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

Pinelopi Koujianou Goldberg
Department of Economics
Yale University and NBER

Rebecca Hellerstein
Research Department
Federal Reserve Bank of New York

September 2005

PRELIMINARY AND INCOMPLETE

PLEASE DO NOT QUOTE

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 at each stage of the distribution chain. 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 nominal price rigidities, local-cost non-traded components, and markup adjustment by manufacturers and retailers in the incomplete transmission of exchange-rate changes to prices. The approach is applied to micro data from the beer market.

*Prepared for the "Microeconomic Sources of Real Exchange Rate Behavior" conference at Carnegie Mellon University, September 16-17, 2005. We are grateful to Charles Engel for very useful conversations at an early stage of this project. The responsibility for any errors contained in this draft is of course our own.

1 Introduction

The incomplete transmission of exchange-rate shocks to consumer prices has been the focus of a substantial amount of theoretical and empirical research. 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 sources of this price inertia within structural models of particular industries (Goldberg (1995), Goldberg and Verboven (2001), Hellerstein (2005)). The basic 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. 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 another potential source of the incomplete price response: the existence of nominal price rigidities. 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 rigidities that would cause firms to (temporarily) deviate from their optimal behavior. The second reason is a practical one. Because the data used in 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 nominal price rigidities, which by nature imply that prices should remain fixed. Hence, to the extent that any 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.

The current paper attempts to extend the previously proposed framework by explicitly introducing price rigidities in 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 nominal price rigidities. In this context we define price rigidities 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 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, etc..

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 magnitude of the price adjustment costs is interesting in its own right as the nature and size of these costs have been the subject of a considerable amount of research in the past.[1] However, 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 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.

[1] 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. Carlton (1986) and Kashyap (1995) find that firms change prices frequently and in small increments, which is not consistent with a menu-cost explanation of price rigidity.

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, and Section 5 presents some preliminary results.

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.

2.1 Market

The imported beer market first developed in the U.S. in the nineteenth century. The Dutch brand *Heineken,* was imported to the U.S. beginning in 1894. The invention of the metal beverage can in 1935 enabled domestic brewers to build national brands without bearing the high fixed costs of maintaining local centers to collect deposit-return glass bottles. Such brands dominated U.S. consumption from that point until quite recently. 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 – resulting in a market share of over seven percent by the end of the decade. Beer is an example of one type of imported good: 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 ten and fifteen percent); both non-traded local costs and price stickiness due to adjustments costs are *a-priori* plausible; and 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 enable us to isolate the role of local nontraded costs and of fixed adjustment costs in firms' incomplete transmission of exchange-rate shocks to their 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.[2] Beer shipments in this market are handled by independent wholesale distributors. The model 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.[3] This practice makes the analysis of pricing behavior along the distribution chain 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.[4] Supermarkets sell approximately 20 percent of all beer consumed in the U.S.[5]

We aggregate data from each *Dominick's* store into one of two price zones. For more details about this procedure, see Hellerstein (2005).[6] 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

[2] The data can be found at http://gsbwww.uchicago.edu/kilts/research/db/dominicks/.

[3] Features of the *Dominick*s' 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. 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.

[4] Canadian Trade Commissioner (2000).

[5] As 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.

[6] The 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.

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 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 to measure the performance of the structural model. We estimate three price equations:

$$\ln p_{jt}^r = \alpha \ln e_{jt} + \beta \ln w_{jt}^d + \gamma \ln w_{jt}^f + \varepsilon_{jt} \tag{1}$$

$$\ln p_{jt}^w = \alpha \ln e_{jt} + \beta \ln w_{jt}^d + \gamma \ln w_{jt}^f + \varepsilon_{jt} \tag{2}$$

$$\ln p_{jt}^r = \ln p_{jt}^w + \varepsilon_{jt} \tag{3}$$

where the subscripts j and t refer to product j in market t where a market is defined as a week and price-zone pair, p^r is the log of a product's retail price, p^w is the log of a product's wholesale price, e is the log of a bilateral nominal exchange rate (domestic-currency units per unit of foreign currency), w^d is the log of a measure of local costs, w^f is the log of a measure of foreign costs, and ε is a random error term. Local- and foreign-cost variables are included in the pricing equation to control for supply and demand shocks other than exchange-rate changes that may affect a brand's price.

Table 3 reports results from estimation of the pricing equations. The first column of the table reports the coefficients from an *OLS* estimation of equation (1). The average pass-through elasticity (α) is 0.12 and is significant at the one-percent level. The domestic- and foreign-cost variables are positive, as one would expect, and significant at the one-percent level. The regression establishes a 12-percent benchmark for the retailer's pass-through elasticity, to be compared with the results from the structural model. The second column of Table 3 reports similar results from estimation of the wholesale-price pricing equation, equation (2): Its pass-through elasticity is 12 percent, and the coefficients on the other variables are positively signed and significant. The coefficient on the foreign-cost variable is higher than in the retail-price regression: Foreign-cost shocks are passed through at a slightly higher level at the wholesale than at the retail stage of the distribution chain, as one would expect. This stage of the distribution chain is closer to the original shock: There is no intermediating optimization of an upstream firm to buffer the shock, as is the case for the retailer. The third 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 close to one and significant, which is consistent with the results from the first two 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. It finds that 12 percent of an exchange-rate change is transmitted to a beer's retail price. Where does the other 88 percent go? This incomplete transmission could be caused by the local costs, markup adjustments, or fixed adjustment costs of manufacturers or retailers, as we examine further in the structural model.

2.4 Patterns of Price Adjustment in the Data

A rough first idea of the timing and frequency of price changes in these data can be obtained from Figure 1, which plots the retail and wholesale prices for a six-pack of the Dutch brand *Heineken*. The figure covers the full sample period, from the middle of 1991 to the middle of 1995. The plot illustrates several important points. First, the figure shows the great advantage of observing price data at a weekly frequency. These data are ideal for analyzing the role of price stickiness, because as the figure illustrates, we clearly see prices remaining constant for several weeks, and then jumping (in a discrete step) to a new level.[7] In other words, the patterns in the price data are exactly those one would expect with price stickiness. This could indicate the existence of adjustment costs, or it could be that prices do not change simply because nothing else changes. The other advantage of the data set is that we observe both retail and wholesale price data for several firms so we can examine the role of price stickiness at both stages of the distribution chain. Retail prices always adjust when wholesale prices change in Figure 1. So it seems that the main reason retail prices do not change in this market is that there is no reason for them to change (the cost facing retailers as measured by the wholesale price does not change).

[7] The role of price stickiness cannot be analyzed within the framework Goldberg (1995, 2005) uses to analyze the auto market or that Hellerstein (2005) uses to study the beverage market. Because the frequency of the data used in these projects is either monthly (Hellerstein) or annual (Goldberg) the econometrician observes prices changing every period given price observations averaged over time. Thus, any price stickiness that may exist is not immediately apparent, or - put differently - cannot be identified from the data.

Now at the wholesale level, it is an entirely different story. There we know that we have enormous variation in exogenous factors affecting costs (namely exchange-rate fluctuations), yet the wholesale prices clearly remain unchanged for long periods of time. Still, wholesale prices do change from time to time, either because there are sales, or because the regular prices eventually adjust.

The price inertia documented in the figure leaves little doubt that both retail and wholesale prices are "sticky" in the sense of not moving in every time period, or even in most. As the source of this price stability does not seem to be an issue that can be settled through simple plots or statistical exercises, we now turn to a more systematic investigation.

3 Model

This section describes the supply and demand sides of the model we use to assess the importance of price rigidities in explaining incomplete exchange-rate pass-through.

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 (2005) 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[8].

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 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

[8] The 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.

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:

$$\Pi_{jt}^{r} = \left(p_{jt}^{r} - p_{jt}^{w} - ntc_{jt}^{r}\right) s_{jt}(p_{t}^{r}) - A_{jt}^{r} \tag{4}$$

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 nontraded[9] 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 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:

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

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. However, one can potentially impose more structure by assuming that adjustment costs are constant over time, and/or constant across products.

[9] We use the term "nontraded" to indicate that these costs are paid in dollars no matter what the origin of the product is. Hence, nontraded costs will not be affected by exchange rate shocks.

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 $p^r_{jt} \neq p^r_{jt-1}$.

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

$$s_{jt} + \left(p^r_{jt} - p^w_{jt} - ntc^r_{jt}\right) \frac{\partial s_{jt}}{\partial p^r_{jt}} = 0, \text{ for } j = 1, 2, ..., J_t. \tag{6}$$

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 $S_{jj} = \frac{\partial s_{jt}(p^r_t)}{\partial p^r_{jt}}$ $j = 1, ..., J$, that is the marginal change in the jth product's market share given a change in the jth product's retail price. The stacked first-order conditions can be rewritten in vector notation:

$$s_t + \Omega_{rt}(p^r_t - p^w_t - ntc^r_t) = 0 \tag{7}$$

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

$$p^r_t = p^w_t + ntc^r_t - \Omega^{-1}_{rt} s_t \tag{8}$$

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

The presence of the adjustment costs A^r_{jt} 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:

$$\left(p^r_{jt} - p^w_{jt} - ntc^r_{jt}\right) s_{jt}(p^r_t) - A^r_{jt} \geq \left(p^r_{jt-1} - p^w_{jt} - ntc^r_{jt}\right) s^c_{jt}(p^r_{jt-1}, p^r_{kt}), \; k \neq j \tag{9}$$

where $s^c_{jt}(p^r_{jt-1}, p^r_{kt})$ denotes the *counterfactual* market share that product j *would* have, if the retailer had kept the price unchanged to p^r_{jt-1}, and p^r_{kt} 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^r_{jt}. By rearranging terms we can use the above inequality to derive an *upper* bound $\overline{A^r_{jt}}$ for the price adjustment costs of product j:

$$A^r_{jt} \leq \overline{A^r_{jt}} = \left(p^r_{jt} - p^w_{jt} - ntc^r_{jt}\right) s_{jt}(p^r_t) - \left(p^r_{jt-1} - p^w_{jt} - ntc^r_{jt}\right) s^c_{jt}(p^r_{jt-1}, p^r_{kt}),\ k \neq j \tag{10}$$

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

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:

$$\left(p^r_{jt-1} - p^w_{jt} - ntc^r_{jt}\right) s_{jt}(p^r_{jt-1}, p^r_{kt}) \geq \left(p^{rc}_{jt} - p^w_{jt} - ntc^r_{jt}\right) s^c_{jt}(p^{rc}_{jt}, p^r_{kt}) - A^r_{jt},\ k \neq j \tag{11}$$

where p^{rc}_{jt} denotes the *counterfactual* price the retailer *would have* charged if he behaved according to the optimality conditions, and $s^c_{jt}(p^{rc}_{jt}, p^r_{kt})$ 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 $\underline{A^r_{jt}}$ for the adjustment costs:

$$A^r_{jt} \geq \underline{A^r_{jt}} = \left(p^{rc}_{jt} - p^w_{jt} - ntc^r_{jt}\right) s^c_{jt}(p^{rc}_{jt}, p^r_{kt}) - \left(p^r_{jt-1} - p^w_{jt} - ntc^r_{jt}\right) s_{jt}(p^r_{jt-1}, p^r_{kt}),\ k \neq j \tag{12}$$

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 $s_{jt}(p^r_t)$ as well as the own and cross price derivatives $\frac{\partial s_{jt}}{\partial p^r_{jt}}$ and $\frac{\partial s_{kt}}{\partial p^r_{jt}}$ can be treated as known. Next we exploit the first-order conditions for each product j (6) to estimate the nontraded 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 nontraded costs are not identified based on the first-order conditions; however, we can derive estimates of the nontraded costs for these periods by imposing some additional structure on the problem, e.g., by modeling nontraded costs parametrically as a function of observables along the lines described in the next section. Once estimates of nontraded costs for these periods have been derived, we can calculate the counterfactual price p^{rc}_{jt} 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 $s^c_{jt}(p^{rc}_{jt,}p^r_{kt})$. In the final step, we can exploit inequalities (10) and (12) to derive upper and lower bounds of the adjustment costs A^r_{jt}.

Note that in the above framework price rigidities as captured by the adjustment cost A^r_{jt} 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. Section ? discusses these simulations.

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^w_{jt} taking the retailer's anticipated behavior into account. Manufacturer w's profit function is:

$$\Pi^w_{jt} = \left(p^w_{jt} - c^w_{jt}(tc^w_{jt}, ntc^w_{jt})\right) s_{jt}(p^r_t\,(p^w_t)) - A^w_{jt} \tag{13}$$

where c^w_{jt} 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^w_{jt},and destination-market specific nontraded costs ntc^w_{jt}. The distinction between traded and nontraded costs is based on the currency in which these costs are paid. Traded costs are by definition incurred by the manufacturer in her home country. As such, they are subject to shocks caused by variation in the nominal exchange rate when they are expressed in dollar terms. In contrast, dollar denominated nontraded costs are (by definition) incurred in the destination market (U.S.), and will not be affected by exchange-rate changes. The term A^w_{jt} 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:

$$
\begin{aligned}
A^w_{jt} &= 0 \text{ if } \; p^w_{jt} = p^w_{jt-1} \\
A^w_{jt} &> 0 \text{ if } p^w_{jt} \neq p^w_{jt-1}
\end{aligned} \tag{14}
$$

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 $p^w_{jt} \neq p^w_{jt-1}$.

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^w_{jt} must satisfy the first-order profit-maximizing conditions given that it has been adjusted from the previous period:

$$
s_{jt} + (p^w_{jt} - c^w_{jt}) \frac{\partial s_{jt}}{\partial p^w_{jt}} = 0 \text{ for } j = 1, 2, ..., J_t. \tag{15}
$$

This gives us another set of J equations, one for each product. Let Ω_{wt} be the manufacturer's reaction matrix with elements $\frac{\partial s_{jt}(p^r_t(p^w_t))}{\partial tc^r_{jt}}$, the change in each product's share with respect to a change in each product's traded marginal cost to the manufacturer. 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 (2004) and Villas-Boas and Hellerstein (2004)[10].

The manufacturers' marginal costs (which are a function of the traded and nontraded costs, tc^w_t and ntc^w_t respectively) are then recovered by inverting the manufacturer reaction matrix Ω_{wt}

[10] To 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^w_f.

according to:

$$p_t^w = c_t^w - \Omega_{wt}^{-1} s_t \tag{16}$$

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

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 changed 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:

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

This conditions is similar to inequality (6) for the retailer, with a slight complication: the counterfactual market share s_{jt}^c that the manufacturer would fact 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:

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

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:

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

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

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

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 conceivable that the retailer adjusts the retail price in periods when the wholesale price remains unchanged. However, in practice we never 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}:

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

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}:

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

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 (**??**) 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:

$$(21) \qquad \left(p^r_{jt-1} - p^{wc}_{jt} - ntc^r_{jt}\right) s^c_{jt}(p^r_{jt-1}, p^r_{kt}) = \left(p^{rc}_{jt} - p^{wc}_{jt} - ntc^r_{jt}\right) s^c_{jt}(p^{rc}_{jt}, p^r_{kt}) - A^r_{jt}, \;\; k \neq j$$

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^{rc}_{jt}, where the latter is determined based on the retailer's first-order condition $s^c_{jt} + \left(p^{rc}_{jt} - p^{wc}_{jt} - ntc^r_{jt}\right)\frac{\partial s^c_{jt}}{\partial p^{rc}_{jt}} = 0$. To find the wholesale price p^{wc}_{jt} the manufacturer would charge in this case, equation (21) can be solved simultaneously with the retailer's first-order condition for $p^{wc}_{jt}$and p^{rc}_{jt}.

The final step in determining the counterfactual optimal wholesale price p^{wc}_{jt}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^{wc}_{jt} 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 random-coefficients model described in Hellerstein (2005), 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 Hellerstein (2005) for a more detailed discussion.

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

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

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:

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

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 $\mu_{ijt}= [\Pi D_i \ \Sigma v_i] * [p_{jt} \ \ x_{jt}]$:

$$u_{ijt} = \delta_{jt} + \mu_{ijt} + \varepsilon_{ijt} \tag{24}$$

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.[11] 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:

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

Let A_j be the set of consumer traits that induce purchase of good j. The market share of good j

[11] The 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.

in market t is given by the probability that product j is chosen:

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

where $P^*(d\zeta)$ is the density of consumer characteristics $\zeta = [D \quad \nu]$ 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:

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

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:

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

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.

4 Empirical Approach

Our empirical approach has two components: estimation and simulation. At the estimation stage, we estimate the demand parameters, the traded and nontraded 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 (2005). We model the mean utility associated with product j at time t as follows[12]:

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

   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 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 $\Delta\xi_{jt}$. Input prices 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 cost variables of two important inputs, the cost of barley in the country of origin of each beer brand in our sample, and the cost of electricity in the Chicago MSA, with indicator variables for each brand; this allows each product's price to respond differently to a given supply shock.

2. **Back out the nontraded retail costs $\mathbf{ntc}_{jt}^r$ 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

[12] The 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.

retailer's first-order conditions for each product j (6) to estimate the nontraded retail costs of product j.

3. **Model these nontraded 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 nontraded 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 nontraded costs for these periods we employ the following procedure:

   First, we collect the data on the nontraded costs ntc^r_{jt} 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:

$$ntc^r_{jtz} = \gamma_{jz} d_z + \gamma_{jw} w^d_t + \eta_{jtz}$$

   where d_z are price zone dummies, and w^d_t denote local wages. We let the coefficients be product-specific, as price changes are not synchronized so that we cannot pool data across products. We run the above regression using data from the periods we observe price adjustment, and then use the parameter estimates to construct the predicted nontraded 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 $\mathbf{A}^r_{jt}$.**

   With the demand parameter and nontraded cost estimates in hand, we employ (10) and (12) to derived the upper and lower bounds of the retailer adjustment costs A^r_{jt}. The computation of the upper bound is straightforward: in (10) all variables are observed, except for the counterfactual market share $s^c_{jt}(p^r_{jt-1}, p^r_{kt})$ 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^{rc}_{jt} that the retailer would charge if she changed the retail price from the previous period, and the associated market share $s^c_{jt}(p^{rc}_{jt}, p^r_{kt})$. These are computed using (8) which reflects the first-order condition of the profit maximizing retailer.

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

   The procedure here is similar as the one we employ to derive the nontraded 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 $\mathbf{c}^{w}_{jt}$.

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 nontraded 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^{w}_{jt} takes the form:

$$c^{w}_{jt} = \exp(\theta_j d_j + \theta_{jt} d_t + \omega_{jt})(w^{d}_{t})^{\theta_{jdw}}(e_{jt} w^{f}_{t})^{F_j * \theta_{jfw}} \tag{29}$$

   or, in log-terms:

$$\ln c^{w}_{jt} = \theta_j d_j + \theta_{jt} d_t + \theta_{jdw} \ln w^{d}_{t} + F_j * \theta_{jfw} \ln(e_{jt} w^{f}_{t}) + \omega_{jt} \tag{30}$$

   where w^{d}_{t} and w^{f}_{t} denote local domestic and foreign wages respectively, e_{jt} is the bilateral exchange rate between the producer country and the U.S., and F_j is a dummy that is equal to 1 if the product is produced by a foreign supplier, and zero otherwise. All variables and parameters are indexed by j to indicate that we estimate a separate equation for each product, as the asynchronous adjustment of wholesale prices implies that we cannot pool data across products given that we confine the estimation to the periods for which we observe price changes, and these periods are different for each product. For the function to be homogeneous of degree 1 in factor prices, we require $\theta_{jdw} + F_j * \theta_{jfw} = 1$. 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 nontraded 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 nontraded costs. The latter are captured in the above specification by the term $(w_t^d)^{\theta_{jdw}}$ 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 nontraded cost component is nonzero (so that $\theta_{jfw} < 1$[13]), 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 nontraded 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 $\mathbf{A}_{jt}^{w}$.**

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

[13] Given 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 $\theta_{jfw} < 1$.

evaluated at this retail price, to the profits that the retailer would make if she kept the retail price unchanged at p^r_{jt-1}. The retailer will choose the price associated with the higher retail profits. Once the counterfactual retail price p^{rc}_{jt} has been determined this way, the associated counterfactual market share $s^c_{jt}(p^{rc}_{jt}(p^w_{jt-1}, p^r_{kt})), k \neq j$, 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^{rc}_{jt} and p^{wc}_{jt} 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^{wc}_{jt} and p^{rc}_{jt} are then determined according to the manufacturer and retailer first-order conditions, equations (16) and (6) respectively, with the inverted manufacturer reaction matrix Ω^{-1}_{wt} reflecting the optimal pass-through of the wholesale price change onto the retail price. Let π^{wc}_1 denote the manufacturer profits associated with the so-computed prices p^{wc}_{jt} and p^{rc}_{jt}.

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^r_{jt-1}. The optimal wholesale price will then be determined based on equation (21) along the lines discussed in the previous section. Let π^{wc}_2 denote the manufacturer profits associated with the prices p^{wc}_{jt} and p^r_{jt-1} in this case.

If $\pi^{wc}_1 > \pi^{wc}_2$, the manufacturer will set the wholesale price anticipating that the retailer will adjust her price too. Hence, the counterfactual prices p^{wc}_{jt} and p^{rc}_{jt} will satisfy the conditions described under the first case above. If $\pi^{wc}_1 < \pi^{wc}_2$, 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^r_{jt-1}.

Once the counterfactual wholesale and retail prices have been determined, evaluation of the adjustment cost lower bound based on (20) is straightforward.

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

(to come)

5 Results

This section first discusses results from the estimation of the demand system. It then describes preliminary estimates of brand-level markups, nontraded costs, and upper and lower bounds on both the retailer's and the manufacturers' adjustment costs.

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.

Table 4 suggests the instruments may have some power. The first-stage F-test of the instruments, at 39.47, 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 -1.49 in the IV estimation. Note that the 95-percent confidence interval of the latter coefficient does not include the value of the former.

5.2 Demand: Random-Coefficients Model

Table 5 reports results from estimation of the demand system. We allow consumers' age and income to interact with their taste coefficients for price, maltiness, bitterness, hoppiness, 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.

The coefficients on the characteristics generally appear reasonable. As consumers' income rises, they become less price sensitive. The random coefficient on income, at 1.48, is significant at the five-percent level. The mean preference in the population is in favor of a malty and/or hoppy taste in beer: Both characteristics have positive and significant coefficients. The mean coefficient on a bitter flavor is negative, though the random coefficient that reflects the influence of unobserved heterogeneity in the population on the mean coefficient is positive and significant, indicating a taste for bitterness in some parts of the population. Finally, as the percent alcohol rises, the mean utility in the population also rises, an intuitive result. The minimum-distance weighted R^2 is .50 indicating these characteristics explain the variation in consumer demand fairly well.

5.3 Retail Markups and Nontraded Costs

Table 6 reports retail and wholesale prices and markups for selected imported brand. The markups are derived using firms' first-order conditions. Under the assumption of no adjustment costs, the markups are 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. The markups appear reasonable. Table 8 compares the retailer's nontraded costs estimated in periods when a brand's price adjusted with the fitted values from a regression of these nontraded costs on their determinants for those periods in which a brand's price did not adjust. It appears that the regression produces fitted values that match the means of each brand's non-traded cost series fairly well.

5.4 Results: Adjustment Costs

Table 9 reports the median estimates by brand of the upper and lower bounds on the retailer's and manufacturers' adjustment costs as a share of their total revenue from that brand. 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 *Guinness* to 5.3 percent of revenue for *Becks*, with a median upper bound across the foreign brands of 2 percent. The upper bounds on adjustment costs for wholesale prices range from 1.6 percent of revenue for *Heineken* to 6.3 percent of revenue for *Beck's*, with a median upper bound across foreign brands of 4.4 percent, more than twice the median upper bound on the retailer's adjustment costs.

Figures 2 and 3 illustrate the relationship of the mean market share to the mean upper bounds on adjustment costs in dollar terms and on adjustment costs as a share of total revenue for both the retailer and the manufacturer. Figure 2 clearly illustrates that larger upper bounds on retailer and manufacturer adjustment costs are positively associated with brands' market shares. Figure 3 indicates that the upper bounds on retailer and manufacturer adjustment costs, as a share of their total revenue, are negatively associated associated with brands' market shares.

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] Berry, S., J. Levinsohn, and A. Pakes. 1995. "Automobile Prices in Market Equilibrium," *Econometrica*, vol. 63, no. 4 (July), pp. 841-90.

[4] Berry, S. 1994. "Estimating Discrete-Choice Models of Product Differentiation," *RAND Journal of Economics*, vol. 25, no. 2, pp. 242-62.

[5] Bils, M. and P.J. Klenow. 2005. "Some Evidence on the Importance of Sticky Prices." *Journal of Political Economy.*

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

[7] Carlton, D. 1986. "The Rigidity of Prices." *The American Economic Review*, vol. 76 (September), pp. 637-58.

[8] 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 (Nov), pp. 683-703.

[9] Engel, C. 2003. "On the Relationship between Pass-Through and Sticky Nominal Prices." Working Paper.

[10] 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), ppl. 891-951.

[11] 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.

[12] 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.

[13] Hellerstein, R. 2005. "A Decomposition of the Sources of Incomplete Cross-Border Transmission." Federal Reserve Bank of New York Working Paper.

[14] Kashyap, A. 1995. "Sticky Prices: New Evidence from Retail Catalogs." *The Quarterly Journal of Economics*, vol. 110 (February), pp. 245-74.

[15] 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*, August.

[16] Slade, M. 1998. "Optimal Pricing with Costly Adjustment: Evidence from Retail-Grocery Prices." *The Review of Economic Studies*, vol. 65 (January), pp. 87-107.

[17] Villas-Boas, S.B. and R. Hellerstein. 2005. "Identification of Supply Models of Retailer and Manufacturer Oligopoly Pricing." *Economics Letters*, forthcoming.

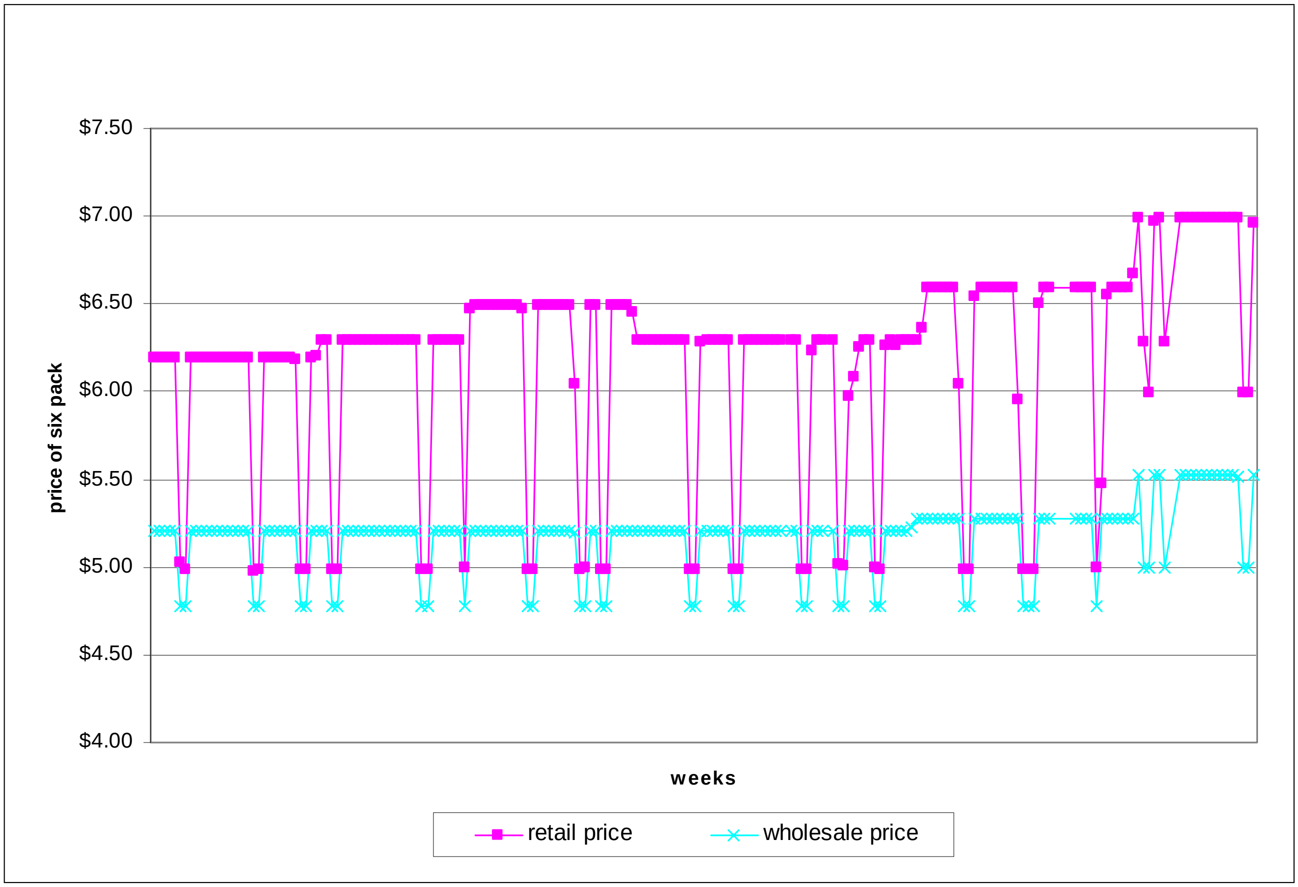

Figure 1: *Weekly retail and wholesale prices for Heineken beer.* Prices are for a single six-pack and are from Zone 1. 202 observations. Source: *Dominick's.*

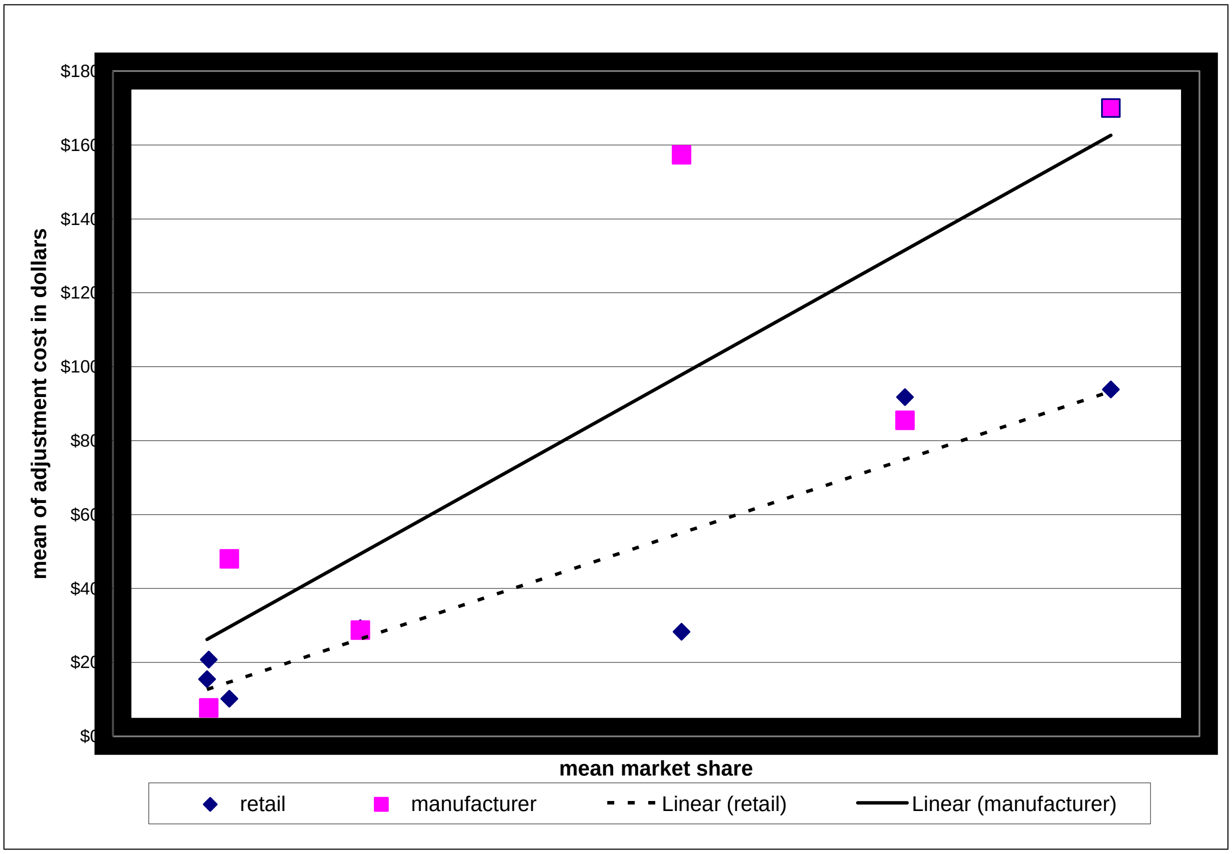

Figure 2: *Larger upper bounds on retailer and manufacturer adjustment costs are associated with brands with larger market shares.* Source: authors' calculations.

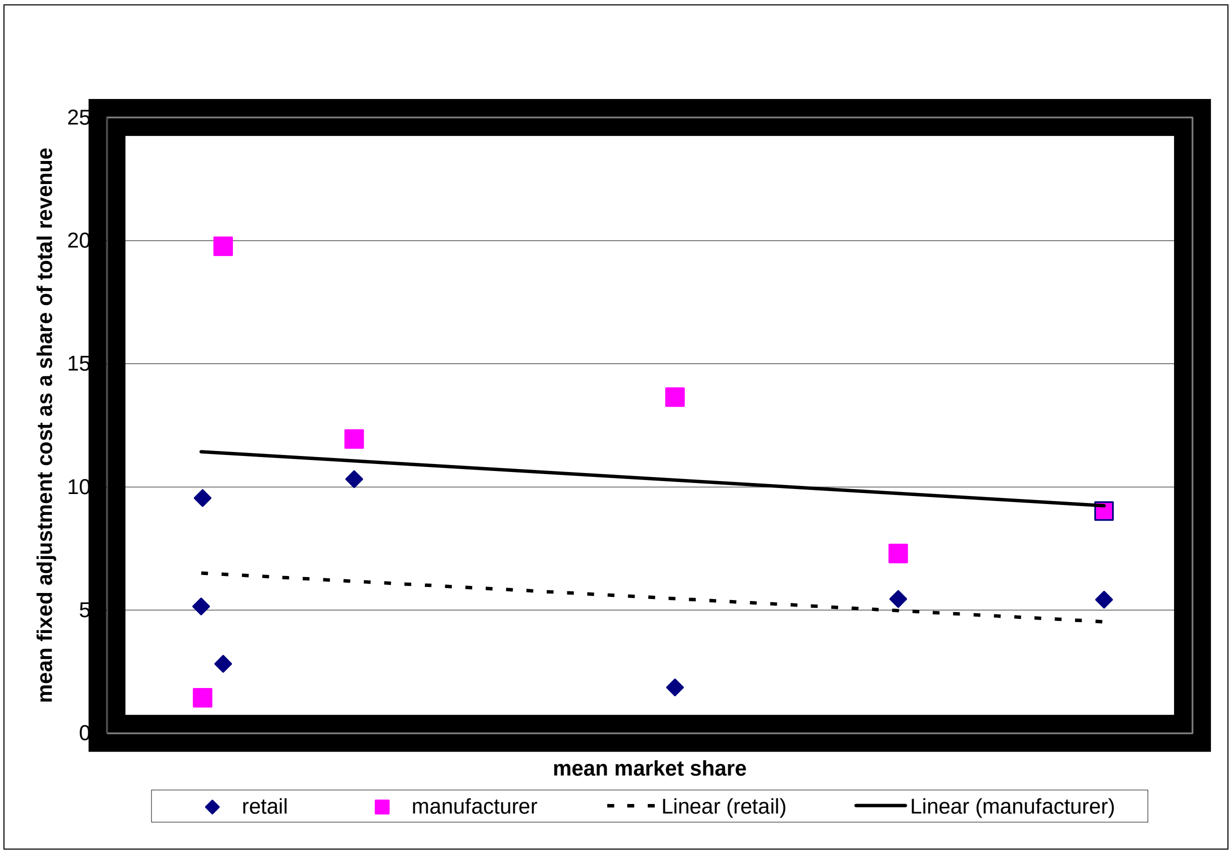

Figure 3: *The upper bounds of retailer and manufacturer adjustment costs, as a share of their total revenue, are negatively associated with brands' market shares.* Source: authors' calculations.

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)	.37	0	.49
Dummy for wholesale-price change (=1 if yes)	.11	0	.32

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.

	Wholesale price	Retail price	Retail price
Exchange rate	.12	.12	
	(.002)**	(.003)**	
Foreign costs	.01	.004	
	(.001)**	(.002)**	
Domestic costs	.28	.35	
	(.02)**	(.07)**	
Wholesale price			.99
			(.01)**
Constant	.36	1.06	.18
	(.09)**	(.16)**	(.07)**
R^2	.20	.18	.72

Table 3: *Some preliminary descriptive results.* The dependent variable is the retail or wholesale price for a six-pack of each imported brand. Foreign costs are monthly prices of barley in the country of origin of each brand. Domestic costs are monthly hourly wages in supermarkets in the Chicago MSA for the retail price and hourly wages in the distribution sector for the manufacturer's price. The exchange-rate is the weekly average of the previous week's spot rate. Source: Authors' calculations.

Variable	OLS	IV	
Price	-.93	-1.49	.
	(.01)**	(.12)**	
Measures of Fit			
R^2	.86		
First-Stage Results			
F-Statistic		39.47	
Observations	6464	6464	
Instruments		costs	

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 costs proxied by domestic energy use and foreign costs proxied by malt and barley prices which are described in the text. Source: Authors' calculations.

Variable	Mean in Population	Interaction with: Unobservables	Age	Income
Constant	-2.68*			
	(.33)			
Price	-2.21*			1.49*
	(.59)			(.24)
Maltiness	2.11*		-1.22	
	(.33)		(1.22)	
Bitterness	-.28*	.44*		
	(.09)	(.22)		
Hoppiness	4.95*	-.93		
	(.28)	(.78)		
Percent Alcohol	7.09*			-1.32*
	(.19)			(.26)
GMM Objective	182.94			
M-D Weighted R^2	.50			

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	6.91	5.84	.77	.62
Beck's	5.74	4.40	.92	.48
Corona	5.68	3.65	1.06	.38
Guinness	7.44	5.69	.97	.71
Heineken	6.18	4.70	.97	.44
Molson Golden	5.09	3.62	.47	.38
St. Pauli Girl	6.22	4.61	.94	.44

Table 6: *Median 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	Observed	Derived
Bass	28	40
Beck's	19	19
Corona	22	18
Guinness	52	55
Heineken	33	24
Molson Golden	83	102
St. Pauli Girl	54	61
Overall	42	46

Table 7: *Observed and derived non-traded costs incurred by the retailer by brand* Each entry reports the mean across markets of the observed or derived measure of a brand's non-traded cost to the retailer in cents per six-pack.

Brand	Bass	Becks	Corona	Guinness	Heineken	Molson	St. Pauli
wages	.50	.06	.23	.35	.54	.32	.07
	(.19)**	(.09)	(.12)	(.26)	(.21)**	(.16)**	(.11)
feature	-.54	-.22		-.51	-.33	-.92	-.29
	(.09)**	(.09)**		(.20)**	(.12)**	(.11)**	(.11)**
transport costs	3.08						
	(1.31)**						
R^2	.55	.10	.10	.26	.29	.56	.21
Observations	32	47	25	27	23	37	33

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

	Retail		Manufacturer	
Brand	Upper Bound	Lower Bound	Upper Bound	Lower Bound
Bass	3.0%	0.0%	4.6%	0.0%
Beck's	5.3%	0.0%	6.3%	0.0%
Corona	0.3%	0.0%	2.6%	0.0%
Guinness	0.1%	0.0%	4.8%	0.0%
Heineken	1.5%	0.0%	1.6%	0.0%
Molson Golden	0.6%	0.0%	5.8%	0.0%
St. Pauli Girl	3.4%	0.0%	·	0.0%
Overall	2.0%	0.0%	4.4%	0.0%

Table 9: *Upper and lower bounds for the retailer's and manufacturers' adjustment costs by brand.* Each entry reports the median across markets of the estimates of a brand's price-adjustment cost as a share of its total revenue.