is indicated by the number of fixed effects. As indicated in the table, if $\beta < 1$ (i.e. producers want to offset the default effect of exchange rate movements on relative markups given destination 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, so we cannot reject the hypothesis that $\beta = 1$. This is consistent with an extreme form of pricing-to-market, where even conditional on adjusting some prices, producers allow relative markups to move one-for-one with movements in exchange rates.[10]

Pricing, product and plant characteristics

We now document what happens when we condition on aspects of pricing behavior and product and plant characteristics. Theory suggests that producers' incentives to realign

[10]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.

relative markups when they change prices may be related to how long they expect to wait before changing prices again.[11] We do not find strong evidence of differences in pricing-to-market behavior along this dimension of price characteristics. The first panel of Table 5 reports the results from estimating (4) and (5), splitting the data by median frequency of price adjustment at the level of the plant-product pair and by the time interval between the current period and the previously synchronized price changes. In all but one case, the estimated coefficients are not significantly different from zero, consistent with $\beta = 1$.

Theory also suggests that, conditional on market segmentation, the desired degree of pricing-to-market may vary systematically with the elasticity of the residual demand curve. This in turn may vary by type of product. We do not find strong evidence in favor of this type of variation. 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.[12] The second panel of Table 5 reports the results from estimating equation (4) and (5), splitting the data by type of product for these two classifications. For the Vermeulen classification, in all but one case, the coefficients on the exchange rate change are not significantly different from zero, consistent with $\beta = 1$. For the Rauch classification, a priori one might expect market integration to be a stronger constraint on pricing-to-market in the organized exchange and reference priced groups. We do find one coefficient significantly different from zero at the 5% level, for the reference priced group, in the case of price increases. The sign is consistent with $\beta < 1$. All the other coefficients are not significantly different from zero, consistent with $\beta = 1$. Having looked in detail at the products assigned to each category, we conclude that the Rauch classification does not capture the intended product attributes in our data. This may be due to errors in concording the original classification, which applies to trade data, with our 4-digit production classification at the plant level.

The incentive to engage in pricing-to-market may differ by measurable plant characteristics. For example, the elasticity of residual demand may be lower for large than for small plants. We investigate this possibility by splitting the data by plant size. The results are

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

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

reported in the first panel of table 6. The evidence is mixed. Most of the coefficients on the exchange rate change are not significantly different from zero, consistent with $\beta = 1$. For price increases, we have one marginally significant coefficient which goes in the "wrong" direction, suggesting that small plants engage in an even more extreme version of pricing-to-market than large plants, with $\beta > 1$. For price decreases, for the two smallest size categories, the coefficients are significantly different from zero at the 5% and 10% levels, and both go in the "right" direction, consistent with $\beta < 1$ and less extreme pricing-to-market for smaller plants.

It is also possible that pricing strategies may differ between plants for which the UK market is a core market, and those for which it is not. We investigate this possibility by splitting the data by quartiles of the share of sales exported to the UK. The results are reported in the second panel of Table 6. We do not find systematic evidence that pricing differs along this dimension.

We also split the data by quartiles of price-cost margins, which may capture some of the variation in the level of markups. In panel 3 of Table 6 we report the results under this split. In all but one case, the coefficients on the exchange rate change are not significantly different from zero.

Under some circumstances, a producer's ability to hedge along different dimensions could 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.[13] 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 (calculated as materials plus energy plus wage bill). In all cases, we cannot reject $\beta = 1$. 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, and consistent with $\beta = 1$ in all cases.

One issue of concern is the fact that we do not know from the PPI whether the prices we observe are for transactions between related parties. Luckily, the CIP collects information on

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

the fraction of sales that is to related parties. We estimate (4) and (5) using only domestic-owned plants, splitting the data by whether some or no sales to related parties are reported. The results are reported in the third panel of Table 6. For price increases, the coefficient on the exchange rate change is not significantly different from zero in both cases, but for price decreases, the coefficient is significantly positive for plants that do not sell to related parties. This is consistent with $\beta < 1$.

To summarize, the evidence from the extensive margin generally points to a value of β 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 significantly less than 1 or significantly greater than 1.

Intensive margin

Tables 7 and 8 report the results from estimating (6) conditioning on consecutive synchronized price changes. The first row of each table gives the baseline estimate, based on pooling all the identifying observations in the data. The number of plants used to identify the coefficient on exchange rate changes is indicated by the number of clusters, while the number of plant-product-months is indicated by the number of fixed effects. The coefficient on the exchange rate change is our estimate of β. The baseline estimate is positive and significantly different from zero, and not significantly different from one. This is consistent with the extensive margin results.

In the case of the intensive margin, it is straightforward to illustrate the variation that identifies β. Figure 3 is a scatter plot of 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.[14] Note that if $\beta = 1$, based on rearranging (6), the log change in the Sterling price in the UK should on average be equal to the log change in the home currency price in Ireland. The data does indeed line up around the 45^o line. This figure also illustrates the fact that $\hat{\beta} = 1$ for positive

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

and negative price changes, and for price changes of different sizes. A different aspect of the variation identifying the coefficient is illustrated in Figure 4 which plots histograms of the log changes in destination currency prices and the log changes in exchange rates for the estimation sample. The standard deviation of price changes is large. 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 the next is on average a good deal smaller than the size of the corresponding price changes.[15]

Pricing, product and plant characteristics

Just as in the case of the extensive margin, we document what happens when we condition on aspects of pricing behavior and product and plant characteristics. We do not find conclusive evidence of differences in pricing-to-market behavior along any of these dimensions. For all cuts of the data, the estimated β is significantly different from zero, and with only one exception, is not significantly different from 1. We now describe these results in more detail.

The second panel of Table 7 reports the results from estimating equation (6), 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 time interval between the first and second synchronized price changes. In all cases, the estimated β is significantly different from zero, and not significantly different from 1. If anything, the point estimates suggest a β closer to zero at short horizons than at long horizons, which goes against what one might expect.

The third panel of Table 7 reports the results from estimating equation (6), splitting the data by type of product for the Vermeulen and Rauch 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. For the Rauch classification, the estimated β is always significantly different from zero and not significantly different from 1.

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

The first panel of Table 8 reports the results from splitting 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.

In the second panel of Table 8, we report the results from splitting the data by quartiles of the share of sales exported to the UK. For the upper three quartiles, the estimated β is significantly different from 0, not significantly different from 1. For the lowest quartile, the coefficient is less than 1, and significantly different from both 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 Table 8, we report the results from splitting the data by quartiles of price-cost margins. We find no significant differences in the estimated β across these groups. In the second panel of Table 8, we report the results from splitting the data by quartiles of the share of materials imported from the UK in variable cost. We find no significant differences in the estimated β across these groups. In the third panel, we report the results splitting the data by ownership. The estimates of β for domestic, foreign and UK-owned plants are statistically indistinguishable.

Finally, we estimate (6) using only domestic-owned plants, splitting the data by whether some or no sales to related parties are reported. The results are reported in the third panel of Table 8. The estimate of β is significantly different from 0 and not significantly different from 1 in both cases.

To summarize, the evidence from the intensive margin points to an extreme form of pricing-to-market, where conditional on adjusting prices in both markets, producers allow relative markups to move one-for-one with movements in exchange rates.

Potential sources of bias

Our identification strategy requires us to make use of synchronized sets of price changes to identify desired markup variation at both the extensive and intensive margins. It is possible that synchronized price changes are somehow special, and that selection along this dimension may in some way drive our results. We are reassured on this point by the following observations. First, synchronization of price changes within plant-product pairs is very common in our data, as documented in Table 4. As a result, together the extensive and intensive margin exercises make use of 59% of the price changes in the sample of plant-product pairs with simultaneous home and Sterling-invoiced export quotes. Second, both the extensive and intensive margins point to similar values of β. These two different exercises make use of different aspects of the variation in the data and non-identical though overlapping sets of observations. Third, we split the data along a wide range of observable pricing, product and plant characteristics, and find no systematic differences in the results across these groups.

The fact that producer price data is collected using a survey is also a potential source of bias. Participants are instructed to provide data drawn from invoices rather than list prices, and are encouraged on their initial participation to report prices for a representative selection of products. But it is almost guaranteed that their responses ultimately violate these guidelines in some way. On the first point, we take some comfort from the fact that by focusing on episodes of price changes, our results are robust to the errors that arise when the previous period's price is filled in in the absence of a transaction in a particular month. On the second point, if producers choose to disproportionately report prices for quote-lines they expect to last for a long time rather than those they anticipate will be short-lived, or for quote-lines where they expect there to be relatively little price variation, this may influence our results. This particular problem is common to producer price data from all countries, and is something we cannot do much about.

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 extensive and intensive margin exercises examine responses to contemporaneous exchange rate changes from one price change to the next. We have already addressed this issue to some extent by splitting the data along the dimension of the length of the interval between the first price change and the second. If there is a dynamic component to adjustment, 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 are statistically indistinguishable.

To further investigate the possibility that there is dynamic adjustment of relative markups, we perform two additional exercises. Both of these follow the intensive margin approach, as our extensive margin strategy cannot be generalized to the case with intermediate price changes. First, we add a second lag of the exchange rate change to (6), where the lag length is the same as for the original independent variable:

$$\Delta_{s_k^i(t)} \ln p_{kt}^i = \alpha + \theta_{t,s_k^i(t)}^i + \beta_1 \Delta_{s_k^i(t)} \ln e_{kt} + \beta_2 \Delta_{s_k^i(t)} \ln e_{kt-s_k^i(t)} + \varepsilon_{kt,s_k^i(t)}^i \tag{7}$$

The results from estimating this equation are presented in Panel 1 of Table 9. 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. We regress the log change in home currency price over this horizon on the contemporaneous log change in the relevant exchange rates. This allows us first, to examine price changes 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 sample of price changes are reported in Panel 2 of Table 9. 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 suggestive. 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 is weak evidence that the combination of price stickiness and strategic complementarities may be partly responsible for the extreme degree of pricing-to-market we document in the shorter run.

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.

One question that we have not so far addressed is how pricing-to-market behavior 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 the invoice currency is a choice variable, one might expect pricing-to-market to differ along this dimension. We cannot test this for the full sample, since we do not know the destination of exports invoiced in home currency. 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 responding plant. 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 along with sales to the Irish market. We restrict this exercise to the period 2003-2005, as the further back we go, the smaller and more selected the set of prices.

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. It is equal to β in the case of price increases, and $-\beta$ in the case of price decreases. The results from estimating (4) and (5) in the case of destination currency invoicing, and its counterparts in the case of home currency invoicing are reported in Table 10. In the destination currency invoicing case, the estimated coefficients on exchange rate changes are not significantly different from zero, so we cannot reject the null hypothesis that $\beta = 1$. 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 $\beta = 0$.

Table 11 reports the results from estimating (6) separately for the destination currency sub-sample and the home currency sub-sample. β is relatively imprecisely estimated in both cases. In the destination 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. We conclude that there is evidence that pricing-to-market behavior differs by choice of invoice currency in the direction that selection would imply. This obviously begs the question of what determines invoice currency choice. We document in the data appendix that there is a good deal of heterogeneity both within and across similar plants in the choice of invoice currency for sales to the UK market. In view of the substantial within-plant heterogeneity (33% of plants invoice exports to the UK in a mixture of home and foreign currencies), it appears that information on the counterparty to a transaction may be necessary to fully understand invoice currency choice, and this is something we do not have available in our data.

Another 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 (4), (5) and (6), 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, if sales to Japan are invoiced in US dollars. The results (reported in the data appendix) are very similar to the baseline. On the extensive margin, we cannot reject the null hypothesis that $\beta = 1$. On the intensive margin, the point estimate of β is very close to 1, not significantly different from 1, and significantly different from zero.

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 charged by producers in Ireland to buyers in Ireland and the UK, two markets segmented by variable exchange rates. Assuming that relative marginal cost across markets is constant over time within the same plant-product pair allows us to use fixed effects to estimate the response of relative markups across the two markets to movements in exchange rates. We make use of episodes where prices change in order to isolate desired relative markup variation from default variation in relative markups arising purely out of the stickiness of prices in destination currency. Our results suggest that for producers invoicing in destination currency, desired relative markups move one-for-one with exchange rates.

As already mentioned, by conditioning on destination currency invoicing, we are selecting the case where it is most likely that there is pricing-to-market, so it is not so surprising that this is what we observe. We show that in cases where home currency invoicing is chosen, desired relative markups do not appear to co-move with exchange rates. This implies that understanding the determinants of invoice currency choice is crucial to fully understanding pricing-to-market behavior. As we note, the pattern of invoicing within and across Irish plants for sales to the UK market suggests that information on counterparties to a transaction may be necessary to fully understand these determinants.

Although it is not surprising that we find evidence of pricing-to-market, the response of relative markets to exchange rate changes that we estimate is more extreme than that found by the previous literature. There are three key differences between our work and the previous literature. First, our data is on prices invoiced in destination currency. Much of the previous literature uses unit value data where the invoice currency is unknown. If pricing-to-market behavior is as strongly correlated with invoice currency as it appears from our results, the findings of the previous literature may be accounted for by the use of unit values based on sales invoiced in a mixture of home and foreign prices.

Second, we observe prices at the plant-product level, and our results are based on holding contract terms and product characteristics constant over time. Price is not the only margin along which producers can respond to exchange rate shocks. Quantities may adjust. Producers may also change terms of the contract besides the price, or they may adjust the quality and characteristics of products in a way that differs across markets. Responses along these dimensions may show up in unit values, or in prices, if these other dimensions are not held constant.

Third, in contrast to the previous literature, we estimate desired pricing-to-market conditional on price changes. Because of this, we focus on shorter horizons than the previous literature, since on average, prices change more than once a year. Partly in consequence, the size of the exchange rate changes we use to identify pricing-to-market is on average quite small. It is over these short intervals, and in response to these small exchange rate changes that we find producers desire a one-for-one comovement betweeen relative markups and exchange rates.

In interpreting our results, the following points are important. First, as we have just reiterated, we identify one-for-one desired relative markup responses to relatively small exchange rate changes. These exchange rate changes are smaller than the size of the price changes we use to isolate desired markup responses, and failure to respond optimally may lead to relatively small losses in profits. Second, price-cost margins for the plants in our data are large. This suggests that competitive pressures may be insufficient to eliminate market segmentation, while the risks of producers being driven to negative profits by small movements in exchange rates are slim. In this sense, 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. 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 market.

Finally, it is crucial to remember that the plants in our data operate in a sticky-price environment. In such an environment, strategic complementarities in pricing behavior may lead producers to desire limited short-run adjustment of relative markups to small movements in exchange rates, though adjustment may eventually take place, through prices, or along some other dimension. It is noteworthy that over particularly long horizons, allowing for multiple intervening price changes, evidence on pricing-to-market is less stark. It remains to be seen what combination of market structure and type of price stickiness can match the moments we document. But it is likely that strategic complementarity in pricing behavior, or adjustment along non-price dimensions will be necessary.

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 (2010), "A Structural Approach to Identifying the Sources of Local-Currency Price Stability," mimeo.

[13] Goldberg, P. and R. Hellerstein (2009), "How Rigid Are Producer Prices," mimeo.

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

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

[16] 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.

[17] Gopinath, G., P.-O. Gourinchas, C.-T. Hsieh and N. Li (2010), "International Prices, Costs and Markup Differences." *American Economic Review*, forthcoming.

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

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

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

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

[22] Klenow, P. and B. Malin (2010), "Microeconomic Evidence on Price Setting," mimeo.

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

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

[25] Nakamura, E. and D. Zerom (2010), "Accounting for Incomplete Pass-Through," *Review of Economic Studies*, 77 (3), 1192-1230.

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

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

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

[29] 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
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: 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 time period. Information on the export currency is based on the roughly 95% of the population of plants 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: These statistics are for the matched plant-price data set. Each column reports the number of distinct values of each category observed in the relevant time-frame. For example, over the course of 1995, we observe prices for 669 distinct plants, though not all of those plants may be observed in any given month of that year.

Table 3: 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 4: 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	51
Irl & UK sample	20	28	45	28

Note: The first column reports the fraction of plant-product-months where there is more then one price quote, and at least one of thos prices 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 5: Extensive Margin of Price Adjustment I

	Probability of a price increase						Probability of a price decrease					
	$\Delta_{s_t^{ik}} \ln e_t^k$		ps-R^2	N	f.e.	clust	$\Delta_{s_t^{ik}} \ln e_t^k$		ps-R^2	N	f.e.	clust
	$\beta - 1$						$1 - \beta$					
all	0.59	(3.61)	0.00	4873	921	129	1.62	(4.99)	0.00	4564	875	103
Panel 1												
Median frequency (f) of price adjustment of plant-product pair												
$f < 0.3$	3.05	(4.79)	0.00	2517	442	111	-5.53	(5.90)	0.00	2187	378	84
$0.3 \leq f < 0.5$	-1.28	(12.05)	0.00	693	172	10	15.73	(7.26)**	0.01	723	187	10
$0.5 \leq f < 0.7$	-0.90	(6.00)	0.00	1160	174	6	5.61	(2.95)	0.00	1129	171	7
$0.7 \geq f$	-13.97	(12.63)	0.00	503	133	4	21.50	(15.00)	0.01	525	139	4
Age of price at second synchronized price change												
1 mth	-4.53	(7.01)	0.00	2537	495	48	8.50	(7.10)	0.00	2424	478	44
2-5 mths	9.47	(8.19)	0.00	1206	232	65	1.70	(8.06)	0.00	1371	258	56
6-11 mths	2.14	(8.26)	0.00	574	98	59	1.67	(9.05)	0.00	450	82	49
12+ mths	0.29	(5.39)	0.00	556	96	62	-4.50	(7.11)	0.01	319	57	40
Panel 2												
Type of product (Vermeulen et al., 2007)												
(1)	-1.97	(6.57)	0.00	1408	331	27	-5.35	(6.63)	0.00	1380	321	21
(2)	23.14	(24.79)	0.08	101	26	11	54.75	(47.53)	0.12	86	16	10
(3)	-1.44	(9.41)	0.00	357	46	14	26.49	(15.11)*	0.08	178	29	10
(4)	1.20	(6.75)	0.00	1904	376	55	3.71	(9.39)	0.00	1840	370	45
(5)	2.20	(6.55)	0.00	1101	141	22	-0.48	(7.99)	0.00	1078	138	19
Type of product (Rauch, 1999)												
org. exch.	1.24	(11.68)	0.00	1173	274	11	3.02	(7.60)	0.00	1213	280	9
ref. priced	-19.08	(4.81)**	0.02	372	112	13	12.15	(14.27)	0.01	352	105	11
diff.	4.83	(5.43)	0.00	1918	312	66	2.99	(5.65)	0.00	1797	284	54

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. 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. Key to Vermeulen categories: (1) Consumer food products, (2) Consumer non-food non-durables, (3) Consumer durables, (4) Intermediates, (5) Capital goods.

Table 6: Extensive Margin of Price Adjustment II

	Probability of a price increase						Probability of a price decrease					
	$\Delta_{s_t^{ik}} \ln e_t^k$		ps-R^2	N	f.e.	clust	$\Delta_{s_t^{ik}} \ln e_t^k$		ps-R^2	N	f.e.	clust
	$\beta - 1$						$1 - \beta$					
all	0.59	(3.61)	0.00	4873	921	129	1.62	(4.99)	0.00	4564	875	103
Panel 1												
Plant size												
<20	-12.82	(16.92)	0.01	142	41	13	3.91	(1.89)**	0.00	72	22	7
20-49	4.47	(2.59)*	0.00	1047	164	35	19.96	(10.37)*	0.02	976	155	31
50-249	3.15	(5.07)	0.00	2808	496	68	-2.87	(5.81)	0.00	2586	469	53
250-499	-4.44	(8.03)	0.00	753	187	9	-6.11	(8.15)	0.00	809	196	10
500+	-6.20	(17.20)	0.00	123	33	4	censored to maintain confidentiality					
Panel 2												
Quartiles of share of sales exported to the UK												
Q1	-26.65	(18.32)	0.14	108	21	5	-2.98	(9.63)	0.00	95	18	6
Q2	15.03	(27.09)	0.03	117	22	10	-68.26	(17.52)*	0.08	92	17	6
Q3	-2.85	(4.95)	0.00	2054	336	50	4.89	(5.53)	0.00	1866	321	34
Q4	2.91	(5.70)	0.00	2594	542	86	0.04	(7.58)	0.00	2511	519	78
Quartiles of share of variable cost imported from the UK												
Q1&Q2	-5.30	(7.42)	0.00	1354	321	49	9.70	(8.46)	0.01	1465	335	45
Q3	2.18	(6.11)	0.00	907	160	46	-5.07	(6.94)	0.00	769	141	36
Q4	2.46	(5.63)	0.00	2585	429	61	0.20	(7.50)	0.00	2289	382	48
Panel 3												
Ownership												
home	0.38	(4.29)	0.00	3356	664	95	5.53	(4.42)	0.00	2999	618	74
foreign	0.89	(6.31)	0.00	1517	257	35	-2.18	(8.79)	0.00	1565	257	31
UK	-7.43	(7.83)	0.02	307	44	8	-14.41	(21.83)	0.01	290	39	8
Quartiles of price-cost margins												
Q1	6.71	(9.39)	0.00	1689	376	52	-0.70	(6.75)	0.00	1756	391	46
Q2	-2.68	(6.80)	0.00	1225	233	51	5.66	(8.81)	0.00	1230	222	42
Q3	2.49	(4.64)	0.00	1591	251	54	5.80	(5.31)	0.00	1248	212	41
Q4	-2.62	(9.04)	0.00	355	58	26	-12.52	(6.11)**	0.03	317	48	24
Share of sales to related parties												
None	-3.98	(3.94)	0.00	1975	392	62	8.44	(3.66)**	0.00	1621	339	45
Some	-5.48	(10.87)	0.00	521	76	19	11.96	(13.42)	0.02	494	75	17

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. 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 7: Intensive Margin of Price Adjustment I

	$\Delta_{s_t^{ik}} \ln e_t^k$		R^2-adj.	N	f.e.	clust
	β					
Baseline						
All	1.01	(0.09)**	0.67	4212	1047	86
Panel 1						
Median frequency (f) of price adjustment of plant-product pair						
$f < 0.3$	1.07	(0.12)**	0.64	1226	293	68
$0.3 \leq f < 0.5$	0.93	(0.15)**	0.61	1089	313	8
$0.5 \leq f < 0.7$	1.03	(0.33)**	0.67	1269	242	7
$0.7 \geq f$	1.00	(0.39)**	0.69	628	199	5
Time interval between synchronized price changes						
1 mth	0.77	(0.38)**	0.66	748	381	42
2-5 mths	0.93	(0.12)**	0.80	1149	516	46
6-11 mths	0.96	(0.04)**	0.88	1377	578	71
12+ mths	1.03	(0.28)**	0.64	938	441	64
Panel 2						
Type of product (Vermeulen et al. 2007)						
consumer food products	0.91	(0.17)**	0.56	1300	393	20
cons non-food non durab	1.86	(0.60)**	0.43	29	13	6
cons durables	1.06	(0.19)**	0.68	147	28	8
intermediates	1.11	(0.23)**	0.81	1481	404	31
capital goods	1.05	(0.10)**	0.60	1251	207	22
Type of product (Rauch 1999)						
org. exch.	1.03	(0.42)**	0.53	985	292	10
ref. priced	0.87	(0.17)**	0.85	378	146	7
diff.	1.14	(0.15)**	0.56	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 (coefficient and s.e. not reported) 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 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 8: Intensive Margin of Price Adjustment II

	$\Delta_{s_t^{ik}} \ln e_t^k$		R^2-adj.	N	f.e.	clust
	β					
Baseline						
All	1.01	(0.09)**	0.67	4212	1047	86
Panel 1						
Number of employees						
<20	1.25	(0.12)**	0.86	283	98	10
20-29	1.07	(0.23)**	0.67	894	188	24
50-249	1.06	(0.12)**	0.55	2129	477	44
250-499	0.95	(0.17)**	0.61	752	227	7
Panel 2						
Quartiles of share of sales exported to the UK						
Q1	0.73	(0.03)**	0.43	68	25	4
Q2	0.94	(0.05)**	0.78	117	25	11
Q3	1.07	(0.22)**	0.58	1865	378	34
Q4	1.01	(0.11)**	0.71	2162	619	57
Quartiles of share of variable cost imported from the UK						
Q1 & Q2	1.03	(0.26)**	0.63	1334	413	32
Q3	1.03	(0.13)**	0.70	617	170	38
Q4	1.00	(0.07)**	0.67	2213	445	40
Panel 3						
Ownership						
domestic	1.04	(0.13)**	0.57	2842	791	64
foreign	0.96	(0.08)**	0.82	1370	256	24
UK	0.99	(0.01)**	0.96	175	27	8
Quartiles of price-cost margins						
Q1	0.98	(0.17)**	0.61	1757	466	39
Q2	0.92	(0.13)**	0.70	1178	294	38
Q3	1.61	(0.55)**	0.81	951	208	34
Q4	0.99	(0.11)**	0.40	326	79	22
Share of sales to related parties						
None	1.11	(0.26)**	0.58	1731	466	38
Some	1.11	(0.11)**	0.67	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 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 9: Intensive Margin of Price Adjustment: Dynamics

	$\Delta_{s_t^{ik}} \ln e_t^k$		$\Delta_{s_t^{ik}} \ln e_{t-s_t^{ik}}^k$		R^2-adj.	N	f.e.	clust
Panel 1								
Two lags of exchange rate changes								
	0.93	(0.08)**	-0.05	(0.11)	0.67	3794	960	67
Panel 2								
First and last synchronized price changes, all observations								
	0.78	(0.51)			0.39	677	116	96
First and last synchronized price changes, interval between first and last								
<6 months	1.59	(5.03)			0.16	43	23	20
6-11 months	1.38	(0.88)			0.44	65	27	22
12-23 months	1.81	(1.34)			0.47	141	40	35
24+ months	0.74	(0.61)			0.36	428	75	66
First and last synchronized price changes, number of intervening price changes								
<2	0.86	(0.46) *			0.33	400	58	49
≥2	0.76	(0.66)			0.30	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 10: Robustness: Extensive margin by choice of invoice currency

		$\Delta_{s_t^{ik}} \ln e_t^k$		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.

Table 11: Robustness: Intensive margin by choice of invoice currency

Invoice curr.	$\Delta_{s_t^{ik}} \ln e_t^k$		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. 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.

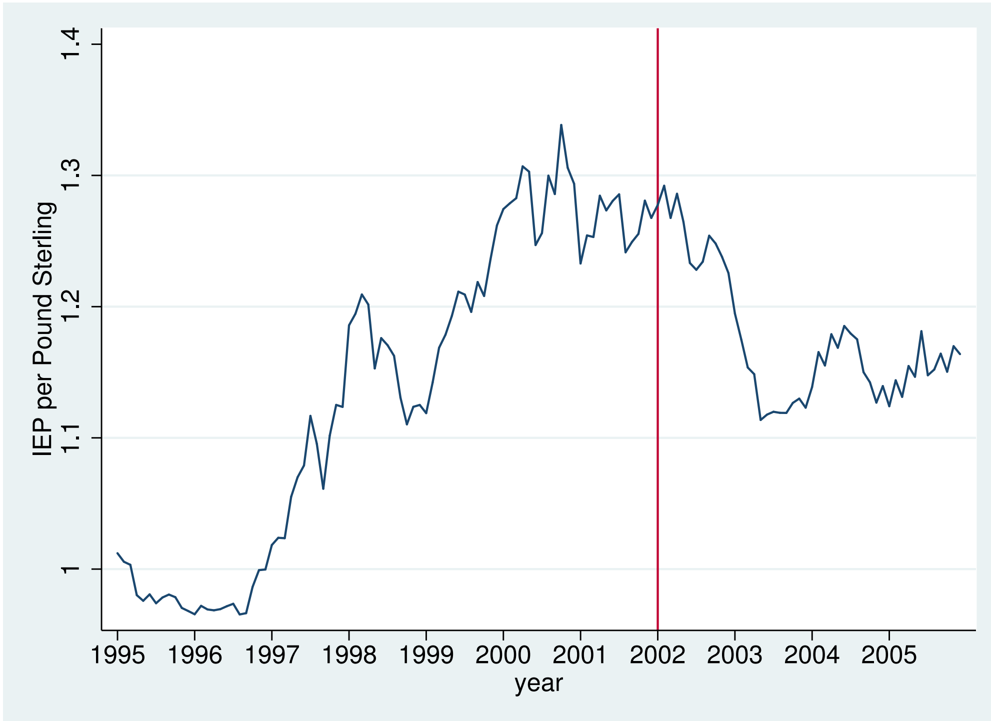

Figure 1: IEP/EUR per Pound Sterling over the sample period

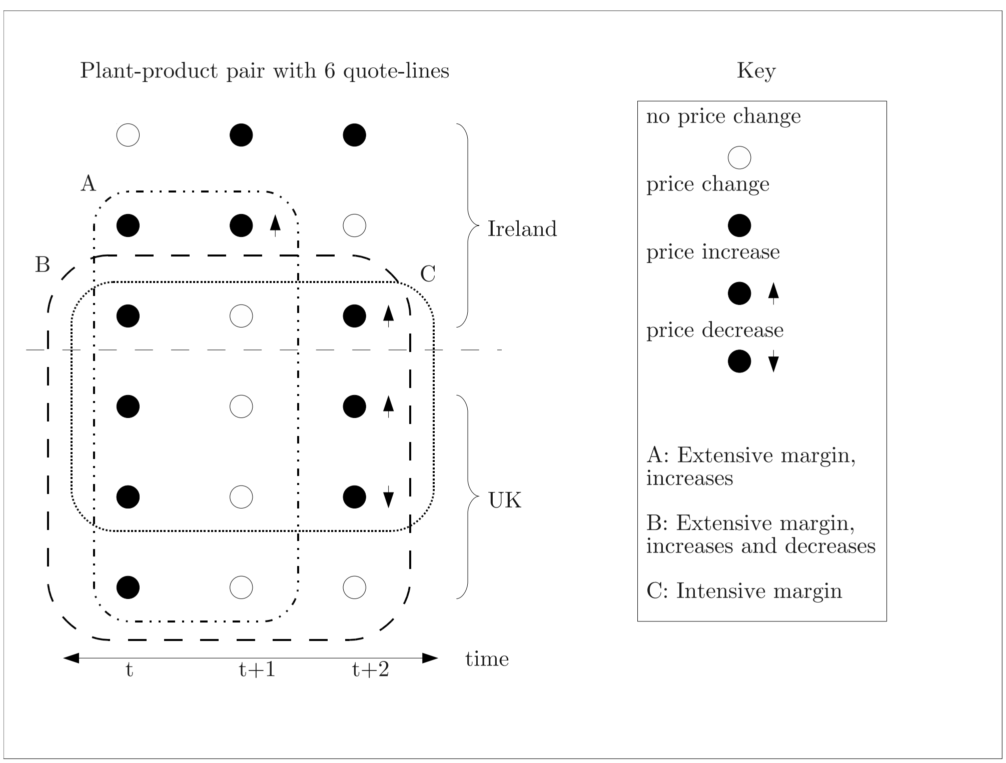

Figure 2: Illustration of variation used to identify the extensive and intensive margins

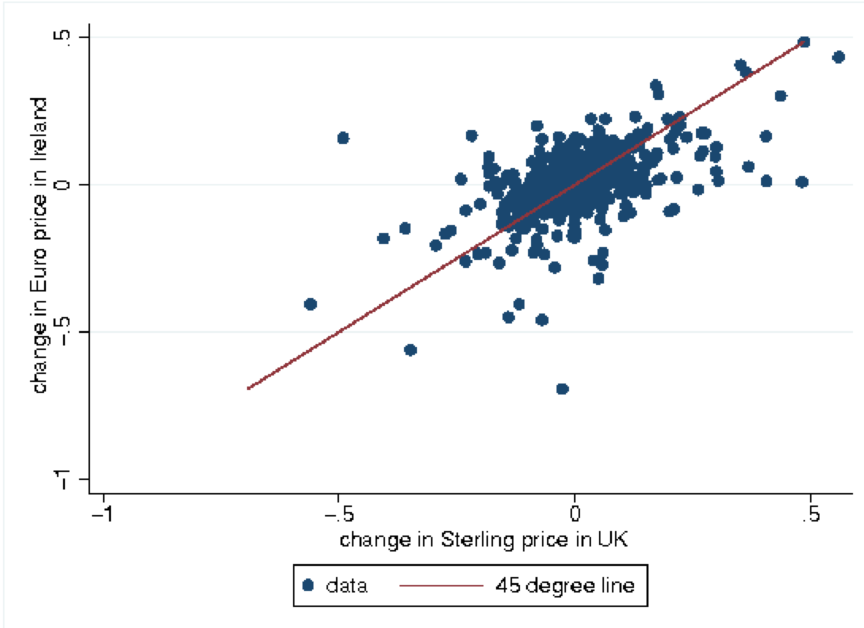

Figure 3: Scatter plot of data used to identify the intensive margin
Note: Figure plots the log change in the Sterling price in the UK against the log change in the home currency price in Ireland. An observation is the change in prices for a plant-product-month where prices change in both markets and 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.

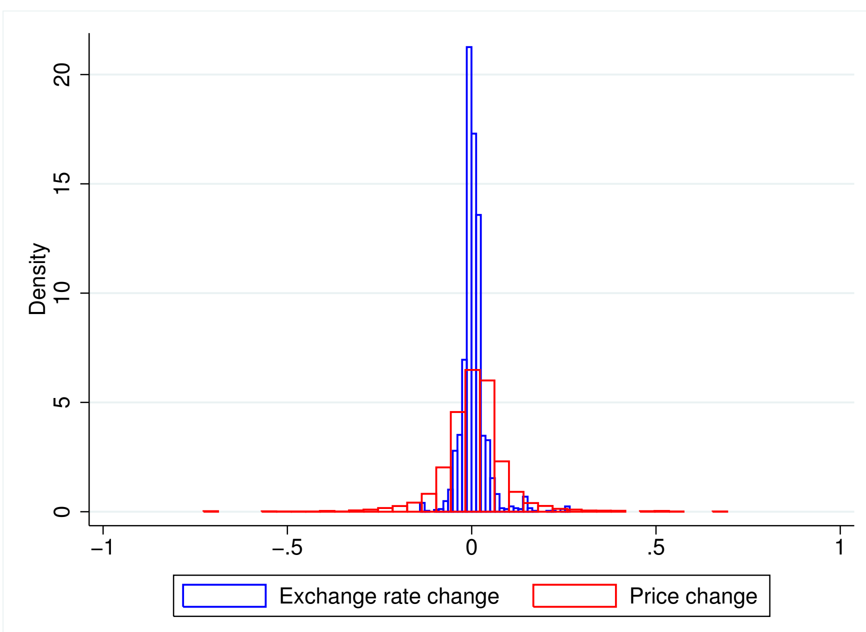

Figure 4: Size of price and exchange rate changes used to identify the intensive margin
Note: Figure plots histograms of the log changes in destination currency prices and the log changes in exchange rates for the sample used in the intensive margin estimation.