

A Framework for Identifying the Sources of Local-Currency Price Stability with an Empirical Application*

Pinelopi Koujianou Goldberg
Department of Economics
Princeton University, NBER and BREAD

Rebecca Hellerstein
Research Department
Federal Reserve Bank of New York

March 2008

Abstract

The inertia of the local-currency prices of traded goods in the face of exchange-rate changes is a well-documented phenomenon in International Economics. This paper develops a framework for identifying the sources of local-currency price stability. The empirical approach exploits manufacturers' and retailers' first-order conditions in conjunction with detailed information on the frequency of price adjustments in response to exchange-rate changes, in order to quantify the relative importance of markup adjustment by manufacturers and retailers, local-cost non-traded components, and nominal price rigidities, in the incomplete transmission of exchange-rate changes to prices. The approach is applied to micro data from the beer market. We find that on average, 54.1% of the incomplete exchange rate pass-through is due to local non-traded costs; 33.7% to markup adjustment; and 12.2% to the existence of price adjustment costs.

*We are grateful to Charles Engel for very useful conversations at an early stage of this project. We have also benefited from the comments of seminar participants at Harvard, NYU, Princeton, the University of Minnesota, the Society for Economic Dynamics, the "Microeconomic Sources of Real Exchange-Rate Behavior" conference at Carnegie Mellon University, September 16-17, 2005, the "Firms' Price Choices: Exchanging Insights between Industrial Organization, Marketing Science and Macroeconomics" conference at the University of Chicago on December 15, 2006, and the International Finance and Macroeconomics meeting of the NBER Summer Institute in 2007. This draft was prepared while Goldberg was visiting the Federal Reserve Bank of New York, which she thanks for hospitality and support. The responsibility for any errors contained in this draft is of course our own. The views expressed in this paper are those of the authors, and do not necessarily reflect the position of the Federal Reserve Bank of New York or the Federal Reserve System.

1 Introduction

The incomplete transmission of exchange-rate shocks to the prices of imported goods has been the focus of a substantial amount of theoretical and empirical research. In his 2002 article in the *NBER Macroeconomics Annual*, Engel extensively discusses this research and identifies three potential sources for the incomplete exchange-rate pass-through: the existence of local costs (e.g., costs for non-traded services) even among goods that are typically considered to be “traded”; markup adjustment on the part of retailers and/or manufacturers; and pure nominal price rigidities (at times also referred to as “menu costs”) that lead to what Engel has labelled “local-currency-pricing”. Despite the significant amount of work and interest in this topic, evidence on the relative importance of each of the contributing factors remains mixed, in part because some of the key variables needed to identify these factors, such as markups or local costs, are not directly observable, especially not in aggregate data. Yet, in an era characterized by a continuing devaluation of the dollar against other major currencies, concerns about the impact of China’s exchange-rate policy on domestic prices, and general uncertainty about the effect of exchange rates on the unwinding of current imbalances, it is more important than ever to understand why import prices do not fully respond to exchange-rate changes, especially since different explanations have very different implications for exchange rate policy.¹

Aided by the increased availability of micro data sets, a set of recent studies has focused on the microeconomics of the cross-border transmission process, trying to identify the relative contribution of each of the sources of this price inertia within structural models of particular industries. The advantage of these studies is that the institutional knowledge of the industry can be used to inform modeling assumptions, which, applied to detailed consumer or product-level data, can deliver credible estimates of markups and local costs. The disadvantage is that the results are not generalizable without further work on other markets. Still, as we show below, the few studies available to date have been able to identify interesting empirical patterns that are surprisingly robust across markets, time, and specific modeling assumptions.

The general structure of the approach proposed in this strand of the literature is as follows. The starting point is an empirical model of the industry under consideration. The model has three elements: demand, costs, and equilibrium conditions. The demand side is estimated first, independently of the supply side, using either consumer level data on individual transactions, or product level data on market shares and prices. On the supply side, the cost function of a producer selling in a foreign country is specified in a way that allows for both a traded, and a non-traded, local (i.e., destination-market specific) component in this producer’s costs. The

¹A recent speech that summarizes policymakers’ concerns about these issues can be found at: <http://www.federalreserve.gov/newsevents/speech/mishkin20080307a.htm>.

distinction between traded and non-traded costs is based on the currency in which these costs are paid. Traded costs are by definition incurred by the seller in her home country. As such, they are subject to shocks caused by variation in the nominal exchange rate when they are expressed in the destination market currency. In contrast, non-traded costs are defined as those costs that are not be affected by exchange rate changes. Costs are treated as unobservable. Assuming that firms act as profit maximizers, the market structure of the industry in conjunction with particular assumptions regarding firms' strategic behavior imply a set of first-order conditions. Once the demand side parameters are estimated, these first-order conditions can be exploited to back out the marginal costs and markups in the industry. Based on the specified cost function, marginal costs are further decomposed into a traded and non-traded component.

With this decomposition in place, one then examines how the particular components of prices (traded cost component, non-traded cost component, and markup) respond to exchange-rate changes. The lack of price response is accordingly attributed to either markup adjustment, or to the existence of a local, non-traded cost component. While the results of this decomposition naturally vary by industry, it seems that existing studies are in agreement that markup adjustment is a big part of the story. The observed exchange-rate pass-through is however too low to be explained by markup adjustment alone; accordingly, the role attributed to non-traded costs in explaining the incomplete price response is non-trivial.

While the above framework allows one to evaluate the relative contributions of markup adjustment and non-traded costs in explaining incomplete exchange-rate pass-through, it is inherently unsuitable to assessing the role of the third potential source of the incomplete price response: the existence of fixed costs of repricing. There are two reasons for this inadequacy. The first reason is a conceptual one. A key element of the framework described above is the premise that the firms' first-order conditions hold every period. Given that *by assumption* firms are always at the equilibrium implied by their profit maximizing conditions, there is no role in this framework for price adjustment costs that would cause firms to (temporarily) deviate from their optimal behavior. The second reason is a practical one. Because the data used in most previous studies are either annual or monthly, and because they are often the outcome of aggregation across more disaggregate product categories, we observe product level prices changing in every period. But with prices adjusting every period, it is inherently impossible to identify potential costs of repricing, which by nature imply that prices should remain fixed. Hence, to the extent that such price rigidities are present, these may be masked by the aggregation across different product lines, and across shorter time periods (e.g., weeks), over which nominal prices may exhibit inertia. This may lead one to overstate the role of non-traded services: whatever portion of incomplete pass-through cannot be accounted for by markup adjustment, will by construction be attributed to non-traded costs, when in reality (and in a more general framework) it could be due to the existence of price

adjustment costs.

The current paper attempts to overcome this shortcoming by explicitly introducing price rigidities into the model and suggesting an approach for quantifying their importance in explaining the documented incomplete cross-border transmission of exchange-rate shocks. To this end, we introduce two new elements.

The first one is to modify the standard framework of profit maximization to allow firms to deviate from their first-order conditions due to the existence of fixed costs of repricing. In this context we define costs of repricing in the broadest possible sense as all factors that may cause firms to keep their prices constant, and hence potentially deviate from the optimum implied by static profit maximization. Such factors may include the small costs of re-pricing (the so-called “menu-costs”) as well as the more substantive costs associated with the management’s time and effort in figuring out the new optimal price, the additional costs of advertising and more generally communicating the price change to the consumers, and – to the extent that one wants to incorporate dynamic considerations in the analysis – the option value of keeping the price unchanged in the face of ongoing uncertainty.

The second innovation of the paper is on the data side. In order to identify the potential role of nominal price rigidities we propose using higher frequency (weekly or bi-weekly) data on the prices of highly disaggregate, well-defined product lines. The advantage of using high-frequency data is that we observe many periods during which the price of a product remains utterly unchanged, followed by a discrete jump of the price to a new level. It is this discreteness in the price adjustment that we exploit in order to identify the role of nominal price rigidities.

The basic idea behind our approach is as follows. First, even with nominal price rigidities, we can estimate the demand and cost parameters of the model along the lines described in earlier papers by constraining the estimation to the periods for which we observe price adjustment; the underlying premise is that once a firm decides to incur the adjustment cost associated with a price change, it will set the product’s price according to the first-order conditions of its profit maximization problem. This of course does not imply that this firm’s behavior will be unaffected by the existence of price rigidities. Such rigidities may still have an indirect effect on the pricing behavior of firms that adjust their prices, as in any model of oligopolistic interaction firms take the prices (or quantities) of their competitors into account; if the competitor prices do not change in particular period (possibly because of price rigidities), this will affect the pricing behavior of the firms that do adjust prices. The estimation procedure takes this indirect effect into account.

Once the model parameters are estimated, we exploit information from both the periods in which prices adjust and periods in which prices remain unchanged to derive bounds on the adjustment costs associated with a price change. Our approach is based on the insight that in periods in which prices change, it has to be the case that the costs of price adjustment are lower

than the additional profit the firm makes by changing its price; we can use this insight to derive an upper bound of this price adjustment cost. Similarly, in periods in which prices do not change, it has to be the case that the costs of adjustment exceed the extra profit associated with a price change; based on this insight, we can derive a lower bound for the price adjustment cost.

The costs of price adjustment are a concept that has a precise meaning *within the context of our model*; they are defined in the broadest possible sense as everything that prevents a firm from adjusting its price in a particular period, including concerns about losing customers in the future, rather than the literal labor or material costs a firm has to pay to change prices. As such, they are not directly comparable to estimates obtained in earlier studies using different methods (e.g., direct measurement or firm surveys).² More importantly, the adjustments costs alone do not allow a full assessment of the impact of nominal price rigidities on exchange-rate pass-through; because such rigidities have both a direct and an indirect (operating through the competitor prices) effect on firms' pricing behavior, it is possible that very small rigidities induce significant price inertia. To provide an overall assessment of the impact of price adjustment costs we therefore perform simulations that compare the pricing behavior with price rigidities to the one that would prevail with fully flexible prices. The differential response of prices in the two scenarios is attributed to the effect of nominal price rigidities. In the same step we also identify the role of markup adjustment and non-traded costs in generating incomplete pass-through.

We apply the framework described above to weekly, store-level data for the beer market. The beer market is well suited for investigating questions related to exchange-rate pass-through and price rigidities for several reasons: (1) a significant fraction of brands are imported and hence affected by exchange-rate fluctuations; (2) exchange-rate pass-through onto consumer prices is low, on the order of 7-10%; (3) there exist highly disaggregate, weekly data on both wholesale and retail prices; this allows us to examine how prices respond at each stage of the distribution chain; (4) both non-traded local costs and price rigidities are a-priori plausible; in particular, weekly data reveal that both wholesale and retail prices remain constant over the course of several weeks, suggesting the existence of price rigidities. The framework we propose is however not tailored to the beer market, and can be more generally applied to any market for which high frequency data are available so that the points of price adjustment can be identified.

Our analysis yields several interesting findings. First, at the descriptive level, we document

²There have been several attempts in the literature to measure price adjustment costs. Levy et al (1997) find menu costs to equal 0.70 percent of supermarkets' revenue from time-use data. Dutta et al (1999) find menu costs to equal 0.59 percent of drugstores' revenue. Levy et al have four measures of menu costs: 1. the labor cost to change prices; 2. the costs to print and deliver new price tags; 3. the costs of mistakes; 4. the costs of in-store supervision of the price changes. Some detailed microeconomic studies have cast doubt on the importance of menu cost in price rigidity. Blinder et al (1998) find in a direct survey that managers do not regard menu costs as an important cause of price rigidity. Both Carlton (1986) and Midrigan (2006) find that firms change prices frequently and in small increments, which is not consistent with a menu-cost explanation of price rigidity.

infrequent price adjustment both at the wholesale and retail level. However, this price inertia seems to be primarily driven by the infrequent adjustment of wholesale rather than retail prices. In our data, there is not a single instance where a product’s retail price remains unchanged in response to a wholesale price change. In contrast, there are several instances in which the retail price changes, while the wholesale price does not. Hence it seems that the primary reason that retail prices do not change from period to period is that there is little reason for them to change, as the underlying wholesale prices remain fixed.

Perhaps not surprisingly given the above documented price adjustment patterns, we estimate price adjustment costs at the retail level to be low relative to such costs at the wholesale level. In particular, employing a procedure analogous to Levy et al (1997) to calculate menu costs as a percentage of firm revenue³, we estimate that price adjustment costs represent on average 0.1% of the revenue at the retail level, while they account for approximately 0.5% of the revenue at the wholesale level.

As we discussed above, nominal price rigidities may affect the pricing decisions of a particular producer in two ways. First, they may prevent this producer from adjusting her price, because her *own* costs of repricing exceed the benefits, even when all other competing producers adjust their prices (direct effect). Second, such costs may induce *other* competing producers to keep their prices fixed, which may make price adjustment less profitable for the producer under consideration (indirect/strategic effect). Our simulations indicate that the direct effect is significant at the wholesale level, accounting for 10.5% on average for the incomplete pass-through. Interestingly, there is substantial variation in this estimate across brands; the *own* costs of price adjustment appear to be more important for brands with low market shares, such as Bass and Beck’s than for brands with large market shares such as Corona and Heineken. In contrast, we find that at the retail level the *own* costs of repricing have no effect; this is perhaps not surprising given the small magnitude of repricing costs we estimate at that level. There is however an *indirect/strategic* effect at this stage of the distribution chain that accounts for approximately 0.1% of the incomplete pass-through. Overall it seems that the *direct* effect of repricing costs is only present at the wholesale level, while the *indirect/strategic* effect plays a role both at the wholesale and retail levels. Our final decomposition attributes 54.1% of the incomplete pass-through to local non-traded costs (52.5% at the wholesale and 1.6% at the retail level); 33.7% to markup adjustment (33.5% at the wholesale and 0.2% at the retail level); and 12.2% to the existence of price adjustment costs, 1.7%

³This procedure involves summing up the estimated upper bounds for the price adjustment costs for those periods within a year in which prices *did* change, and then dividing this sum by the firm’s revenue over the entire year. We should emphasize that because our “price adjustment costs” are defined in the most general sense to include *all* factors that may prevent firms from changing their nominal prices (and *not just the literal labor and material costs* of changing prices), our numbers are not directly comparable to Levy et al’s. Still, it is interesting that despite these differences the two sets of numbers are of similar order of magnitude.

of which represent the indirect/strategic effect of such costs. As suspected at the beginning, costs of price adjustment appear to be substantially more important at the wholesale than retail level.

The approach we just described is static. While estimating a full dynamic model would of course be desirable, the complexity of our setup and model does not allow us to do so (at least not without simplifying parts of the model that we consider essential to our analysis, such as the strategic interactions between domestic and foreign firms as well as the retailer and manufacturers, the flexibility of the demand substitution patterns, and the consistent treatment of the cost side). However, we believe that there are two particular features of this market that allow us to abstract from dynamics without sacrificing the realism and relevance of our model. First, neither consumers nor sellers hold inventories for beer.⁴ Second, the primary source of cost variation facing beer sellers and focus of our analysis is the *exchange rate*, which is highly persistent.⁵ This implies that the premise of static expectations is not unreasonable here. If the exchange rate were the only source of cost variation, and exchange rate shocks were perceived as permanent, the static model would be equivalent to a dynamic one up to a discount factor: firms would be maximizing the sum over the (discounted) profits of each period, which in the absence of any cross-period price effects or inventories would be equivalent to maximizing period-by-period profits (Dixit 1991). If the exchange rate is the primary (but not only) source of cost variation and it is highly persistent, our model provides a reasonable approximation to a dynamic setup. We discuss these issues formally in Section 3.3.

The remainder of the paper is organized as follows. To set the stage, we start by providing a brief description of the market and the data in the next section; in the same section, we also provide some descriptive statistics and discuss the price adjustment patterns evident in the retail and wholesale price data. Section 3 discusses the model and shows how it allows us to derive bounds for the price adjustment costs. Section 4 discusses the steps of the empirical implementation of the model in detail. Section 5 presents the estimation and simulation results, and Section 6 concludes.

2 The Market and the Data

In this section we describe the market our data cover. We present summary statistics and some preliminary descriptive results to build intuition for the results from the structural model. We then discuss some of the price-adjustment patterns in the data.

⁴The role of inventories has been emphasized in several recent papers by Hendel and Nevo (2006a and 2006b), Pesendorfer (2002), Aguirregabiria (1999), and Alessandria, Kaboski and Midrigan (2008). We refer the reader to section 3.3 for a detailed discussion of how the arguments put forth in these papers affect our analysis.

⁵The consensus among International economists is that the exchange rate follows a random walk.

2.1 Market

The imported beer market first developed in the U.S. in the nineteenth century. As late as 1970, imported beers made up less than one percent of the total U.S. consumption of beer. Consumption of imported brands grew slowly in the 1980s and by double digits for each year in the 1990s – on average by 11 percent per year from 1993 to 2001. Beer is an example of one type of imported goods: packaged goods imported for consumption. Such imports do not require any further manufacture before reaching consumers and make up roughly half of the non-oil goods imports to the U.S. over the sample period.

The beer market is well suited for an exploration of the sources of local-currency price stability for the reasons discussed in the introduction: a significant fraction of brands are imported; exchange-rate pass-through to prices is generally low (between eight and ten percent); both non-traded local costs and price stickiness due to adjustments costs are *a-priori* plausible; last but not least, we have a rich panel data set with weekly retail and wholesale prices. It is unusual to observe both retail and wholesale prices for a single product over time. These data enable us to separate us to isolate the role of local non-traded costs and of fixed adjustment costs in firms' incomplete transmission of exchange-rate shocks to prices.

2.2 Data

Our data come from *Dominick's Finer Foods*, the second-largest supermarket chain in the Chicago metropolitan area in the mid 1990s with a market share of roughly 20 percent. The data record the retail and wholesale prices for each product sold by *Dominick's* over a period of four years. They were gathered by the *Kilts Center for Marketing* at the University of Chicago's Graduate School of Business and include aggregate retail volume market shares and retail prices for every major brand of beer sold in the U.S.⁶ Beer shipments in this market are handled by independent wholesale distributors. The model we develop in the next section of the paper abstracts from this additional step in the vertical chain, and assumes distributors are vertically integrated with brewers, in the sense that brewers bear their distributors' costs and control their pricing decisions. It is common knowledge in the industry that brewers set their distributors' prices through a practice known as *resale price maintenance* and cover a significant portion of their distributors' marginal costs.⁷ This practice makes the analysis of pricing behavior along the distribution chain

⁶The data can be found at <http://gsbwww.uchicago.edu/kilts/research/db/dominicks/>.

⁷Features of the *Dominicks'* wholesale-price data confirm that brewers control distributors prices to the supermarket. Across individual *Dominicks'* stores, which may each be served by a different distributor, each with an exclusive territory, the variation in *UPC*-level wholesale prices is less than one cent. Asker (2004) notes that one cannot distinguish distributors by observing the wholesale prices they charge to individual *Dominicks* stores. This supports the industry lore that distributors pricing is coordinated by brewers and is not set separately by each distributor to each retail outlet.

relatively straight-forward, as one can assume that distributors are, *de facto*, vertically integrated with brewers.

During the 1990s supermarkets increased the selection of beers they offered as well as the total shelf space devoted to beer. A study from this period found that beer was the tenth most frequently purchased item and the seventh most profitable item for the average U.S. supermarket.⁸ Supermarkets sell approximately 20 percent of all beer consumed in the U.S.⁹

We aggregate data from each *Dominick's* store into one of two price zones. For more details about this procedure, see Hellerstein (2008).¹⁰ We define a product as one six-pack serving of a brand of beer and quantity as the total number of servings sold per week. We define a market as one of *Dominick's* price zones in one week. Products' market shares are calculated with respect to the potential market which is defined as the total beer purchased each week in supermarkets by the residents of the zip codes in which each *Dominick's* store is located. We define the outside good to be all beer sold by other supermarkets to residents of the same zip codes as well as all beer sales in the sample's *Dominick's* stores not already included in our sample. We have a total of 16 brands in our sample (5 domestic and 11 imported), each with 404 observations (202 weeks spanning the period from June 6, 1991 to June 1, 1995 in each of two price zones). We supplement the *Dominick's* data with information on manufacturer costs, product characteristics, advertising, and the distribution of consumer demographics. Product characteristics come from the ratings of a *Consumer Reports* study conducted in 1996. Summary statistics for the price data are provided in Table 1. Table 2 reports summary statistics for the characteristics data used in the demand estimation.

2.3 Preliminary Descriptive Results

We begin the analysis by documenting in several simple regressions whether *Dominick's* imported-beer prices are systematically related to movements in bilateral nominal exchange-rates. These results can provide a benchmark against which we can measure the performance of the structural

⁸Canadian Trade Commissioner (2000).

⁹As our data focus on one metropolitan statistical area, we do not need to control for variation in retail alcohol sales regulations. Such regulations can differ considerably across states.

¹⁰The zones are defined by *Dominick's* mainly on the basis of customer demographics. Although they do not report these zones, we identify them through zip-code level demographics (with a few exceptions, each *Dominick's* store in our sample is the only store located in its zip code) and by comparing the average prices charged for the same product across stores. We classify each store according to its pricing behavior as a low- or high-price store and then aggregate sales across the stores in each pricing zone. This aggregation procedure retains some cross-sectional variation in the data which is helpful for the demand estimation. Residents' income covaries positively with retail prices across the two zones.

model. We estimate three price equations:

$$(1) \quad \ln p_{jzt}^r = c_j + \zeta_z + \theta_t + \alpha \ln e_{jt} + \beta \ln co_{jt} + \varepsilon_{jzt}$$

$$(2) \quad \ln p_{jzt}^w = c_j + \zeta_z + \theta_t + \alpha \ln e_{jt} + \beta \ln co_{jt} + \varepsilon_{jzt}$$

$$(3) \quad \ln p_{jzt}^r = c_j + \zeta_z + \theta_t + \alpha \ln p_{jzt}^w + \varepsilon_{jzt}$$

where the subscripts j , z , and t refer to product, zone, and week respectively; p^r is the product's retail price; p^w is the product's wholesale price; c_j , θ_t , and ζ_z are product, week and zone dummies respectively that proxy among other things for demand shocks that may affect a brand's price independent of exchange rates; e is the bilateral nominal exchange rate (domestic-currency units per unit of foreign currency); co_{jt} denotes a set of variables that proxy for cost shocks that again may affect prices; such variables include measures of domestic (U.S.) wages, the price of barley in each country producing beer in our sample, the price of electricity in the Chicago area, and - for foreign brands - wages in each beer exporting country in our sample; ε is a random error term. All variables are specified in levels, and not first differences, as our focus is on the long-run pass-through of exchange rate changes, and not the short-term dynamics.

Table 3 reports results from OLS estimation of the pricing equations. Columns 1 and 3 report results from specifications that include the full set of controls specified above, while in columns 2 and 4 the cost controls are omitted (since the latter do not vary at the weekly level). The results across the two specifications are remarkably similar. The average pass-through elasticity α for the retail price is - based on column 2 - 6.7 percent and is significant at the one-percent level. The regression establishes a roughly 7-percent benchmark for the retailer's pass-through elasticity, that we will try to explain within the framework of the structural model. The fourth column of Table 3 reports similar results from estimation of the wholesale-price pricing equation, equation (2): Its pass-through elasticity is 4.7 percent, and the coefficient is again highly significant. Finally, the fifth column of Table 3 reports the results from an *OLS* regression of each brand's retail price on its own wholesale price. The coefficient on the wholesale price is not significantly different from 100, which is consistent with the results from the other columns: Exchange-rate shocks that are passed on by manufacturers to the retailer appear to be immediately and almost fully passed on to consumer prices.

This preliminary analysis reveals that local-currency price stability is an important feature of this market: only around 7 percent of an exchange-rate change is transmitted to a beer's retail price. Where does the other 93 percent go? Existing literature on exchange rate pass-through has

identified three potential sources of this incomplete transmission: a non-traded cost component in the manufacturing of traded goods, variable markups, and nominal price rigidities. The goal of our paper is to quantify the relative contribution of each of these sources in explaining incomplete pass-through.

2.4 Patterns of Price Adjustment in the Data

A rough idea of the timing and frequency of price changes in the beer market can be obtained from Figure 1, which plots the retail and wholesale prices for a six-pack of the British brand *Bass Ale*. The figure covers the full sample period, from the middle of 1991 to the middle of 1995. The plot serves to illustrate several interesting points.

First, the figure demonstrates the advantage of observing price data at a weekly frequency. Such data are ideal for analyzing the role of price stickiness, since we clearly see prices remaining constant for several weeks, and then jumping up (in a discrete step) to a new level. This pattern in the price adjustment process is exactly the one we would expect with price stickiness. That said, the infrequent adjustment of prices is by itself no definitive proof that price rigidities exist, as it is in principle possible that prices do not change simply because nothing else changes.

Second, a substantial fraction of the price variation reflects temporary price reductions (sales). To investigate if there was a systematic pattern in sales, we estimated sales determinant specifications similar to those in Pesendorfer (2002). Interestingly, in our sample of beer brands, sales appear to be random, in the sense that we did not find *anything* that could predict the timing of a sale. In particular, we did not find that the time that had elapsed since a previous sale, holiday dummies, or sales for other brands, could predict a current sale for a particular brand. This apparent randomness of sales in our sample is consistent with a Varian-type explanation of temporary price reductions. A potential problem with such an explanation is that in Varian's model firms randomize prices each period, so that the notion of a "regular price" does not exist. However, in the Varian (1980) model there are no costs of price adjustment. Introducing costs of price adjustment (as we do in this paper) can explain the existence of "regular prices" that is evident in Figures 1 and 2. We should point out that we by no means claim that the patterns we observe in our data for beer generalize to other product categories. In fact, there is substantial evidence, based on Pesendorfer (2002) and recent papers by Hendel and Nevo (2006a and 2006b), that intertemporal considerations, such as those emphasized in models by Sobel (1984) and Pesendorfer (2002), are important in determining sales in other products that are more storable. Our findings regarding the timing of sales for beer are consistent with our premise throughout the paper that consumers do not store beer, so that intertemporal considerations, while potentially important in other markets, are not a first order concern here. We discuss this issue in more

detail in Section 3.3.

Third, a striking feature of Figure 1 is that retail prices always adjust when wholesale prices adjust. So it seems that the main reason retail prices do not change in this market is that there is little reason for them to change (the cost facing retailers as measured by the wholesale price does not change). This is to be contrasted with the pattern we observe at the wholesale level: despite enormous variation in exogenous (to the industry) factors affecting manufacturer costs (i.e., exchange-rate fluctuations), wholesale prices remain unchanged for long periods of time.

A final point that Figure 1 together with similar plots for other brands illustrates is that price adjustment is not synchronized across brands. Given the strategic interactions between firms, this asynchronous price adjustment can generate significant price inertia, even if the nominal price rigidities facing each individual manufacturer or retailer are estimated to be small.

3 Model

This section describes the supply and demand sides of the model we use to identify the sources of incomplete exchange rate pass-through, and in particular the role of price rigidities.

3.1 Supply

We model the supply side of the market using a linear-pricing model in which manufacturers, acting as Bertrand oligopolists with differentiated products, set their prices followed by retailers who set their prices taking the wholesale prices they observe as given. Thus, a double margin is added to the marginal cost of the product before it reaches the consumer. Our framework builds on Hellerstein’s (2008) work on the beer market, but makes two modifications to her model: First, we introduce price rigidities both at the wholesale and retail level; the effect of these price rigidities is to cause firms to potentially deviate from their first-order conditions. Second, to keep the framework as simple and transparent as possible, we model both retailers and manufacturers as single-product firms. While this assumption may be hard to defend, especially in the context of the retailers, it is not essential for the approach we propose in order to identify price rigidities, and can be relaxed in the future.¹¹

The strategic interaction between manufacturer and retailer is as follows. First, the manufacturer decides whether or not to change the product’s price taking into account the current period’s observables (costs, demand conditions, and competitor prices), and the anticipated reaction of the retailer. If she decides to change the price, then the new price is determined based on

¹¹The assumption of single-product retailers would however be valid if manufacturers were able to enact vertical restraints, hence exercising control over retailers’ brand-level pricing and promotional decisions. In this case, retailers will act as if they were single-product firms with respect to each brand.

the manufacturer’s first-order conditions. Otherwise the wholesale price is the same as in the previous period. Next, the retailer observes the wholesale price set by the manufacturer and decides whether or not to change the product’s retail price. If the retail price changes, then the new retail price is determined according to the retailer’s first-order conditions. Otherwise the retail price is the same as in the previous period. To characterize the equilibrium we use backward induction and solve the retailer’s problem first.

3.1.1 Retailer

Consider a retail firm that sells all of the market’s J differentiated products. Let all firms use linear pricing and face constant marginal costs. The profits of the retail firm associated with selling product j at time t are given by:

$$(4) \quad \Pi_{jt}^r = (p_{jt}^r - p_{jt}^w - ntc_{jt}^r) s_{jt}(p_t^r) - A_{jt}^r$$

The first part of the profit expression is standard. The variable p_{jt}^r is the price the retailer sets for product j , p_{jt}^w is the wholesale price paid by the retailer for product j , ntc_{jt}^r are local non-traded¹² costs paid by the retailer to sell product j , and $s_{jt}(p_t^r)$ is the quantity demanded of product j which is a function of the prices of all J products. The new element in our approach is the introduction of the second term, A_{jt}^r , which captures the fixed cost of changing the price of product j at time t . This cost is zero if the price remains unchanged from the previous period, but takes on a positive value, known to the retailer, but unknown to the econometrician, if the price adjusts in the current period:

$$(5) \quad \begin{aligned} A_{jt}^r &= 0 \text{ if } p_{jt}^r = p_{jt-1}^r \\ A_{jt}^r &> 0 \text{ if } p_{jt}^r \neq p_{jt-1}^r \end{aligned}$$

We interpret the adjustment cost A_{jt}^r as capturing all possible sources of price rigidity. These can include the management’s cost of calculating the new price; the marketing and advertising expenditures associated with communicating the new price to customers; the costs of printing and posting new price tags, etc...The particular interpretation of A_{jt}^r is not important for our purposes. What is important is that this cost is independent of the sales volume; it is a discrete cost that the retailer pays every time the price adjusts from the previous period. The indexing of A by product j and time t in our notation corresponds to the most flexible specification, in which the price adjustment cost is allowed to vary by product and time. One could potentially impose

¹²We use the term “non-traded” to indicate that these costs are paid in dollars no matter what the origin of the product is. Hence, non-traded costs will not be affected by exchange rate shocks.

more structure by assuming that adjustment costs are constant over time, and/or constant across products.

The implication of the adjustment cost in the profit function is that it can cause firms to deviate from their first-order conditions, even if the retailer acts as a profit maximizer. Specifically, in the data we will observe one of two cases:

Case 1: The price changes from the previous period, that is $p_{jt}^r \neq p_{jt-1}^r$.

In this case the retailer solves the standard profit maximization problem to determine the new optimal price, and the observed retail price p_{jt}^r will have to satisfy the first-order profit-maximizing conditions:

$$(6) \quad s_{jt} + (p_{jt}^r - p_{jt}^w - ntc_{jt}^r) \frac{\partial s_{jt}}{\partial p_{jt}^r} = 0, \text{ for } j = 1, 2, \dots, J_t.$$

This gives us a set of J equations, one for each product. One can solve for the markups by defining a $J \times J$ matrix Ω_{rt} , called the retailer reaction matrix, with all off-diagonal elements equal to zero, and the diagonal elements equal to $S_{jj} = \frac{\partial s_{jt}(p_{jt}^r)}{\partial p_{jt}^r}$ $j = 1, \dots, J$, that is the marginal change in the j th product's market share given a change in the j th product's retail price. The stacked first-order conditions can be rewritten in vector notation:

$$(7) \quad s_t + \Omega_{rt}(p_t^r - p_t^w - ntc_t^r) = 0$$

and inverted together in each market to get the retailer's pricing equation, in vector notation:

$$(8) \quad p_t^r = p_t^w + ntc_t^r - \Omega_{rt}^{-1} s_t$$

where the retail price for product j in market t will be the sum of its wholesale price, non-traded costs, and markup.

The presence of the adjustment costs A_{jt}^r in the profit function implies that for the retailer to change her price in the current period, it will have to be the case that the extra profits associated with the new price are at least as large as the adjustment cost:

$$(9) \quad (p_{jt}^r - p_{jt}^w - ntc_{jt}^r) s_{jt}(p_{jt}^r) - A_{jt}^r \geq (p_{jt-1}^r - p_{jt}^w - ntc_{jt}^r) s_{jt}^c(p_{jt-1}^r, p_{kt}^r), \quad k \neq j$$

where $s_{jt}^c(p_{jt-1}^r, p_{kt}^r)$ denotes the *counterfactual* market share that product j *would* have, if the retailer had kept the price unchanged to p_{jt-1}^r , and p_{kt}^r denotes the prices of the other products k that may or may not have changed from the previous period. The above inequality simply states that the profits the retailer makes by adjusting the price of product j in the current period have to be greater than the profits the retailer *would have* achieved, if she had not changed the price

(in which case the first-order condition of profit maximization would have been violated, but the retailer would have saved on the adjustment costs A_{jt}^r). By rearranging terms we can use the above inequality to derive an *upper* bound $\overline{A_{jt}^r}$ for the price adjustment costs of product j :

$$(10) \quad A_{jt}^r \leq \overline{A_{jt}^r} = (p_{jt}^r - p_{jt}^w - ntc_{jt}^r) s_{jt}(p_t^r) - (p_{jt-1}^r - p_{jt}^w - ntc_{jt}^r) s_{jt}^c(p_{jt-1}^r, p_{kt}^r), \quad k \neq j$$

Case 2: The price remains unchanged from the previous period, that is $p_{jt}^r = p_{jt-1}^r$.

In this case the first-order conditions of profit maximization do not necessarily hold. If the retailer does not adjust the price of product j in period t , it must be the case that the profits she makes from keeping the price constant are at least as large as the profits the retailer would have made if she had adjusted the price according to the first-order condition minus the adjustment costs associated with the price change:

$$(11) \quad (p_{jt-1}^r - p_{jt}^w - ntc_{jt}^r) s_{jt}(p_{jt-1}^r, p_{kt}^r) \geq (p_{jt}^{rc} - p_{jt}^w - ntc_{jt}^r) s_{jt}^c(p_{jt}^{rc}, p_{kt}^r) - A_{jt}^r, \quad k \neq j$$

where p_{jt}^{rc} denotes the *counterfactual* price the retailer *would have* charged if he behaved according to the optimality conditions, and $s_{jt}^c(p_{jt}^{rc}, p_{kt}^r)$ is the *counterfactual* market share that would correspond to this optimal price holding the prices of the competitor products at their observed levels. Just like in Case 1, we can rewrite the above inequality to derive a *lower* bound $\underline{A_{jt}^r}$ for the adjustment costs:

$$(12) \quad A_{jt}^r \geq \underline{A_{jt}^r} = (p_{jt}^{rc} - p_{jt}^w - ntc_{jt}^r) s_{jt}^c(p_{jt}^{rc}, p_{kt}^r) - (p_{jt-1}^r - p_{jt}^w - ntc_{jt}^r) s_{jt}(p_{jt-1}^r, p_{kt}^r), \quad k \neq j$$

The essence of our empirical approach to quantify the adjustment costs can be described as follows. First, we estimate the demand function. Once the demand parameters have been estimated, the market share function $s_{jt}(p_t^r)$ as well as the own and cross price derivatives $\frac{\partial s_{jt}}{\partial p_{jt}^r}$ and $\frac{\partial s_{kt}}{\partial p_{jt}^r}$ can be treated as known. Next we exploit the first-order conditions for each product j (6) to estimate the non-traded costs and markups of product j , but contrary to the approach typically employed in the Industrial Organization literature, we use *only the periods in which the price of product j adjusts*, to back out costs and markups. In periods when the price does not adjust, the non-traded costs are not identified based on the first-order conditions; however, we can derive estimates of the non-traded costs for these periods by imposing some additional structure on the problem, e.g., by modeling non-traded costs parametrically as a function of observables along the lines described in the next section. Once estimates of non-traded costs for these periods have been derived, we can calculate the counterfactual price p_{jt}^{rc} that the retailer would have charged

if there were no price rigidities and she behaved according to the profit maximization conditions, as well as the associated counterfactual market share $s_{jt}^c(p_{jt}^c, p_{kt}^c)$. In the final step, we can exploit inequalities (10) and (12) to derive upper and lower bounds of the adjustment costs A_{jt}^r .

Note that in the above framework price rigidities as captured by the adjustment cost A_{jt}^r affect pricing behavior in two ways. First, there is a direct effect: price rigidities may prevent the retailer from adjusting the price of any particular product if the adjustment cost associated with *this product's* price change exceeds the additional profit. Second, there is an indirect effect that operates through the effect that price rigidities may have on the prices of competing products. When our retailer sets the price of product j , she conditions on the prices of the other products with which product j competes. If these prices remain constant (potentially because of the existence of price rigidities), then the price change of product j may be smaller than the one we would have observed if price rigidities were altogether non-existent. The existence of this indirect effect implies that relatively small adjustment costs can potentially lead to significant price inertia. Accordingly, the magnitude of the adjustment costs cannot by itself provide a measure of the significance of price stickiness in explaining incomplete pass-through. To assess the overall impact of price adjustment costs it is necessary to perform simulations to compare the pricing behavior we observe to the one that would prevail with fully flexible prices.

3.1.2 Manufacturers

Let there be M manufacturers that each produce one of the market's J_t differentiated products. Each manufacturer chooses its wholesale price p_{jt}^w taking the retailer's anticipated behavior into account. Manufacturer w 's profit function is:

$$(13) \quad \Pi_{jt}^w = (p_{jt}^w - c_{jt}^w(tc_{jt}^w, ntc_{jt}^w)) s_{jt}(p_t^r(p_t^w)) - A_{jt}^w$$

where c_{jt}^w is the marginal cost incurred by the manufacturer to produce and sell product j ; this cost is in turn a function of traded costs tc_{jt}^w , and destination-market specific non-traded costs ntc_{jt}^w . As noted above, the distinction between traded and non-traded costs is based on the currency in which these costs are paid; traded costs are by definition incurred in the manufacturer's home country currency, and are subject to exchange rate shocks, while (dollar-denominated) non-traded costs are not. The term A_{jt}^w denotes the price adjustment cost incurred by the manufacturer. The interpretation of this cost is similar to the one for the retail adjustment cost; it is a discrete cost that is paid only when the manufacturer adjusts the price of product j :

$$(14) \quad \begin{aligned} A_{jt}^w &= 0 \text{ if } p_{jt}^w = p_{jt-1}^w \\ A_{jt}^w &> 0 \text{ if } p_{jt}^w \neq p_{jt-1}^w \end{aligned}$$

Given this structure, we can use the same procedure as the one we applied to the retailer's problem in order to derive upper and lower bounds for the manufacturer adjustment cost. The derivation of the manufacturer bounds is however more complicated as the manufacturer needs to take into account the possibility that the retailer does not adjust her price due to the existence of the retailer adjustment cost.

As with the retailer, in the data we will observe one of two cases:

Case 1: The wholesale price changes from the previous period, that is $p_{jt}^w \neq p_{jt-1}^w$.

Due to the existence of the retail adjustment cost, it is – in principle – possible in this case that the retail price does not adjust, while the wholesale price does adjust. However, in our data we do not observed a single instance of this happening. We therefore concentrate our discussion on the case where the retail price adjusts when the wholesale price adjusts.

Assuming that manufacturers act as profit maximizers, the wholesale price p_{jt}^w must satisfy the first-order profit-maximizing conditions given that it has been adjusted from the previous period:

$$(15) \quad s_{jt} + (p_{jt}^w - c_{jt}^w) \frac{\partial s_{jt}}{\partial p_{jt}^w} = 0 \text{ for } j = 1, 2, \dots, J_t.$$

This gives us another set of J equations, one for each product. Let Ω_{wt} be the manufacturer's reaction matrix with elements $\frac{\partial s_{jt}(p_t^r(p_t^w))}{\partial p_{jt}^w}$, the change in each product's share with respect to a change in each product's wholesale price. The manufacturer's reaction matrix is a transformation of the retailer's reaction matrix: $\Omega_{wt} = \Omega'_{pt} \Omega_{rt}$ where Ω_{pt} is a J -by- J matrix of the partial derivative of each retail price with respect to each product's wholesale price. Each column of Ω_{pt} contains the entries of a response matrix computed without observing the retailer's marginal costs. The properties of this manufacturer response matrix are described in greater detail in Villas-Boas (2007).¹³

The manufacturers' marginal costs (which are a function of the traded and non-traded costs, tc_t^w and ntc_t^w respectively) are then recovered by inverting the manufacturer reaction matrix Ω_{wt} according to:

$$(16) \quad p_t^w = c_t^w - \Omega_{wt}^{-1} s_t$$

For product j , the wholesale price is the sum of the manufacturer traded costs, non-traded costs, and markup function. The manufacturer of product j can use her estimate of the retailer's non-traded costs and reaction function to compute how a change in the manufacturer price will affect the retail price for the product.

¹³To obtain expressions for this matrix, one uses the implicit-function theorem to totally differentiate the retailer's first-order condition for product j with respect to all retail prices and with respect to the manufacturer's price p_j^w .

For the manufacturer to have changed her price from the previous period, it has to be the case that the profits she makes from having changing the price (net of the price adjustment cost A_{jt}^w) exceed the profits that the manufacturer *would have* made if she had left the wholesale price unchanged at p_{jt-1}^w :

$$(17) \quad (p_{jt}^w - c_{jt}^w) s_{jt}(p_t^r(p_t^w)) - A_{jt}^w \geq (p_{jt-1}^w - c_{jt}^w) s_{jt}^c(p_{jt}^{rc}(p_{jt-1}^w, p_{kt}^r)), \quad k \neq j$$

This condition is similar to inequality (6) for the retailer, with a slight difference: the counterfactual market share s_{jt}^c that the manufacturer would face if she left the price of product j unchanged is a function of the counterfactual retail price p_{jt}^{rc} that the retailer would charge when faced with an unchanged wholesale price p_{jt-1}^w . But given the existence of the retail adjustment cost, this counterfactual price can follow one of two scenarios: the first one is that the retailer does not change the price from the previous period, so that $p_{jt}^{rc} = p_{jt-1}^r$; the second possibility is that the retailer adjusts her price according to the retailer's first-order conditions (6). Hence, before one can use the above inequality to infer the upper bound of the manufacturer's adjustment cost, it is necessary to solve the retailer's problem to determine how the retailer's price response. Specifically, if:

$$(p_{jt-1}^r - p_{jt-1}^w - ntc_{jt}^r) s_{jt}^c(p_{jt-1}^r, p_{kt}^r) \geq (p_{jt}^{rc} - p_{jt-1}^w - ntc_{jt}^r) s_{jt}^c(p_{jt}^{rc}(p_{jt-1}^w), p_{kt}^r) - A_{jt}^r, \quad k \neq j$$

the retailer will leave her price unchanged. Otherwise, she will adjust her price to p_{jt}^{rc} , where p_{jt}^{rc} is itself determined according to the first-order condition:

$$s_{jt}^c + (p_{jt}^{rc} - p_{jt-1}^w - ntc_{jt}^r) \frac{\partial s_{jt}^c}{\partial p_{jt}^{rc}} = 0$$

Once the optimal pricing behavior of the retailer, conditional on the wholesale price being equal to p_{jt-1}^w has been determined, the upper bound of the *manufacturer's* adjustment cost $\overline{A_{jt}^w}$ can be derived based on the inequality:

$$(18) \quad A_{jt}^w \leq \overline{A_{jt}^w} = (p_{jt}^w - c_{jt}^w) s_{jt}(p_t^r(p_t^w)) - (p_{jt-1}^w - c_{jt}^w) s_{jt}^c(p_{jt}^{rc}(p_{jt-1}^w, p_{kt}^r)), \quad k \neq j$$

where p_{jt}^{rc} is either equal to p_{jt-1}^r or determined according to the retailer's first-order condition, and s_{jt}^c is evaluated accordingly.

Case 2: The wholesale price does not change from the previous period, that is $p_{jt}^w = p_{jt-1}^w$.

The lack of price adjustment in this case implies that the wholesale price is not necessarily

determined based on the manufacturer first-order condition. Regarding the retail price, it is again possible that the retailer adjusts the retail price in periods when the wholesale price remains unchanged. However, in practice we rarely observe this case in the data. Hence, we concentrate on the case where both wholesale and retail prices remain unchanged, that is $p_{jt}^w = p_{jt-1}^w$ and $p_{jt}^r = p_{jt-1}^r$.

Given that the manufacturer does not adjust the wholesale price, it has to be the case that the profits she makes at p_{jt-1}^w are at least as large as the profit she *would have* made if she had changed the price to a counterfactual wholesale price p_{jt}^{wc} according to the profit maximization condition and paid the associated adjustment cost A_{jt}^w :

$$(19) \quad (p_{jt-1}^w - c_{jt}^w) s_{jt}(p_{jt-1}^r(p_{jt-1}^w), p_{kt}^r) \geq (p_{jt}^{wc} - c_{jt}^w) s_{jt}^c(p_{jt}^{rc}(p_{jt}^{wc}), p_{kt}^r) - A_{jt}^w, \quad k \neq j$$

As with the case of the retailer, we can exploit this insight to derive a lower bound $\underline{A_{jt}^w}$ for the price adjustment cost A_{jt}^w :

$$(20) \quad A_{jt}^w \geq \underline{A_{jt}^w} = (p_{jt}^{wc} - c_{jt}^w) s_{jt}^c(p_{jt}^{rc}(p_{jt}^{wc}), p_{kt}^r) - (p_{jt-1}^w - c_{jt}^w) s_{jt}(p_{jt-1}^r(p_{jt-1}^w), p_{kt}^r), \quad k \neq j$$

The determination of the counterfactual optimal wholesale price p_{jt}^{wc} and the associated counterfactual market share s_{jt}^c is however more involved in this case, as the manufacturer has to take into account the reaction of the retailer, who may or may not adjust her price in response to a wholesale price change.

To find the price p_{jt}^{wc} the manufacturer would set if she were willing to incur the adjustment cost, we proceed as follows. First, we consider the case in which the retail price would have changed in response to the wholesale price change. In this case p_{jt}^{wc} would be determined according to equation (16) which reflects the manufacturer's first-order condition; the inverted manufacturer reaction matrix Ω_{wt}^{-1} in this equation incorporates the optimal pass-through of the wholesale price change onto the retail price.

Next we consider the case in which the retailer does not adjust her price in response to the wholesale price change. Even though as noted above we never observe this case in the data, the possibility that the wholesale price change does not get passed through by the retailer is factored in when manufacturers set prices. If the manufacturer anticipates an equilibrium in which the retailer does not adjust her price, the optimal manufacturer behavior will be to change the wholesale price up to the point where the retailer is just indifferent between changing the retail price and leaving it the same as in the previous period, that is:

$$(21) \quad (p_{jt-1}^r - p_{jt}^{wc} - ntc_{jt}^r) s_{jt}^c(p_{jt-1}^r, p_{kt}^r) = (p_{jt}^{rc} - p_{jt}^{wc} - ntc_{jt}^r) s_{jt}^c(p_{jt}^{rc}, p_{kt}^r) - A_{jt}^r, \quad k \neq j$$

The left hand side of the above equation denotes the profits the retailer would make if she did not pass-through the change in the wholesale price. The right hand side represents the profits the retailer would make if she changed the retail price to p_{jt}^r , where the latter is determined based on the retailer's first-order condition $s_{jt}^c + \left(p_{jt}^r - p_{jt}^{wc} - ntc_{jt}^r \right) \frac{\partial s_{jt}^c}{\partial p_{jt}^r} = 0$. To find the wholesale price p_{jt}^{wc} the manufacturer would charge in this case, equation (21) can be solved simultaneously with the retailer's first-order condition for p_{jt}^{wc} and p_{jt}^r .

The final step in determining the counterfactual optimal wholesale price p_{jt}^{wc} that the manufacturer would choose if she changed the wholesale price from the previous period is to compare the manufacturer profits for the case where the retailer adjusts the price, to the manufacturer profits for the case where the retailer does not pass-through the wholesale price change, in which case the wholesale price will be set according to (21). The manufacturer will pick the p_{jt}^{wc} that corresponds to the higher profits. Once the wholesale price is found, the optimal retail price response and associated market share can be determined as well, and inserted in (20) in order to infer the manufacturer adjustment cost lower bound.

3.2 Demand

The estimation of costs, markups, and adjustment costs requires consistent estimates of the demand function as a first step. Market demand is derived from a standard discrete-choice model of consumer behavior. Given that the credibility of all our results will ultimately depend on the credibility of the demand system, it is imperative to adopt as general and flexible a framework as possible to model consumer behavior. We use the BLP random-coefficients model described in Hellerstein (2008), as this model was shown to fit the data well, while imposing very few restrictions on the substitution patterns. In the following we provide a brief overview of the model, and refer the reader to Nevo (2001) and Hellerstein (2008) for a more detailed discussion of the implementation.

Let the indirect utility u_{ijt} that consumer i derives from consuming product j at time t take the quasi-linear form:

$$(22) \quad u_{ijt} = x_{jt}\beta_i - \alpha_i p_{jt} + \xi_{jt} + \varepsilon_{ijt} = V_{ijt} + \varepsilon_{ijt}, \quad i = 1, \dots, I., \quad j = 1, \dots, J., \quad t = 1, \dots, T.$$

where ε_{ijt} is a mean-zero stochastic term. The utility from consuming a given product is a function of a vector of product characteristics (x, ξ, p) where p are product prices, x are product characteristics observed by the econometrician, the consumer, and the producer, and ξ are product characteristics observed by the producer and consumer but not by the econometrician. Let the

taste for certain product characteristics vary with individual consumer characteristics:

$$(23) \quad \begin{pmatrix} \alpha_i \\ \beta_i \end{pmatrix} = \begin{pmatrix} \alpha \\ \beta \end{pmatrix} + \Pi D_i + \Sigma v_i$$

where D_i is a vector of demographics for consumer i , Π is a matrix of coefficients that characterize how consumer tastes vary with demographics, v_i is a vector of unobserved characteristics for consumer i , and Σ is a matrix of coefficients that characterizes how consumer tastes vary with their unobserved characteristics. Conditional on demographics, the distribution of consumer unobserved characteristics is assumed to be multivariate normal. The demographic draws give an empirical distribution for the observed consumer characteristics D_i . Indirect utility can be expressed in terms of mean utility $\delta_{jt} = \beta x_{jt} - \alpha p_{jt} + \xi_{jt}$ and deviations (in vector notation) from that mean $\mu_{ijt} = [\Pi D_i \ \Sigma v_i] * [p_{jt} \ x_{jt}]$:

$$(24) \quad u_{ijt} = \delta_{jt} + \mu_{ijt} + \varepsilon_{ijt}$$

Finally, consumers have the option of purchasing an “outside” good; that is, consumer i can choose not to purchase any of the products in the sample. The price of the outside good is assumed to be set independently of the prices observed in the sample.¹⁴ The mean utility of the outside good is normalized to be zero and constant over markets. The indirect utility from choosing to consume the outside good is:

$$(25) \quad u_{i0t} = \xi_{0t} + \pi_0 D_i + \sigma_0 v_{i0} + \varepsilon_{i0t}$$

Let A_j be the set of consumer traits that induce purchase of good j . The market share of good j in market t is given by the probability that product j is chosen:

$$(26) \quad s_{jt} = \int_{\zeta \in A_j} P^*(d\zeta)$$

where $P^*(d\zeta)$ is the density of consumer characteristics $\zeta = [D \ \nu]$ in the population. To compute this integral, one must make assumptions about the distribution of the error term ε_{ijt} . Assuming that ε_{ijt} is i.i.d. with a Type I extreme-value distribution, the market share function becomes:

$$(27) \quad s_{jt} = \int_{\mu_{it}} \frac{e^{\delta_{jt} + \mu_{ijt}}}{1 + \sum_k e^{\delta_{kt} + \mu_{ikt}}} f(\mu_{it}) d\mu_{it}$$

¹⁴The existence of an “outside” good means that the focus on a single retailer (*Dominick’s*) does not imply that this retailer has monopoly power in the retail market; consumers faced with a price increase at *Dominick’s* have the option of switching to beer sold in other supermarkets, which represents the “outside” good in our framework.

The integral is approximated by the smooth simulator which, given a set of N draws from the density of consumer characteristics $P^*(d\zeta)$, can be written:

$$(28) \quad s_{jt} = \frac{1}{N} \sum_{i=1}^N \frac{e^{\delta_{jt} + \mu_{ijt}}}{1 + \sum_k e^{\delta_{kt} + \mu_{ikt}}}$$

Given these predicted market shares, we search for the demand parameters that implicitly minimize the distance between these predicted market shares and the observed market shares by using a generalized method-of-moments (GMM) procedure.

3.3 Discussion

The framework we just described is static. In this section we explore the implications of dynamic considerations for our analysis.

Dynamic considerations may affect the analysis in several ways. First, to the extent that consumers and/or retailers hold inventories of beer, the demand and supply side parameter estimates obtained by the static approach may be biased. Specifically, on the demand side, Hendel and Nevo (2006a, 2006b) have shown that when consumers stockpile in response to temporary price reductions, static demand estimates may overstate the long-run price elasticities of demand by a factor of 2 to 6.¹⁵ On the supply side, Aguirregabiria (1999) has analyzed the pricing behavior of a monopolistically-competitive retailer that holds inventories in a central store and delivers goods from this store to individual outlets. He shows that in the presence of fixed ordering costs and nominal price rigidities inventory dynamics have a significant effect on the retailer’s decision to change a brand’s price; ignoring such dynamics may hence lead to biased estimates of the importance of nominal price rigidities.¹⁶ Fortunately, these concerns that both build on the importance of inventories appear to be less relevant in our case. The industry wisdom is that consumers typically consume beer within a few hours after its purchase, so that consumer stockpiling is not a first-order concern.¹⁷ On the supply side, state and local regulations concerning the distribution of all alcohol, including beer, in Illinois stipulate that it is illegal for the central store of a retail chain to maintain inventories of beer and to deliver them to individual outlets.¹⁸ This must be done by firms exclusively licensed to be distributors. It is also illegal for beer to be transported from one outlet to another by the central store. So from the point of view of the central store

¹⁵Their numbers refer to laundry detergents, however, which consumers stockpile much more intensively than they do beer.

¹⁶The role of inventories is also emphasized in a recent paper by Alessandria, Kaboski and Midrigan (2008), which investigates exchange rate pass-through directly.

¹⁷See Anheuser-Busch, Inc. “Beer Shopper Poll,” 2005.

¹⁸For more on Illinois liquor regulations, see the homepage of the Illinois Liquor Control Commission at <http://www.state.il.us/LCC/>.

or the individual outlet, there is no inventory problem associated with beer, unlike most other products which are distributed by the central store. As the central store does not keep inventories of beer (indeed cannot by law), there is no relationship between inventory decisions and prices.¹⁹ And there is no incentive for individual outlets to maintain inventories over a period of longer than a week, as they can get a new shipment each week from the distributor, rather than bearing the costs of holding inventories themselves. An additional piece of evidence that stores do not hold beer inventories comes from the wholesale price series in Figures 1 and 2. In the Dominick's data, reported wholesale prices are based on average acquisition costs. If inventories were large and moved slowly, we would not expect to see the sharp "round-trip" drops we see in our wholesale price graphs.

A second way in which dynamics can enter the analysis is through intertemporal effects on the demand side, stemming for example from habit formation. The implications of such dynamics are explored in Slade (1998) and in a recent theoretical paper by Ravn, Schmitt-Grohe and Uribe (2007). While we believe that such dynamics are potentially important, it is not feasible to incorporate them into our current framework in a way that would allow us to estimate the model. To our knowledge, the only paper that has succeeded in estimating a model similar to ours but with dynamics on the demand side is the one by Slade (1998), which employs a linear demand specification. In comparison to Slade's paper, our approach puts more emphasis on the modelling of the demand side by allowing for a high degree of product differentiation and consumer heterogeneity and flexible substitution patterns. We believe that this is important in the context of the question we are trying to address, since one of the explanations for incomplete pass-through is markup adjustment by retailers and manufacturers, and there is a close link between markups and demand structure. However, the price we pay for the complexity of the demand system is that we have to abstract from demand side dynamics²⁰.

Finally, dynamic considerations may impact the analysis through the modeling of the price adjustment costs. As pointed out by Dixit (1991), under the assumption of rational expectations there is an option value in not adjusting prices, which magnifies the effects of even small costs of adjustment. The static approach we adopt, which is most similar to the models of Akerlof and Yellen (1985) and Mankiw (1985), fails to model this option value and hence may result in estimates of adjustment costs that are biased upwards. Nakamura (2007) in a recent paper adopts

¹⁹As Aguirregabiria argues, "There are some brands for which the central store does not keep inventories. Some of them are very perishable goods which are delivered daily from wholesalers to outlets (e.g. fresh vegetables, fish, some types of bread, etc.) In other cases, they are brands from manufacturers with efficient distribution networks that allow them to deliver their brands to individual outlets. From the point of view of the company's central store, there is not any inventory problem associated with those brands. Since we are interested in the relationship between price and inventory decisions we only consider those brands for which the central store keeps inventories" (Aguirregabiria 1999, p. 286).

²⁰We should note however that our static approach is standard in the Industrial Organization literature.

a dynamic framework to estimate menu costs, which are modeled as random i.i.d. variables. While the menu costs are derived within a dynamic framework, computational constraints force her to estimate the cost parameters based on static profit maximization conditions that do not take menu costs into account, so that there is a tension between the model used to obtain menu costs and conduct simulations, and the static model that is used to obtain the cost estimates. Our approach has the advantage that it enables us to maintain a theoretically consistent approach throughout the analysis. Still, we are concerned about the implications of Dixit's arguments for our analysis, so we turn to those next.

Suppose that instead of modeling the manufacturer's problem²¹ according to equation (13), we adopted a dynamic framework, in which the firm seeks to maximize the expected discounted value of future profits:

$$\sum_{t=0}^{\infty} E_t \delta^t \Pi_{jt}^w = \sum_{t=0}^{\infty} E_t \delta^t \{ p_{jt}^w - c_{jt}^w (tc_{jt}^w, ntc_{jt}^w) s_{jt} (p_t^r (p_t^w)) - A_{jt}^w \}$$

where δ denotes the discount rate and as before:

$$\begin{aligned} A_{jt}^w &= 0 \text{ if } p_{jt}^w = p_{jt-1}^w \\ A_{jt}^w &> 0 \text{ if } p_{jt}^w \neq p_{jt-1}^w \end{aligned}$$

Given that there are no dynamics on either the demand or cost side (the market shares depend only on current and not past prices), it follows that if the firm were faced with cost shocks that were considered permanent, the dynamic problem would be equivalent to the static one up to a discount factor: firms would be maximizing the sum over the (discounted) profits of each period, which in the absence of any cross-period price effects or inventories would be equivalent to maximizing period-by-period profits. In contrast, if shocks are perceived as temporary, firms will have an incentive to wait to adjust prices, so that the dynamic solution will be different from the static one.

The exchange rate shocks that move the traded component of costs (tc_{jt}^w) around, and are the focus of our analysis, are highly persistent; empirical work in International Economics has consistently failed to reject the unit root hypothesis for exchange rates. We therefore believe that the premise that such shocks are perceived as permanent is not unrealistic. Of course, exchange rates are not the only source of cost variation. Beer producers also face cost variation caused by changes in the prices of inputs such as barley, glass and energy. However, there is nothing that

²¹The argument for the retailer is similar.

comes close to the variation caused by exchange rate volatility! In Figure 3, we plot the exchange rate facing beer producers over our sample period along with the prices for glass and energy (all series are normalized to start at 1 at the beginning of our sample). While there is some variation in the prices of energy and glass, the exchange rate variation trumps everything else.

In sum, the static approach we adopt provides a reasonable approximation to a dynamic setup if the main source of cost variation is exchange rates, and if exchange rates are highly persistent. We believe that both premises are supported by the data.

Finally, we should point out that even though the static model is equivalent to a dynamic model with static expectations in terms of their implications for pass-through, the exact magnitudes of the price adjustment costs will differ depending on which of the two models we use to derive them. In particular, if we use the dynamic model (with static expectations) to derive adjustment cost bounds, inequality (18) becomes:

$$A_{jt}^w \leq \overline{A_{jt}^w} = \sum_{t=0}^{\infty} \delta^t \{ (p_{jt}^w - c_{jt}^w) s_{jt}(p_t^r(p_t^w)) - (p_{jt-1}^w - c_{jt}^w) s_{jt}^c(p_{jt}^{rc}(p_{jt-1}^w, p_{kt}^r)) \}, \quad k \neq j$$

which, assuming a constant discount rate, would reduce to:

$$A_{jt}^w \leq \overline{A_{jt}^w} = \frac{\delta}{1-\delta} \{ (p_{jt}^w - c_{jt}^w) s_{jt}(p_t^r(p_t^w)) - (p_{jt-1}^w - c_{jt}^w) s_{jt}^c(p_{jt}^{rc}(p_{jt-1}^w, p_{kt}^r)) \}, \quad k \neq j$$

That is the firm would compare the fixed cost of adjusting its price today, to the present discounted value of all future additional profits that it would make if it changed its price. A similar argument applies to the lower bound $\underline{A_{jt}^w}$. This implies that if we adopted a dynamic framework, in order to obtain the bounds for the price adjustment costs consistent with the dynamic model we would have to multiply the static estimates by the factor $\frac{\delta}{1-\delta}$, which is of the same order of magnitude as the inverse of the interest rate. This would “blow up” the static estimates for the price adjustment costs. Note however, that this would leave our simulations and decomposition results at the end of the paper *completely unchanged*: in the dynamic setup, *both* sides of the inequalities (10) and (18) that determine whether the firm will adjust its price in response to an exchange rate shock or not, would be multiplied by the same factor ($\frac{\delta}{1-\delta}$), so that this factor cancels out. This is why we emphasize that the term “price adjustment costs” is meaningful only within the context of the particular model we use as the basis of the empirical analysis, and that our results regarding pass-through are robust as long as this term is used consistently throughout the analysis.

4 Empirical Approach

Our empirical approach has two components: estimation and simulation. At the estimation stage, we estimate the demand parameters, the traded and non-traded costs and markups of the retailer and manufacturers, and the upper and lower bounds for the price adjustment costs. As noted above, these bounds are not by themselves informative regarding the role of price rigidities in explaining the incomplete cross-border cost shock transmission. To see why, suppose we estimate the adjustment cost of changing the price of a particular product j to be very small at the retail level. Still, as long as the adjustment cost is nonzero, it will cause the price of product j to remain unchanged in some periods. This in turn will affect the pricing of competing products: if the price of j does not change, then the prices of the products that *do* change may change by less than they would if *all* prices adjusted. Similarly at the wholesale level, the presence of a small adjustment cost at the retail level may cause the manufacturer to keep the wholesale level price constant if she anticipates that the retailer will not pass-through the change. Hence, a small adjustment cost may cause significant price inertia at both the retail and wholesale levels.

To assess the overall impact of adjustment costs on pricing behavior we employ simulation. In particular, we compute the industry equilibrium that would emerge if the dollar appreciated (depreciated) and prices were fully flexible, that is all adjustment costs were set to zero. Next we compare this equilibrium to the one that prevails in the presence of price rigidities. We interpret the differential response of prices across the two cases as a measure of the overall impact of nominal price rigidities. In the following we describe each step of our empirical approach in more detail.

4.1 Estimation

The estimation stage consists of the following steps:

1. Demand Estimation

The estimation of the demand system follows Hellerstein (2008). We model the mean utility associated with product j at time t as follows²²:

$$\delta_{jt} = \beta d_j - \alpha p_{jt} + \Delta \xi_{jt}$$

where the product fixed effects d_j proxy for both the observed characteristics x_{jt} in the term in equation (22) and the mean unobserved characteristics. The residual $\Delta \xi_{jt}$ captures deviations of the unobserved product characteristics from the mean (e.g., time-specific local

²²The demand model is also indexed by price zone z . In each period we have observations for two separate price zones. To keep the exposition simple, we omit the subscript z from our notation.

promotional activity) and is likely to be correlated with the price p_{jt} ; for example, an increase in the product’s promotional activity may simultaneously increase the mean evaluation of this product by consumers and a rise in its retail price. Addressing this simultaneity bias requires finding appropriate instruments, that is a set of variables z_{jt} that are correlated with the product price p_{jt} but are orthogonal to the error term $\Delta\xi_{jt}$. Factor prices and exchange rates satisfy this condition as they are unlikely to have any relationship to promotional activities while they are by virtue of the supply relation correlated with product prices. To construct our instruments we interact hourly wages in each country’s beverage industry with weekly bilateral exchange rates and indicator variables for each brand; this allows each product’s price to respond differently to a given supply shock.

2. Back out the non-traded retail costs ntc_{jt}^r and retail markups using data only for the periods in which retail prices adjust.

Once the parameters of the demand system have been estimated, we compute the market share function $s_{jt}(p_t^r)$ as well as the own and cross price derivatives $\frac{\partial s_{kt}}{\partial p_{jt}^r}$. Then we use the retailer’s first-order conditions for each product j (6) to estimate the non-traded retail costs of product j .

3. Model these non-traded costs parametrically as a function of observables (e.g., zone dummies, month dummies, local wages), and estimate the parameters of this function using data from the periods for which we observe retail price adjustment.

The procedure described under Step 2 allows us to back out the retailer’s non-traded costs for the periods for which we observe the price of a product adjusting, so that we can reasonably assume that the retailer sets the new price according to the first-order conditions. However, this approach does not work in periods in which the price does not change. To get estimates of the non-traded costs for these periods we employ the following procedure:

First, we collect the data on the non-traded costs ntc_{jt}^r in Step 2 for the periods in which the price of product j adjusted. Then we model these costs parametrically as a function of observables:

$$ntc_{jtz}^r = c_j + \gamma_z d_z + \gamma_w w_t^d + \eta_{jtz}$$

where c_j are brand fixed effects, d_z are price zone dummies, and w_t^d denote local wages. We run the above regression using data from the periods we observe price adjustment, and then use the parameter estimates to construct the predicted non-traded costs for the periods for which we do not observe price adjustment.

4. **Derivation of upper and lower bounds for the retailer price adjustment costs A_{jt}^r .**

With the demand parameter and non-traded cost estimates in hand, we employ (10) and (12) to derive the upper and lower bounds of the retailer adjustment costs A_{jt}^r . The computation of the upper bound is straightforward: in (10) all variables are observed, except for the counterfactual market share $s_{jt}^c(p_{jt-1}^r, p_{kt}^r)$ that product j would have if the retailer did not change her price from the previous period. This counterfactual share can however be easily evaluated once the demand parameters are estimated, given that the market share function is known.

The computation of the lower bound based on (12) requires the derivation of the counterfactual optimal price p_{jt}^{rc} that the retailer would charge if she changed the retail price from the previous period, and the associated market share $s_{jt}^c(p_{jt}^{rc}, p_{kt}^r)$. These are computed using (8) which reflects the first-order condition of the profit maximizing retailer.

5. **Back out the manufacturer marginal cost c_{jt}^w using data only for the periods in which wholesale prices adjust.**

The procedure here is similar as the one we employ to derive the non-traded costs for the retailer. In periods when the wholesale price changes, manufacturers behave according to their first-order conditions. Hence, we can use equation (16) to back out the manufacturer marginal cost c_{jt}^w .

6. **Model the manufacturer marginal cost parametrically as a function of observables (e.g., time dummies, local and foreign wages), and estimate the parameters of this function using data from the periods for which we observe wholesale price adjustment.**

The manufacturer first-order conditions we utilize under Step 5 allow us to back out the *total* marginal cost of the manufacturer; however they do not tell us how to decompose this cost into a traded and non-traded component. Furthermore, it is not possible to back out the marginal manufacturer costs for the periods when wholesale prices do not adjust based on this procedure, given that the first-order conditions do not necessarily hold then. To accomplish the above tasks, we model the total manufacturer costs parametrically as a function of observables, and estimate this function using data from the periods of wholesale price adjustment *only*. Specifically, we assume that the manufacturer marginal cost c_{jt}^w takes the form:

$$(29) \quad c_{jt}^w = \exp(\theta_j + \omega_{jt})(w_t^d)^{\theta_{dw}}(e_{jt}w_t^f)^{F_j*\theta_{fw}}(p_{bjt})^{D_j*\theta_{dp}}(e_{jt}p_{bjt})^{F_j*\theta_{fp}}$$

or, in log-terms:

$$(30) \ln c_{jt}^w = \theta_j + \theta_{dw} \ln w_t^d + F_j * \theta_{fw} \ln(e_{jt} w_t^f) + D_j * \theta_{dp} \ln(p_{bdt}) + F_j * \theta_{fp} \ln(e_{jt} p_{bjt}) + \omega_{jt}$$

where w_t^d and w_t^f denote local domestic and foreign wages respectively, e_{jt} is the bilateral exchange rate between the producer country and the U.S., p_{bjt} is the price of barley in the country of production of brand j , F_j is a dummy that is equal to 1 if the product is produced by a foreign supplier, and zero otherwise, and D_j is a dummy that is equal to 1 if the product is produced by a domestic supplier, and zero otherwise. For the function to be homogeneous of degree 1 in factor prices, we require $\theta_{dw} + F_j * \theta_{fw} + D_j * \theta_{dp} + F_j * \theta_{fp} = 1$. Equation (30) can be easily estimated by Least Squares.

The estimation of the above equation for the manufacturer marginal cost serves two purposes. First, it allows us to decompose the total marginal cost into a traded and a non-traded component. Recall that by definition the traded component refers to the part of the marginal cost that is paid in foreign currency and hence is subject to exchange-rate fluctuations. For domestic producers the traded component will be (by definition) zero. Foreign producers selling in the U.S. will generally have both traded and local non-traded costs. The latter are captured in the above specification by the term $(w_t^d)^{\theta_{dw}}$ that indicates the dependence of foreign producers' marginal costs on the local wages in the U.S.. The specification in (29) can be used to demonstrate two important facts regarding foreign suppliers' costs. First, foreign producers selling to the U.S. will typically experience substantially more volatility than domestic producers due to their exposure to exchange-rate shocks. Second, as long as the local non-traded cost component is nonzero (so that $\theta_{fw} + \theta_{fp} < 1$ ²³), the dollar denominated marginal cost of foreign producers will change by a smaller proportion than the exchange rate. This incomplete marginal cost response may partially explain the incomplete response of exchange-rate changes on to prices.

Estimation of the marginal cost equation (30) furthermore allows us to use the parameter estimates to construct predicted values for the manufacturer traded and non-traded costs for the periods in which wholesale price adjustment is not observed.

7. Derivation of upper and lower bounds for the wholesale price adjustment costs

A_{jt}^w .

The final step is to use all parameter estimates obtained in the previous steps to compute the

²³Given the assumption of $\theta_{jdw} + F_j * \theta_{jfw} = 1$ which guarantees homogeneity of degree 1 of the marginal-cost function in factor prices, if local nontraded costs are zero, then $\theta_{jdw} = 0$ and $\theta_{jfw} = 1$. In contrast, with positive nontraded costs we will have $\theta_{jfw} < 1$.

upper and lower bounds of the manufacturer price adjustment costs based on (18) and (20). Consider inequality (18) first that determines the adjustment cost upper bound. Once steps 1-6 are completed, all variables in this inequality are known, except for the counterfactual retail price p_{jt}^{rc} that the retailer would charge if the manufacturer did not change her price in that period. The counterfactual price p_{jt}^{rc} can take on one of two values: it is either equal to p_{jt-1}^r , or it is determined according to the retailer's first-order condition, *conditional on the retailer observing the wholesale price p_{jt-1}^w* . To determine which of the two prices the retailer will choose, we first solve for the optimal price that the retailer would pick if she behaved according to her profit maximization condition. Then we compare the retail profits evaluated at this retail price, to the profits that the retailer would make if she kept the retail price unchanged at p_{jt-1}^r . The retailer will choose the price associated with the higher retail profits. Once the counterfactual retail price p_{jt}^{rc} has been determined this way, the associated counterfactual market share $s_{jt}^c(p_{jt}^{rc}(p_{jt-1}^w, p_{kt}^r)), k \neq j$, can easily be evaluated.

Next consider inequality (20) that determines the adjustment cost lower bound. Again, all variables in this inequality can be treated as known once steps 1-6 are completed, except for the counterfactual retail and wholesale prices, p_{jt}^{rc} and p_{jt}^{wc} respectively, which we would observe if the manufacturer changed her price from the previous period. To determine those, we consider two cases. In the first case the retail price changes from the previous period; the optimal prices p_{jt}^{wc} and p_{jt}^{rc} are then determined according to the manufacturer and retailer first-order conditions, equations (16) and (6) respectively, with the inverted manufacturer reaction matrix Ω_{wt}^{-1} reflecting the optimal pass-through of the wholesale price change onto the retail price. Let π_1^{wc} denote the manufacturer profits associated with the so-computed prices p_{jt}^{wc} and p_{jt}^{rc} .

Next, consider the case in which the retail price does not change, even though the wholesale price does. As noted earlier, the optimal manufacturer pricing behavior in this case will involve changing the wholesale price up to the point where the retailer is just indifferent between changing her price and keeping it constant at p_{jt-1}^r . The optimal wholesale price will then be determined based on equation (21) along the lines discussed in the previous section. Let π_2^{wc} denote the manufacturer profits associated with the prices p_{jt}^{wc} and p_{jt-1}^r in this case.

If $\pi_1^{wc} > \pi_2^{wc}$, the manufacturer will set the wholesale price anticipating that the retailer will adjust her price too. Hence, the counterfactual prices p_{jt}^{wc} and p_{jt}^{rc} will satisfy the conditions described under the first case above. If $\pi_1^{wc} < \pi_2^{wc}$, the manufacturer will price the product anticipating that the retailer will not adjust her price. The resulting counterfactual wholesale price will then satisfy the indifference condition discussed under the second case, while the

retail price will remain unchanged at p_{jt-1}^r . Once the counterfactual wholesale and retail prices have been determined, evaluation of the adjustment cost lower bound based on (20) is straightforward.

4.2 Simulations and Decomposition of Incomplete Exchange-Rate Pass-Through

To assess the overall impact of adjustment costs on pricing behavior we employ simulations. First, we compute the industry equilibrium that would emerge if a particular firm faced an exchange rate shock and prices were fully flexible, that is, all adjustment costs were equal to zero. In a second set of simulations, we derive the industry equilibrium under the presence of nominal rigidities. We interpret the differential response of prices across the different cases as a measure of the impact of nominal price rigidities.

To be more specific, to identify the channels through which nominal rigidities affect prices, we conduct three separate simulations that we briefly describe below.

Simulation 1: Simulate the effect of a 1% exchange-rate change when there are no price rigidities. To recover manufacturer and retailer pass-through coefficients we derive the effect of a shock to foreign firms' marginal costs (e.g., an exchange rate shock) on all firms' wholesale and retail prices by computing a new Bertrand-Nash equilibrium. Suppose that an exchange rate shock hits the traded component of the j th product's marginal cost (that is the component that is denominated in foreign currency). To compute the transmission of this shock to wholesale prices, we substitute the new vector of traded marginal costs, tc_t^{w*} , into the system of J nonlinear equations that characterize manufacturer pricing behavior, and then search for the wholesale price vector p_t^{w*} that will solve the system:

$$(31) \quad p_{jt}^{w*} = c_{jt}^w(tc_{jt}^w, ntc_{jt}^w) - \sum_{k \in \Gamma_{mt}} (S_{wt} * T_w)^{-1} s_{kt} \text{ for } j = 1, 2, \dots, J_t.$$

To compute pass-through coefficients at the retail level, we substitute the derived values of the vector p_t^{w*} into the system of J nonlinear equations for the retail firms, and then search for the retail price vector p_t^{r*} that will solve it:

$$(32) \quad p_{jt}^{r*} = p_{jt}^{w*} + ntc_{jt}^r - \sum_{k \in \kappa^r} (S_{wt} * T_w)^{-1} s_{kt} \text{ for } j, k = 1, 2, \dots, J_t.$$

The transmission of the original marginal-cost shock to the retail price, is given by $\left(\frac{dp^r}{dp^w}\right)' \frac{dp^w}{dtc_f^w}$.

This first simulation gives us the benchmark pass-through elasticities that would emerge in the absence of nominal rigidities. Using our estimation results on the traded and non-traded

components of manufacturer and retailer costs respectively, we can further decompose these pass-through elasticities into the part that is due to the presence of traded costs and the part that reflects markup adjustment. Given the Cobb-Douglas specification for the manufacturer marginal cost described earlier, the contribution of the traded costs will be given by the sum of the coefficients $\theta_{fw} + \theta_{fp}$ (or, in other words, the contribution of local costs to generating incomplete pass-through will be captured by the coefficient on domestic wages θ_{dw}). The difference between the derived wholesale pass-through and the one attributed to traded costs reflects markup adjustment on the part of the manufacturer. Similarly at the retail level, we can use our estimates of the retailer non-traded costs to compute the effect that such costs have in generating incomplete pass-through of wholesale to retail prices; this effect will be given by $(d \ln(p_j^w + ntc_j) / d \ln p_j^w)$. The difference between the derived retailer pass-through and the one attributed to retailer non-traded costs captures the markup adjustment on the part of the retailer.

Next we consider the case where nominal price rigidities are present. Because firms in our framework are not symmetric, and price changes will not be synchronized, characterizing the equilibrium in this case becomes extremely involved. To keep the problem tractable and get a sense of how price rigidities affect prices, we confine our discussion to two extreme cases; one in which the firm facing the exchange rate shock assumes that all other competitors will adjust their prices, and one in which the firm under consideration assumes that competitor prices will remain fixed as a result of nominal rigidities. These cases correspond to the following two simulations:

Simulation 2: Simulate the effect of a 1% exchange-rate change assuming that the foreign-brand producer facing the exchange rate shock also faces price adjustment costs but that all other prices adjust freely. In this case, the new industry equilibrium is computed taking into account the fixed adjustment costs for the foreign brand affected by the exchange-rate shock included in equations (31) and (32). As discussed earlier, our approach does not allow us to pin down these adjustment costs, but only derive upper and lower bounds for these costs. In all simulations involving price adjustment costs we employ the estimated upper bounds. As we show below, our lower bound estimates are not significantly different from zero.

By comparing the results from Simulation 2 to the results from Simulation 1 we can get a sense of the additional effect that price adjustment costs have on pass-through behavior. However, given that we let competitor prices adjust freely, this simulation captures only the *direct* effect of repricing costs, in other words, the effect that a firm's *own* adjustment costs have on its pricing behavior. It ignores the effect that *other firms'* repricing costs have on a firm's decision whether and by how much to change its price in response to an exchange rate shock. To address this *indirect* effect, we conduct the third simulation.

Simulation 3: Simulate the effect of a 1% exchange-rate change assuming that the foreign-brand producer facing the exchange rate shock faces price adjustment costs and also assumes that competitor prices will remain fixed. In this case, the new industry equilibrium is computed as in Simulation 2, but by additionally imposing that all other product prices remain unchanged. This simulation captures the *indirect* or *strategic* aspect of repricing costs; even if nominal rigidities do not prevent a particular firm from adjusting its price, this adjustment may be smaller if the firm assumes that nominal rigidities will keep competitor prices fixed compared to the case without any rigidities. A comparison between the results from Simulation 3 and Simulation 2 allows us to capture precisely this effect.

5 Results

This section first discusses results from the estimation of the demand system. It then describes estimates of brand-level markups, non-traded costs, and upper and lower bounds for the retailer and the manufacturer adjustment costs. Finally, it reports the results from the simulations that allow us to decompose the incomplete transmission of exchange rate shocks into its sources in order to quantify the relative contribution of local costs, markup adjustment, and repricing costs in generating the documented local currency price stability.

5.1 Demand Estimation: Logit Model

Table 4 reports results from estimation of demand using the multinomial logit model. Due to its restrictive functional form, this model will not produce credible estimates of pass-through. However, it is helpful to see how well the instruments for price perform in the logit demand estimation before turning to the full random-coefficients model.²⁴

Table 4 suggests that the instruments have power. The first-stage F-test of the instruments, at 34.45, is significant at the one-percent level. The consumer's sensitivity to price should increase after we instrument for unobserved changes in characteristics. That is, consumers should appear more sensitive to price once we instrument for the impact of unobserved (by the econometrician, not by firms or consumers) changes in product characteristics on their consumption choices. It is promising that the price coefficient falls from -0.93 in the *OLS* estimation to -2.43 in the IV estimation. Note that the 95-percent confidence interval of the latter coefficient does not include the value of the former. Results reported in columns 2 and 4 of the table show that including holiday dummies does not affect the demand coefficients in either the OLS or the IV estimation.

²⁴An appendix (available online at <http://www.ny.frb.org/research/economists/hellerstein/papers.html>) reports the first-stage results for demand. Most of the coefficients have the expected sign: as hourly compensation increases, the retail price of each product should increase. T-statistics calculated using Huber-White robust standard errors indicate that most of the coefficients are significant at the 5-percent level.

5.2 Demand: Random-Coefficients Model

Table 5 reports results from estimation of the demand system. We allow consumers' income to interact with their taste coefficients for price and percent alcohol. As we estimate the demand system using product fixed effects, we recover the mean consumer-taste coefficients in a generalized-least-squares regression of the estimated product fixed effects on product characteristics (maltiness, bitterness, hoppiness and percent alcohol).

The coefficients on the characteristics generally appear reasonable. As consumers' income rises, they become less price sensitive. The random coefficient on income, at 0.85, is significant at the five-percent level. The mean preference in the population is in favor of a bitter taste in beer, which has a positive and significant coefficient. The mean coefficient on a malty or a hoppy flavor is negative. As the percent alcohol rises across brands, the mean utility in the population also rises, an intuitive result. There is heterogeneity in the population with respect to this characteristic: Those with higher incomes get less utility from a high percent of alcohol in their beer, that is, prefer light beers. This is consistent with industry lore: Higher income individuals tend to prefer light and imported beers.²⁵

5.3 Retail Markups and Non-Traded Costs

Table 6 reports retail and wholesale prices and markups for selected imported brands. The markups are derived using firms' first-order conditions. Under the assumption of no adjustment costs, the markups would be derived using the first-order conditions of every product in every period. Under the alternative assumption of some adjustment costs, the markups are derived in each period by using the first-order conditions of only those products whose prices adjust from the previous period. As discussed earlier, many of the price changes in our data reflect promotions, during which the price of a particular brand is reduced for a few weeks (see also Figures 1 and 2). A striking characteristic of these promotions is that product prices return to their exact pre-promotion level once the promotion is over. In theory, the transition from the discount price to the pre-promotion level is a price change that could be handled in the same manner as a level change in price (after all, firms do incur some cost every time they change the posted price); yet, given that firms seem to charge *exactly* the same price that they were charging before the promotion, we were skeptical about the plausibility of the assumption that the post-promotion prices are determined based on firms' first order conditions. To be safe, we conducted the empirical analysis both ways, first applying the FOC's to all periods in which the price changed (including

²⁵Our demand side estimates are similar to those obtained in Hellerstein (2008), who estimated a random-coefficients demand model with the same data, but after aggregating the data up to the monthly frequency. This makes us confident that our price elasticities of demand are not affected by the type of arguments made in Hendel and Nevo (2006b).

changes associated with promotions), and then excluding those time periods during which firms charged the same price as before the promotion. The results did not differ in any significant manner across the two approaches, but the second approach significantly reduces the number of observations associated with a price change than we can exploit in the empirical analysis. Still, in the remainder of the paper we report results based on this second, more conservative approach, as we are more comfortable with the assumption that FOC's hold only in those periods during which a firm charges a price that is genuinely different from the price charged in earlier periods. In general, the markups that are derived based on this approach appear reasonable and consistent with industry wisdom as evident from Table 6. The same applies to the retailer's non-traded costs that were derived based on the regression reported in Table 8.

Using a similar procedure we derived the marginal costs facing each beer manufacturer, and then employed the regression described in equation (30) in Section 4.1 to obtain an estimate of the "local content" of foreign manufacturers' marginal cost. This local content is reflected in the magnitude of the "domestic U.S. wages" coefficient that captures the cost share accounted for by domestic labor. As discussed earlier (see p. 8, in particular footnote 6), because distributors pricing is coordinated by brewers, we treat the manufacturer and distributor as one entity, so that the "local" manufacturer costs include the marketing and distribution costs incurred by the distributor. With a highly significant coefficient of 0.52 the share of local costs appears to be substantial. It implies that a big part of foreign manufacturers' costs of selling in the U.S. market are not affected by exchange rate fluctuations. Hence it comes as no surprise that foreign producers do not fully adjust their U.S. dollar prices in response to exchange rate changes. This finding implies that even without menu costs, the existence of local, non-traded costs can generate a significant degree of inertia in local currency prices. However, whether local costs can by themselves fully explain the observed patterns in the price data remains an open question. We therefore turn to an examination of the price adjustment costs next.

5.4 Adjustment Costs

Table 11 reports the mean estimates by brand of the upper and lower bounds on the retailer's and manufacturers' adjustment costs. The entries in the first and third columns report the mean of each brand's price-adjustment cost as a share of its total revenue from that brand in that week. The bounds generally are consistent for each brand as well across the brands. The lower bound is indistinguishable from zero across brands for both retail and wholesale prices. The upper bounds on adjustment costs to retail prices range from 0.1 percent of revenue for *Heineken* to 0.9 percent of revenue for *St. Pauli Girl*, with a mean upper bound across foreign brands of 0.4 percent. Manufacturer adjustment costs are generally larger as a share of revenue than retail adjustment

costs: Their upper bounds range from 0.0 percent of revenue for *Heineken* to 4.3 percent of revenue for *Beck's*, with a mean of 3 percent of revenue across all foreign brands. The second and fourth columns of Table 11 report the sum of the upper bounds for each brand's menu costs divided by the total retail (or manufacturer revenue) for that brand, both computed over the full sample period. These numbers are more comparable to those of the Levy et al (1997) and Dutta et al (1999) studies, which follow a similar approach of dividing costs of repricing calculated for only those periods when prices change divided by the revenue earned by the firm across all periods, whether prices change or not²⁶. The sum of repricing costs across all foreign brands is 0.1 percent of total revenue for the retailer and 0.5 percent for the manufacturers. As pointed out earlier, the particular estimates of price adjustment costs are not of interest here, as these numbers are meaningful only in the particular context of our model²⁷, and are accordingly meant to be used only in the context of our model in order to perform simulations. Nevertheless, we find it interesting that the order of magnitude of our estimates is similar to the one obtained in other studies that used a completely different methodology.

Finally, Table 12 reports the results of a fixed-effects panel regression of the derived retail menu costs as a share of revenue on a dummy for a level change in a brand's price (as opposed to a sale, that is, a temporary reduction). Adjustment costs appear to be significantly higher for level changes in prices. This finding provides empirical support for a recent paper by Kehoe and Midrigan (2007) that is based on the premise that the fixed cost of changing a regular price is larger than the fixed cost of a temporary reduction.

5.5 Simulations

Using the full random-coefficients model and the derived measures of traded, nontraded, and repricing costs, we conduct the counterfactual experiments described earlier to analyze how firms react to exchange-rate shocks. This subsection presents and discusses the results from these experiments. We consider the effect of a one-percent foreign currency appreciation on foreign brands' prices in three scenarios, each with a different assumption about the nature of the repricing costs faced by foreign brands. We use the results from all three scenarios to assess the relative importance of foreign firms' local non-traded costs, markup adjustment, and repricing costs in their incomplete pass-through of exchange-rate fluctuations.

²⁶This comparison is subject to the same caveats mentioned in the Introduction, p. 6, footnote 2.

²⁷For example, if we interpreted these numbers in the context of the dynamic model outlined at the end of Section 3.3, then we would have to divide the current estimates by the interest rate to obtain the price adjustment costs corresponding to the dynamic setup. If we used an interest rate of 10%, the resulting price adjustment cost estimates would be around 1% and 5% of revenue for retailers and manufacturers respectively. Levy et al and Dutta et al's estimates for *retailer* costs, which were based on the sum of the labor and material costs associated with a price change, were 0.70% and 0.59% of the revenue respectively.

The counterfactual experiments consider the effect of a one-percent appreciation of the relevant foreign currency on the prices of a British, German, Mexican, and Dutch brand (*Bass*, *Beck's*, *Corona*, and *Heineken*, respectively) in twelve exercises reported in Table 13. There are three panels in the table, each one corresponding to one of the simulations we described above. The first column of the table reports for each simulation the manufacturer pass-through elasticity of the original shock that is due to local dollar-denominated costs incurred by the manufacturer. The second column reports the pass-through of the original shock to the wholesale price that is attributable to manufacturer markup adjustment. The third column reports the (incomplete) pass-through of the original shock to the retail price due to the presence of a local component in retail costs. The last column reports the incomplete pass-through of the original shock to the retail price due to the retailer's markup adjustment.

Simulation 1: Simulate the effect of a 1% appreciation of the relevant foreign currency when there are no repricing costs. The first counterfactual experiment examines the manufacturers' and the retailer's pass-through following a 1-percent appreciation of the relevant foreign currency when they face no repricing costs. Its results are reported in the top panel of Table 13. The median pass-through of the exchange rate shock on manufacturer's total marginal cost is 50 percent, which is determined by averaging down the coefficient on local wages (0.52) from the regressions reported in Table 10: The average local non-traded cost incurred by a foreign manufacturer is approximately 50 percent of her total costs. Thus, a nontrivial amount of non-traded value is added at this stage of the distribution chain. Next, manufacturer markup adjustments are large and vary quite a bit across brands in this counterfactual: With markup adjustments taken into account, the median pass-through elasticity of the exchange rate shock to the wholesale price ranges from 13.1 percent for *Bass* to 33.4 percent for *Heineken*. It is 18.3 percent across all brands. With retailer local costs taken into account, the median pass-through becomes 15.4 percent and ranges from 11.2 percent for *Bass* to 27.1 percent for *Heineken*. Finally, the retailer appears to adjust her markup significantly for some brands and only marginally for others: The median retailer pass-through elasticity across all brands is 14.3 percent and ranges from 11.3 percent for *Bass* to 26.7 percent for *Heineken*.

Simulation 2: Simulate the effect of a 1% appreciation of the relevant foreign currency assuming that only the foreign brand affected by the exchange-rate shock faces fixed repricing costs, and that all other brands' prices adjust freely. The second counterfactual experiment considers how manufacturers and the retailer adjust their prices following a 1-percent appreciation of the relevant foreign currency if they must incur fixed repricing costs to alter their prices. Its results are reported in the middle panel of Table 13. The median

pass-through of the exchange rate change on manufacturer total marginal cost is again 50 percent as the share of non-traded costs is unaffected by the nature of the counterfactual. But final manufacturer pass-through elasticities now vary significantly across brands depending on whether repricing costs are large enough at this level to prevent manufacturers from changing wholesale prices. Their median ranges from 0 percent for *Bass* and *Becks*, both brands with significant own-brand repricing costs, to around 30 percent for both *Corona* and *Heineken*, the two imports with the highest market share. The retailer's median pass-through elasticities are naturally 0 percent for the two brands that do not change their manufacturer price. For *Corona* and *Heineken*, the pass-through elasticities are around 26 and 27 percent respectively, once the retailer local costs are taken into account, and slightly lower once retailer markup adjustment is accounted for. The retail traded pass-through elasticity is 6 percent for all four brands. Thus taking into account a brand's own price adjustment costs reduced the average pass-through elasticity from 14.3% in Simulation 1 to 6% in Simulation 2. Note however that this reduction is due to the zero transmission of the exchange rate shock on the *wholesale* prices of *Bass* and *Becks* due to these two brands' *manufacturer* repricing costs. In contrast, retail repricing costs do not contribute to any further reductions in the pass-through elasticities. This is consistent with the patterns we documented earlier that suggest that retail prices always adjust whenever wholesale prices adjust.

Simulation 3: Simulate the effect of a 1% appreciation of the relevant foreign currency assuming that the foreign firm affected by the exchange-rate shock faces fixed repricing costs and also assumes that competitor prices will remain fixed. The third counterfactual experiment considers whether manufacturers and the retailer adjust their prices following a 1-percent appreciation of the relevant foreign currency if they take their competitor prices to be fixed. As before, the own manufacturer repricing costs of *Bass* and *Becks* prevent the wholesale prices of these two brands from adjusting, which translates to zero pass-through elasticities. But what is interesting in this counterfactual is the fact that the pass-through elasticities in those cases where prices do adjust (*Corona* and *Heineken*) are now lower compared to Simulation 2. This additional reduction in the pass-through elasticities captures the indirect or strategic effect of repricing costs: because each brand assumes that repricing costs will prevent its competitors from changing their prices, the brand's own response to the exchange rate shock is less pronounced than it would be with flexible prices. Overall, this effect accounts for the reduction of the manufacturer pass-through elasticity from 8.3% in Simulation 2, to 6.8% in Simulation 3, and the reduction of the retail pass-through elasticity from 6% in Simulation 2, to 5% in Simulation 3.

5.6 Decomposition of the Incomplete Transmission

Table 14 decomposes the sources of the incomplete transmission of the exchange-rate shock to retail prices that is documented in Table 13. The first column of the table reports the share of the incomplete transmission that can be attributed to a local dollar-denominated cost component in manufacturers' marginal costs. The second column reports the share that can be attributed to markup adjustment by manufacturers following the shock (separate from any costs of repricing faced). Columns three and seven report the shares of the incomplete transmission attributable to the fixed costs of price adjustment incurred by the manufacturer and retailer, respectively, when they change their own prices (the direct effect of repricing costs). Columns four and eight report the shares of the incomplete transmission attributable to the effect that the fixed costs of repricing faced by competitors have on the manufacturer and retailer's pricing behaviors (the indirect or strategic effect). The fifth column reports the share attributable to a local-cost component in the retailer's marginal costs, and the sixth column the share attributable to the retailer's markup adjustment, separate from any markup adjustment associated with repricing costs.

Manufacturers' local non-traded costs play the most significant role in the incomplete transmission of the original shock to retail prices. Following a 1-percent appreciation of the relevant foreign currency, it is responsible for roughly half, or 52.5 percent, of the observed retail-price inertia. Manufacturers' markup adjustment accounts for 33.5 percent of the remaining adjustment, and their own repricing costs for another 10.5 percent. At the retail level, we attribute roughly 1.6 percent of the incomplete pass-through to local non-traded costs. In contrast, the retailer's markup adjustment and own repricing costs have a negligible role in explaining the incomplete transmission. Finally, the competitive effects of rival brands' repricing costs account for 1.6 and 0.1 percent of the incomplete transmission by manufacturers and retailers, respectively. These results support the initial intuition conveyed by Figures 1 and 2 that the effects of fixed repricing costs are most evident in the infrequent adjustment of wholesale prices, while such costs play only a minor role in explaining the inertia of retail prices.

Overall, local-cost components account for 54.1 percent of the observed price inertia following a currency appreciation, firms' markup adjustments account for 33.7 percent, manufacturer repricing costs for 12.1 percent, and retailer repricing costs for 0.1 percent.

6 Conclusions

This paper set out to build a framework that can be used to identify the determinants of local currency price stability in the face of exchange rate fluctuations. The empirical model we developed incorporates the three main potential sources identified in the literature: local non-traded costs; markup adjustment; and fixed costs of repricing. Our analysis yields several interesting

findings.

First, at the descriptive level, we document that while both wholesale and retail prices do not change every period, retail prices always respond to changes in wholesale prices. Hence, it appears that infrequent price adjustment is primarily driven by the behavior of wholesale prices. Second, when we use our model to derive upper and lower bounds for the fixed costs of price adjustment facing retailers and manufacturers, we find that these costs are substantially higher for manufacturers than for retailers, both in absolute terms and as a share of revenue. Third, the counterfactual simulations we conduct in order to decompose the incomplete transmission of exchange rate shocks into its sources suggest that both local non-traded costs and markup adjustment are important in generating local currency price stability. However, these two factors alone cannot completely explain the incomplete pass-through of exchange rate changes to consumer prices. To generate the pass-through observed in the data we need to allow for price adjustment costs. To be more specific, local costs and markup adjustment put together generate an exchange rate pass-through of approximately 18%. The pass-through observed in our data is between 6 and 7%. The missing 11-12% represent the overall impact of costs of repricing. These rigidities affect primarily the adjustment of wholesale prices; their effect on retail prices is very minor, and operates only indirectly through the strategic interaction of competing brands. Why nominal price rigidities are present primarily at the wholesale but not retail level is - in our opinion - an intriguing question worth further exploration. One possible explanation is that wholesale prices are set through long-term contracts and are therefore less responsive to changes in economic conditions. We hope that future research can shed more light into this issue.

References

- [1] "Can you Judge a Beer by Its Label? Beer Ratings," 1996. *Consumer Reports*, June, pp. 10-19.
- [2] "The Alcoholic Beverage Market in the United States," 1998. Canadian Trade Commissioner, June.
- [3] Aguirregabiria, V. 1999. "The Dynamics of Markups and Inventories in Retailing Firms," *Review of Economic Studies*, vol. 66, no. 2, pp. 275-308.
- [4] Akerlof, G. and J. Yellen. 1985. "A Near-Rational Model of the Business Cycle with Wage and Price Inertia", *Quarterly Journal of Economics*, vol. 100, pp. 823-38.
- [5] Alessandria, G., J. Kaboski and V. Midrigan. 2008. "Lumpy Trade and Large Devaluations" NBER Working Paper 13790.
- [6] Asker, J. 2004. "Measuring Advantages from Exclusive Dealing," Working Paper, NYU Stern.
- [7] Berry, S., J. Levinsohn and A. Pakes. 1995. "Automobile Prices in Market Equilibrium," *Econometrica*, vol. 63, pp. 841-90.
- [8] Blinder, A., E. Canetti, D. Lebow, and J. Rudd. 1998. *Asking About Prices: A New Approach to Understanding Price Stickiness*. Russell Sage Foundation.
- [9] Burstein, A.T., J.C. Neves and S. Rebelo. 2003. "Distribution Costs and Real-Exchange-Rate Dynamics during Exchange-Rate-Based Stabilization," *Journal of Monetary Economics*, vol. 50, pp. 1189-1214.
- [10] Carlton, D. 1986. "The Rigidity of Prices," *The American Economic Review*, vol. 76 (September), pp. 637-58.
- [11] Corsetti, G., Dedola, L. and S. Leduc. 2006. "DSGE models with High Exchange-Rate Volatility and Low Pass-through," International Finance Discussion Papers of the Federal Reserve Board, No. 845; CEPR dp 5377.
- [12] Dixit, A. 1991. "Analytical Approximations in Models of Hysteresis," *Review of Economic Studies*, vol. 58, pp. 141-51.
- [13] Dutta, S., M. Bergen, D. Levy, and R. Venable. 1999. "Menu Costs, Posted Prices, and Multiproduct Retailers," *Journal of Money, Credit, and Banking*, vol. 31, no. 4 (November), pp. 683-703.

- [14] Engel, C. 2002. "Expenditure Switching and Exchange-Rate Policy," *NBER Macroeconomics Annual*, pp. 231-72.
- [15] Engel, C. 2003. "On the Relationship between Pass-Through and Sticky Nominal Prices." Working Paper.
- [16] Goldberg, L. and J. Campa. 2006. "Distribution Margins, Imported Inputs, and the Sensitivity of the CPI to Exchange Rates," NBER Working Paper No. 12121.
- [17] Goldberg, P.K. 1995. "Product Differentiation and Oligopoly in International Markets: The Case of the U.S. Automobile Industry," *Econometrica*, vol. 63, no. 4 (July), pp. 891-951.
- [18] Goldberg, P. and M. Knetter. 1997. "Goods Prices and Exchange Rates: What Have We Learned?" *Journal of Economic Literature*, vol. 35 (September), pp. 1243-73.
- [19] Goldberg, P.K. and F. Verboven. 2001. "The Evolution of Price Dispersion in the European Car Market," *Review of Economic Studies*, vol. 67, pp. 811-48.
- [20] Goldberg, P.K. and F. Verboven. 2005. "Market Integration and Convergence to the Law of One Price: Evidence from the European Car Market," *Journal of International Economics*, vol. 65, pp. 49-73.
- [21] Hellerstein, R. 2008. "Who Bears the Cost of a Change in the Exchange Rate? Pass-through Accounting for the Case of Beer," *Journal of International Economics*, forthcoming.
- [22] Hendel, I. and A. Nevo. 2006a. "Sales and Consumer Inventory," *Rand Journal of Economics*, vol. 37, no. 3, pp. 543-561.
- [23] Hendel, I. and A. Nevo. 2006b. "Measuring the Implications of Sales and Consumer Inventory Behavior," *Econometrica*, vol. 74, no. 6, pp. 1637-73.
- [24] Kehoe, P. and V. Midrigan. 2007. "Sales and the Real Effects of Monetary Policy," Federal Reserve Bank of Minneapolis Working Paper 652.
- [25] Krugman, P. 1987. "Pricing to Market when the Exchange Rate Changes," in J.D. Richardson and S. Arndt (eds.), *Real-Financial Linkages among Open Economies*, Cambridge, MA: MIT Press, 1987.
- [26] Levy, D, Bergen, M., S. Dutta, and R. Venable. 1997. "The Magnitude of Menu Costs: Direct Evidence from Large U.S. Supermarket Chains," *The Quarterly Journal of Economics*, vol. 112, pp. 791-825.

- [27] Mankiw, N.G. 1985. "Small Menu Costs and Large Business Cycles," *The Quarterly Journal of Economics*, vol. 100, pp. 529-39.
- [28] Midrigan, V. 2006. "Menu Costs, Multi-Product Firms, and Aggregate Fluctuations", mimeo, New York University.
- [29] Nakamura, E. 2007. "Accounting for Incomplete Pass-Through", mimeo.
- [30] Nevo, A. 2001. "Measuring Market Power in the Ready-To-Eat Cereal Industry", *Econometrica*, vol. 69, no. 2, pp. 307-42.
- [31] Pesendorfer, M. 2002. "Retail Sales. A Study of Pricing Behavior in Supermarkets", *Journal of Business*, 75(1), 33-66.
- [32] Ravn, M., S. Schmitt-Grohe and M. Uribe. 2007. "Incomplete Cost Pass-through Under Deep Habits," NBER Working Paper 12961.
- [33] Slade, M. 1998. "Optimal Pricing with Costly Adjustment and Persistent Effects: Empirical Evidence from Retail-Grocery Prices", *Review of Economic Studies*, Vol. 65, No. 1, 87-107.
- [34] Sobel, J. 1984. "The Timing of Sales", *Review of Economic Studies*, 51, 353–368.
- [35] Varian, H. R. 1980. "A Model of Sales", *American Economic Review*, 70, 651–659.
- [36] Villas-Boas, S.B. 2007. "Vertical Relationships Between Manufacturers and Retailers: Inference with Limited Data," *Review of Economic Studies*, vol. 74, no. 2, pp. 625-52.




Figure 1: *Weekly retail and wholesale prices for Bass Ale.* Prices are for a single six-pack and are from Zone 1. 202 observations. Source: *Dominick's.*




Figure 2: *Weekly retail and wholesale prices for Corona.* Prices are for a single six-pack and are from Zone 1. 202 observations. Source: *Dominick's.*

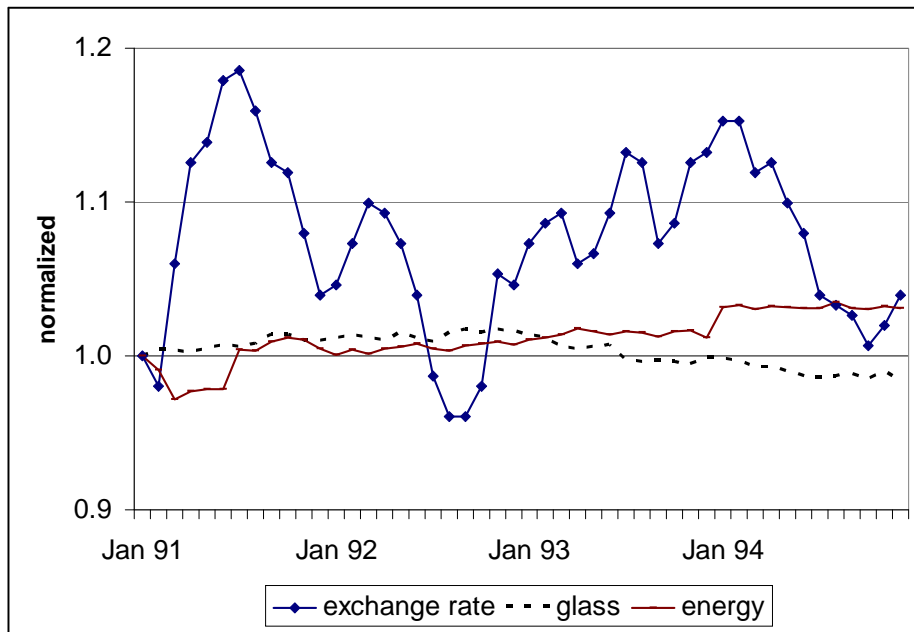


Figure 3: *The nominal exchange rate varies by more than do typical input prices for German brewers.*

Description	Mean	Median	Standard Deviation
Retail prices (\$ per six-pack)	5.44	5.79	1.28
Wholesale prices (\$ per six-pack)	4.50	4.92	1.09
Dummy for retail-price change (=1 if yes)	.18	0	.24
Dummy for wholesale-price change (=1 if yes)	.06	0	.16

Table 1: *Summary statistics for prices for the 16 products in the sample. 6464 observations.*
Source: *Dominick's*.

Description	Mean	Median	Std	Minimum	Maximum
Percent Alcohol	4.52	4.60	.68	2.41	6.04
Bitterness	2.50	2.10	1.08	1.70	5.80
Maltiness	1.67	1.20	1.52	.60	7.10
Hops (=1 if yes)	.12	–	–	–	–
Sulfury/Skunky (=1 if yes)	.29	–	–	–	–

Table 2: *Product characteristics.* Source: "Beer Ratings." *Consumer Reports*, June (1996), pp. 10-19.

	Retail price	Retail price	Wholesale price	Wholesale price	Retail price
Exchange rate	5.96 (1.50)**	6.72 (1.56)**	4.27 (1.50)**	4.74 (1.52)**	
Wholesale price					105.37 (2.53)**
Constant	1.83 (.02)**	1.79 (.16)**	1.66 (.08)**	1.67 (.10)**	.08 (.04)
Observations	3636	3636	3636	3636	3636
R^2	.65	.65	.81	.81	.80

Table 3: *Some preliminary descriptive results.* The dependent variable is the retail or the wholesale price for a six-pack of each brand of beer. The exchange-rate is the average of the previous week’s bilateral spot rate between the foreign manufacturer’s country and the United States (is the number of dollars per unit of foreign currency). All regressions include brand, price-zone, and week fixed effects. The second and fourth columns of the table report results from regressions with controls for domestic and foreign costs. The coefficients on the control variables, and their sources, are reported in the appendix. Robust standard errors in parentheses: Those starred are significant at the * 5-percent or ** 1-percent level. Source: Authors’ calculations.

Variable	OLS	OLS	IV	IV
Price	-.93 (.01)**	-.92 (.01)**	-2.43 (.35)**	-2.43 (.35)**
Holiday		.06 (.02)**		.001 (.01)
First-Stage Results				
F-Statistic			34.45	34.24
Observations	6464	6464	6464	6464
Instruments			wages	wages

Table 4: *Diagnostic results from the logit model of demand.* Dependent variable is $\ln(S_{jt}) - \ln(S_{ot})$. Both regressions include brand fixed effects. Huber-White robust standard errors are reported in parentheses. Costs are domestic wages in the beverage industry interacted with weekly nominal exchange rates for foreign brands. Source: Authors’ calculations.

Variable	Mean in Population	Interaction with Income
Constant	-18.94* (.33)	
Price	-2.55* (.01)	.85* (.05)
Maltiness	-1.16* (.10)	
Bitterness	1.40* (.09)	
Hoppiness	-1.57* (.03)	
Percent Alcohol	4.91* (.05)	-1.08* (.09)

Table 5: *Results from the full random-coefficients model of demand.* Based on 6464 observations. Asymptotically robust standard errors in parentheses. Starred coefficients are significant at the 5-percent level. Source: Authors' calculations.

Brand	Retail price	Manufacturer price	Retail markup	Manufacturer markup
Bass	\$6.92	\$5.76	\$0.43	\$0.37
Beck's	\$5.61	\$4.61	\$0.42	\$0.37
Corona	\$5.72	\$4.75	\$0.39	\$0.36
Guinness	\$7.46	\$5.86	\$0.69	\$0.50
Heineken	\$6.16	\$5.06	\$0.39	\$0.37

Table 6: *Mean prices and price-cost markups for selected brands.* The markup is price less marginal cost with the marginal costs derived from the structural model and with units in dollars per six-pack. Source: Authors' calculations.

Brand	Backed-out	Fitted
Bass	.24 (.078)**	.19 (.076)**
Beck's	.40 (.055)**	.38 (.054)**
Corona	.40 (.064)**	.40 (.066)**
Guinness	.48 (.162)**	.50 (.162)**
Heineken	.47 (.086)**	.36 (.087)**
Molson Golden	.30 (.056)**	.30 (.060)**
St. Pauli Girl	.78 (.092)**	.72 (.082)**
Overall	.37 (.060)**	.35 (.063)**

Table 7: *Backed-out and fitted non-traded costs incurred by the retailer by brand.* Each entry reports the mean across weeks of the backed-out or fitted measure of a brand's non-traded cost to the retailer in cents per six-pack. Standard errors from bootstrap simulations with 200 draws reported in parentheses under each coefficient. Those starred significant at the 1-percent level.

Brand	Bass
wages	.51 (.11)**
constant	-4.33 (.99)**
R^2	.08
Observations	261

Table 8: *Results from regressions of backed-out retailer non-traded costs on determinants.* Dependent variable is retailer's non-traded cost which varies by week. Huber-White robust standard errors are reported in parentheses. Source: Authors' calculations.

Brand	Backed-out	Fitted
Bass	5.30 (.06)**	5.26 (.05)**
Beck's	4.03 (.04)**	4.03 (.03)**
Corona	3.46 (.06)**	3.47 (.06)**
Guinness	5.19 (.07)**	5.16 (.07)**
Heineken	4.47 (.06)**	4.41 (.06)**
Molson Golden	3.29 (.04)**	3.31 (.04)**
St. Pauli Girl	4.22 (.00)**	4.28 (.10)**

Table 9: *Backed-out and fitted total costs incurred by foreign manufacturer by brand.* Each entry reports the mean across weeks of the backed-out or fitted measure of a brand's total cost to the manufacturer in cents per six-pack. Standard errors from bootstrap simulations with 200 draws reported in parentheses under each coefficient. Those starred significant at the 1-percent level.

Brand	Bass
Domestic U.S. wages	.52 (.12)**
Price foreign barley	.22 (.05)**
Foreign wages	.27 (.07)**
Observations	133

Table 10: *Results from constrained linear regression of foreign manufacturer total backed-out costs on determinants.* Dependent variable is manufacturer's total marginal cost for periods when the wholesale price changes which varies by week. Includes brand fixed effects. Source: Authors' calculations.

Brand	Retailer		Manufacturer	
	Mean	Total	Mean	Total
Bass	0.302% (.215)	0.079% (.064)	2.456% (.499)**	0.313% (.091)**
Beck's	0.379% (.332)	0.208% (.179)	3.196% (.060)**	1.127% (.208)**
Corona	0.078% (.054)	0.020% (.022)	0.269% (.128)**	0.060% (.039)
Guinness	0.483% (.666)	0.224% (.324)	0.909% (.772)	0.051% (.060)
Heineken	0.096% (.057)	0.029% (.019)	0.306% (.386)	0.21% (.024)
Molson Golden	0.413% (.232)	0.158% (.091)	2.275% (.611)**	0.688% (.190)**
St. Pauli Girl	0.866% (.299)**	0.395% (.140)**	.	.
Overall	0.370% (.232)	0.121% (.009)**	2.157% (.373)**	0.443% (.077)**

Table 11: *Upper bounds for the retailer's and manufacturers' adjustment costs as a share of revenue by brand.* The entries in the first and third columns report the mean across markets of the estimates of a brand's price-adjustment cost as a share of its total revenue. The second and fourth columns report the sum of the upper bounds on each brand's menu costs divided by the total retail or manufacturer revenue for that brand over the full sample period. Source: Authors' calculations.

Dummy for level change in retail price	5.23 (.1690)**
Constant	.13 (.0615)*
Overall R^2	.21
Observations	3636

Table 12: *Regression of retailer's fixed adjustment costs as a share of its revenue on a dummy for a level retail-price change.* The regression includes brand fixed effects. Source: Authors' calculations.

	Manufacturer:		Retailer:	
	Traded	Markup Adjustment	Traded	Markup Adjustment
<hr/>				
No repricing costs				
Bass	50.1 (.000)**	13.1 (.0995)	11.2 (.0878)	11.3 (.1251)
Beck's	50.1 (.000)**	16.3 (.1108)**	13.5 (.1006)**	13.4 (.1243)**
Corona	50.1 (.000)**	29.9 (.0728)**	26.3 (.0632)**	25.5 (.1236)**
Heineken	50.1 (.000)**	33.4 (.0641)**	27.1 (.0544)**	26.7 (.1160)**
All	50.1 (.000)**	18.3 (.1040)**	15.4 (.0906)**	14.3 (.1326)
<hr/>				
Own-brand repricing costs				
Bass	50.1 (.000)**	0 (.000)	0 (.000)	0 (.000)
Becks	50.1 (.000)**	0 (.000)	0 (.000)	0 (.000)
Corona	50.1 (.000)**	29.9 (.0728)**	26.3 (.0632)**	25.5 (.1236)**
Heineken	50.1 (.000)**	33.4 (.0641)**	27.1 (.0544)**	26.7 (.1160)**
All	50.1 (.000)**	8.3 (.0632)	6.2 (.0344)	6.0 (.0566)
<hr/>				
Competitor-brand repricing costs				
Bass	50.1 (.000)**	0 (.000)	0 (.000)	0 (.000)
Becks	50.1 (.000)**	0 (.000)	0 (.000)	0 (.000)
Corona	50.1 (.000)**	29.4 (.0728)**	26.1 (.0632)**	25.2 (.1236)**
Heineken	50.1 (.000)**	32.1 (.0641)**	26.6 (.0544)**	25.9 (.1160)**
All	50.1 (.000)**	6.8 (.0510)	5.3 (.0302)	5.0 (.0310)

Table 13: *Counterfactual experiments: median pass-through of a 1-percent appreciation of the relevant foreign currency.* Median over 404 markets. Retailer's incomplete pass-through: the retail price's percent change for the given percent change in the exchange rate, attributed to the presence of local dollar-denominated costs or to the retailer's markup adjustment. Manufacturer's incomplete pass-through: the manufacturer price's percent change for a given percent change in the exchange rate, attributed to the share of local dollar-denominated costs in the manufacturer's total costs or to the manufacturer's markup adjustment. Source: Authors' calculations.

	Manufacturer				Retailer			
	Local Costs	Markup Adj.	Costs of Repricing: Own	Other	Local Costs	Markup Adj.	Costs of Repricing: Own	Other
Bass	49.9	37.0	13.1	0	0	0	0	0
Beck's	49.9	33.8	16.3	0	0	0	0	0
Corona	66.7	27.0	0	0.7	4.4	1.1	0	0.1
Heineken	67.3	22.5	0	1.8	7.4	0.5	0	0.4
All	52.5	33.5	10.5	1.6	1.6	0.2	0	0.1

Table 14: *Counterfactual experiments: Decomposition of the incomplete transmission of a 1-percent appreciation of the relevant foreign currency to consumer prices.* Median over 404 markets. Local costs: the share of the incomplete transmission explained by the presence of a local dollar-denominated component in foreign manufacturers' or the retailer's marginal costs. Markup adjustment: the share of the incomplete transmission explained by the retailer or manufacturer's markup adjustment excluding markup adjustment due to fixed costs of price adjustment. Repricing costs own: Fixed costs of price adjustment incurred by the manufacturer or retailer to change its own price. Repricing costs, other: The effect of competitors' costs of price adjustment on the manufacturer or retailer's own price adjustment behavior. Source: Authors' calculations.